Multi-Step Reasoning for IoT Devices

José Miguel Blanco[1] [a] and Bruno Rossi[2] [b]
[1] *Escuela Técnica Superior de Ingenieros de Telecomunicación, Universidad Politécnica de Madrid, Spain*
[2] *Faculty of Informatics, Masaryk University, Brno, Czech Republic*

Keywords: Internet of Things (IoT), Web of Things (WoT), Semantic Reasoning.

Abstract: Internet of Things (IoT) devices are constantly growing in numbers, forecasted to reach 27 billion in 2025. With such a large number of connected devices, energy consumption concerns are a major priority for the upcoming years. Cloud / Edge / Fog Computing are critically associated with IoT devices as enablers for data communication and coordination among devices. In this paper, we look at the distribution of Semantic Reasoning between IoT devices and define a new class of reasoning, *multi-step reasoning*, that can be associated at the level of the edge or fog node in the context of IoT devices. We conduct an experiment based on synthetic datasets to evaluate the performance of *multi-step reasoning* in terms of power consumption, memory, and CPU usage. Overall we found that *multi-step reasoning* can help in reducing computation time and energy consumption on IoT devices in the presence of larger datasets.

1 INTRODUCTION

Internet of Things (IoT) refers to a large number of physical devices, sensors, software artifacts, all with processing capabilities that can interact over the more disparate communication network to provide smart services to users (Dorsemaine et al., 2015; Atzori et al., 2017). The number of IoT devices deployed globally in the world has constantly been growing, increasing 9% in 2021, representing 12.3 billion connected devices with 27 billion IoT devices forecasted for 2025 (Hasan, 2021).

Natural support to IoT was the emergence of Edge and Fog Computing – bringing the computation and data storage closer to the devices that are generating, interacting, and integrating data (Ai et al., 2018; Satyanarayanan, 2017). As such, with Edge Computing, the data is processed directly on the device/sensor at the edge without any data transfer, while in Fog Computing, data can be processed in an intermediary node close to the edge, differently to Cloud Computing in which data needs to be propagated back to the Cloud for processing (Ai et al., 2018). The advantage in the context of IoT is that the large amount of data generated by different IoT devices can be processed directly at the edge of the network – on the other hand, more processing power is required (Ai et al., 2018).

While Edge Computing can bring services with faster response time, reasoning at the level of edge devices still needs to take into account the knowledge generated by peer devices to grant Semantic Reasoning based on the knowledge derived from all the connected IoT devices (Jara et al., 2014). Furthermore, energy consumption patterns need to be considered when moving reasoning from the Cloud to edge devices (Mocnej et al., 2018; Cui et al., 2018).

In this paper, we look at the concept of Semantic Reasoning in distributed and connected IoT devices. In particular, we look at how different distributions of reasoning at divergent moments in time can influence computational performance and energy consumption during the reasoning process. We set up an experiment with synthetic datasets looking at different reasoning rules and evaluating the performance of alternative reasoning types.

We have the following contributions in this paper:

1. Definition of the concept of multi-step reasoning for Semantic Reasoning of IoT devices at the Fog or Edge level;
2. An experiment with synthetic generated data, studying the different properties of multi-step reasoning – the experimental package is available on Figshare (see Section 5);

The paper is structured as follows. In Section 2 we review the background about IoT, Edge and Fog Com-

[a] https://orcid.org/0000-0001-9460-8540
[b] https://orcid.org/0000-0002-8659-1520

Blanco, J. and Rossi, B.
Multi-Step Reasoning for IoT Devices.
DOI: 10.5220/0011772700003464
In *Proceedings of the 18th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2023)*, pages 330-337
ISBN: 978-989-758-647-7; ISSN: 2184-4895

puting and the connection with Semantic Reasoning. In Section 3 we provide the fundamental definitions of multi-step reasoning. In Section 4 we provide a scenario in the context of a smart energy domain with charging station, smart homes, and electric vehicles – all needing reasoning capabilities for the most efficient utilization of resources. In Section 5 we delve into the experimental design, explaining how the fundamentals of multi-step reasoning were instantiated in terms of implementation and the datasets used for the evaluation of computational performance and energy consumption. In Section 6 we provide the main results from the experiment together with the threats to validity. In Section 7 we provide a discussion about the result in relation to the proposed research questions. Section 8 concludes the paper.

