fastvpn – Secure and Flexible Networking for Industry 4.0
fastvpn – Sichere und flexible Vernetzung für Industrie 4.0

Andreas Bluschke[1], Wolfgang Büschel[2], Michael Hohmuth[3], Frieder Jehring[1], Rainer Kaminski[1], Konstantin Klamka[2], Stefan Köpsell[4], Adam Lackorzynski[3], Tim Lackorzynski[4], Michael Matthews[1], Philipp Rietzsch[1], Alexander Senier[4], Peter Sieber[1], Volker Ulrich[1], Rainer Wiggers[5], Jean Wolter[3]

[1] Teleconnect GmbH, Am Lehmberg 54, 01157 Dresden, blua@teleconnect.de
[2] Interactive Media Lab, Faculty of Computer Science, Technische Universität Dresden, 01062 Dresden, first name.surname@tu-dresden.de
[3] Kernkonzept GmbH, Buchenstraße 16b, 01097 Dresden, info@kernkonzept.com
[4] Chair of Privacy and Data Security, Faculty of Computer Science, Technische Universität Dresden, 01062 Dresden, first name.surname@tu-dresden.de
[5] Vattenfall Europe Netcom GmbH, Köpenicker Str. 73, 10179 Berlin, rainer.wiggers@vattenfall.de

Abstract

This article presents the *fastvpn* project, which has been handled by a consortium of five partners since 2016 as a part of the fast cluster. The goal of *fastvpn* is the enhancement of existing industrial IT infrastructures with respect to IT security aspects addressing industrial SMEs, and taking into account the minimization of the additional network delay due to the IT security measures applied.

Kurzfassung

Im Beitrag wird das Projekt *fastvpn* vorgestellt, welches im Rahmen des fast-Clusters seit 2016 von einem Konsortium aus fünf Partnern bearbeitet wird. Das Ziel von *fastvpn* ist die Erweiterung bestehender industrieller IT-Infrastrukturen hinsichtlich IT-Sicherheitsaspekten für industrielle KMUs unter Berücksichtigung einer Minimierung der zusätzlichen Verzögerung aufgrund der eingesetzten IT-Sicherheitsmaßnahmen.

1 Introduction

Real-time is an enabler for many advanced and novel applications. The fast (fast actuators sensors & transceivers) cluster [1], [2] focuses on business segments where (a) real-time is a unique selling point, (b) the consortium is strong, and (c) promising markets are expected. Accordingly, the fast project focuses on four areas: (1) Connectivity, (2) Manufacturing, (3) Health and (4) Traffic. For these areas estimations indicate revenues in the three-digit billion € range in 2020.

fastvpn (VPN – Virtual Private Network) is one of the fast projects with the following partners: Teleconnect GmbH, Kernkonzept GmbH, Technische Universität Dresden Chair of Privacy & Data Security and Chair of Multimedia-Technology, and Vattenfall Europe Netcom GmbH as associated partner.

The goal of *fastvpn* is the enhancement of existing industrial IT infrastructures with respect to IT security aspects addressing industrial SMEs, and taking into account the minimization of the additional delay due to the IT security measures used. The project *fastvpn* addresses the core problems of Industry 4.0 systems: the reliable fulfilment of real-time requirements while at the same time considering increased security requirements and high data rates with optimal use of already existing infrastructure. It was not until the discussions of recent years about industrial espionage and sabotage (NSA revelations, Stuxnet, Wannacry ...) that awareness about previously neglected IT security has grown. Since IT security, especially in case of SMEs, is often still seen as a cost factor without immediate benefits, it was largely ignored by providers of large, integrated systems. At the same time, the pressure on SMEs to use increasingly powerful, IP-based data transmission systems (Industry 4.0, real-time) to modernize their existing legacy systems is growing. As a result, complexity and the associated management overhead increase. *fastvpn* addresses this problem at several levels.

Industry 4.0 requires a real-time, broadband and easily expandable data transmission infrastructure. This should, use the existing wired networks as far as possible, in the ideal case simultaneously with existing systems and without affecting them (“any media”). New infrastructure should only be installed if there are local shortcomings. This helps massively to reduce the total cost of ownership and the installation effort, and to preserve existing investments.

Furthermore, it is important to secure the data transmission. This can be done, for example, directly on the endpoints of the communication between the individual network subscribers (end-to-end security). However, any device in the network is, in principle, directly accessible, even if this is not necessary for the function of the entire system. At the same time, the devices involved must support secure communication protocols, which is often not the case due to the integration of legacy devices with different data transmission standards. We are looking for a solution that provides a comprehensive security concept for the entire heterogeneous network infrastructure. Our concept is intended to enable fine-grained access control with respect to individual devices and functions. Therefore, we cryptographically secure all data streams independently of the application layer protocols providing confidentiality and integrity. Our approach can be applied to local area networks (LANs) as well as more widely distributed networks (WANs).

