

Industrial Agents in the Era of Service-Oriented Architectures and Cloud-based Industrial Infrastructures

Armando Walter Colombo^{1,2}, Stamatios Karnouskos³, João Marco Mendes², Paulo Leitão^{4,5}

¹ University of Applied Sciences Emden/Leer, Germany, awcolombo@technik-emden.de

² Schneider Electric Automation GmbH, Germany, marco.joao@schneider-electric.com

³ SAP, Germany, stamatis.karnouskos@sap.com

⁴ Polytechnic Institute of Bragança, Portugal, pleitao@ipb.pt

⁵ LIACC - Artificial Intelligence and Computer Science Laboratory, Portugal

1. Introduction

The umbrella paradigm underpinning novel collaborative industrial systems is to consider the set of intelligent system units as a conglomerate of distributed, autonomous, intelligent, pro-active, fault-tolerant and reusable units, which operate as a set of cooperating entities (Colombo & Karnouskos, 2009). These entities are forming an evolvable infrastructure, entering and/or going out (plug-in / plug-out) in an asynchronous manner. Moreover, these entities, having each of them their own functionalities, data and associated information, are now connected and able to interact. They are capable of working in a pro-active manner, initiating collaborative actions and dynamically interacting with each other in order to achieve both local and global objectives. New emergent behaviors resulting from the co-operations arise and need to be managed in a smart manner.

Service-oriented architectures (SoA) principles and technologies are considered as an adequate backbone to enable the industrial implementation of such collaborative industrial automation and management system corresponding to the e.g., ISA'95 standard, from the sensor/actuator level through the control devices (CNC, PLC, Robot Controls) and SCADA (Supervisory Control and Data Acquisition) to the Manufacturing Execution System (MES) levels and above. Another very important result of the implementation of the SoA paradigm in the collaborative industrial environment is associated to the digitalization (virtualization) of the physical environment, i.e. (1) "things in the real world" may get a digital address (get connected to the Internet) and expose their own data and information, and (2) the Internet "things in the cyber world" get real (physical) world aware.

A first consequence of the digitalization of the industrial environment is that "Services" are having direct physical impact and real physical world integrates part of the cyber world. A second major consequence is the big amount of machine processable data, servitized functions, generated by heterogeneous data sources located both in physical and cyber world. Both function and data, but also information derived from the data processing and intelligent decision making systems are offered / exposed as services in both worlds, i.e.,

physical world by devices and systems, and cyber world by a cloud of services. Each entity, located in the physical and/or in the cyber (cloud-based) world, connected into the Cyber-Physical Systems (CPS)-network, is then able to access and to consume those services, and also to use these services for generating new ones.

Smartness is intrinsically embedded in an immense set of distributed but networked physical and cyber entities; products, solutions and services. Major challenges arise when this smartness of such collaborative industrial infrastructures needs to be mastered, i.e., mastering the inherent autonomy of each of the entities and mastering the co-operation capabilities of the networked entities.

The application of the Industrial Agents paradigm is well-fit to act as an enabler for mastering such collaborative industrial systems. Physical agents following the “Holon Control” principles (Leitão et al., 2005) are capable of using the information exposed as services in an autonomous manner to perform their own functions and are able to negotiate among them to achieve common goals as e.g., controlling emergent behaviors of the multi-agent community by processing, combining, orchestrating, composing that data. In summary, both kind of data source in a digitalized industrial environment, i.e., physical and cyber (cloud) entities need the support of the agents for fulfilling many of their collaborative behaviors, for achieving their “common goals”.

Although the adoption of service-oriented cyber-physical systems is increasingly getting industrial consensus, it should not be underestimated that this kind of systems needs the connectivity and interoperability with real-time decision systems responsible for supporting the management of the emergent behaviors and timely assessment of the big amount of reachable digital data. On the one side, multi-agent based real-time decision systems that have been designed for managing emergent behaviors need access to the information/data exposed by the components of the industrial environment; they need the SoA-based cyber-physical infrastructure. On the other side, the functionality and usability of SoA-based industrial CPS need to be enriched by multi-agent decision making system.

In this work a brief overview of the Service-oriented Architectures paradigm and related technologies that are currently used as backbone to implement industrial cloud-based cyber-physical systems is discussed. Additionally founded arguments to consider the industrial agents as an unavoidable complementary automation and management system in that CPS industrial environment is analyzed.

The chapter is organized as follows: First key concepts such as service-oriented architectures (SOA), cloud systems and the way they can be used in industrial automation systems is discussed. Subsequently it is investigated how multi-agent systems (MAS) and SOA principles can be combined to extract the best of the two worlds. Some example use cases are then analyzed with the first related to a cyber-physical simulation infrastructure using agents and services, and the second one related to a prototype industrial implementation of service-oriented industrial automation system. The last section rounds up the chapter with some considerations.

2. Technologies

2.1. Towards service oriented architectures

In the last years, significant efforts have been done (Colombo et al., 2014) to investigate the benefits as well as the impact of emerging technologies, such as SOA, cloud computing, CPS, etc. Industrial Agents have long been recognized as a key approach (Leitão et al., 2013) for developing intelligent solutions, e.g., for simulating behaviors, monitoring and autonomously taking decisions in the field, as well as acting as a glue among disparate systems and functionalities.

The SOA paradigm is a way of building distributed systems, originally designed for business systems and electronic commerce, but progressively adopted in other domains. SOA is based on the concept of providing and requesting services. Basically, a service is a software piece that encapsulates the control logic or functionality of an entity that responds to a specific request. In such systems, a provider entity hides its internal structure and functionalities by encapsulating them as services and offer to the other entities (requesters) by publishing them in a service registry central repository, as illustrated in Figure 1.




