

The Internet of Things in an Enterprise Context

Stephan Haller¹, Stamatis Karnouskos², Christoph Schroth³

¹SAP (Schweiz) AG, SAP Research CEC Zürich,
Kreuzplatz 20, CH-8008 Zürich, Switzerland

²SAP AG, SAP Research CEC Karlsruhe,
Vincent-Priessnitz-Strasse 1, D-76131 Karlsruhe, Germany

³SAP (Schweiz) AG, SAP Research CEC St.Gallen,
Blumenbergplatz 9, CH-9000 St.Gallen, Switzerland
{stephan.haller, stamatis.karnouskos, christoph.schroth}@sap.com

Abstract. This paper puts the Internet of Things in a wider context: How it relates to the Future Internet overall, and where the business value lies so that it will become interesting for enterprises to invest in it. Real-World Awareness and Business Process Decomposition are the two major paradigms regarding future business value. The major application domains where the Internet of Things will play an important role and where there are concrete business opportunities are highlighted. But there are also many technical challenges that need to be addressed. These are listed and it is shown how they are tackled by existing research projects with industrial participation.

Keywords: Internet of Things, Future Internet, Real-World Visibility, Business Process Decomposition, Business Value, SOA, Enterprise Services

1 Introduction

The Internet of Things is a term that has been around for several years. It was first introduced by the MIT Auto-ID Center, the precursor to the current EPCglobal organisation, and at that time stood for the vision of a world where all physical objects are tagged with an RFID transponder with a globally unique ID – the EPC or electronic product code. RFID easily allows tracking the objects, and the EPC serves as a link to data which can be queried over the Internet about each individual object. Since then, the meaning of the Internet of Things has expanded. Using sensors or sensor networks, additional information about the objects or the environment that they are in can be recorded as well. Software embedded in the objects enables data processing directly on the item, and in combination with actuators, local control loops can be implemented.

The Internet of Things is a key part of the Future Internet. Many new opportunities can be foreseen for citizens as well as for businesses and other organisations, but also for the society as a whole.

2 The Internet of Things and the Future Internet

There are currently many terms flying around when trying to characterise the future development of the Internet: In addition to the Internet of Things, there is the Internet of Services, 3D Internet, Internet of Content, and Next-Generation Networks, just to name a few. It is important to note that these terms should not be regarded as different “Internets” that will exist in parallel, but rather as different aspects of a common Future Internet. The European Commission has understood this and is therefore taking concerted action and clustering the research projects it is funding in the 7th Framework Programme into what it calls the *Future Internet Assembly*. Furthermore, collaborations are on-going and likely will be intensified with similar efforts in the USA and Japan.

From an enterprise and economic perspective, the Future Internet is the basis for a *web-based service economy* [1]. There will be service platforms and a multitude of services available over the Internet, hence the term Internet of Services. The granularity of these services will be very different, ranging from high-level business services to low-level sensor services provided by the Internet of Things. The role of the Internet of Things is to bridge the gap between the physical world and its representation in information systems. This leads us to our definition of the Internet of Things:

“A world where physical objects are seamlessly integrated into the information network, and where the physical objects can become active participants in business processes. Services are available to interact with these ‘smart objects’ over the Internet, query their state and any information associated with them, taking into account security and privacy issues.”

It is noteworthy that in this definition, we don’t talk about technologies. RFID, sensor networks, embedded systems etc. are just enabling technologies, and we will see the technologies change over the years, but the main concept behind the Internet of Things will remain. Furthermore, the objects can be passive, as is the case with RFID-tagged objects, or active as in the case of machines with embedded process logic. Key is though the seamless integration into the business processes.

3 The Business Value of the Internet of Things

For the Internet of Things to become reality and not just stay at the buzz-word and concept level, investments will be needed: to solve current research challenges, to develop the necessary hard- and software, and to deploy the infrastructure required. This will only happen if there is a clear economic benefit. We see two major paradigms from which business value can be derived, which we term *real-world visibility* and *business process decomposition*. In the following, we define and explain both these terms before looking at some of the application areas that stand to profit most from the possibilities in the Internet of Things.

