Interactive Visual Analysis of a Timing Chain Drive
Using Segmented Curve View and other Coordinated Views

Zoltán Konyha
VRVis Research Center, Vienna

Konyha@VRVis.at

Krešimir Matković
VRVis Research Center, Vienna

Matkovic@VRVis.at

Denis Gračanin
Virginia Tech, Blacksburg

gracanin@vt.edu

Mario Duras
AVL Croatia, Zagreb

mario.duras@avl.com

Josip Jurić
VRVis Research Center, Vienna

Juric@VRVis.at

Helwig Hauser
University of Bergen

Helwig.Hauser@UiB.no

Abstract

A timing chain drive transfers motion from the engine's crankshaft to the camshaft that operates the valves. The design process of timing chain drives involves computer simulation of many design variants in order to find an optimum. Most of the simulation results can be represented as families of function graphs (data series). Previously, the analysis of those results was based on static 2D diagrams and animated 3D visualizations. They were suitable for the detailed analysis of a few simulation variants, but not for the comparison of many cases. In this paper we propose a new approach to the analysis based on coordinated linked views and advanced brushing features. Our proposed method supports the interactive analysis of many design variants. We introduce a novel view, called segmented curve view, which can display distributions in families of function graphs. The segmented curve view combines individual function graphs where for a fixed value of the independent variable, a bar extends from minimum to maximum values across the family of function graphs. Each bar is divided into segments (bins) with a color that represents the number of function graphs with the value in that segment. In the case study, we demonstrate that the new view combined with "traditional" views provides a strong support for the interactive visual exploration and analysis of a real world timing chain design problem.

Keywords— **Visual exploration, coordinated views, segmented curve view, aggregated views, composite brushing.**

1 Introduction

Durability and low noise levels are some of the key design goals for car engines. One of the factors affecting those design criteria is the quality of the motion transfer from the engine's crankshaft to the camshaft. The camshaft rotation must be synchronized to the piston's motion since the camshaft actuates the valves. Most automotive engines use chain or belt drives for this motion transfer. Various factors such as high engine speeds, dynamic and inertial phenomena may cause undesired vibrations that increase noise levels and mechanic wear. This behavior can be simulated in software thus simulation tools are used extensively in the design of timing chain drives.

The design process involves running many simulation variants with various values of input (control) parameters. However, the analysis of many simulation runs poses new challenges. Previous analysis methods included static diagrams and animated 3D visualizations [14, 6] suitable for the detailed analysis of one or very few simulations. However, they do not facilitate the interactive comparison of many different simulations variants. We demonstrate that using multiple linked views of timing chain simulation data simplifies and accelerates the analysis.

In a timing chain simulation data set, the control parameters constitute a subspace that we call independent variables. The output (response) parameters are the dependent variables. Many response parameters are data series, often represented as function graphs. We use the term *family of function graphs* to denote a set of functions that take the independent variables and an additional variable like time as input. A significant part of the simulation results of timing chain drives are families of function graphs — one function graph for each simulation run. In our previous work [5, 8] we described an approach and a related tool for interactive visual analysis with emphasis on the function graph characteristic of the data. A family of function graphs was presented as an aggregated 2D plot view, the *curve view*, of all the function graphs.

While this approach proved to be very effective, the curve view fails to reflect certain properties of some data sets containing families of function graphs. Accordingly, we propose a novel view, the *segmented curve view*, which is related to histograms or histogram-like frequency-based

visualization techniques. For a fixed value of the independent variable, a bar extends from minimum to maximum values across the family of function graphs. Each bar is divided into segments (bins) where color represents the number of function graphs with the value in that segment. The new view provides a valuable insight into chain simulation data and offers, in some cases, significant advantages over the curve view.

The remainder of the paper is organized as follows: Section 2 describes related work on the visualization of data series, segmenting and grouping of function graphs for analysis. Section 3 discusses the new view and the related analysis techniques. Section 4 provides a short introduction to the rigid body based simulation of timing chain drives. Section 5 gives a case study that illustrates analysis and exploration of simulation results using multiple linked views. Section 6 concludes the paper.

2 Related Work

Visualizing families of function graphs is part of a larger topic of high dimensional and time-dependent data visualization [15]. The new segmented curve view is based on the curve view [5, 8] and several other visualization techniques including binning, box plot views and methods for the visualization of time-dependent data.

The process of binning [13] partitions the data space into intervals (bins) and assign an occupancy value to each bin indicating the number of data items that belong to the bin. In other words, the "new data" represents a count of the original data items in each bin.

A histogram is one example of binning. The width of each bin and the offset of where the bins start determine how the histogram appears. There are many other applications of binning, for example, the identification of outliers in parallel coordinates visualization of large data sets [11].

The box plot view [16] is a graph that summarizes the statistical measures like median, upper and lower quartiles, and minimum and maximum data values. It provides a graphical display of a variable's location and spread and allows quick examination of a data set by presenting a summary of the data distribution. Unlike histograms and probability density functions, it does not require assumptions of the statistical distribution.

Box plots can be modified to provide additional information like the sample size, density information and further statistical properties [12]. The variations include histplot, box-percentile plot, violin plot, and section density plot [12]. For example, the section density plot uses occlusion and color intensity variation to create an implied 3D display of distribution.

The box plot can also be used for higher dimensional data. The main problem is how to represent the summary values using some simple visual metaphors with meaningful spatial positions. The variant of the bivariate box plot include the rangefinder box pot, two-dimensional box plot, bagplot, relplot and quelplot [12].

Müller and Schumann provide an overview of the visualization methods for time-dependent data [10]. The related visualization techniques can be classified into two distinct groups based on whether or not the visual representation itself is time-dependent.

The simplest approaches use mapping of time on a quantitative scale. That includes a *Sequence Graph* for the one-dimensional case and a *Time Series Graph* for the two- and three-dimensional cases. Similarly, we can leverage characteristics of the data set to use a *Point Graph* (point data), a *Line Graph* (continuous data), a *Bar Graph* (cumulative data), and a *Circle Graph* (cyclic data) [10].

Special visual metaphors like *Calendar View* [17] or *Lexis Pencils* [3] are also used. In *Calendar View*, van Wijk et al. use a very intuitive calendar view to display clusters of time series data [17]. Lexis pencils [3] display various time-dependent variables to the faces of a pencil. Pencils can depict several variables and can be positioned in space to indicate the spatial context of the data.

Extensions of traditional visualization techniques to time-dependent data include *Wormplot* that provides a scatter plot for each time step. *Time Wheel* leverages parallel coordinates so that the time axis is in the center and other axes are located circularly around it. *MultiComb* arranges those other axes in a star-shaped pattern [10].

Brushing the time axis to display details of the selected time frame is a very common and useful interaction technique used with static representations. *Timebox* widgets by Hochheiser et al. [4] can be used to brush both the time axis and the attribute axis of graphs.

