

SignalLens: Focus+Context Applied to Electronic Time Series

Robert Kincaid




Figure 1. SignalLens view of an amplitude modulated electronic signal. The lens area exposes an anomalous glitch in the otherwise uniform oscillations. The tracks below the signal depict the location of filtered measurements and indicate a periodicity to the occurrences of similar glitches. See Section 5.1 for more details.

Abstract—Electronic test and measurement systems are becoming increasingly sophisticated in order to match the increased complexity and ultra-high speed of the devices under test. A key feature in many such instruments is a vastly increased capacity for storage of digital signals. Storage of 10^9 time points or more is now possible. At the same time, the typical screens on such measurement devices are relatively small. Therefore, these instruments can only render an extremely small fraction of the complete signal at any time. SignalLens uses a Focus+Context approach to provide a means of navigating to and inspecting low-level signal details in the context of the entire signal trace. This approach provides a compact visualization suitable for embedding into the small displays typically provided by electronic measurement instruments. We further augment this display with computed tracks which display time-aligned computed properties of the signal. By combining and filtering these computed tracks it is possible to easily and quickly find computationally detected features in the data which are often obscured by the visual compression required to render the large data sets on a small screen. Further, these tracks can be viewed in the context of the entire signal trace as well as visible high-level signal features. Several examples using real-world electronic measurement data are presented, which demonstrate typical use cases and the effectiveness of the design.

Index Terms—Focus+Context, Lens, Test and Measurement, Electronic Signal, Signal Processing.

1 INTRODUCTION

Modern electronic devices increasingly depend on complex systems of high-speed interconnected components. These systems require sophisticated test and measurement platforms (oscilloscopes, logic analyzers, spectrum analyzers, etc.) to aid in the design as well as debugging of such systems. Modern electronic designs often leverage features such as increased clock rates, high speed serial data interconnects and multi-core CPUs and GPUs. These trends push the limits of current measurement technology. One response is to provide increasingly larger storage capacity for test and measurement devices. This increased capacity enables longer signal traces to be captured and subsequently analyzed for behaviour over longer time periods. High-end devices now typically record up to a billion or more data points per trace. However, due partly to the necessity of convenient portability for such devices and partly for historical and traditional reasons, screen sizes have remained relatively small. For example a typical *large* oscilloscope display is on the order of only 1024x768 pixels, which includes both the signal

trace and the UI elements for control of the device. Based on these observations, we realized there was a critical need to better manage the visual display of very large signal traces that would enable a compact representation of the entire context of the trace as well as inspection of low-level details of interest.

To explore these issues we developed SignalLens, a research prototype for the visual analysis of large electronic time series. We use a Focus+Context approach to display the entire time series with simultaneous inspection of low-level signal details via a lens distortion in the region of interest. Unlike typical two-dimensional Focus+Context lens implementations, we are only interested in distorting the very long and visually problematic time axis.

Any attempt to display large data sets with a small set of pixels requires data overlay or pixel merging, which leads to loss of potentially important visual details. Recognizing this fact we also introduced a flexible ad hoc scheme for computing time-aligned properties of the signal and presenting the results as an aligned track below the original signal. These computed tracks can be further combined and/or filtered to reveal specific low-level features of the signal that might otherwise be hard to visualize on the small screen. These computed and filtered properties can be viewed in the context of the entire signal trace and used for navigating the lens to regions of potential interest. We will demonstrate that combining these tracks with the full trace and Focus+Context approach yields an effective means to quickly isolate and visually analyze subtle signal features within the constraints of typical instrument displays.

- Robert Kincaid is with Agilent Laboratories,
E-Mail: robert_kincaid@agilent.com.

Manuscript received 31 March 2010; accepted 1 August 2010; posted online 24 October 2010; mailed on 16 October 2010.

For information on obtaining reprints of this article, please send email to: tvcg@computer.org.

This paper describes the following contributions:

- A high-magnification lens-based visualization for viewing large electronic time series.
- Extension of the signal context via computed measurement tracks to aid feature finding and navigation.
- Demonstration of the design by analyzing real electronic signals.
- Observations on benefits and limitations of a lens-based approach for viewing such time series.

2 PROBLEM DESCRIPTION

The specific problem we are trying to address is to design a compact visualization that provides both global context and local details for long duration signal traces. In addition to the specific analysis tasks we wish to support, we must also consider aspects of the underlying data characteristics as well as the target users of the prototype solution.

2.1 Target Users

During the early information gathering phase of this project we noted an interesting characteristic of our target users. Most of the visual interaction between user and instrument is based on ad hoc real-time visualizations of live electronic signals. The typical strategy is to manipulate the visualization (e.g. an oscilloscope trace) to reveal salient features in the signal. As such, electronic test and measurement can be viewed as a largely dynamic ad hoc information visualization task with the same requirements for a high degree of interactivity and the same dependency on leveraging pre-attentive human visual perception to detect relevant patterns, trends and outliers in the data. These users are also highly visually oriented in their analysis tasks and therefore (in our opinion) likely to accept information visualization approaches. We note that this target audience is largely untapped by the information visualization community and represents a valuable domain for further information visualization research.

Paradoxically, at the same time most test equipment as well as users are extremely tied to legacy designs and paradigms, even when they are no longer physically required. For example, modern digital oscilloscopes still maintain displays that mimic the fixed etched reticules found on cathode-ray displays used in early oscilloscopes as well as their physical knob-based controls. This observation leads to the added requirement that whatever solution we provide must try to be minimally disruptive to the pervasive culture and expectation of our users. So while we believe users of test and measurement equipment would be receptive to information visualization approaches, designs must be sensitive to the domain and its historical prejudices.