The described usage of cryptography allows a logical separation of different traffic flows, even if a shared physical medium is used. As legacy machines will not support this, we introduce so-called *fastvpn*-nodes, which are controlled by a so-called *fastvpn*-gateway (see **Figure 1**). A *fastvpn*-node is responsible for applying the necessary cryptographic operations.

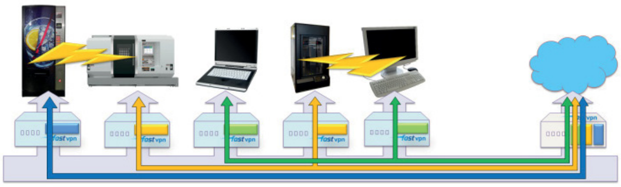

Figure 1 Example use case and basic idea: isolation/containment of different devices in a common network.

An important point here is that this security concept is based on a system that itself adds only a minimal additional potential attack surface. To ensure this, we implement an architecture that combines state-of-the-art cryptography with separation based on virtualization and microkernels.

The described security functions typically increase the latency and decrease the goodput of the overall system. It is the core goal of *fastvpn* to optimize latency and goodput through appropriate architectural decisions by selecting existing algorithms and hardware as well as proprietary developments so that the real-time requirements are met despite the strong security guarantees.

SMEs in particular often lack the necessary IT expertise to perform complex security-relevant configuration and maintenance work on their IT infrastructure effectively and efficiently. In order to promote the viability of the resulting solution, especially in the SME environment, it is the goal of *fastvpn* to develop suitable user interfaces, while taking into account aspects of usability from the outset so that the installation and maintenance costs for the secure networking can be reduced to a minimum.

We currently develop a comprehensive demonstration platform, which combines complex components of hardware and software as well as algorithms and usage concepts, which allows the validation of the technological fundamentals and their performance individually and in complex interaction.

At the 11. ITG-Fachkonferenz “Breitbandversorgung in Deutschland” in 2017 results about „any media“ [3] and security [4] were presented.

2 Problem Statement & Requirements

As already stated in the introduction, the enrichment of industrial production plants with powerful IT infrastructures becomes more and more important due to the demands and advantages related to Industry 4.0, predictive maintenance, production optimization etc. This pressure does not only affect large factories but even SMEs.

Looking at the industrial reality one can still find many situations where the necessary IT networks do not exist at all or are a patchwork of many different network/ communication technologies. One reason for this is that industrial machines usually have a very long lifetime – at least compared to the rather short innovation cycles, which we experience in the IT sector. Even if high-performance networks exist, they often do not integrate the necessary IT security measurements. Many studies have shown that IT security in the industrial context is still an open issue. Finally, even if security measurements are implemented, they often follow a concept of so called “protected automation cells”. The basic idea is, that the network traffic between cells is restricted, e.g. by using firewalls and other access control mechanisms – but within an automation cell basically no IT security measurements are applied. The underlying assumption is, that attacks mainly originate from the outside, while the own machines are treated as trustworthy. Given that these machines are produced by many different manufactures, which sometimes even have remote access due to maintenance reasons, we believe that the assumption of “all my machines are trustworthy” is too strong. On the other hand, we assume that the security risks associated with a malicious machine will increase within the next years due to the increasing demands on flexibility and digitalization (Industry 4.0). Imagine a malicious machine of an Industry 4.0 production plant, which spies on the network cable (shared medium) and therefore gets access to a huge number of different plans related to all the different goods, which are manufactured. We therefore want to develop a solution that offers security even in case of arbitrary malicious machines.

Besides securing the local network, our solution should also support remote entities. These entities could be either whole networks (e.g., connection of two production plants over the Internet) or single machines (e.g. secure access for remote maintenance).

As mentioned in the introduction, industrial machines have a rather long lifetime. Additionally, it is very often impossible to apply any hard- or software changes (due to terms of warranty, compliance, certification etc.). Therefore, we treat the machines as “black-boxes” which we cannot change but for which our solution has to maintain the necessary connectivity. Nevertheless we assume this connectivity is realized by means of (industrial) Ethernet, e.g. PROFINET or EtherCAT. Our hardware nodes are powerful and flexible enough to act as interface converter allowing, e.g., the transportation of serial communication over our security solution.

One important goal of our project is to support SMEs. While large companies often have an IT department, sometimes even a specialized security team, SMEs very often cannot afford to have dedicated IT personnel – not to mention security specialists. Therefore, our solution should be as easy to use as possible. Novel user interfaces and configuration technologies will support this.

Finally, we assume that providing a solution solely dedicated to network security might not be sufficient from a business / market point of view. Therefore, the software stack running on our *fastvpn*-nodes should allow the execution of third party applications. This applications run close to the machine and have access to the interfaces (field busses etc.) offered by a given machine. They could therefore execute many useful data processing tasks, like preprocessing of collected sensor data etc.