Figure 1: SOA concepts

The list of provided services must be published, so they can be discovered by the service requester. Using discovery mechanisms, e.g., UDDI (Universal Description, Discovery and Integration), service requesters can find the services they need. After getting information about the available services, the service requester can invoke the execution of those services. More complex services may be created by aggregating the functionalities provided by simpler (atomic) ones. This functionality is referred as service composition and the aggregated service becomes a composite service (Chafle et al., 2004). The composition of services requires mechanisms for coordination and synchronization and shares many common features with workflow systems. However, service composition requires additional functionalities for discovery and checking interoperability of the services (Karakoc et al., 2006).

Other concepts, such as service orchestration and choreography, are important for the coordination and composition, and particularly to determine how the services “play” together. Orchestration is the practice of sequencing and synchronizing the execution of services,

which encapsulate business or manufacturing processes (Jammes et al., 2005). An orchestration engine implements the logic for the workflow-oriented execution and sequencing of atomic or composed services, and provides a high-level interface for the composed process. The service choreography is a complementary concept, which considers the rules that define the messages and interaction sequences that must occur to execute a given process through a particular service interface.

Despite the possibility of using other implementation strategies, SOA is commonly implemented using web services (OASIS, 2006). A web service, as defined by the World Wide Web Consortium (W3C), is a software system that supports interoperable machine-to-machine interaction over a network (W3C, 2004). The use of the service-oriented paradigm, implemented through web services technologies, enables the adoption of an unifying technology for all levels of the enterprise, from sensors and actuators to enterprise business processes (Bepperling et al., 2006, Karnouskos et al., 2010).

2.2. Towards web service enabled devices: DPWS, REST, OPC-UA

Current industrial monitoring and control applications are facing many challenges as the complexity of systems increases and the systems evolve from synchronous to asynchronous. When hundreds of thousands of devices and service-oriented systems are asynchronously interconnected and share and exchange data and information, i.e., services, for monitoring, controlling, and managing the processes, key challenges such as interoperability, real-time performance constraints, among others, arise and need to be addressed. Several Internet based technologies and concepts have found their way to the industrial automation and especially on the integration of devices (Bangemann et al., 2014). Some of the most widely-used constitute of Devices Profile for Web Services (DPWS), Representational State Transfer (REST) and OPC Unified Architecture (OPC-UA) .

A standard dealing with ubiquitous device integration is DPWS as described in OASIS (2009) standard, which is a collection of web service standards. Initially, DPWS was conceived as a successor of UPnP (Universal Plug and Play) for home automation scenarios, but recent works have shown its applicability to the automation world (Karnouskos et al., 2010). DPWS advances previous dynamic discovery concepts, such as Jini (www.jini.org) and UPnP (www.upnp.org) to integrate devices into the networking world and make their functionality available in an interoperable way. DPWS is an effort to bring web services to embedded devices taking into consideration their constrained resources. Several implementations exist in Java and C (e.g., www.ws4d.org, www.soa4d.org), while Microsoft has also included a DPWS implementation (WSDAPI) by default in Windows Vista onwards and Windows Embedded CE operating systems. DPWS exists in a number of devices today, and basically brings the SOA world down to the devices, hence extending a fully service oriented infrastructure down to the physical world and resource constrained networked embedded systems.

An alternative integration approach is REST, as described by Fielding (2000), which is the architectural principle that lies at the heart of the Web and shares a similar goal with integration techniques, such as WS-* web services, that is increasing interoperability for a looser coupling between the parts of distributed applications. However, the goal of REST is to achieve this in a more lightweight and simpler manner; therefore it focuses on resources,

and not functions as is the case with WS-* web services. In particular, REST uses the Web as an application platform and fully leverages all the features inherent to HTTP such as authentication, authorization, encryption, compression, and caching. This way, REST brings services “into the browser”, i.e. resources can be linked and bookmarked and the results are visible with any Web browser. There is no need to generate complex source code out of WSDL (Web Services Description Language) files to be able to interact with the service.

Finally, OPC-UA (Mahnke, 2009) was developed with the goal to provide a path from the traditional OPC communications model to a SOA. OPC-UA supports a binary protocol for high performance and a web service protocol (e.g., SOAP (Simple Object Access protocol)) which is firewall friendly using standard http/https ports. IEC 62541 is a standard for OPC Unified Architecture.

OPC-UA, DPWS and REST constitute some of the “emerging” technologies and blend with many other traditional ones in the shop floor (Bangemann et al., 2014). The selection of the best fit technology depends on the scenario and the requirements posed, as at this stage all of them have benefits but also drawbacks (Jammes et al., 2014). Lighthouse projects, such as SOCRADES (www.socrades.eu) and IMC-AESOP (www.imc-aesop.eu), have developed and tested prototypes in industrial settings that use a mix of these technologies to integrate industrial systems (Colombo et al., 2014) as well as couple them with information and business systems (Karnouskos et al., 2010). There are also ongoing efforts, e.g., to further enhance the performance in DPWS with the introduction of Efficient XML Interchange (EXI), as well as integrate more lightweight protocols, such as the IETF Constrained Application Protocol (CoAP) as well as fusion of DPWS and OPC-UA (Colombo et al., 2014; Jammes et al., 2014).

All of these efforts that promote modularization and easy integration over heterogeneous infrastructures act as enabler for Industrial Agents. The latter can be realized both within the device itself as well as externally and interact with the devices via well defined services as will be analyzed later in this chapter.

2.3. Cloud-based industrial systems

The future industrial automation systems are expected to be complex system of systems that will empower a new generation of today hardly realisable applications and services. The rapid advances in technology during the last years have given rise to virtualisation and cloud systems. Virtualisation addresses many enterprise needs for scalability, more efficient use of resources and lower Total Cost of Ownership (TCO) to name a few. Cloud computing has emerged powered by the widespread adoption of virtualization, service-oriented architecture and utility computing. IT services are accessed over the Internet and local tools and applications (usually via a web browser) offer the feeling that they were installed locally. However, the important paradigm change is that the data are computed in the network but not in a priori known places. Typically, the physical infrastructure may not be owned and various business models exist that consider access-oriented payment for usage (Karnouskos et al, 2014).