2 BACKGROUND

The emergence of Cloud Computing was one of the cornerstones for the evolution and growth of IoT (Satyanarayanan, 2017). Initially, IoT devices would represent devices with limited resources, thus with reduced possibility of in-place reasoning algorithms implementations and data aggregation (Lea, 2018). However, as the added value of IoT devices is in the aggregation and integration of data from heterogeneous and numerous devices (such as thousands of sensors, cameras, smartphones, etc...), the Cloud constituted the first step towards added capabilities — in terms of data ingestion, aggregation, analytics, machine learning, among others (Fig. 1).

As such, Edge Computing, placing computation at the edge of the network in proximity to mobile devices and sensors, emerged as a way to reduce latency issues and move the data processing part closer to the devices themselves (Satyanarayanan, 2017). Conversely, Fog Computing emerged as an extension to the Cloud to support devices with a layer between the edge and the Cloud (Fig. 2). Both Edge / Fog Computing have advantages over traditional IoT Cloud-based topologies: more responsive services, support of edge analytics for scalability, enforcement of privacy policies, or the possibility to mask Cloud outages (Satyanarayanan, 2017).

However, one critical turning point for the decision about the network topology is that communication from the Cloud infrastructure to edge gateway and edge devices is subject to disparate communication latency: from the real-time performance at the level of smart devices to milliseconds latency at the edge gateway and potentially milliseconds to seconds at the level of the Cloud (Lea, 2018).

Furthermore, recently *Green Computing* concerns have emerged: green computing has been defined as *"the environmentally responsible and eco-friendly use of computers and their resources"* (Salama, 2020). With the increasing number of connected devices and computational units, the goal of efficient power consumption has become a key target for the IT domain, making considerations about power consumption one of the key aspects of network and computation optimizations— leading to the so-called Green IoT (Zhang et al., 2018; Lyu et al., 2018). This brings considerations about the power consumption of computations and communications related to edge devices (Mocnej et al., 2018; Cui et al., 2018).

Additionally, IoT devices have become pervasive in most of our everyday life aspects: from granting smart sensors to providing remote access to different systems (Arfaoui et al., 2020). Nevertheless, this has also prompted some problems regarding the interoperability of these devices. One of the key points for the extended connectivity of IoT devices, and solving the interoperability problem, is their interaction with the Semantic Web through the standard known as the Web of Things (WoT). WoT provides standardized metadata and other re-usable technological building blocks to enable easy integration across IoT platforms and application domains[1]. This allows for processing the generated data through Semantic Reasoning to power the inference capabilities of devices (Kisliuk et al., 2022) or increase the security for knowledge bases queries (Krishnasamy-Sivaprakasam and Slutzki, 2021). In particular, the extension of IoT into WoT allows for developing a formal space with all the advantages that can be extracted from it.

The techniques used to extract knowledge from the data generated by IoT devices, such as sensors, is by semantic reasoning (Maarala et al., 2017). This method is based on the use of logical rules to derive conclusions from the annotations that the semantic web framework has provided (Hitzler and Van Harmelen, 2010). These rules, applied by using a semantic web reasoner and sometimes strengthened by the structure provided by an ontology, help to present the user with an extended representation of the situation so any measures can be adopted. In particular, it is necessary to mention that the deployment of the reasoner and rules, with regards to the architecture, its absolutely critical in time constrained situations. For that matter some effort has been made to ensure that the required tasks are performed as fast as possible (Wang et al., 2018).

