An Application of Matching-Pursuit for the Characterization of Glitch Noise in Gravitational Wave Interferometers

Andrew Altheimer, Rice University, in collaboration with Prof. I.M. Pinto, University of Sannio

1. Introduction

The Ligo and Virgo Gravitational Wave (GW) Interferometers are currently the three most sensitive GW observatories in the world. While all three observatories have seen some operation over the past several years, they are currently undergoing significant upgrades and are scheduled resume operation in 2009. This downtime provides us with a prime opportunity to undergo a more detailed study of the significant background noise seen in these detectors.

In addition to a constant Gaussian background, there has also been observed a significant impulsive noise component, commonly referred to as glitch noise. Glitches are believed to be caused by either equipment malfunction or environmental disturbances and have been observed at a high rate, some several thousand per day per detector. While in some cases auxiliary channels (seismic detectors, etc) may be used to veto glitches, the source of the majority of these glitches is unknown. Distinguishing many of these surviving glitches from Gravitational Wave Bursts (GWBs) may be nearly impossible using data from a single detector.

Much attention has been given to the use of coherent analysis techniques to distinguish between noise and GWBs by examining the correlation between multiple detector outputs in a redundant network of three or more detectors [1,2,7]. However, these techniques still have some significant difficulties in distinguishing between glitches and GWB's. While these methods are effective at vetoing single glitches, the high glitch rate means that there remains a significant risk for glitches to occur simultaneously in multiple detectors. In some cases, such simultaneous glitches may be nearly indistinguishable from GWB's, especially in the case in which both the waveform and the direction of origin of the GWB are unknown.

Particularly problematic is the fact that for any given detector there exist regions of the sky and possible polarizations for which the expected response to a GWB is below the detection threshold. As a consequence, in a three detector network, it becomes very difficult to distinguish between pairs of simultaneous glitches occurring in the two remaining detectors and a GWB of unknown waveform originating from these directions [7].

We believe that a better understanding of glitch noise is essential to resolving these problems. Over the past couple of months, we have begun an attempt to characterize glitch noise through decomposition into time-frequency atoms. In particular, we are interested in applying a matching pursuit (MP) algorithm of the type originally proposed by Mallat and Zhang [3]. It is our hope that in doing so, we will be able to obtain a better model for glitch noise which will prove invaluable in future attempts to distinguish between glitches and GWBs. While other attempts have been made to represent glitches using time frequency atoms [6], we believe we are the first to attempt decomposition via MP and that this algorithm may offer a significant improvement over previous approaches.

2. Modeling Glitches

We began our study with the hypothesis that glitches can be conveniently modeled as a composition of very simple time-frequency atoms. Time-frequency atoms are defined as functions with a nearly compact support in the time frequency plain, the simplest example being a Sine-Gaussian. Our hypothesis stems partly from observation and partly from reasoning about the nature of glitches. We assert that a glitch acts as an impulse or series of impulses at some point or points in our detector. Before it may be observed, it must travel from its initial location(s) to one of the mirrors or other crucial component of our interferometer. The shape of the observed glitch may then depend primarily on the nature of the transfer functions between detector components, rather than the initial impulse itself, resulting in simple waveforms.

Experimental evidence suggests that there are two primary categories of glitches, which essentially appear either as Sine-Gaussians (SGs) or as a Ring-Downs (RDs). In some cases a number of such functions may be observed in very rapid succession. SGs and RDs have the following respective forms

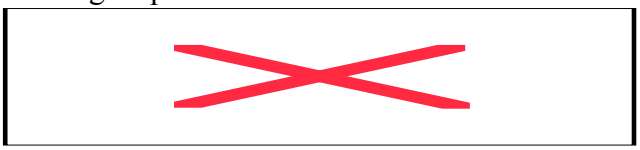

Both of these functions have time parameter u, frequency center ω and phase φ. However, the SGs have a single width parameter, the standard deviation S, while the RDs are more complicated with stretch factor α and shape parameter k. The actual width of a RD depends upon both parameters, with a standard deviation of $S = \sqrt{k+1}/\alpha$. Likewise, the actual peak of a RD is located at u+k/α.

As a result of this complication, and of the larger parameter space (five parameters instead of four), RDs are substantially more difficult to analyze, as we will discuss later. For this reason, we chose to focus our attention primarily on SGs for the time being. We will return to RDs in a later section of this report.

3. Matching Pursuit Algorithms

Over the next few sections I will describe MP as it has been implemented in our analysis. I will not go into many of the more technical details of MP, but will instead focus on details which are directly related to our project. For a more thorough treatment, I will refer readers to the original paper by Mallat and Zhang [3].

3.1 Overview

It will be useful to begin by briefly outlining the basic structure of our MP algorithm. The principle behind MP is a simple one. Given a time-sampled signal, we wish to decompose our signal into a summation of simple functions (henceforth dictionary elements), chosen from a given family (henceforth, a dictionary). The most typical real-valued dictionary and of course the one in which we are primarily interested is the family of SGs.

In general and certainly in our case, dictionaries will be over-complete and thus necessarily non-orthogonal and linearly dependent. As a consequence, this decomposition is not unique, as would be the case if we were decomposing our signal into an orthogonal basis (such as sinusoids in a Fourier Transform.) The problem at the heart of MP algorithms is that of selecting the decomposition that requires as few significant terms as possible to represent our signal. This is commonly referred to as the principle of scarcity.

While algorithms for finding globally optimal decompositions do exist [4], they are too computationally expensive for our purposes. Rather the MP algorithm operates by iteratively selecting and removing the element with the strongest correlation. This is what is termed as a 'locally optimal' solution.

In order to make MP practical, we must first select a finite subset of our dictionary, referred to as our subdictionary. This is easily done by choosing a discreet subset of the parameters in our parameter space. Once this selection has been made, we then proceed as outlined below:

1. Coarse Stage: Compute the correlation between our signal and each element of our subdictionary and store the results in a table.

2. Fine Stage: Select the subdictionary element with the maximal correlation, and perform a local search in our parameter space to find the nearby dictionary element with maximal correlation.

3. Adjust our correlation table to account for the removed element, and repeat fine stage.

In addition to computational efficiency and strong performance, MP decompositions are unusual in that they also satisfy a sort of energy conservation. The energy of a signal or signal component is generally defined to be the square of the norm. Because we are working with a non-orthogonal basis of elements, we do not necessarily expect to see energy conservation, i.e. we do not expect for the energy of a signal to be

equal to the sum of the energies of its components. However, because MP optimizes the decomposition of each removed component individually, it is easy to show that each residual must in fact be orthogonal to the most recently removed atom. As a result, a simple inductive argument will show that energy is in fact conserved in such a decomposition.

A more detailed description of our MP algorithm and of the differences between our algorithm and that originally proposed by Mallat and Zhang follows in subsequent sections. For the time being we have assumed that signals have a sampling rate of 2^14 samples per second and that glitches may be contained in a signal ½ seconds in length for the purposes of analyzing our algorithms performance.

3.2 Complex Subdictionaries

Before describing our algorithm in more detail, it will be necessary to cover a slight digression: the very close connection between real and complex sinusoids. Using this connection, we will devise a means to find the correlations between elements in a real subdictionary with a real signal while only considering elements of a smaller complex subdictionary of the form $g * e^{iwt}$. This will allow us to find near-optimal real dictionary elements by considering only a complex subdictionary with one fewer parameter, resulting in a significant reduction in the necessary computational resources.