Carlis et al. introduce spiral graphs [1] to facilitate exploration of serial and cyclic data. Spiral arrangement of data provides easy visual cues to both serial and periodic aspects of the data.

A combination of the curve view and conventional views [5, 9] can be used for the analysis of similar data sets. The curve view can show all function graphs at once using transparency to depict the density of the curves. When combined with linked views and brushing techniques it can be used to display curves in focus and those forming the context. The curve view was used to successfully display and explore data sets containing more than 40,000 function graphs per dimension.

A similar approach is described in [19]. This approach uses extruded parallel coordinates, linking with wings and three-dimensional parallel coordinates integrated in a single rendering system that visualize trajectories of higher dimensional dynamical systems [19].

3 Proposed Approach

In general, a data set contains m independent variables and n dependent variables. The independent variables can be expressed as $\mathbf{x} = [x_1, \ldots, x_m] \in I$. A dependent variable can be either a scalar or a data series (function graph). For a given dependent variable $f_j(\mathbf{x}, t)$ that is a function graph, a family of function graphs is a set of function graphs for each possible value of $\mathbf{x}$, $\{f_j(\mathbf{x}_i, t) | \forall \mathbf{x}_i \in I\}$.

The combination of function graphs and scalars in the data set is particularly challenging. In our previous work we used a combination of the curve view and standard views (e.g. histogram or parallel coordinates) [5, 9] for the analysis of such data sets (Figure 1).
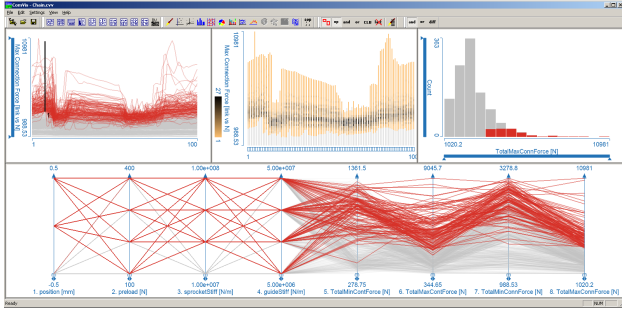

Figure 1: An example of coordinated multiple views (curve view, segmented curve view, histogram, and parallel coordinates) and a composite brush.

We use multiple linked views and advanced interactive brushing to assess the data using iterative visual analysis. The key factor is the ability to create complex, composite brushes that span multiple views. The brushes are constructed using various combination schemes based on AND, OR, and NOT operations. The first operand is always the result of the previous composition. The next iteration takes the state of the previous brushing and the new brush as the only two arguments in the new operation. This approach provides for an intuitive and iterative workflow.

All the variables are treated the same so it is possible to brush in views of both dependent and independent variables to study their relationship in both directions. For example, brushing some output variables allows us to find suitable inputs. For more detailed explanation of analytical procedures and brushing mechanism see [5].

The curve view allows us to display a whole family of function graphs and select a subset of graphs. Figure 2 shows the maximum connection forces between neighboring chain links from a timing chain simulation. Since there are 100 chain links in the chain, the horizontal axis is labeled from 1 to 100. There are 1,152 different simulation variants with different settings of design parameters. The

curve view shows the family of 1,152 function graphs, one function graph represents each simulation variants.
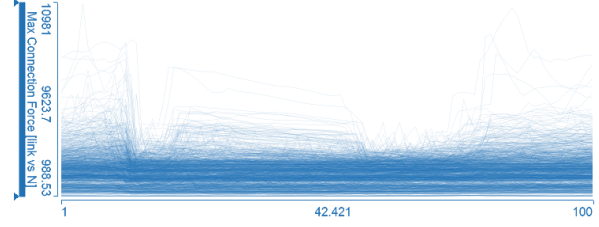

Figure 2: Curve view of a family of function graphs. The x axis represents the chain link index (1–100). The y axis indicates the maximum connection force at the given link. In this curve view there is one function graph for each of the 1,152 simulation variants, producing an overlay of 1,152 individual function graphs of maximum connection force.
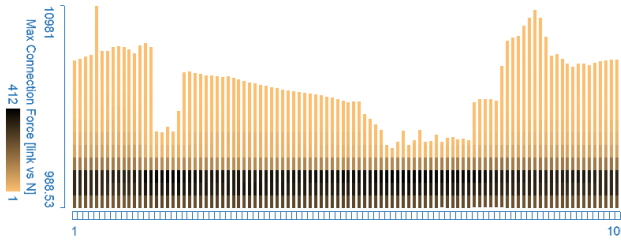
However, there are several issues with the curve view. The continuous curve view suggests continuity in the displayed data series. If the actual data is continuous in nature then connecting the sample points with continuous lines in a function graph plot is a valid and meaningful approach. However, if the data is not continuous then the visualization is semantically incorrect. Choosing the right transparency factor in an alpha-blended curve overlay is always a difficult task. A transparency setting that reveals the distribution in densely populated regions makes outliers barely visible. Conversely, a transparency setting that preserves outliers makes dense regions indiscernible because of overdraw. We address those issues by suggesting a novel view, the segmented curve view, for the visualization of families of function graphs. We will use the data set in Figure 2 to introduce the segmented curve view so that direct comparisons between the two views can be made.
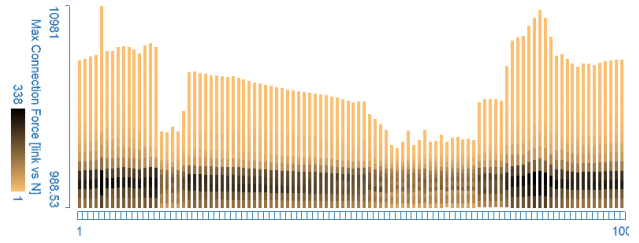
3.1 Segmentation and Binning

The segmented curve view provides two-fold segmentation of the displayed data. Individual bars for each value of t serve to avoid the possibly false suggestion of continuity and to provide segmentation along the independent axis.

The bars depict minimum and maximum values of the dependent variable for each independent variable. This idea is not new and similar approaches have been used [2]. However, we extended it and instead of showing only the minimum to maximum range, we also show the distribution of dependent variables in a bar.

In order to depict the distribution we introduce segmentation along the vertical axis. Each bar is divided into segments (bins) of uniform height where each segment is colored to indicate the number of curves passing through the segment. We use the same color mapping (the orange-

(a) Global binning creates equal size bins in all bars, which makes direct comparison of bars easier, but fails to show fine details of distribution in shorter bars.

(b) Local binning creates the same number of bins in all bars, which reveals more information in densely populated areas.

