PhD Proposal
Regulation in Hybrid Communities
of People and Agents on the Web

Olivier Boissier, Serena Villata

April 2020

1 Context

The Web is pervasive, increasingly populated with interconnected data, services, people and things [Kov+20]. To empower people through sophisticated *social machines* [HM16], the Web has to enlarge its scope to address hybrid communities [Gan+13] in which intelligent agents (e.g., crawlers, bots, recommenders) collaborate with people as peers to help them cope with the growing number of available resources. As the Web continues to grow, it will continue to unlock practical uses for increasingly autonomous, cooperative and long-lived agents [Cio+19]. In this context, hypermedia will enable flexible autonomous interaction among agents, people, and their environment. To ensure social order and global coherent behaviour, it is of first importance to complement this powered engine for hypermedia-driven interaction with regulation mechanisms to enforce and constrain autonomous behaviour in desired directions while keeping alive autonomy.

2 Objective

In this context, the objective of this PhD project is to address the regulation of hybrid communities in the context of the Web, which is the most scalable, flexible, and human-centric software system deployed ever. The Web was specifically designed to be an Internet-scale and long-lived system in which components can be deployed and can evolve independently from one another at run time [FT02]. Thus, the high degrees of *openness*, *heterogeneity*, *long-livedness*, and the *immense scale* of the Web pose unique challenges that have to be considered when designing the regulation of hybrid communities on the Web.

Regulation is a multi-faceted concept, which is studied in several domains. In the Web domain, regulation has generally been addressed by *policies* targeting the individual level, and in multi-agent systems [Boi+13], by *norms*, a kind of agreed-upon policies targeting the social level.

Policies on the Web have been discussed for a long time now. The developed approaches are mainly targeting representations that support static verification and automated reasoning. As stated in [Dag18] such representations have different meanings in relation to security and privacy [FG06; KFJ03], access and control [CVG12; KMD17], adaptable and context-aware systems, for instance in order to control behaviour [Ton+03] or to express legal knowledge (terms and conditions of licenses) [Ben+05] (e.g., ODRL[Ian+15]). In the multi-agent domain, regulation has been studied to large extent in research on normative approaches [And+13], where regulative norms [Oss12] or prescriptions [Bal+13] are generally used to specify who does what, in what context and as subject to what deontic modality (e.g., obligation, prohibition, permission). Regulative norms affect the agent behaviour in an indirect manner so that agents can autonomously decide to conform or not to the prescriptions or the norms for instance balancing internal motivation and desires versus the external consequences of violating a regulative norm. While the representations are certainly less rich than those proposed in the Web communities, these works have laid out architectures and mechanisms for the monitoring and enforcement of the agents' behaviour.

To the best of our knowledge, there doesn't exist yet a model that combines both approaches to address the challenges raised by the development of hybrid communities on the Web so that agents and people are regulated in the right *way* and *amount*, with a proper balance between empowerment and constraints in their interactions on the Web.

3 Main activities

We aim to address the above mentioned issues through an innovative approach for designing regulation mechanisms that are domain- and application- independent, that support and advertise heterogeneous representations of regulations, that consider overlapping and possibly conflicting regulations (assuming a central authority no longer holds in the context of the Web), and that support the decentralized enactment and enforcement of regulations.

The PhD thesis proposal takes place in the context of the ANR-SNF *HyperAgents* project that brings together internationally recognized researchers actively contributing to research on autonomous agents and MAS, the Web architecture, Semantic Web, and to the standardization of the Web, in order to define a new class of MAS that uses *hypermedia* as a general mechanism for uniform interaction among the agents and people on the Web.

The main tasks of this PhD thesis will consist in the creation of hypermedia-based mechanisms for the regulation of hybrid communities on the Web by defining the models and related mechanisms to:

- *Advertise* and *discover* regulation entities in Hypermedia MAS: the definition of the mechanisms to allow people and autonomous agents to discover *regulation entities* that they can interact with (e.g., policies, regulative in-

stitutions of agents and/or people, tools, or services[1]), the possible means of regulations that may apply to their current context. These discovered entities will carry semantic descriptions in order to support automated reasoning in open, heterogeneous and Internet-scale systems.

- *Monitor* and *enforce* regulation in Hypermedia MAS: the creation of the mechanisms to enact regulations in Hypermedia MAS so that the monitoring and enforcement of expected behaviour of people and autonomous agents are possible in the open, Internet-scale context of the Web.

- *Support the evolution* and *manage the life-cycle* of regulations: the definition of the models and mechanisms to instrument the life cycle of regulations. The aim is to define the proper tools that would allow regulation in Hypermedia MAS to "self-evolve". That is to say that people (e.g., developers, users) and autonomous agents will be able to publish regulations and to modify them to better fit their goals so that regulations adapt and evolve in systems that are long-lived.

Cross-cutting the design of these mechanisms, we will consider the challenging issue of balancing transparency with the necessity to enforce security and preserve privacy.

4 Required skills

- Master degree in Computer Science or Computer Engineering is required.

- General knowledge of software engineering and programming skills.

- Basic knowledge of Multi-Agent Systems, Web architecture and/or Semantic Web is preferred.

- Fluent English required, both oral and written.

