

# A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web

Blue Sky Ideas Track

Andrei Ciortea  
University of St. Gallen  
St. Gallen, Switzerland  
Inria, Université Côte d’Azur, CNRS  
Sophia Antipolis, France  
andrei.ciortea@unisg.ch

Simon Mayer  
University of St. Gallen  
and ETH Zürich  
St. Gallen, Switzerland  
simon.mayer@unisg.ch

Fabien Gandon  
Inria, Université Côte d’Azur, CNRS  
Sophia Antipolis, France  
fabien.gandon@inria.fr

Olivier Boissier  
MINES Saint-Étienne, CNRS  
Saint-Étienne, France  
olivier.boissier@emse.fr

Alessandro Ricci  
University of Bologna  
Cesena, Italy  
a.ricci@unibo.it

Antoine Zimmermann  
MINES Saint-Étienne, CNRS  
Saint-Étienne, France  
antoine.zimmermann@emse.fr

## ABSTRACT

The World Wide Web has evolved drastically over the past decade – and the proliferation of Web APIs has turned it into the middle-ware of choice for most distributed systems. The recent focus on hypermedia-driven APIs together with initiatives such as the Web of Things and Linked Data are now promoting and advancing the development of a new generation of *dynamic*, *open*, and *long-lived* systems on the Web. These systems require agent-based solutions to the point that Web researchers have started to build autonomous systems on their own. It is thus both timely and necessary to investigate and align the latest developments in Web research and multi-agent systems (MAS) research. In this paper, we analyze in hindsight the factors that hindered the widespread acceptance of early Web-based MAS. We argue that the answer lies equally in a lack of practical use cases as well as the premature development and alignment of Web and agent technologies. We then present our vision for a new generation of autonomous systems on the Web, which we call *hypermedia MAS*, together with the research opportunities and challenges they bring.

## KEYWORDS

Hypermedia; Multi-Agent Systems; Web of Things; Semantic Web

### ACM Reference Format:

Andrei Ciortea, Simon Mayer, Fabien Gandon, Olivier Boissier, Alessandro Ricci, and Antoine Zimmermann. 2019. A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, Montreal, Canada, May 13–17, 2019, IFAAMAS, 5 pages.

## 1 WHERE ARE THE INTELLIGENT AGENTS?

The above question was posed by James Hendler in an editorial with the same title published in IEEE Intelligent Systems in 2007 [18]. After having spent the better half of a decade working on the

Semantic Web vision [4], Hendler was looking back to conclude that most ideas in the original article were already seeing widespread deployment on the Web, except for agent-based systems. One month later, McBurney and Luck [31] published a reply in the same journal in which they argued the question was misplaced and defended agent technologies as a future paradigm for mainstream computing – once *dynamic* and *open* systems become the norm in computing.

In this paper, we revisit Hendler’s question and analyze in hindsight – after another decade of Web and MAS research – why early Web-based MAS did not see widespread acceptance. We conclude that the answer lies equally in a lack of practical use cases as well as the premature development and alignment of Web and agent technologies. To support this claim, we contrast our findings with more recent results and developments that can help bridge the conceptual gap between MAS and the Web. We speculate that MAS are now at a turning point in their development and have all the prerequisites to move towards mainstream computing. The Web can provide a vessel for this transition, but a number of research and engineering challenges still have to be recognized and addressed.

## 2 A DECADE IN HINDSIGHT

The Web is already populated with billions of agents: we, as *human agents*, achieve most of our everyday goals by *browsing* and *acting on* Web resources. In pursuit of our goals, we are often assisted by invisible *software agents*, such as crawlers used by search engines to navigate and index the Web [6], agents that curate online content produced by people (e.g., Wikipedia’s content agents [33]), and recommender systems used all over the Web to generate more links and navigation paths (e.g., suggestions of books on an Amazon page). Although some of these agents may already use different AI methods (learning, reasoning, etc.), we have yet to see *more autonomous, cooperative, and long-lived* agents on the Web – the “intelligent agents” in Hendler’s question.

