



HAL
open science

Insights about user-centric contextual online adaptation of coordinated multi-agent systems in smart homes

Iago Felipe Trentin, Olivier Boissier, Fano Ramparany

► To cite this version:

Iago Felipe Trentin, Olivier Boissier, Fano Ramparany. Insights about user-centric contextual online adaptation of coordinated multi-agent systems in smart homes. Rencontres des Jeunes Chercheurs en Intelligence Artificielle 2019, Jul 2019, Toulouse, France. pp.35-42. hal-02160421

HAL Id: hal-02160421

<https://hal.science/hal-02160421>

Submitted on 21 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Insights about user-centric contextual online adaptation of coordinated multi-agent systems in smart homes

Iago Felipe Trentin^{1,2}

Olivier Boissier²

Fano Ramparany¹

¹ Orange Labs, Meylan, France 38240

² Univ. Lyon, MINES Saint-Étienne, CNRS Lab Hubert Curien UMR 5516, Saint-Étienne, France F-42023

iagofelipe.trentin@orange.com / iago.felipe.trentin@gmail.com

Résumé

Avec le développement des maisons intelligentes, il est de plus en plus nécessaire de déployer des systèmes capables de s'adapter à des contextes fortement dynamiques ainsi qu'aux besoins des utilisateurs. Ces maisons sont de plus en plus équipées d'objets connectés qui manquent souvent de l'intelligence nécessaire à leur adaptation et coordination en ligne. Cet article a pour objectif d'analyser des approches susceptibles de répondre à ces défis et d'identifier des directions possibles de travaux futurs pour développer des systèmes multi-agents capables de combiner planification et action, en étant sensibles au contexte et dirigés par des buts, et qui pourront inspirer les générations futures de maisons intelligentes.

Mots Clef

Systèmes multi-agent, Planification, Action, Coordination, Adaptation, Maison intelligente

Abstract

Smart homes are becoming a reality, and with that the need to deploy systems capable of adapting to highly dynamic contexts in order to satisfy user needs. These homes are increasingly permeated by internet of things objects, but often they are not embedded with the artificial intelligence necessary for their adaptation and coordination at run time. This paper aims to present promising approaches for these challenges, and identify possible directions for future work in developing a coordinated responsible goal-driven context-aware planning and acting multi-agent system, which may inspire future generations of smart homes.

Keywords

Multi-agent systems, Planning, Acting, Coordination, Adaptation, Smart home

1 Introduction

With the advance of ubiquitous technologies, the physical objects and environments of our daily life are becoming more intelligent, capable of perceiving and acting in

the world. But for the time being, such technologies have difficulties adapting to highly dynamic environments such as smart homes, and users also struggle to make different objects work together. This paper is motivated by a current smart home challenge: the ability to adapt to context in order to satisfy user evolving needs. Such complex behavior can only be achieved by embedding Artificial Intelligence (AI) techniques into Internet of Things (IoT) objects.

To tackle the aforementioned challenge, there are solutions in the literature that allow embedding intelligent entities (a.k.a. agents) with: acting capabilities, i.e. act in the world and react to predicted changes; and also planning capabilities, i.e. reason upon current context to find new solutions to a given need. In a smart home, for maximum user satisfaction, agents should be able to timely adapt to user needs, with minimum conscious and direct input from occupants. This online adaptation may be possible by endowing agents with planning and acting functions designed to execute together at run-time.

In this paper, we aim to bring up and renew questionings about adaptation capabilities of IoT objects, specifically focusing on online adaptation of plans and actions in a user-centric manner that respects user goals and user values. We analyze current state-of-the-art solutions, organizing them according to common properties we find in the literature. The understanding and exploring of the state-of-the-art approaches presented may allow future work to create a decentralised and coordinated planning and acting system. It would be: *user-centric*, gravitating towards enhancing human experience; *goal-driven*, intelligently reasoning about and pursuing user objectives; *data-driven*, adapting at run-time to environmental and user inputs; and *trustworthy*, being consistent with user values and expectations.

Based on these challenges we analyze the state of the art as follows: Section 2 presents the challenges we are interested in, with special focus on the smart home domain; Section 3 presents well-established solutions, especially in the planning and acting domains, that have potential to help addressing our challenges; Section 4 shows the interest in uniting planning and acting to ultimately have a synergistic

solution; Section 5 identifies complementary domains, the goal and value reasoning, that add valuation to such a solution; Section 6 expands the envisioned solution to consider interaction and coordination in a decentralised system.