Let either $\langle g, f \rangle = \int g f^* \, dt$ or $\langle g, f \rangle = g \bullet f^*$ be our inner product. Given a real valued function g(t) and real valued signal f(t), there exists a φ on [0, π) such that  is maximized. Furthermore, it can be shown that for this particular φ:

Provided that  is small the following holds :

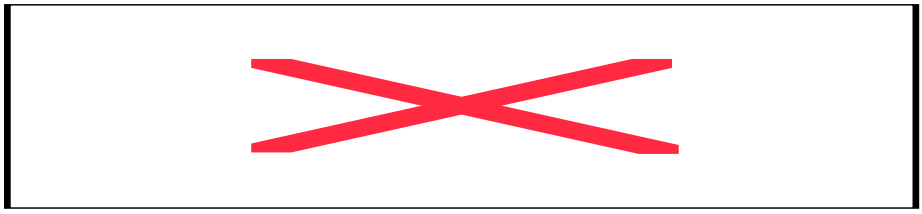

Thus, 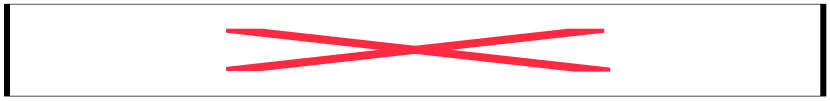 .

Furthermore, if the supremums are realized, then both must occur at the nearly the same

values of S, u, and ω. Lastly, the supremum φ occurs very nearly at  . This is of course true regardless of whether the parameter space for S, u, and ω is discreet or continuous. In the case in which the supremums are not realized, then analogous arguments may be made about points arbitrarily close to the supremum.

The above atoms, however, are unnormalized, whereas we are really interested in their normalized counterparts. In this case, the arguments given above essentially still hold, as  . However, the relevant normalization constant N is that for the real valued function rather than its complex counterpart and this normalization depends upon the phase. The consequence is that if we wish to work only with a complex dictionary, then normalization can be completed only *after* the correlations are calculated and the optimal values of φ have been found. This is easy enough to do however, since:

 .

Thus we now have a method for finding the near-supremum correlation over our desired real subdictionary by considering only a smaller complex dictionary of one fewer parameters!

It is worth noting, however, that the φ which we have found is the optimal value only over our unnormalized dictionary. Since our normalization constant itself depends on φ, this is not quite the optimal value over a normalized dictionary. This error can in turn slightly skew our optimal values for the other parameters. However, this error should be very small, and is generally more than made up for by the fact that a finite complex subdictionary is effectively continuous in φ, whereas any real subdictionary must contain only discreet values of φ.

Additionally, we must consider our claim that 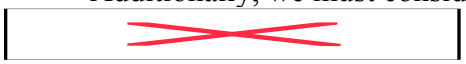 is very small. This has been tested to hold for both SGs and RDs, with an error in the correlation on the order of several parts in a thousand or less. This error may become larger with extremely low frequencies or narrow glitches, but has not been observed to be a problem.

The coarse stage analysis does not need to produce exact values for our parameters, only close enough to be successfully optimized in the fine stage. The accuracy lost by switching from a real to a complex subdictionary are relatively small. This error was in fact dominated by the loss in precision resulting from our use of a finite dictionary of discreet parameters rather than our original continuous dictionary. In some cases the best match produced in the coarse stage even improved under a complex dictionary due to a better resolution of the phase parameter.

3.3 Subdictionary Selection and Coarse Analysis

Now that we have shown how a complex subdictionary may be used in place of a real one, we may address more specifics about our subdictionary. The selection of a subdictionary is not surprisingly a trade off between computational efficiency and performance, and finding the correct balance is still an open question. In the meantime, we have selected a subdictionary which seems both computationally manageable and effective, finding atoms whose correlations are within 5% of the maximum over our entire SG dictionary.

As can be gathered from the previous section, the subdictionary in which we are interested is actually not that of real SG's, but rather complex SG's, more commonly referred to as complex Gabor functions. These, of course, have the form  . This allows us to temporarily ignore the phase parameter in our real dictionary, resulting in a very significant reduction in the required computational resources.

Our subdictionary is created as follows. The standard deviation is varied on a logarithmic scale from S=32 to S=1024 time samples (~.0020-.0625 seconds), with a resolution of half powers of two. The frequency is then varied from 0 to 1638 hz such that at each step, the wave number (ie, the number of wavelengths in one standard deviation) increases by one fourth. The position is likewise varied over the length of the signal in steps of one quarter the length of the standard deviation. We thus have increasing frequency resolution and decreasing position resolution as S increases.

A complex Gabor for each width-frequency combination is computed out to 3 standard deviations in either direction and is stored in a look up table. For obvious reasons, functions which are merely time translated versions of each other are not stored. As the program currently stands, this lookup table requires a very manageable 135mb of memory (an order of magnitude less than achieved with a real dictionary.)

It should be noted, however, that the memory and processing power required by our program grow very quickly with the largest tested width parameter S, and that our algorithm becomes increasingly expensive if one desires to look for glitches with standard deviation greater than a tenth of a second or so. This is of course despite the logarithmic scale.

Once our complex subdictionary is generated, computing the correlations is a straightforward task. The inner product is taken between our signal and time shifted versions of every element in our table. The phase is then determined from the angle of this correlation, and the correlation is adjusted by the relevant normalization factor. If the resulting correlation is of sufficient magnitude, the 3 parameters of our Gabor along with our phase are passed along as parameters for an SG in the next stages of the analysis.

If one desires a significantly finer resolution in the position parameter, it may become worthwhile to convolve the signal with the complex Gabor rather than compute each translated inner product explicitly. However, our current resolution is course enough that this actually results in a significant increase in computation time.

It should be noted that there is an additional difference between our implementation and that originally proposed by Mallat and Zhang. For computational reasons, we found

it convenient to store only the top matches in our correlation table (up to 25,000 entries or so), rather than the entire subdictionary. Since each stored correlation must be adjusted after every removed element, this significantly speeds up the decomposition. MP may then continue until the correlation of extracted elements fall below a certain threshold, at which point we must compute a new correlation table if we wish to continue.

3.4 Fine Stage Analysis

Once a correlation table has been computed, we select the maximal entry and the associated parameters. This entry should be close enough to the global maximum correlation over our entire parameter space that if a nearby local maximum is found, we may hope that it is in fact a global maximum. In their paper, Mallat and Zhang give a proof that even if this is not the global maximum, it must be very nearly as strong a match as the global maximum [4].

While Mallat and Zhang proposed using Newton's method to perform the fine stage analysis, we found it was more practical to perform a binary search on the parameter space describing our dictionary. While this method may require more iterations to converge, it bypasses the need for finding derivatives and second-derivatives of our correlation, and is therefore less computationally expensive.

For an SG dictionary in particular we discovered that it is possible to do this on each parameter individually. Each parameter is optimized in turn, and the process is repeated until our position in parameter space doesn't move over a complete iteration through each parameter. Surprisingly, if the parameters are optimized in the correct order (phase, position, frequency, and finally width), the algorithm generally concludes after only 2-3 iterations.

Once this is completed, this element is extracted from our signal, the correlation table is recomputed, and a new set of parameters are passed into the fine stage. This may be continued until a desired energy threshold is reached or until a desired number of time-frequency atoms have been removed.

4. Performance

4.1 Matching Simulated Glitches

While the paper by Mallat and Zhang devoted significant discussion the performance of their algorithm [4], there are rather important differences between the original intended use of the algorithm and our proposal. Mallat and Zhang intended their algorithm to be a means of approximating any arbitrary signal in its entirety as a sum of several hundred elements. Rather than this, our hope is to be able to adapt this algorithm to approximate a very simple signal such as a single glitch by picking out either one or a very small number of specific atoms, while leaving the Gaussian background intact.

For this reason, while we are able to ignore several issues handled by Mallat and Zhang (especially a process called back projection), greater care is required in our selection of our dictionary and subdictionary and in our performance of the fine stage. Furthermore, some verification of the algorithms performance is necessary. More precisely, we desire to confirm that if given a signal containing Gaussian noise and one or

more SG's, MP will successfully find the inputted SG's and correctly determine their parameters.