2.2 Data

Modern test and measurement systems generally sample voltages at extremely precise fixed intervals. For example, high-end oscilloscopes have sample rates as high as 40GHz or more, but this sample interval can be adjusted lower depending on the source being analyzed. At high sample rates, even a Gigapoint trace (10^9 points) may only involve a small fraction of a second. Such traces can be thought of as simply time series data. The form in which such data is usually represented on the measurement device can be thought of as viewing a small portion of a much larger smoothed or interpolated line graph of the voltage over time. Navigation is typically accomplished with zoom-and-pan techniques.

Aigner et al. [1] explored taxonomies for time series data and visualizations. Using their terminology the temporal primitive consists of equally spaced time points and the time structure is linear. The data measured is abstract and univariate (usually voltage).

2.3 Target Analysis Tasks

There are a multitude of analysis tasks performed by electrical engineers testing and debugging electronic circuits. It is beyond the scope of this paper to discuss all of these in detail. However, the goal of SignalLens is to address the specific issue of analyzing very large signal traces. Thus, we focused our attention primarily on finding specific signal features or anomalies buried within the complexity of a large sample. This task is generally a hard problem to solve with existing instrument software, and not easily subject to standard ad hoc techniques. We can break this down into two main tasks:

Anomaly Detection: which may consist of simple, obvious outliers or more subtle cases of “a pattern surprising if the frequency of its occurrence differs substantially from that expected by chance” (Keogh et al. [13]). In many cases, *extreme* outliers may be detectable visually. More subtle cases often require computational and/or statistical approaches.

Motif Discovery: which involves searching a time series for a specific time-dependent pattern (Lin et al. [20]). Typically, the search will begin with a specific pattern of interest. In the case of electronic signals this pattern or waveform may have been initially found during a preceding anomaly detection task.

These signal processing tasks are generally based on computing appropriate signal measurements and filtering by ad hoc criteria. Existing instruments typically require the tedious use of pan-and-zoom interfaces to navigate to these measurements of interest. Instrument analysis software does not typically support visualizing multiple measurements and the signal in a global context while at the same time viewing low level signal details.

3 RELATED WORK

In designing SignalLens we leveraged several key topics from information visualization research:

3.1 Context-based Visualization Systems

There exists a well known body of research in context-based visualization frameworks. Two common approaches are Overview+Detail (O+D) [25] and Focus+Context (F+C) [7]. Each approach attempts to provide a contextual overview simultaneously with a detailed view for inspecting details of interest. Several user studies have provided evidence that providing an overall context along with details of interest provides a more effective analytical framework. Whether O+D is more effective than F+C (or vice versa) is an area of some ambiguity in the current literature and may be somewhat task specific [5]. For example, Gutwin and Skopik [11] found distortion lenses best for steering tasks while Pietriga et al. [24] found zoomable O+D interfaces more efficient for multi-scale search tasks. Evidence generally seems to indicate that for many tasks both are similarly effective with slightly different tradeoffs. Cockburn et al. [5] provide a recent review of this topic including a discussion of relevant user studies.

However, the high magnifications we will typically require for electronic signals make Overview+Detail somewhat problematic. For useful navigation beyond 100X, one or more intermediate views are likely to be required [25]. For example, assuming the user could accurately select a single pixel on an overview trace of 10^6 points on a display 10^3 pixels wide, the selection of a single pixel would correspond to a range on the order of 10^3 points. Scaling to larger time series the problem becomes even more problematic. This inability to select precise ranges in the overview makes navigation with high magnification difficult in a scenario consisting of an overview and a single detail view. Further, any selection indicators in the overview will be of such a tiny size (~1 pixel wide) that they become difficult to select and manipulate easily. In contrast, while the initial context selection is still problematic in our Focus+Context design, we will demonstrate examples where the lens actually facilitates rapid navigation to the region of interest.

3.2 Lens Frameworks

A distortion lens is a frequent design pattern in Focus+Context implementations and traces its origins to Furnas' Fisheye Views [6]. Several approaches to lens construction have been described [3, 18, 19]. In this work we follow the mathematical formalism for lens distortions provided by Carpendale and Montagnese's Elastic Presentation Space [3]. This framework provides for a variety of lens shapes. A further discussion of how this framework is used within SignalLens is described in more detail in section 4.1. Carpendale et al. have also studied the issues associated with high magnification [4] and those findings are also leveraged in this work.

Recently Pietriga and Appert have published work on Sigma Lenses [23] based in part on earlier work by Gutwin [10]. Sigma Lenses replace standard optical distortions with dynamic translucence to transition between focus and context areas. While for some use cases Sigma Lenses have been shown to be more effective than traditional lens implementations, they have not been tested under conditions of extreme magnification. Further, Sigma Lenses go beyond the simple mental metaphor of an optical magnification lens. Since we desired to keep our implementation intuitive and as close to a standard instrument display as possible (for the reasons described in Section 2.1) we chose a more traditional distortion-based implementation.

