Agent-Based Orchestration on a Swarm of Edge Devices

Banani Anuraj
banani.anuraj@hevs.ch
University of Applied Sciences and Arts Western Switzerland HES-SO
Sierre, Switzerland
Supervised by: Jean-Paul Calbimonte

ABSTRACT

The proliferation of smart devices, sensors, autonomous robots, drones, and other similar instruments have profoundly changed the way of implementing and deploying systems in industrial and home environments, for diverse scenarios such as smart agriculture, healthcare, or manufacturing. Devices in these settings are not limited to simply observe and acquire data for monitoring, but they are also equipped with actuation capabilities, as well as the possibility of autonomously processing the incoming data through various techniques. However, given their intrinsic limitations regarding the capacity to store and process computations, it is often necessary to delegate some of these processing tasks to intermediary edge nodes in the network. These nodes, given their unique position can act as orchestrators guiding the decentralized work of the interconnected autonomous devices. Beyond static and pre-defined organization structures, in this work we propose the usage of agent and multi-agent-based models for designing and implementing swarms of edge nodes, conceived to dynamically orchestrate other devices, while meeting quality of service conditions. Allowing the control of intelligent edge nodes as conveyors and orchestrators on swarms of devices, we aim at providing intelligence to the self-organization of edge nodes, which may interchange streaming data, and represent their own capabilities through semantic models. Swarm-inspired behavioral patterns would guide the collaborative distribution of their computational tasks. Finally, we will implement and demonstrate the proposed technologies in an elderly home environment powered with a host of edge computing, sensing, and actuating devices.

CCS CONCEPTS

• **Computer systems organization** → **Embedded and cyber-physical systems**; **Sensors and actuators**; **Real-time system architecture**; • **Information systems** → **Semantic web description languages**; • **Software and its engineering** → **Interoperability**.

KEYWORDS

edge computing, sensor networks, actuators, streaming data, agents, swarms, distributed architecture

DEBS '23, June 27–30, 2023, Neuchatel, Switzerland

ACM ISBN 979-8-4007-0122-1/23/06.
https://doi.org/10.1145/3583678.3603285

ACM Reference Format:
Banani Anuraj. 2023. Agent-Based Orchestration on a Swarm of Edge Devices. In *The 17th ACM International Conference on Distributed and Event-based Systems (DEBS '23), June 27–30, 2023, Neuchatel, Switzerland*. ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3583678.3603285

1 PROBLEM STATEMENT

A key challenge in smart applications including heterogeneous sensors and actuators is the timely and efficient processing of incoming data streams, with the goal of producing actionable outputs consumable and interpretable by both end users and other applications. Although cloud-based solutions offer powerful computational and storage resources to answer to smart applications' demands, latency incurred by these services remains prohibitive, while the available resources of nodes closer to the domain of observation are often underutilized.

Nowadays edge devices are equipped not only with sensing and actuating functionalities, but also with a wide range of computational, networking, and storage capacities. Deployments of these smart devices and edge nodes have now the potential of working in a decentralized manner, dynamically distributing the work among participating nodes. This type of deployment can bring low latency and efficient processing in diverse use-cases, including autonomous vehicles, distributed robotic environments in manufacturing, or smart healthcare.

Building intelligent applications on top of these IoT devices and platforms requires detailed knowledge and expertise for understanding, supporting and configuring their protocols, communication standards, architectures, data schemas, etc. In this context, it becomes essential to be able to model not only the capabilities and characteristics of devices, but also the way they interact with each other when they need to achieve a certain goal.

In this PhD project, we propose to address the problem of orchestrating sensing and actuating edge devices, which participate in a decentralized manner towards a common objective. This orchestration should allow devices to find efficient ways of self-organizing, leveraging the capabilities of the participating nodes. As it will be discussed below, the proposed work will explore the usage of concepts from swarm computing and agent-based systems in combination with semantic technologies for machine interpretable description of devices' capabilities, goal, and behaviors.

The results of this project, given the wider and wider availability of smart devices, has a potential applicability in multiple use cases including: self driving autonomous vehicles, smart healthcare, industrial automation, etc. In particular, this PhD project will use a healthcare tele-rehabilitation scenario as a means to implement and validate the proposed work, so to showcase the practical application of the proposed technologies in a concrete use-case.