3.1 Real-World Visibility

The term real-world visibility denotes the fact that through the use of automated identification and data collection technologies like RFID and sensors it will be possible for a company to better know what actually is happening in the real world, i.e., how its operations are performing and what the status and whereabouts of its assets and products in the supply chain are. This enables what Fleisch et al. call *high-resolution management* [2]. The increased accuracy and timeliness of information about the business processes provides competitive advantages in terms of process optimisation [3]. The deeper insights gained into the processes allows a better understanding of them at the operational level and leads to their optimisation. Low data entry costs offset the considerable investment currently required, enabling organizations to benefit from automated entry and, largely, from continuous, near real-time measurement. Optimized shelf replenishments serve here as a well-known example where ‘out-of-shelf’ situations are avoided, thereby providing optimal product availability for the customer. Shop floor control can similarly be optimized with lot and equipment tracking functionalities that provide MES and ERP systems with more accurate information about the physical state and conditions of manufacturing inventories, thereby potentially increasing the efficiency and availability of production plants [4].

These example applications make use of the monitoring of the real world and the derived data in order to better control and manage business processes that deal with products and assets, the environment or persons. Sensory information increases the accuracy of real world checks and thus is the foundation of event-driven management. Automated sensing enables a new and much finer granularity of management, since it helps to control what was before uncontrollable. With a network architecture that supports the easy discovery, communication, and usage of events across the enterprise, events enriched with contextual business information will improve business data quality on all levels.

Additionally, (near) real-time analytics combines the usage of enterprise data with incoming sensory information to discover patterns and derive more timely and accurate business insight. New strategies and predictions can be generated based on certain pre-configured criteria.

3.2 Business Process Decomposition

Real-world visibility as described above is what basically is done today with RFID and sensing. Data is collected and sent to a central application, where it is processed and decisions about actions are taken. Business process decomposition takes this a step further, as shown in Fig. 1: A business process is decomposed into process steps, some of which are executed in a distributed manner, even at the edges of the network and on physical items themselves. Hence sometimes the term edge processing is used, and the physical items that can process some business logic are called *smart items*.

Current enterprise strategies already acknowledge a few interfaces to smart items, but with increased computational and communication capabilities of these items, the power shifts towards the edges of the network. Intelligent mechanisms for data

aggregation, filtering, fusion and conversion can be deployed to and executed at the network edge, or within the network, as appropriate. The decomposition and decentralization of existing business processes increases scalability and performance, allows better decision making and could even lead to new revenue streams through entitlement management of software products deployed on smart items. Software is already the key innovation driver in many industries and many new business models of the future will heavily rely on the use of such items [5].




Fig. 1. Traditional vs. decomposed and distributed business processes

Edge processing and business process decomposition allows applications to make (part of their) decisions locally in a decentralized manner and act accordingly. It thereby extends the real world visibility concept with real world interaction.

In addition to a (central) business application, data is processed locally in a distributed fashion. Actuators, which directly influence the real world, can be integrated as well. Smart items become thereby active participants in the overall business process.

Utilizing the computing power and intelligence on smart items or other edge devices allow enterprise systems to propagate data and services to these new, potentially underutilised, run time environments. It allows local decision-making and early aggregation and filtering of raw sensor data. However, it also requires innovation in distributed computing and network utilization.

Based on this paradigm shift, a new need arises for design time environments that allow modelling of business applications through a decentralized optimization strategy, taking into account all business objects, business processes, services, as well as processing, sensing and communication capabilities of smart items. New and optimized business processes can be modelled as collections of these 'artefacts' based on specific policies. Self-organisation in this context describes the adaptability of the model during deployment. Changes in the environment (e.g. location change, connectivity outage, reconfiguration of business processes etc.) require reorganization of the deployed components during run time to meet given Quality of Service (QoS) constraints.

In summary, edge processing can bring significant benefits, but this comes at a cost of increased management complexity. The decomposition of business processes can be seen as an evolutionary continuation of *enterprise service-oriented architectures*

(E-SOA). E-SOA enables the efficient decomposition of large segments of functionality into largely decoupled components. In essence, E-SOA is a set of architectural components for building autonomous yet interoperable composites. The convergence of computing and networking technologies and the trends to increased intelligence of smart items opens the opportunity of applying these concepts to the real world. Because of the increased management complexity however, it has to be determined for each application on a case-by-case basis if distributed business logic is appropriate. The following are criteria that indicate an advantage for a decentralized approach:

- Responsiveness of the overall system, since unnecessary ‘expensive’ communication round-trips are eliminated
- Scalability, since the execution of business logic is distributed
- Network independence: The system will also work when there is no connection to backend systems, e.g., during transportation or in temporary storage areas.

3.3 Major Application Areas

In the past, the benefits of RFID and related technologies have mainly been applied to supply chain logistics applications in the consumer products and retail industries. While these will remain an important cornerstone also in the future, we see with the expanded definition of the Internet of Things many other interesting application domains. Some of the most promising ones are discussed below.