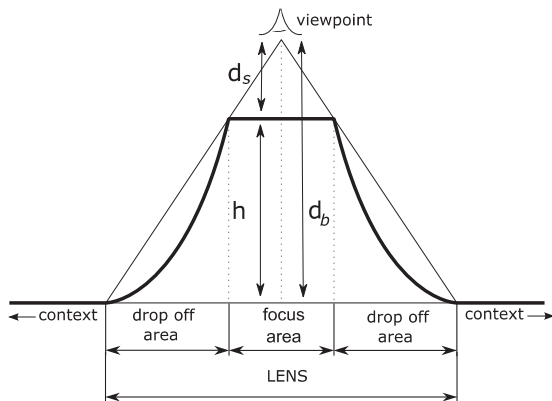


Figure 2. Schematic of lens parameters. The heavy line denotes the hypothetical viewing surface.

3.3 Time Series Visualizations and Analytics

A large body of work exists for visual analysis of time series. Aigner et al. [1] take a systematic view of the field and provide a framework for categorizing various approaches.

LiveRAC [22] provides semantic zooming of a reorderable matrix of system management time series. Time-series bitmaps [17] uses a bitmap encoding to efficiently represent a large collection of time series. Line Graph Explorer [16] supported visual analysis of large collections of time series, but was not designed to support very large time series.

TimeSearcher [2, 12] uses an Overview+Detail approach for querying and visualizing features in time series data. TimeSearcher is perhaps most relevant, since it works directly with line graphs of time series data and provides facilities for pattern matching similarly shaped motifs. However, it is primarily designed toward motif searching versus a general purpose signal analysis tool. Also as mentioned previously, our screen and magnification requirements make Overview+Detail a poor choice for our design.

Lopez-Hernandez et al. [21] recently extended oscilloscope eye diagrams using an interactive 2.5D layer-based technique.

4 DESIGN

The dominant feature of SignalLens is obviously the signal trace, which consists of a relatively straightforward two-dimensional line chart of the signal. No lens distortion is apparent until the user clicks

the mouse at a region of interest. Our choice to take a Focus+Context approach is driven primarily by issues of limited screen real estate combined with the requirement for high magnification. For prototyping purposes SignalLens was developed as a standalone application. However, our intent was to design an effective visualization that could function embedded within the constrained dimensions of an instrument display.

4.1 Lens

Our lens design follows the distance-based model described in Carpendale and Montagnese's framework for Elastic Presentation Space (EPS) [3]. Due to the high magnification typically required, we also leverage further work by Carpendale et al. that is specifically related to high magnification [4]. Readers should refer to these references and related EPS work for complete details on computing lens distortions. We review briefly the features of this framework that are relevant to our discussion of SignalLens. Figure 2 defines the relevant variables used in the distance-based (vs. optical) magnification scheme. From the viewpoint looking downward at the presentation surface we define d_s as the distance to the surface and d_b as the distance to the base. For the region of interest we, in effect, stretch the viewing surface toward the viewpoint where d_s is much smaller than d_b resulting in a magnification of d_b/d_s . In the surrounding context area, d_s is identical to d_b and no magnification is achieved. Between the region of interest and the unmagnified context area we provide a smooth height function h corresponding to d_b-d_s . Equation 1 summarizes the relationship between these distances and the magnification m .

$$m = \frac{d_b}{d_s} = \frac{d_b}{d_b - h} \quad (1)$$

To enable continuously variable magnification between the region of interest and the unmagnified context area, we allow h to be a variable function dependent on the distance away from the region of interest. This height function is also sometimes called a "drop off" function. Further, we define our height functions based on a standard lens radius of 1. This simplifies the computations and the resulting displacements can simply be multiplied by the lens radius to get the final displacement.

SignalLens implements a variety of lens functions for comparison (linear, cubic, quadratic, hyperbolic, spherical and Gaussian). As was found previously [4], high magnification generally requires "concave" drop-off functions to avoid occlusion. For the signal traces we have analyzed we have found the quadratic height function (for a unit radius):

$$h(t) = (1 - t)^2 \quad (2)$$

to be the most practical since it gives a useful distortion field and the transition to the context area seems relatively natural. It is this lens shape that we use for all figures presented in this paper.

A focus area of fixed magnification is provided in the center of the lens corresponding to the region of interest. We found this design advantageous for accurately interpreting signal details as any lens distortion also distorts the signal shape, making accurate analytical comparisons or interpretation difficult. As a further aid we allow an optional grid or reticle to be displayed, so that the user can clearly see the boundary where the non-linear distortion begins. Another optional bounding box is provided to indicate the full extent of the lens. This feature is often useful in training or initial usage, but once users become familiar with the lens concept they generally turn this feature off. We also restrict the focus area to a fixed percentage of the plot size as this restriction was previously found to effectively avoid visual occlusion at the transition between the focus and drop-off areas when using high magnification [4].

Settings are also provided for the magnification level, the size of the lens (as a percentage of the full signal trace) and the lens shape (by selecting the corresponding drop-off function).

4.2 Measurement Tracks

For large data sets the visual density of the display can obscure fine details and make it difficult to find interesting visual features. To address this problem we chose to use computational assistance in finding features of interest. The user can choose from an extensible list of time-aligned computations based on the original signal. Following the terminology used with test and measurement instruments we refer to these computations as “measurements” and thus we are creating “measurement tracks”. In effect we are making post-acquisition measurements of the already acquired signal. Such measurements are common features of devices such as oscilloscopes and logic analyzers, so our measurement tracks also provide a mechanism for incorporating these computations into the Focus+Context design.