New industrial systems and architectures are being developed to take advantage of the cloud and its services (Karnouskos et al., 2014b). Figure 2 illustrates such an effort carried

out within the IMC-AESOP project (Colombo et al., 2014). There we see the emergence of an information-based infrastructure that is built in a complementary fashion to the traditional automation “pyramid” as defined in ISA-95. The ever increasing need for rapid development and deployment of applications and services has taken advantage of the modularization of functionalities and availability of services at the different traditional automation levels (Level 0 up to Level 4) and combines them in lightweight application specific manner.




Figure 2: Industrial automation evolution: complementing the traditional ISA-95 automation world view (pyramid on the left side) with a flat information-based infrastructure for dynamically composable services and applications (right side)

Hence, although the traditional hierarchical view is left untouched, hooks in the form of services enable now the emergence of a flat information-based architecture. Next-generation industrial applications can now rapidly be composed by selecting and combining the new information and capabilities offered (as services in the cloud) to realise their goals. The envisioned transition to the future cloud-based industrial systems is depicted in Figure 2.

For Industrial Agents, such visions and technology trends, signals a new era. Industrial agents can very well act as enablers for the servicification of the traditional ISA-95 infrastructure by capturing key functionalities and providing them as services. In addition they could play coordination roles by orchestrating the integration of various services in the cloud while hosting the intelligence needed.

3. Bridging Agents and SOA-enabled devices

The Internet of Things is prevailing in the industrial domain where devices are acquiring increasingly sophisticated computing and communication capabilities. As such these are envisioned to play active role in emerging collaborative infrastructures and systems. Hence we witness efforts to migrate advanced functionality previously hosted in powerful static back-end systems, towards more lightweight mobile distributed embedded devices. Web services nowadays can be implemented directly on devices, providing them with the

necessary technology abstraction and making them easily integratable in heterogeneous environments. Additionally, intelligence can also be realized in various forms including in the forms of agents. In such systems agents can be integrated within the intelligent device or as an orchestrator at a higher level. Therefore coupling agents and devices for industrial purposes could yield several benefits.

3.1. Agents and services commonalities

Service-oriented principles can be integrated with MAS to enhance some functionalities and to overcome some limitations, namely in terms of interoperability, legacy system integration and IT-vertical integration. In spite of being based on the same concept of providing a distributed approach to the system, MAS and SOA present some important differences, namely in terms of computational requirements and interoperability, as illustrated in Table 1 (Ribeiro et al. (2008) provide a deeper study of these differences).

Table 1: Differences between MAS and SOA (adapted from Ribeiro et al. (2008))

Multi-agent systems	Service-oriented Architectures
Well established methods to describe the behavior of an agent	Focus is on detailing the public interface rather than describing execution details
Agents denote social ability regulated by internal or environmental rules	Social ability is not defined for SoA
Most implementations are optimized for LAN use, but Internet is also possible	Supported by Web related technologies and can seamlessly run on the internet
Reactive to changes in the environment	Reconfiguration often requires reprogramming
Interoperability heavily dependent on compliance with FIPA-like standards	Interoperability assured by the use of general purpose Web technologies
Heavy computational requirements	High performance without significant interoperability constraints

These differences highlight the complementary aspects of the two paradigms, suggesting the benefits of combining them to extract the best of both worlds: the intelligence and autonomy provided by MAS solutions and interoperability offered by SOA solutions (Huhns, 2002). This suggestion is not new since services are already part of the agents' specification, e.g., already included in the Foundation for Intelligent Physical Agents (FIPA) specification (FIPA, 2002), and agents are also present in standard documents of SOA methodologies, e.g., in the OASIS (2006) standard. Thereof, the under-considered elements (services in MAS and agents in SOA) are vaguely defined and have a more passive and customized role.

3.2. Approaches to combine agents and services

Traditionally, the combination of MAS and SOA paradigms can be performed in different ways, as illustrated in Figure 3 (Mendes et al., 2009). The first traditional option, illustrated in Figure 3.a), considers gateways to translate the semantics from the agent world to the

services world. According to the FIPA specifications, this task is basically performed by translating:

- Service registration: DF (Directory Facilitator) ↔ UDDI (Universal Description Discovery and Integration)
- Service description: agent service ↔ WSDL (Web Services Description Language)
- Message: ACL (Agent Communication Language) ↔ SOAP




Figure 3: Common approaches for integrating SOA and MAS

An example is the Web Services Integration Gateway (WSIG) plug-in provided by the Java Agent Development (JADE) framework to offer an implementation of the concept of gateway agent (Bellifemine, 2007). This plug-in, in form of gateway agent, was implemented by Whitestein Technologies and allows transparent and bidirectional transformations between FIPA compliant services and web services employing the WSDL / SOAP / UDDI stack (i.e. publishing agents' capabilities as Web services used in a SOA environment). The communication between the WSIG Gateway Agent and the other agents use FIPA-ACL, as illustrated in Figure 4, and the service discovery is performed by using two repositories: DF (for the agents world) and UDDI (for the services world). The discovery transformation performed by the gateway agent allows agents to perform service discovery in web services registry using the UDDI and web service clients to perform service discovery in MAS registry using the DF.

Other similar examples are the WS2JADE (Nguyen and Kowalczyk, 2007) and AgentWeb Gateway (Shafiq et al., 2005). Several applications combining MAS and SOA principles employing the concept of gateway agent are reported in the literature. For example, Jacobi et al. (2010) use a model-driven approach that combines SOA and MAS to model a segment of a production chain in the steel industry and Fayçal et al. (2010) propose the integration of legacy systems by the encapsulation of its features by agents. Another idea is to join the subscribing directories from the agent side (DF) from the web services side (UDDI) in just one common place named UD³ (Cheaib et al., 2008).