Manufacturing. Due to the rapid advances in the embedded systems domain, the manufacturing domain is undergoing significant changes as ubiquitous computing is applied to the shop floor. An entirely new dynamic network of cooperating devices can be created which effectively shows that the Internet of Things can reshape the manufacturing domain. Currently, shop floor intelligent systems based on distributed embedded devices concentrate the programming of the behaviour and intelligence on a handful of large monolithic computing resources accompanied by large numbers of “dumb” devices. The intelligence and behaviour are tailored and individually programmed for each application.

As SOA concepts are becoming the de-facto standard to connect to enterprise applications, there has been a move also in the automation domain towards putting web services on the devices themselves and giving them the capability of providing their functionality as a service. This will create eventually cross-layer web service mash-ups, with services hosted at the enterprise, middleware and device level [6].

The key idea is to provide the same interoperability and ease of integration of devices as is the case with service mash-ups. All the devices would offer their functionality as a web service. Device integration thus means service integration, focusing on the functionality a device offers and not on the particular device technology. This not only creates a new paradigm on the shop floor, but it also would encourage the development of new devices in the automation industry that offer embedded web services. Furthermore, it would kick-start collaboration at the lowest level, i.e., among the devices themselves, as well as offer new opportunities by effectively connecting the vendor-locked isolated islands of today.

Supply Chain Integrity. Internet of Things technologies allow the tracking of the location and the state of an object throughout the full product life-cycle and throughout the supply chain. Nowadays this is already used to detect diversions into illicit or grey markets as well as the introduction of counterfeit products [7]. But this is just one part of ensuring complete integrity of the supply chain. Complete integrity though includes many other aspects. First of all, there is the physical integrity of the product itself. Sensors can be used to ensure that the product was never exposed to potentially damaging environmental conditions, e.g., that safe temperature or shock levels never were exceeded. Secondly, there is the integrity of the transportation routes, i.e. that the product never was in an area where it was not allowed to be. For example, hazardous goods may not be transported through heavily populated or environmentally sensitive areas because of the severe effects of an accident. For the same reason, some means of transportation may be forbidden, and other areas might be off-limits due to fears of unauthorized access to a product. And thirdly, there is the integrity of a product and all its subcomponents regarding the means of production. Here in particular environmental compliance is important. It must be ensured that the final product doesn't exceed certain levels regarding emissions and carbon footprint. With the problem of global warming, we believe that especially this last point will be a driving factor for research and development. Internet of Things technologies will enable the recording of all the emissions that were generated in the production and transportation of every subcomponent across multiple levels of a complex supply chain. In general, the new technologies will be used to ensure full *compliance* of all participants of the supply chain to a set of agreed-upon rules: legal regulations, internal policies and service-level agreements. Time is important in most cases: The violation of certain temporal boundaries would mean a possible breach of integrity.

Energy. In the quest for providing abundant clean, secure and affordable energy, Europe and the rest of the world are investing heavily in ICT in the energy domain. The usage of modern technologies, coupled with concepts coming from the Internet of Things and the Internet of Services will lead to a paradigm change. Innovative new technologies and concepts will emerge as we move towards a more dynamic, service-based, market-driven infrastructure, where energy efficiency and savings can be better addressed through interactive distribution networks. One example from the energy area that shows the importance of the Internet of Things is the creation of an advanced metering infrastructure (AMI).

AMI refers to systems that measure, collect, and analyze energy usage from advanced devices such as electricity, gas, and water meters. Communication can either be on request or on a pre-defined schedule. These devices are usually referred to as *smart meters*. Smart meters can be either considered as sensors with wireless capabilities themselves, or they are expected to cooperate with (wireless) sensor networks at home in order to deliver their data. Driven by new regulation and the global energy crisis, in the United States advanced metering with embedded wireless sensor networks has grown from a few thousand units in 2004 to 1.5 million smart meters this year, according to a survey by ON World. ON World projects that the global wireless sensor networking smart meter and demand response market will be worth \$1.6 billion in 2011.

Smart meters empower an advanced metering infrastructure which is able to react almost in real time, provide fine-grained energy production or consumption info and adapt its behaviour proactively. These smart meters will be multi-utility ones, managing not only electricity but also gas and heat, and they will depend on multiple (wireless) sensors. New information-dependent intelligent energy management systems will be needed for an infrastructure capable of supporting the deregulated energy market. Smart meters will have to be installed for millions of households and companies and get connected to transaction platforms.

Smart meters provide new opportunities and challenges in networked embedded system design and electronics integration. They will be able not only to provide (near) real-time data, but also process them and take decisions based on their capabilities and collaboration with external services. That in turn will have a significant impact on existing and future energy management models. Decision and policy makers will be able to base their actions on real-world, real-time data and not just on simple predictions. Households and companies will be able to react to market fluctuations by increasing or decreasing consumption or production, thus directly contributing to increased energy efficiency.