2 Motivation

This section elaborates the above stated challenges and interest in developing a system capable of adapting itself to context and users, in IoT environments such as smart homes. We consider publicly available solutions, their current state and limitations, that partially fulfill identified smart home requirements. It also shows aspects that can be improved or redesigned to better address our requirements. Since it explores a large domain, this paper is not expected to be an exhaustive study of the available literature.

2.1 Current IoT solutions

Current smart homes are becoming increasingly more complex, permeated by ubiquitous and heterogeneous elements designed by different vendors [1]. Some solutions help with discovery of objects and composition of new services mixing available ones [1, 2], other popular solutions allow the development of IoT mashups¹. For a great smart home user experience, the composition and delivery of user-centric services are important.

Although there is interest in ad-hoc mashups, often temporary and non-scalable, for small tactical and specific needs [3], such applications do not present reasoning nor adaptive behavior. They simply perform tasks previously programmed by humans. Also, with discovery and composition of services but no intelligent entities to demand and use them, user experience tends to become cumbersome – due to constant need for manual actions, like system re-configuration, user input, etc.

An IoT sub-domain, the growing Web of Things (WoT), allows objects to connect by means of the existing Web infrastructure [4]. Nevertheless, the resulting system has a priori no coordinating abilities, therefore provided services may not be optimal, or can even conflict with each other.

2.2 Intelligent agents

In the previous section, we discussed the benefits in embedding intelligence into objects, giving them adaptive, reasoning, coordinating and self-managing capacities, which can be done by the use of AI techniques. In the AI field, an intelligent agent is a goal-driven entity capable of perceiving and acting upon its environment [5]. An agent must have acting capabilities in order to interact with its environment, and this connection may have many forms.

An agent with direct and tight coupling to its environment performs reactive acting, i.e. a data-driven behavior where actions follow predefined models and plans.

One with a more flexible coupling performs adaptive acting, i.e. contextual changes trigger refinement² of avail-

able actions to find possible alternatives to continue current plans, individually or collectively with other agents [6, 7].

Going further, an agent with a freely and relaxed coupling has a proactive behavior, presenting planning and reasoning capabilities. It is then able to reason upon available context information in order to find new solutions to a given need, allowing it to create new plans to achieve goals, which improves its adaptability.

Numerous solutions for industrial scenarios, such as manufacturing lines [8], use intelligent agents, but these environments are predictable by design, being composed of agents that perform regular, repetitive and controlled processes. In highly dynamic scenarios, like smart homes, reasoners need to be both timely and context-aware in order to answer user needs fast and autonomously.

For an agent to be context-aware, it needs to perceive and understand the world. In order to do so, context collection, modeling and reasoning techniques [9] are needed, specially in the IoT domain where sensors generate enormous amounts of raw data [10].

We are interested in dynamic context modeling techniques because they present open and expandable models that can fit wide ranges of IoT environments [11], and we know smart homes are highly heterogeneous and dynamic ones. In this paper we will consider that context is given ready-to-use by an external function.

Based on the analysis of this section, the agent of interest in a smart home is one capable of adaptive acting and context-aware planning.

2.3 Challenges and requirements

An advanced smart home is expected to perform responsive contextual adaptation in a user-centric manner. The use of agents to tackle this challenge allows us to identify some smart home requirements.

As said previously, such environments are highly dynamic, and the system's main objective is to satisfy user needs. Its final architecture must also take into account heterogeneous objects, while allowing flexible interaction among autonomous agents, creating a social IoT ecosystem [12], which we can call a *multi-agent based smart home*.

To identify smart home requirements the following scenario is presented: a user watching a movie receives a call.

In this case, the smart home has to be *context-aware* to perceive and understand the situation, and *adaptive* to change its plan at run-time, allowing the user to answer the call. It may, for example, mute the TV and activate subtitles. In this case the smart home needed no user input, being *autonomous* to pursue goals to satisfy user needs.

In the beginning of the scenario, when the user asked to watch a movie, it is interesting that the smart home is capable to perform *goal reasoning* and infer secondary related goals, such as turning of lights and closing windows.

Another aspect we tackle is the fact that a smart home is constantly interacting with humans, therefore it must be *responsible*, i.e. in accordance with user values. In our

¹Such as <https://nodered.org>, and <https://ifttt.com>

²Process of finding alternative actions applicable in the current context

scenario, if the user privileges their leisure time more than their work, the smart home may decide to ignore the call if it was from the user's boss. But if family is of utmost importance, the system should never ignore a call from user's mom. This property goes beyond defining rules and constraints for system behavior [13], it dynamically reasons on the impact of user values on actions in current context.

The last property is *coordination*. The TV and the phone presented in the scenario must interact in order to perform any of the described adaptations.