Figure 3: (a) Global and (b) local binning in the segmented curve view using 16 bins. Orange bins contain few curves while black ones contain many. In (a) black bins contain 412 curves while in (b) they contain 336.

black color gradient) throughout the paper to provide a consistent presentation of views. Bins with the smallest curve count are shown in orange and the ones with the largest curve counts are black. Although a detailed discussion of color scales is not in the scope of this paper, we must mention that there are perceptual issues related to color mapping [18]. The human eye distinguishes some hues and gradients better than others. The orange-black gradient was chosen because it seems to convey the density of curves well.

When segmenting a bar we can take the bar's actual minimum and maximum and uniformly partition this range. We term this *local binning*. Alternatively, we can partition the range defined by the overall minimum and maximum of all bars. We term this *global binning*. Figure 3 illustrates both binning strategies.

Global binning produces bins in different bars that have the same limits. This makes the direct comparison of distributions of various bins possible. However, when the minima and maxima of various bars are very different, many bins will be out of that range, leaving relatively few within the bars. That does not allow a detailed display of distribution within the bars. We could create more bins to have a better resolution. The number of curves in the most populated bin will decrease and the color gradient will have fewer entries. That fails to show the distributions properly.

On the contrary, local binning creates the same number of bins in each bar. This provides a finer resolution for bars whose minimum and maximum values are closer. Since the bin limits in different bars are different, comparing distributions in different bars is less straightforward. Consequently, global binning works better with more uniform minimum and maximum values, while local binning shows distribution in very non-uniform data using relatively few bins.

Some bins between the minimum and maximum of a bar may have no curves in them so they should not be visible. Bins beside empty ones are not drawn using original

bins limits. If a non-empty bin has an empty bin as a neighbor, then the limits of the non-empty bin are moved to the actual minimum or maximum values of the bin. Figure 4 shows the procedure and Figure 5 shows the result.
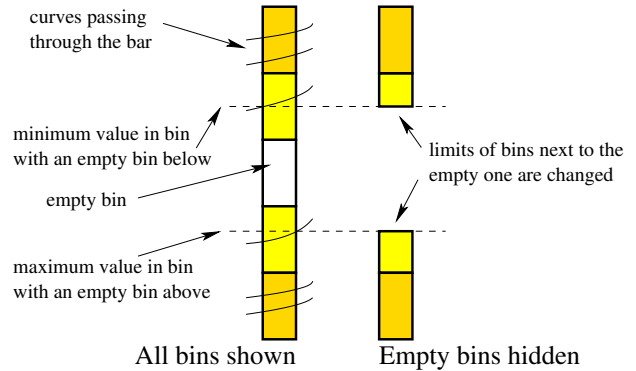

Figure 4: The process of removing empty bins. The limits of the bins next to the removed one are set to the actual minimum and maximum values of the curves in them.
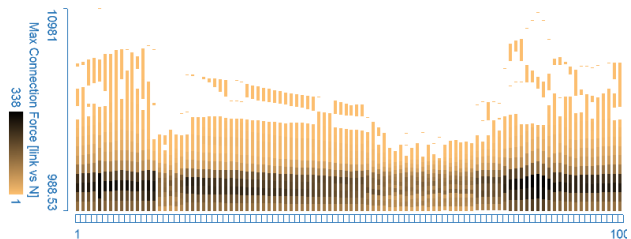

Figure 5: Empty bins in Figure 3(b) are hidden using the idea shown in Figure 4.

We propose two possibilities to add frames to the individual bars in order to preserve their integrity (Figure 6).
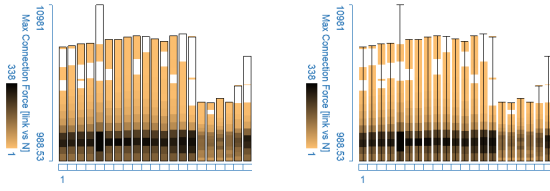
Figure 6: Framing the bars with rectangles (left) or boxplot-like lines (right) helps preserve the integrity of bars when empty bins are hidden.
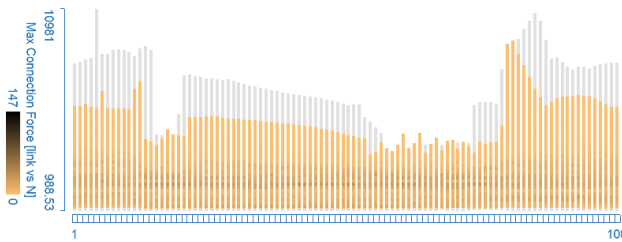
3.2 Color Mapping Strategies and Linking

By default, bins that contain one curve are drawn using the first color of the gradient. The most populated bin of all is drawn using the last color of the gradient. This is termed *global color scale*.

The segmented curve view provides focus+context visualization. Figure 7 illustrates that the focus set is visualized using colored bins and the context is shown in uniform gray. Local binning is performed on the range between the minimum and maximum values of the currently brushed items in each bar, as opposed to the minimum and maximum values in each bar, Consequently, more details of the distribution are displayed in the focus.

When the segmented curve view is linked to some other view, the most populated bin in the focus set can contain considerably fewer curves than the absolute maximum. This means that only the first few colors of the color gradient are used for the displayed bins. Figure 7(a) shows that this renders the color gradient and the distribution less visible. We can overcome this problem by mapping the last color in the gradient to the number of curves in the most populated bin in the actual focus set. This *local color scale* is shown in Figure 7(b). It helps to reveal the details of the distribution in the focus.

However, using a global color scale has an important advantage: the same color represents the same number of curves regardless of how many items are actually brushed. This preserves the temporal coherence of the visualization.
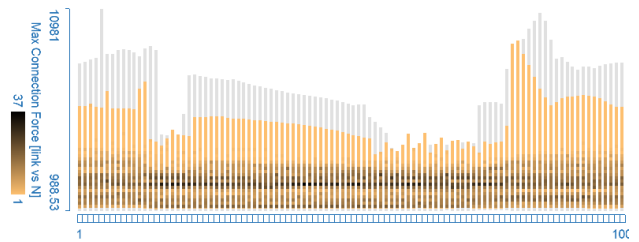
3.3 Brushing in the Segmented Curve View

We can brush a specific range in the segmented curve view to reveal finer details of the distribution. Local binning is performed on the brushed range to reveal more details and the brushed items are also highlighted in the linked views. We offer a rectangular brush that can span multiple bars. There are two ways to interpret a brush that spans multiple bars. A curve can be selected if it enters and leaves the rectangle through its vertical edges (Figure 8(a)). That is the intersection (logical AND) of the sets of curves selected in the individual bars. Alternatively, a curve can be selected if it enters and leaves the rectangle through any edge (Figure 8(b)). This is the union (logical OR) of the curves selected in the individual bars.