Tracks can be filtered with the dynamic query sliders shown in Figure 3D. The two black “carets” can be dragged independently with the mouse to set minimum and maximum values. This setting specifies the range of data being included by the filter and is indicated by the green coloring between the carets. The red regions are values excluded by the filter. Additionally, the user can type the values directly into the flanking text boxes. Text entry facilitates entering exact values that would be tedious to select via mouse drag. Alternatively the user can click the histogram button (Figure 3F) to invoke a separate histogram window as shown in Figure 5E. As the user zooms or pans the histogram, the visible region of the time axis is used to filter the display and only those time points with measurement values passing the user-set filter are displayed.

Color is used to distinguish between tracks enabling the user to visually connect the track with the similarly colored query slider. We originally tried to encode the measurement value in a manner similar to Line Graph Explorer [16] by mapping a color gradient to represent the measurement value. However, the density of the resulting data was such that this method was seldom of any analytic value. The brightness and frequency of large values tended to overwhelm and dominate the unmagnified display while smaller values, even when filtered suffered from visibility issues. We chose instead to not encode the measurement value and merely use the tracks as binary markers indicating where there are filter-passing values. In practice this implementation seems a simpler and more tractable approach, since the task is generally one of trying to isolate a relatively small number of outlier features via filtering, and navigate to those, versus an attempt to view the entire population of measurement values.

This paper is primarily concerned with the visualization and interaction design of SignalLens and a full description of all measurement algorithms is beyond the scope of this paper. Briefly summarized, the currently supported measurements include:

- Savitsky-Golay filtering [26] to produce smoothed 1st or 2nd derivatives of the signal.
- A simple signal motif finder based on computing a Pearson correlation between a target waveform and each matching window of the signal [15].
- Rise time and fall time (duration of leading and trailing edges of pulsed signals).

Measurements based on other measurements are also supported such as:

- Logical and/or between tracks. This operation takes two tracks and generates a third based on the union (or) or intersection (and) of measurements with corresponding time intervals.
- Distance between measurements in a single track. For example, we can compute the distance between the occurrences of motif matches.

- Distance between measurements in two tracks. For example one can compute rise time in one track, fall time in a second and pulse width in a third track by computing the distance between rise time and fall time of the previous two tracks.

The intent is to provide a limited set of basic operations (e.g. rise time/fall time) and produce more complex measurements (e.g. pulse width) as combinations of these basic primitives.

4.3 Additional Features

The “goto” button allows quick navigation to a specific time point and the “clr” button merely clears the display to the state when the data was first opened, thus closing the lens. The user can also select to “animate” navigation. In this case any operation that navigates to a new time point will be animated from the old time point to the new one. We observed that for dense signal plots with very repetitive signal pulses there were minimal visual cues that navigation was taking place. However, by animating intermediate steps through the time dimension, one can see the cursor actually moving in time and the movement forward or backward is much more apparent and natural to the user. A scratchpad area (Figure 3G) allows the user to save and/or overlay the contents of the focus area for later reference or comparison.

4.4 Evaluation by Experts

While no formal user studies have yet been performed, there have been a number of informal as well as formal reviews of the working prototype by experts in test and measurement devices. These reviews included instrument hardware as well as instrument software designers and also field engineers directly supporting users of test and measurement instruments. Experts were also provided access to the software for their own trials. The response from these interactions was generally positive. Review of the first prototype raised several issues which led to further refinement of the prototype. These subsequent design modifications were subject to further expert review. Two key issues raised were:

There was generally an initial aversion to distorting the time axis due to a cultural bias that viewed the visual time increment as inviolate. This finding led us to implement the fixed magnification focus area which makes the focus area seem more consistent with previous user expectation, even though the drop-off areas still maintain distortion. The lens grid also reinforces the correspondence to a traditional reticule found on real instruments.

Initial reviews also revealed the gaps in the current methods of measurement visualization, which led to the design of the current track and filtering mechanisms in order to support a more efficient and visual means for manipulating and navigating to salient measurement features in the context of the full signal.

5 EXAMPLE RESULTS

In this section we present several example analyses to illustrate how SignalLens is effective in visually analyzing large signal traces.

5.1 Amplitude Modulated Signal

Figure 3 displays an amplitude modulated (AM) signal from an oscilloscope demonstration board used for training and sales purposes. This signal consists of a relatively modest 65,534 data points. The carrier frequency of this signal is 2 MHz with a modulating frequency of approximately 440 kHz. However, the amplitude of every 2nd envelope has slightly different amplitude when compared to the previous envelope due to an intentional discontinuity in the carrier at approximate repetitive rate of 220 kHz. This discontinuity is synchronous with the overall modulating frequency, not the carrier signal, and hence each occurrence of the discontinuity will appear differently depending upon the phase relationship of the modulating and the carrier signals. Therefore, this