It is important that all above properties are preserved during run-time, because smart home and ubiquitous systems in general must be constantly online and responsive.

3 Planning and Acting Agents

The previous section presented a global overview of the main motivation of this article, with special attention to smart homes. In these environments, we expect IoT objects to be intelligent agents, capable of acting upon and perceiving a dynamic world. This section will investigate the deliberation functions necessary to plan and act, and some state-of-the-art implementations. It is important to note that the word "deliberation" in this paper has the same meaning as in [14], which can sometimes differ from the common meaning we see in multi-agent systems literature.

The *adaptability*, *context-awareness* and *autonomy* of a smart home depends on the sophistication of acting and planning functions. The acting function receives the context from the external world (see Figure 1). The planning function above gives plans to be used by the acting one.

3.1 Acting

The "Acting" shown in Figure 1 is a deliberation function, the one that has direct contact to the external world, responsible for following plans (a set of ordered actions built to achieve goals), and react to context (with every input it checks if plans are still feasible and chooses appropriate available actions to be executed), all driven by previously chosen goals.

Among existing agent architectures, a popular one is the BDI (belief-desire-intention) agent model [15]. BDI agents have 3 components: *beliefs*, i.e. information about the current context; *desires*, i.e. goals the agent would like to achieve; *intentions*, i.e. agent's chosen goals to be achieved by executing a plan. BDI agents have a basic deliberation function that may be mapped to the "Acting" function described. The procedural plans building the plan library of BDI agent may be mapped to the "Know-How" appearing in Figure 1 in the sense that they define methods with their preconditions that can be used by the "Acting" function to execute and refine its execution based in context [14].

The BDI architecture inspired several agent-oriented programming languages, such as Jason [15], an implementation of AgentSpeak [16], technologies actively studied and expanded [17]. The BDI model has also been extended in many ways. For example, the suspension and resuming

of goals, not present in the original model, is presented in [18]. Also, a common goal notation called CAN (Conceptual Agent Notation) is proposed in [19].

Planning capabilities are not part of the BDI model. In Figure 1, the "Planning" step is executed before the agent becomes active – it means the agent will have access to a pre-compiled library of plans, but will not be able to create new plans during execution.

This behavior is similar to the Refinement Acting Engine (RAE) [14]. RAE-enabled agents present adaptive acting capabilities – they do not perform search algorithms to build new plans, but simply execute available plans applicable in the current context to achieve desired goals.

There are similar approaches in the AI field to create adaptive systems, such as [20], and others that use machine learning techniques, like neural networks, etc. We will focus on intelligent agents because they have an explainable process, whose reasoning is often easier to follow and comprehend than other AI techniques.

3.2 Planning

"Planning" is a deliberation function which complements "Acting" by providing it a set of plans customized to achieve previously defined user goals.

Several planning techniques exist, such as first-principles planning (FPP, informally, the creation of new plans based upon lower-level actions to achieve a goal) [21, 22], and markov decision process (MDP) [23]. Hierarchical Task Network (HTN) is a well-known planning technique in which abstract actions are refined into lower-level actions, until tangible actions that can be sent to execution are obtained [24]. In Figure 1, the HTN planner would be the "Planning" function. Its inputs are a set of user goals and an operational model (the "Know-What"), composed of more abstract actions than the commands in the "Know-How" set, which will be combined by the planner to create plans capable of achieving goals given a context. Its outputs are new plans. In this process, goals are the guiding elements of the planning process.

The difference between "Know-What" (*descriptive model* in [14]) and "Know-How" (*operational model* in [14]) in Figure 1 is: the former is a set of abstract actions used by the "Planning" function to create plans to achieve goals; the latter is a set of low-level commands used by the "Acting" to execute the abstract actions present in plans.

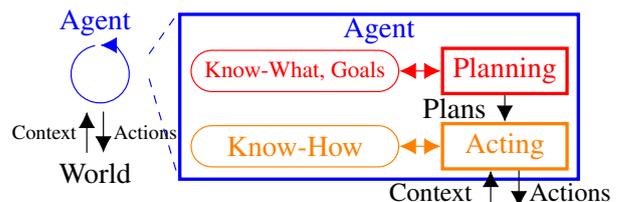


Figure 1: A planning and acting agent

4 Online Adaptation

Using planning and acting functions, described in the previous section, one can implement a basic agent, capable of planning by using known beliefs to create plans to react to predicted changes to achieve goals. But its "Planning" is not context-aware, therefore it provides plans compiled off-line, and does not adapt at execution since it has no feedback from "Acting" as shown in Figure 1.