3.4 Comparison with the Curve View

The curve view suggests that the displayed data series is continuous while the segmented curve view does not. For example, the time scale of function graph data acquired via measurements or produced in simulation is inherently discrete. The segmented curve view may be a better choice for such discreet data, particularly when connecting the sample points with continuous lines may be misleading. An obvious example is the Fourier spectrum of time series data.
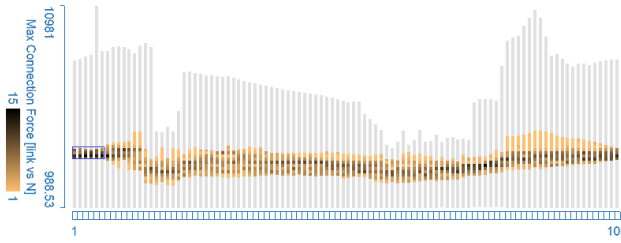
Furthermore, the analysts are often interested in the minima, maxima and distribution of function graphs of a family. They want to see and compare how many graphs fall within given ranges and discover patterns or outliers in the distribution. More specifically, they want to find out if there is any correlation between the distribution for different values of t. In Figure 8(a) for example, the relatively narrow and uniformly distributed range expands but still has a roughly uniform distribution around the middle of the horizontal axis. However, near the right end of the axis most curves are concentrated in the lower ranges.

One could argue that the segmented curve view is nothing more than a discretized version of the alpha-blended curve overlay view, especially if many bins are created with global binning. However, there is a very important

(a) Global color scale. The last color of the gradient is mapped to the most populated bin in the entire family.
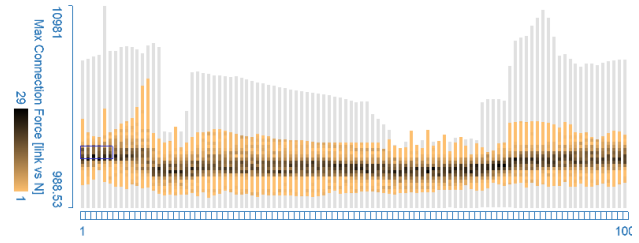

(b) Local color scale. The last color of the gradient is mapped to the most populated bin in the current focus set.

Figure 7: (a) Global and (b) local color scales.

(a) Curves that never leave the brush are selected by taking a logical AND of the brushed items in individual bars.

(b) Curves that enter the brush anywhere are selected by taking a logical OR of the brushed items in individual bars.

Figure 8: Comparison of the brushes spanning multiple bars. In both (a) and (b) middle ranges of the first six bars are brushed by the blue rectangle. The difference is how the sets of items selected in the individual bars are combined.

advantage of the segmented curve view compared to alpha-blended curve overlays: the colors assigned to the segments can be chosen from an arbitrary color scale.
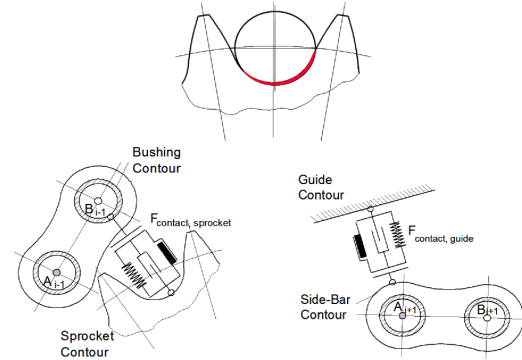
This avoids one drawback of alpha-blended curve overlays. When many curves are displayed in alpha-blended curve overlays, their transparency has to be increased to avoid clutter because of overdraw. However, it is often difficult to discern fine details of the distribution, especially in areas where only few curves pass through. The individual curves are simply too transparent to be visible. With the segmented curve view one can use a color scale for the segments that is clearly visible against the background and the visualization preserves both outliers and fine details in the distribution.
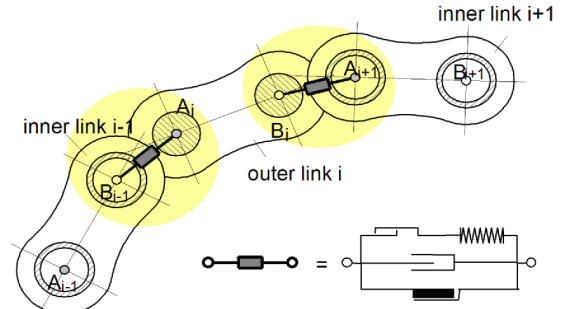
4 Simulation of Timing Chain Drives

A timing chain drive (or a belt drive) is found in most production car engines. The purpose of the timing chain drive is to transfer motion from the engine's crankshaft to the camshaft. The camshaft actuates valves that must be synchronized to the piston's motion. Therefore, precisely two crankshaft rotations are needed for a single camshaft rotation. The chain's motion deviates from its ideal kinematic path especially at high engine speeds. Dynamic and inertial phenomena cause vibrations that increase noise levels and mechanic wear. The vibration causes the accuracy of the motion coupling to deteriorate. The camshaft's rotational velocity does not remain constant and undesired high frequency components are added. This induces rougher and less controlled valve operation which can reduce fuel economy and power output.

This behavior can be simulated in software thus simulation tools are used extensively in the design of timing chain drives. The basic approach is to model each chain link as a single rigid mass connected by spring/damper units to neighboring chain links and simulate the resulting multi-body system based on the Newton-Euler laws. The motion of the chain is computed only in the y-z plane. The motion perpendicular to the plane can be ignored without

influencing the validity of the simulation. Therefore, each chain link has three enabled DOFs (degrees of freedom) — translations in y and z directions and rotation around x-axis. Sprockets are modeled as generic mass elements with only one DOF — rotational around x-axis.

(a) A model of the contact forces between chain links and sprockets. The red region indicates the overlap area of the contact contours. Each chain link has smaller circular contours at both ends that model the bushing which comes into contact with the sprockets. The side-bar contour is the outer edge of the chain link that slides on guides.

(b) A model of the connection forces. Each chain link is connected to its neighbors via stiff spring and damper elements.

Figure 9: Contact and connection forces in the chain simulation.

In a simulation the objects are described using their contact contours. The contours of sprockets and guides reflect their actual shapes. However, the contours of chain links consist of circles around the pins that interconnect the chain links, because these are the surfaces in contact with the sprockets and guides. There are two circular contours at both pins of the chain links. The smaller circle is the surface that comes into contact with the sprocket. The larger circle is the outer surface of the link sliding on guides.

The simulation computes (among other attributes) forces between elements. There are two classes of forces: (a) contact forces between chain links and the sprockets/guides, and (b) connection forces between neighboring chain links. The contact forces act when a link comes into contact with a sprocket or guide. The algorithm for computing contact forces is based on evaluating the size of the overlap area between the contact contours of the two objects, the stiffness of materials, the relative velocities and the damping properties of the materials. The connection forces act between neighboring chain links and are computed in a similar way. Figure 9 illustrates the corresponding connection and contact force models. For a more detailed explanation about the simulation of chain drives with rigid body dynamics please refer to [14].

