Using Visual Analytics and Information Visualization to Investigate In-Car Communication Processes

Michael Sedlmair*
BMW Group Research and Technology

Andreas Butz†
University of Munich

ABSTRACT

Modern cars contain a wide spectrum of functionality, which is implemented by many interconnected electronic control units (ECUs). Overlooking all details of these increasingly complex in-car communication networks is a major challenge for developers. In our work, we have designed a number of analysis tools for in-car communication networks to enable developers to trace errors better and faster. By observing current working practices of automotive analysis experts, we found that the tools in use are mostly text-based and often fail to provide sufficient insight into correlations and overview aspects. They lack sophisticated visualization, navigation and data reduction techniques. Our research goal is to find novel and adapt existing methods of visual analytics (VA) and information visualization (InfoVis) to support the process of analyzing in-car communication networks. With a set of prototypes and their evaluation, we managed to provide concrete solutions and verify how in-car communication analysis can benefit form research in VA and InfoVis.

1 INTRODUCTION

During the last few years the functionality in automobiles has increased enormously. A large number of advanced functions, such as ACC (adaptive cruise control), rear seat entertainment systems or automatic start/stop engines were integrated into the car. They enabled step-by-step safer, more efficient and enjoyable driving. Thereby the functions became more and more distributed and several communication networks had to be installed. Current in-car communication networks consist of several bus systems (CAN, MOST, FlexRay, Lin, Ethernet) interconnecting a great number of electronic control units (ECUs) and distributing up to 15.000 messages per second. Analyzing the flow of messages in this network and tracking down problems has become a major challenge for automotive engineers. Current analysis tools are based on purely textual representations. In consequence, a lot of time and experience is needed to understand the processes and their complex correlations, to detect the sources of errors, and to take adequate action. By applying VA/InfoVis to this area we see a large potential to: (a) substantially increas working speed (i.e., find more errors in less time), (b) gain more insight and a better understanding, (c) discover novel aspects within the data and (d) provide different perspectives on the data. Our approach is based on two phases: First, it is necessary to get a clear understanding of the underlying data, the engineers' technical background, their current practices, tasks, problems, demands and challenges. This approach has been applied to other domains (cf. [1]), and it has become clear that only if these aspects are understood properly, effective visualization applications can be developed [7]. Second, based on these observations, prototypical VA/InfoVis applications have to be designed, built and evaluated. Some of our designs are adaptations of existing visualization techniques based on the findings from similar problem areas (e.g., [2]). Others are novel, domain-specific visualization methods with suitable interaction techniques.

2 DATA, TASKS, PRACTICES AND PROBLEMS

We conducted several empirical studies (user observations, workshops, focus groups, interviews and informal discussions) with our target group of automotive analysis experts. In addition, we analyzed existing data and tested current tools. In this process, we found that analysis experts mostly work with so-called *trace files*, a time-stamped log of recorded messages from the in-car network. Trace files contain the plain message content with exact timestamps, as well as information about the sending ECUs in hexadecimal form. They also contain special fields for both automatically detected errors (i.e., detected by an ECU), and manually detected errors. In the latter case, the driver detects a malfunction, presses a specific button, and annotates the marked entry in the trace later on. Considering that there are up to 15.000 messages per second, and that traces are recorded on weekend test runs, this leads to enormous files with up to 10GB of raw data.

Traces are currently handled, interpreted and analyzed with specific analysis tools (e.g., Canalyzer[1]). With the help of these tools, engineers try to locate the source of errors and subsequently inform the responsible developers to fix it. However, locating error sources is not trivial for several reasons: a) In most cases, erroneaous entries detected by an ECU are not the actual source of the error. b) There are multiple reasons for misbehavior, e.g., message displacement, timing problems, or hardware defects. c) Manually detected errors are enormously delayed within the trace.

An analysis expert's typical procedure to locate an error is as follows: After an initial attempt at getting a general overview, he or she starts locating points of interest in the trace and tries to get a clear understanding of their entire context. From this detailed information, they then try to understand the complete situation at a more abstract level and thereby deduce the error source.

However, this workflow is not properly supported by the current tools. Most tools are based on text and lists, and therefore fail to provide overview and correlation information. Navigation within the data is slow, and sophisticated data reduction and data mining techniques are absent. The tools often fail to provide multiple different perspectives of the data, and finally, collaboration – be it distributed or co-located – is not supported at all.

3 VISUAL ANALYTICS APPROACHES AND PROTOTYPES

We have built a number of different VA and InfoVis prototypes to investigate potential improvements to the tool chain for automotive engineers. All of these prototypes were then evaluated by the target group and partly led to the subsequent development of actual tools.

3.1 Providing Visual Overview and Navigation

The *AutobahnView* [5] (see fig. 1-a) is a visualization prototype to support overview and navigation in bus traces. It is (ironically) based on the metaphor of a crowded highway and represents all messages of a trace horizontally in chronological order. Each bus

---

*e-mail:michael.sedlmair@bmw.de
†e-mail:andreas.butz@ifi.lmu.de

[1] www.canalyzer.com

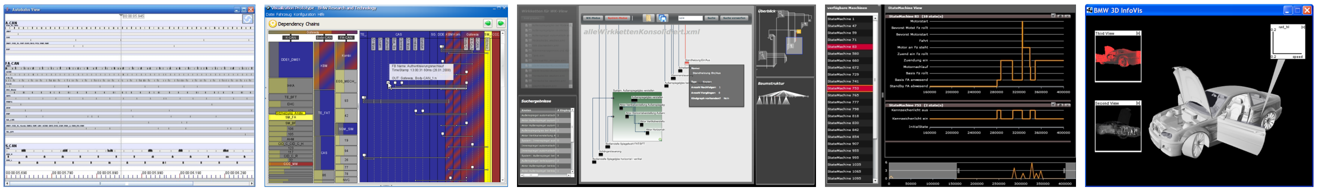

(a) Autobahn View (b) Dependency Chains 1 (c) Dependency Chains 2 (d) Statemachines (e) 3D Model Vis.