Utilizing the described approach, the design of truly service-oriented multi-agent systems are far from the real expected potential and benefits, since the combination is only focused in the communication perspective offered by SOA approaches, and not fully exploring the potential of designing the system using service-orientation. Another option, illustrated in Figure 3.b), was introduced by Mendes et al. (2009) and is characterized by the use of a set of

autonomous agents that use the SOA principles, i.e. oriented by the offer and request of services, to fulfil industrial systems goals. The achieved service-oriented multi-agent systems (SOMAS) approach is different from the traditional MAS mainly because agents are service-oriented, i.e. according to Mendes et al. (2009):

- Agents share services as the major form of communication among agents.
- Individual goals of agents may be complemented by services provided by other agents.
- The internal functionalities of agents can be offered as services to others agents.

An important note is that these service-oriented agents do not only share services as their main form of communication, but also complement their own goals with externally provided services.

An example of using the SOMAS approach is illustrated in Figure 4, where devices represent conveyors (transporting pallets) and pallets, and have associated service-oriented agents that are responsible of part of their environment (Leitão, 2012). The conveyor agent provides a service, called transfer pallet, which encapsulates its internal functionality of transferring the pallet from the input location to the output location. Therefore, it has the ability to read the sensors, execute the embedded logic control and send commands to the actuators of the conveyor. This service is published in the Service Registry to be discovered by other agents representing devices, e.g., conveyors or pallets.




Figure 4: Example of a service-oriented multi-agent system

Other neighbor devices, e.g., a pallet agent that needs this transfer service to accomplish its goals, may request the service to the conveyor agent. However, to complete the execution of the service and also to respect global objectives, the conveyor must request an availability service from the next transport unit or workstation connected to its output, using the SOAP protocol. This can be seen as the form of collaboration among the service-oriented agents in the system.

3.3. Enterprise Service Bus based solutions

SOA based systems can be realized by an Enterprise Service Bus (ESB) that provides a layer on top of an implementation of an enterprise messaging system (Ziyaeva et al., 2008), acting as backbone for supporting the interoperability among the connected software applications. Typically desirable capabilities of ESBs include, without being exhaustive, process orchestration (typically via WS-BPEL), protocol translation, hot deployment, versioning, lifecycle management and security. The use of an ESB constitutes an alternative way to implement the integration of MAS and SOA following the SOMAS concept, where software applications are MAS-based systems that are interacting through the use of the ESB, by exposing and consuming services.

An example of the use of this approach to integrate MAS and SOA paradigms is provided by the Adaptive Production Management (ARUM) project (arum-project.eu) that addresses the development of solutions to handle emergent challenges in ramp-up production of complex and highly customized products, such as aircraft industry, and particularly mitigation strategies to respond faster to unexpected events and intelligent decision support systems for planning and operation (Marin et al., 2013).

Aiming to achieve a full interoperability across the entire ARUM solution, traditional ESBs, e.g., the open source JBoss ESB (Jboss, 2014) and the proprietary TIE Smart Bridge (TSB, 2014), are enriched with a plethora of advanced modules and functionalities that support the tools lifecycle from creation time until they are unplugged from the system, resulting on an intelligent enterprise service bus (iESB). Examples of such modules are the Ontology Service, Data Transformation Service, Sniffer, Node Management and Life-Cycle Management. The iESB provides a common infrastructure for the integration of heterogeneous agent-based planning and scheduling tools, and legacy systems using the services principles, as illustrated in Figure 5.




Figure 5: Integration of MAS and SOA using an Enterprise Service Bus

The pluggability of the agent-based tools is facilitated by the exposition of their functionalities as services and by the use of the ontology services for the representation of the shared knowledge, improving the interoperability in such distributed and heterogeneous systems.

4. Use Case: Cyber-Physical Infrastructure Simulation by coupling Software Agents and Physical Devices

Today, we see the emergence of cyber-physical infrastructures composed from high number of heterogeneous devices. The latter may as well be SOA-enabled devices on the basis of technologies such as OPC-UA, DPWS and REST as we have already discussed. However, in order to study large-scale systems, the development of real testbeds with hundreds or thousands of such devices is costly. Hence, a compromise might be to simulate their behavior as realistic as possible. Simulating an infrastructure populated by a high number of web service enabled devices is not trivial, but it could provide a very useful tool in the hands of enterprise application developers.

Coupling agents with such physical devices could provide an interesting approach for investigating some of the aspects including management and network aspects. An architecture for such a simulation is depicted in Figure 6 (Karnouskos & Tariq, 2009).




Figure 6: A Simulator of CPS infrastructures relying on agent-driven integration

The devices at the lowest layer make available their functionality via web services, while a subscription can be made to their services. The device layer consists of devices that directly implement web services, e.g. via the DPWS protocol, and/or via the DPWS gateway (due to resource constraints, etc). Typical examples of such devices that implement web services (SOA-ready) are programmable logic controllers (PLCs), robots, advanced sensors, e.g. SunSPOTs etc., and example of devices connected via a DPWS gateway could be RFID tags that connect via an RFID reader that acts as a DPWS gateway.

At execution layer, the mobile agent system hosts several agents that not only cooperate but also control the created virtual devices. One layer higher relies the logic, which describes the scenarios that users run within the simulator. The scenarios range from simple ones running standalone up to complex ones which may start other simpler scenarios first. Finally at enterprise layer, various services and applications can communicate via web services with the devices, both real and simulated ones.

For the implementation, the JADE multi-agent platform (Bellifemine et al., 2007) is used to create the agents representing DPWS devices. Each agent represents one DPWS device which needs to be created using the DPWS toolkit (www.soa4d.org). This integration has been achieved by creating two types of agents interacting with the DPWS toolkit, i.e. (i) a DPWS Client Agent (DC-Agent), and (ii) a DPWS Server Agent (DS-Agent), as analyzed in detail in Karnouskos & Tariq (2008).

