Toward a Web of Systems

Web and semantic technologies will form the foundation for ecosystems of machines that interact with each other and with people as never before.

By *Florian Michahelles and Simon Mayer*

DOI: 10.1145/2845163

The concept of the Internet reaching out into the real world has been around for more than a decade. This development started by providing real-world things with a machine-readable identifier, such as an RFID tag or a barcode, and associating digital content with these IDs. Driven by the technological advances in processing, sensing, and communication technologies and their decreasing cost, physical things have become ever more interconnected and pervasive, allowing them to collect real-time data and take control of other connected devices that affect the real world.

On top of this development, we can identify an emerging trend to develop solutions based on physical devices and interconnected digital services both in the consumer and industrial

domains. This will result in connected value chains of energy systems, public infrastructure, automated manufacturing lines, and healthcare solutions. To leverage the potential borne by this development, every connected system and subsystem needs to be capable to make sense out of the data it receives, processes and shares.

This “digitalization”—a term that has become popular to describe the process of moving to a digital business—is leading to an ever-greater use of information and communication technologies. Thus, application areas are starting to overlap and interact. For example, smart grids can interact with decentralized energy genera-

tors—factories, buildings, electric vehicles, and the urban infrastructure—and consumers. This could help to flatten the electricity demand curve and match renewable energy sources with demand peaks. Previously, each of these subsystems had its own norms and standards. Now, their products and solutions must be

Figure 1. Semantic core components of our approach. Our “Web of Systems Semantic Framework” enables horizontal interoperability while pluggable knowledge packs are responsible for linking concepts vertically within a domain. We propose to base as much of the machine-readable information as possible on readily available human-readable standards documents.

mutually compatible, across different domains. However, norms and standards of many different and previously independent application areas must now become integrated with each other and work together seamlessly, across vertical domains.

Vertical industry suppliers, like Siemens, have a strong tradition of commitment to standardization in their vertical businesses, such as industrial automation, energy generation and transmission, building control technology, and mobility and medical technology. However, industry needs to actively tackle the described development toward applications that act across verticals. As an example, in the electric mobility domain, electric vehicles, charging stations, storage batteries, urban infrastructure, and power networks have to exchange information about battery charge state, power availability, and energy system stability in real time. To perform together in a meaningful way, these subsystems need to truly collaborate. This collaboration goes beyond “syntactic” agreement about exchanged messages. Instead, a common understanding of meaning of shared information will be established on a semantic level. Only then will the system be capable of decentralized coordination tasks such as using idle vehicle batteries for leveling peak loads and stabilizing power grids.

EMPOWERING MACHINES TO SHARE MEANING

As machines will not only communicate within, but also across domains, new mechanisms of describing meaning to machines have to be developed. Unlike the human-created content on the web, much of the terminology and many processes in industrial domains already are partially structured thanks to standards, norms, agreements, and process descriptions.

Therefore, when aiming to interconnect systems vertically within a field and horizontally across domains, one main research question for industry is: How to translate domain-specific information represented in human-readable form into machine-readable shared understanding?

To approach this problem, we propose to build upon established technologies from the web and Semantic Web domains: We add meaning to machine-to-machine communication by establishing an ontology of interlinked terms, entities, and relationships. For example, this ontology would contain a machine-readable definition of what an “electric generator” is and how it relates to the unit “volts.” We believe this approach is viable in many domains that are relevant for industry since these terms, entities, and relationships can be based on existing domain-specific standards. Doing this will allow for communication across

hierarchies (vertically) and domains (horizontally), and yields a paradigm that we refer to as the “Web of Systems.”

USABLE SEMANTICS FOR INTEROPERABILITY

We propose to use semantic technologies to add shared meaning to information that is exchanged between components of a system and between systems across domains. The basic technologies to codify meaning in an appropriate way have been available for several decades. They have been applied to the World Wide Web in the context of the Semantic Web, but have yet to achieve breakthrough outside isolated application domains. We believe we can utilize semantic technologies to connect agents within and across several domains in a pragmatic way. Instead of “inventing” concepts and models for application domains, we propose to translate existing domain standards into machine-readable representations. Furthermore, although we are working with technologies that require huge up-front investments for large future payoffs, we are attempting to apply an iterative development process where we already demonstrate (limited) added value early on. In the rest of this section, we will first introduce the core components of our approach from a more technical perspective, and then give a few examples of how these components are used within several of our current projects.