Figure 1: Screenshots of current prototypes. For more detail please see referenced papers

system corresponds to a separate group of lanes - the highway. Every bus line transports messages from different ECUs represented by a lane of the highway. The lanes in turn contain black rectangles – the cars – which each represent a message sent by the ECU through the bus to another ECU. The AutovbahnView provides zoom and pan interaction, allows fast searching and browsing in the data, and enables direct access to the raw data in an additional list view at various points. Qualitative user studies showed that the AutobahnView was usful for fast navigation, detection of message bursts and identification of repetitive patterns. Additionally, the close coupling of raw data and visualization was very useful for the analysis engineers.

3.2 Reducing the Data

A continuous problem in trace analysis is the flood of information with millions of messages, even in short traces. We propose two different approaches to substantially reduce the data and to provide simplified and condensed (visual) perspectives to the data.

By introducing the concept of so-called *dependency chains* we reduced the trace data and showed causal correlations in the data. On the one hand, this allows to exclude causal redundant information by condensing frequent messages which are not distracted by any other communication. On the other hand, network areas could be excluded by revealing that they were independent of a certain point of interest. Our prototype for visualizing this reduced information was based on a dual view approach (see fig. 1-b). One view presents a treemap-like approach showing physically dependent elements (ECUs). The second view is a specialization of a message sequence chart showing the reduced trace information on a functional level. Upon selecting a functional block, both views highlight the transitive hull of this functional block (yellow=selected element, blue=predecessors, red=successors) and eliminate unnecessary information [4].

A second prototype based on the concept of dependency chains uses hierarchically clustered, semantically zoomable node-link diagrams and represents all dependencies within an in-car communication network (see fig. 1-c). For this prototype we moved one step further. We completely skipped time dependent information and just represented causal dependencies between functional and physical elements. This helped the analysis experts to trace dependency and propagation paths without being overwhelmed by the time-dependent trace information.

Another approach to reducing traces are *state machines*. State machines (e.g., a Door state machine with the two states ”open” and ”closed”) are defined upfront by analysis experts and can be applied to the trace. Obviously, real world state machines are much more complex and are often dedicated to system functions and not to observable behavior. A timeline-based, zoomable visualization (cf. fig. 1-d) then allows the user to get an overview, browse, search and examine the behavior of and the correlations between the defined machines.

3.3 Relating Electronic and Mechanical Information

Finally, we also conducted several experiments with *3D model visualizations* to bridge the gap between electronic and mechanical information. Fig. 1-e[2] shows a configurable 3d model view that can freely be coupled and coordinated with other applications. The coordination can either be linking and brushing elements, e.g., hovering over a message and highlighting its sending ECU in the 3D model view, or semantic linking, e.g., navigating through time and replay ”real” behavior in the 3d model [6]. Qualitative user studies showed that especially the combination of 3D visualizations with abstract data representation can add value to the analysis process.

4 FUTURE WORK

In our future work, we plan to design and implement further prototypes in related automotive application areas, and refine the existing designs. In these further design cyvles, we will conduct more profound user studies to evaluate the benefits of our prototypes in real world scenarios. Our current studies were often restricted in time and transferability to daily practice. Reasons for this are the scarce availability and time restrictions of experts and in the detached character of our prototypes. Therefore, future iterations will be integrated tighter with the existing tools to allow users an easy and time-independent access to our prototypes. In doing so, we will study our concepts in a real world, day-to-day environment. This will allow more quantitative evaluations, such as close observation and interviews, but also provides a solid basis for quantitative studies of the tools in real use. In doing so, we might be able to contribute to the challenge of properly evaluating visual analytics and information visualization applications [3], especially in an industrial environment, such as BMW. On a more general level, we’d like to derive general guidelines and a more comprehensive understanding from our specific solutions. We’d like to outline how solutions could be integrated in a homogeneous environment and investigate what fundamental hardware and software setups will support this.

REFERENCES

[1] M. Büscher, J. O’Neill, and J. Rooksby. Designing for diagnosing: Introduction to the special issue on diagnostic work. *Comput. Supported Coop. Work*, 18(2-3):109–128, 2009.

[2] G. Namata, B. Staats, L. Getoor, and B. Shneiderman. A dual-view approach to interactive network visualization. In *Proceedings of the 16th Conference on information and knowledge management*, pages 939–942. ACM New York, NY, USA, 2007.

[3] C. Plaisant. The challenge of information visualization evaluation. In *Proceedings of AVI 2004*, pages 109–116. ACM New York, NY, USA, 2004.

[4] M. Sedlmair, W. Hintermaier, K. Stocker, T. Bring, and A. Butz. A dual-view visualization of in-car communication processes. In *Proceedings of the 12th International Conference on IV*, July 2008.

[5] M. Sedlmair, B. Kunze, W. Hintermaier, and A. Butz. User-centered Development of a Visual Exploration System for In-Car Communication. In *Proceedings of 9th International Symposium on Smart Graphics*, May 2009.

[6] M. Sedlmair, K. Ruhland, F. Hennecke, A. Butz, S. Bioletti, and C. O’Sullivan. Towards the Big Picture: Enriching 3D Models with Information Visualisation and Vice Versa. In *Proceedings of 9th International Symposium on Smart Graphics*, May 2009.

[7] J. van Wijk. Bridging the Gaps. *IEEE Computer Graphics and Applications*, pages 6–9, 2006.