3 Architecture

The basic concept of our solution is to add transparent encryption to all communication between all communication participants (machines, local and remote computers etc.). This encryption will ensure confidentially and integrity of the communication. This concept is similar to the well-known approach of a virtual private network (VPN). But besides the usual solution, where exactly one VPN connects all participants, we span multiple VPNs over the same physical infrastructure. Thereby only the entities which need to communicate with each other (with respect to a given use case) participate in a given VPN. Assume, e.g., that we have five machines M1, …, M5 all connected over the same physical medium. Furthermore, M1, M2 and M3 need to communicate with each other in order to produce a certain product, while M1, M4, and M5 need to communicate in order to produce some other good. In this case, we would have two VPNs, one containing M1, M2, M3 and the other containing M1, M4, M5. Assuming that M2 is under the control of an attacker, the attacker will not be able to get access to the content of the communication of the second VPN.

As stated above, we cannot modify existing machines. Therefore, our solution is based on additional components (*fastvpn*-nodes) which are placed between the machines and the existing network. Each *fastvpn*-node, which comprises a self-developed hardware and a customized software stack, has two Ethernet interfaces: one is meant for the communication towards the machine and the other is meant for connecting to the existing network (LAN). All data from a given machine is encrypted by the *fastvpn*-node of that machine before it is sent over the existing network to the *fastvpn*-node of the receiving machine. The corresponding *fastvpn*-node will decrypt the data and forward it to the final destination (machine). A *fastvpn*-node prototype is shown in **Figure 2**.

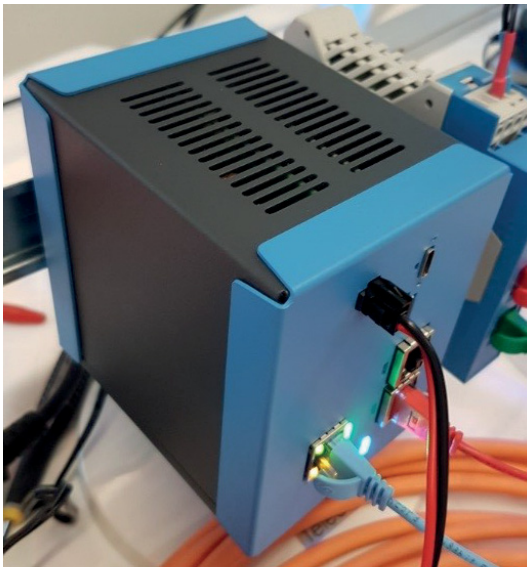

Figure 2 *fastvpn*-node (Box M).

To orchestrate and manage the *fastvpn*-nodes and the VPNs, we created an additional component, which we call *fastvpn*-gateway. *fastvpn*-nodes register with the *fastvpn*-gateway and receive control commands from the gateway[1]. The foundation of our software stack is a microkernel architecture. A microkernel is a minimalistic kernel, which only provides the bare minimum necessary for an operating system. Many of the components, which are usually part of a monolithic kernel, as device drivers, filesystems etc. are not part of a microkernel but run as services in user space. This leads to a very small (in terms of lines of code) trusted computing base (TCB). On top of such a microkernel processes are executed. A process can cover a whole operating system. In our case, we run a tailored version of Linux on top of our microkernel. In fact, we run multiple instances of Linux to separate the different software components even more (e.g., those responsible for the cryptographic operations from the ones responsible for packet forwarding or management functionality).

The different aspects of this architecture are shown in **Figure 3**, whereby the LAN is seen as the untrusted internal network, based on any available medium (Ethernet, wireless, G.hn, SHDSL). The *fastvpn* boxes spread a managed network as overlay on this LAN allowing communication between selected peers and offering services. These can be for example the VPN-service, which allows trusted communication between PC 1&3 and Productioncell 1&3. Besides this, service offload to a “Custom Server” is provided in this example. Even the management is a component, which can in principle be provided by any *fastvpn*-equipment.

[1] For understandability reasons we only mention one gateway, but we can have multiple gateways, each on responsible only for a subset of all available nodes.