To this end, signals of ½ seconds in length were generated containing between one and three SG's. While the observed glitches about which we are most worried have typical SNR's ranging from around 2 to 20, we tested the algorithm with SNR's starting at around 1/3.

Even with an SNR of 1/3, the algorithm had no difficulty in finding a single glitch. In this case, the programs output clearly reveals the presence of a single high energy atom in the signal followed only by much lower energy components. In the majority of cases, even when multiple components are present in the signal, MP had no trouble finding them as distinct elements, distinguishing between them, and again returned only low energy elements after removing them. However, the exception to this not surprisingly occurs in cases where two glitches of very similar frequency and position are both present in the same signal, in which case the algorithm can not resolve them into separate atoms.

Getting a good match on the parameters is a bit trickier in cases where the SNR ratio is very low, as the Gaussian background significantly distorts the parameters of the highest correlating atom. For an SNR of 1/3, the discovered parameters showed significant deviations from those of the simulated glitch, while an SNR of 2 produced closely matching results.

Some of the parameters are more difficult to match than others. Finding the correct phase, for instance, is essentially impossible. Since the phase is measured relative to the peak of the SG, even a slight error in the estimated position center has a significant impact on the phase. However, this is a fundamental problem and is not unique to our algorithm. Among the remaining parameters, the width is the most difficult to match. In the case when the SNR is low, at about 1, this value is generally only correct to about +/- 10%, but improving quite significantly as the SNR increases. Figures 1-4 in the appendix show a distribution of the correlations and relative errors in each of the parameters between the discovered atom and the simulated glitch for 600 simulated SG's with an SNR of 2. One can see that even at this relatively low SNR, the MP algorithm is generally successful at matching input parameters.

4.2 'Real' Data?

We have thus shown that our MP algorithm can successfully decompose signals which contain actual SGs. However, real glitches are unlikely to be true SGs, and we must therefore test the performance of our algorithm on real data. While we do not have independently determined parameters to compare our results to, we may confirm that our algorithm successfully approximates a glitch using only a small number of elements. Unfortunately, due to administrative complications, we are still in the process of obtaining data directly from Ligo archives. We hope that this data will be available to us soon, and that we will be able to conduct more extensive testing of our algorithm at that point.

In the meantime, we have managed to obtain a few glitches from a library available at [5]. However, in addition to being very small, this library is very limited for

several reasons. First, the selection criterion for the glitches contained in the library is unknown to us and the available glitches may not be a good representation. Additionally, we do not have the filters which have been applied to this library available to us and do not know exactly what filtering process has been used. Lastly, the signals have been stored as .wav files intended to allow interested parties to 'listen' to glitches. The resulting resolution in the signal amplitude is a bit low.

For the time being, simple signals containing only a single apparent SG were selected from the database and analyzed by our program. In order to analyze the performance, we have made a crude attempt to estimate the energy of each glitch in the library. This was done by finding the energy of a section of signal containing a glitch and subtracting out the energy of a section immediately adjacent containing only Gaussian white noise. However, this assumes that the background is completely stationary and may result in errors in our energy estimates. It also assumes that the signal residual produced by our algorithm should be similar to adjacent sections of Gaussian noise. In many cases, however, the energy of this residual was significantly *smaller* than adjacent Gaussian noise, especially after removing multiple atoms.

Using this method, the first removed atom seems to have an energy .7-1.05 times that which we estimate for the glitch itself. Removing additional atoms improves the approximation significantly, generally starting at around an additional 5% of the estimated energy. The significant drop in energy between the first and second atoms suggests that our glitch contains a single large component which is only a moderate match with a SG. Additional removed components are either corrections to the imperfect initial match, or matches for much smaller signal components.

Figures 5 and 6 in the attached appendix show a typical glitch signal and parameters of the first 6 matched SG atoms. Figures 7-12 show the resulting approximations and residuals of this signal using after removing one, three, or six SGs. One can see that the approximation converges to our glitch rapidly, and after removing 6 atoms the residual appears to have only a very small non-Gaussian component. The energy of the original signal was around 4.178, while the residual after removing all 6 atoms is around .136. The energy of adjacent sections of supposedly Gaussian background noise is .115, suggesting that we have in fact produced a very close approximation for our glitch.

Our algorithm seems to be equally effective at approximating more complicated glitches. Figures 13-15 show a more complicated glitch along with the resulting approximation and residual from a 10 atom decomposition.

There is also one additional uncertainty which we have yet to seriously test, which is the effect which filtering has on our glitches. It is possible that pre-filtered glitches more closely resemble SGs than their filtered counterparts. Shourov Chatterji of Caltech has suggested that double-whitening may be used to filter our data without distorting the shape of the original glitch, but we have yet to test this. Additional approaches such as using a dictionary of filtered SGs may also be considered, if necessary. However, due to technical/compatibility issues, we have not been able to obtain a working whitening filter.

While our current testing has been very limited, we hope to soon have access to additional data from Ligo archives as well as working whitening filters. We also hope to

compare the results of our MP algorithm to that used by Shourov Chatterji et al. and the Omega Pipeline [6].

5. Ring Downs

The success of our attempts to extend our algorithm to a Ring Down dictionary has been mixed and depends largely on ones criteria for performance. RD dictionaries require more computational resources and have significant difficulties matching a couple of the parameters in particular, but nevertheless produce approximations of inputted signals which are comparable to or slightly better than the SG case.

Both the coarse and fine stages of the MP analysis over an RD dictionary requires significantly more computational resources than performing the MP using only SGs. On the coarse stage, this is evidently due to the larger parameter space. This has been improved through the implementation of a complex subdictionary and of look up tables and should now be manageable. However, the computer which I have been working on does not have sufficient RAM to run this code on the full subdictionary efficiently and I have had to restrict my subdictionary to included only RDs of about half the desired maximum width. We hope that a computer with slightly more RAM (about 2 gigabytes) will be available to us soon.

The problems with the fine stage analysis are related to the relationship between RD parameters. As has already been mentioned, the width of a RD depends both upon the parameters α and k. As a result of this, it is possible to adjust these two parameters together in a way which has only a minimal effect on the overall shape of the RD, especially for larger values of k ($k>1$.) Initial attempts to apply MP over a RD dictionary were therefore able to produce strong energy matches to simulated RDs, but with α and k both deviating from the true values by as much as a factor of 2 even in the absence of Gaussian noise. Furthermore, this search is very slow to converge on its final value.

Several modifications to the fine stage binary search algorithm have been attempted with limited success. Currently, while the error in α and k remain high, the correlation between the original RD and the approximation generally differ from one by less than 1 part in 10^4. This suggests that we are seeing fundamental difficulties with this type of correlation-optimizing approach.

Decomposition into RDs via this method may therefore be possible and computationally reasonable if one accepts significant uncertainty in these two parameters. This uncertainty will, however, have only a minimal impact on the strength of our approximation. Furthermore, restricting our dictionary to only elements with $k<1$ may solve this problem. This may be particularly appealing given that RDs with larger values of k correlate very strongly with SGs, and thus may offer little improvement over a SG dictionary. As an alternative, at some point in the future we may consider implementing another algorithm altogether (Prony), with the hopes that it will be more efficient and better able to resolve these parameters.

While a significant error was observed in the fitted parameters, the overall fit of the RD remains good. The correlations of fitted atoms and the rate of convergence to an inputted signal are generally comparable or slightly better than in the SG case. Graphs of this have not been included, since they do not differ visibly from those produced using a SG dictionary.

Interestingly, fitting a RD to an apparently SG glitch generally results in a match as good as or in some cases better than the best fit SG. Likewise, if one tries to match an (apparently) SG glitch using a combined SG and RD dictionary, the algorithm will predominately select only RD components.

The reason for this is unclear. However, if one attempts to match a RD to a simulated SG (without noise), it is actually possible to obtain a correlation of around .98 or better. The energy difference between a SG and a RD with a relatively high value for k (~2 or more) is therefore quite small. It is possible that the extra parameter may give a RD dictionary more flexibility to fit an imperfect SG signal, compensating for the small energy loss which results from switching dictionaries.