To overcome the smart home challenge of being capable of adapting to context to satisfy user evolving needs, we thus seek to connect planning and acting capabilities. The expected synergistic result is an agent capable of performing online contextual and user-centric adaptation.

In our scenario from Section 2.3, if the smart home system never encountered the situation before (a call while playing a movie), it may need to perform planning in order to adapt to it. It may, for example, create a plan to pause the film and resume it when the call ends. By doing so, it performed online adaptation, a feature that highly improves system *adaptability*, *context-awareness* and *autonomy*.

4.1 Combining planning and acting

The original BDI model can be extended to perform lookahead planning. The integration of planning algorithms and agent reasoning was surveyed in [25]. An example is the embedding of FPP in BDI agents [21, 22], giving it some planning capabilities. Another is the use of HTN libraries as guides for BDI reasoning [26] aiming at combining HTN domain and BDI, allowing the latter to have access to the information contained in the former. A different approach fuses HTN planner with mental attitudes (beliefs, desires, and intentions) of the BDI architecture [27].

In a smart home, agents should do their best to autonomously adapt at run-time to user needs and environmental changes. Therefore we have interest in merging planning and acting capabilities, and even further, in a solution that perform both in a synergistic and online manner.

4.2 HTN planning and BDI acting

BDI agents and HTN planners are both well-known solutions. The former is capable of adaptive acting by interleaving refinement and execution of actions, and the latter performs planning in order to create new plans to achieve goals. The previous section showed that BDI agents can use the outputs (plans) from HTN planners, but Figure 1 clearly shows there is no feedback from the "Acting" function to the "Planning" one. It will prevent "Planning" from being context-aware, and the resulting agent will not be able to adapt to non-predicted situations.

If we establish that connection, as we can see in Figure 2, we create a feedback loop that allows the agent to perform online adaptation, i.e. online planning and acting. Now "Planning" has access to the current context, and can provide new plans to the "Acting" whenever it encounters a new situation, performing planning at run-time.

The Refinement Engine for Acting and Planning (REAP)

[14] and the BDI agent programming language named CANPlan [28, 29] are solutions that aim to incorporate HTN-style planning into BDI-like agents, ultimately allowing agents to perform online planning and acting.

Another solution is HTN Acting [30, 31], whose approach also combines HTN planning with BDI behavior, i.e. performing interleaved deliberation, acting and failure recovery. By adapting HTN planning semantics, HTN Acting does the opposite of REAP and CANPlan, that is adapting BDI agents.



Figure 2: An online planning and acting agent

5 Responsible Goal-driven Adaptation

Agents in a smart home are subject to environmental changes, and the agent architecture presented in the previous section is capable of online adaptation to those changes. In this section, we identify that agent adaptation must also take into account user goals and values, two concepts that deserve specialized treatment during agent deliberation by the performing of goal- and value-reasoning.

These two reasoners are able to give a smart home the *goal-driven* and *responsible* properties we are looking for.

5.1 Beyond planning and acting

In the previously presented scenario in Section 2.3, we saw the system chose to achieve secondary related goals. The agent capabilities were beyond online planning and acting, its behavior were augmented to perform goal-reasoning. It gave the system a higher level of autonomy, improving user experience since it needed minimum to no user input.

The smart home from our scenario also presents value-reasoning capabilities, which increased system autonomy by helping the choice to answer or ignore the call from user's boss or mom. This function directly influences the goal-reasoning, the planning and the acting processes.

In the following sections we better elaborate the goal- and value-reasoning.

5.2 Goal-reasoning automation

In the literature we find two main types of goals [32]: *goals-to-be* (a.k.a. declarative goals) – i.e. states the agent wants to achieve or maintain [33]; and *goals-to-do* (a.k.a. procedural goals) – i.e. sets of actions the agent wants to successfully execute.

Goals are the guiding elements of the planning process. Most systems expect users to provide the goals to be pursued. In this case, they have automated planning and act-

ing, but they do not deliberate about which goals should be pursued, i.e. they do not reason upon goals, nor create new goals, nor choose which ones to pursue, nor drop existing ones (except in cases of failure or inability to achieve).

In our example scenario, the creation and activation of "turn off lights" and "close window" goals can be performed by a goal-reasoner as shown in Figure 3, which using its current beliefs (e.g. lights are on, windows are open, user started watching a movie) would reason and decide to pursue those goals to maximize user comfort, bettering user experience.

The architecture presented in [34, 35] is an example of a basic goal reasoning, which autonomously creates goals to redirect agent behavior if a user value is at stake. The work from [36] gives operational semantics in the CAN [19] language that enables goal reasoning features necessary to manage goals and plans, such as abortion, suspension and resuming. Another approach presents a goal deliberation strategy called Easy Deliberation [37], which allows activation and deactivation of goals in BDI agents, as well as defining relationships among goals to enable conflict free pursuing.