CORE COMPONENTS

From a technical perspective, our proposed approach rests on three core components: First, we propose a method of facilitating the interconnection of heterogeneous devices and services that builds on the emerging activity streams format. Second, using a common semantic framework, we enable these distributed agents to share a common understanding of the world, thus enabling them to “speak the same language.” Third, to ensure the created systems are accessible for human users, we propose the use of augmented reality (AR) technologies to elicit semantic relationships and dependencies in the real world.

Activity streams for bridging Internet of Things (IoT) silos. Activity streams (AS) originated in the social web domain as a simple format that

allows to link walled gardens such as Facebook or Twitter. It makes information about (user) activities, which take place in the scope of one of those platforms, available to outside partners. We propose to take this one step further, using AS as the foundation for a connective fabric IoT silos. To test this approach, we created ASbase, a custom-built AS broker that allows clients to easily implement distributed AS-based scenarios [1]. The main purpose of this platform is to collect activities in the AS format and distribute them to interested parties, where it supports both a request-response and publish-subscribe pattern. For instance, clients can subscribe to streams that concern a specific patient in a healthcare setting, and also select the concrete data streams they are interested in (such as heart-rate data). To enable this, ASbase implements a filtering language that is based on the query language of the widespread MongoDB NoSQL database system. The openness and extensibility of AS allow users of the format to increase the richness of concepts that can be represented by them. When describing several of our current projects in the next section, we present a scenario where AS are used for integrating heterogeneous devices that collect health data, analysis services, and a visualization tool for health professionals within a single vertical silo. Furthermore we showed their applica-

As machines will not only communicate within, but also across domains, new mechanisms of describing meaning to machines have to be developed.

bility to bridging such vertical silos in an experiment where multiple project groups used AS to interconnect initially independent implementations and form cross-vertical mashups [1].

Semantic framework. After laying the foundation for seamless machine-to-machine (M2M) communication using AS, the next crucial step is to make this communication meaningful by enabling distributed agents to share a common understanding about exchanged information. We propose to tackle this problem with technologies known from the Semantic Web domain that are applied in a pragmatic way to real-world problems. In particular, we propose the creation of a set of core ontologies, referred to as the “Web of Systems Semantic Framework” (WSF) that captures cross-do-

main concepts such as information about units and dimensions (see Figure 1a). The WSF also contains information about other abstract concepts that are reusable across domains, such as machine-readable definitions of what constitutes a problem, how such problems relate to states and so forth. For usage in concrete applications, the WSF is extended with knowledge packs (KPs) that encode domain-specific information (see Figure 1b). Thus, the domain-specific KPs enable vertical interoperability between agents within a domain (for instance, an electric car and a charging station; see Figure 1c), and their integration with the WSF ensures horizontal interoperability across domains (see Figure 1d). To facilitate the transition to this system and mitigate interoperability problems, we propose to base WSF and KP concepts on agreed-upon industrial standards. In these cases, it is thus sufficient to translate standards documents into a machine-readable language, rather than inventing new concepts on a clean slate (see Figure 1e).

AR to uncover hidden relationships. One crucial factor to support the adoption of semantic technologies by subject matter experts is to make ontologies more tangible, and maybe even allow (limited) changes to the employed domain models at run time. We believe this could be achieved by using AR sys-

Figure 2: Visualization of wearable health sensor data. Our visualization dashboard displays data streams from heterogeneous wearable health devices in a way that enables doctors to quickly gain an overview of the patient’s health-related activities.

tems that overlay semantic information on real-world scenes and thereby enable humans to directly observe which background information a system is processing. Beyond displaying such rather static information, the usability of semantics-driven systems could be increased even further by showing runtime data on the AR overlay, such as communication between devices.

CURRENT PROJECTS

We will demonstrate the value and power of our approach in the context of a selection of three projects: the healthcare, smart grid, and industrial automation domains.

HealthViz: Professional visualization of wearable health sensor data.