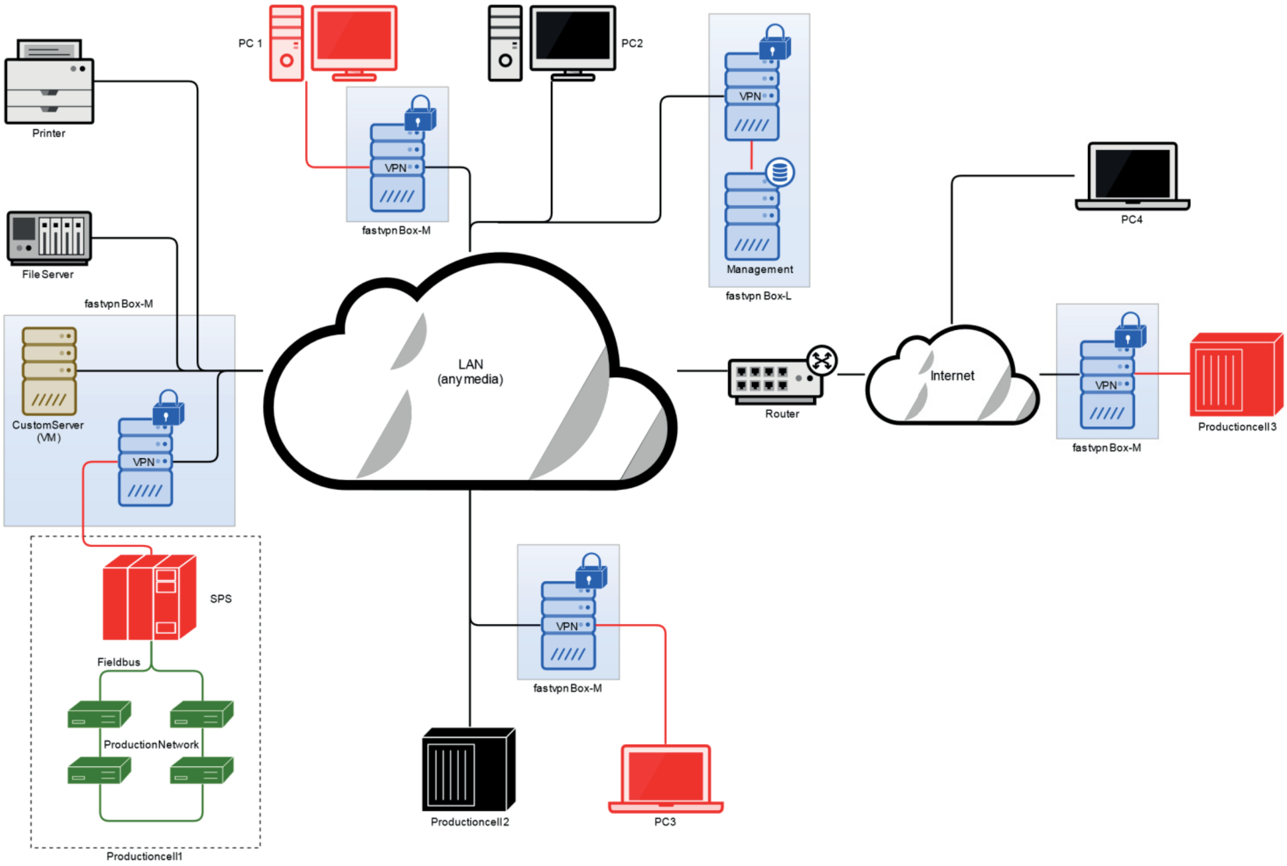

Figure 3 Generic *fastvpn* network structure.

Major part of the project besides the development of the hardware, the adaption of the microkernel operating system and the security solution is the management of the entire *fastvpn*-network.

This allows, in addition to the described VPN use-case, the possibility to embed even 3rd party applications, which are separated with the help of virtual machines. The idea is to allow the best possible isolation of these applications, which will enable completely new applications in the field of generic IT (provide different custom services), security (e.g., specialized solutions like firewalls or deep-packet-inspection) and industrial automation (e.g., SPS-as-a-service, predictive maintenance, automated data and pattern acquisition).

The different components involved in the management of the *fastvpn*-nodes are depicted in **Figure 4**. The NetworkManager is responsible for configuring the VPNs. Therefore, one part of this component is running on the gateway and another part is running on each *fastvpn*-node. The NetworkManager on the gateway receives higher-level commands like: “create a VPN comprising the following *fastvpn*-nodes”.

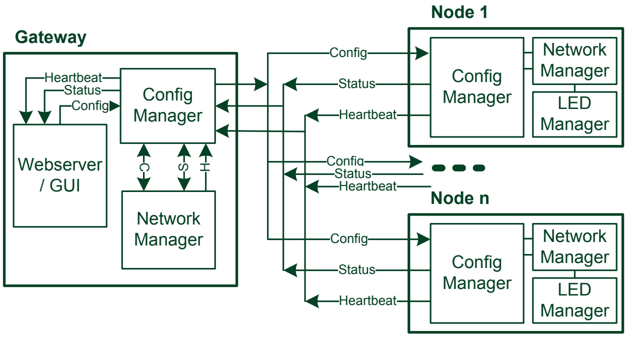

Figure 4 Interaction of management components running on the gate-way and the nodes.

The NetworkManager will translate this to a set of individual commands it will sent to the NetworkManagers running on the involved *fastvpn*-nodes. The NetworkManager on the gateway receives its commands from the ConfigManager. The ConfigManager on the one hand stores the current configuration persistent in a database and on the other hand, it is the only component, which exposes an interface to the “outside world”. This interface is utilized by a Web application, which provides a Web based GUI for the user.

The communication between all these components is done using ZeroMQ[2], a high-performance asynchronous and platform independent messaging library.

