

AgileChains: Agile Supply Chains through Smart Digital Twins

Barbara Pernici, Pierluigi Plebani

DEIB, Politecnico di Milano, Italy. E-mail: [name.lastname]@polimi.it

Massimo Mecella, Francesco Leotta

DIAG, Sapienza Università di Roma, Italy E-mail: [lastname]@diag.uniroma1.it

Federica Mandreoli, Riccardo Martoglia, Giacomo Cabri

FIM, Università di Modena e Reggio Emilia, Italy E-mail: [name.lastname]@unimore.it

Currently, production and logistics performance of a single organization are only partially dependent on the internal resources, but more and more often, they also depend on the interactions that happen across the so-called supply chain, that is, the interactions between the organization and its customers and suppliers. In particular, the production and logistics coordination between actors in the supply chain is often a difficult activity which draws significant resources. Also, such coordination requires continuous revisions and updates to be performed. In Industry 4.0, the digital twins paradigm is currently adopted to represent, simulate and test the behavior of one or more machines and production plants belonging to an organization. This paper introduces the AgileChains paradigm, extending the digital twin paradigm to supply chains and the dynamics of their participants. This extension also positively affects the reactivity and resilience of the internal processes in case the supply chain has to be reconfigured. We propose a novel conceptual framework that combines Service Oriented Architectures (SOA) with Cyber-Physical Systems (CPS), in order to create service oriented systems suited for exchanging data in a dynamic and adaptive way. In addition, we propose a novel data management mechanism capable of finding the right balance between the internal needs of each organization when handling their data and the need to securely and efficiently export data in the supply chain (cf. smart data movement). Finally, we plan to define governance tools to model and manage the supply chain that treat agility as a first-class citizen. These tools will allow users to dynamically and predictively change the involved actors, as well as the nature of the exchanged data and the data exchange policies, focusing in particular on adverse, risk-prone events, so to minimize the risk and to optimize the supply chain performance both in terms of efficiency and effectiveness.

Keywords: Digital twins, Data spaces, Industry 4.0, reliability, resilience, Big data, IoT.

1. Introduction

Currently, production and logistics performance of a single organization are only partially dependent on the internal resources, but more and more often, they also depend on the interactions that happen across the so-called supply chain, that is, the interactions between the organization and its customers and suppliers. In particular, the production and logistics coordination between actors in the supply chain is often a difficult activity which draws significant resources. Also, such coordination requires continuous revisions and updates to be performed. In Industry 4.0, the digital twin paradigm illustrated in Rosen et al. (2015) is currently adopted to represent, simulate and test the behavior of one or more machines and production plants belonging to an organization. This paper introduces the AgileChains paradigm, extending the digital twin paradigm to supply chains and the dynamics of their participants. This extension also

positively affects the reactivity and resilience of the internal processes in case the supply chain has to be reconfigured. We propose a novel conceptual framework that combines Service Oriented Architectures (SOA) with Cyber-Physical Systems (CPS), in order to create service oriented systems suited for exchanging data in a dynamic and adaptive way. In addition, we propose a novel data management mechanism capable of finding the right balance between the internal needs of each organization when handling their data (cf. data sovereignty, as discussed in Jarke et al. (2019)) and the need to securely and efficiently export data in the supply chain (cf. smart data movement). Finally, we plan to define governance tools to model and manage the supply chain that treat agility as discussed in Christopher (2000) as a first-class citizen. These tools will allow to dynamically and predictively change the involved actors, as well as the nature of the exchanged data and the

data exchange policies, focusing in particular on adverse, risk-prone events, so to minimize the risk and to optimize the supply chain performance both in terms of efficiency and effectiveness.

The structure of the paper is as follows. In Section 2. we discuss different aspects in related work which are the basis for the proposed approach. In Section 3 we illustrate the service-oriented approach to digital twins that is the basis of the present work. Flexible data management is illustrated in Section 4, and finally in Section 5. we discuss issues for supporting agile design for risk-prone events.