[1] https://www.w3.org/WoT/

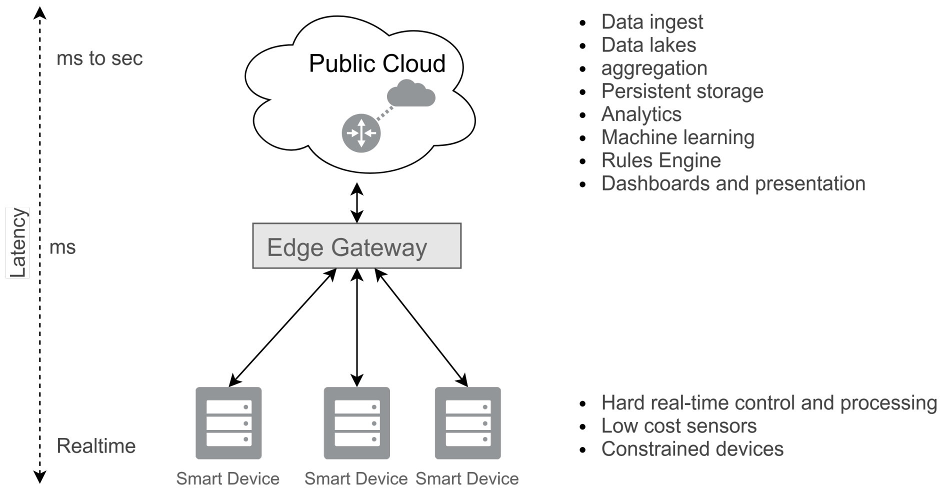

Figure 1: Cloud, Edge Gateway, and Smart Devices (adapted from (Lea, 2018)).

3 FUNDAMENTALS

In this section, we detail the main fundamentals about Semantic Reasoning in the context of IoT devices and introduce the concept of multi-step reasoning.

Definition 1 (Reasoning)**.** *Reasoning is the automatic procedure to generate new knowledge and relationships based on provided data and a set of additional rules. Generally speaking, this procedure executes the schema provided by the rules to obtain new data at once.*

Definition 2 (Multi-step reasoning)**.** *Multi-step reasoning is a method that acts as a subclass of Reasoning as defined above. This subclass is characterized by executing the rules' schema one at a time instead of all of them at once.*

Now, we will consider the following rules:

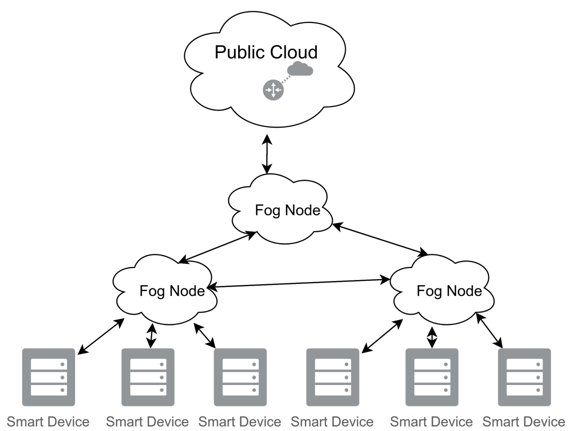

Figure 2: Fog Computing and smart devices (adapted from (Lea, 2018)).

R1. [ruleConjunction: (?a :and ?b) (?b :and ?c) → (?a :and ?c)]

R2. [ruleTransitivity: (?a :then ?b) (?b :then ?c) → (?a :then ?c)]

Both rules share the same structure: they have two different triples in the first term that share the same predicate, *:and* and *:then*, and both have as a conclusion one triple in which elements of both of the previous triples are combined. This means that the rules are comparable in terms of depth of the triples, elements of the terms and general structure. One might argue that both rules are indeed just different instances of the same general rule. A general rule that could be read as follows:

[ruleGeneral (?a :predicate ?b) (?b :predicate ?c) → (?a :predicate ?c)]

And while the argument holds some value, it is necessary to understand that these rules have not been selected by their compliance with the Resource Description Framework (RDF), one of the most common building blocks of the semantic web, structure, but rather for their logical value. Because of this, we could understand the rules as Hilbert-style rules formalized as follows:

R1. $(A \wedge B) \,\&\, (B \wedge C) \Rightarrow (A \wedge C)$

R2. $(A \rightarrow B) \,\&\, (B \rightarrow C) \Rightarrow (A \rightarrow C)$

As we can see, there are no common connectives between the two rules aside from the meta-connectives & and ⇒. Therefore, if we are defining both connectives of the rules, conjunction (∧) and conditional (→) as primitives from a certain algebra, we should be in the clear as to why both rules are of interest and cannot be inferred one from the other.

4 IoT SCENARIO

To showcase the concept of *multi-step reasoning* in the cyber-physical context, we can provide a scenario with Smart Meters, data aggregators/concentrators, and central servers in a typical smart energy application (e.g., (Chren et al., 2016)), for which there has already been semantic work done (Blanco et al., 2023).