[2] http://zeromq.org

As stated above, we want to provide transparent encryption for local and wide area network communication. More specifically, we focus on Ethernet-based local and IP-based wide area communication. We selected the IEEE 802.1AE standard (MACsec) for securing the layer 2 (Ethernet) communications. Basically MACsec encrypts each Ethernet frame (confidentiality) and appends a mes-sage authentication tag (integrity). We utilize the MACsec implementation of the Linux kernel. The Linux kernel creates a virtual MACsec device for every MACsec Secure Connectivity Association. Such a MACsec Secure Connectivity Association corresponds to one of our VPNs. Therefore, we have one virtual MACsec device per VPN running on a *fastvpn*-node.

Although MACsec provides transparent layer 2 encryption in principle, there is a specific issue with respect to our requirements: being transparent implies that we can transport Ethernet frames (sent by machines) up to the maximum transmission unit (MTU). In standard Ethernet, this is 1500 bytes. But appending the authentication tag (and some additional MACsec header) will enlarge each Ethernet frame by 32 bytes. Therefore, we could not send such an “oversized” packet using the existing Ethernet. One solution could be to utilize so-called Jumbo Frames, which in fact allow sending Ethernet frames with a larger MTU. But Jumbo Frames are not a standard feature and we cannot assume that every deployed Ethernet in our industrial setting will support them. Another approach would be to tell the machines to only send Ethernet frames with a smaller MTU. In fact there exist techniques like Path MTU Discovery (PMTUD) to allow this. But not all layer 3 protocols actually support PMTUD. Especially in our industrial context, we must even assume that machine manufacturers utilize proprietary layer 3 protocols, which simply require a standard MTU of 1500 bytes. Therefore, we finally implemented a third approach: we split large Ethernet frames into two MACsec frames and reassemble them on the receiving side.

We use Open vSwitch[3] to manage the packet flows on the *fastvpn*-node. Open vSwitch is a virtual (software) switch, which is part of the Linux kernel. With the help of flow rules one can configure how packets are forwarded, e.g. which device (port) should be used. We use this functionality to ensure that packets dedicated for a given remote machine are send over the corresponding MACsec device (which in turn corresponds to a given VPN as described above). One advantage of Open vSwitch is that it supports the OpenFlow protocol. Although at the moment (to keep the prototype development simple) we use our own protocol to distribute the configuration information related to VPNs to the involved *fastvpn*-nodes, we plan to investigate how our *fastvpn*-nodes can be integrated in existing Software-defined networking (SDN) controllers, which speak OpenFlow.

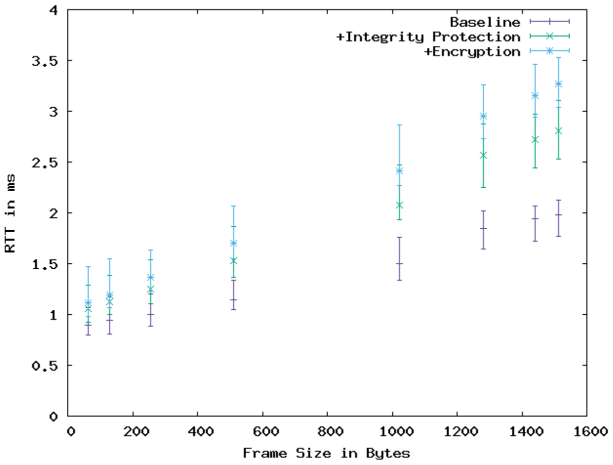

Figure 5 Measurements of round trip time (RTT). Shown are the minimum/average/maximum values of 300 runs for three different settings: without any encryption (baseline), with integrity protection, and with encryption and integrity protection.

We have done some initial performance analysis of our solution using Raspberry Pi3 computers instead of our final hardware, which was not ready during that time. We measured the roundtrip time of packets send between two Raspberry Pi3 computers, which were connected over Ethernet. For our measurements, we issued a simple ping command with different packet sizes. Although the Raspberry Pi3 is less performant than our *fastvpn*-node hardware, we already got some first promising results (see **Figure 5**). On the one hand, one can clearly see that there is an overhead in terms of increased latency due to the applied security mechanisms. On the other hand, the increase in latency is less than a factor of two. Especially for small sized packets, the additional latency is quite low. These packets are of special importance because in industrial communication most packets are small.

We are currently investigating how we can extend our solution beyond the borders of a local network. We therefore do performance analysis of different network security protocols, which work on layer 3. This includes IPsec, OpenVPN, WireGuard and tunneling MACsec e.g., using the Generic Routing Encapsulation (GRE).