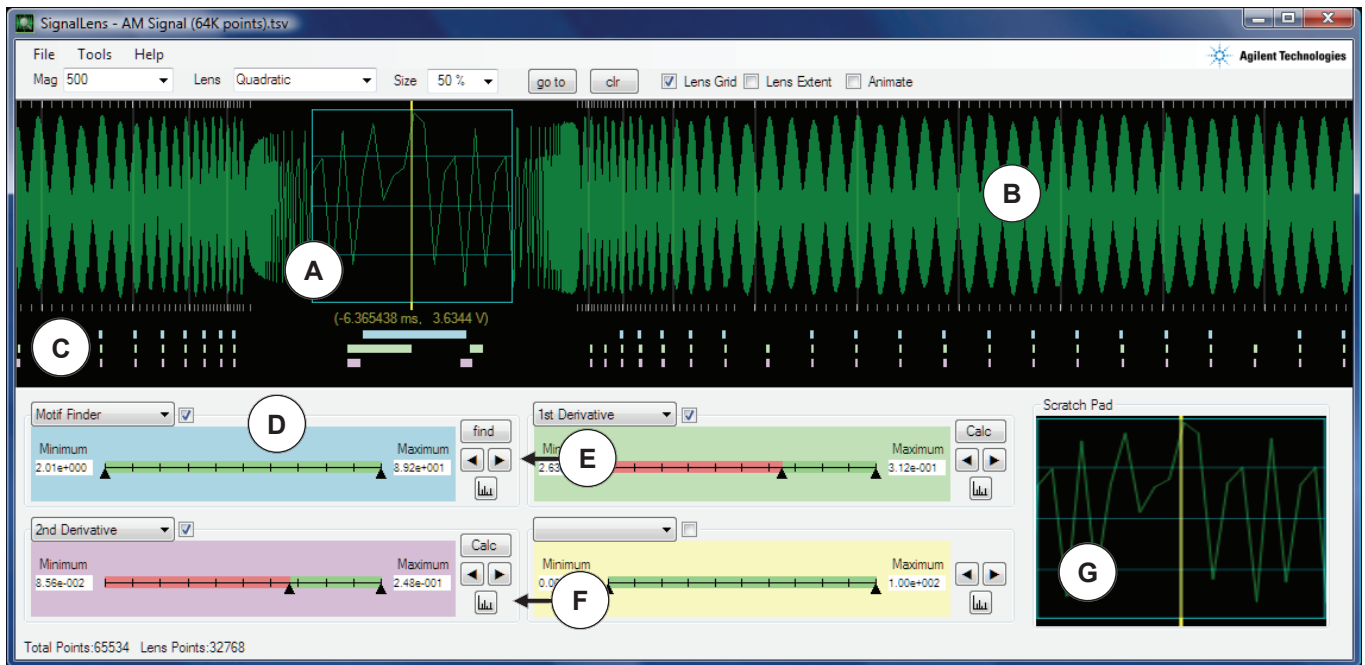


Figure 3. SignalLens User Interface. This is a view of an AM test signal with a lens magnification of 500X. (A) Fixed magnification focus area bounded by lens grid, showing a signal “glitch”. (B) Undistorted context area showing global scope oscillations. (C) Filtered measurement tracks (see text for details). (D) Measurement panel where measurements can be selected and filtered using the dynamic query sliders. (E) Forward/Backward buttons for directly navigating to the next or previous measurement which passes the filter criteria. (F) Histogram button for invoking histogram used for filtering. (G) Scratch Pad for storing/overlying waveforms of interest.

signal represents a useful test case where we know there is a repetitive “glitch” within the signal and we can attempt to find these anomalies using SignalLens.

We begin by observing that the overall signal appears to be quite uniform except for a few instances where a slight “whisker” appears to protrude from the otherwise uniform shape. Figure 4 shows an example of such a whisker and the sequence of navigation events that narrow into the corresponding signal glitch. Figure 3A shows the lens positioned over this feature. We can use the signal motif finder to look for a similar wave shape as this glitch in Figure 3A. After filtering for only statistically significant matches we plot the top track show in Figure 3C. At this stage we notice that the repetition of the glitch generally appears every other carrier oscillation except for a few instances which seem to skip this pattern. We can investigate further by examining 1st and 2nd derivatives. Filtering for just the most extreme slopes and curvatures yields the additional two tracks shown in Figure 3C. In this case we notice that the pattern is uniform and inspection of the indicated locations show that even when the motif finder missed a glitch, the filtered derivatives indicate its location.

If this exercise had been on a completely unknown signal, it

would be reasonable to follow this work flow by first noticing the obvious anomalies and then try to pattern match against an example waveform as we have done here. Noticing that there is an apparent pattern, we can perform further calculations to dig deeper into the signal and find the remaining glitches. From a task perspective we can describe this process as (1) a classic outlier detection followed by (2) motif searching and finally (3) detection of a more subtle anomaly consisting of a surprising frequency of occurrence of the motif (cf. Keogh et al. [13]).

5.2 PCI Express

Peripheral Component Interconnect Express (PCI Express or PCI-E) is a widely used standard for computer expansion card interfaces as well as motherboard-level interconnects. It is a high-speed point-to-point protocol with speeds up to 1GB per second per lane. Motherboard designers frequently need to measure these high speed signals to establish compliance with PCI-E specifications as well as debug transient signal issues.

Figure 5 shows a sample of PCI-E data generated from PCI-E demo board. In this case the overview context is relatively featureless and provides no real clues as to where any anomalies




Figure 4. An example navigation sequence using the AM signal. (A) We initially find an anomalous “whisker” protruding from the standard waveform (inset shows magnified view for clarity). (B) After clicking the mouse on the “whisker” we are nearer the anomaly, which can now be seen more clearly in the drop-off area of the lens. Clicking on this larger target navigates precisely to (C) the point of interest. The arrow indicates the same target of interest in all views.

