

Towards Hypermedia Environments for Adaptive Coordination in Industrial Automation

Ganesh Ramanathan
Siemens AG
Zug, Switzerland
ganesh.ramanathan@siemens.com

Simon Mayer
University of St. Gallen
Switzerland
simon.mayer@unisg.ch

Andrei Ciortea
University of St. Gallen
Switzerland
Wimmics, Inria, Université Côte d'Azur,
CNRS, I3S, Sophia Antipolis, France
andrei.ciortea@unisg.ch

Abstract—Electromechanical systems manage physical processes through a network of inter-connected components. Today, programming the interactions required for coordinating these components is largely a manual process. This process is time-consuming and requires manual adaptation when system features change. To overcome this issue, we use autonomous software agents that process semantic descriptions of the system to determine coordination requirements and constraints; on this basis, they then interact with one another to control the system in a decentralized and coordinated manner. Our core insight is that coordination requirements between individual components are, ultimately, largely due to underlying physical interdependencies between the components, which can be (and, in many cases, already are) semantically modeled in automation projects. Agents then use hypermedia to discover, at run time, the plans and protocols required for enacting the coordination. A key novelty of our approach is the use of hypermedia-driven interaction: it reduces coupling in the system and enables its run-time adaptation as features change.

Index Terms—hypermedia environments, multi-agent systems

I. INTRODUCTION

Systems in installations like chemical processing plants, factories, and buildings are composed of electro-mechanical components that progressively modify substances and manipulate forces and energy forms towards desired states. Ultimately, each of the components involved in such processes manages a set of physical mechanisms controlled by a software program. We use the term *agent* to refer to such programs that are responsible for controlling the functioning of one or more components. However, the functioning of the components is often also dependent on their peers, which are located up- and downstream in the process. Coordination between the agents is hence required to orchestrate the physical mechanisms and ensure safety conditions at the startup, operation, and shutdown of the process.

In deployed systems, coordination is typically implemented in one of two broad architectural schemes: (1) through the use of hierarchical supervisory control, or (2) by distributing the responsibility amongst the participants (e.g., in complex large-scale deployments). For both cases, the *coordination plans* (i.e., when and how the agents interact) are synthesised by a combination of top-down design decomposition and bottom-up integration of component features. In current systems, these plans are *hard-coded in the agents* and any change

in the system features requires re-engineering [3], [15]. This hard coding affects especially distributed process control in scenarios that face reconfiguration during the system's lifetime (e.g., flexible manufacturing or changing usage patterns in building control systems) [3], [6], [15].

Though several researchers have advocated the use of Multi-agent Systems (MAS) for industrial automation that have to deal with dynamic environments [6], [7], [17], [21], the question of how agents would determine their *coordination responsibilities* so that they discover and interact with collaborators at run time is still largely under-explored. Our past research has demonstrated that current engineering ontologies based on Semantic Web technologies can be integrated to assist agents in choosing control strategies for local functions [12]–[14]. In this work, we show that agents can also use such ontologies to infer coordination responsibilities.

To enable coordination to happen in a loosely coupled and interoperable manner, we motivate the adoption of hypermedia-driven interaction where agents utilize a Web-based environment to publish their profiles, navigate using semantic links, and discover technical operations (on the electro-mechanical equipment) to set up and enact the coordination dynamically. We emphasise that *technical operations* can indeed be both inferred from and linked to the *physical actions* (i.e., the causation of flow of energy, matter, or effort [2]).

II. RELATED WORK

Engineering of coordination is based on the knowledge of system requirements, technical design, models of physical processes, and the control strategies [10], [20], [22]. One interoperable and extensible way to represent these descriptions is through ontologies such as BRICK [1] and OntoCAPE [8] for topologies, PhySys [2] and Elementary [13] for physical processes, CTRLOnt [11] for control programs, and GORE [9] for requirements. Our work [14] has shown that these ontologies can be *integrated* using high-level concepts and that, in this way, control agents can know their context and strategy for operation (cf. [12]). This knowledge can be extended further to describe the physical effects of actions on a component (cf. [13]). [16] has proposed an ontology to model finite state machines which can also be used for modeling coordination

program of a supervisory controller [11]. In a distributed control system where coordination goals are known *a priori*, agent-to-agent *interaction* can be achieved by adopting one of well-known interaction protocols such as those proposed by FIPA¹. But a method to allow *run time* selection and enactment of a protocol based on dynamically changing needs is missing, and so is a way to embed *physical process-related semantics* into the coordination messages. For systems that need frequent reorganization, [4], [6] have shown the benefits of placing industrial software agents in a hypermedia environment, introducing hypermedia MAS as a new class of Web-based MAS; and [18], [19] have shown that agents can use standardized W3C Web of Things *Thing Descriptions*² (WoT TDs) to reveal their affordances. We seek to combine the benefits of these individual approaches to achieve adaptive coordination in distributed process control.