One concrete example of AS and an ASbase platform is the HealthViz project, in which we investigate the integration of wearable sensors from the consumer wellness domain into a professional health IT platform [2]. In this scenario, individuals use wearables or services on portable devices (such as activity tracker applications on smartphones) that allow them to collect information about their level of exercise and other physical activities. Health coaches or doctors can then use this information to better assess a patient's lifestyle, and to monitor and suggest lifestyle changes. Here, AS and ASbase are used to decouple the data collection from concrete devices and services from

the use of the collected data in further stages of the data processing pipeline, such as during aggregation and analysis, and also from the final consumption by user interfaces that are targeted at health coaches and doctors.

Our visualization interface (see Figure 2) consumes AS activities and data representations that are linked from them and displays them in a dashboard-style way, which gives the health coach or doctor a quick overview of the most relevant analyses of the recent past. If they choose to do so, they can drill down into the data, in which case more detailed representations will be accessed and visualized. The ASbase subscription API can also be used to have the interface dynamically update itself while being used by the health professional.

Intelligent secondary substation for smart grids. While the HealthViz system focuses strongly on the classical integration of heterogeneous services in an IoT context, we demonstrate the value of our approach of combining a core WSF with domain-specific KP in the smart grid domain. Here, we support the creation of an intelligent secondary substation prototype that comes with the capability of deploying additional functions at run time via an “app store.” Our primary contribution to this project is a domain-specific KP that describes which applications can be deployed on a secondary substation together with metadata such as application depen-

dencies. For instance, a voltage regulation application requires access to voltage data streams. This enables substations to semantically resolve such dependencies and automatically install required additional applications. Using cross-domain knowledge from the WSF, our smart grid KP furthermore relates application functionality to potential grid problems—for instance, the KP encodes the fact that a voltage regulation application resolves voltage band violation problems. Together, the WSF and the domain-specific KP thus deliver a semantic layer that greatly facilitates the operation of a secondary substation, and even holds the potential to fully automate the resolution of common issues in electric grids.

Goal-driven manufacturing. In industrial manufacturing, our approach of combining semantic technologies with seamless device communication is valuable with respect to a novel way of controlling manufacturing devices. In contrast to process-driven approaches where relationships between manufacturing devices (such as a welding machine and a robot arm) are statically defined, we make use of embedded semantic descriptions of device functionality to dynamically create mashups that fulfill the production's specified goal (such as “produce a red car door,” see Figure 3). The main advantage of our system is its high degree of flexibility, as service mashups can adapt to dynamic environments. Additionally, they are fault-tolerant with respect to individual devices becoming unavailable, for instance because they undergo maintenance [3]. The goal-driven control of manufacturing devices is particularly valuable to reduce machine tooling times that are an important factor especially when producing small batch sizes. Our approach holds the potential to have manufacturing lines reconfigure themselves at runtime, based on descriptions of the functionality of individual devices, and even considering “non-functional” properties that influence the process indirectly, such as the required time or monetary cost of a process.

CHALLENGES

By establishing a Web of Systems, we can enable machines to collaborate

Figure 3. Based on a production goal [a] and functional descriptions of individual devices [b], a Reasoner [c] derives a production plan that can immediately be implemented, [d] and takes into account dynamic context factors such as individual devices being unavailable.

horizontally within domains as well as vertically across domains. Web technologies and semantic technologies are key enablers for this development. However, when implementing this paradigm on a larger scale, a variety of challenges have to be solved.

Engineering: Supporting the efficient creation of ontologies from standard documents. One challenge lies in facilitating the process of “translating” today’s merely document-based human-readable to machine-readable ontologies. In order to deploy semantics on a larger scale, tools and methods have to be developed that overcome the limits of a manual translation process and can support domain experts in the codification of knowledge for its usage by computer systems.

Lifecycle management: Accounting for upgrades and extensions. Once a domain-specific ontology has been created, it has to be kept up to date to changes and updates. Thus, efficient ways of updating and extending ontologies, as well as guaranteeing the consistency of distributed knowledge bases, are crucial to drive adoption.

Virtual affordances: Informing users about virtual cause-effect relationships. Humans are adept to spotting cues that convey how physical and virtual objects should be used, for instance whether a door can be opened by pushing or pulling. However, as physical devices get more connected and become parts of virtual systems, human actions will trigger side effects across virtual and physical spaces. We therefore require new and intuitive ways to inform users about the effects of their actions and the virtual, hidden, causal relationships. Perhaps, technologies such as AR are well suited to tackle this challenge.