2. Related work

In the present section we discuss the state of the art related to the main underlying components of the proposed framework: service composition, service discovery and data management and interoperability in a supply chain context.

2.1. Service composition

The supply chain management requires the possibility to connect several, and often heterogeneous, information systems related to the involved companies. This generates the need to govern complex “systems of systems” which, as indicated by the recommendation of NESSI (Networked Software and Services Initiative), presented in NESSI Europe (2018), requires the study of solutions that ensure the service continuity also in case of failures, as well as the creation of “digital trust” combining both the human and technological aspects as required by both industrial and supply chains processes.

Related to the processes, the problem of making them agile, that in AgileChains is related to both internal and supply chain processes, has been considered in terms of flexibility by Reichert and Weber (2012) and risk awareness by Suriadi et al. (2014). Conversely, about resilience studies are limited and only on a preliminary stage Marrella et al. (2019). Several approaches have been proposed to dynamically compose services. For instance, adaptive services frameworks have been proposed in Ardagna et al. (2007) and in Ardagna et al. (2011), where the flexibility of business processes is based on a dynamic composition, of available services, that can be reconfigured in case of changes in the execution context. A flexible approach based on service-oriented architectures has been introduced for digital twins in cyber-physical systems in Catarci et al. (2019). There is a need to further evolve such approaches in order to manage the ecosystem of available services and adapt their utilization to changing circumstances.

In the SAPERE project, the aim was to define support for *dynamic service composition* Zambonelli et al. (2015), based on the concept of dis-

tributed pervasive service *ecosystem*. The pervasive ecosystem is modeled as a virtual *spatial environment*, built on top of the physical network of devices infrastructure. The basic concept is that of building a shared space in which all service components are situated, and the environment has the role of mediating all interactions among services, so that different types of services indirectly interact and combine with each other. A limited set of basic interaction laws (also called “eco-laws”, due to their nature-inspired origins) regulates the interaction among services and typically takes into account the spatial and contextual relationships between services. In SAPERE, each *service component* populating the ecosystem has an associated semantic representation which called “LSA” (*Live Semantic Annotations*), describing a list of properties and characteristics for each services, available in the spatial environment as it were a sort of shared spatial memory and providing the basis for enabling dynamic environment-mediated interactions between services. In this way it is possible to support semantic and context-aware interactions both for service aggregation/composition and for data/knowledge management using the eco-laws as a basis. Data and services (as represented by their associated LSAs) are considered as a sort of chemical reagents, and interactions and compositions will occur via chemical reactions, relying on semantic matching between LSAs.

Adaptivity in SAPERE is not in the capability of individual services, but in the overall self-organizing dynamics of the service ecosystem as a whole. In particular, such an approach to adaptivity in AgileChains will be adopted so that changes in the system (as well as any change in its services or in the context of such services, as reflected by dynamic changes in their LSAs) can be represented also in their interactions, which can imply the firing of new eco-laws, thus possibly leading to the establishment of new compositions or aggregations, and/or in the breaking of some existing service compositions.

2.2. Services Discovery

Service discovery is an important aspect in digital twins, in particular when participants in supply chains must exploit services provided by others, since it provides a mechanism which allows automatic detection of services offered by any component in the system. The goal of a service discovery mechanism is to develop a highly dynamic infrastructure where requestors would be able to seek particular services of interest, and service providers offering those services would be able to announce and advertise their capabilities. Furthermore, service discovery should minimize manual intervention and allows the system to be self-healing by automatic detection of services which have become unavailable. Once services

