

Hypermedia to connect them all: Autonomous hypermedia agents and socio-technical interactions

Simon Mayer^{1,2}  | Andrei Ciortea³ | Alessandro Ricci⁴ | Maria Ines Robles⁵ |
Matthias Kovatsch⁶ | Angelo Croatti⁴

¹Pro2Future GmbH, Austria

²Institute for Technical Informatics, Graz
University of Technology, Austria

³CNRS Laboratoire Hubert Curien UMR 5516,
MINES Saint-Étienne, Université de Lyon,
Saint-Étienne, France

⁴Department of Computer Science and
Engineering, University of Bologna, Cesena, Italy

⁵Ericsson Research, Jorvas, Finland

⁶Siemens AG, Corporate Technology, Munich,
Germany

Correspondence

Simon Mayer, Pro2Future GmbH, Inffeldgasse
25F/I, 8010 Graz, Austria.

Email: simon.mayer@pro2future.at

Funding Information

Österreichische Forschungsförderungsgesellschaft,
COMET Centre Pro2Future.

Current standardization efforts of the Web of Things provide a unique opportunity to integrate technologies from the research domain of multiagent systems and the human-computer interaction field that could take us closer to creating intuitively usable distributed artificial intelligence. In this paper, we discuss what bridges are missing between these largely detached research communities: we propose to use the Hypermedia As The Engine Of Application State principle that is at the core of the Web architecture to underpin systems that integrate multiagent planning and acting with semantic technologies and with interoperable mixed reality interfaces. This should enable the creation of highly augmented environments in private as well as commercial and industrial environments where physical and digital things coexist and interact with one another.

KEYWORDS

HATEOAS, hypermedia, Internet of Things, mixed reality, multiagent systems, Web of Things

1 | INTRODUCTION

The World Wide Web is the most scalable and versatile information system deployed on the Internet to date. Since its creation in the early 1990s, the Web emerged as *the application layer of choice* for the Internet: it now pervades the everyday lives of almost 4 billion Internet users (according to the International Telecommunication Union's *ICT Facts and Figures 2017*), and it has become a de facto industry standard for exposing digital services on the Internet. The Web's huge success is justified by its architectural properties, which were specifically designed to meet the needs of an *Internet-scale* and *long-lived* system in which components can be deployed and can evolve independently from one another at runtime.¹ To this end, a central feature that distinguishes the Web architecture from other software architectures is the use of *hypermedia* to reduce coupling between components, formulated as an architectural constraint known as *Hypermedia As The Engine Of Application State (HATEOAS)*.¹ Hypermedia is thus used to create a homogeneous information fabric that provides an abstraction layer on top of component implementation details, allowing clients to discover at runtime how to interact with servers.

Over the past decade it has also been shown that even *class 1 devices** with as little as 100 KiB of Read-Only Memory (ROM) and 10 KiB of Random Access Memory (RAM) can be integrated into the Web,² thereby extending the Web as an application layer for the Internet of Things (IoT), commonly referred to as the Web of Things (WoT).³ Once IoT devices are integrated into the Web, developers can use standard Web technologies to mash-up digital services that perceive and act on the physical world⁴—the WoT thus fosters innovation and rapid prototyping in the IoT. However, in recent years it became apparent that current programming paradigms for Web development have shortcomings when it comes to engineering IoT systems in scientifically and economically important domains, such as flexible

manufacturing planning for mass-customized products or when integrating and configuring personal and industrial IoT environments that are made up of heterogeneous devices⁵: static service mash-ups cannot adapt to dynamic environments in which services can become available or unavailable at runtime, and manually mashing up services in the IoT does not scale.⁶

Many of the research questions underlying the engineering of autonomous systems have already been explored to a large extent in the scientific literature on artificial intelligence and, in particular, on autonomous agents and multiagent systems (AAMAS). AAMAS research provides multiple models and technologies that can be readily applied for prototyping autonomous WoT systems.⁵ At the same time, the AAMAS community has done extensive research on integrating multiagent systems (MAS) and Web services^(7,8)—with more recent efforts turning toward RESTful Web services,⁹ linked data,¹⁰ and hypermedia.¹¹

