Elastic Smart Contracts across Multiple Blockchains

Schahram Dustdar
Distributed Systems Group
TU Wien
Austria
dustdar@dsg.tuwien.ac.at

Pablo Fernandez
Applied Software Engineering Group
Universidad de Sevilla
Spain
pablofm@us.es

José María García
Applied Software Engineering Group
Universidad de Sevilla
Spain
josemgarcia@us.es

Antonio Ruiz-Cortés
Applied Software Engineering Group
Universidad de Sevilla
Spain
aruiz@us.es

[Vision short paper]

ABSTRACT

In this paper, we deal with questions related to blockchains in complex Internet of Things (IoT)-based ecosystems. Such ecosystems are typically composed of IoT devices, edge devices, cloud computing software services, as well as people, who are decision makers in scenarios such as smart cities. Many decisions related to analytics can be based on data coming from IoT sensors, software services, and people. However, they typically are based on different levels of abstraction and granularity. This poses a number of challenges when multiple blockchains are used together with smart contracts. This paper proposes to apply our concept of elasticity to smart contracts and thereby enabling analytics in and between multiple blockchains in the context of IoT.

Categories and Subject Descriptors

• Computer systems organization~Architectures

Keywords

Elastic Smart Contracts; Internet of Things; Blockchain; Virtual Chains, Smart Cities.

1. INTRODUCTION

Cities are complex ecosystems, and their effective and efficient functioning has enormous impact on the quality of life of their citizens and society as a whole. However, building smart cities is probably one of the most difficult challenges our society faces today. Among the variety of problems that need to be solved, the question of how to leverage existing ICT technologies to develop foundations for smart city analytics in a transparent and trustworthy form greatly concerns all stakeholders in today's smart cities.

As of today, we have observed several technologies enabling the connection between social and technical subsystems for smarter city analytics. A huge number of Internet of Things (IoT) devices as well as human participation have been introduced to provide various types of data about urban mobility and transportation systems, electricity grid, smart buildings, manufacturing, intelligent logistics systems, and critical infrastructures. Cloud systems have been introduced and used to store and analyze these big "volume, variety, velocity and veracity" streaming things-based and social data through complex middleware for various analytics needed for the operation and optimization of cities. Human capabilities have been invoked in the loop to design and monitor cities together with software. All of these data, analytics capabilities, and domain knowledge in smart cities are involved by a large number of stakeholders, ranging from individual citizens, corporates, to government agencies for both vertical and horizontal problems (such as energy consumption analytics or human mobility analytics). In this view, one needs to understand that analytics of smart cities are far from just "big data analytics" and IoT data analytics. Smart cities analytics have an inherent ecosystem requirement, leading to different paradigm shifts in big data analytics from transactions to ecosystem perspectives as well as in the involvement of multiple, not necessarily trusted stakeholders besides ICT sensors, networks and analytics.

Key city analytics often require data, analytics, and capabilities from both vertical and logical domains (e.g., related to energy consumption) in a complex ecosystem of things, software services, and people with multiple stakeholders, with varying trustworthiness degrees. Complexities in these analytics can be viewed by multiple stakeholders from different angles: **(i)** ***physical (space) view*****:** city analytics can be carried out for a single block, a street, or a house, **(ii)** ***logical domain view***: city analytics are needed for various vertical domains (e.g., building management, intelligent transportation management, and infrastructure maintenance) and horizontal domains (e.g., energy policy and governance, social wellbeing, and urban planning), and **(iii)** ***time view*****:** city analytics can be performed at different time-scales, e.g., online (with near real-time streaming data), offline (with historical data), as well as a combination of both near real-time and historical

Second International Symposium on Foundations and Applications of Blockchain (FAB '19)
April 5, 2019, Los Angeles, California, USA.