5   Case Study: Interactive Visual Analysis of a Timing Chain Drive

One of the authors is an engineer working on valve train and timing drive design projects at AVL-List GmbH. This section documents our analysis procedure of a real timing chain drive model. We used a model of a basic type chain drive system which consists of two sprockets, the camshaft sprocket (38 teeth) and the crankshaft sprocket (19 teeth). Additionally, two fixed guides lead a bushing chain along the chain path to reduce vibrations. The constant load is applied at the camshaft sprocket and a constant rotation is prescribed at the crankshaft sprocket. Lash in the chain and friction in contacts between the chain and sprockets are not considered. A commercially available software package, EXCITE Timing Drive from AVL [7], was used for the timing drive simulation. It is a multi-body simulation software tool for the simulation of engine valve train components. Figure 10 shows the model in the simulation software at two different levels of abstraction.

5.1   Simulation Parameters

Numerous control parameters can be defined in the simulation software. In the scope of this case study we varied engine speed and four design parameters: sprocket stiffness, guide stiffness, chain preload and camshaft sprocket offset from the designed position in the z direction (see Figure 10). A positive offset means that the camshaft sprocket is further away from the crankshaft than its kinematically ideal position, therefore the chain becomes rela-

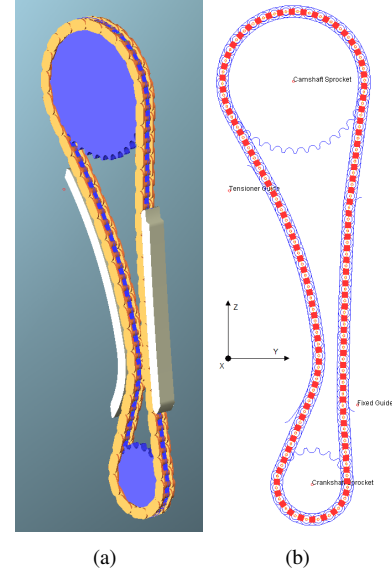tively short. Table 1 shows the ranges of variation of the parameters.

Figure 10: (a) 3D model of the chain and (b) the model simplified to simulation contact contours. Red boxes indicate the centers of gravity of chain links.

Table 1: Control parameters

Parameter	Values	Unit
sprocket stiffness	1.0E+7, 4.0E+7, 7.0E+7, 1.0E+8	N/m
guide stiffness	5.0E+6, 2.0E+7, 3.5E+7, 5.0E+7	N/m
chain preload	100, 200, 300, 400	N
sprocket offset	-0.5, 0.0, +0.5	mm
engine speed	1000, 2000, 3000, 4000, 5000, 6000	rpm

There are 1,152 possible combinations of these parameters or 1,152 simulation cases. A simulation run is performed for each case with a simulation period that is sufficient for all chain links to complete a full revolution. Five response parameters were computed for each case:

- Maximum contact forces [N] for each chain link. This is a function graph where the independent variable represents a chain link (1 to 100) and the dependent variable represents the maximum of the contact forces that act on the link in the simulated time span.
- Maximum connection forces [N] for each chain link. This is a function graph similar to the previous one.
- Overall maximum contact force [N]: scalar.
- Overall maximum connection force [N]: scalar.
- Fourier Transform of the camshaft sprocket's rotational velocity [rad/s vs. orders]: function graph.

(a) Contact forces for each chain link.   (b) Parallel coordinates of the parameter space.   (c) Connection forces between chain links.
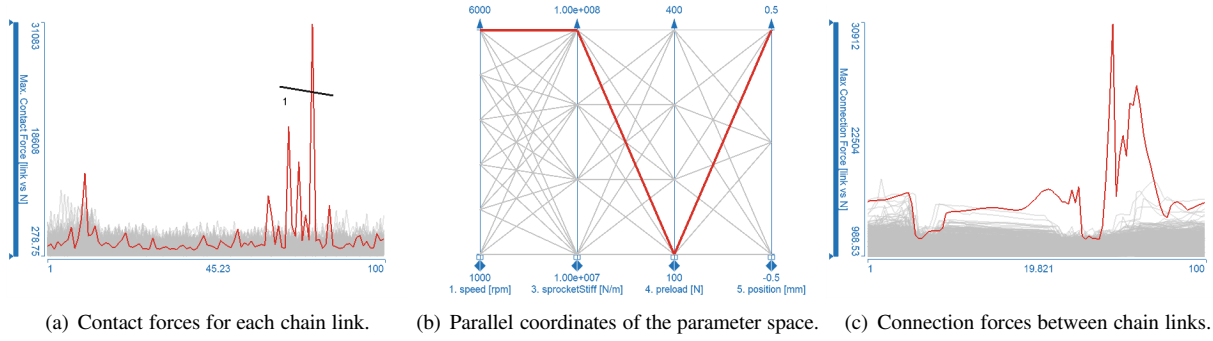
Figure 11: Invalid parameter combinations investigation. One graph of very unusual shape is brushed in red in (a). Other graphs are shown in gray to provide context. This graph shows that a number of chain links suffer extremely high contact forces. The parallel coordinates reveal that this happens at 6000 rpm when the camshaft sprocket is offset by 0.5 mm, the sprocket material is very stiff and the preload is small. The corresponding connection forces (c) are also very high.

5.2 Visual Analysis Goals

Standard tools for the analysis of chain drive simulation data include 2D charts and animated 3D representations, as described in [14, 6]. Those methods are sufficient if engineers want to analyze one or only a few simulation runs. With increasing computer power it is possible to run multiple simulations (several thousands of them) in order to fine tune design parameters for optimal results. However, the analysis of many simulations requires a different approach. For example, studying one thousand 2D charts describing connection forces is not practical, nor is an elaborate 3D rendering of one thousand cases with many geometric details. We have developed a new visualization approach for this application and propose a new view to analyze some attributes of such a large number of simulations.

Engineers generally face four main tasks in the analysis of chain drive simulations:

- finding invalid parameter combinations
- parameter sensitivity evaluation
- reducing chain noise
- keeping contact and connection forces within a range

5.3 Invalid Parameter Combinations

The first task engineers are interested in is checking if some of the control parameter combinations result in invalid responses. These combinations must be excluded from further analysis. This is an investigation process.

In order to check results, the curve view is used to investigate maximum contact forces. One clear outlier curve is visible in Figure 11(a). The outlier is selected by the black line brush in and parameters that lead to such system behavior are shown in the parallel coordinates Figure 11(b). This outlier comes from a simulation case where the very stiff camshaft sprocket is pushed 0.5 mm away from the crankshaft, the preload is small and the engine speed was

high. A quick investigation of this case in the animated 3D visualization [6] revealed that the chain is vibrating wildly, as shown in Figure 12. We suspect that the aforementioned combination of parameters produces a resonance in the chain that generates extremely high forces. The behavior is actually so out of control that we assume that there may be an error in the other model parameters, perhaps incorrect initial conditions for this case.
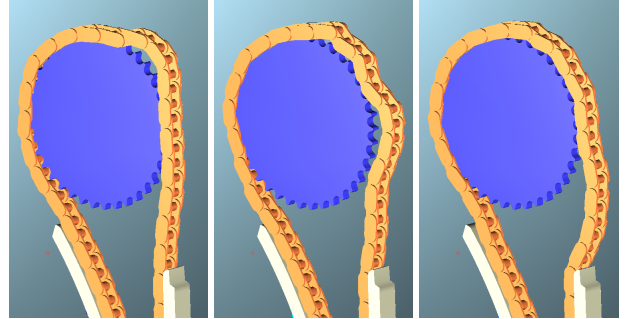

Figure 12: Three frames of the animated 3D view of the extreme chain vibration.