In yet another largely separate research area, the human-computer interaction (HCI) community has early-on identified mixed reality (MR) as a key technology for interleaving the digital and physical worlds.¹² More recently, MR received remarkable attention both in mainstream computing and WoT research as a means to enable the intuitive interaction of people with digital services in smart home and factory environments,¹³ and to create highly augmented WoT environments where physical and digital things coexist and interact with one another.¹⁴

We believe that synergies among these (largely detached) research communities—WoT, AAMAS, and HCI—can fundamentally transform Internet-based systems in the years to come, which holds great promise for advancing the development of information systems for augmenting human intellect,¹⁵ with wide-ranging consequences for systems that are—or will be—connected to the Web. We envision people and autonomous agents coexisting and augmenting one another in a worldwide sociotechnical system that spans across the physical-digital boundary. The potential applications of such an Internet-scale system would cut across society and could potentially revolutionize domains from industrial manufacturing to health care, academic research, and public administration. In this paper, we propose a pathway to achieve this vision and elaborate on the main challenges that must be overcome by integrating research across the mentioned communities.

2 | A HYPERMEDIA FABRIC FOR PEOPLE AND AUTONOMOUS AGENTS

In our vision, people and autonomous software agents—henceforth we refer to both as *agents*—coexist in a worldwide hypermedia environment composed of interconnected physical and digital *things*. In pursuit of their individual or common goals, agents interact with one another, take decisions, and consume and/or compose services exposed by things in their environment: they observe, reason about, and act on physical and digital things, but they can also create higher-level services by mashing-up those already available in their environment. Physical places are populated by both physical and digital things, where digital things can be exposed in the physical world through holographic representations and extensions. Agents can then interact with things using their physical, digital, or holographic interfaces.

The backbone of the envisioned sociotechnical system is the *hypermedia fabric* of the World Wide Web. This fabric has proven itself both in terms of scalability and flexibility¹—in our vision, it is used to create a homogeneous information fabric composed of interrelated resources that are distributed across heterogeneous servers and can have physical, digital, and holographic embodiments. Coupling between components is reduced to a minimum by relying on hypermedia and the HATEOAS constraint of REST, which enables their coevolution. Agents use the hypermedia fabric to discover and interact with one another, and with physical and digital things in their environment. Also Semantic Web standards (eg, RDF,¹⁶ OWL¹⁷) provide the technical means to enable agents to reliably process knowledge embedded in the hypermedia.

AAMAS research provides various models and technologies that can be applied to engineer large-scale systems of autonomous agents able to exploit the hypermedia fabric.⁵ At the agent level, goal-directed and human-based models of behavior (eg, the belief-desire-intention model¹⁸) and various agent programming languages (eg, JACK,¹⁹ Jadex,²⁰ Jason,²¹ 2APL²²) are available to ease the design, programming, but also inspection and debugging of *cognitive agents*: agents are able to deliberate on their own internal state (eg, their beliefs, desires, and intentions), the state of their environment, and the actions of other agents. At the system level, advanced AAMAS research topics—such as agent coordination, regulating autonomous behavior, or distributed decision making—investigate challenges relevant for enabling the development and deployment of worldwide, open systems of autonomous agents and people.

Finally, to support people in interacting with autonomous agents and/or physical and digital things, MR technologies provide the means to effectively dissolve the physical-digital boundary by bridging from the digital world into the physical world. Holograms can provide embodiments for digital entities, but they can also help physical things and even people overcome their physical limitations, for instance by enabling their dynamic extension with reconfigurable holographic components or by providing spatial and temporal synchronization of (remote) agents.

2.1 | Use cases

The practical value of an implementation of this vision is immediately evident in domains where services that are provided by heterogeneous digital and physical things need to be flexibly composed—given the increasing penetration of our environment with interconnected devices as well as the increasing reliance on digital services in the “cloud,” this is true for more and more aspects of our lives at home, on the road, and at the workplace.