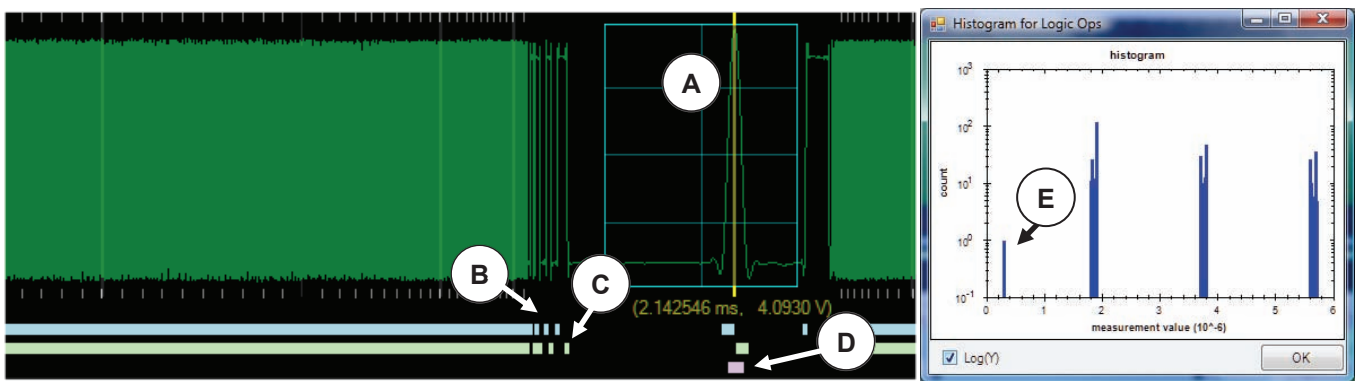


Figure 5. PCI-E Signal with spurious signal spike. (A) The non-compliant signal spike shown at magnification 250X. (B) Rise time track. (C) Fall time track. (D) The anomalous pulse width which was computed as the difference between rise time and fall time, and then filtered for any anomalous values. (E) Histogram view of the anomalous pulse width used for filtering the visualization.

might lie. We could manually scroll through the entire data set visually looking for anything suspicious, but this procedure would be quite time consuming and tedious. Instead, we can use the measurement tracks to aid us in discovering any hidden features.

We begin by calculating rise and fall times of the signal pulses. These calculations provide us the location of the leading and trailing edges of the PCI-E pulses. We calculate a third track from the first by calculating the distance from the rise to fall time, essentially giving us the pulse width. If we then view the histogram of pulse widths (Figure 5E) we notice they mostly follow the expected equally spaced distribution of PCI-E widths corresponding to the different width multiples in the specification. However, we notice that there is a very narrow outlier peak. Using the histogram, we can filter for just this peak. Navigating to the single outlier, we find an extremely sharp spike at 2.1426 ms as shown in Figure 5A.

Users of this demo board had never seen this anomaly before due to the single instance of a very narrow spike. However, SignalLens found the spike quite easily with just a few simple operations. In this case the outlying pattern is actually discovered in the histogram of pulse widths and the visualization is used to navigate to and view this anomaly in context.

5.3 ECG

Not all electronic time series are measured from semiconductor based systems. Physiological measurements such as electrocardiograms (ECG) are particularly well suited for analysis by SignalLens. Figure 6 shows one channel of an ECG spanning four continuous hours of heart monitoring. This data of a 53 year old male is taken from dataset chf01 from the BIDMC Congestive Heart Failure Database obtained from PhysioBank [8]. As in the previous AM signal, we navigate to a suspiciously protruding feature to find that it looks different from the typical repetitive waveform. Referring to the literature we find that the wave shape corresponds to a

premature ventricular contraction [9]. Using the motif finder we can locate a number of similar features in the data set, and can quickly navigate to each one using the next/previous buttons (Figure 3E).

Another aspect we can quickly recover from the ECG is a rough approximation of the heart rate. If we calculate the rise time, we capture the leading edge of most heart beats. If we then calculate the distance between these edges and inspect the corresponding histogram (not shown) we find that the distribution of time between heartbeats centers around 0.9 seconds corresponding to a heart rate of 67 beats per second which seems plausible. Similar distance measures could be used to characterize the distribution of time spans between PVC's or any other anomalies of interest.

While considerable work would be required for SignalLens to be an FDA-approved tool for ECG analysis, the results with this data set are encouraging that the general approach has broader potential application to other time series domains. This example also demonstrates that the methods are relatively robust against signals which are less uniform and well behaved than those generated by well designed electronics.

6 DISCUSSION

During the development of SignalLens several important issues arose. We discuss these here with the hope that these insights may be of value to others interested in implementing similar designs.

6.1 Magnification, Context and Levels of Detail

One might think that extreme magnifications of 500x or more would be impractical or at least problematic. However, in this paper we have demonstrated real (not hypothetical) examples where such magnifications are clearly tractable and useful. The AM signal example in Section 5.1 shows how even the drop-off regions of the lens can be very useful for iteratively navigating to an otherwise difficult-to-find feature (see Figure 4).




Figure 6. A four hour ECG Signal with a lens magnification of 1000X. (A) An instance of a premature ventricular contraction (PVC). (B) Other PVC's detected by the motif finder and shown in the track display. (C) Inset illustrating how the scratch pad can overlay two or more waveforms, here showing a comparison between two PVC signals.




Figure 7. NTSC Video Signal with a lens magnification of 500X. Multiple levels of signal detail are visible: (A) vertical sync (B) horizontal sync (C) color burst (D) single line of video. The motif finder has been used to find the rising edge of each video scan line.