III. ACHIEVING ADAPTIVE COORDINATION

In our research, we show that by being part of a hypermedia environment, agents in industrial automation are enabled to discover their collaboration partners and autonomously align themselves to fulfill coordination responsibilities. To this end, in Section III-A we show how existing machine-understandable knowledge about the system can be used to infer physical process dependencies between the components and how this can be used to establish the *need for coordination* with their dependent counterparts (in Section III-B). On this basis, the enactment of coordination requires an agent to interact with others (e.g., to query and update states or invoke actions) according to a plan. However, today, the interfaces between programs are coupled using communication protocols and *design-time* knowledge about information resources—again, if these interfaces change, manual re-engineering would be required. To overcome this, in Section III-C we propose the use of semantically described *coordination protocols* together with the Web of Things paradigm to achieve loose coupling of agents and interaction opportunities, and thereby run-time adaptation of the agents.

Our work therefore connects existing research in control and process engineering, the Semantic Web, and hypermedia MAS to show that adaptive coordination of process control in large dynamic systems can be achieved.

Throughout this section, we use the example of a chilled water production system shown in Figure 1 as a running example³. In many places, the listings in this paper have been simplified to convey the gist of the approach more easily.

A. Inferring Coordination Responsibilities

Figure 2 shows the inter-connected high-level concepts in the current ontologies for representing the different aspects of system knowledge [12]. We introduce the concept of *coordination responsibility* (CR) to capture the knowledge about *why* an agent needs to coordinate with others and show




Fig. 1. A chilled water refrigeration plant where agents have to coordinate flow control and heat-extraction processes

that this can be obtained through semantic queries on the integrated system description. Specifically, given the process goals of the system and the set of components that an agent manages, CR is inferred in terms of *process dependencies* between the components within an agent's scope and those outside it.




Fig. 2. We use the integrated system description to infer coordination needs and its enactment (highlighted box).

The inferred CR also allows the agent to list its own *affordances* that may be relevant for its collaborating partners. In section III-C we describe how WoT TD enables machine-understandable description of such affordances. The agent combines these to create an *agent profile* (AP).

In the example of the chilled system, the CR of the agent managing the pumps can be inferred using a semantic query (listed using SPARQL and the result in RDF):

```
1 SELECT ?comp ?rel ?dep WHERE{
2   :agent-p elem:manages ?subsystem.
3   ?subsystem elem:hasComponent* ?comp.
4   ?comp ?rel ?dep.
5   ?rel a elem:processRelation.
```

Part of the results to the query, as shown below, expresses that pump-1, which is managed by agent-p, is physi-

¹<http://www.fipa.org/>

²<https://www.w3.org/TR/wot-thing-description11/>

³Can be accessed here <https://autonomous-buildings.github.io/hmas>

cally dependent on the valves downstream and influences the chillers upstream:

```
1 :pump-1 a elem:Component, brick:Pump.
2 elem:affectedVariable :flow-rate;
3 elem:influences :chlr-1,chlr-2,chlr-3;
4 elem:influencedBy :vlv-1,vlv-2,vlv-3;
```

Therefore, agent-p has to coordinate with the agents responsible for the chillers and valves—but it does not know these agents upfront and has to discover them. Consequently, agent-p, in its profile, expresses that other agents can ask it to achieve a particular flow rate or query its current state:

```
1 :profile-p a elem:AgentProfile;
2 elem:hasCR :cr-p-process-dependencies;
3 td:hasActionAffordance :change-flowrate;
4 td:hasPropertyAffordance :current-flowrate;
5 ...
```

In the next step, we show how an agent uses its profile to navigate the hypermedia environment to discover other agent(s) with which it needs to establish coordination interaction.

B. Discovering Coordination Partners

Traditionally, the deployment scheme of agents in the automation system would be known upfront from the design. But we are aiming to cater for a system that can be dynamically reconfigured. For this purpose, we argue that hypermedia environments are valuable in industrial automation due to their ability to integrate diverse types of data and media formats, such as knowledge graphs, TDs, functional mock-up units for simulation, into a cohesive framework [5]. This integration facilitates more efficient information retrieval because agents are redirected to resources in a contextualized manner. For example, an agent may submit its profile to a registry service, which could extract the TD and publish it to a repository or query the knowledge graph to determine recipients likely to be interested in the profile. If the profile is not of interest to a recipient, but the recipient knows a further context, it may simply redirect the profile there. Therefore, agents can navigate complex systems more intuitively, accessing relevant information quickly and effectively. Furthermore, hypermedia environments support dynamic updates and notification of changes, which are crucial when the system's state or even its design changes at run time.