For instance, in the *industrial domain*, it is becoming more and more important that *heterogeneous devices* can be interconnected and can *flexibly cooperate* to achieve higher-level goals⁶—the integration of technologies from MAS, the WoT, and MR has the potential to bridge the deployment gap for agent-based manufacturing systems and could even enable globally distributed autonomous manufacturing systems²³ that can include Web-enabled legacy components such as Programmable Logic Controllers.^{6†} These systems will need to implement manufacturing plans in a dynamic environment, requiring integrated planning and acting systems,²⁴ will need to be able to integrate newly discovered services, and will need to take their users into account as well: production engineers and machine operators will require an interface that enables them to configure, monitor, and control missions of the manufacturing system and that can be *used with little prior training*. They will also need to have an interface that informs them—in situ, *intuitively, and swiftly*—about events on the shop floor and actions that should be taken in response, and they should have a means of finding out *why devices behave in specific ways*.

Similar requirements hold true for commercial and private *smart buildings* where we are currently witnessing an exploding heterogeneity of connected devices, and across diverse fields ranging from *logistics* to *transportation* and *health care*. The described vision would furthermore have broad impact beyond technical systems, for instance in psychology and the social sciences: here, autonomous agents that can interact with the *Social Web* could help explore and raise awareness on social phenomena and social media dynamics that get reflected in the hypermedia structure, such as emotional contagion, self-reinforcing spirals²⁵ and group polarization (commonly referred to as *echo chambers*²⁶)—similar effects are at work in economics (eg, chaotic dynamics such as stock market bubbles and crashes²⁷). Finally, enabling agents to autonomously use the hypermedia fabric could revolutionize online goods and service markets by enabling customers to circumvent lock-in measures by providers: autonomous agents in the form of personal assistants could discover alternative offerings and dynamically and seamlessly enrich a customer’s online experience, thereby increasing consumer choice—and, crucially, competition between service providers.

2.2 | Requirements and challenges

Technology that implements this vision will need to fulfill a number of requirements that are implicitly or explicitly evident in the above use cases: it will need to provide an interoperability layer for globally distributed heterogeneous devices, services, and users on an Internet scale, and to support autonomous and flexible service mash-ups that integrate planning and the acting-out of inferred plans. The technology also needs to be usable: interacting with it should be intuitive and should not overload users; tools for monitoring, programming, and managing smart environments must be available and straightforward to use; and services that act autonomously must be able to efficiently explain to users why certain actions are taken. In the following, we give an overview of the main research challenges that arise in the context of goal-directed and flexibly interacting agents, and elaborate on how we believe they could be overcome within a combined effort of the WoT, AAMAS, and HCI communities.

2.2.1 | Hypermedia environments for autonomous agents

There has been extensive research on using the Web as an infrastructure for distributed MAS. However, most of the existing approaches use the Web merely as a transport layer and are thus misaligned with the Web architecture¹—this includes all MAS platforms that implement the FIPA specification[‡] for using HTTP as a message transport protocol as well as those approaches that implement the WS-* standards (SOAP, WSDL, UDDI, etc.). More recent approaches for Web-based MAS have turned to RESTful Web services,^{9,28} but they generally do not use hypermedia or HATEOAS. Recent exceptions include work by Dikinelli et al¹⁰ and the authors’ own work^{11,23}: the former uses the linked data principles (which partly reflect HATEOAS) and the latter investigate the use of hypermedia and HATEOAS for designing *agent environments*, making them first-class abstractions in MAS and promoting it as a conceptual bridge for deploying MAS on the Web. Furthermore, the use of hypermedia gives autonomous agents a means of traversing *local*—perhaps statically integrated—service mash-ups, and HATEOAS supports the flexible interaction between agents and their environment.[§] Research in this direction is, however, in a very early phase and further investigation is required, in particular on the efficient navigation of hypermedia environments (eg, for finding resources), or on the regulation of autonomous agents in hypermedia environments (eg, for enabling reasoning about nonfunctional aspects of Web services such as API rate limiting policies or data licensing information).