In this section, we search for new answers to Hendler’s question – a decade later. We start from the practical usefulness of MAS on the Web, and then continue with the evolution of Web and MAS research.

## 2.1 Availability of Practical Use Cases

Two key factors for the recent success of deep learning were the availability of large amounts of data and the advent of general-purpose computing on GPUs, which brought a speed-up of an order of magnitude when training neural networks [27]. Together, these two factors unlocked uses cases for (deep) neural networks that were simply not practical a couple of decades earlier.

We believe that recent developments in Web research are about to unlock similar opportunities for MAS – as *dynamic*, *open*, and *long-lived* systems become the norm on the Web. These properties of Web systems have already been identified as motivating factors for MAS research by Singh and Huhns [37] in the context of service-oriented computing. Still, in practice, Web systems flourished just fine without MAS, relying instead on developers to manually integrate Web APIs across service providers. This practice, however, is now becoming a pressing bottleneck given the explosive growth in number of Web APIs<sup>1</sup>. Researchers and practitioners are already looking for means to reduce the development and maintenance costs of Web systems, for instance by designing more autonomous, general-purpose clients that can cope with evolvable hypermedia-driven APIs [1], or by creating new paradigms and languages for programming more flexible Web systems (e.g., [9]) – some of which are based on reactive rule-based programming (e.g., [24, 39]).

In terms of application domains, McBurney and Luck argued in their reply to Hendler’s question that “the most compelling applications of agent technologies were always, are now, and will always be in business-to-business (B2B) domains, out-of-sight of most consumers” [31]. Even though business-to-consumer (B2C) applications of agent technologies might have been out of reach a decade ago, the recent *Web of Things* (WoT) initiative<sup>2</sup> unlocks a whole new area of practical uses cases for MAS *across* the B2C and B2B domains: as devices able to sense and act on physical environments are integrated into the Web’s architecture – from home and office appliances [29] to industrial robots [30] – the dynamics of Web systems increase to a point where existing programming paradigms become impractical and agent-based systems become *a necessity*. The motivation is so strong that WoT researchers have started to build autonomous systems on their own [26, 30]. A concrete application of autonomous WoT systems to the manufacturing domain was developed at Siemens and is presented in detail in [8].

In addition to becoming increasingly dynamic, Web systems are also becoming increasingly open – transcending the data silos exposed by various providers. Central to this evolution, the *Linked Data* initiative [5] promotes the uniform representation and access of data on the Web, and the discovery of data via hyperlinks – as suggested by the five-star Linked Data publishing principles<sup>3</sup>. As of June 2018, the Linked Open Data Cloud contains 1,229 datasets and 16,125 links.<sup>4</sup> More recent W3C standards such as the Linked Data Platform [38] and Linked Data Notifications [7] go beyond the mere exposure of datasets: they not only allow clients to *browse* Linked

Data, but also to reliably *observe* and *act on* it. This growing source of linked structured data on the Web now provides a breeding ground for MAS research that was not available a decade ago.

Another evolutionary trend that makes the Web more amenable to MAS is the growing interest in peer-to-peer interactions on the Web as an alternative to the traditional client-server interactions. Among the most prominent, such initiatives include WebRTC [3], the decentralization of the Semantic Web [40] (see also SOLID [28]), the WoT Architecture currently being standardized in the W3C WoT Working Group (see the WoT Servient in [25]), or the IPFS [2] initiative to build a versioned, peer-to-peer Web.

In summary, the above evolutionary trends promote and advance the development of a new generation of *dynamic*, *open*, and *long-lived* systems on the Web. These systems unlock practical use cases that stress the need for *autonomy* more than ever before. It is thus both timely and necessary to create thorough conceptual and technological bridges between MAS and Web research. To this end, in the next two sections we analyze the evolution of these two fields over the past decade in search of strategies for their integration.

## 2.2 Juxtaposition of MAS and the Web

Early attempts to integrate MAS and the Web would often just juxtapose the two – for instance, using MAS to manage Web resources or using Web standards as message formalisms in MAS [14], or using Web languages for limited implementations of autonomous Web agents [15]. Around the early 2000s, service-oriented computing took off – and the AAMAS community turned to integrating MAS and Web services [16, 20, 21, 37].