have been discovered, devices in the system could remotely control each other by adhering to some standard of communication. Over the past years, many organizations and major software vendors have designed and developed a large number of service discovery protocols. Among different proposals, we can mention some approaches. *SLP - Service Location Protocol (SLP)* Guttman (1999) is an open, simple, extensible, and scalable standard for service discovery developed by the IETF (Internet Engineering Task Force). SLP addresses only service discovery and leaves service invocation unspecified. The SLP architecture consists of three main components: *User Agent (UA)*, *Service Agent (SA)*, *Directory Agent (DA)*. *Jini* Arnold et al. (1999) is a distributed service discovery system developed by Sun-Microsystems (now Oracle) in Java. The goal of the system is the federation of groups of clients/services within a dynamic computing system. A Jini federation is a collection of autonomous devices which can become aware of one another and cooperate if need be. To achieve this goal, Jini uses a set of lookup services to maintain dynamic information about available services and specifies how service discovery and service invocation is to be performed among Java-enabled devices. The heart of Jini is a trio protocols called: *discovery*, *join*, and *lookup*. Jini does not allow the evaluation of complex queries with Boolean operators or comparators such as SLP. *UPnP^a* is a Microsoft-developed service discovery technology aimed at enabling the advertisement, discovery, and control of networked devices and services. It is built upon IP that is used for communication between devices, and uses standard protocols like HTTP, XML, and SOAP for discovery, description, and control of devices. UPnP uses a non-directory based approach for service discovery where each device hosts a device description document. Control points can search only for: all services, specific service type, or specific device type since SSDP does not support attribute-based querying for services. *UDDI - Universal Description, Discovery, and Integration (UDDI)^b* is an XML-based registry for business internet services. Publishing a Web service involves creating a software artifact and making it accessible to potential consumers. Web service providers augment a Web service endpoint with an interface description using the Web Services Description Language (WSDL) so that a consumer can use the service. Optionally, a provider can explicitly register a service with a Web Services Registry such as UDDI or publish additional documents intended to facili-

tate discovery such as Web Services Inspection Language (WSIL) documents. The service users or consumers can search Web Services manually or automatically. In a business solution, it is very normal to search multiple UDDI registries or WSIL documents and then aggregate the returned result by using filtering and ranking techniques.

Also in service discovery the semantic representation of services and their adaptive match-making play an important role. A light-weighted semantic approach such as the one proposed in URBE (Plebani and Pernici (2009)) can be adopted.

2.3. Data management and interoperability

A proper supply chain management necessarily requires the interconnection among the different - and often heterogeneous - information systems managed by the several organizations involved. Although the Service Oriented Computing paradigm is providing methods and tool to make such an interconnection easier, an agile supply chain requires even more flexible solutions to make the data flow easier and quicker. In such a scenario, a proper data management in a supply chain becomes fundamental (Michelberger et al. (2013)), with particular emphasis on the delivery of information at the right time, the right place, and with the right quality and format to the user.

Three data management solutions, acting at different levels, can be considered to this end. Polystores have been recently proposed as very flexible data management solution that extends the concept of federated and parallel databases to simultaneously work with disparate database engines and programming/data models while supporting complete functionality of underlying DBMSs (Gadepally et al. (2016)). This architecture can thus provide support to the processing of queries on heterogeneous data stores, to the transparent transformation of just-in-time data and to multiple query interfaces. A data lake (LaPlante and Sharma (2016)) is a central architecture to store data regardless of formats and sources. Specifically, in a data lake data are stored in their native format and there is no need to design a common data structure. A dataspace is an abstraction to integrate data which enable the coexistence of heterogeneous data sources by providing basic functionalities to all of them (Franklin et al. (2005)).

Digital twins are characterized by a high production rate of monitoring data, which can overload the IT infrastructure with data transmission and also the computational infrastructure for performing predictive analysis on available data. As a consequence dynamic monitoring infrastructures such as in Andreolini et al. (2015), adaptive monitoring and data movement Vitali et al. (2019) and a logistic approach to data management as proposed

^a<https://openconnectivity.org/developer/specifications/upnp-resources/upnp/>

^b<http://www.oasis-open.org/committees/uddi-spec/doc/tcpspecs.htm#uddi3>

in the DITAS project (Plebani et al. (2017)) are needed.