5.4 Parameter Sensitivity Evaluation

This is also an investigation process where the goal is to identify the main parameters and understand how changing those parameters influences the defined results. If one considers the simulation process as a black box, then this procedure can be called a black box reconstruction step [5].

First, we investigate connection forces at various engine speeds. The engine speed is not a design parameter that we can control but we want to understand how the results depend on engine speed. We brush various engine speeds and study the linked segmented curve view displays. Figure 13 shows connection forces at 1000 rpm and 4000 rpm.

(a) Connection forces at 1000 rpm.
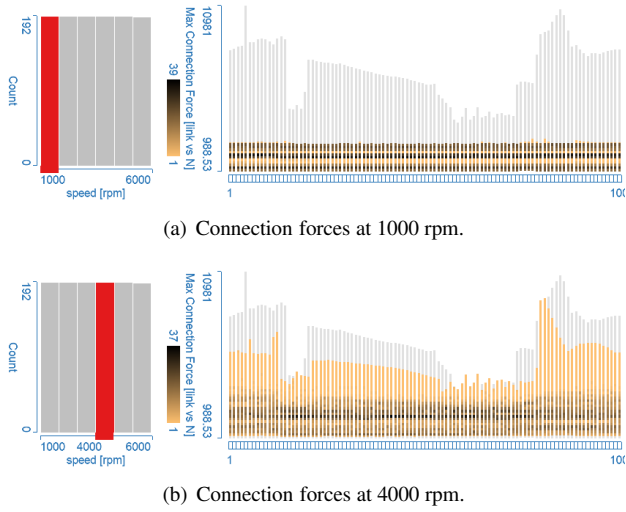

(b) Connection forces at 4000 rpm.

Figure 13: Engine speeds of (a) 1000 rpm and (b) 4000 rpm are brushed and the corresponding connection forces are highlighted in the segmented curve view that uses 64 bins, global binning and a local color scale. The connection forces are rather varied at 4000 rpm, but at 1000 rpm there are three clearly visible clusters.
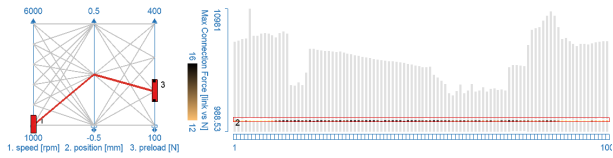

Figure 14: First, engine speed of 1000 rpm was brushed in the parallel coordinates. This highlighted three clusters in the segmented curve view of connection forces (Figure 13(a)). An AND brush (see Section 3.3) spanning all bars was used to select the middle cluster. The corresponding sprocket offset is shown in the parallel coordinates. Brush number 3 narrows the focus to a specific preload setting in order to examine the additional influence of preload on the forces. By moving brush 3 we discover that higher preload causes slightly higher connection forces.

There are three clearly visible clusters of connection forces at 1000 rpm in Figure 13(a). We assume that there is a control parameter causing this clustering. This will be investigated later. We also note that at 4000 rpm there are no clearly visible clusters. This indicates that the individual graphs exhibit more variety. We infer that it is probably more complicated to keep the connection forces in a specified range at higher engine speeds.

We examine the clusters noted in Figure 13(a). This time we brush the engine speed in the parallel coordinates and then brush the three clusters of connection forces

(Figure 14, brushes number 1 and 2). The corresponding sprocket offsets are highlighted in the parallel coordinates. The low, middle and high clusters correspond to -0.5 mm, 0 mm and $+0.5$ mm sprocket offsets, respectively. Cases with high connection forces are not invalid, but they are not desired, since they cause increased wear.

We examine the consequences of various preload parameters by brushing the corresponding axis in parallel coordinates (Figure 14, brush number 3). Larger preload values generate slightly larger connection forces, but the influence is far more subtle than that of the sprocket offset.

(a) Low guide and sprocket stiffness produces low contact forces. The distribution is very even, which means similar forces act on all chain links.

(b) Higher sprocket stiffness creates higher contact forces. Very few chain links have exceptionally high forces. The reason for three-way clustering is again the different sprocket offset parameter like in Figure 13(a).

(c) Stiffer guide material also increases the contact forces, but it also makes their distribution more varied. Some chain links often suffer considerably larger than others.
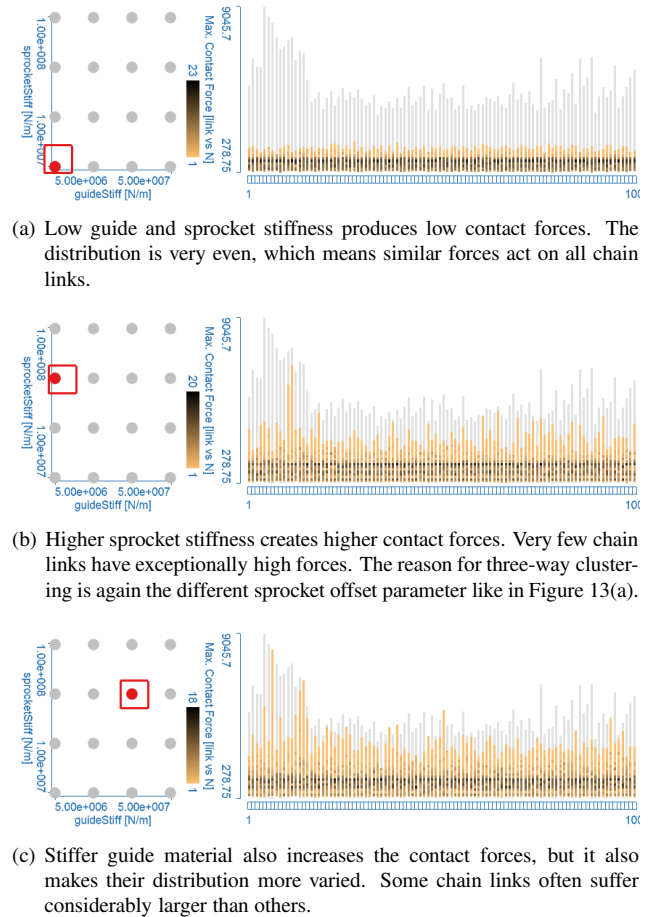
Figure 15: Three snapshots showing the influences of various stiffness parameters on contact force.