Web services have evolved drastically over the past decade: from an RPC-style, control-oriented paradigm to a REST-style, data-oriented paradigm [34]. The RPC-style paradigm would typically use the Web merely as a transport layer – for instance, to transport serializations of procedure calls via HTTP. Most prominently, this approach was used by the WS-\* standards (SOAP, WSDL, UDDI, etc.), which used HTTP to transport SOAP messages [32]. Today, it is widely accepted that using the Web merely as a transport layer entails important drawbacks in *open* and *long-lived* Web systems [12, 34]: systems that adhere to this paradigm are in juxtaposition with the Web and thus fail to inherit its architectural properties (e.g., scalability, loose coupling), make limited use of the existing Web infrastructure (e.g., for load balancing, caching) and its future extensions. In fact, the SOAP specification also defined a binding for sending messages via SMTP [32] completely outside of the Web, which is a clear sign of juxtaposition with the Web architecture.

The evolution of Web services has also influenced MAS research. The much closer relationship between the AAMAS and Web communities in the early 2000s, which was in turn due to the upcoming Semantic Web, led to a strong uptake of WS-\* services in MAS research [16, 20, 37]. FIPA also proposed a specification for using HTTP as a transport protocol for messages exchanged among agents [13], which was implemented by several MAS platforms (e.g., [10, 11, 17]). Many researchers in the AAMAS community still view the Web merely as a transport layer for messages in MAS, thereby remaining in juxtaposition with core principles of the Web architecture as it is conceived of today. This view seems to be the natural fit when considering MAS as composed only of agents that

<sup>1</sup><http://programmableweb.com/news/research-shows-interest-providing-apis-still-high/research/2018/02/23>, accessed: 16.11.2018.

<sup>2</sup>The WoT is currently being standardized through combined efforts of the W3C (<https://www.w3.org/WoT/WG/>), IETF (<https://datatracker.ietf.org/group/core/about/>), and IRTF (<https://datatracker.ietf.org/rg/t2trg/about/>).

<sup>3</sup><https://www.w3.org/DesignIssues/LinkedData.html>, accessed: 16.11.2018.

<sup>4</sup><https://lod-cloud.net/>, accessed: 16.11.2018.

exchange messages; however, when considering also the *environment* as a first-class abstraction in MAS [43], the Web’s natural fit is as an *application architecture* for MAS that is used not only for message transport, but for many types of interactions between agents and resources in their environment (see Section 2.3).

The idea that dynamic, open, and long-lived systems require the principled Web architecture as their foundation took over a decade to become mainstream (even within the Web community). However, today it is clear that this is the most promising pathway. Central to the Web architecture, *hypermedia* is now increasingly used for designing highly scalable, dynamic, open, and interoperable systems – for instance, in the W3C Linked Data Platform [38], the W3C WoT Thing Description (TD) [23], or in open specifications such as Hydra<sup>5</sup>. The research done in the Web community over the past decade is therefore highly relevant for MAS, and a low-level architectural integration between MAS and the Web could entail a trove of benefits for the AAMAS community.

### 2.3 MAS – More Than Just Agents

Turning our discussion from the Web to MAS, over the past decade agent-oriented software engineering has grown to recognize that MAS consist of more than just agents and thus should be designed on multiple dimensions. Communities such as *Environment for Multiagent Systems (E4MAS)*<sup>6</sup> [42] and *Coordination, Organizations, Institutions, and Norms in Agent Systems (COIN)*<sup>7</sup> have advanced the *environment* and the *organization* (respectively) as first-class abstractions in MAS. These conceptual dimensions open new perspectives on the conceptual integration of MAS and the Web.

If we conceive of MAS as being composed of just agents, then when designing Web-based MAS we obscure large parts of the systems: Web resources (and services) remain hidden behind agents. Furthermore, as already pointed out in Section 2.2, this perspective leaves little room for considering the Web as anything more than just a transport layer for messages exchanged between agents – which has been the conventional view in MAS research.