5.3 Value-driven responsibility

Smart home systems are designed for human users, and therefore they should be responsible systems by taking into account user values, which may include for example the user safety and privacy. The value-reasoning function shown in Figure 3 reasons based on current available beliefs, choosing to prioritize one value over another. In the scenario where a user receives a call while watching a movie, the system could prioritize work-related value over leisure-related value if, for example, the user calendar showed an important work meeting the next day, therefore if the boss was calling their call would not be ignored.

Figure 3 illustrates that the system reasoning process happens in the following order: first, it ranks the values based on the current context; second, it chooses goals that respect those values; third, it chooses or creates plans to achieve the chosen goals that respect the values; finally, it performs acting to realize those plans.

Value-based reasoning research have been highly influenced by The Theory of Basic Human Values [38], that identifies ten motivationally distinct values. For example, the article [39] performs value-reasoning by creating a constraint satisfaction problem in which all applicable plans are ranked using human values, and the plan that better suits user values is chosen. Note that the order of reasoning is the opposite of the one presented in Figure 3: here the agent first chose goals and create plans, then later ranks the available plans based in user values.

The architecture in Figure 3 has similarities to the emotional BDI agent architecture proposed in [34, 35], which includes not only values but also emotional appraisal into the agent reasoning, but whose architecture works differently: its goal reasoning is not affected by the values, in-

stead once the system interprets that the current context endangers user values, it will create a goal that tells the system to satisfy the endangered values.

Another approach use cultural values and rules to not only influence agent individual behavior, but to help coordinating it with other agents that, being in the same environment, have also similar cultural values and rules [40].

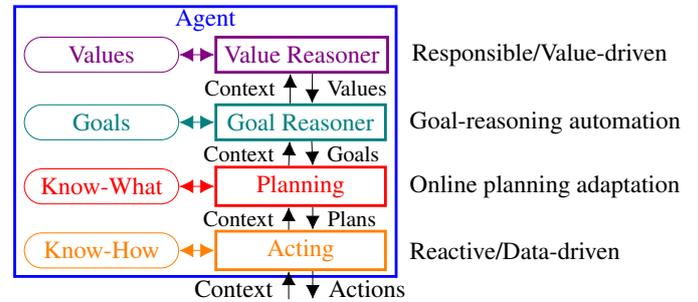


Figure 3: A responsible goal-driven online planning and acting agent

6 Multi-agent Adaptation

An agent with the architecture presented in the previous section is capable of adapting itself in an online, goal-driven and value-based fashion. But its behavior is completely ignorant of the presence of other intelligent agents, ultimately simply sensing other agents actions in the world and considering them simple changes to be adapted to. Therefore we are interested in allowing interaction and coordination among decentralised agents, as shown in Figure 4. We discuss in this existing solutions, their features and limitations for this coordination among agents.

6.1 Multi-agent system

Multi-agent systems (MAS) are compositions of multiple interacting intelligent agents. A smart home can be modeled as one. As proposed in [41], such systems are characterized by the presence of four dimensions: Agents, Environment, Interaction and Organisation.

The Agents dimension is composed of the intelligent agents themselves, i.e. entities capable of reasoning, perceiving and acting on their environment.

The Environment dimension is composed by artifacts the agents can interact with, i.e. real objects, digital services, etc. that can be used by agents in order to change the world or to have access to needed information.

The Interaction dimension constitutes the direct link among agents, allowing them to share messages.

The Organisation dimension establishes system norms and agents roles. The agents may or may not obey to system norms, but in a cooperative setting they would normally follow norms. Also, agents may have specific roles in the system, which gives them some responsibilities and consequently affects agents behaviors.

The exploration of these dimensions may help design a cooperative smart home comprising decentralised agents.

6.2 Coordination approaches

The architecture presented in [8] shows a possible approach using the organisation dimension as a mean of coordinating agents: the organisation has social schemes (sets of collective goals), which are used to allocate available agents that are capable of executing such goals. The allocation is done by giving roles to agents. In their system, the environment is a first-class abstraction [42], i.e. a component designed and programmed with clear responsibilities, such as providing agents with mechanisms for interaction and coordination and with access to available real objects.

Another approach, the Decentralised Online Multi-Agent Planning (DOMAP) presented in [43, 44], uses the environment dimension to coordinate agents. They have three special artifacts: "task board", "contract net board" and "social laws". The goals are declared in the "task board" artifact, and for each goal a "contract net board" is created, in which agents bid their "cost" for executing the goal. In the end, the agent with the best cost is chosen. The "social laws" are special artifacts that impose rules in order to resolve conflicts at run-time, and that all agents obey to.