2 RESEARCH APPROACH/METHODOLOGY

The research approach to be adopted is based on principles of action research, adapted to computer science research work [8]. In practice it will follow a pattern of iterations, each consisting of steps, namely: planning, acting, observing, reflecting. Given the exploratory nature of the research work to be studied in this PhD project, we will need to first analyze existing works in the areas of distributed edge computing, swarm self-organizing systems, and decentralized (semantic) knowledge management. A systematic analysis of these topic and especially their intersections points will be essential to have a comprehensive view on the topic, as well as to identify the metrics that can be used later to evaluate the success of the proposed approach with respect to existing works.

After planning and executing this survey work, and reflecting on their results, the next step will consist in the planning of the semantization of edge nodes behaviors, goals and plans, following a hybrid model that includes the use of semantic technologies for representation and reasoning, over agent-based models like Beliefs-Desires-Intentions (BDI). After the formalization of these approaches and implementation, again a validation of the proposed models will be performed and subsequently we will reflect on those results before taking them as a basis for the swarm computing part.

In this stage, we will prepare and plan for the design and implementation of the self organization aspects of the project using swarm-inspired concepts. For this we will use a combination of multi-agent systems and their high-level representation of goals, as the basis for configuring behaviors expected by the swarm as a whole. This approach will be first implemented using a simulated environment, and later we will provide a full fledged implementation and validation with a real prototype.

For the prototyping part of the PhD project, we will take a smart healthcare use case from the tele-rehabilitation domain. The setting consists on a multi-patient environment where elderly people at home (or nursing home) perform personalized exercise targeting post-surgical physiotherapy treatments. They will wear motion sensing devices to capture multi-joint exercise movements, and will also be equipped with a tablet for instantaneous feedback and monitoring. In addition, an anthropomorphic robot will be available to enhance user engagement and quality of interactions. From the methodological point of view, we will perform lab-only experiments to observe metrics related to the technical implementation, including the swarm formation and node management, scheduling of the agent behaviors, as well as data flow metrics.

3 RESEARCH SETTING

Within this context, the PhD thesis will investigate the following research questions:

- Can we use semantic models as a means to share edge computing devices' capabilities for dynamic formation of an edge swarm?
- Can we use agent-based execution models for effectively representing the goals, behaviors and knowledge of swarms of edge devices?
- Can we use self-organizing swarm computing principles for orchestrating the tasks and interactions of decentralized edge devices, towards a common goal?

These research questions try to address the problem of orchestrating edge devices, as mentioned earlier, taking into account the potential resources that these nodes have and that can potentially handle certain computational tasks, and implement at least some degree of intelligent behavior. Beyond the existing pre-planned configuration of IoT devices, we envision a setting where different nodes may even negotiate their participation towards a given goal, exposing their capabilities and dynamically participating in networks of swarm edge devices.

In this context, we will investigate the following hypotheses: (i) ontologies and other semantic models emerging from the IoT and Web of Things communities can be extended in order to serve as the basis for discoverable edge nodes in decentralized swarms; (ii) agent models such as BDI can be enacted in edge devices to effectively reflect the common goals of a swarm, as well as to formally model, and implement their behaviors in pursuance of the aforementioned goals; and (iii) modeling edge devices as swarm computing entities can be used effectively as an execution model for dynamic formation of device "communities" that cooperatively work towards a common objective. Moreover, we will study if the technical hypotheses described above can be applied into a concrete use case in the realm of smart healthcare and tele-rehabilitation. This would target showing the feasibility and implementability of the theoretical and technical contributions.

The research questions and hypotheses imply certain assumptions, including the scope of the devices to be included in the study. These will include both brownfield devices (legacy devices with little or no intelligent behavior) and green field devices (those where the swarm and agent based techniques can be deployed). This means that not all devices can necessarily implement all the features included in the PhD proposal. Also, the selection of devices will be driven by the use case, at least in the prototyping phase, even if for simulation evaluations a broader choice can be made. In addition, we will assume the existence of an orchestrator for setting up the organization schemes. This does not necessarily means centralization, given that even the orchestrator could change depending on the given settings.