If in addition to agents we also conceive of the *environment* as a first-class abstraction in MAS – i.e., a key component designed and programmed with clear-cut responsibilities [35] – then we can achieve a *conceptual* integration of MAS and the Web. The Web is no longer just a hidden transport layer in MAS, but a visible application layer and a place for stigmergic interactions – a world-wide environment that can be designed and programmed for agents. Web resources regain their status as first-class abstractions situated in the agents’ environment, where agents can share, observe, reason, and act upon them. At the same time, MAS no longer remain outside of the Web (in juxtaposition), but they become part of the Web – and in the process inherit its architectural properties as a world-wide, open, and long-lived system.

The *organization* dimension is equally relevant to achieving a proper conceptual integration of MAS and the Web. The various normative and organizational models proposed by the COIN community over the past decade can now be used to represent norms, organizational structures, and social relations in the environment




Figure 1: Hypermedia MAS: agents are situated in an environment designed as a distributed hypermedia application.

(externally to agents). Governance and social enforcement mechanisms then allow to control and regulate the autonomous behavior of agents in Web-based MAS. The alternative of regimenting autonomous behavior by hard-coding rules into agents is impractical in a complex, world-wide, and open environment such as the Web.

## 3 RETHINKING WEB-BASED MAS

Following our analysis, it seems that MAS and Web research have evolved to a point where their conceptual integration is mostly a matter of alignment – one that comes naturally and to the benefit of both fields. In what follows, we first present our vision on how this alignment could be realized, and then discuss the research opportunities and challenges it brings.

### 3.1 Our Vision: Hypermedia MAS

We hypothesize that the *missing bridge* between MAS and the Web is a conceptual integration between the *environment dimension* in MAS and the *Web architecture*. This conceptual integration would then allow MAS to be seamlessly distributed across the Web, and in the process to inherit its architectural properties (e.g., scalability, evolvability). The key point, however, is that agent environments should not be layered on top of the Web (technological integration), but instead they should be integrated into the *hypermedia fabric* of the Web (conceptual integration).

We envision *hypermedia MAS* as socio-technical systems composed of *people* and *autonomous agents* – henceforth *agents* – situated in a shared *hypermedia environment* that is distributed across the open, world-wide Web. The environment is considered here as a first-class abstraction – an *information layer* that provides agents with various functionalities: it mediates access to the external environment (e.g., devices, digital services); it provides an abstraction layer for modeling, representing, and programming non-autonomous entities; it mediates interaction, communication, and coordination among agents. The environment is thus populated not only with agents, but also with non-autonomous entities (devices, digital services, knowledge repositories, organizations, datasets, etc.) that agents can discover and use in pursuit of their goals. In this vision, the Web provides the underpinning that interconnects *all entities* – within and across MAS.

Figure 1 illustrates our vision: the *agent environment* is designed as a *hypermedia application* that is distributed across the Web. All

<sup>5</sup><http://www.hydra-cg.com/spec/latest/core/>

<sup>6</sup><https://distrinet.cs.kuleuven.be/events/e4mas/>, accessed: 16.11.2018.

<sup>7</sup>[https://www2.pcs.usp.br/~coin/coin\\_springer.html](https://www2.pcs.usp.br/~coin/coin_springer.html), accessed: 16.11.2018.

entities (autonomous and non-autonomous) in the MAS are represented in the environment as *Web resources*. Note here that a Web resource is anything that might be identified by an IRI [22], which includes physical entities (people, things, etc.) and abstract entities (e.g., organizations). Web resources in agent environments could thus have one or more embodiments (physical, digital, and/or holographic<sup>8</sup>), and they could be interrelated through hyperlinks (cf. Figure 1) to support discoverability in open hypermedia MAS – similar to how hyperlinks enable the discovery of Web pages on the open Web. Agents could then navigate the hypermedia themselves to discover resources in the MAS, or they could use hypermedia search engines (cf. Figure 1) – an approach that has proven very successful for the open, world-wide Web.