Provenance: Tracing causes and their effects in the Web of Systems. On a more abstract level, we are very interested in tracing cause-effect relationships ex post to gain a deeper understanding of a specific system. For instance, this could allow tracing fluctuating plant back to specific machines operating below capacity or undergoing maintenance. Beyond industrial automation, a similar system could allow doctors to trace analysis results about

We can utilize semantic technologies to connect agents within and across several domains in a pragmatic way.

a patient’s health back to the specific measurements that gave rise to these findings, which might be valuable during diagnosis processes.

Business perspective: Creating a business model that supports sharing ontologies. Ontologies will be valuable if they can be shared across partners and re-used over and over again. For instance, ontologies such as QUDT (www.qudt.org)—which contains information about quantities, units, and data types—are relevant in most domains, and therefore should constitute standard building blocks. Thus, organizational structures, business incentives, ways of protecting IP, and managing the distributed semantic knowledge base have to be identified that permit sharing the cost and benefits of these initial investments.

DISCUSSION AND CONCLUSION

The digitalization of industries has become a reality that requires new concepts of managing the information exchange between machines. To fully leverage the potential of coupling services to products, machines will need to collaborate horizontally across domains. This calls for developing mechanisms to add shared meaning to data that build upon established domain-specific standards, but make them usable in machine-comprehensible ways.

AS provides a good mechanism to coordinate and synchronize collaboration across devices in an emerging (industrial) IoT. Within that space, the development of shared vocabularies is valuable for bridging applications across IoT domains, which should support the creation of a common AS

description language. We suggested using semantic technologies to encode domain-specific standards and connect them to established and existing cross-domain ontologies. Finally, we proposed to apply AR technology to provide access to the invisible digital properties of smart devices and to help maintain and curate the corresponding semantic models. We also presented various examples and first results of our efforts to digitalize processes in the healthcare, smart grid and industrial manufacturing domains.

We believe enabling machines to “speak the same language” is pivotal for transforming physical products into digitally connected solutions. Although we have presented our first steps toward a Web of Systems, there is still a lot to do. We invite you to join us, and together we can make the digitalization of industry happen.

ACKNOWLEDGEMENTS

This article is based on joined work of the Siemens Web of Things research group in Berkeley, CA. The authors want to thank particularly Jack Hodges, Mareike Kritzler, Ralf Mosshammer, and Dan Yu.

References

- [1] Mayer, S., Wilde, E., and Michahelles, F. A Connective Fabric for Bridging Internet of Things Silos. In *Proceedings of the Fifth International Conference on the Internet of Things* (Seoul, South Korea, Oct. 26–28). IEEE, Washington D.C., 2015.
- [2] Ryokai, K., Michahelles, F., Kritzler, M., and Syed, S. Communicating and Interpreting Wearable Sensor Data with Health Coaches. In *Proceedings of the Ninth International Conference on Pervasive Computing Technologies for Healthcare* (Istanbul, Turkey, May 20–23). IEEE, Washington D.C., 2015.
- [3] Mayer, S., Inhelder, N., Verborgh, R., Van de Walle, R., and Mattern, Friedemann. Configuration of Smart Environments Made Simple. In *Proceedings of the Fourth International Conference on the Internet of Things* (Cambridge, Oct. 6–8). IEEE, Washington D.C., 2014.

Biographies

Florian Michahelles heads the Siemens Web of Things research group. Having worked in the fields of ubiquitous and wearable computing for more than a decade, Michahelles’ focus at Siemens is to leverage the web architecture and semantic technologies for enabling new business opportunities, especially in the fields of wearable sensing and human–robot interaction.

Simon Mayer has a strong background in distributed systems and Web of Things research; his focus at Siemens is to apply his knowledge on (Semantic) Web technologies, functional modeling, and human–computer interaction to enable interoperability and self-configuration across devices in factory and building automation scenarios. Mayer also co-chairs the international Web of Things (WoT) workshop series.

© 2015 Copyright held by Owner(s)/Author(s).