In summary, using the proposed hypotheses we will study how semantic models can be used as lingua franca so that the agent-based devices can be orchestrated following swarm computing self-organization techniques. According to the research approach, we would follow an interactive process, first evaluating the feasibility of using the semantic models, then analyzing, designing, and implementing the agent-driven behaviors directed towards achieving their corresponding goals, and finally implementing the swarm-based orchestration, In all cases, the formalization will be followed by a design and implementation evaluated first using simulated environments where multiple configurations can be tested and assessed. Afterwards, the prototype implementation targeting smart healthcare will be architectured, implemented and tested, to show the general feasibility in a closer to real-conditions environment.

4 RELATED WORK

Swarm-based techniques for task orchestration have been studied in the context of cloud computing. One such work is QRAS [15], which proposes to use an ant colony optimization to provide more predictable results. Other studies propose cuckoo search algorithms [14] and modified dragonfly [16] to optimize resource management and VM placement respectively. Already in these settings, these approaches show promising results, especially regarding task allocation.

Swarm-inspired approaches have been applied to an ample spectrum of application domains including edge computing devices. Among these we can mention its use in unmanned aerial vehicles (UAVs) assisted by edge computing nodes [11], where energy savings are of primordial importance. Also, in the field of smart manufacturing [6], this type of technologies has been applied, most notably through group learning techniques, with the aim of improving processing efficiency. Industry 4.0 applications as the one described in [13] also show the potential of swarm-based organization of shared knowledge among edge devices. This line of work showcases the potential for knowledge sharing nodes that can influence not only the scheduling but also the (dynamic) behaviors of participating nodes. Similarly swarm optimization algorithms have been employed in edge-powered vehicles, for instance benefiting from GPU architectures for task offloading [1].

Considering the potential that these swarms can have for the self-organization of edge nodes and smart devices, we identified relevant works regarding the management of heterogeneous agent knowledge. In fact, multiple schemas and ontologies exist for representing sensor and IoT devices and data [2, 3, 7]. Based on these works we will evaluate if these models are enough to allow self-organization, or what extensions would be required.

Regarding the mediation, we will investigate if existing swarm-inspired models [5] can be propagated to smart edge agents, so that participating nodes in a swarm can expose their capabilities. Existing mediation and agency patterns have not been yet used in edge environments in swarm-inspired architectures.

Also, in the context of the Web there is the possibility of social coordination among agents in a swarm environment. There have been little efforts in this direction, but there is a good potential of using the Web as an environment for establishing knowledge exchanges [12].

One particular aspect in this respect is the negotiation among participating entities. Negotiation in the context of swarm agents can bring both advantageous optimization patterns, as well as delays and potential conflicts. It may become crucial to identify ways of reducing the risks, and to find compromises that adapt to changing situations and redefinition of goals, beyond existing approaches investigated in the literature [9, 10].

In this context the risks of manipulation and mischievous behaviors is more than plausible. The conception of integrity, confidence, and transparency mechanisms need to be integrated into this research road-map. Edge and sensing devices are often required to acquire and handle sensitive data, which should be subject to strict privacy constraints, while keeping the computation and processing goals [4, 17].

5 RESEARCH APPROACH

In the context of the research problems and questions raised, this PhD proposal aims at addressing these challenges with the combination of semantic modeling, swarm computing and agent-based systems.

This approach will provide a novel concept for creation of swarm intelligent apps, which can be instantiated from existing templates and matched to available (swarms of) devices. Through this approach, powered by semantic models, we would allow a declarative way of specifying application objectives and behaviors. These semantic templates will define machine interpretable specifications of high-level business requirements, which can be easily instantiated as a swarm intelligent application. This ensures an easy matching of templates inputs and outputs with interfaces exposed by devices. This ensures they can be easily discovered and configured for a new business goal. They will enable semantic-based online discovery, matchmaking, and configuration for embedded AI apps, as well as orchestration and monitoring of swarm intelligence apps.

To achieve this the project will specify goals and behaviors using a declarative semantic model proposed based on the Web of Things and Things Description ontologies. The capabilities specified using these models would be matched with the needs of smart applications, thus enabling the orchestration of the app, and initiating the interactions among the swarm devices. The templates themselves will be specified using semantic descriptors, and the data flows will be enacted using the data communication features of the devices. An expert developer will be able to use these abstractions rather than dealing with the complexity of low level configuration and programming of interactions among devices. The semantic representation of the nodes' capabilities, goals and behaviors will therefore constitute a common and machine understandable description of what the swarm should do. The developer may not necessarily configure these semantic representations, but rather use high-level environments to produce them. One such high-level environment we may use is Mendix Development Kit and its extension modules. One may imagine that these semantic documents can be reused and extended, for instance through a search and discovery system.