Health. In the health sector, RFID is now being used by some hospitals like the Jena University hospital in Germany [8] to optimise both the logistics processes as well as to provide better care, to track equipment, patients – in particular, new-born babies – and medications. This benefits both the health care providers in the form of cost savings as well as the patients, since false treatments can be avoided. The expected usage of wireless sensors as well as the emerging interconnection of all devices in the hospitals will increasingly change the landscape. Even bigger potential comes though with the expected ageing of society with the associated rise of health care costs: The Internet of Things will be essential in realising the vision of ambient assisted living.

Automotive. In the automotive industry, sensors and embedded systems already play a large role. These will become even more important when these are integrated into a future “Internet of Vehicles”: *Car-to-X Communication*, as it is more commonly known, denotes communication between vehicles as well as between vehicles and certain road-side infrastructure. This has turned into a very important research area over the last few years: Several different use-cases have been proposed and analyzed which range from safety-related warning systems over information and entertainment applications to Car-to-Business scenarios. In case the vehicles are wireless enabled (car manufacturers are currently investing into the definition of a common standard), feature a communication module and sensors for detecting road conditions and local dangers, they may autonomously connect to each other and exchange information (we refer to these temporary and autonomously organized connections as Vehicular Ad hoc Networks). When detecting hazards or blockings on the road as visualized, these vehicles can generate appropriate messages containing a description, their geographical position and many more application-dependent pieces of information. Messages may be immediately broadcasted to all other cars within communication range, which in turn can store, evaluate and forward them.

Besides such applications with a safety improvement focus, other scenarios encompass the communication of vehicles with backend business services. Based on

car-to-infrastructure communication technologies, cars may consume services such as remote diagnosis in case of break-downs, software version management, and other applications in the field of *Vehicle Relationship Management*. These exemplary scenarios show that vehicles are about to become smart items interacting over partly decentralized (ad-hoc) networks and partly over stable infrastructure-based Internet connections. The shift away from “mute and autistic” cars towards intelligent and proactive sensor nodes allows for both the above-mentioned real-world visibility and for business process decomposition: Firstly, cars are enabled to gather fine-granular information about their environments and their own status in order to facilitate better services (e.g., for central traffic management, anti-theft-, or smart insurance pricing applications). Secondly, vehicles become truly autonomous, intelligent items which are capable of performing event-based business logic in a decentralized fashion. Safety-related applications rely on cars being able to independently sense road conditions, create proper warning messages, disseminate them appropriately in temporary, ever-changing ad hoc networks and continuously evaluate them for their current relevance. Significant intelligence is thus allocated to the edges of vehicular networks, following the vision of an Internet of Vehicles.

Insurance. Insurance companies are assessing the consequences and potential benefits of new technologies and their applications. The considered technology trends lead potentially to a “high-resolution world”, which will allow insurance companies to better estimate and manage risks, enabling the setting of individual insurance rates as well as better decisions what to underwrite. It will also allow them to minimise losses by preventing accidents and damage events or at least reducing their severity thanks to early detection and intervention, and by helping to locate missing persons, animals or things. Furthermore, costs can be reduced by automating the claims management processes. Besides direct benefits (e.g., less costs due to less damage events), the main opportunity for insurance companies is to use the technologies to differentiate their offering from the competition – innovative offerings – and to increase the number of contacts with their customers. While in the past insurance companies only dealt with customers when signing a policy and when settling a claim, now they can leverage technology to frequently discuss how individual risks could be measured or reduced and how this could influence existing contracts or initiate new ones. In the case of loss, the technology also enables the insurance company to provide immediate and better assistance to their customers.

4 Major Technical Issues

The Internet of Things promises many opportunities and socio-economic benefits. However, there still are a number of technical as well as governance issues that need to be solved in order to realize the vision. In the following, we highlight some of what we believe to be major issues in the years ahead.

4.1 Internet Scalability

The Internet currently consists of more than 540 million hosts [9] and 1.4 billion users [10]. The Internet of Things will possibly have trillions of things connected. As an example of scale, let's assume that each thing will have its own IP address. IPv6 could accommodate 2^{128} (about 3.4×10^{38}) things, which raises the number of possible participants in the Future Internet by several orders of magnitude. As a result of connectivity, the data volumes will significantly increase; and in addition to discrete data we will see more and more streaming data, e.g., from sensors and cameras being transmitted over the Internet. Traditional client-server paradigms and methods of data processing where all data is collected and processed in a central instance will not suffice. A paradigm change needs to be undertaken. Considering the advances in the computing capabilities that are now available within the network as well as on its edges, new models and algorithms will need to be developed that can depend on "in-network" and "on-edge" processing with the help of information fetched from enterprise systems.