Both presented approaches use a multi-agent programming framework called JaCaMo [45]. This framework allows the deployment of multi-agent systems by combining three platforms: Jason for the agents dimension, CArtaGo for the environment dimension, and Moise for the organisation dimension.

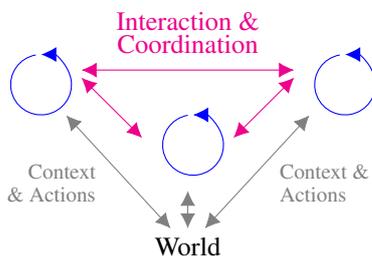


Figure 4: Decentralized agents coordinating, and perceiving and acting upon the world

7 Conclusions and Perspectives

The presented state-of-the-art insights help delimiting the current level of sophistication of each smart home identified requirements.

We saw that there are theoretical and concrete solutions that allow agents to perform online adaptation, i.e. online planning and acting, to a greater or lesser extent. For instance, the cited Refinement Engine for Acting and Planning (REAP), an augmented version of the Refinement Acting Engine (RAE), is a lead that is worth exploring to build a *context-aware*, *adaptive* and *autonomous* agent that plans and acts at run-time.

The BDI model has been shown to be highly expandable, e.g. incorporating planning capabilities. In future work we aim to develop the presented architecture in Figure 3, expanding even further the BDI model with *goal-driven* and

responsible behavior, allowing agents to interpret and take into account user goals and values.

However, as previously said, smart homes are composed of heterogeneous objects. Therefore, the next step following the design of an agent capable of responsible goal-driven online adaptation is to tackle the *coordination* challenge of decentralised agents, in a variety of IoT environments, including the smart home.

Such a multi-agent system may inspire future smart homes to be better designed to human users, by being not only intelligent, adaptive and coordinated systems, but also autonomous and trustworthy. The home of the future truly smart is hopefully just a few years away, and with help from the advancements in research, specially the emergence of AI and IoT solutions, the next human generation may start seeing their home as an integral part of their very own family.

References

- [1] M. Kovatsch, Y. N. Hassan, and S. Mayer, "Practical semantics for the internet of things: Physical states, device mashups, and open questions", in *5th International Conference on the Internet of Things, IOT 2015, Seoul, South Korea, 26-28 October, 2015*, IEEE, 2015, pp. 54–61.
- [2] F. Lécué, "Composition de services web: Une approche basée liens sémantiques. (web service composition : Semantic links based approach)", PhD thesis, École nationale supérieure des mines de Saint-Étienne, France, 2008.
- [3] D. Guinard, V. Trifa, T. Pham, and O. Liechti, "Towards physical mashups in the web of things", in *2009 Sixth International Conference on Networked Sensing Systems (INSS)*, Pittsburgh, PA, USA: IEEE, Jun. 2009, pp. 1–4. (visited on 01/14/2019).
- [4] M. Blackstock and R. Lea, "IoT mashups with the WoTKit", p. 8,
- [5] S. J. Russell and P. Norvig, *Artificial intelligence: a modern approach*, Third edition, Global edition, in collab. with E. Davis and D. Edwards, ser. Prentice Hall series in artificial intelligence. Boston Columbus Indianapolis: Pearson, 2016, 1132 pp.
- [6] A. Ciortea, O. Boissier, A. Zimmermann, and A. M. Florea, "Responsive decentralized composition of service mashups for the internet of things", in *Proceedings of the 6th International Conference on the Internet of Things, IOT 2016, Stuttgart, Germany, November 7-9, 2016*, S. Schneegass, A. Schmidt, F. Michahelles, M. Kritzler, A. Ilic, and K. Kunze, Eds., ACM, 2016, pp. 53–61.
- [7] M. Vallée, "Un intergiciel multi-agent pour la composition flexible d'applications en intelligence ambiante. (a multi-agent middleware for flexible application composition in ambient intelligence)", PhD

thesis, École nationale supérieure des mines de Saint-Étienne, France, 2009.