Using the semantically enriched models, the system can instantiate them, i.e., the different roles specified can be assumed by concrete and specific nodes in the swarm. The instantiation in a first step will consist in the passing of the semantic model to the orchestrator, who will be in charge of executing it on run-time. This will include the coordination of the operations to be executed by different nodes in the swarm. Therefore, the orchestrator will also include the instantiation of the semantic description of the tasks, goals, sub-tasks, and skills established beforehand. With this information the orchestrator will be able to know, for example, that it needs the participation of (smart) edge nodes with certain skills (e.g., stream reasoning over sensor measurements), and then it will need to find and discover which nodes comply with these requirements. In certain cases, the orchestrator may not find the necessary resources to achieve the required goal, and it could either fail or latently wait until the necessary resource can be scheduled. In case of a successful node discovery, then the orchestration itself will be organized.

The orchestrator may often coincide with the swarm coordinator, and as such will organize the forming and management of the swarm, with different nodes joining or leaving when needed. In certain cases, the edge nodes may have the capability of holding their own semantic descriptions (e.g., skills or affordances) so that they can be found and incorporated at runtime in the Swarm. In any case the orchestrator will need to perform a matching operation between the set of available nodes and the different roles established in the model, and respecting the specification of skills, goals, and tasks. Once this has been established, the different nodes will be able to start operating and exchanging the different data/messages, potentially using stream data operations components in the background.

6 EVALUATION PLAN

The evaluation will be performed using different settings for different parts of the project. Regarding the evaluation of the semantic discovery and formation of the swarm, first a simulation environment will be designed and implemented. This environment will allow the configuration of multiple scenarios representing diverse situations in which participating nodes may need to interact in order to form the swarm. This will include first the instantiation of the orchestrator, and the dynamic join of nodes into the swarm after a discovery phase. This will allow evaluating certain metrics related to the compliance with swarm formation specifications, compliance with the predefined goals, finding of available/suitable nodes for the given tasks, performance of the swarm formation, etc.

Next, an evaluation of the agent-based mechanisms will be performed, most notably regarding the usage of behaviors and goals declared for both the swarm as a whole and the edge devices individually. Again, the evaluation will be performed first on a simulated environment in order to maximize the topologies and configurations of the swarms and the participating agents. On a second step, the agent-based mechanisms will be implemented in the healthcare prototype, in order to assess the feasibility of the approach.

Finally, the orchestration and self organization strategies will be implemented and evaluated again using both the simulator and the prototype, focusing on the efficacy of the formation of the swarm, and the adaptability to dynamic node changes, discovery of nodes through capabilities, and effective execution of the proposed goals.

7 CONCLUSIONS AND REFLECTIONS

At the moment this PhD work is at its very beginning, so there are still a number of aspects to be refined and further explored. In any case, addressing the proposed challenges, swarm intelligent agents for smart edge may open a number of research opportunities some of which will be covered in the scope of this project. Some of the challenges which we may face during the next few years, are included the following.

Semantic representation of shared goals, organization patterns, and behaviors have the potential to facilitate agent coordination and negotiation in swarm smart edge environments. however it is still to be evaluated in these can effectively be used by low-code solutions can help abstracting the complexity of smart edge agents, and their deployment under heterogeneous conditions. Also, smart edge agents necessarily need to go well beyond the Web as a means to interact. Nevertheless, the Web offers a solid and standard mechanism to enable the interoperability, discovery, and accessibility among different swarms. Fully decentralized swarms of agents on the Web can make use of existing approaches for self-organization, although it is yet unclear which ones can really be adapted to the specific conditions of edge nodes. Finally, the issues related to privacy, which have not really been considered, may also limit the possibilities of knowledge sharing among agent-based edge nodes.

ACKNOWLEDGMENTS

This PhD project is partially supported by the Horizon Europe project SmartEdge under grant agreement No. 101092908.