6. Conclusions

While significantly more testing will be required, MP algorithm seems to be generally successful at extracting SG components from noisy signals. While it is clear that the glitches available to us are not true SGs, it may be effective to approximate them in this manner. More work is needed to compare the performance of our algorithm using SG and RD dictionaries on real data.

We hope that in the very near future we will have more data available to us and will be able to begin more extensive testing of our algorithm and of the suitability of our chosen dictionaries for modeling glitches. We also hope soon to investigate the effect which whitening and high pass filters have, and the degree to which the process of filtering is responsible for distorting our data. Lastly, we may make further attempts to extend our analysis to also include a RD dictionary in a more efficient manner than the current implementation.

Once we feel confident in the performance of our MP algorithm on real data, we hope to begin the process of determining the parameter distributions of these glitches. We hope to determine the number of components generally needed to model a single glitch and the distributions of the parameters of these components. Once this is done, we hope to be able to accurately model glitch noise and that this model will be useful in designing GWB detection criteria.

Appendix 1: Charts and Graphs

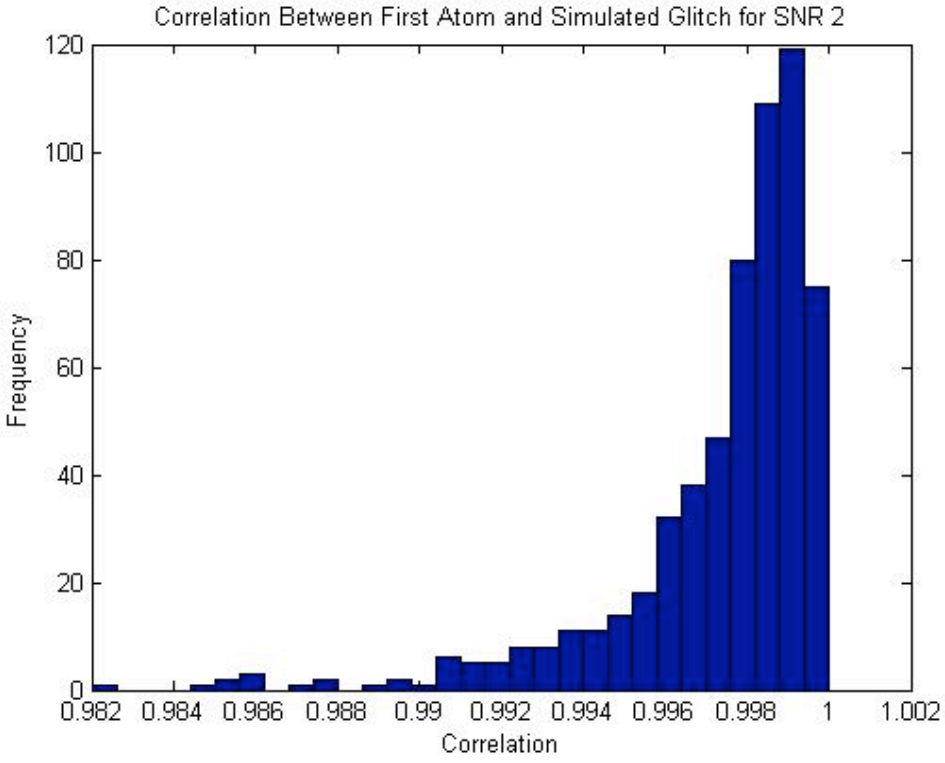

Fig. 1

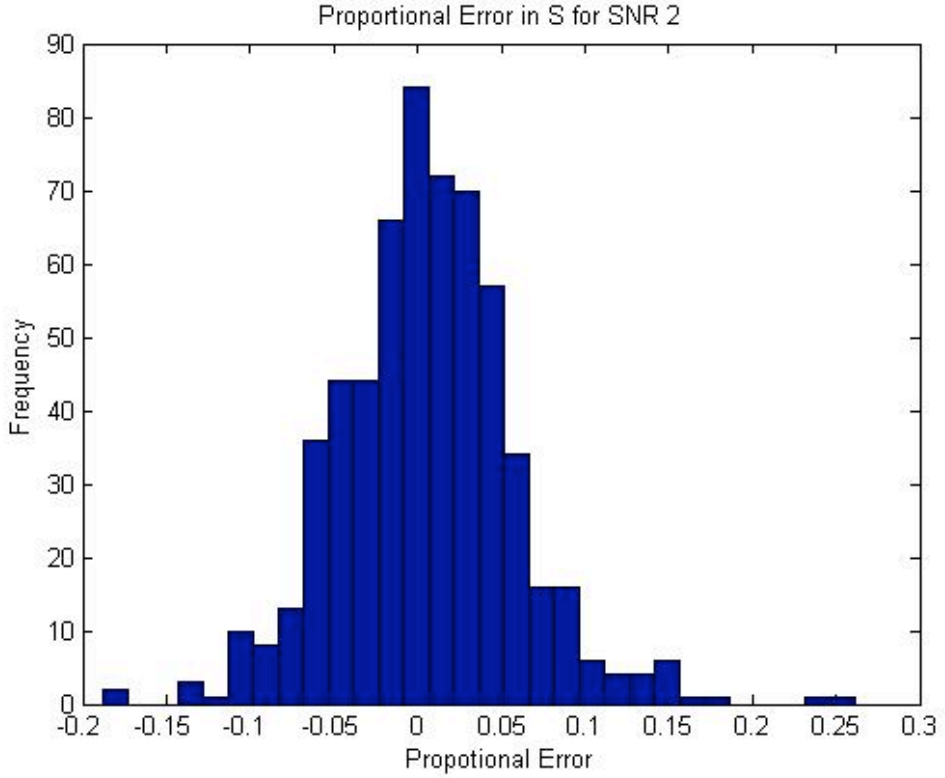

Fig. 2

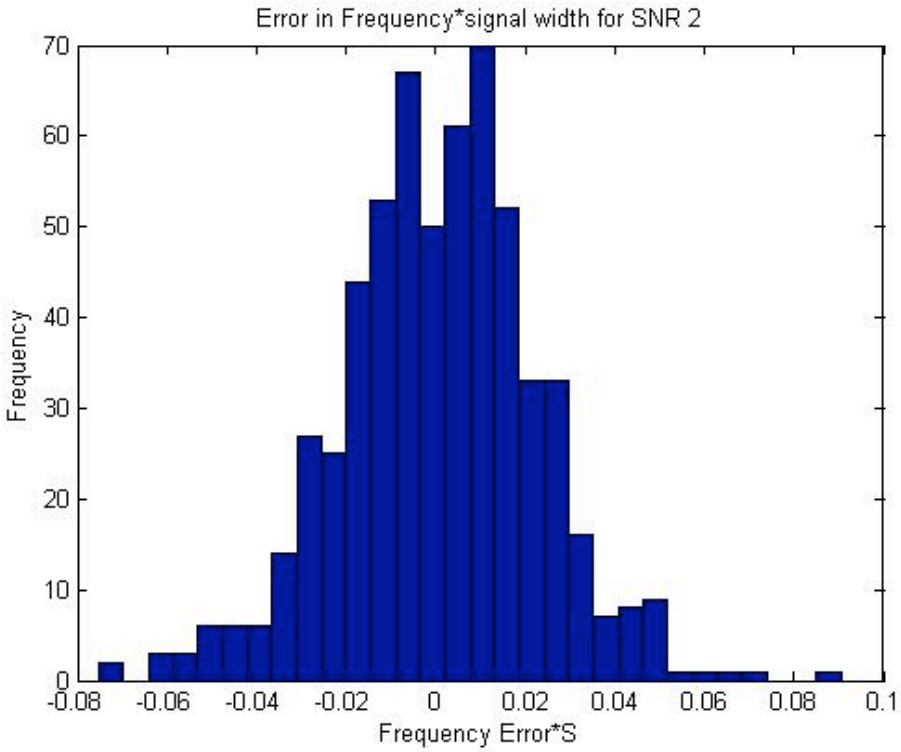

Fig. 3

*This is a measure of the error in the number of wavelengths contained in 1 standard deviation of our SG.