4 Novel User Interfaces for the Maintenance and Configuration

SMEs often do not have the financial resources to employ full-time security experts to handle complex security-relevant configuration and maintenance work on their IT infrastructure. The ease of use is therefore an important key requirement for the success of secure and flexible networking solutions for many SMEs. Therefore, we investigate how traditional workflows (e.g., paper forms), currently available mobile devices as well as emerging technologies

[3] www.openvswitch.org

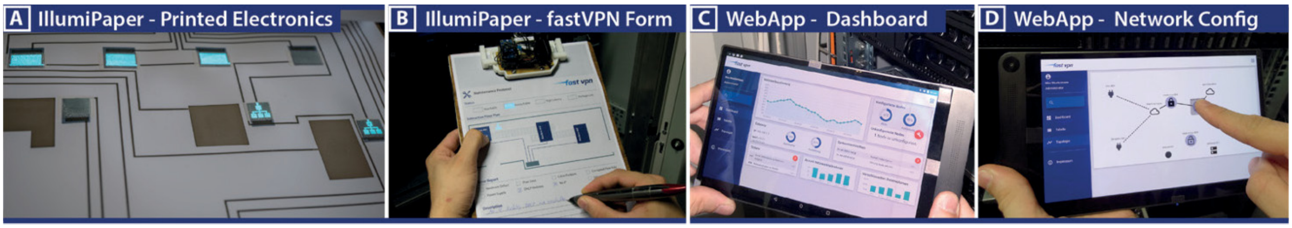
A IllumiPaper - Printed Electronics
B IllumiPaper - fastVPN Form
C WebApp - Dashboard
D WebApp - Network Config

Figure 6 Illuminated Interactive Paper based on printed electronics (**A**) allows digitally enhanced form filling applications for maintenance tasks (**B**). Responsive web application to configure IT infrastructure (**C+D**).

(e.g., AR glasses or smart E-Textiles) can be used to support maintenance and configuration tasks.

4.1 Digital-Enhanced Paper

Due to their unique qualities and affordances, physical paper forms are still widespread in many companies and constitute an important basis to handle maintenance tasks. To digitize paper notes seamlessly, pen-and-paper solutions have been introduced that are able to digitally capture and convert handwriting on standard forms. Unfortunately, such systems lack dynamic visual feedback and thereby restrain advanced digital functionalities. In order to digitally enhance paper applications, we developed a research platform called IllumiPaper [5] that provides new forms of paper-integrated visual feedback and enables multiple input channels (e.g., pen, touch, bend gestures). For the realization of our IllumiPaper sheets, we use novel printed electronics, such as conductive inkjet-printing [6], and thin-film display technologies like screen-printed electroluminescence displays [7] (see **Figure 6**, **A**). Based on our IllumiPaper system, we fabricate an industrial form application to support network maintenance tasks and reports [8]. Therefore, the upper region of a maintenance form gives an overview of the factory's floor plan with all machines that the engineer is responsible for (see **Figure 6**, **B**). Illuminated pictograms highlight devices for which issues have been detected. By simply touching or pen tapping a machine on the overview map, further information on the type of error is shown. This is an example of the underlying feedback design of a smart request that allows a user to immediately get current status information of the associated machine in real time. In this way, we support the maintenance worker by visualizing the network status dynamically. After checking the device, a maintenance engineer can write his or her report on the paper. The handwritten text is digitized using Anoto technology. After completing the report, a digital copy of the form can be sent via e-mail or synchronized with a database, e.g., to inform an expert about a specific hardware problem.

4.2 Responsive Web Application

In addition, we implemented a web application using state-of-the-art web technologies to support maintenance engineers in the configuration and maintenance of their IT infrastructure. To ensure a high degree of flexibility, our responsive web application is designed to be accessible for mobile contexts, e.g., in factory halls with mobile devices, as well as for stationary applications, such as desktop working places. For our design, we aim to reduce complexity and increase usability by providing familiar workflows, interactions and easy-to-understand views. Therefore, we first introduce a dashboard for quick overview of the current system states, such as current workloads, registered nodes, tickets or other problems (see **Figure 6**, **C**). Further, we developed an interactive network graph view that allows an IT engineer to visually set up and easily maintain the network system (see **Figure 6**, **D**). In this graph view, an administrator can add or remove both machines and *fastvpn*-nodes to create a secured VPN. In addition, it is possible to connect machines or nodes by simply linking two entities with each other. For touch devices, this can be done by activating the connection mode and then tapping both entities in sequence. Further, we implemented an additional drag-and-drop technique for desktop systems. Every defined VPN connection has its own color and can be activated or deactivated at any time. Furthermore, we distinguish between an edit mode and a live mode in our network graph view. While the live view presents further real-time status information, the edit view aims to avoid accidental misconfiguration by requiring an explicit mode switch.

4.3 Spatially-aware Mobile Devices and a Display Wall for Maintenance Tasks

The visualization of comprehensive network infrastructures on mobile devices and traditional desktop systems can be very challenging since their display sizes are limited. To address this issue, one well-known approach in administration and control rooms is the usage of large display walls that are able to show complex network infrastructures, machine states, and possible problems at the same time. The interaction with such display walls, however, is yet limited since they are mostly used as central displays that are only controlled by a single traditional desktop working place. We think that the future of Industry 4.0 application will be characterized by a rich set of new interaction techniques that work together in concert, simplify collaboration, and seamlessly support cross-device interactions. Therefore, we investigate how spatially-aware mobile displays and a large display wall can be coupled to support graph visualization and interaction (see **Figure 7**, **A**).