Finally, we examine the influences of the different guide and sprocket stiffness parameters on the contact forces. We brush those two parameters in a scatterplot, move the brush to various combinations and study the contact forces. This would be a difficult task with the alpha-blended curve view. It is practically impossible to choose an alpha value which is transparent enough to reveal details in densely populated regions but at the same time opaque enough to preserve

graphs in less crowded parts of the view.

The segmented curve view succeeds in displaying the details of the distribution in dense areas and the outlier graphs at the same time. Figure 15 captures three snapshots of this interactive exploration. We find that using stiffer material for sprockets increases the maximum contact force on chain links. Stiffer guide material has similar implications, but it also makes the distribution of contact forces less even. A very stiff guide results in some of the chain links suffering extremely high contact forces, while forces on others links is kept within a tighter range. It is obvious that this uneven distribution of forces causes extra wear of the more loaded chain links and it is more preferable to have an even distribution.

5.5 Optimization

We investigate two optimization goals. We try to reduce chain noise and keep the magnitude and fluctuation of maximum forces low in order to reduce dynamic load on the chain that causes extreme wear.

The level of chain noise is in correlation with the magnitudes of the contact forces. The noise spectrum is related to the higher order components of the camshaft sprocket's rotational velocity. Therefore we want to minimize peaks in the Fourier transform of the rotational velocity. We first examine what conditions lead to high peaks in the spectrum. Figure 16 shows that these peaks appear at higher engine speeds only. However, even at 6000 rpm there is a gray part of the histogram, which means there are parameter combinations which avoid those high peaks in the spectrum at 6000 rpm.

There is an interesting peak at the 19th order (located at the right end of the spectrum). This is related to the chain drive specific polygon effect. Chain links engage and disengage to/from the sprocket as it rotates. That produces a whining noise with frequency of $n * (u/60)$ where n is the number of crankshaft sprocket teeth and u is the engine speed in rpm.

There is another important consequence. The sprocket's moment is distributed onto a different number of chain links depending on the sprocket's angular position. This can be considered as a discrete process due to high contact stiffness. These abrupt changes in load lead to oscillations in the rotational velocity and chain connection forces with an order that equals the number of crankshaft sprocket teeth. In Figure 17 we brush this peak in the segmented curve view and discover that all simulation cases at 1000 rpm are highlighted in the histogram. That means we are not able to remove the polygon effect completely with any combination of parameters used here. We will try to minimize it, though.

We now try to find the parameter values that produce optimal chain behavior. The ideal case has low contact and

connection forces and the fluctuation of forces should also be minimal. We want flat and low force diagrams. The camshaft sprocket's rotational velocity spectrum should not have high amplitude components, especially at higher orders, in order to reduce high frequency noise.
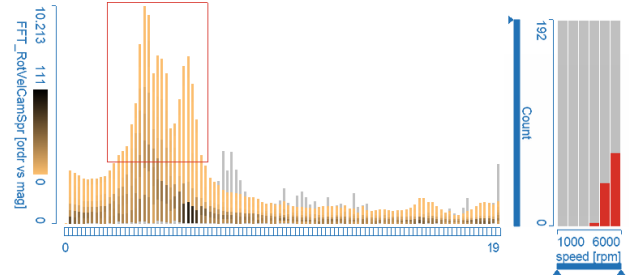

Figure 16: High peaks are brushed with the red rectangle in the Fourier transform of the camshaft sprocket's rotational velocity. The OR brush (see Section 3.3) selects all graphs that enter the rectangle. The histogram shows that these peaks appear at high engine speeds. Note that in this segmented curve view we use local binning because it is more useful for the visualization of the wide variations in the maxima.
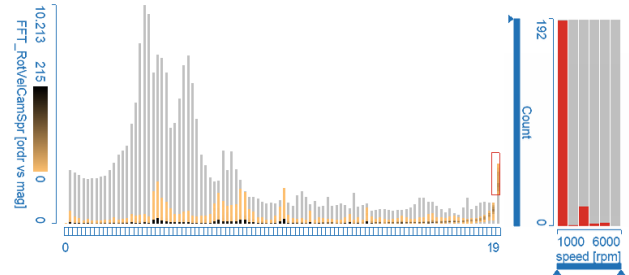

Figure 17: We brush the peak at the 19th order of the FFT of the sprocket's rotational velocity (red rectangle at the right end of the segmented curve view). That peak is related to the polygon effect specific to chain drives. All simulation cases of 1000 rpm engine speed are highlighted in the histogram. That means the polygon effect at 1000 rpm cannot be removed by any combination of the parameters.

We now have an overview of how the parameters influence the result. We want to progressively narrow the set of design parameter values to a single combination which provides the best compromise. This is a highly interactive process where the analyst creates and moves brushes and examines the linked views of the response parameters.

For this procedure we will brush various ranges of control parameters in a parallel coordinates view of the design parameters and observe the linked function graph views of
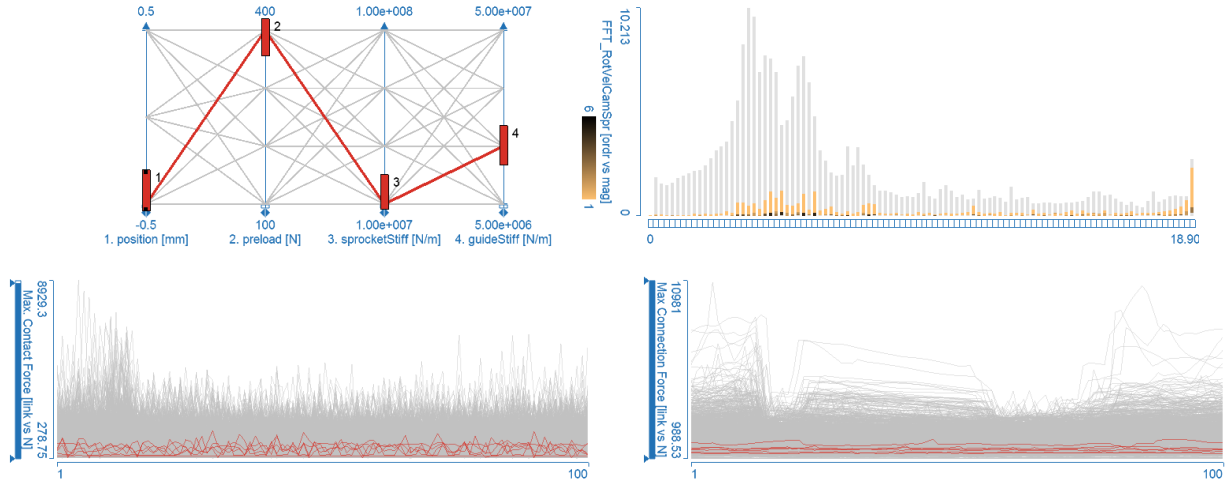
Figure 18: This snapshot captures the last iteration of the optimization process and shows the optimum design parameters. Top left: The design parameters are brushed in the parallel coordinates. Top right: The segmented curve view shows that the Fourier transform of the sprocket's rotational velocity has few peaks and they have small amplitude. Bottom left: Maximum contact forces are small at all chain links. Bottom right: All chain links have very small maximum connection forces.

connection force and contact force diagrams as well as a segmented curve view of the spectrum of the sprocket's rotational velocity to see its distribution. Figure 18 captures the final snapshot of this process.