The DC-Agent implements the client part of the DPWS toolkit, acting as client for consuming services offered by devices as well as for services offered by DS-Agents. This agent acts as a bridge between a device and a DS-agent offering service(s) to applications. Tasks assigned to the DC-Agent include discovery of other in-network DPWS-enabled devices, acquisition of services and data offered by those devices, processing of data, and exposition of data to other applications via the DPWS protocol.

The DS-Agent implements the server part of the DPWS toolkit and is more complex as it consists of two distinct components, i.e. a server and a service. The server part instantiates the services, registers them and listens at specified port for the client requests. The service part is exposed to the external world and handles all the client requests.




Figure 7: DC-/DS-Agents and DPWS devices discoverable in Windows

As it can be seen in Figure 7, simulated and real DPWS-enabled devices can be discovered by third party DPWS clients. These appear as normal devices (distinguishable only by their name), and coexist with other devices such as a robotic arm, a SunSPOT sensor (www.sunspotworld.com) and a windows computer. This makes obvious that the simulator

created devices can at least be discovered/used by other infrastructure actors in an agnostic, non-intrusive way.

The simulation environment consists of a basic set of agents, each of which has its goals and internal logic (Karnouskos & Tariq, 2009):

- *Management Agent*: Tasks of this agent include evaluation of user arguments, creation of other agents and other management functions, e.g. logging.
- *Device Explorer Agent*: This agent is based on the concept of DC-Agent which is a DPWS Client Agent and its job is to discover all the DPWS enabled devices in the network based with a specific scope.
- *Device Generator Agent*: The core function of this agent is to receive and execute requests towards creating and initializing service agents that simulate a specific service.
- *Scenario Agent*: This agent is specific for each scenario as it executes its strategy/logic.
- *Service Agent(s)*: Design of a service agent is based on the DS-Agent model. Such types of agents simulate a DPWS service and are visible to the external world via the DPWS communication.

Using the capabilities of the simulator, thousands of DPWS devices were instantiated and investigated (Karnouskos & Tariq 2009). However, limitations in the hosting computer(s) played a role and potentially these results can be revisited with more powerful hardware, larger distribution of the agents, e.g. in the cloud, and more efficient implementations of the DPWS toolkit.

The agents played various key roles in this system. Firstly, they acted as “glue” that servicified physical devices and exposed their capabilities via web services and more specifically the DPWS protocol. As such, any “legacy” or other non-SOA devices could now be easily integrated via web services. The agents also acted as simulation scenario orchestrators, holding the intelligence needed to execute the simulation. As such, we can witness a diverse utilisation of their capabilities and some potential roles they can play in industrial settings.

5. Use Case: Service-oriented industrial automation system

The European research project SOCRADES had explored the application of service-orientation and web services for the next generation of industrial automation systems. In particular, an engineering framework for the development of service-oriented automation systems was introduced by Mendes et al. (2008), using the Petri nets formalism as a unified tool for the specification, modeling, analysis and execution of service-based automation systems. Petri nets are also exploited as the form of orchestration and composition in service-oriented automation systems.

The application scenario used to demonstrate the SOA approach is based on a customized Prodatec/FlexLink DAS 30 – Dynamic Assembly System – as shown in Figure 8. The DAS 30 system is a modular factory concept platform for the light assembly, inspection, test, repairing and packing applications.




Figure 8: Prodatec/FlexLink DAS 30 used for the demonstration (located at Schneider Electric Automation GmbH in Seligenstadt, Germany)

The DAS 30 system comprises a flexible production system with two work stations (that can be used by operators and robots), several conveyors that route production pallets into/out of the system and to the workstations, and also two lifters. The central part of the transfer system (C1-C9) is made of nine transfer units (conveyors) of unidirectional and cross types. The unidirectional transfer unit provides an input and an output port and the cross transfer unit provides transfers not only in the longitudinal but also in transversal axis. The lower transfer units (C10, C11) have the same behavior as the normal unidirectional transfer units (such as unit C5), but are physically longer. Lifter units (L1 and L2) are responsible for the interface between the upper and lower part of the system, and also for transferring pallets into and out of the automation system.

The pallets enter in the system via the unit C4 and are conveyed using alternative paths to the two workstations W1 and W2. The routing is done at the transfer units based on the required production operations needed by the product mounted on a particular pallet and based on the location and availability of production services in the system (at W1 and W2). A workstation can provide more than one type of production operations and one kind of production operation could be provided by more than one workstation. The units C4, C6, C2 and C8 are equipped with RFID (Radio-Frequency IDentification) that are able to read/write information from/to tags attached to the pallets.

The approach for creating complex, flexible and reconfigurable production systems is based on a network of modular, reusable entities that expose their production capabilities as a set of services. Data and information associated to industrial equipment, i.e. physical entities like a robot, as shown in Figure 11, are digitalized by smart embedded devices and exposed as services into a cyber-infrastructure such as a “Service Bus” (cloud of services).

A composition approach applies to most levels of the factory floor; simple devices compose complex devices or machines, which in turn are composed to build cells or lines of a production system and so on. The same applies to concept of service-oriented production systems and composing complex services from simpler services, complemented with orchestration engines as illustrated in Figure 9. As a matter of fact, the orchestration engines will be located (embedded) into selected devices and their orchestration/composition functionalities exposed from the devices or directly from the service bus considered here as

the service recipient of the service cloud. Note: orchestration engines appear where atomic services discovered in the service bus have to be composed or orchestrated to generate new services or to manage and control results of service compositions.




Figure 9: Important elements of the service-oriented automation system

The service ecosystem for the case study automation system is represented in Figure 10. The atomic services are exposed by the transfer units (Transfer), lifters (Lifting) and RFID devices (RFID). These services are the building blocks for the more advanced engineering of this system and can be associated and composed depending on the requirements and objectives of the application.