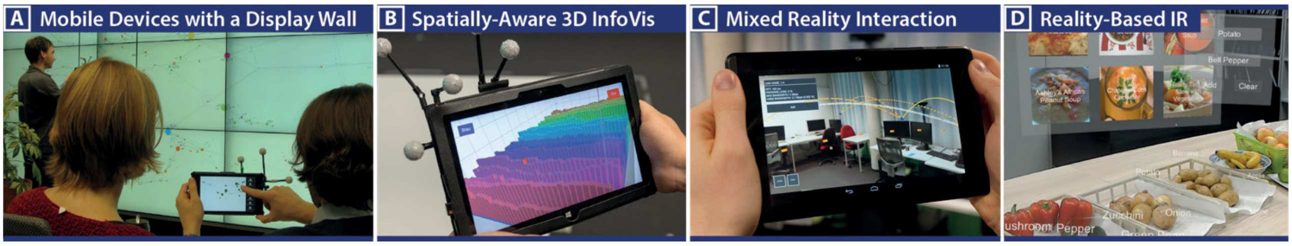
A Mobile Devices with a Display Wall
B Spatially-Aware 3D InfoVis
C Mixed Reality Interaction
D Reality-Based IR

Figure 7 Combining spatially-aware mobile devices and a display wall for infrastructure network exploration (**A**), spatial 3D information visualization using handheld mobile devices (**B**), Mixed Reality network visualization on a tablet (**C**), Reality-Based Information Retrieval supporting just-in-time retrieval and in-situ result visualization (**D**).

For that purpose, we distribute typical visualization views of classic node-link and matrix representations between the displays [9]. The focus of our work lies in novel interaction techniques that allow engineers to work with mobile devices in combination with the wall-sized control room displays. Therefore, we devised and implemented a comprehensive interaction repertoire that supports basic and advanced graph exploration and manipulation tasks, including selection, details-on-demand, focus transitions, interactive lenses, and data editing that could be used for network maintenance tasks. A qualitative study has been conducted to identify strengths and weaknesses of our techniques. Feedback showed that combining mobile devices and a wall-sized display is useful for diverse graph-related tasks. We also gained valuable insights regarding the distribution of visualization views and interactive tools among the combined displays.

4.4 Mixed Reality Interfaces

Today's AR headsets and current generation smartphones begin to support affordable Augmented Reality applications. Specifically, they allow in-situ visualizations, which have the potential to facilitate ad-hoc analysis processes [10], [11].

While these 3D visualizations typically suffer from general problems of 3D representations, such as occlusion, misleading perspective, and readability, they are useful when dealing with data that has a strong connection to physical locations, e.g., data linked to objects. Furthermore, they support spatial interaction, a form of mobile input that uses the movement of handheld, spatially-aware mobile devices. To gain a better understanding of this input modality for in-situ visualization, we conducted a user study [12] in which we compared spatial interaction with classic touch input (see **Figure** 7, **B**). In this study, we were able to show that spatial interaction is easy to understand, requires little training, and improves performance in navigation tasks for 3D visualizations.

One example for this class of visualizations, particularly interesting in the context of our project, are 3D node-link diagrams of computer networks.

We built an interactive prototype (see **Figure** 7, **C**) to illustrate the management and maintenance of small-scale local area networks, typically found in small and medium-sized businesses. Virtual nodes are placed at the physical, real-world location of network devices. Connections between these devices can then be shown as spatially located, augmented reality node-link diagrams. In an initial study [13], we examined different link attribute encodings to get an overview of how to visualize typical network parameters such as link health, bandwidth usage, or VPN status.

Furthermore, we were also interested how Augmented Reality can support Information Retrieval processes. We coined the term Reality-Based Information Retrieval (RBIR) [14] to describe a new class of IR interfaces that make use of the environment to satisfy information needs based on real-world stimuli. We developed a series of prototypes to highlight RBIR use cases such as searching for recipes on a farmer's market or querying an image database (see **Figure** 7, **D**). Users build their search queries with the help of labels extracted from the environment via computer vision techniques. For example, in the recipe search prototype, fruits and vegetables are labeled and a user can select several of them to find recipes containing the selected items. This general concept can also be applied to the management of computer networks, e.g., by allowing a maintenance engineer to access status information from network devices with the help of an AR interface or to generate service reports for a selection of devices in the vicinity.

Finally, we aim to investigate new interaction techniques that are well suited to fulfil the increasing needs of unobtrusive and powerful mobile controls. Therefore, we explore the promising combination of garment-integrated E-Textile cords [15] and the dynamic visual output capabilities of emerging AR glasses [16]. This novel combination of simultaneous input and output on a cord has the potential to create rich AR user interfaces that seamlessly support direct interaction and reduce cognitive burden by providing visual and tactile feedback. Therefore, we propose a set of cord-based interaction techniques for browsing menus, selecting items, adjusting continuous values & ranges and solving advanced tasks in AR. For example, a maintenance engineer with AR glasses could use our techniques to carry out maintenance work with functional clothing in order to get quick access to machine status information. To evaluate our approach, we realized a series of touch-enabled cords, its data transmission, and a basic AR visualization. Finally, we discussed the future role of visual feedback regarding its social accessibility for unobtrusive E-Textile interfaces along the dimensions of its type, position, time and visibility [17].