Figure 14 shows that sprocket offset settings above −0.5 mm produce undesirably high connection forces so we brush −0.5 mm sprocket offset. Various preload parameters are then brushed and logical AND of the two brushes is shown in the linked views of the response parameters. The preload of 400 N is selected because the spectrum of the rotational velocity shows the smallest peaks in the segmented curve view, while the maximum contact and connection forces are kept in the lower range.

Identifying the optimal stiffness of sprockets and guides is the next goal. We achieve that by adding new brushes to the parallel coordinates for the varied parameters and moving them while observing the other three linked views. The smallest sprocket stiffness creates the smoothest contact force diagrams. However, while the lowest guide stiffness also produces low forces, the spectrum of the rotational velocity has a more expressed peak at the 7th order. The selected guide stiffness value in Figure 18 reduces that 7th order amplitude. The connection and contact forces are still relatively low and with little fluctuation which is desirable. Unfortunately, the peak at the 19th order in the spectrum is still there but its amplitude is reduced. As we can see in Figure 17 we cannot remove this peak completely with this range of design parameters.

Considering all these factors, we find the optimal set of design parameters that provides the best compromise for

noise, motion coupling accuracy and durability. Figure 18 shows the optimal parameters and the corresponding simulation response. Note that we use the alpha-blended curve view to display the contact and connection forces in this figure. The reason is that we are interested in the actual exact shapes of those graphs rather than their distributions since we narrowed down the focus to the optimum set. The alpha-blended curve view works better for this purpose than the segmented curve view. The optimal design parameters are summarized in Table 2.

Table 2: Optimum parameters

Parameter	Value	Unit
Sprocket stiffness	1.0E+7	N/m
Guide stiffness	2.0E+7	N/m
Chain preload	400	N
Sprocket offset	-0.5	mm

6 Conclusion

Data sets from many application domains can be represented by a data model consisting of independent and dependent variables that can either be scalars or function graphs. The analysis of distributions in families of function graphs is a common task in those application domains. We introduced a visualization method based on coordinated linked views and composite brushing that supports such interactive analysis. The novel segmented curve view can visualize fine details in the distribution while also preserving outlier graphs. The case study on the analysis of a timing

chain drive simulation data set demonstrates how the use of linked views and the novel segmented curve view can provide an insight into the analyzed data set that otherwise would be very difficult. The segmented curve view features (global/local binning, color mapping and brushing) proved to be extremely useful.

Acknowledgments

References

[1] John V. Carlis and Joseph A. Konstan. Interactive visualization of serial periodic data. In *UIST '98: Proceedings of the 11th annual ACM symposium on User interface software and technology*, pages 29–38, 1998.

[2] Jerding Dean F and John T. Stasko. The information mural: A technique for displaying and navigating large information spaces. *IEEE Transactions on Visualization and Computer Graphics*, 4(3):257–271, July-September 1998.

[3] Brian Francis and John Pritchard. Visualisation of historical events using Lexis pencils. Advisory Group on Computer Graphics, 1997. http://www.agocg.ac.uk/reports/visual/casestud/contents.htm.

[4] Harry Hochheiser and Ben Shneiderman. Dynamic query tools for time series data sets: timebox widgets for interactive exploration. *Information Visualization*, 3(1):1–18, 2004.

[5] Zoltán Konyha, Kresimir Matković, Denis Gračanin, Mario Jelović, and Helwig Hauser. Interactive visual analysis of families of function graphs. *IEEE Transaction on Visualization and Computer Graphics*, 12(6):1373–1385, November-December 2006.

[6] Zoltán Konyha, Krešimir Matković, and Helwig Hauser. Interactive 3D Visualization Of Rigid Body Systems. In *Proceedings IEEE Visualization 2003*, pages 539–546. IEEE Computer Society, 2003.

[7] AVL List GmbH. *EXCITE Timing Drive - Users Manual (Version 7.0)*. Graz, 2005.

[8] Krešimir Matković, Mario Jelović, Josip Jurić, Zoltán Konyha, and Denis Gracanin. Interactive Visual Analysis and Exploration of Injection Systems Simulations. In *Proc. of the IEEE Visualization 2005*, pages 391–398, 2005.

[9] Krešimir Matković, Josip Jurić, Zoltán Konyha, Jürgen Krasser, and Helwig Hauser. Interactive visual analysis of multi-parameter families of function graphs. In *CMV '05: Proc. of the Coordinated and Multiple Views in Exploratory Visualization (CMV'05)*, pages 54–62, 2005.

[10] Wolfgang Müller and Heidrun Schumann. Visualization for modeling and simulation: visualization methods for time-dependent data - an overview. In *WSC '03: Proceedings of the 35th conference on Winter simulation*, pages 737–745. Winter Simulation Conference, 2003.

[11] Matej Novotný and Helwig Hauser. Outlier-preserving focus+context visualization in parallel coordinates. *IEEE Transactions on Visualization and Computer Graphics*, 12(5):893–900, 2006.

[12] Kristin Potter. Methods for presenting statistical information: The box plot. *Visualization of Large and Unstructured Data Sets*, S-4:97–106, 2006.

[13] Bernard W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC, April 1986.

[14] Martin Sopouch, Wolfang Hellinger, and Hans Herwig Priebsch. Prediction of vibroacoustic excitation due to the timing chains of reciprocating engines. *Proceedings of the Institution of Mechanical Engineers, Part K: Journal of Multi-body Dynamics*, 217(3):225–240, 2003.

[15] Edward R. Tufte. *The Visual Display Of Quantitative Information*. Graphics Press, 1983.

[16] John W. Tukey. *Exploratory Data Analysis*. Addison-Wesley, Reading, MA, 1977.

[17] Jarke J. van Wijk and Edward R. van Selow. Cluster and calendar based visualization of time series data. In *INFOVIS '99: Proceedings of the 1999 IEEE Symposium on Information Visualization*, page 4, Washington, DC, USA, 1999. IEEE Computer Society.

[18] Colin Ware. *Information Visualization*. Morgan Kaufmann, 2000.

[19] Rainer Wegenkittl, Helwig Löffelmann, and Eduard Gröller. Visualizing the behavior of higher dimensional dynamical systems. In *Proceedings of the IEEE Visualization (VIS '97)*, pages 119–125, 20–21 October 1997.