2.2.2 | Global orchestration and choreographies of hypermedia agents

Given MAS that are integrated with hypermedia systems, semantic technologies can provide the means to enable autonomous agents to reliably interpret, manipulate, and reason on hypermedia. On top of guidance regarding their local actions that is supported by HATEOAS, this ensures that agents are able to strive for *global* goals that involve mashing-up systems that are not already interconnected—this also includes the potential to mash-up agent *environments* themselves. To enable this, agreement is required on the terminology and semantics that provide functional and nonfunctional descriptions of the APIs (on top of agreement on the low-level API interactions that are subject of the aforementioned standardization efforts by W3C, IETF, and IRTF). To accomplish this, we recommend a bottom-up approach where this idea is applied to scenarios and projects with limited scope, but which feature—and expose—information models that are (1) based as much as possible on published models, (2) reuse established domain terminology (eg, from industrial standards), and (3) are equipped with open interfaces to preserve the potential to bridge across domains.³⁰

2.2.3 | Explainable agents

Especially when considering globally integrated service chains, autonomous agents will require a way of conveying *why* they selected a specific (past or future) action, not least for enabling people to build trust in autonomous distributed systems: in the same way that machine operators need to be able to retrieve information about the provenance of events in an industrial setting, inhabitants of smart homes should be able to find out why their devices are behaving in a certain way. A similar challenge has been identified for machine learning systems in general and researchers are in particular targeting the explainability of deep learning systems. Existing approaches already allow autonomous agents to log their (inter) actions, and to provide proofs and explanations regarding their behavior³¹; furthermore, since most people are familiar with *activity streams* in online social networks, we expect that it would be practical to extend this metaphor to *sociotechnical networks* that include hypermedia agents in order to provide a more natural way of interacting with people.^{32,33}

2.2.4 | Intuitive interoperable holograms and end-user programming

People will also require ways to efficiently and intuitively interact with autonomous agents and their hypermedia environments. For in situ monitoring and control, MR technologies are the obvious choice, as users should be equipped with a single interface to interact with physical as well as digital services. This interface shall enable users to specify goals for agents in their environment, to monitor agent behavior and interactions, and to combine provided functionality within higher-level service mash-ups—and the same interface can visualize mash-ups that have been derived by the systems automatically.

To this end, we need to ensure the interoperability of virtual environments across holographic renderers: users should not need to carry MR devices that define and visualize their private holographic mini-worlds; rather, head-worn, handheld, and ambient MR projectors should only be concerned with the *visualization itself*—and the content should be provided by a person's physical and digital environment, ideally reusing mechanisms that have proven successful in the Web. Efforts such as the Augmented Reality Markup Language (ARML)[¶] should be extended to create an HTML for MR: holographic browsers should display information based on agreed-upon protocols, ensuring that browsers and content can evolve independently; holograms should be able to link to one another and to other hypermedia content (for instance, based on their interactions within service mash-ups); and—similar to other content in the Web of documents, services, things, and people—holographic visualizations could be generated *automatically and dynamically* based on the underlying virtual and physical things and their metadata.

3 | THE ROAD AHEAD

The broader vision of information systems for augmenting human intellect has been articulated in various forms and under various terms, from Engelbart's seminal work in the mid-1960s¹⁵ to the Semantic Web movement and more recent work.^(5,27,34,35)

In this paper, we present our vision of sociotechnical systems of people and autonomous software agents whose interactions with each other and with their environment are coordinated through the use of hypermedia, and specifically through the HATEOAS principle that is one of the core tenets of the Web architecture. We observe that implementing this vision requires concerted action of, at the very least, the WoT, the AAMAS, and the HCI research communities: to enable autonomous agents to make full use of hypermedia for their interaction and regulation, core Web principles and novel standards for the WoT need to be considered by the AAMAS community and applied to agent environments²³; implementations, including industrial and research prototypes, need to adhere to common vocabularies whenever possible, and keep their underlying information models open to enable bridging scenarios, and domains; efforts that strive to give autonomous systems the ability to explain their actions need to interact with experts in the HCI domain to also make these explanations intelligible for end users; and the rendering

of holographic content should be decoupled from its generation and management, ideally reusing techniques that have proven successful for the global interlinking and rendering of documents.

We believe that the possibility of using hypermedia as the underpinning for sociotechnical systems of people and autonomous agents can have broad impact on a range of research areas within and beyond computer science—and on society as a whole, as people around the planet interact with hypermedia systems everyday: to access documents and online services, to interact with one another through online social platforms, and more recently to interact with physical devices in their surroundings.