Scalability has been recognized as an important topic since the days the term "Internet of Things" was first used. Several methods exist and will have to be used in combination to achieve the desired scalability: filtering of irrelevant data, aggregating data to a higher semantic level, event-driven architectures and complex event processing, as well as the execution of business logic at the edges and on physical items themselves. The key is to reduce the number of messages and the amount of data transmitted throughout all the layers of the system; to that end data should be evaluated locally and where it makes sense, while propagation should be done only if a network-wide usage is expected or where there are clear application requirements for doing so.

Regarding the management of networks and devices, the focus currently is on the self-* research that includes self-configuration (automatic configuration of components), self-healing (automatic discovery and correction of faults), self-optimization (automatic monitoring and control of resources to ensure the optimal functioning with respect to the defined requirements) and self-protection (proactive identification and protection from arbitrary attacks).

As an example of testing scalability issues, within the IST FP6 SOCRADES [11] project we have embedded web services on a variety of devices. These devices provide their functionality as a service to which other entities can subscribe and get the status of the device and/or its sensor results. In our preliminary research we have successfully managed to simulate approx. 25'000 such web-service enabled devices, generated and hosted within a single Java virtual machine. As expected however, the overhead on the communication channel increased tremendously (as every second each one of these devices would both try to propagate its values to the subscribers, as well as to answer new discovery requests). In the Future Internet such operations will be much more common, requiring the right architectural decisions to be made. Event-Driven Architectures and Complex Event Processing (CEP) will gain significance, as this will be one major tool to manage complex data evaluations and extract meaningful and business relevant correlations.

Continued development and testing for scalability is needed. The IST FP7 project SENSEI [12] is addressing this by setting up a large-scale distributed Pan-European test bed for sensor and actuator based ambient intelligence.

4.2 Identification and Addressing

In order to be able to address the billions of entities in the Internet of Things, we first need to be able to identify them with a unique ID. That ID can also be used to find other information about the entity of interest. The EPC Network [13] as proposed by EPCglobal is one possible implementation of such a scheme. The EPC is used as a unique ID, and via the ONS the EPCIS server of the manufacturer of the product can be found where further information on the product can be looked up. However, for the unique identification of objects other approaches exist: In Japan, the ucode [14] is gaining popularity, and other groups – especially when we are talking about smart items – propose to use IPv6 addresses for this purpose. Furthermore, there are many well-established industry-specific IDs that at least so far have not been mapped to EPCs.

Also the ONS has not gained the popularity yet that was originally envisioned. The reason is that there are concerns regarding security and privacy [15] as well as reliability, since currently a single company, VeriSign, is in charge of operating the ONS root server. To achieve the necessary level of confidence globally, a more distributed model is required with a set of geographically distributed, but synchronized root servers, as is the case with DNS today. A first step has been taken in that direction with the awarding of a contract to Orange to operate a (still regional) ONS server by GS1 France [16].

Tracking and tracing objects as they are moving along the supply chain is one of the most important basic functions of the Internet of Things. It provides the foundation for product authentication, anti-counterfeiting and other supply chain integrity applications. To implement this functionality, discovery services are required that allow dynamically finding all information about a specific object. In an EPC context, this means a service that can find all EPCIS instances that contain some information about a certain object. One of the key difficulties here is to be able to find all relevant information in a world with multi-level adaptive supply chains where business relationships change quickly, and to still keep the necessary confidentiality. Such discovery services are currently being developed in the BRIDGE project; and a joint requirements group has also started to work on this within EPCglobal. Like with ONS, discussions centre here around the issue of centralisation vs. decentralisation. A centralised discovery service is certainly easier to implement and manage, but it raises concerns regarding confidentiality as well as availability. A decentralised or even peer-to-peer model therefore seems more acceptable in industry.

The ID used to uniquely identify an object can also be used when we want to communicate with the object. In the case of IPv6 addresses, the ID and the technical address are the same. In the case of EPCs, some mapping is required. Implementing such a mapping is trivial though, since the EPC can be used to look up the physical address; the technical address is simply another property of the object.