We have also seen that some signals (e.g. the PCI-E signal in Section 5.2) can have a relatively featureless uniform context area. However, even for those instances where the context area is relatively uninformative, we can still rely on the measurement tracks to provide assistance in detecting and navigating to interesting features or outliers in the data. The temporal position and spacing of detected features is shown in relation to the global context and impart useful contextual information even when the signal itself is relatively featureless. Put another way, since the measurement tracks are subject to the same Focus+Context techniques, they provide supplemental contextual information at the overview level, which aids in navigation and feature finding in the same way that visible signal features would.

Finally, we wish to show one final striking example of how rich signal details can be found in large signal traces, even when viewed in their entirety. Figure 7 shows slightly more than a full frame of an NTSC video signal [14]. What is striking about this figure is that one can see rich signal detail at *all* levels. One can clearly see two flanking vertical sync regions (Figure 7A) in the context area. As expected, one can see clear details in the focus region of the lens (here the color burst signal, Figure 7C). What is perhaps not expected is that the drop off areas (because of their intermediate magnification ranges) still enable significant interpretation of the signals flanking the focus region. Here we can clearly see the horizontal sync area (Figure 7B) as well as the beginning of the scan line signal (Figure 7D). While some signals such as the PCI-E data in Section 5.2 are so uniform that little useful detail is available in the context area, there are examples such as the NTSC signal where significant detail at multiple levels of magnification is discernable. In these cases the focus, drop-off and context regions work together to provide useful landmarks for navigation and feature selection.

We speculate that many modern *analog* electronic signals will be similarly amenable to Focus+Context techniques, since they typically require multiple levels of modulation and/or multiplexing leading to a similar degree of multi-level detail. In contrast, purely *digital* signals such as the PCI-E signal have relatively uniform wave shapes leading to featureless overviews. Focus+Context applied to such signals will have limited value unless supplemented with filterable measurement tracks for navigation and feature finding.

6.2 Scalability and Performance

Signal data is read as simple two column tab-delimited text in order to support a variety of data sources. Each line of text corresponds to a time/signal value pair. No specific pre-processing is required and all pre-computations and caching are performed as they are needed. The current implementation of SignalLens is capable of reasonable performance on datasets up to 2-3 Megapoints (10^6). To achieve this performance with standard hardware we implemented several strategies:

Lens calculations are pre-computed at runtime for the unit radius and fetched as needed from a lookup table. We use a lookup table containing 10^6 entries which provides more than sufficient accuracy for our needs. For the hardware described in Table 1 the lookup table can be computed in approximately 160 ms for the quadratic lens. The resulting runtime performance is very fast and consistent across lens

shapes regardless of the complexity of the lens shape (e.g. Gaussian vs. polynomial vs. linear).

We aggregate those line segments in the signal plot which will map within the same horizontal pixel and cache these as a single vertical line segment. Additionally we cache the first and last time points for each line segment to allow drawing the appropriate connection between adjacent line segments. This aggregation enables the undistorted areas to be plotted as a relatively small number of line segments ($\sim 10^3$) rather than several million segments, while still accurately representing the signal's appearance. From the time span of each cached line segment we can determine if a segment will still fall within a single horizontal pixel after lens distortion. In such cases we also plot the cached line segments rather than the lines connecting the individual points.

Table 1 presents typical performance data for three of the datasets discussed in this paper using a nominal off-the-shelf consumer grade computer. For smaller data sets such as the amplitude modulated signal (AM) the frame rate during navigation animation is quite high. For the larger more complex data sets the frame rate declines but is still well within a usable range for visual navigation cues. Frame rates depend on a variety of factors including the size of the dataset as well as the magnification level. Even though the ECG data is larger than the PCI-E data set, the higher magnification results in more use of the cached aggregated line segments and fewer draws of individual time points.

Table 1. Typical Rendering Performance
(2.70 GHz AMD Athlon II X2 235e, 4GB RAM)

Data	Time Points	Magnification	Frames Per Second
AM	65534	500	49
PCI-E	1023983	250	17
ECG	1799998	1000	20

6.3 Lens Navigation

As discussed previously, high magnification creates some difficulties in selecting targets since a pixel of context corresponds to a potentially large number of underlying time points, and the user intent becomes somewhat ambiguous. As with direct target selection, positioning the lens by dragging has the similar difficulties since a single pixel movement corresponds to potentially large time increments which may pass over many time points. Hence, simply dragging the lens is incapable of positioning to a specific target.

For most navigation, direct selection via mouse clicks (e.g. the sequence shown in Figure 4) is more effective than dragging the lens into position. In our design, keyboard actions (arrow keys) can be used to micro-position the lens to the precise time point required. Navigation is further supplemented by the ability to navigate to a specific time or individual features in the filtered measurement tracks. The combination of these navigation aids largely overcomes the issues related to lens movement.

Despite these limitations we maintain the capability to drag the lens in order to maintain consistency with the metaphor of a movable lens, and to allow for coarse positioning.

7 CONCLUSION

We designed SignalLens to explore the possibility of using Focus+Context techniques to address the issues of long electronic signal traces within the constraints of the limited screen real estate typical of current instruments. Our preliminary results using the tool as well as reviews by experts indicate that it is a promising direction for further study. However, a formal user study is required before any definitive claims can be made. In lieu of such studies, we believe the included examples illustrate that the SignalLens design is tractable and useful for conditions constrained by small screens.