From the agents’ viewpoint, hypermedia can enable a seamless distribution of MAS on the Web: similar to how people use hyperlinks and HTML forms to discover and interact with Web pages regardless of their location, autonomous agents could use *hyperlinks* and *hypermedia controls* to discover and interact with other entities (other agents, tools, etc.) regardless of their location. Furthermore, the use of hypermedia controls would allow autonomous agents to discover at runtime the affordances of entities in their environment, such as the operations exposed by a light bulb (cf. Figure 1). A similar approach is taken by the W3C WoT TD [23], which uses hypermedia controls to allow software clients to discover at runtime the affordances of WoT devices. Hypermedia can thus reduce coupling in MAS, which enhances the openness and evolvability of long-lived systems.

Hypermedia could also be used to advertise in the environment various resources that agents could discover and consume at runtime, such as specifications of interaction protocols (e.g., using BSPL [36]), specifications of organizations (e.g., using MOISE [19]), data licensing policies [41], etc. (cf. Figure 1). Such resources could be designed into the environment to further reduce coupling in MAS and thus to further enhance the openness and evolvability of the systems. Engineers in different parts of the world could then develop and deploy autonomous agents and other entities independently from one another, and old and new implementations could co-exist in one system.

## 3.2 Research Opportunities and Challenges

Large-scale systems of people and autonomous agents on the open Web raise a broad range of technical, social, and ethical challenges. In this section, we identify immediate challenges raised when weaving hypermedia MAS into the Web.

**Scaling up to the size of the Web:** The Web is already used by more than 3 billion people. Web-scale hypermedia MAS would also count billions of agents and other entities in world-wide environments. This perspective emphasizes topics of research such as reasoning, planning, interaction and coordination in large-scale MAS. But before anything can happen, an immediate problem is to enable autonomous agents to search Web-scale hypermedia MAS for resources required to achieve their goals. Search on the Web is still an active research domain, and hypermedia MAS raise challenges that are not addressed by traditional search engines: they are populated with non-textual resources that cannot be simply

indexed and ranked based on term frequency; autonomous agents would conduct searches for a broader range of purposes than mere information needs and would require structured query and result capabilities to fulfill their needs; hypermedia MAS are expected to evolve much more rapidly than the documentary Web (in particular in the context of the WoT); and agents can shape hypermedia environments by weaving unlimited overlays of hyperlink structures. It is then necessary to investigate *semantic, goal-directed, real-time* search in Web-scale hypermedia MAS. We need to consider the searching capacity required by agents on the Web and, inversely, the solutions agents can bring to the problem of searching the Web.

**Coping with the open Web:** In environments as open as the Web, enabling autonomous agents to “arrive and operate” raises new challenges, such as planning and acting on explicit or learned models of hypermedia environments, or the specification, evolution, and enactment of interaction protocols in open and long-lived hypermedia MAS. An immediate problem is the support of agent-oriented programming for semantic hypermedia environments. For instance, the W3C WoT TD [23] allows MAS engineers to use WoT ontologies and hypermedia controls to program their agents against semantic models of devices (rather than against individual devices). Such use cases stress the importance of integrating agent programming languages with methods for knowledge representation and reasoning on the Web, which can further reduce coupling between agents and their hypermedia environments – therefore enabling agents to better cope with open and long-lived hypermedia MAS.

**Governing MAS on the open Web:** In open MAS, little prior assumptions can be made about agents, their intentions and behavior. Governing hypermedia MAS thus raises challenges such as the identification and authentication of agents on the open Web, the representation and advertising of norms in hypermedia environments, the decentralization and distribution of mechanisms for regulating and coordinating the agents’ behavior, or explaining the behavior of autonomous agents. An immediate problem is to enable autonomous agents to reliably interpret and reason on Web-specific norms, data licensing policies, or terms of services that are imposed when entering and using hypermedia environments. Another pressing problem in hypermedia MAS is *open organization*: while it is common to assume agents in a given MAS are part of a single organization (and fall under a single coordination model), in open and long-lived hypermedia MAS it is reasonable to expect agents would participate in multiple organizations throughout their lifetime, and possibly in multiple organizations at the same time – where liaisons could emerge among some organizations, while other organizations could work towards cross-purposes.