In contrast to this technical, direct addressing, from a business application perspective another method of addressing that we call *logical addressing* is needed as well. This denotes a method to address a group of objects or any random object within that group based on certain properties. For example, all drums in storage area A that contain isopropyl alcohol. Or, an arbitrary temperature sensor in meeting room B. This issue has been recognized in the research community. For example, in the IST FP6 CoBIs project [17], a reliable multicast dissemination protocol was developed and used to transmit business rule updates to a group of sensor nodes [18]. The SENSEI project is currently developing a service framework to interact with wireless sensor & actuator networks and semantically annotated service interfaces as well as context and actuation models, where this issue will also need to be addressed. However, a lot of work still needs to be done to develop both appropriate methods as well as globally accepted standards for logical addressing.

4.3 Heterogeneity

For connecting and integrating all the objects into the Internet of Things, there are and will be many different technologies. RFID is one of the most prominent ones, but sensors, wireless sensor & actuator networks and embedded systems will also play an important role.

In all these areas standards are in place or standardisation efforts are underway. For example, EPCglobal as well as ISO offer a family of standards for RFID, ZigBee [19] and 6LoWPAN [20] are gaining popularity in the wireless sensor networking area, and OPC [21] is well accepted in factory automation. But the technologies are too different to expect any standard to be able to cover them all. Furthermore, as shown in the section above regarding unique identification, there are also competing standards in the same field.

For all these reasons, it is clear that we will have to deal with heterogeneity when building the software infrastructures for the Internet of Things. Standards are helpful, but they alone cannot be the solution. Rather asked for here is *interoperability*. The key is to separate the functionality from its technical implementation. Service-oriented architectures are ideally suited for this, since they encapsulate functionality in services with a common interface, abstracting from the underlying hardware and protocols. This has already been understood in several domains, e.g., in the home automation with the usage of DPWS or in factory automation with the usage of OPC-UA.

Having infrastructures that allow connecting and integrating a diverse set of technologies is in our opinion not just a “necessary evil”, but rather a strength, since it offers two key benefits. Firstly, it allows applying different solutions to different applications. Depending on the application requirements, the best-fitting technology can be used. For example, in the CoBIs project three different sensor network platforms – Particles, μ Nodes and Sindrion, all exhibiting different properties [22] – were used. μ Nodes are best suited regarding energy efficiency, Particles offer more flexible and powerful programming capabilities and Sindrion is superior regarding standards support and commercialisation. All these platforms were integrated into the

same CoBIs framework using a service-oriented concept based on UPnP [23] on the lower levels and web services towards the applications.

Secondly, an infrastructure where diverse technologies can easily be integrated into will be future-proof. Especially in the area of wireless sensor networks, the technical developments are not complete and we can expect that new technologies, protocols and standards will arise. These will have to be interoperable with existing and already deployed devices and networks. An infrastructure built with heterogeneity in mind will easily allow this. The framework under development in the SENSEI project is a good example of such a system: It not only defines a plug & play service interface to connect heterogeneous wireless sensor & actuator network islands to the system, but also offers a framework to orchestrate different services provided by the islands into higher-level context and actuation services.

4.4 Service Paradigms

Service-Oriented Architecture (SOA) has been recognized as *adequate architectural style for the organization of large-scale and distributed business logic* since a number of years. Many definitions of SOA have already been published [24]. The normative OASIS Reference Model for SOA [25] defines SOA as "...a paradigm for organizing and utilizing distributed capabilities that may be under the control of different ownership domains. It provides a uniform means to offer, discover, interact with and use capabilities to produce desired effects consistent with measurable preconditions and expectations". According to this model, the major components of a basic SOA and their possible interactions are: a service provider publishes his service interface via a service registry where a service requester/consumer can find it and subsequently may bind to the service provider. The central concept of the SOA reference model is the existence of services which provide access to capabilities by well-defined interfaces to be exercised following a service contract with constraints and policies. This enables a loose coupling of services (thereby minimizing mutual dependencies) and complies with some of the probably most-known principles in software-engineering, information-hiding and modularisation [26]. Services are provided by entities, the service provider, and are to be used by others, the service consumers. Services may be composed on the basis of other, existing services, thereby adhering to the principle of reuse. They are autonomous in the sense that they solely control the logic they encapsulate, uniformly described and publicly retrievable via certain discovery mechanisms. The concept of supporting loosely coupled, business-aligned and networked services as introduced above can be realized with the help of numerous different technologies such as the Web Services stack. This stack comprises WSDL as a uniform format for service interfaces, UDDI as standard for specifying publicly available service registries, SOAP as data exchange protocol and BPEL as language for orchestrating services according to specific business logic.

However, to enable the seamless and agile interoperation of smart items in the Internet of Things, a number of challenges still exist with respect to the organization and implementation of services.