data, also considering accountability aspects. While current data gathering techniques are able to collect various types of data, state-of-the-art analytics techniques isolate data produced by technical systems (e.g., from sensors) and social systems (e.g., from people) and then centralize the data in centers (e.g., in clouds) to carry out analytics at centralized places (although utilizing parallel and distributed computing resources). Such approaches rely entirely on software capabilities to deal with big data captured through distributed hierarchical networks of computing elements. In city analytics, data, information, knowledge, and computational capabilities from software services, things, and people are distributed in deep, interwoven distributed ICT architectures. Therefore, state-of-the-art approaches are not adequate as they collect data at the edge of the city where things and people reside, bring the data to the root of the hierarchy (e.g. cloud), and perform analytics based on data provided by predefined settings. First, it does not support time-scale because fine-scale and coarse-scale data analytics are not interoperable, as either we miss a lot of data (in coarse-scale data) or we have to deal with lots of data (in fine-scale data). Second, this also makes the filtering and pre-processing data challenging for supporting complex logical domains, which must deal with different logical horizontal and vertical scales. Finally, we also have severe problems with physical scale: as most of the time we centralize data in one cloud data center so we don't have enough information to cover all physical spaces with sufficient quality to guarantee time-aware analytics, e.g., subjects to be analyzed change rapidly in physical world and we lack up-to-date information in the centralized computing environment.

We believe we need *flexible and elastic mechanisms to support city analytics by harnessing collective capabilities of things, people, and software to carry out timely, quality-aware, and elastic analytics spanning both horizontal and vertical domains*. Given the huge number of things, people, and software services easily to be found and utilized without the need of centralized control, we should investigate a fundamental paradigm shift in utilizing collective capabilities that are distributed across the city infrastructure to enable coordinated analytics in a flexible and elastic manner. Such analytics must be provided with adjustable quality of results for multiple stakeholders where complex, transparent, and trusted collaboration between things, software, and people is needed to understand and address past, current, and future problems of smart cities based on historical, current, and predicted data.

In this paper, we discuss to what extent blockchain technologies are adequate to support complex analytics in these ecosystems. We first introduce a concrete motivating scenario in smart cities analytics (Sec. 2) and analyze how existing approaches to smart contracts and virtual chains can be applied to carry out the relevant analytics (Sec. 3). Our vision, described in Sec. 4, further develops the smart contract notion towards an *elastic smart contract*, which considers elasticity concerns, while providing a framework to horizontally and vertically integrate data and its associated analytics capabilities by promoting the idea of *glue contracts*. In Sec. 5 we conclude that our envisioned proposal will provide a comprehensive support for the different capabilities required in complex scenarios like smart cities.

2. MOTIVATING SCENARIO

As depicted in Figure 1, in this paper we consider smart city infrastructures consisting of (a) IoT sensors, (b) edge devices, which perform computational tasks such as analytics tasks, (c) more "powerful" edge servers (aka fog computing nodes), and (d) cloud computing data centers as the fundamental architectural building blocks for sensing and processing IoT data in smart cities.

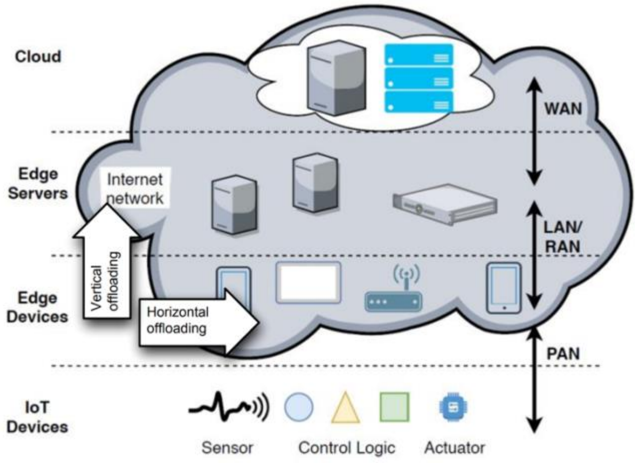

Figure 1. Vertical offloading of analytics computation [1].

At the lowest level of the current smart city infrastructure, we see that data flows from the edges to the data center. From the infrastructure perspective, at the edge (e.g., buildings or districts) we can identify numerous capabilities offered by things, software services, and even people. At (and through) the data centers, several types of software services and people (from the crowds, professional groups, etc.) are available to perform data management and analysis. Although various types of infrastructures connecting people, IoT, and software services are distributed, current city analytics processes are mainly performed in the cloud using software services to provide results to humans. In principle, analytics processes can be carried out in multiple places within the city infrastructure by leveraging the collective capabilities of units of IoT, people, and software services. However, with today's techniques, such units cannot be collectively composed and provisioned on the fly for subsequent distribution throughout the city infrastructure. This prevents us from providing timely and elastic analytics to support non-functional concerns, such as cost, security, and privacy.