REFERENCES

[1] Mohamed A Alqarni, Mohamed H Mousa, and Mohamed K Hussein. 2022. Task offloading using GPU-based particle swarm optimization for high-performance vehicular edge computing. *Journal of King Saud University-Computer and Information Sciences* 34, 10 (2022), 10356–10364.
[2] Maria Bermudez-Edo, Tarek Elsaleh, Payam Barnaghi, and Kerry Taylor. 2016. IoT-Lite: a lightweight semantic model for the Internet of Things. In *International IEEE conferences on ubiquitous intelligence & computing*. IEEE, 90–97.
[3] Victor Charpenay and Sebastian Käbisch. 2020. On modeling the physical world as a collection of things: The w3c thing description ontology. In *The Semantic Web: 17th International Conference, ESWC 2020, Heraklion, Crete, Greece, May 31–June 4, 2020, Proceedings 17*. Springer, 599–615.
[4] Lijie Chen, Shaojing Fu, Liu Lin, Yuchuan Luo, and Wentao Zhao. 2022. Privacy-preserving swarm learning based on homomorphic encryption. In *Algorithms and Architectures for Parallel Processing: 21st International Conference, ICA3PP 2021, Virtual Event, December 3–5, 2021, Proceedings, Part III*. Springer, 509–523.
[5] Christophe Guéret, Stefan Schlobach, Kathrin Dentler, Martijn Schut, and Gusz Eiben. 2012. Evolutionary and swarm computing for the semantic web. *IEEE Computational Intelligence Magazine* 7, 2 (2012), 16–31.
[6] Jianli Guo and Miguel Martinez-Garcia. 2021. Key technologies towards smart manufacturing based on swarm intelligence and edge computing. *Computers & Electrical Engineering* 92 (2021), 107119.
[7] Krzysztof Janowicz, Armin Haller, Simon JD Cox, Danh Le Phuoc, and Maxime Lefrançois. 2019. SOSA: A lightweight ontology for sensors, observations, samples, and actuators. *Journal of Web Semantics* 56 (2019), 1–10.
[8] Vikash Jugoo and Vimolan Mudaly. 2016. The use of action research in a computer programming module taught using a blended learning environment. *International Scientific Researches Journal* 72 (2016), 52–69.
[9] Najwa Kouka, Raja Fdhila, and Adel M Alimi. 2017. Multi objective particle swarm optimization based cooperative agents with automated negotiation. In *Neural Information Processing: 24th International Conference, ICONIP 2017, Guangzhou, China, November 14–18, 2017, Proceedings, Part IV 24*. Springer, 269–278.
[10] Fabiana Lorenzi, Daniela Scherer dos Santos, and Ana LC Bazzan. 2005. Negotiation for task allocation among agents in case-base recommender systems: a swarm-intelligence approach. In *2005 International workshop on multi-agent information retrieval and recommender systems*. 23–27.
[11] Yiming Miao, Kai Hwang, Di Wu, Yixue Hao, and Min Chen. 2022. Drone Swarm Path Planning for Mobile Edge Computing in Industrial Internet of Things. *IEEE Transactions on Industrial Informatics* (2022).
[12] Shahab Mokarizadeh, Alberto Grosso, Mihhail Matskin, Peep Kungas, and Abdul Haseeb. 2009. Applying semantic web service composition for action planning in multi-robot systems. In *2009 Fourth International Conference on Internet and Web Applications and Services*. IEEE, 370–376.
[13] Xiaoyang Rao and Xuesong Yan. 2022. Particle swarm optimization algorithm based on information sharing in Industry 4.0. *Wireless Communications and Mobile Computing* 2022 (2022).
[14] B Reshmi, Kerala Palakkad, and P Poongodi. 2021. Bigdata and HPC Convergence with Locality Based Cuckoo Search Method. *International Journal of Innovative Science and Research Technology* (2021).
[15] Harvinder Singh, Anshu Bhasin, and Parag Ravikant Kaveri. 2021. QRAS: efficient resource allocation for task scheduling in cloud computing. *SN Applied Sciences* 3 (2021), 1–7.
[16] Atul Tripathi, Isha Pathak, and Deo Prakash Vidyarthi. 2020. Modified dragonfly algorithm for optimal virtual machine placement in cloud computing. *Journal of Network and Systems Management* 28 (2020), 1316–1342.
[17] Bowen Zhao, Ximeng Liu, An Song, Wei-Neng Chen, Kuei-Kuei Lai, Jun Zhang, and Robert H Deng. 2022. PRIMPSO: A Privacy-Preserving Multiagent Particle Swarm Optimization Algorithm. *IEEE Transactions on Cybernetics* (2022).