An example of the proposed network is the one established by a Smart City in which we would have smart devices and smart sensors at the edge that require total availability to capture the multiple streams of data in real-time (e.g., (Schleicher et al., 2016)). On the other hand, the data center acting as a central server would be used to process the data received from the edge devices, ensure that all the services are working properly, and make possible predictions on future events and requirements. The final third element is a data aggregation point. This point would collect the data generated by the smart devices and smart sensors and pre-process it so the work that needs to be done in the central server would be less and could be carried out in a more manageable and efficient way.

This network would go on to generate certain data in the shape of triples. Some of these triples could be the following:

T1. :smartcar :in :chargestation

T2. :chargestation :in :mainstreet

T3. :smartcar :charges :batterycar

T4. :batterycar :charges :1.5Kw/h

The first two triples, T1 and T2, could be processed with the *ruleConjunction* that we have included before. In the same way, T3 and T4 could be processed with *ruleTransitivity*. The outcomes would be, respectively, as follows:

T5. :smartcar :in :mainstreet

T6. :smartcar :charges :1.5Kw/h

This would allow the main data center of the Smart City to know about the situation of the car as well as the task that is being performed. In this case, we would be able to infer that the car is being charged at a rate of 1.5Kw/h on the main street. Information that should be fundamental when considering the energy required by a certain part of the city and the traffic management of said part of the city. This obviously is not a problem when dealing with just one element (in this case, the smart car) but could be a potential overload when dealing with thousands of vehicles and the streams of data that are constantly generated. Therefore, it is necessary to find a way to improve the performance of reasoning when dealing with datasets that can contain thousands, if not millions, of elements.

5 EXPERIMENTAL EVALUATION

We run an experiment with the following goal: to evaluate the *multi-step reasoning* (Definition III.2) in terms of power consumption and computational performance of IoT devices when compared with traditional reasoning (Definition III.1).

Graphs of the results discussed in this and further sections and all the models and experiment's scripts are available in a Figshare repository (Blanco and Rossi, 2023)[2] where the reader might consult them at any given time.

For multi-step reasoning, we adopted *Apache Jena* and used N3 files. In a production environment, the usual flow is to load the rules after one reasoning model is created. The traditional technique for implementing a reasoning model (Def III.1) consists of creating the model and loading it with a set of rules that encompasses all the rules that are to be used or that are plainly considered, even if they are not used. In the multi-step reasoning paradigm we are proposing (Def III.2), this will happen as different instances of loading sets of rules, one for each rule. The loading of these sets excludes the rules that are not used in further pursuit of efficiency and performance as, even if it is just a fraction of a second, some time, energy, and other resources would be saved.

5.1 Research Questions

We have the following Research Questions (RQs):

Research Questions

RQ1. What is the energy consumption when the reasoning is done as multi-step reasoning compared when it is done traditionally?

RQ2. Can the performance of the reasoning (time and memory / CPU usage) be improved when it goes through the means of multi-step reasoning as opposed to the traditional technique?

5.2 Datasets

For the experimental evaluation, we used datasets that have been synthetically generated. These datasets

[2]https://doi.org/10.6084/m9.figshare.19493996.v1

have been created sequentially with a number n, for $n \in \mathbb{N}$, of elements per rule thanks to a *Python* script. The datasets contain an increasing number of cases that are constructed in such a way that the rules introduced are forced to process all the elements of the datasets and all the outcomes. This means that, at the end of the reasoning, the reasoner would deal with a relationship between the first and last elements. An excerpt of any of these datasets is as follows:

```
...
:n-1 :and :n-2
:n-1 :then :n-2
:n :and :n-1
:n :then :n-1
```

These datasets have been designed so there is no data overload on the size of the triples, but rather can test the performance of the rules engine. Furthermore, they can provide a scenario in which the model is asked to perform a scalable number of tasks, thus showing the performance of the reasoner against multiple, totally different scenarios. The different datasets are, by name and total number of elements: DS1(20), DS2(200), DS3(500), DS4(1000), and DS5(2000).

Table 1: Average Energy Usage.