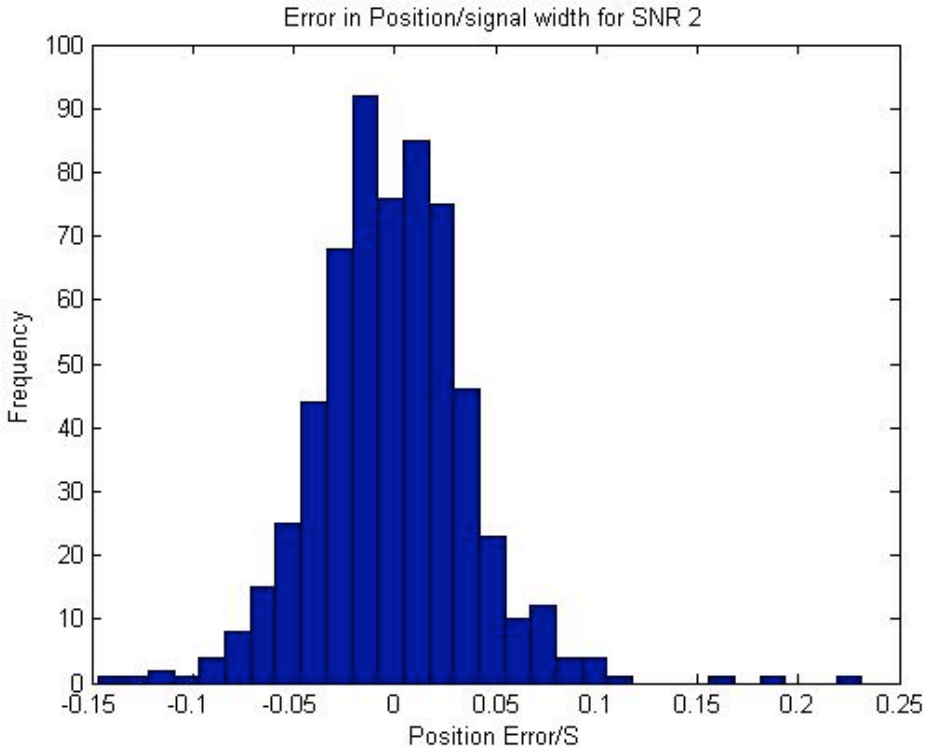

Fig. 4

The Time-Center error expressed as a number of standard deviations

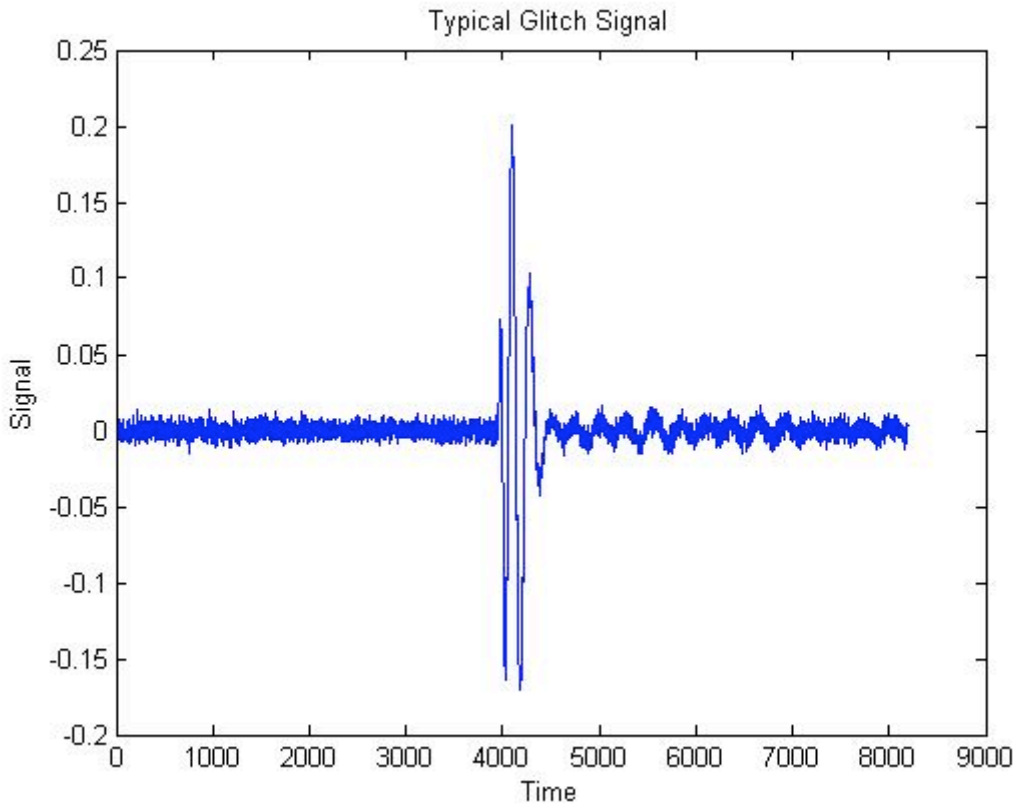

Fig. 5: Example Glitch

Width (S) (time-samples)	Frequency (ω) (Hz)	Time Center(u) (time-samples)	Correlation/ Norm	Energy/ Norm^2	Energy (% of estimated total glitch energy)
141.44	192.79	4137.0	1.8857	3.556	87.5%
120.78	325.92	4023.9	.51416	.264	6.5%
89.23	168.01	4339.4	.36757	.135	3.3%
141.93	140.53	4056.9	.18167	.033	0.8%
1434.00	119.85	5848.1	.18316	.034	0.8%
105.21	521.29	3992.6	.13952	.020	0.5%

Fig. 6

Parameters of first 6 found SG glitches using MP algorithm. The energy can be expressed as the Correlation^2 or equivalently as the Norm^2. The energy of the original signal is 4.178, including Gaussian background, with an estimated 4.063 belonging to the glitch. The energy of the residual is .136, only slightly larger than the estimated energy of the Gaussian background of .115. Note that the energy of the signal is the sum of the energies of each extracted atom along with the residual.