For complex problems, city analytics processes are logically divided into a set of sub-analytics processes that cover a set of concerns in distinct horizontal and vertical domains, as shown in Figure 1. Computational tasks can be structured in a “vertical” way or in a “horizontal” way. Given the exemplified city analytics process for policy and regulation of sustainable environments, let us consider an analysis for a city block. Sub-analytics process concerns could be energy consumption of buildings and infrastructure, citizen wellbeing and opinions, environmental impacts of regulations, or incentive policies for green businesses, to name just a few. These sub-analytics processes belong to different vertical and horizontal domains and we need to correlate them and their results in order to understand how to create policy and how to regulate sustainable environments. In principle, such sub-analytics processes are also complex and some of them will be carried out in the cloud (such as, environmental impacts, and incentive policies) whereas others can be performed at the edge where things and people reside (e.g., building energy consumption, and citizen wellbeing and opinions). They also require different algorithms, data, and knowledge from different stakeholders. Among them, there are different ways to exchange analytics results and requests to ensure the final result of the city analytics to be delivered. To the best of our knowledge, state-of-the-art techniques just focus on centralized analytics for single domains. This leads to a severe problem for city analytics: as the scope of current analytics processes is limited to isolated domains and problems are either solved by software services or people, the results may not be adequate and substantial in the overall context of a city. We argue that smart city analytics must be researched from the perspective of ecosystems in which capabilities to contribute to analytics processes are based on hybrid resource types composed of software, people, and things. Moreover, different stakeholders from multiple vertical and horizontal domains impose requirements on analytics processes due to the associated ecosystem of people, technology, and institutions.

Analytics processes in smart city applications can therefore be performed along two dimensions: horizontally, e.g., monitoring and controlling across a number of different domains (and edge or IoT infrastructures) and vertically, e.g., performing analytics processes for particular domains such as buildings in a particular street, as depicted in Figure 2.

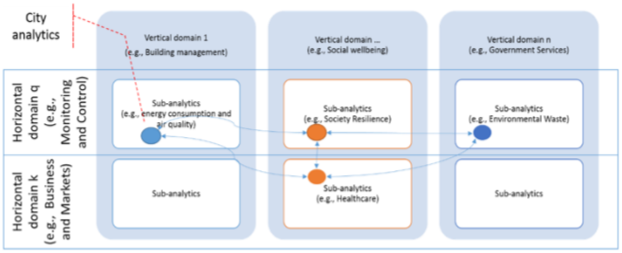

Figure 2. City analytics – logical horizontal and vertical sub-analytics, domains, and stakeholders.

In this scenario it is crucial that the shared data used to perform decentralized analytics in any dimension comes from trusted sources. However, considering the number of agents and stakeholders participating in a smart city ecosystem, trustworthiness cannot be assumed. Furthermore, certain stakeholders, such as public administrations, usually require transparency and tamper resistance to the data they use to analyze and provide services to other agents. For instance, a local administration may enact a contract with an external company to provide street cleaning services, using data from IoT sensors and possibly edge devices located on the streets to plan the optimal cleaning routes. Both the input data and the cleaning routes derived from its analysis should be publicly accessible in a transparent and immutable form, so that the local administration or even citizens can check whether the street cleaning company adheres to the contract in place and the quality level of the provided service, while providing flexibility and adaptability to changes in the ecosystem.

3. RELATED WORK

3.1 Smart Contracts

In a complex scenario like the introduced before, a variety of stakeholders have to collaborate, sharing information between them and allowing each party to carry out analysis and provide decentralized services over the shared data. Trust issues become fundamental in this setting, since parties have to continuously agree on the validity of the data and services they need to integrate. Blockchain technologies are a natural fit, providing transparency and non-tampering to the data shared in a trustless network [2]. In addition to these features, privacy and rights management can be considered by using different blockchain implementations, ranging from permissioned blockchains [3] to specific solutions tailored to IoT-based ecosystems [4].