## 4 CONCLUSIONS: A CALL TO ACTION

All the elements required to design and deploy hypermedia MAS are already available. Even simple designs of hypermedia MAS could already solve problems (e.g., in the WoT) – they would not yet have to be Web-scale or fully open, and regimentation can still be useful in many practical cases. However, the most important next step we see is to build mature open-source MAS technology for the (modern) Web. Such technology could encourage innovation and rapid prototyping of MAS within the AAMAS community, and would help promote MAS to other communities.

<sup>8</sup><https://www.w3.org/community/immersive-web/>, accessed: 16.11.2018.

## REFERENCES

- [1] M. Amundsen. 2017. *RESTful Web Clients: Enabling Reuse Through Hypermedia*. O'Reilly Media.
- [2] Juan Benet. 2014. IPFS - Content Addressed, Versioned, P2P File System. (2014). arXiv:1407.3561 <http://arxiv.org/abs/1407.3561>
- [3] Adam Bergkvist, Daniel C. Burnett, Cullen Jennings, Anant Narayanan, Bernard Aboba, Taylor Brandstetter, and Jan-Ivar Bruaroey. 2018. *WebRTC 1.0: Real-time Communication Between Browsers, W3C Candidate Recommendation 27 September 2018*. W3C Recommendation. World Wide Web Consortium (W3C). <https://www.w3.org/TR/2018/CR-webrtc-20180927/>
- [4] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The Semantic Web. *Scientific american* 284, 5 (2001), 34–43.
- [5] Christian Bizer, Tom Heath, and Tim Berners-Lee. 2011. Linked data: The story so far. In *Semantic Services, Interoperability and Web Applications: Emerging Concepts*, Amit Sheth (Ed.), IGI Global, 205–227. <https://doi.org/10.4018/978-1-60960-593-3.ch008>
- [6] Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems* 30, 1 (1998), 107–117. [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X) Proceedings of the Seventh International World Wide Web Conference.
- [7] Sarven Capadisi and Amy Guy. 2017. *Linked Data Notifications, W3C Recommendation 2 May 2017*. W3C Recommendation. World Wide Web Consortium (W3C). <https://www.w3.org/TR/2017/REC-ldn-20170502/>
- [8] Andrei Ciortea, Simon Mayer, and Florian Michahelles. 2018. Repurposing Manufacturing Lines on the Fly with Multi-agent Systems for the Web of Things. In *Proceedings of the 17th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '18)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 813–822. <http://dl.acm.org/citation.cfm?id=3237383.3237504>
- [9] Olivier Corby, Catherine Faron-Zucker, and Fabien Gandon. 2017. LDScript: A Linked Data Script Language. In *The Semantic Web – ISWC 2017*, Claudia d'Amato, Miriam Fernandez, Valentina Tamma, Freddy Lecue, Philippe Cudré-Mauroux, Juan Sequeda, Christoph Lange, and Jeff Heflin (Eds.), Springer International Publishing, Cham, 208–224.
- [10] Oğuz Dikenelli. 2008. SEAGENT MAS Platform Development Environment. In *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems: Demo Papers (AAMAS '08)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 1671–1672. <http://dl.acm.org/citation.cfm?id=1402744.1402758>
- [11] Jose Exposito, Joan Ametller, and Sergi Robles. 2010. Configuring the JADE HTTP MTP. <http://jade.tilab.com/documentation/tutorials-guides/configuring-the-jade-http-mtp/>. (2010). Accessed: 03.03.2019.
- [12] Roy Thomas Fielding. 2000. *Architectural styles and the design of network-based software architectures*. Ph.D. Dissertation. University of California, Irvine.
- [13] Foundation for Intelligent Physical Agents. 2002. FIPA Agent Message Transport Protocol for HTTP Specification. <http://www.fipa.org/specs/fipa00084/SC00084F.html>. (2002). Document number: SC00084F.