Figure 10: Service landscape of the automation system

Since services aren't isolated entities exposed by the intervenient software components, a kind of logic that is responsible for the interaction is needed. The model-based orchestration engine is able to interpret a given work-plan made of services (an orchestration) and execute it. The work-plan can be defined in Business Process Execution Language (BPEL) as defined in OASIS (2007), Petri nets formalism e.g. Hamadi and Benatallah (2003) and Bing and Huaping (2005), or even in adapted IEC 61131-3 languages, beside others.

The modeling language used in this work derives from Petri net specifications, including time considerations, property system and customizable token game engine. The developed Petri net orchestration engine needs to know how and when to respond to services and to represent them in the model. This is done by describing transitions in the Petri net model. A transition willing of sending a request/response or an event must be enabled, and the action is done when it fires. In the other hand, a transition receiving a message from a request, response or event, will only fire if it is enabled and the message is there.

The information to be used by transitions is gathered by an imported WSDL file that contains the description of the service. Depending on the operation, transitions can be part of a client request/response, server request/response, client event and server event. The first two types require two transitions: one for initializing the request and one for the response. It is also possible to test responses by their return parameters, implying the use of one response transition for each test. The difference of an operation being a server or client is obvious: a server waits for the request and then gives a response, and a client makes a request and waits for a response. Events are possible as client and server, but only require one single direction (and consequently, one transition).

A major task at this stage is to fit the automation bot, including the orchestration engine and web service technology into an automation device. The resulting smart embedded device host the most of the services exposed in the system and also responsible for the coordination and control of the mechanical parts of demonstrator, as represented in Figure 11. These devices use the Telemecanique Advantys STB (Small Terminal Box) NIP2311 prototype devices, which provide two main interfaces: mediating the automation equipment via input/output modules and managing the access to the service bus by exposing and requesting services (using the Ethernet network interface module). Atomic services representing resources and functions of the connected equipment are provided by the device interface. Some of them may include an orchestration engine to "link" services together and to create new composite services. An internal decision support system is responsible to sustain the engine for decisions, e.g., selecting the best process based on decision criteria.

For this prototype implementation, the controller of the Ethernet module is used to host the service infrastructure, based on the SOA4D implementation of DPWS (forge.soa4d.org), allowing the deployment of user-defined applications as DPWS-compliant service components. The services are implemented by the STB with an embedded IEC-61131 engine. The ControlBuild prototype developed by Geensys (www.geensys.com) is used to specify the logic and services offline and then to deploy those into STBs. Another STB prototype has been implemented that provides an embedded service orchestration engine based on the Continuum Bot Framework with Petri nets kernel (Mendes et al., 2009) and the DPWS stack with the same deployment mechanisms as for the STB with IEC-61131 engine.

The orchestration engines run on their own STBs and provide composed services to the system.




Figure 11: Structure of a smart embedded device using Advantys STB NIP2311

The control logic is managed by the Petri Nets Kernel module that interprets a given Petri net model (Mendes et al., 2009b). During the execution of the behavioral models, some decision nodes may appear requiring their real-time resolution. In case of Petri nets to represent the system behavior, this detection is performed with the identification of marked places that can evolve into more than one alternative ways, i.e. the marked places that have connected more than one enabled transition. As illustrated in Figure 12, the place p_1 constitutes a decision node because there are three alternatives to evolve the model, i.e. the operation service can be performed by three distinct machines. The decision point is translated in the Petri net model as a conflict, being necessary that someone, in this case a decision support system, resolves the conflict, i.e. select one of the machines depending on various criteria.




Figure 12: Petri net based orchestration with decision support system

The degree of complexity associated to the decision-support system can range from simple algorithms to complex cognitive systems, being the use of agents a natural option to provide intelligence during the orchestration process. After selecting the best option to evolve, the

achieved decision is translated to the PN model by increasing the priority associated to the selected transition, in this case transition t3. Analyzing the priority of alternative transitions, the logic controller will evolve the system by firing the transition with higher priority, being activated the correspondent web services and sent a message to the machine.

The orchestration models can be connected together via the ports of the models, using two alternative ways:

- *Offline composition*, which permits to generate a new model based on the connection of individual ones. For this connection, the information has to be setup in the Petri net models and an XML connection file must be defined to describe which models will be connected and via which ports.
- *Online composition*, which permits the intercommunication of two engines and their respective models via the exposition and request of services (this is already part of the information of the models designed before).

At the time of the experimentation, there were only three available STB devices embedding Petri net orchestration engines, which one able to run one model at a time. The solution was using the offline composition to generate only three composed models (one for each orchestration device) and let them work together in real-time using the online composition. Afterward, the decision was to split the system into 3 clusters of units, resulting in one model for C1-C3, one model for C4-C5, L1, L2, C10 and C11, and another model for C7-C9, ending up in three composed Petri nets models. The generated models communicate via each other (for inter transfer operation of pallets) using service invocation (i.e. the "TransferIn/TransferOut" mechanism).

The composition application shows that it is possible to design individual models without knowing the availability and disposability of the final orchestration devices. The experiment shows one possible way to compose the system using three devices and a defined distribution, but it could also be done with a different number of devices and other ways of division. Offline composition is used to limit the use of devices, network traffic, but introduces more complex models to be orchestrated (considering the limitations of embedded devices). On the other hand, online composition is focused more on the distributed orchestration and the synchronization thereof. The correct division and use of the composition types depends always on the available resources, the optimization strategies and the layout of the system, but orchestration models can be individually developed without knowing this information.