5 Conclusion

As mentioned above, Industry 4.0 requires a real-time, secure, broadband, and easily expandable data transmission infrastructure. As a result of our work during the last two years it was shown that such an infrastructure can be realized based on usage of key technologies, for instance, MACsec for security, G.hn for “any media”, microkernel for virtualization, ZeroMQ for configuration and management in combination with novel user interfaces.

6 Acknowledgement

The authors acknowledge the Bundesministerium für Bildung und Forschung for financial support within the framework of the zwanzig20 program.

zwanzig20
PARTNERSCHAFT FÜR INNOVATION

GEFÖRDERT VOM
Bundesministerium
für Bildung
und Forschung

7 Literature

[1] Ellinger, F., et. al.: Project FAST - fast actuators sensors & transceivers. IEEE MTT-S Latin America Microwave Conference (LAMC 2016), Puerto Vallarta, Mexico 12-14 December 2016; ISBN (Online) 978-1-5090-4287-6

[2] https://de.fast-zwanzig20.de/

[3] Bluschke, A.; Matthews, M.; Rietzsch, P.; Schäfer, A.: G.hn – offroad. 11. ITG-Fachkonferenz “Breitbandversorgung in Deutschland”. 29.-30.03.2017, Berlin. ITG-Fachbericht 270, pp. 61-65

[4] Lackorzynski, T.; Köpsell, St.: “Hello Barbie” – Hacker Toys in a World of Linked Devices. 11. ITG-Fachkonferenz “Breitbandversorgung in Deutschland”. 29.-30.03.2017, Berlin. ITG-Fachbericht 270, pp. 41-47

[5] Klamka, K.; Dachselt, R.: IllumiPaper: Illuminated interactive paper. In Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems. ACM, pp. 5605-5618

[6] Kawahara, Y., Hodges, S., Cook, B. S., Zhang, C., and Abowd, G. D.: Instant inkjet circuits: Lab-based inkjet printing to support rapid prototyping of ubicomp devices. In Proc. of UbiComp ’13, ACM (New York, NY, USA, 2013), pp. 363-372

[7] Olberding, S., Wessely, M., and Steimle, J.: Printscreen: Fabricating highly customizable thin-film touch-displays. In Proc. of UIST ’14, ACM (New York, NY, USA, 2014), pp. 281-290

[8] Klamka, K.; Büschel, W.; Raimund Dachselt, R.: Illuminated Interactive Paper with Multiple Input Modalities for Form Filling Applications. In Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces. ACM, pp. 434-437

[9] Kister, U.; Klamka, K.; Tominski, C.; Dachselt, R.: GraSp: Combining Spatially-aware Mobile Devices and a Display Wall for Graph Visualization and Interaction. Computer Graphics Forum 36, 3 (2017)

[10] ElSayed, N.; Thomas, B.; Marriott, K.; Piantadosi, J.; Smith, R.: Situated Analytics. In Proceedings of Big Data Visual Analytics (BDVA), Hobart, TAS, 2015, pp. 1-8

[11] Ens, B.; Irani, P.: Spatial Analytic Interfaces: Spatial User Interfaces for In Situ Visual Analytics. IEEE Comput. Graph. Appl. 37, 2 (March 2017), pp. 66-79

[12] Büschel, W.; Reipschläger, P.; Langner, R.; Dachselt, R.: Investigating the Use of Spatial Interaction for 3D Data Visualization on Mobile Devices. In Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces. ISS '17, Brighton, United Kingdom. ACM, pp. 62-71

[13] Büschel, W.; Vogt, S.; Dachselt, R.: Investigating Link Attributes of Graph Visualizations in Mobile Augmented Reality. In Proceedings of the CHI 2018 Workshop on Data Visualization on Mobile Devices. MobileVis ’18, Montreal, QC, Canada

[14] Büschel, W.; Mitschick, A.; Dachselt, R.: Here and Now: Reality-Based Information Retrieval. In Proceedings of the Conference on Human Information Interaction and Retrieval. CHIIR '18. New Brunswick, USA, 2018

[15] Schwarz, J.; Harrison, C.; Hudson, S.; and Mankoff, J.: Cord Input: An Intuitive, High-accuracy, Multi-degree-of-freedom Input Method for Mobile Devices. In Proceedings. of CHI ’10. ACM, New York, NY, USA, pp. 1657-1660

[16] Klamka, K.; Dachselt, R.: ARCord: Visually Augmented Interactive Cords for Mobile Interaction. In Proceedings of the 2018 CHI Conference Extended Abstracts on Human Factors in Computing Systems (CHI EA ’18), ACM

[17] Klamka, K.; Dachselt, R.: The Future Role of Visual Feedback for Unobtrusive E-Textile Interfaces. ACM CHI 2018 Workshop: (Un)Acceptable!?! – Re-thinking the Social Acceptability of Emerging Technologies