	Traditional Reasoning (J)	Multi-Step (J)
DS1 - CPU cores	48.74	56.48
DS2 - CPU cores	127.17	122.71
DS3 - CPU cores	723.02	520.65
DS4 - CPU cores	9 723.97	6 493.61
DS5 - CPU cores	157 939.74	105 357.26
DS1 - RAM	17.53	18.20
DS2 - RAM	24.49	24.95
DS3 - RAM	125.57	90.51
DS4 - RAM	1 659.19	1 084.57
DS5 - RAM	32 385.26	18 921.56
DS1 - PKG	114.07	126.12
DS2 - PKG	214.90	213.58
DS3 - PKG	1 209.13	866.76
DS4 - PKG	16 616.42	10 862.94
DS5 - PKG	282 784.75	183 546.67

5.3 Analysis Method

The energy usage has been measured using the *perf* tool and is in Joules (J). The machine used is a *Intel CPU Core i7-4790@3.60GHz, 16Gb RAM DDR3 @ 1600 MHz*, and running *Ubuntu 20.04.3 LTS*. The IDE utilized for the implementation and testing is *IntelliJ Idea Ultimate 2021.2*. Each benchmark was performed five (5) times, considering the average. Finally, the reasoning engine was implemented using *Apache Jena 4.4.0*.

Table 2: Average Measurements for Traditional Reasoning.

	Time (ms)	Max Memory (KB)	Max CPU (%)
DS1	256	13 054	6.96
DS2	1 833	120 450	39.10
DS3	40 178	157 445	32.94
DS4	667 629	222 322	33.16
DS5	13 025 579	353 993	34.90

6 EXPERIMENTAL RESULTS

The initial measurements, performed with the smallest dataset, DS1 of 20 elements, show that the performance of traditional reasoning is better than the new alternative introduced by multi-step reasoning (481ms vs 225.9ms on average, Tables 2,3). The results show an increase of almost a 100% in time to perform the reasoning when comparing the multi-step reasoning to the traditional way. Despite that, the maximum memory and CPU usage are similar. In a similar sense, only the energy consumption of the CPU cores is above a 15% in multi-step reasoning, while the energy consumption of the RAM and the PKG remains stable all through the different techniques (Table 1).

Table 3: Average Measurements for Multi-Step Reasoning.

	Time (ms)	Max Memory (KB)	Max CPU (%)
DS1	481	12 902	6.56
DS2	1 942	30 484	34.38
DS3	26 543	150 845	33.63
DS4	403 304	190 399	31.78
DS5	7 358 757	245 303	38.4

Similarly, the results for DS2, of 200 elements, are comparable in time spent and Maximum CPU usage. Also, the energy consumption across CPU cores, RAM, and PKG is almost the same in traditional and multi-step reasoning (Table 1). The main difference resides in the max memory usage, where traditional reasoning requires almost four (4) times the amount that multi-step reasoning does (Tables 2,3).

By the time that we reach DS3, of 500 elements, the multi-step reasoning starts to show improvements. Time spent reasoning is almost cut in half by multi-step reasoning in comparison to traditional methods (26 543ms vs 40 178ms, Tables 2,3). Also, the energy consumption of the CPU cores, the RAM and the PKG in traditional reasoning are increased by a third once compared to multi-step reasoning (Table 1). One trend that we see at this point and is maintained in further datasets is that the Maximum Memory and CPU usage are similar in traditional and multi-step reasoning. Nevertheless, the overall CPU usage is lower in multi-step reasoning with lower spikes over time.