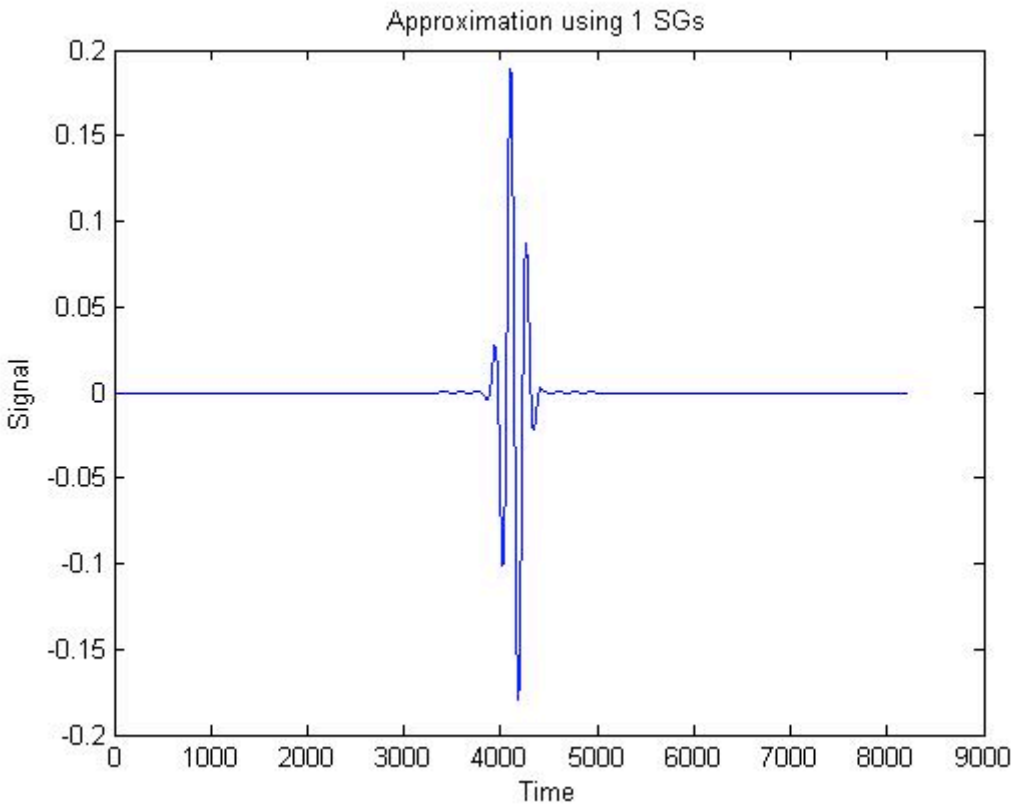

Fig. 7: Approximation Using a Single SG

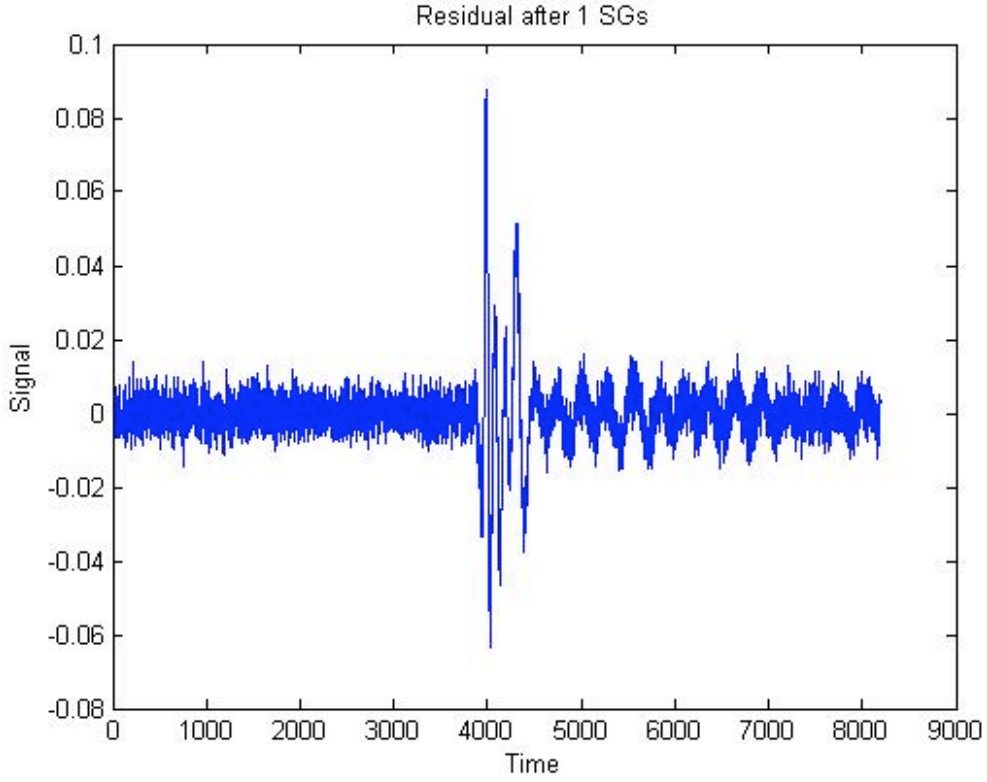

Fig. 8: Residual After Single SG Approximation

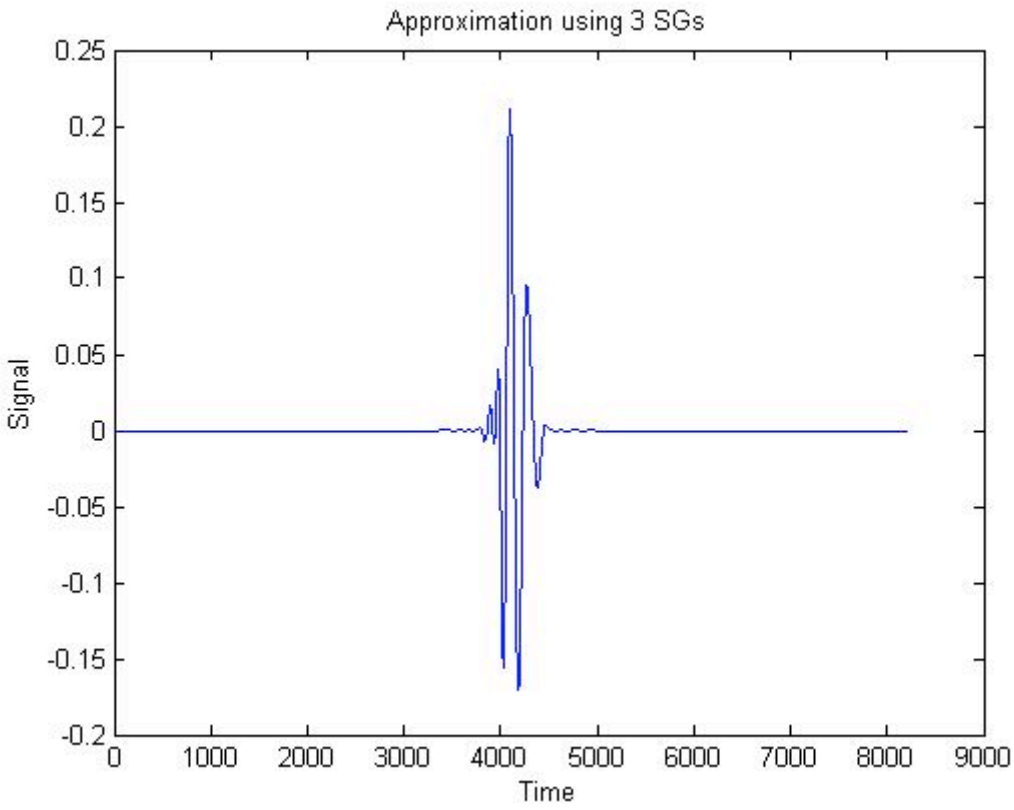

Fig. 9: Approximation Using Three SGs

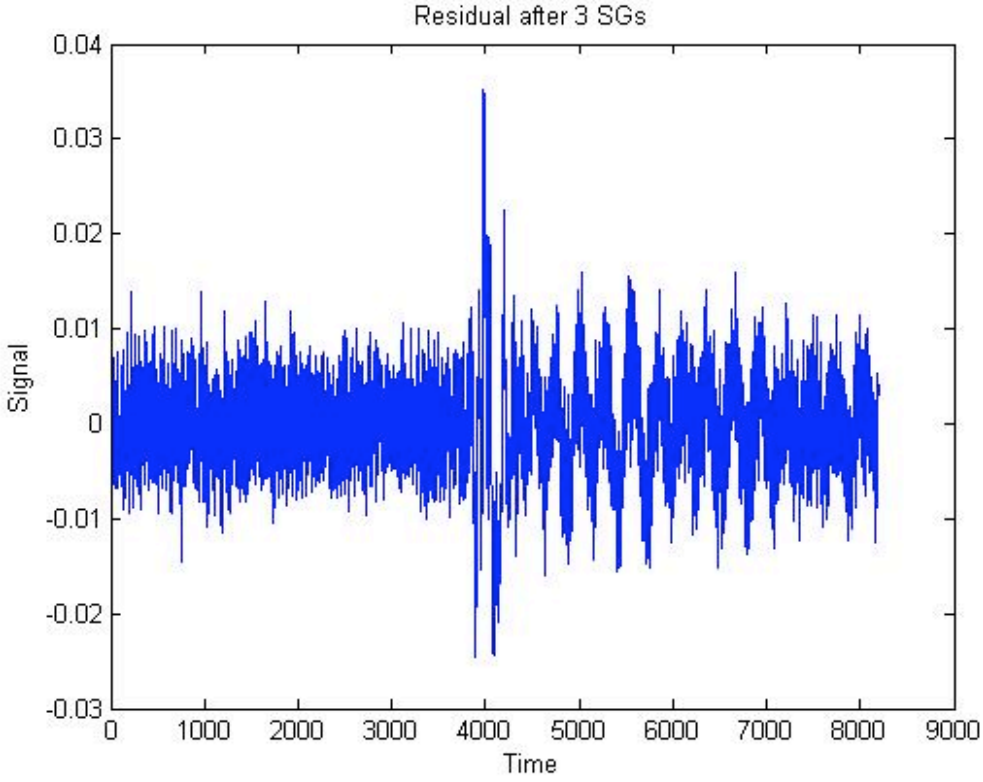

Fig. 10: Residual After 3 SG Approximation.