First of all, the lack of comprehensive, trustworthy and widely accepted *service intermediaries* in the Internet of today still prevents the establishment of a Future

Internet as described above. The few globally available UDDI-based registries are mostly only usable for technical experts, do not offer any additional functionality such as service performance monitoring and only feature a small number of references to different web services. As a part of the so called Web 2.0 movement, novel and richer kinds of intermediaries have already been established [27]. These allow for improved navigation by facilitating the efficient retrieval of services which match a given user need, service performance transparency through permanent monitoring of service quality and availability, and governance, i.e., controlling and enforcing the delivery of service levels as defined in Service-Level Agreements (SLA). Although promising, these intermediaries are still at an early stage of maturity and also rarely reach into the world of physical items; mostly they focus on simple information services.

Secondly, the *lack of means for composing diverse services* towards higher-order services meeting actual user demands imposes additional challenges to the emergence of the Future Internet. A phenomenon sometimes referred to as “corporate household” problem today prevents web services-based SOAs from being set up quickly and easily. Companies and their respective systems are subject to rigid governance mechanisms and adhere to different conventions with respect to business documents and also internal business processes. For this reason, the information objects which need to be exchanged between services are virtually guaranteed to have inconsistency in both syntax and semantics. Novel technologies will be required to annotate the business meaning of services and their input as well as output objects in a comprehensible way. Semantically organized services are believed to improve retrievability and seamless composability significantly. The EU-funded project FAST [28] deals with these challenges as it aims to develop a method as well as a tool enabling human users to quickly find electronic resources in the Web, and compose them with other resources in order to model a service-based scenario.

Thirdly, services in the Internet of today can be considered as *mute and autistic*. In the Future Internet, services are expected to become more open and reactive to their respective environments particularly in two respects: First, existing service interfaces (e.g., WSDL-based) are not designed to be interpreted and used by humans but rather by machines. This frequently prevents the long tail of Internet users from easily discovering and interacting with them. In the course of the above mentioned FAST project, novel methods are currently being developed which allow designers to “put a face” on services and to thus improve user-service interaction. Second, today’s Internet-based services are rarely context-sensitive. However, in order to realize the business process decomposition which has been discussed as a fundamental advantage of the Future Internet, services need to be aware of different aspects of the environment they are acting in [29]. Context services can provide such information to the application services that implement local control loops and that trigger specific actions. Such context services will be composed of several lower-level services which deliver individual sensor readings. However, this is not enough. For the application services to base their decisions on context data, some notion of *Quality of Information* is needed. Quality of Information provides meta information to what degree the information that a service provides is accurate – i.e., how well the measured value reflects the true value in the real world –, as well as timely, reliable and trustworthy. The EU-funded project SENSEI has taken on this challenge and will provide novel means for modelling, quantifying, representing and processing of Quality of

Information of individual sensor services as well as the aggregated context information itself.

5 Outlook

While we believe that there are many business opportunities ahead, industry is still reluctant to widely adopt the new technologies that will lead to the realisation of the vision set by the Internet of Things. The reason behind this is many-fold: As outlined above, many technical challenges still need to be resolved. But more importantly, often the real business cases are not clear yet. While many papers on this exist, the business cases have not been proven yet in the real world on a large-scale basis. Trust in the technology – and thereby also wider adoption – will only come with real-world deployments that operate reliably and productively and have a real impact on a company's financial bottom line. Furthermore, there is a need for open global standards, in particular on the topics of unique identification and object addressing, service interfaces, decentralised object discovery services and service infrastructures.

A combined – and focused on real-world needs – effort of industry and academia is needed to overcome the challenges. The technical challenges need to be addressed, and the solutions have to be tested and proven in real settings. In addition, new business models and quantifiable business cases need to be developed. The 7th Research Framework Programme of the European Union provides the means to fund such collaborative efforts; however, open collaboration with initiatives in other regions of the world – in particular Asia and the Americas – will be essential.

Acknowledgments. This paper describes work undertaken in the context of the SENSEI project, 'Integrating the Physical with the Digital World of the Network of the Future' (www.sensei-project.eu). SENSEI is a Large Scale Collaborative Project supported by the European 7th Framework Programme, contract number: 215923. The authors would also like to thank the other colleagues working on Internet of Things topics at SAP Research.