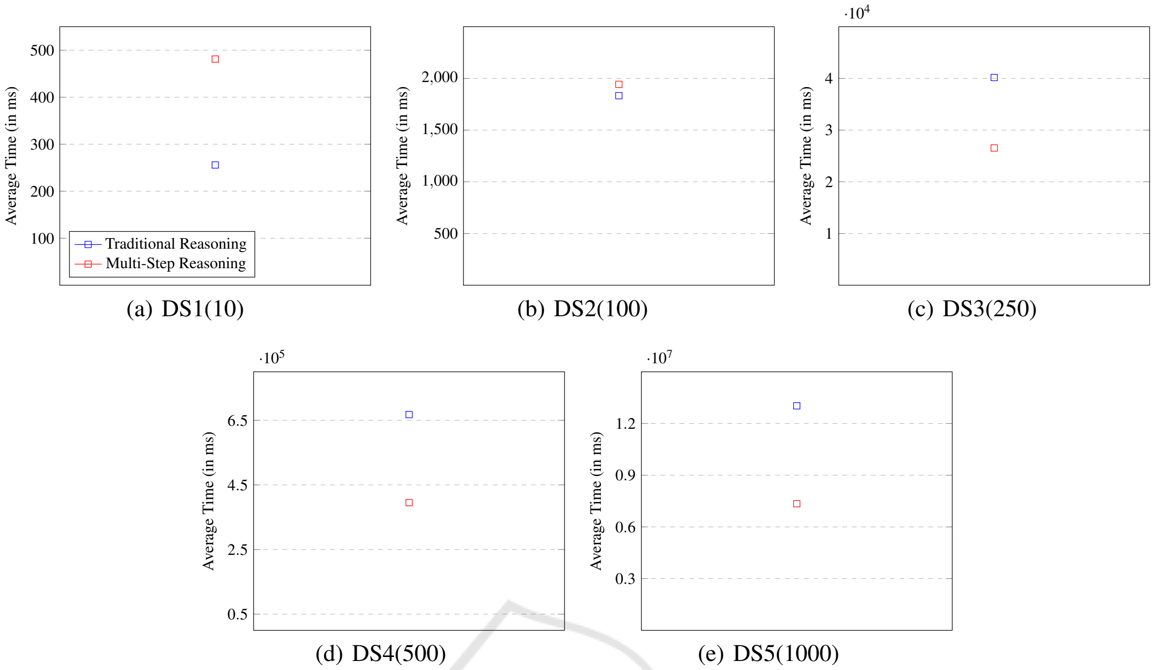

Figure 3: Average execution times of different datasets (note different scale).

The most significant differences are perceived in the test runs on DS4 and DS5. DS4, of 1000 elements, shows that the energy consumption of traditional reasoning is 1.5 times bigger than that of multi-step reasoning in all the measurements made: CPU cores, RAM and PKG (Table 1). Both maximum memory and CPU usage are similar as we pointed in the previous paragraph. Nevertheless, the time of reasoning, when done traditionally, is more than 1.5 longer, which, at this point, accounts for more than 4 minutes of difference (404 304ms (6.7min) vs 667 629ms (11.1min), Tables 2,3).

In case of DS5 of 2000 elements, the energy consumption almost doubles in PKG and RAM, while it is up to a 1.5 times more than multi-step reasoning in CPU cores (Table 1). Both maximum memory and CPU use is comparable in both approaches. However, the time spent reasoning is where we see a difference of double the amount in traditional against multi-step reasoning. If in DS4 the time difference accounted for 4 minutes of difference, here we have a difference of up to two hours on the averages (7 358 757ms (∼122min) vs 13 025 579ms (∼217min), Tables 2,3).

6.1 Threats to Validity

We have some threats to validity of the experiment.

First of all, in this paper, we deal with the *reasoning* that is performed on the edge device or at the fog node by looking at the different reasoners' rules that can be involved. We do not deal with the network aspects, such as delays and transmission power requests, that can have a relevant impact. Other research has dealt with measuring an extra amount of energy and time required for the transmission and reception of tasks dived among edge and fog nodes – so-called offloading (Bozorgchenani et al., 2018). Having a view that includes Semantic Reasoning and networking considerations can be considered the next step for research.

The experiment has only been performed by looking at two different rules. The results might vary when the set of rules reaches a certain size or that, after a certain point, we could see a diminishing returns situation. We also considered synthetic data generated according to typical scenarios. We plan to extend the evaluation to the reasoning on IoT devices.

Furthermore, we can point out that our testing was done utilizing Apache Jena's reasoner. Other reasoning models might perform differently. Nevertheless, this should be a limited concern as the most extended and adopted reasoning model is the one of Apache Jena. The configuration details and experimental package are available in Section 5.

Additionally, we performed the experiment considering abstract rules instead of an ontology. We intended to test with synthetically generated datasets with the idea that techniques should be first tried out in an *in-vitro* environment before including other variables such as ontologies, as those can affect performance as seen in (Blanco et al., 2021). We plan to adopt more complex models in future works.