6. Conclusions and Future Directions

Although agents in general as well as Industrial agents have been investigated for several years, their productive use in industrial settings has been demonstrated but is limited. Other technologies and approaches that complement them have been used as we have already discussed. However, with the prevalence of a new high-tech infrastructure driven by cyber-physical systems as defined in the Industrie 4.0 vision, Industrial Agents come again into the forefront of realizing key features needed. As such, we see a renewed interest in the practical applications of Industrial Agents, especially in conjunction with Cyber Physical systems, SOA and cloud computing. Their roles can vary from delivering intelligence to the infrastructure, acting as "glue" for legacy systems, negotiating or mediating functionalities and services, etc.

To achieve large portions of the Industrie 4.0 vision, further research is required, with focus on the usage of modern Internet technologies and services but for Industrial production use. The latter assumes a good understanding of the challenges and limitations posed in real-world industrial systems, as well as the optimization of agent systems to make them sustainably operational in such environments.

Acknowledgements

The authors would like to thank for their support the European Commission, and the partners of the EU FP7 projects SOCRADES (www.socrades.eu) and IMC-AESOP (www.imc-aesop.eu) for the fruitful discussions.

References

Bangemann, T., Karnouskos, S., Camp, R., Carlsson, O., Riedl, M., McLeod, S., Harrison, R., Colombo, A. W., Stluka, P. (2014), "State of the Art in Industrial Automation", Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach, A. W. Colombo, T. Bangemann, S. Karnouskos, J. Delsing, P. Stluka, R. Harrison, F. Jammes, J. L. Martínez Lastra (eds.), Springer, pp. 23-47. URL: http://dx.doi.org/10.1007/978-3-319-05624-1_2

Bellifemine, F., Caire, G. and Greenwood, D. (2007), "Developing Multi-Agent Systems with JADE", Wiley, April 2007.

Bepperling, A., Mendes, J.M., Colombo, A.W., Schoop, R. and Aspragathos, A. (2006), "A Framework for Development and Implementation of Web Service-Based Intelligent Autonomous Mechatronics Components", Proceedings of the IEEE International Conference on Industrial Informatics, Singapore, pp. 341-347, 2006.

Bing, L., Huaping, C. (2005), "Web Service Composition and Analysis: A Petri-net Based Approach", First International Conference on Semantics, Knowledge and Grid (SKG '05), November.

Chafle, G. B., Chandra, S., Mann, V. and Nanda, M. G. (2004), "Decentralized Orchestration of Composite Web Services", Proceedings of the 13th International World Wide Web Conference on Alternate track papers & posters, ACM Press, pp. 134-143.

Cheaib, N., Otmane, S., and Mallem, M. (2008), "Combining FIPA Agents and Web Services for the Design of Tailorable Groupware Architecture", Proceedings of the 10th International Conference on Information Integration and Web-based Applications & Services, pp. 702–705.

Colombo A.W. & Karnouskos S. (2009), "Towards the factory of the future: a service-oriented cross-layer infrastructure", ICT shaping the World: a Scientific View, European Telecommunications Standards Institute (ETSI), Wiley, New York, pp 65–81

Colombo, A. W., Bangemann, T., Karnouskos, S., Delsing, J., Stluka, P., Harrison, R., Jammes, F. & Lastra, J., eds (2014), "Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach", Springer. ISBN: 978-3-319-05623-4.

Fielding, R.T. (2000), "Architectural Styles and the Design of Network-based Software Architectures", PhD thesis, University of California, Irvine, Irvine, California, USA.

Fayçal, H., Habiba, D., Hakima, M. (2010), "Integrating Legacy Systems in a SOA using an Agent based Approach for Information System Agility", Proceedings of the International Conference on Machine and Web Intelligence (ICMWI'10), pp.338-343, 2010.

FIPA (2002), FIPA Abstract Architecture Specification, Standard of the Foundation for Intelligent Physical Agents, 2002. URL: <http://www.fipa.org/specs/fipa00001>

Hamadi, R., Benatallah, B. (2003), "A Petri net-based model for web service composition", Proceedings of the 14th Australasian database conference, Darlinghurst, Australia, pp. 191-200.

Huhns, M.N. (2002), "Agents as Web services", Internet Computing, IEEE, vol. 6, n° 4, pp.93–95.

Jacobi, S., Hahn, C. Raber, D. (2010), "Integration of Multiagent Systems and Service Oriented Architectures in the Steel Industry", Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT'10), vol.2, pp.479-482, 2010.

Jammes, F., Karnouskos, S., Bony, B., Nappey, P., Colombo, A. W., Delsing, J., Eliasson, J., Kyusakov, R., Stluka, P., Tilly, M., Bangemann, T. (2014), "Promising Technologies for SOA-based Industrial Automation Systems", Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach, A. W. Colombo, T. Bangemann, S. Karnouskos, J. Delsing, P. Stluka, R. Harrison, F. Jammes, J. L. Martínez Lastra, (eds.), Springer, pp. 89-109. URL: http://dx.doi.org/10.1007/978-3-319-05624-1_4

Jammes, F., Smit, H., Martinez Lastra, J.L. and Delamer, I. (2005), "Orchestration of Service-Oriented Manufacturing Processes", Proceedings of the 10th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA'05), vol. 1, pp. 617-624.

JBOSS (2014), JBOSS middleware, URL: <http://www.jboss.org> (available on 23rd September 2014).

Karakoc, E., Kardas, K. and Senkul, P.A. (2006), "Workflow-based Web Service Composition System", Proceedings of the 2006 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology, IEEE Computer Society, pp. 113-116.

Karnouskos, S., Colombo, A. W. and Bangemann, T. (2014), "Trends and Challenges for Cloud-based Industrial Cyber-Physical Systems", Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach, A. W. Colombo, T. Bangemann, S. Karnouskos, J.