References

1. ISTAG Working Group Report on "Web-based Service Industry", February 2008, ftp://ftp.cordis.europa.eu/pub/ist/docs/web-based-service-industry-istag_en.pdf
2. Fleisch, E., Sarma, S., Subirana, B.: High-Resolution Management. IESE Alumni Magazine, Juli/September 2006, pp. 8-13
3. Spiess, P., Karnouskos, S.: Maximizing the Business Value of Networked Embedded Systems through Process-Level Integration into Enterprise Software. In: Second International Conference on Pervasive Computing and Applications, pp. 536-541 (2007)
4. Karnouskos, S., Baecker, O., Moreira Sa de Souza, L., Spiess, P.: Integration of SOA-ready Networked Embedded Devices in Enterprise Systems via a Cross-Layered Web Service Infrastructure. 12th IEEE Conference on Emerging Technologies and Factory Automation, pp. 293-300 (2007)

5. Odenwald, T.: Putting M2M into Business Context, M2M Pioneer Magazine, Volume 5, Number 2, March/April 2007
6. Moreira Sa de Souza, L., Spiess, P., Koehler, M., Guinard, D., Karnouskos, S., Savio, D.: SOCRADES: A Web Service based Shop Floor Integration Infrastructure. Internet of Things 2008 Conference, pp. 50-67 (2008)
7. Lehtonen, M., Oertel, N., Vogt, H.: Features, Identity, Tracing, and Cryptography in Product Authentication. 13th International Conference on Concurrent Enterprising (2007)
8. SAP Press Release: Jena University Hospital Leads Healthcare Innovation with RFID Solution from SAP, May 2006, <http://www.sap.com/industries/healthcare/newsevents/Press.epx?PressID=6329>
9. ISC Domain Survey: Number of Internet Hosts, <http://www.isc.org/ds/host-count-history.html>
10. World Internet Users, March 2008, <http://www.internetworldstats.com/stats.htm>
11. IST FP6 Project SOCRADES, <http://www.socrades.eu/>
12. IST FP7 Project SENSEI, <http://www.sensei-project.eu/>
13. EPCglobal Architectural Framework, <http://www.epcglobalinc.org/standards/architecture/>
14. Sakamura, K.: Ubiquitous ID Technologies 2008, uID Center booklet, http://www.uidcenter.org/pdf/UID910-W001-080226_en.pdf
15. Fabian, B., Günther, O., Spiekermann, S.: Security Analysis of the Object Name Service. In: First International Workshop on Security, Privacy and Trust in Pervasive and Ubiquitous Computing, pp. 71-76 (2005)
16. GS1 France Press Release: GS1 lance le premier service de l'Internet des Objets en Europe, May 2008, http://www.gs1.fr/gs1_fr/actualites__1/informations__1/communiqués_de_presse/communiqués_de_presse_2008/gs1_lance_le_premier_service_de_l_internet_des_objets_en_europe
17. IST FP6 Project CoBIs, <http://www.cobis-online.de/>
18. Marin-Perianu, M., Havinga, P.: RMD: Reliable Multicast Data Dissemination within Groups of Collaborating Objects. In: 31st IEEE Conference on Local Computer Networks, pp. 656-663 (2006)
19. ZigBee Alliance, <http://www.zigbee.org/>
20. N. Kushalnagar, N., Montenegro, G., Schumacher, C.: IPv6 over Low-Power Wireless Personal Area Networks (6LoWPANs): Overview, Assumptions, Problem Statement, and Goals. RFC 4919, IETF, Aug 2007, <http://www.ietf.org/rfc/rfc4919.txt>
21. OPC Foundation, <http://www.opcfoundation.org/>
22. CoBIs Project Final Report, http://www.cobis-online.de/files/Deliverable_D104V2.pdf
23. UPnP Forum, <http://www.upnp.org/>
24. Alonso, G., Casati, F., Kuno, H., Machiraju, V.: Web Services Concepts, Architectures and Applications. Springer, Berlin (2004)
25. MacKenzie, M., Laskey, K., McCabe, F., Brown, P.F., Metz, R.: OASIS - Reference Model for Service Oriented Architecture 1.0, http://www.oasisopen.org/committees/tc_home.php?wg_abbrev=soa-rm,2006
26. Parnas, D.L.: On the criteria to be used in decomposing systems into modules. In: Communications of the ACM, Vol.15, No.2, ACM Press (1972)
27. Hiero, J.J., Janner, T., Lizcano, D., Reyes, M., Schroth, C., Soriano, J.: Enhancing User-Service Interaction Through a Global User-Centric Approach to SOA. In: Fourth International Conference on Networking and Services, pp. 194-203 (2008).
28. IST FP7 Project FAST, <http://fast.morfeo-project.eu/>
29. Schroth, C., Christ, O.: Brave New Web: Emerging Design Principles and Technologies as Enablers of a Global SOA. In: The 2007 IEEE International Conference on Services Computing, pp. 597-604 (2007).