- [14] Fabien Gandon. 2002. *Distributed Artificial Intelligence And Knowledge Management: Ontologies And Multi-Agent Systems For A Corporate Semantic Web*. Theses. Université Nice Sophia Antipolis. <https://tel.archives-ouvertes.fr/tel-00378201>
- [15] Fabien Gandon. 2003. Combining reactive and deliberative agents for complete ecosystems in infospheres. In *IEEE/WIC International Conference on Intelligent Agent Technology (IAT)*. Halifax, Canada. <https://doi.org/10.1109/IAT.2003.1241082>
- [16] Nicholas Gibbins, Stephen Harris, and Nigel Shadbolt. 2003. Agent-based Semantic Web Services. In *Proceedings of the 12th International Conference on World Wide Web (WWW '03)*. ACM, New York, NY, USA, 710–717. <https://doi.org/10.1145/775152.775251>
- [17] Miguel Escrivá Gregori, Javier Palanca Cámara, and Gustavo Aranda Bada. 2006. A Jabber-based Multi-agent System Platform. In *Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS '06)*. ACM, New York, NY, USA, 1282–1284. <https://doi.org/10.1145/1160633.1160866>
- [18] J. Hendler. 2007. Where Are All the Intelligent Agents? *IEEE Intelligent Systems* 22, 3 (May 2007), 2–3. <https://doi.org/10.1109/MIS.2007.62>
- [19] Jomi F. Hübner, Jaime S. Sichman, and Olivier Boissier. 2007. Developing Organised Multiagent Systems Using the MOISE+ Model: Programming Issues at the System and Agent Levels. *Int. J. Agent-Oriented Softw. Eng.* 1, 3/4 (Dec. 2007), 370–395. <https://doi.org/10.1504/IJAOS.2007.016266>
- [20] Michael N Huhns. 2002. Agents as Web services. *IEEE Internet computing* 6, 4 (2002), 93.
- [21] Michael N Huhns and Munindar P Singh. 2005. Service-oriented computing: Key concepts and principles. *IEEE Internet computing* 9, 1 (2005), 75–81.
- [22] Ian Jacobs and Norman Walsh. 2004. *Architecture of the World Wide Web, Volume One, W3C Recommendation 15 December 2004*. W3C Recommendation. World Wide Web Consortium (W3C). <http://www.w3.org/TR/2004/REC-webarch-20041215/>
- [23] Sebastian Kaebisch and Takuki Kamiya. 2018. *Web of Things Thing Description, W3C Working Draft 21 October 2018*. W3C Recommendation. World Wide Web Consortium (W3C). <https://www.w3.org/TR/2018/WD-wot-thing-description-20181021/>
- [24] Tobias Käfer and Andreas Harth. 2018. Rule-based Programming of User Agents for Linked Data. In *Proceedings of the 11th International Workshop on Linked Data on the Web (LDOW) at the 27th Web Conference (WWW)*.
- [25] Kazuo Kajimoto, Matthias Kovatsch, and Uday Davuluru. 2018. *Web of Things (WoT) Architecture, W3C First Public Working Draft 14 September 2017*. W3C Recommendation. World Wide Web Consortium (W3C). <https://www.w3.org/TR/2017/WD-wot-architecture-20170914/>
- [26] M. Kovatsch, Y. N. Hassan, and S. Mayer. 2015. Practical semantics for the Internet of Things: Physical states, device mashups, and open questions. In *2015 5th International Conference on the Internet of Things (IOT)*. 54–61. <https://doi.org/10.1109/IOT.2015.7356548>
- [27] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521 (may 2015), 436. <https://doi.org/10.1038/nature14539><http://10.0.4.14/nature14539>
- [28] Essam Mansour, Andrei Vlad Sambra, Sandro Hawke, Maged Zereba, Sarven Capadisi, Aburrahman Ghanem, Ashraf Aboulnaga, and Tim Berners-Lee. 2016. A Demonstration of the Solid Platform for Social Web Applications. In *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 223–226. <https://doi.org/10.1145/2872518.2890529>