3. Service-oriented approach to digital twins

AgileChains operates at both data and process level inspired by the digital twin concept. The main properties of a digital twin are: (i) connectivity: ability to communicate with external, entities and other digital twins; (ii) autonomy: ability to run regardless of the other digital twins; (iii) homogeneity: ability to reuse the same digital twin in different contexts; (iv) customizability: ability to modify the behavior of a physical entity by operating through the functions exposed by the digital twin, and (v) traceability: ability to have effects to the related physical entities.

A digital twin exposes a set of services which are directly connected to the machines in the shopfloors to execute operation and to produce data which are related to the performed activities useful to query and manipulate the status of the system and to support diagnostics and prognosis.

In particular, we suppose digital twins deployed following the architecture proposed in Catarci et al. (2019). In Figure 1 an adaptation of this architecture to AgileChain is proposed.

Digital twins wrap physical entities involved in the process. These physical entities can be manufacturing machines or human operators. A digital twin exposes a Web API consisting, in general, of three parts: the synchronous one, the query interface and the asynchronous one. The synchronous interface allows one to give instructions to the physical entity. These instructions may, for example, produce a state change in a manufacturing machine (in case the twin is over a machine) or ask a human operator to perform a manual task (in case the twin is over a manufacturing worker). The query interface allows for asking information to the physical entity about its state and related information; noteworthy, these latter can be obtained by applying diagnostic and prognostic functions results of machine learning. The asynchronous interface generates events available to subscribers. Available technologies for digital twins easily allow to define such interfaces.

It is important to note how manufacturing machines and human actors can be considered identical from the point of view of the offered API, e.g., human actors produce asynchronous events as well, for example generating alarms.

A coordinator, which is a specific information system, takes as input a set of twin specifications and a process plan, and orchestrate digital twins, and underlying machines, using exposed interfaces.

4. Flexible Data management

AgileChains aims to overcome the shortcomings of traditional information system integration solutions, which are too rigid to support the dynamism of an agile supply chain. The goal is an effective and efficient data management including a safe, efficient, quality-aware and reconfigurable exchange of data useful for both internal processes and data export.

4.1. Smart data lakes

Also considering that the digital twin paradigm is currently under investigation, there is still a gap in the companies between the management level and the production level (shopfloor). Hence, it is important that the adoption of digital twins focuses on the data management, especially in this era in which companies are collecting an enormous amount of data related to any aspects of the daily routines and storing them in different types of systems: from streaming systems, to DBMSs, data warehouses and NoSQL solutions. This resulted in the creation of polyglot solutions, whose innovative management is related to the three concepts of polystore, data lake, and dataspace.

In such a context, AgileChains aims to provide solutions to orchestrate data sources in a smart way. Data will be organized in a smart dataspace, where the heterogeneous data sources exchange data through mappings, i.e. declarative specifications describing how to convert data from one or more data sources into the format of one or more data sources. The dataspace will adhere to the polystore model, where each data source belongs to an island of information with a unique data access model that describes the kind of managed data, e.g., streaming data vs. static data, and the supported operators. Information of different types will be considered: numerical monitoring data, numerical data from other actors, information derived from open data or from social media (e.g., geolocated events).

Two key issues in this context are mapping specification and entity resolution. As to the first, while most currently available polystore solutions require manual pre-processing and cleansing of raw data before ingestion (no dynamic schema/mapping discovery), this can really become a critical bottleneck in large scale deeply heterogeneous and dynamic integration scenarios such as the one considered. Instead, AgileChains will propose (semi-)automatic mapping discovery techniques combining the strengths of recent approaches and adapting them to the digital twin context: pay-as-you-go mapping discovery with reduced up-front setup cost, thanks to user feedback exploited in mapping annotation and refinement (Belhajjame et al. (2013)), matching validation and confirmation (Jeffery et al. (2008)), and mapping discovery through user information need