- [8] A. Ciorcea, S. Mayer, and F. Michahelles, “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 MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, E. André, S. Koenig, M. Dastani, and G. Sukthankar, Eds., International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018, pp. 813–822.
- [9] A. Sorici, G. Picard, O. Boissier, A. Zimmermann, and A. M. Florea, “CONCERT: Applying semantic web technologies to context modeling in ambient intelligence”, *Computers & Electrical Engineering*, vol. 44, pp. 280–306, 2015.
- [10] C. Perera, A. B. Zaslavsky, P. Christen, and D. Georgakopoulos, “Context aware computing for the internet of things: A survey”, *IEEE Communications Surveys and Tutorials*, vol. 16, no. 1, pp. 414–454, 2014.
- [11] J. Euzenat, J. Pierson, and F. Ramparany, “Dynamic context management for pervasive applications”, *Knowledge Eng. Review*, vol. 23, no. 1, pp. 21–49, 2008.
- [12] A.-N. Ciorcea, “Weaving a social web of things : Enabling autonomous and flexible interaction in the internet of things. (tisser le web social des objets : Permettre une interaction autonome et flexible dans l’internet des objets)”, PhD thesis, École nationale supérieure des mines de Saint-Étienne, France, 2016.
- [13] P. Rust, G. Picard, and F. Ramparany, “Using message-passing DCOP algorithms to solve energy-efficient smart environment configuration problems”, in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, S. Kambhampati, Ed., IJCAI/AAAI Press, 2016, pp. 468–474.
- [14] M. Ghallab, D. S. Nau, and P. Traverso, *Automated Planning and Acting*. Cambridge University Press, 2016.
- [15] R. H. Bordini, J. F. Hübner, and M. Wooldridge, *Programming Multi-Agent Systems in AgentSpeak using Jason*. 2007.
- [16] A. S. Rao, “AgentSpeak(I): BDI agents speak out in a logical computable language”, in *Agents Breaking Away, 7th European Workshop on Modelling Autonomous Agents in a Multi-Agent World, Eindhoven, The Netherlands, January 22-25, 1996, Proceedings*, W. V. d. Velde and J. W. Perram, Eds., ser. Lecture Notes in Computer Science, vol. 1038, Springer, 1996, pp. 42–55.
- [17] A. Ricci, R. H. Bordini, J. F. Hübner, and R. Collier, “AgentSpeak(ER): Enhanced encapsulation in agent plans”, p. 16,
- [18] J. Thangarajah, J. Harland, D. N. Morley, and N. Yorke-Smith, “Suspending and resuming tasks in BDI agents”, in *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, 2008, Volume 1*, L. Padgham, D. C. Parkes, J. P. Müller, and S. Parsons, Eds., IFAAMAS, 2008, pp. 405–412.
- [19] M. Winikoff, L. Padgham, J. Harland, and J. Thangarajah, “Declarative & procedural goals in intelligent agent systems”, in *Proceedings of the Eighth International Conference on Principles and Knowledge Representation and Reasoning (KR-02), Toulouse, France, April 22-25, 2002*, D. Fensel, F. Giunchiglia, D. L. McGuinness, and M.-A. Williams, Eds., Morgan Kaufmann, 2002, pp. 470–481.
- [20] A. Diaconescu, P. Mata, and K. Bellman, “Self-integrating organic control systems: From craysh to smart homes”, p. 8,
- [21] M. Xu, K. Bauters, K. McAreevey, and W. Liu, “A formal approach to embedding first-principles planning in BDI agent systems”, in *Scalable Uncertainty Management - 12th International Conference, SUM 2018, Milan, Italy, October 3-5, 2018, Proceedings*, D. Ciucci, G. Pasi, and B. Vantaggi, Eds., ser. Lecture Notes in Computer Science, vol. 11142, Springer, 2018, pp. 333–347.
- [22] L. d. Silva, S. Sardiña, and L. Padgham, “First principles planning in BDI systems”, in *8th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009), Budapest, Hungary, May 10-15, 2009, Volume 2*, C. Sierra, C. Castelfranchi, K. S. Decker, and J. S. Sichman, Eds., IFAAMAS, 2009, pp. 1105–1112.
- [23] R. Bellman, “A markovian decision process”, *Journal of Mathematics and Mechanics*, vol. 6, no. 5, pp. 679–684, 1957.
- [24] A. Herzig, L. Perrussel, and Z. Xiao, “On hierarchical task networks”, in *Logics in Artificial Intelligence - 15th European Conference, JELIA 2016, Larnaca, Cyprus, November 9-11, 2016, Proceedings*, L. Michael and A. C. Kakas, Eds., ser. Lecture Notes in Computer Science, vol. 10021, 2016, pp. 551–557.
- [25] F. Meneguzzi and L. d. Silva, “Planning in BDI agents: A survey of the integration of planning algorithms and agent reasoning”, *Knowledge Eng. Review*, vol. 30, no. 1, pp. 1–44, 2015.