NOTES

- * Terminology for Constrained-Node Networks, <http://www.ietf.org/rfc/rfc7228.txt>
- † For integrating legacy devices in more complicated situations, *reverse proxies* can provide a blanket solution (cf. <https://tools.ietf.org/html/draft-irtf-t2trg-rest-iot-00>)
- ‡ The *Foundation for Intelligent Physical Agents Message Transport Protocol for HTTP Specification*, <http://www.fipa.org/specs/fipa00084/SC00084F.html>
- § This is the subject of standardization efforts by the W3C (<https://www.w3.org/WoT/WG/>), IETF,²⁹ and IRTF (<https://datatracker.ietf.org/rg/t2trg/documents/>)
- ¶ See <http://www.opengeospatial.org/projects/initiatives/arpilot2017>

ORCID

Simon Mayer  <http://orcid.org/0000-0001-6367-3454>

REFERENCES

1. Fielding RT, Taylor RN. Principled design of the modern web architecture. *ACM Trans Internet Technol.* 2002;2(2):115-150.
2. Kovatsch M, Duquenooy S, Dunkels A. A low-power CoAP for Contiki. Paper presented at: 2011 IEEE 8th International Conference on Mobile Adhoc and Sensor Systems (MASS); 2011; Valencia, Spain: 855–860.
3. Wilde E. Putting things to REST. UCB iSchool Report 2007-015, Berkeley, CA: School of Information, UC Berkeley; 2007.
4. Guinard D, Trifa V, Pham T, Liechti O. Towards physical mashups in the web of things. Paper presented at: 2009 Sixth International Conference on Networked Sensing Systems (INSS); 2009; Pittsburgh, USA:1–4.
5. Ciortea A, Boissier O, Ricci A. Beyond physical mashups: autonomous systems for the web of things. Paper presented at: Eighth International Workshop on the Web of Things; 2017; Linz, Austria: 16–20.
6. Mayer S, Plangger D, Michahelles F, Rothfuss S. UberManufacturing: a goal-driven collaborative industrial manufacturing marketplace. Paper presented at: IoT'16 Proceedings of the 6th International Conference on the Internet of Things; 2016; Stuttgart, Germany:111–119.
7. Huhns MN, Singh MP. Service-oriented computing: key concepts and principles. *IEEE Internet Comput.* 2005;9(1):75-81.
8. Singh MP, Huhns MN. *Service-Oriented Computing: Semantics, Processes, Agents*. West Sussex, England: John Wiley & Sons; 2006.
9. Mitrović D, Ivanović M, Budimac Z, Vidaković M. Radigost: interoperable web-based multi-agent platform. *J Syst Softw.* 2014;90:167-178.
10. Dikaneli O, Alathi O, Erdur RC. Where are all the semantic web agents: establishing links between agent and linked data web through environment abstraction. Paper presented at: 4th International Workshop on Agent Environments for Multi-Agent Systems IV; 2015; Istanbul, Turkey:41–51.
11. Ciortea A, Boissier O, Zimmermann A, Florea AM. Give agents some REST: A resource-oriented abstraction layer for internet-scale agent environments. Paper presented at: Proc. AAMAS; 2017; São Paulo, Brazil:1502–1504.
12. Milgram P, Kishino F. A taxonomy of mixed reality visual displays. *IEICE Trans Inf Syst.* 1994;77(12):1321-1329.
13. Mayer S, Hodges J, Yu D, Diwold K. HoloMiracle: intuitive in-situ querying for industrial environments. Paper presented at: Proceedings of ISWC; 2017; Vienna, Austria.
14. Croatti A, Ricci A. Mashing up the physical and augmented reality: the web of augmented things idea. Paper presented at: Proceedings of WoT; 2017; Linz, Austria: 4–7.
15. Engelbart DC, English WK. A research center for augmenting human intellect. Paper presented at: AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference; 1968; Las Vegas, USA:395–410.
16. Cyganiak R, Wood D, Lanthaler M. RDF 1.1 concepts and abstract syntax. W3C Recommendation: World Wide Web Consortium (W3C); 2014.
17. W3C OWL Working Group. OWL 2 Web Ontology Language Document Overview. W3C Recommendation: World Wide Web Consortium (W3C); 2012.
18. Rao AS, Georgeff MP. BDI agents: from theory to practice. Paper presented at: Proceedings of the First International Conference on Multi-Agent Systems, vol. 95; 1995; San Francisco, USA:312–319.
19. Howden N, Rönquist R, Hodgson A, Lucas A. JACK intelligent agents - summary of an agent infrastructure. Paper presented at: Proceedings of AGENTS; 2001; Montreal, Canada.
20. Braubach L, Pokahr A, Lamersdorf W. *Jadex: a BDI-agent system combining middleware and reasoning*. Unland, R, Calisti, M, Klusch, M, eds. *Software Agent-Based Applications, Platforms and Development Kits*. Basel, CH: Springer; 2005:143-168.
21. Bordini RH, Hübner JF, Wooldridge M. *Programming Multi-Agent Systems in AgentSpeak Using Jason*. West Sussex, England: John Wiley & Sons; 2007.
22. Mehdi D. 2APL: a practical agent programming language. *Auton Agents Multiagent Syst.* 2008;16(3):214-248.
23. Ciortea A, Mayer S, Michahelles F. Industry 4.0: repurposing manufacturing lines on-the-fly with multi-agent systems for the web of things. Paper presented at: Proceedings of the 17th International Conference on Autonomous Agents and Multiagent System; 2018; Stockholm, Sweden.
24. Ghallab M, Nau D, Traverso P. *Automated Planning and Acting*. New York, USA: Cambridge University Press; 2016.
25. Slater MD. Reinforcing spirals model: conceptualizing the relationship between media content exposure and the development and maintenance of attitudes. *Media Psychol.* 2015;18(3):370-395.