The inclusion of computed measurement tracks to aid in the analysis of an acquired signal significantly improves the usability of the system. These tracks can be viewed as extending the overall context of the data. Filtered tracks improve the visibility of salient features and aid in the navigation of the original data set. As we have seen with the PCI-E data, with appropriate domain knowledge, even minute features can be found visually after computing and filtering characteristics of the signal. Measurement tracks enable rapid analysis and navigation even if the context/overview of the full data set is so dense as to be featureless to the user. While this implementation has proven to be valuable for SignalLens's Focus+Context implementation, the same concept could be applied to an Overview+Detail approach. In this case the tracks should be plotted in parallel with *both* the overview and the detail view. This would enable viewing features at both levels of detail, and offer the same affordances.

While SignalLens has made progress towards handling very large time series, further work and optimization still needs to be made to handle Gigapoint traces with reasonable interactivity. Our current implementation has reasonable performance for dataset as large as 2-3 Megapoints simply using an aggregation and caching strategy. Scaling to larger datasets will likely require graphics hardware acceleration and/or leveraging the GPU as a computational engine.

REFERENCES

- [1] W. Aigner, S. Miksch, and W. Muller, "Visualizing time-oriented data—A systematic view," *Computers & Graphics*, vol. 31, pp. 401-409, 2007.
- [2] P. Buono, A. Aris, C. Plaisant, et al., "Interactive pattern search in time series," *In Proc. SPIE*, pp. 175-186, 2005.
- [3] M. Carpendale and C. Montagnese, "A framework for unifying presentation space," *In Proceedings of the ACM Symposium on User interface Software and Technology*, pp. 70, 2001.
- [4] S. Carpendale, J. Ligh, and E. Pattison, "Achieving higher magnification in context," *In Proceedings of the ACM Symposium on User interface Software and Technology*, pp. 71-80, 2004.
- [5] A. Cockburn, A. Karlson, and B. B. Bederson, "A Review of Overview plus Detail, Zooming, and Focus plus Context Interfaces," *ACM Computing Surveys*, vol. 41, pp. 1-31, 2008.
- [6] G. Furnas, "Generalized fisheye views," *ACM SIGCHI Bulletin*, vol. 17, pp. 23, 1986.
- [7] G. Furnas, "A fisheye follow-up: further reflections on focus+ context," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1008, 2006.
- [8] A. L. Goldberger, L. A. N. Amaral, L. Glass, et al., "PhysioBank, PhysioToolkit, and PhysioNet - Components of a new research resource for complex physiologic signals," *Circulation*, vol. 101, pp. E215-E220, 2000.
- [9] L. Gomella and S. Haist, *Clinician's Pocket Reference*: McGraw-Hill Medical, 2003, p390.
- [10] C. Gutwin, "Improving focus targeting in interactive fisheye views," *In Proceedings of the SIGCHI Conference on Human factors in Computing Systems*, pp. 267-274, 2002.
- [11] C. Gutwin and A. Skopik, "Fisheyes are good for large steering tasks," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 208, 2003.
- [12] H. Hochheiser and B. Shneiderman, "Interactive exploration of time series data," *Lecture Notes in Computer Science*, pp. 441-446, 2001.
- [13] E. Keogh, S. Lonardi, and B. Chiu, "Finding surprising patterns in a time series database in linear time and space," *In Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data mining*, pp. 550-556, 2002.
- [14] W. Kester and A. Devices, *Data Conversion Handbook*: Newnes, 2005, p. 608.
- [15] R. Kincaid and K. Dejgaard, "MassVis: Visual analysis of protein complexes using mass spectrometry," *In Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, pp. 163-170, 2009.
- [16] R. Kincaid and H. Lam, "Line graph explorer: scalable display of line graphs using Focus+ Context," *In Proceedings of the Working Conference on Advanced Visual Interfaces*, pp. 404-411, 2006.
- [17] N. Kumar, N. Lolla, E. Keogh, et al., "Time-series bitmaps: A practical visualization tool for working with large time series databases," *In SIAM 2005 Data Mining Conference*, pp. 531-535, 2005.
- [18] J. Lamping, R. Rao, and P. Pirolli, "A focus+ context technique based on hyperbolic geometry for visualizing large hierarchies," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 408, 1995.
- [19] Y. Leung and M. Apperley, "A review and taxonomy of distortion-oriented presentation techniques," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 1, pp. 126-160, 1994.
- [20] J. Lin, E. Keogh, S. Lonardi, et al., "Finding motifs in time series," *In Proc. of the Workshop on Temporal Data Mining*, pp. 53-68, 2002.
- [21] R. Lopez-Hernandez, D. Guilmaine, M. J. McGuffin, et al., "A layer-oriented interface for visualizing time-series data from oscilloscopes," *In Proceedings of the IEEE Pacific Visualization Symposium (PacificVis), 2010*, pp. 41-48, 2010.
- [22] P. McLachlan, T. Munzner, E. Koutsofios, et al., "LiveRAC: Interactive visual exploration of system management time-series data," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1483-1492, 2008.
- [23] E. Pietriga and C. Appert, "Sigma lenses: focus-context transitions combining space, time and translucence," *In Proceeding of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1343-1352, 2008.
- [24] E. Pietriga, C. Appert, and M. Beaudouin-Lafon, "Pointing and beyond: an operationalization and preliminary evaluation of multi-scale searching," *In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 1224, 2007.
- [25] C. Plaisant, D. Carr, and B. Shneiderman, "Image-browser taxonomy and guidelines for designers," *IEEE Software*, pp. 21-32, 1995.
- [26] W. H. Press, *Numerical recipes in C++*, 2nd ed. Cambridge; New York: Cambridge University Press, 2002.