5 Benefits package

- Working in an international project with international researchers.

- Exchange research stays with the other project partners.

- Ability to shape a possibly emerging field of research around hypermedia-based MAS.

- Subsidized catering service.

- Partially-reimbursed public transport.

- Sport facilities.

---

[1]e.g., `http://licentia.inria.fr`

6 Terms of the PhD position

Duration: 3 years.
Hosting teams:

- Intelligent Systems and Informatics Dpt at Henri Fayol institute , Mines Saint-Etienne & Connected Intelligence research team of the Laboratoire Hubert Curien UMR 5516 CNRS. The research fields of this team are knowledge representation and reasoning, multi-agent systems to model and support inter-connected communities of people, things and digital services in cyber-physical systems.

- WIMMICS is a research team of Université Côte d'Azur - INRIA. The research fields of this team are graph-oriented knowledge representation, reasoning and operationalization to model and support actors, actions and interactions in web-based epistemic communities.

Location: Mines Saint-Etienne, Saint-Etienne, France. Regular visits to the WIMMICS team in Nice.

Salary: 1700 euros/month (net salary).

Applications : a curriculum vitae together with a motivation letter, the Master / engineering school transcripts should be sent to Olivier Boissier (Olivier.Boissier@emse.fr) and Serena Villata (villata@i3s.unice.fr).

Deadline for applications: June 15th, 2020

Keywords: Multi-Agent Systems, Semantic Web Technology, Norms, Policies

References

[FT02]   Roy T. Fielding and Richard N. Taylor. "Principled design of the modern Web architecture". In: *ACM Trans. Internet Techn.* 2.2 (2002), pp. 115–150. DOI: 10.1145/514183.514185. URL: https://doi.org/10.1145/514183.514185.

[KFJ03]   Lalana Kagal, Tim Finin, and Anupam Joshi. "A policy based approach to security for the semantic web". In: *International semantic web conference*. Springer. 2003, pp. 402–418.

[Ton+03] Gianluca Tonti et al. "Semantic Web languages for policy representation and reasoning: A comparison of KAoS, Rei, and Ponder". In: *International Semantic Web Conference*. Springer. 2003, pp. 419–437.

[Ben+05]   V Richard Benjamins et al. "Law and the semantic web, an intro-
            duction". In: *Law and the Semantic Web*. Springer, 2005, pp. 1–
            17.

[FG06]   Carlos Flavián and Miguel Guinalıu. "Consumer trust, perceived
            security and privacy policy: three basic elements of loyalty to a
            web site". In: *Industrial Management & Data Systems* 106.5 (2006),
            pp. 601–620.

[CVG12]   Luca Costabello, Serena Villata, and Fabien Gandon. "Context-
            Aware Access Control for RDF Graph Stores". In: *ECAI - 20th
            European Conference on Artificial Intelligence - 2012*. Montpellier,
            France, Aug. 2012. URL: `https://hal.inria.fr/hal-00724041`.

[Oss12]   Sascha Ossowski. *Agreement technologies*. Vol. 8. Springer Science
            & Business Media, 2012.

[And+13]   Giulia Andrighetto et al. *Normative multi-agent systems*. Vol. 4.
            Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2013.

[Bal+13]   Tina Balke et al. "Norms in MAS: definitions and related concepts".
            In: *Dagstuhl Follow-Ups*. Vol. 4. Schloss Dagstuhl-Leibniz-Zentrum
            fuer Informatik. 2013.

[Boi+13]   Olivier Boissier et al. "Norms, organizations, and semantics". In:
            The Knowledge Engineering Review 28.1 (2013), pp. 107–116.

[Gan+13]   Fabien Gandon et al. "Challenges in bridging social semantics and
            formal semantics on the web". In: *International Conference on En-
            terprise Information Systems*. Springer. 2013, pp. 3–15.

[Ian+15]   R. Iannella et al. *ODRL: Open Digital Rights Language 2.1. Tech-
            nical report*. Tech. rep. W3C Community Group, 2015.

[HM16]   James Hendler and Alice M Mulvehill. *Social machines: the coming
            collision of artificial intelligence, social networking, and humanity*.
            Apress, 2016.

[KMD17]   Sabrina Kirrane, Alessandra Mileo, and Stefan Decker. "Access con-
            trol and the resource description framework: A survey". In: *Seman-
            tic Web* 8.2 (2017), pp. 311–352.

[Dag18]   Enrico Daga. "Knowledge Components and Methods for Policy Prop-
            agation in Data Flows". PhD thesis. The Open University, 2018.

[Cio+19]   Andrei Ciortea et al. "A Decade in Hindsight: The Missing Bridge
            Between Multi-Agent Systems and the World Wide Web". In: *Pro-
            ceedings of the 18th International Conference on Autonomous Agents
            and MultiAgent Systems*. International Foundation for Autonomous
            Agents and Multiagent Systems. 2019, pp. 1659–1663.

[Kov+20]   Matthias Kovatsch et al. *Web of Things (WoT) Architecture, W3C
            Recommendation 9 April 2020*. W3C Recommendation. World Wide
            Web Consortium (W3C), Apr. 2020. URL: `https://www.w3.org/
            TR/2020/REC-wot-architecture-20200409/`.