C. Enacting the Coordination

In Section III-A, we have seen that an agent seeks counterparts to collaborate by advertising its profile, a part of which describes the interaction interfaces it offers to the outside world. Complementary to this, potentially collaborating agents also expose their profiles.

We propose to use WoT TD to describe an agent's contextualized affordances, including ways to observe its state or act on it (or request it to act). In [14], we have shown that TDs

can be embedded with machine-understandable knowledge about physical mechanisms. Based on this, the agent profiles can be semantically matched to determine if the process goals envisaged in the CR are achievable through the use of affordances exposed by the agents. The following example illustrates a TD exposed by agent-cl in response to the profile of the agent-p:

Listing 1. An example TD of an agent that manages a chiller

```
{ "@id": "urn:agent-cl", "@type": "hmas:Agent",
  "links": ["urn:chiller-2", "urn:chiller-3"],

  "properties": { "present-output": {
    "@type": "hvac:CoolingPower",
    "observes": ":variable_energy-output",
    "readOnly": true, "type": "number",
    "forms": [{"href": "/states/output"}]
  }},
  "actions": { "set-flowrate": {
    "@type": "hvac:FlowRate",
    "manipulates": ":variable_water-flow-rate",
    "forms": [{"href": "/actions/set-flow"}] }
  }
}
```

If an agent has nothing to offer towards collaboration, it may result in an empty TD (i.e. without any affordances). In such cases, the caller is often redirected to another agent through the use of hypermedia links embedded in the TD, highlighting their crucial role in the process. The following example query finds affordances that affect the desired physical state of the process (and therefore relevant to the CR):

```
SELECT ?affordance WHERE{
  :agent-p elem:hasProfile ?ap_p.
  :agent-cl elem:hasProfile ?ap_cl.
  ?ap_p elem:desiredState ?process_state.
  ?ap_cl elem:hasTD ?td_cl.
  ?td_cl td:hasActionAffordance ?affordance.
  ?affordance elem:manipulates ?process_state.
}
```

Once the agents infer that affordances required for enacting the CR exist in each other's TDs, they then need to set up the coordination enactment, which comprises two parts: (1) the plan or sequence of actions to take and (2) the protocol to use for communicating the requests (for action or status update).

We see two possible ways in which this can happen: (1) the agents search for a plan and a corresponding interaction protocol created by human experts, which are made available in a library as Web resources, or (2) they use the machine-understandable knowledge of physical mechanisms and simulation models to *creatively synthesize* a coordination plan (and further, an interaction protocol based on the affordances listed in the TDs).

We are exploring the former approach by introducing a novel way of machine-understandable description of *coordination protocols* which enhances message specifications by embedding the semantics of the physical process, a concept that has not been previously explored. We do this by extending the message descriptions with semantics about the *physical* pre- and post-conditions (both on the part of the caller and the

receiver). For example, the pre- and post-conditions referring to physical states can be obtained from the state machine models of the control programs of the agents. An ontology provides this formalization, which we make available on our project website.

Further, instead of describing low-level protocol details and message schemes directly as a part of the coordination protocol, we include only the *semantic type of the affordance* to search for in the TD, which, on the other hand, is more suited for conveying network and application protocol bindings.

Such coordination protocols hold potential by allowing us to model well-known interaction scenarios in various domains while carrying *in-band knowledge* about the physical requirements and consequences of the interactions. For instance, they can be used to create demand-response protocols in energy management and be made readily available as Web resources in the hypermedia environment. The following code snippet shows how a hypothetical message for flow change can be sent by a chiller whose valves are open.

```
1 :request-flowrate-change a intr:Message.
2   intr:hasSender :_sender.
3   :_sender a hvac:ChillerManager.
4   intr:precondition :_valves_open;
5   intr:desiredPostcondition :_has_input_flow.
6   intr:hasReceiver :_reciever.
7   :receiver a hvac:PumpManager
8   intr:precondition :_is_primed;
9   intr:commitment :_flow_changed.
10  intr:hasAffordance :_td_af_flowchange.
11  :_td_action_flowchange a hvac:FlowModulation.
```

IV. CONCLUSION

We have shown how machine-understandable descriptions of the engineering system, inclusive of the model of physical processes, enable automated inference of the agents' coordination responsibilities. While seeking to tackle the challenge of dynamically enacting coordination without requiring manual engineering, we see strong conceptual synergies in hypermedia-driven interactions. Hypermedia environments consisting of knowledge graphs, resources described in a machine-understandable manner, and support for large-scale deployments provide ideal support for automation agents operating in a distributed and dynamic environment. By seeking to make available the semantics about the technical system and the underlying physical mechanisms to the agents, we augment existing approaches in agent-to-agent interaction for coordination. In particular, by describing interaction protocols in terms of the underlying physical processes they influence, we enable agents to discover and use interaction protocols based on their current goals. Our research efforts in the past decade have laid the foundations by enabling the creation of integrated system descriptions and infusing knowledge of physics into digital representations.