Since the introduction of smart contracts [5], blockchains have evolved from mere distributed digital ledgers to distributed computing platforms that can include not only an immutable data repository, but also logical and behavioral information to automatically rule the relationships between stakeholders. Thus, smart contracts can encode functionality needed to provide additional services on top of the data registered in the blockchain. These contracts essentially aggregate some data under certain conditions that will trigger its execution. Although the data used within the contract logic is mostly obtained from the blockchain where the contract is deployed, oftentimes there is a need to consider external data (commonly referred as off-chain data). In order to retain the trustless characteristic of blockchains, an additional agent, namely an oracle, needs to provide the external data in a secured, trusted form [6].

3.2 Virtual Chains

Furthermore, there are scenarios where there is a need to separate nodes and information between different levels, as in our motivating scenario (see Sec. 2). Virtual blockchains provide means to implement specific functionality on top of existing blockchains [7]. They introduce an abstraction layer on top of existing blockchains, so that the different application nodes subscribing to the virtual chain will access data and execute smart contracts tailored to their characteristics, while using a single blockchain as the backbone for recording every transaction within the whole system. Thus, multiple virtual blockchains (or virtual chains for short) comprising the different levels discussed in our motivating scenario can be deployed and integrated using this approach. However, sharing data between different virtual chains and from off-chain sources still needs the introduction of oracles, which could be just rights management systems in case of internal oracles allowing data access between virtual chains deployed on the same regular blockchain.

3.3 Elasticity

As the complexity of the systems grows, the need to adapt to variable flows of information and constraints to develop appropriate outcomes represents an important challenge. To this concern, elasticity is presented as the capabilities to react and accommodate changes in the environment with an autonomous mechanism. In [8], authors provide a formal model of elasticity as a three-dimensional space involving resources, quality, and cost aspects that provide the appropriate framework to define and analyze the elasticity properties of an information system that will be used as a starting point of our conceptual proposal.

4. CONCEPTUAL PROPOSAL

Smart contracts represent an appropriate framework to develop a computational mechanism combining data off-chain with the one present in the blockchain. However, in order to address the analytical challenges discussed in the motivational scenario, the framework should be extended to support a variable and multilevel nature of the actors involved. Specifically, in this section, we outline how the elasticity and integration aspects are fundamental cornerstones to build an appropriate smart contract ecosystem to develop more capable blockchains for complex scenarios such as a smart city.

4.1 Integration Concerns

Separating the information needs in different levels allows organizations to focus on their interests, while regulatory bodies can grant access to those organizations only to specific data. In this context, from a blockchain perspective there are multiple architectural alternatives to implement the level stratification, which can be characterized by analyzing three aspects:

- **Granularity**. Several mapping options could be defined to assign a given blockchain to a single level (fine granularity) or multiple levels (coarse granularity). In addition, there could be some scenarios where the same levels are composed of multiple different blockchains

- **Accessibility**. From this perspective, we refer to the capability to analyze the blockchain content by different agents; i.e., the blockchain represents an open system (public) to any agent or a closed system (private or permissioned) to certain agents.

- **Deployment Model.** In this context, we address the logical implementation and deployment of the chain: existing blockchains that are implemented over a specific technical protocol or virtual chains that are materialized inside a regular existing blockchain.

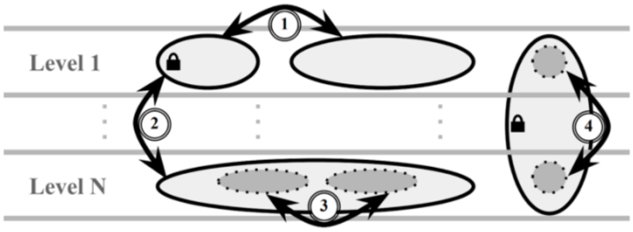

Figure 3. Integration points between blockchains.