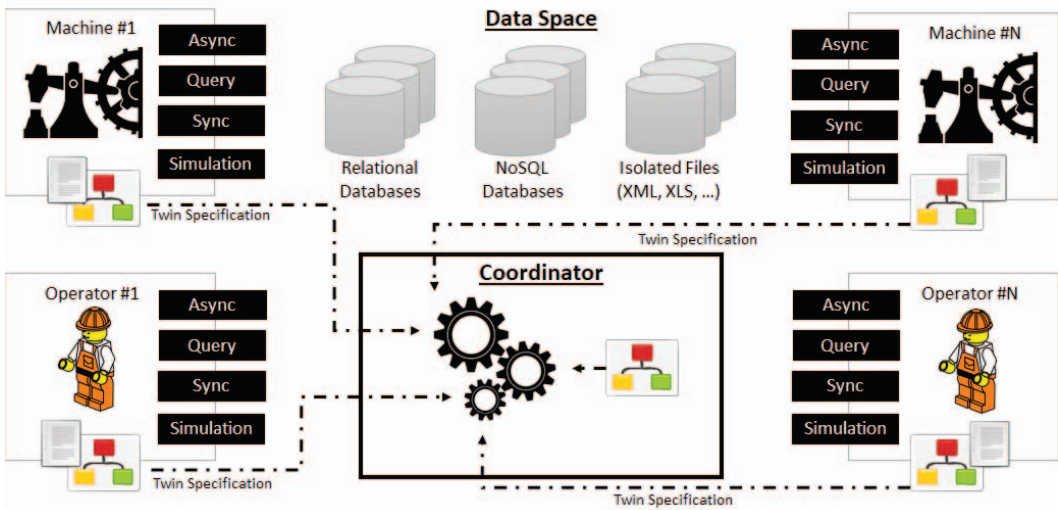


Fig. 1. The proposed architecture.

specification (Mandreoli (2017)). As to entity resolution, agility requirements ask for efficient techniques to resolve large amounts of data in limited time, therefore progressive entity resolution techniques will be considered (Simonini et al. (2018)).

The AgileChains smart data lake solution will thus ultimately include mapping discovery and entity resolution on demand algorithms, allowing to dynamically align and exchange data along the supply chain, also considering the data quality (Ofner et al. (2012)) as well as the respect of the data sovereignty (Jarke et al. (2019)).

4.2. Smart data movements

The concept of smart data movement has emerged in particular in the context of fog and edge computing. The availability of huge quantities of monitoring data enables a rapid detection of possible problems and performing predictions of the impact of decisions of a supply chain at a large scale. On the other hand, it is necessary to manage both available data and analytic computation to avoid system overloads. Several research works have explored the theme of data reduction, in particular with possible adaptations to the context of execution such as discussed in Andreolini et al. (2015), based on variable sampling rates.

More recently, approaches both to data and computation movement have emerged. In AgileChains, the approach proposed in Vitali et al. (2019) will be the basis for an adaptive approach to data and computation movement, based on the evaluation of the dependencies between data and the computations required on them, both reducing data compressing them semantically and allocat-

ing the needed computations and data reconstruction as close as possible to where they are needed, to reduce the global transmission and computation overload. Further research in AgileChains will investigate the issue of visibility of data across the AgileChains based on the above-mentioned approach, with the goal of providing significant, but synthesized data to other partners in the chain to facilitate interactions while preserving data sovereignty by the actors in the chain.

5. Agile design for risk-prone events

One of the goals in AgileChains is to provide methods and tools to make a supply chain robust in facing unforeseen events which could undermine the operation.

While, as seen in Sections 2. and 3., there are several approaches to provide adaptive and context-aware service-oriented system, the problem of operation disruption due to missing data has been studied less. The main focus in the literature has been on issues related to data quality in information systems (Batini and Scannapieco (2016)) and for big data applications (Ardagna et al. (2018)).

On the other hand, adaptivity to be able to cope with unexpected events is not limited to services. Another cause of disruption are missing data needed for the operations to be performed.