7 DISCUSSION

As the results have shown, the traditional way of reasoning has a higher performance, and lower energy consumption than multi-step reasoning, whenever the datasets are below 100-200 elements for each rule. This is related to the time needed to perform the reasoning, as traditional reasoning takes reduced time compared to multi-step reasoning.

Nevertheless, this trend changes whenever the datasets get larger. Once the barrier of 100 elements per rule is passed – in the case of this experiment, once the 500 elements dataset has been reached – the performance, time spent for reasoning, and energy consumption are in favour of multi-step reasoning. At first, it is a discreet improvement, with a difference of only mere seconds. However, once the dataset reaches the size of 2000 elements, we see a consistent improvement.

RQ1 Findings

☞ Data collected in Table 1, shows that multi-step reasoning is more energy efficient way to process data. The disparity between energy consumption gets bigger with the increasing cardinality of the datasets.

Although for the first dataset multi-step reasoning shows that it is not the best option for small datasets, results improve on the other datasets with larger number of elements. This propels multi-step reasoning as a more energy efficient option than traditional reasoning, especially under large workloads.

RQ2 Findings

☞ The performance in terms of memory/CPU usage is better when addressed from a traditional perspective as long as the datasets remain small (up to 100 elements). From this point on, the introduction of multi-step reasoning improves the performance, in terms of memory/CPU and time spent (Tables 1,2,3).

About RQ2, it is worth noting that CPU does not see a difference as the one that can be noted in time and energy consumption. Particularly, if we strictly speak about maximum numbers, we can see that the max CPU usage in multi-step reasoning is higher than the traditional one, but this is remedied by a difference in the usage of the CPU. We can also note different patterns in CPU usage comparing multi-step and traditional reasoning[3] with more spikes in % of CPU usage. The rest of the data shows that multi-step reasoning is much more efficient in both time needed and max memory utilized. While the difference in max memory might not be significant, the difference in time is more prominent. This can be seen as the reduction in time needed to compute the biggest dataset is almost double, with a difference of almost 2 hours. This would appear crucial for developing a reasoning scheme to process data in real-time.

Based on the experimental data, we can suggest that multi-step reasoning is a better technique than traditional reasoning, as it improves all facets: time, max memory, max CPU, and energy consumption. Given the current state of the art and the amount of data generated by the billions of online devices, the performance in larger datasets is more relevant. The difference in the case of small datasets is minimal: less than a quarter of a second and less than 20 Joules overall for datasets of 20 elements and less than one-tenth of a second in datasets of 200 elements.

8 CONCLUSION

With the importance that Edge and Fog Computing have acquired in recent years, it is fundamental to support data processing at the level of IoT devices. One step in this direction is to provide Edge devices with Semantic Reasoning capabilities.

In this paper, we have looked at the concept of Semantic Reasoning in distributed and connected IoT devices. In particular, we proposed *multi-step* reasoning as an approach to improve the performance of reasoning and conducted an experiment with synthetic data to evaluate IoT devices' computational performance and energy consumption.

Overall, we found that multi-step reasoning can help reduce computation time and energy consumption in larger datasets. Our *in-vitro* evaluation will be further expanded in future works considering the deployment of ontologies. It is also crucial to expand the work to include more complex rules that intertwine; this would allow us to showcase the performance of multi-step reasoning in a situation where the rules are not considered as stand-alone entities but rather a complex system that benefits from previous results to get more robust conclusions.

[3]The figures are available in the additional material on Figshare

ACKNOWLEDGEMENTS

The research was supported from ERDF/ESF "Cyber-Security, CyberCrime and Critical Information Infrastructures Center of Excellence" (No. CZ.02.1.01/0.0/0.0/16_019/0000822).

This work has been supported by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with Universidad Politécnica de Madrid in the line Support for R&D projects for Beatriz Galindo researchers, in the context of the V PRICIT (Regional Programme of Research and Technological Innovation).

REFERENCES

Ai, Y., Peng, M., and Zhang, K. (2018). Edge computing technologies for internet of things: a primer. *Digital Communications and Networks*, 4(2).

Arfaoui, A., Cherkaoui, S., Kribeche, A., and Senouci, S. M. (2020). Context-aware adaptive remote access for iot applications. *IEEE Internet of Things Journal*, 7(1).