- [29] S. Mayer, N. Inhelder, R. Verborgh, R. Van de Walle, and F. Mattern. 2014. Configuration of smart environments made simple: Combining visual modeling with semantic metadata and reasoning. In *2014 International Conference on the Internet of Things (IOT)*. 61–66. <https://doi.org/10.1109/IOT.2014.7030116>
- [30] Simon Mayer, Dominic Plangger, Florian Michahelles, and Simon Rothfuss. 2016. UberManufacturing: A Goal-Driven Collaborative Industrial Manufacturing Marketplace. In *Proceedings of the 6th International Conference on the Internet of Things (IoT'16)*. ACM, New York, NY, USA, 111–119. <https://doi.org/10.1145/2991561.2991569>
- [31] Peter McBurney and Michael Luck. 2007. The agents are all busy doing stuff! *IEEE Intelligent Systems* 22, 4 (2007), 6–7.
- [32] Nilo Mitra and Yves Lafon. 2007. *SOAP Version 1.2 Part 0: Primer (Second Edition), W3C Recommendation 27 April 2007*. W3C Recommendation. World Wide Web Consortium (W3C). <http://www.w3.org/TR/2007/REC-soap12-part0-20070427/>
- [33] Sabine Niederer and José van Dijck. 2010. Wisdom of the crowd or technicity of content? Wikipedia as a sociotechnical system. *New Media & Society* 12, 8 (2010), 1368–1387. <https://doi.org/10.1177/1461444810365297> arXiv:<https://doi.org/10.1177/1461444810365297>
- [34] Cesare Pautasso, Olaf Zimmermann, and Frank Leymann. 2008. Restful Web Services vs. "Big" Web Services: Making the Right Architectural Decision. In *Proceedings of the 17th Intl. Conference on World Wide Web (WWW '08)*. ACM, New York, NY, USA, 805–814. <https://doi.org/10.1145/1367497.1367606>
- [35] Alessandro Ricci, Michele Piunti, and Mirko Viroli. 2011. Environment programming in multi-agent systems: an artifact-based perspective. *Autonomous Agents and Multi-Agent Systems* 23, 2 (2011), 158–192.
- [36] Munindar P. Singh. 2011. Information-driven Interaction-oriented Programming: BSPL, the Blindingly Simple Protocol Language. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 2 (AAMAS '11)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 491–498. <http://dl.acm.org/citation.cfm?id=2031678.2031687>
- [37] Munindar P Singh and Michael N Huhns. 2006. *Service-oriented computing: semantics, processes, agents*. John Wiley & Sons.
- [38] Steve Speicher, John Arwe, and Ashok Malhotra. 2015. *Linked Data Platform 1.0, W3C Recommendation 26 February 2015*. W3C Recommendation. World Wide Web Consortium (W3C). <http://www.w3.org/TR/2015/REC-ldp-20150226/>
- [39] Steffen Stadtmüller, Sebastian Speiser, Andreas Harth, and Rudi Studer. 2013. DataFu: A Language and an Interpreter for Interaction with Read/Write Linked Data. In *Proceedings of the 22nd International Conference on World Wide Web (WWW '13)*. ACM, New York, NY, USA, 1225–1236. <https://doi.org/10.1145/2488388.2488495>
- [40] Ruben Verborgh, Tobias Kuhn, and Tim Berners-Lee. 2018. *Proc. 2nd Workshop on Decentralizing the Semantic Web co-located with the 17th International Semantic Web Conference (ISWC 2018)*. Vol. 2165.
- [41] Serena Villata and Fabien Gandon. 2012. Licenses Compatibility and Composition in the Web of Data. In *Proceedings of the Third International Conference on Consuming Linked Data - Volume 905 (COLD'12)*. CEUR-WS.org, Aachen, Germany, Germany, 124–135. <http://dl.acm.org/citation.cfm?id=2887367.2887378>
- [42] Danny Weyns and Fabien Michel. 2015. Agent Environments for Multi-agent Systems – A Research Roadmap. In *Agent Environments for Multi-Agent Systems IV*, Danny Weyns and Fabien Michel (Eds.), Springer International Publishing, Cham, 3–21.
- [43] Danny Weyns, Andrea Omicini, and James Odell. 2007. Environment as a first class abstraction in multiagent systems. *Autonomous agents and multi-agent systems* 14, 1 (2007), 5–30.