Delsing, P. Stluka, R. Harrison, F. Jammes, J. L. Martínez Lastra (eds.), Springer, pp. 231-240. URL: http://dx.doi.org/10.1007/978-3-319-05624-1_11

Karnouskos, S., Colombo, A. W., Bangemann, T., Manninen, K., Camp, R., Tilly, M., Sikora, M., Jammes, F., Delsing, J., Eliasson, J., Nappey, P., Hu, J., Graf, M. (2014b), "The IMC-AESOP Architecture for Cloud-based Industrial CPS", Industrial Cloud-based Cyber-Physical Systems: The IMC-AESOP Approach, A. W. Colombo, T. Bangemann, S. Karnouskos, J. Delsing, P. Stluka, R. Harrison, F. Jammes, J. L. Martínez Lastra (eds.), Springer, pp. 49-88. URL: http://dx.doi.org/10.1007/978-3-319-05624-1_3

Karnouskos, S., Savio, D., Spiess, P., Guinard, D., Trifa, V., Baecker, O. (2010), "Real World Service Interaction with Enterprise Systems in Dynamic Manufacturing Environments", Artificial Intelligence Techniques for Networked Manufacturing Enterprises Management, L. Benyoucef, B. Grabot, Eds., Springer, pp. 423-457. URL: http://dx.doi.org/10.1007/978-1-84996-119-6_14

Karnouskos, S & Tariq, M. M. J. (2008), "An Agent-Based Simulation of SOA-Ready Devices", Proceedings of the 10th International Conference on Computer Modeling and Simulation, Cambridge, England, IEEE Computer Society, 2008, pp. 330-335.

Karnouskos, S & Tariq, M. M. J. (2009), "Using Multi-agent Systems to Simulate Dynamic Infrastructures Populated with Large Numbers of Web Service Enabled Devices", Proceedings of the International Symposium on Autonomous Decentralized Systems (ISADS'09), Athens, Greece, pp. 1-7.

Leitão, P., Colombo, A.W. and Restivo, F. (2005), "ADACOR, A Collaborative Production Automation and Control Architecture", IEEE Intelligent Systems, vol. 20, n° 1, pp. 58-66.

Leitão, P. (2012), "Towards Self-Organized Service-oriented Multi-agent Systems", Service Orientation in Holonic and Multi-agent Manufacturing and Robotics, T. Borangiu, A. Thomas and D. Trentesaux (eds), Springer, pp 41-56.

Leitão, P., Marik, V. & Vrba, P. (2013), "Past, Present, and Future of Industrial Agent Applications", IEEE Transactions on Industrial Informatics, vol. 9, n. 4, pp. 2360–2372.

Mahnke W., Leitner S.H., Damm M. (2009), "OPC Unified Architecture", Springer, Heidelberg, ISBN 978-3-540-68899-0.

Marín, C., Mönch, L., Leitão, P., Vrba, P., Kazanskaia, D., Chepegin, V., Liu, L., Mehandjiev, N. (2013), "A Conceptual Architecture Based on Intelligent Services for Manufacturing Support Systems", Proc. of the IEEE International Conf. on Systems, Man, and Cybernetics (SMC'13), pp. 4749-4754.

Mendes, J.M., Leitão, P., Colombo, A.W. and Restivo, F. (2008), "Service-oriented Process Control using High-Level Petri Nets", Proceedings of the 6th IEEE International Conference on Industrial Information (INDIN'08), Daejeon, South Korea, 13-16 July, pp. 750-755.

Mendes, J.M., Leitão, P., Restivo, F. and Colombo, A.W., (2009), "Service-oriented Agents for Collaborative Industrial Automation and Production Systems", Proceedings of the 4th International Conference on Industrial Applications of Holonic and Multi-Agent Systems (HoloMAS'09), V. Marik, T. Strasser and A. Zoitl (eds), LNAI 5696, Springer, pp. 1-12, 2009.

Mendes, J.M., Bepperling, A., Pinto, J., Leitão, P., Restivo, F. and Colombo, A.W. (2009b), "Software Methodologies for the Engineering of Service-oriented Industrial Automation: The Continuum Project", Proceedings of the 33rd Annual IEEE International Conference on Computer Software and Applications (COMPSAC'09), Seattle, Washington, USA, 20-24 July, pp. 452-459.

Nguyen, X.T., Kowalczyk, R. (2007), "WS2JADE: Integrating Web Service with Jade Agents", Proceedings of the SOCASE 2007 Conference on Service-oriented Computing: Agents, Semantics, and Engineering, J. Huang et al. (eds.), LNCS 4504, Springer-Verlag, pp. 147-159.

OASIS (2006), "Reference Model for Service Oriented Architecture 1.0", October 12, 2006. URL: <http://docs.oasis-open.org/soa-rm/v1.0>

OASIS, "Web Services Business Process Execution Language Version 2.0", OASIS Standard, April 2007.

OASIS (2009), Devices Profile for Web Services (DPWS), URL: <http://docs.oasis-open.org/ws-dd/dpws/1.1/os/wsdd-dpws-1.1-spec-os.html>

Ribeiro, L., Barata, J. and Mendes, P. (2008), "MAS and SOA: Complementary Automation Paradigms", IFIP International Federation for Information Processing, vol. 266, Springer Boston, pp. 259-268, 2008.

Shafiq, M.O., Ali, A., Ahmad, H.F., Suguri, H. (2005), "AgentWeb Gateway - a Middleware for Dynamic Integration of Multi Agent System and Web Services Framework", Proceedings of the 14th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprise, pp. 267-268.

TSB, TIE Smart Bridge,
URL: <http://businessintegration.tiekinetix.com/nl/contact/smartbridge-for-suppliers> (available on 23rd September 2014).

W3C (World Wide Web Consortium), Web Services Glossary, 2004. URL: <http://www.w3.org/TR/ws-gloss/>

Ziyayeva, G., Choi, E., Min, D. (2008), "Content-Based Intelligent Routing and Message Processing in Enterprise Service Bus", Proceedings of the International Conference on Convergence and Hybrid Information Technology (ICHIT '08), pp.245-249.