- [26] L. d. Silva, “BDI agent reasoning with guidance from HTN recipes”, in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, K. Larson, M. Winikoff, S. Das, and E. H. Durfee, Eds., ACM, 2017, pp. 759–767.
- [27] G. Wickler, S. Potter, A. Tate, M. Piechouicek, and E. Semsch, “Planning and choosing: Augmenting HTN-based agents with mental attitudes”, in *Proceedings of the 2007 IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Silicon Valley, CA, USA, November 2-5, 2007*, IEEE Computer Society, 2007, pp. 222–228.
- [28] S. Sardiña, L. d. Silva, and L. Padgham, “Hierarchical planning in BDI agent programming languages: A formal approach”, in *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2006), Hakodate, Japan, May 8-12, 2006*, H. Nakashima, M. P. Wellman, G. Weiss, and P. Stone, Eds., ACM, 2006, pp. 1001–1008.
- [29] S. Sardiña and L. Padgham, “A BDI agent programming language with failure handling, declarative goals, and planning”, *Autonomous Agents and Multi-Agent Systems*, vol. 23, no. 1, pp. 18–70, 2011.
- [30] L. d. Silva, “HTN acting: A formalism and an algorithm”, in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden, July 10-15, 2018*, E. André, S. Koenig, M. Dastani, and G. Sukthankar, Eds., International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM, 2018, pp. 363–371.
- [31] —, “Addendum to “HTN acting: A formalism and an algorithm””, *CoRR*, vol. abs/1806.02127, 2018.
- [32] M. B. v. Riemsdijk, M. Dastani, and M. Winikoff, “Goals in agent systems: A unifying framework”, in *7th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008), Estoril, Portugal, May 12-16, 2008, Volume 2*, L. Padgham, D. C. Parkes, J. P. Müller, and S. Parsons, Eds., IFAAMAS, 2008, pp. 713–720.
- [33] S. Duff, J. Thangarajah, and J. Harland, “Maintenance goals in intelligent agents”, *Computational Intelligence*, vol. 30, no. 1, pp. 71–114, 2014.
- [34] N. Cointe, “Jugement éthique pour la décision et la coopération dans les systèmes multi-agents”, 2017.
- [35] C. Battaglino, R. Damiano, and L. Lesmo, “Emotional range in value-sensitive deliberation”, in *International conference on Autonomous Agents and Multi-Agent Systems, AAMAS '13, Saint Paul, MN, USA, May 6-10, 2013*, M. L. Gini, O. Shehory, T. Ito, and C. M. Jonker, Eds., IFAAMAS, 2013, pp. 769–776.
- [36] J. Harland, D. N. Morley, J. Thangarajah, and N. Yorke-Smith, “Aborting, suspending, and resuming goals and plans in BDI agents”, *Autonomous Agents and Multi-Agent Systems*, vol. 31, no. 2, pp. 288–331, 2017.
- [37] A. Pokahr, L. Braubach, and W. Lamersdorf, “A goal deliberation strategy for BDI agent systems”, in *Multiagent System Technologies*, T. Eymann, F. Klügl, W. Lamersdorf, M. Klusch, and M. N. Huhns, Eds., red. by D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, and G. Weikum, vol. 3550, Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 82–93. (visited on 04/01/2019).
- [38] S. H. Schwartz, “Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries”, in *Advances in Experimental Social Psychology*, vol. 25, Elsevier, 1992, pp. 1–65. (visited on 03/18/2019).
- [39] S. Cranefield, M. Winikoff, V. Dignum, and F. Dignum, “No pizza for you: Value-based plan selection in BDI agents”, in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, Ed., ijcai.org, 2017, pp. 178–184.
- [40] L. Vanhée, *Using Cultures and Values to Support Flexible Coordination*. Utrecht University, 2015.
- [41] O. Boissier, R. H. Bordini, J. F. Hübner, and A. Ricci, “Dimensions in programming multi-agent systems”, *Knowledge Eng. Review*, vol. 34, e2, 2019.
- [42] D. Weyns, A. Omicini, and J. Odell, “Environment as a first class abstraction in multiagent systems”, *Autonomous Agents and Multi-Agent Systems*, vol. 14, no. 1, pp. 5–30, 2007.
- [43] R. C. Cardoso and R. H. Bordini, “A modular framework for decentralised multi-agent planning (extended abstract)”, p. 3, 2017.
- [44] R. C. Cardoso, “A decentralised online multi-agent planning framework for multi-agent systems”, 2018.
- [45] O. Boissier, R. H. Bordini, J. F. Hübner, A. Ricci, and A. Santi, “Multi-agent oriented programming with JaCaMo”, *Sci. Comput. Program.*, vol. 78, no. 6, pp. 747–761, 2013.