26. Del Vicario M, Vivaldo G, Bessi A, et al. Echo chambers: emotional contagion and group polarization on Facebook. *Sci Rep.* 2015;6:37825.
27. Bernstein A, Klein M, Malone TW. Programming the global brain. *Commun ACM.* 2012;55(5):41-43.
28. Abdelkader G, Michael B. *REST-A: an agent virtual machine based on REST framework.* Demazeau, Y, Dignum, F, Corchado Rodríguez, JM, Bajo, J, eds. *Advances in Practical Applications of Agents and Multiagent Systems.* Springer Verlag Berlin Heidelberg: Springer; 2010:103-112.
29. Ishaq I, Carels D, Teklemariam GK, et al. IETF standardization in the field of the internet of things (IoT): a survey. *J Sens Actuat Netw.* 2013;2(2):235-287.
30. Mayer S, Hodges J, Yu D, Kritzler M, Michahelles F. An open semantic framework for the industrial internet of things. *IEEE Intell Syst.* 2017;32(1):96-101.
31. Broekens J, Harbers M, Hindriks K, Bosch K, Jonker C, Meyer J-J. Do you get it? User-evaluated explainable BDI agents. Paper presented at: German Conference on Multiagent System Technologies; 2010; Leipzig, Germany: 28-39.
32. Ciortea A, Zimmermann A, Boissier O, Florea AM. Towards a social and ubiquitous web: a model for socio-technical networks. Paper presented at: 2015 IEEE/WIC/ACM International Conference on Web Intelligence and Intelligent Agent Technology (WI-IAT); 2015; Singapore: 461-468.
33. Mayer S, Wilde E, Michahelles F. A connective fabric for bridging internet of things silos. Paper presented at: 2015 5th International Conference on the Internet of Things (IOT); 2015; Seoul, South Korea:148-154.
34. Gandon F, Buffa M, Cabrio E, et al. Challenges in bridging social semantics and formal semantics on the web. Paper presented at: International Conference on Enterprise Information Systems; 2013; Angers, France: 3-15.
35. Rovatsos M. Multiagent systems for social computation. Paper presented at: Proceedings of AAMAS; 2014; Paris, France:1165-1168.

How to cite this article: Mayer S, Ciortea A, Ricci A, Robles MI, Kovatsch M, Croatti A. Hypermedia to connect them all: Autonomous hypermedia agents and socio-technical interactions. *Internet Technology Letters* 2018;1:e50. <https://doi.org/10.1002/itl2.50>