Therefore, in addition to provide adaptive service-oriented principles and flexible data management infrastructures, in AgileChains we pose a set of basic requirements for dealing with risk situations due to unavailability of data. The requirements are based on the following principles:

- a design time, the critical data needed for business process execution have to be identified.
- Alternative sources of data have to be identified, or possible workarounds designed.

In AgileChains the data resiliency approach developed in Marrella et al. (2019), a methodological approach based on different resiliency levels, is the basis for developing systems resilient to risk-prone events. The four defined levels are the following:

- level 0: no design for resiliency is performed.
- level 1: failure awareness, to be able to detect data-related failures.
- level 2: data resilience, when alternative sources of data are identified for the data which are critical for the business process.
- level 3: milestone resilience, when alternative milestones can be reached for a partial completion of a business process.
- level 4: process resilience, when a resilient execution of the process is considered during the design.

In AgileChains, the goal is to identify the requirements for resiliency in the supply chain, and to design the compensation actions to guarantee that services can be provided, even if at a lower level of performance, also when disruptions in the supply chain occur. The main compensation actions include the availability of alternative sources of data, alternative services, milestones, or workarounds in the process execution.

The relationship among adaptive processes in a service eco-system providing digital twins in combination with a flexible data management and resiliency levels defined by design is still posing open research problems in order to design a supply chain resilient to risk-prone events, anticipate them and design strategies for a preventive approach based on the defined consistency levels, without introducing excessive overhead or instabilities in the system.

6. Concluding remarks

In this paper we illustrated a conceptual framework to support the design and development of resilient agile support chains, able to adapt to varying circumstances and in particular to risk-prone events. The approach is based both on dynamic service-oriented service composition and a flexible data management approach. Finally, some methodological issues related to the design of agile supply chains are discussed. Future work will focus on the evaluation of the different design alternatives and on the development of design

support tools to develop risk-aware resilient chains, able to adapt to different types of adverse circumstances, focusing on the impact of major calamities such as the ones induced by climate changes or disruptive global events.

Acknowledgement

This work was supported by the EU H2020 program under Grant No. 734599 - FIRST project.

References

- Andreolini, M., M. Colajanni, M. Pietri, and S. Tosi (2015). Adaptive, scalable and reliable monitoring of big data on clouds. *J. Parallel Distrib. Comput.* 79-80, 67–79.
- Ardagna, D., L. Baresi, S. Comai, M. Comuzzi, and B. Pernici (2011). A service-based framework for flexible business processes. *IEEE Software* 28(2), 61–67.
- Ardagna, D., C. Cappiello, W. Samá, and M. Vitali (2018). Context-aware data quality assessment for big data. *Future Generation Comp. Syst.* 89, 548–562.
- Ardagna, D., M. Comuzzi, E. Mussi, B. Pernici, and P. Plebani (2007). PAWS: A framework for executing adaptive web-service processes. *IEEE Software* 24(6), 39–46.
- Arnold, K., R. Scheifler, J. Waldo, B. O'Sullivan, and A. Wollrath (1999). *Jini specification*. Addison-Wesley Longman Publishing Co., Inc.
- Batini, C. and M. Scannapieco (2016). *Data and Information Quality - Dimensions, Principles and Techniques*. Data-Centric Systems and Applications. Springer.
- Belhajjame, K., N. W. Paton, S. M. Embury, A. A. A. Fernandes, and C. Hedeler (2013). Incrementally improving dataspace based on user feedback. *Inf. Syst.* 38(5), 656–687.
- Catarci, T., D. Firmani, F. Leotta, F. Mandreoli, M. Mecella, and F. Sapio (2019). A conceptual architecture and model for smart manufacturing relying on service-based digital twins. In *2019 IEEE International Conference on Web Services, ICWS 2019, Milan, Italy, July 8-13, 2019*, pp. 229–236.
- Christopher, M. (2000). The agile supply chain: Competing in volatile markets. industrial marketing. *Management* 29(1), 37–44.
- Franklin, M., A. Halevy, and D. Maier (2005). From databases to dataspace: a new abstraction for information management. *ACM Sigmod Record* 34(4), 27–33.
- Gadepally, V., P. Chen, J. Duggan, A. Elmore, B. Haynes, J. Kepner, S. Madden, T. Mattson, and M. Stonebraker (2016). The bigdawg polystore system and architecture. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pp. 1–6. IEEE.
- Guttman, E. (1999). Service location protocol:

- Automatic discovery of ip network services. *IEEE Internet Computing* 3(4), 71–80.
- Jarke, M., B. Otto, and S. Ram (2019). Data sovereignty and data space ecosystems. *Bus Inf Syst Eng* 61, 549–550.
- Jeffery, S., M. Franklin, and A. Halevy (2008, 01). Pay-as-you-go user feedback for dataspace systems. pp. 847–860.
- LaPlante, A. and B. Sharma (2016). *Architecting Data Lakes*. O'Really.
- Mandreoli, F. (2017). A framework for user-driven mapping discovery in rich spaces of heterogeneous data. In *On the Move to Meaningful Internet Systems. OTM 2017 Conferences - Confederated International Conferences: CoopIS, C&TC, and ODBASE 2017, Rhodes, Greece, October 23-27, 2017, Proceedings, Part II*, pp. 399–417.
- Marrella, A., M. Mecella, B. Pernici, and P. Plebani (2019). A design-time data-centric maturity model for assessing resilience in multi-party business processes. *Inf. Syst.* 86, 62–78.
- Michelberger, B., R.-J. Andris, H. Girit, and B. Mutschler (2013). A literature survey on information logistics. In W. Abramowicz (Ed.), *Business Information Systems*, pp. 138–150. Springer Berlin Heidelberg.
- NESSI Europe (2018). Next generation software technologies empowering the digital transformation of europe.
- Ofner, M. H., B. Otto, and H. Österle (2012). Integrating a data quality perspective into business process management. *Business Process Management Journal* 18(6), 1036–1067.
- Plebani, P., D. García-Pérez, M. Anderson, D. Bermbach, C. Cappiello, R. I. Kat, F. Pallas, B. Pernici, S. Tai, and M. Vitali (2017). Information logistics and fog computing: The ditas* approach. In *Proceedings of the Forum and Doctoral Consortium Papers Presented at the 29th International Conference on Advanced Information Systems Engineering, CAiSE 2017, Essen, Germany, June 12-16, 2017*, pp. 129–136.
- Plebani, P. and B. Pernici (2009). URBE: web service retrieval based on similarity evaluation. *IEEE Trans. Knowl. Data Eng.* 21(11), 1629–1642.
- Reichert, M. and B. Weber (2012). *Enabling Flexibility in Process-Aware Information Systems - Challenges, Methods, Technologies*. Springer.
- Rosen, R., G. Von Wichert, G. Lo, and K. Bettenhausen (2015). About the importance of autonomy and digital twins for the future of manufacturing. *IFAC-PapersOnLine* 28(3), 567–572.
- Simonini, G., G. Papadakis, T. Palpanas, and S. Bergamaschi (2018). Schema-agnostic progressive entity resolution. *IEEE Transactions on Knowledge and Data Engineering* 31(6), 1208–1221.
- Suriadi, S., B. Weiß, A. Winkelmann, A. H. ter Hofstede, M. Adams, R. Conforti, C. Fidge, M. La Rosa, C. Ouyang, A. Pika, et al. (2014). Current research in risk-aware business process management—overview, comparison, and gap analysis. *Communications of the Association for Information Systems* 34(1), 52.
- Vitali, M., X. Peng, and B. Pernici (2019). An adaptive monitoring service exploiting data correlations in fog computing. In *Service-Oriented Computing - 17th International Conference, ICSOC 2019, Toulouse, France, October 28-31, 2019, Proceedings*, pp. 383–389.
- Zambonelli, F., A. Omicini, B. Anzengruber, et al. (2015). Developing pervasive multi-agent systems with nature-inspired coordination. *Pervasive and Mobile Computing* 17, 236–252.