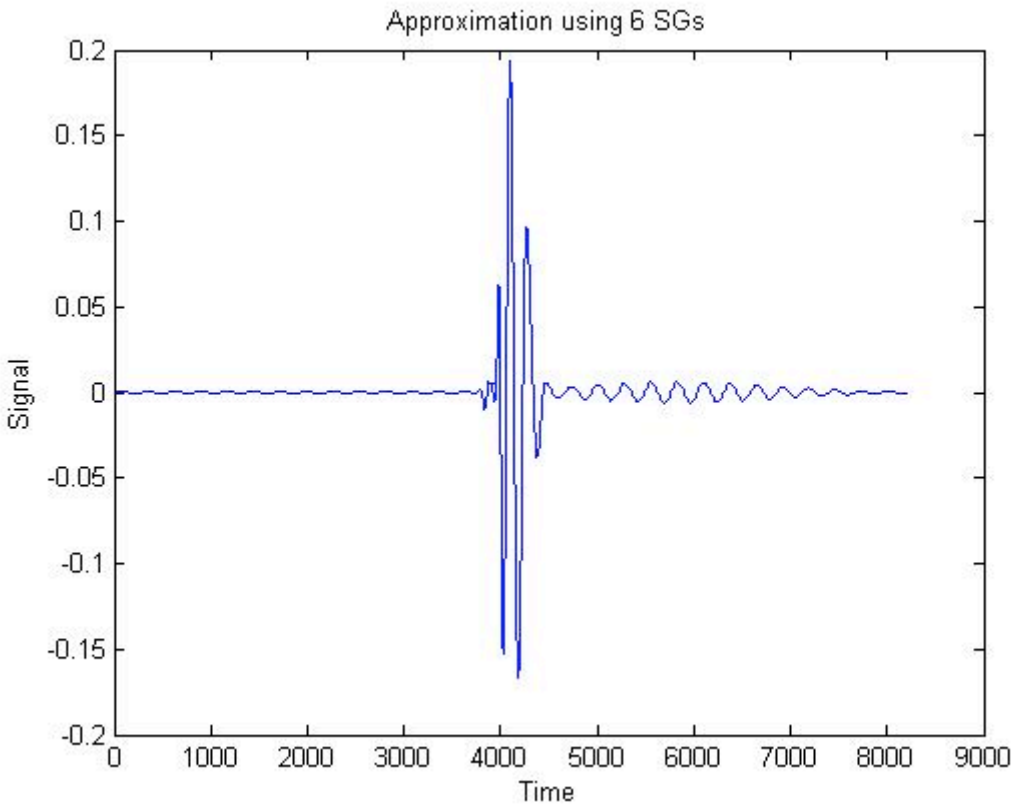

Fig. 11: 6 Atom SG Approximation

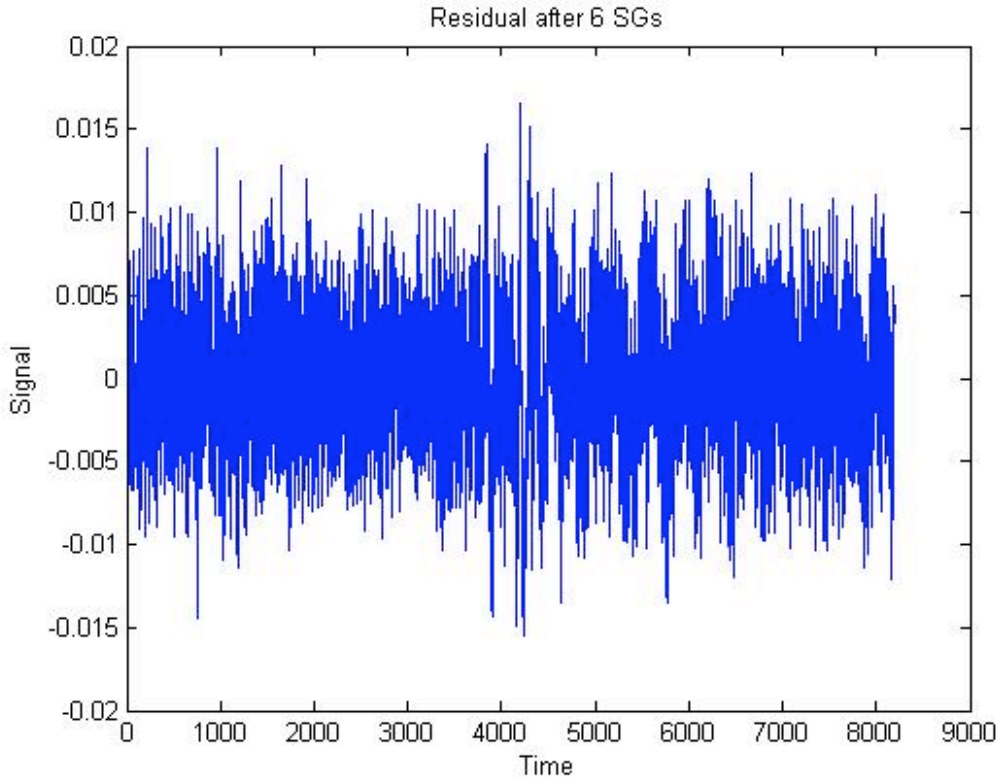

Fig. 12: Residual After 6 SG Approximation

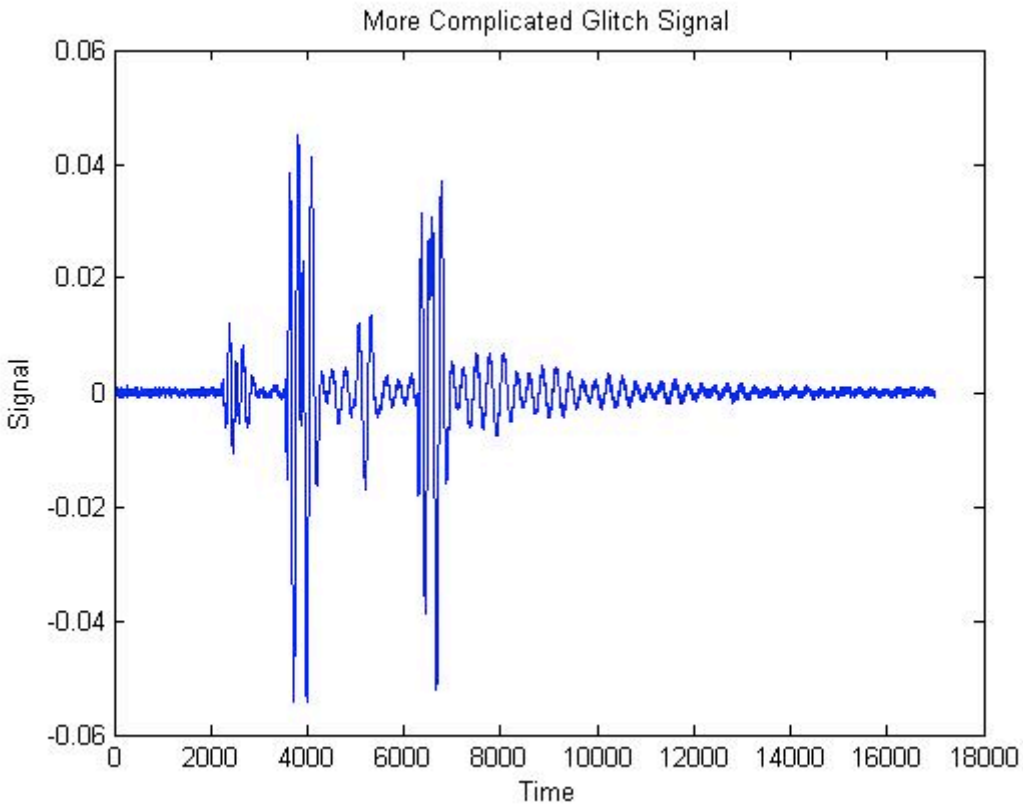

Fig. 13: More Complicated Glitch Signal

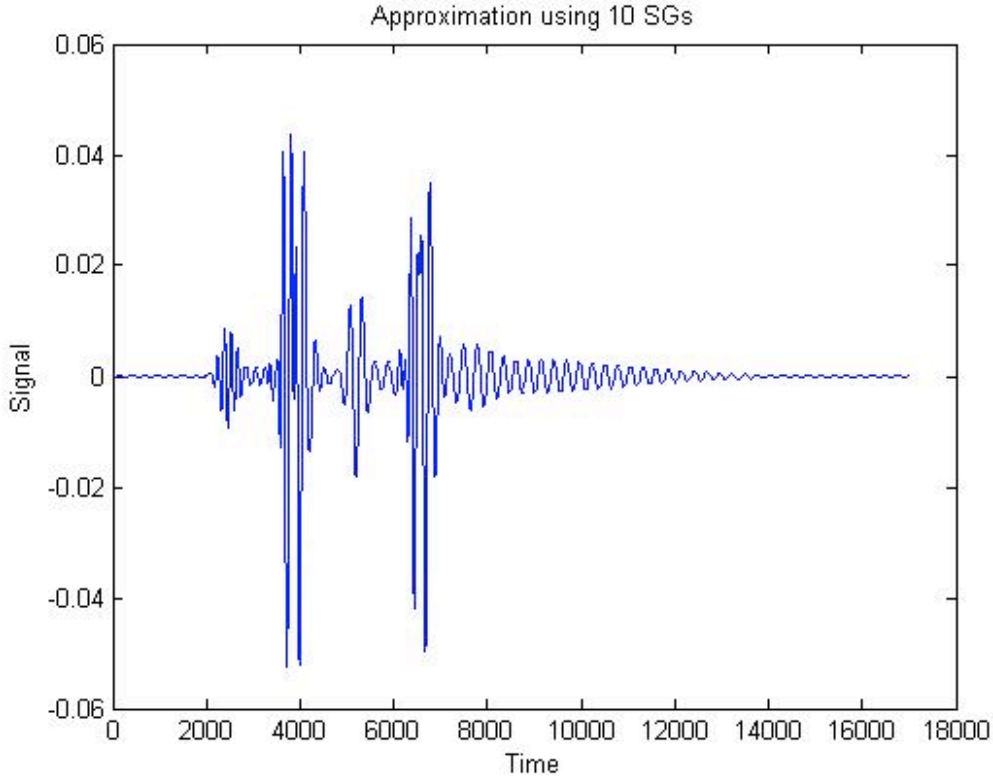

Fig. 14. 10 Atom SG Approximation

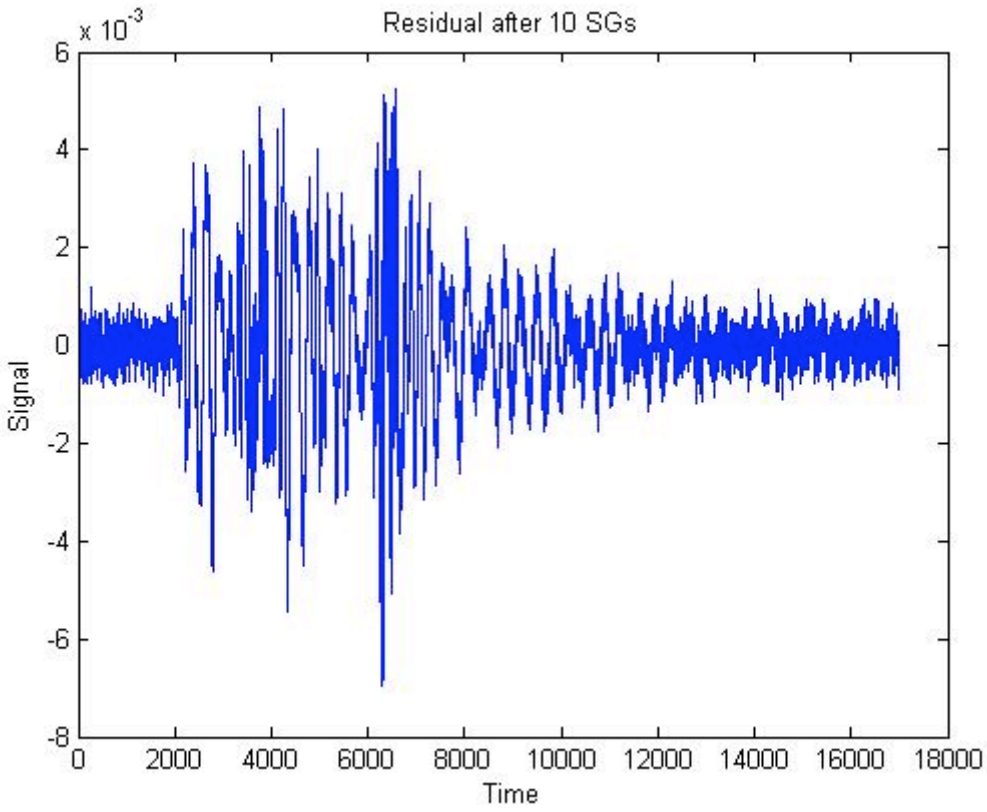

Fig. 15: Residual After 10 SG Approximation

Citations

[1] M Rakhmanov, and S Klimenko, 2005 "A cross-correlationmethod for burst searches with networks of misaligned gravitational-wave detectors" *Class. Quantum Grav.* **22** S1311

[2] M. Principe, and I. Pinto, 2008, "Modeling the impulsive noise component and its effect on the operation of a simple coherent network algorithm for detecting unmodeled gravitational wave bursts" *Class. Quantum Grav.* **25** 075013

[3] S.G. Mallat, and Z. Zhang, "Matching pursuit with time frequency dictionaries," IEEE Trans. Signal Processing, vol.41,no. 12,pp. 3397-3415, Dec. 1993.

[4] R. Coifman and V. Wickerhauser, "Entropy-based algorithms for best basis selection," IEEE Trans. Informar. Theory**,** vol. 38, Mar. 1992.

[5] Limited data archive available at
http://www.phy.syr.edu/research/relativity/ligo/restricted/glitch_catalog/

[6] Omega Pipeline project: http://geco.phys.columbia.edu/omega, run by S. Chatterji et al.

[7] S Chatterji et al. "Coherent network analysis technique for discriminating gravitational-wave bursts from instrumental noise," Physical Review D 74 082005 (2006).