However, our work is far from complete. There is a pressing need for further research to delve into topics such as modeling

coordination plans (e.g., to include temporal logic), norms and constraints, and the crucial aspect of handling failures. By showing how hypermedia-based multi-agent systems could support future methods addressing adaptive automation, we hope to motivate researchers to rethink architectural paradigms which result in tight coupling.

REFERENCES

- [1] Balaji, B., et al.: Brick: Towards a unified metadata schema for buildings. In: Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments (2016)
- [2] Borst, P., Akkermans, J., Pos, A., Top, J.: The physsys ontology for physical systems. In: Working Papers of the Ninth International Workshop on Qualitative Reasoning QR. vol. 95, pp. 11–21 (1995)
- [3] Chokshi, N., McFarlane, D.: A distributed coordination approach to reconfigurable process control. Springer Science & Business Media (2007)
- [4] Ciortea, A., et al.: A decade in hindsight: the missing bridge between multi-agent systems and the world wide web. In: AAMAS 2019-18th International Conference on Autonomous Agents and Multiagent Systems. p. 5 (2019)
- [5] Ciortea, A., Mayer, S., Bienz, S., Gandon, F., Corby, O.: Autonomous search in a social and ubiquitous web. Personal and Ubiquitous Computing pp. 1–14 (2020)
- [6] Ciortea, A., Mayer, S., Michahelles, F.: 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. pp. 813–822 (2018)
- [7] Maturana, F.P., et al.: Distributed multi-agent architecture for automation systems. Expert Systems with Applications **26**(1), 49–56 (2004)
- [8] Morbach, J., Wiesner, A., Marquardt, W.: Ontocape—a (re) usable ontology for computer-aided process engineering. Computers & Chemical Engineering **33**(10), 1546–1556 (2009)
- [9] Negri, P.P., et al.: Towards an ontology of goal-oriented requirements. CIBSE (2017)
- [10] Nielsen, C.B., et al.: Systems of systems engineering: basic concepts, model-based techniques, and research directions. ACM Computing Surveys **48**(2) (2015)
- [11] Preisig, H.A.: Ontology-based process modelling-with examples of physical topologies. Processes **9**(4), 592 (2021)
- [12] Ramanathan, G.: Synthesizing multi-agent system organization from engineering descriptions. In: International Workshop on Engineering Multi-Agent Systems. Springer (2023)
- [13] Ramanathan, G., Mayer, S.: Reasoning about physical processes in buildings through component stereotypes. In: Proceedings of the 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (2023)
- [14] Ramanathan, G., Mayer, S.: A match made in semantics: Physics-infused digital twins for smart building automation. arXiv preprint arXiv:2302.06970 (2024)
- [15] Rasmussen, L.: Decentralized coordination for open distributed systems. SICS Report, October (1997)
- [16] Schneider, G.F., Pauwels, P., Steiger, S.: Ontology-based modeling of control logic in building automation systems. IEEE Transactions on Industrial Informatics **13**(6), 3350–3360 (2017)
- [17] Terán, J., Aguilar, J., Cerrada, M.: Integration in industrial automation based on multi-agent systems using cultural algorithms for optimizing the coordination mechanisms. Computers in Industry **91**, 11–23 (2017)
- [18] Vachtsevanou, D., et al.: Signifiers as a first-class abstraction in hypermedia multi-agent systems. arXiv preprint arXiv:2302.06970 (2023)
- [19] Vachtsevanou, D., et al.: Enabling bdi agents to reason on a dynamic action repertoire in hypermedia environments. In: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems. pp. 1856–1864 (2024)
- [20] Van Ruijven, L.: Ontology for systems engineering. Procedia Computer Science **16** (2013)
- [21] Wang, J., Chen, K., Lewis, F.L.: Coordination of multi-agent systems on interacting physical and communication topologies. Systems & Control Letters **100**, 56–65 (2017)
- [22] Yang, L., Cormican, K., Yu, M.: Ontology-based systems engineering: A state-of-the-art review. Computers in Industry **111** (2019)