Atzori, L., Iera, A., and Morabito, G. (2017). Understanding the internet of things: definition, potentials, and societal role of a fast evolving paradigm. *Ad Hoc Networks*, 56.

Blanco, J., Rossi, B., and Pitner, T. (2021). A comparison of smart grids domain ontologies. In *Proceedings of the 17th International Conference on Web Information Systems and Technologies - Volume 1: WEBIST,*, pages 115–123. INSTICC, SciTePress.

Blanco, J. M. and Rossi, B. (2023). Multi-Step Reasoning for IoT Devices.

Blanco, J. M., Rossi, B., and Pitner, T. (2023). A comparative study of energy domain ontologies. In *Web Information Systems and Technologies*, pages 43–58. Springer International Publishing.

Bozorgchenani, A., Tarchi, D., and Corazza, G. E. (2018). Centralized and distributed architectures for energy and delay efficient fog network-based edge computing services. *IEEE Transactions on Green Communications and Networking*, 3(1).

Chren, S., Rossi, B., and Pitner, T. (2016). Smart grids deployments within eu projects: The role of smart meters. In *2016 Smart cities symposium Prague (SCSP)*. IEEE.

Cui, L., Xu, C., Yang, S., Huang, J. Z., Li, J., Wang, X., Ming, Z., and Lu, N. (2018). Joint optimization of energy consumption and latency in mobile edge computing for internet of things. *IEEE Internet of Things Journal*, 6(3).

Dorsemaine, B., Gaulier, J.-P., Wary, J.-P., Kheir, N., and Urien, P. (2015). Internet of things: a definition & taxonomy. In *2015 9th International Conference on Next Generation Mobile Applications, Services and Technologies*. IEEE.

Hasan, M. (2021). State of IoT 2021: Number of connected IoT devices growing 9% to 12.3 billion globally, cellular IoT now surpassing 2 billion.

Hitzler, P. and Van Harmelen, F. (2010). A reasonable semantic web. *Semantic Web*, 1.

Jara, A. J., Olivieri, A. C., Bocchi, Y., Jung, M., Kastner, W., and Skarmeta, A. F. (2014). Semantic web of things: an analysis of the application semantics for the iot moving towards the iot convergence. *International Journal of Web and Grid Services*, 10(2-3).

Kisliuk, B., Tieben, C., Niemann, N., Bröcker, C., Lingemann, K., and Hertzberg, J. (2022). Online Inference of Robot Navigation Parameters from a Semantic Map. In *14th International Conference on Agents and Artificial Intelligence*.

Krishnasamy-Sivaprakasam, G. and Slutzki, G. (2021). Secrecy-preserving reasoning in acyclic dl-liter knowledge bases in the presence of bcqs. In *Proceedings of the 13th International Conference on Agents and Artificial Intelligence - Volume 2: ICAART,*. SciTePress.

Lea, P. (2018). *Internet of Things for Architects: Architecting IoT solutions by implementing sensors, communication infrastructure, edge computing, analytics, and security*. Packt Publishing Ltd.

Lyu, X., Tian, H., Jiang, L., Vinel, A., Maharjan, S., Gjessing, S., and Zhang, Y. (2018). Selective offloading in mobile edge computing for the green internet of things. *IEEE network*, 32(1).

Maarala, A. I., Su, X., and Riekki, J. (2017). Semantic reasoning for context-aware internet of things applications. *IEEE Internet of Things Journal*, 4(2).

Mocnej, J., Miškuf, M., Papcun, P., and Zolotová, I. (2018). Impact of edge computing paradigm on energy consumption in iot. *IFAC-PapersOnLine*, 51(6).

Salama, M. (2020). Green Computing, a contribution to save the environment.

Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1).

Schleicher, J. M., Vögler, M., Dustdar, S., and Inzinger, C. (2016). Enabling a smart city application ecosystem: requirements and architectural aspects. *IEEE Internet Computing*, 20(2).

Wang, S., Wan, J., Li, D., and Liu, C. (2018). Knowledge reasoning with semantic data for real-time data processing in smart factory. *Sensors*, 18(2).

Zhang, K., Leng, S., He, Y., Maharjan, S., and Zhang, Y. (2018). Mobile edge computing and networking for green and low-latency internet of things. *IEEE Communications Magazine*, 56(5).