Consequently, we can have a wide variety of modelling choices for a given scenario; exemplary, Figure 3 depicts a particular abstract scenario showing several options: level 1 with two fine grain blockchains (one private and one public), level N with one fine grain public blockchain that contains two virtual chains, and a coarse grain private blockchain that spans over all levels and contains a virtual chain for each level. From an analytics perspective, since smart contracts are meant to be executed in the context of a single blockchain, we envision the need for different cross-chain integration mechanisms (as exemplified in Figure 3) depending on three factors: whether integration is done between regular blockchains (Examples labeled with 1 and 2) or virtual chains (3 and 4); between chains in the same level (1 and 3) or different level (2 and 4); or between the same accessibility context (3 and 4) or between a public and a private chain (1 and 2). Taking these challenges into account, we claim the need for a special kind of smart contracts, coined as *glue contracts*, with the special responsibility of making data available across two different chains (virtual or regular) corresponding to the same level (horizontal integration) or different levels (vertical integrations).

In this context, it is important to highlight that integration options presented would represent different types of glue contracts: as an example, in order to integrate two different chains, a possible solution could make use of oracles in order

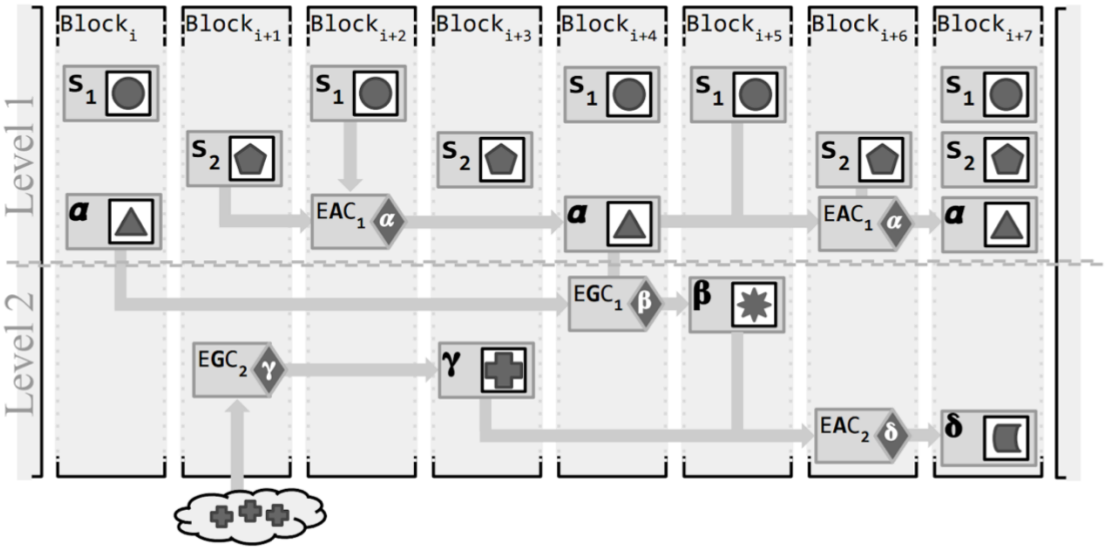

Figure 4. Blockchain fragment example for the use case.

to maintain the trust level of the whole ecosystem; in this particular case, the software oracles are just simple gateways to the accessed blockchains which do not need to add an additional trust method to the already trusted data from the accessed blockchain. Another example of mechanism used by glue contracts to address an integration between accessibility contexts could be the usage of IPFS[1] as the intermediary persistence area for data. In the case of cross-level (or vertical) integrations, glue contracts would be in charge of aggregating the data from inferior levels into new kind of information for higher levels. Furthermore, glue contracts need to address possible divergences between blockchain implementations and protocols of chains to be integrated. There exist alternatives to reconcile these divergences when dealing with crypto currencies [10, 11] that could be extended to allow dealing with complex asset integration.

4.2 Elasticity Concerns

Following the model presented in [8], the envisioned proposal takes into account the elasticity concerns to allow stakeholders to dynamically reconfigure the integration between levels, depending on the horizontal and/or vertical offloading needs (i.e. contract execution), by leveraging elasticity for analytical and glue contracts, correspondingly.

Specifically, in order to incorporate the elasticity dimensions in smart contracts, we need to provide means to elastically define resources, quality properties, and costs associated with a particular contract. To this end, we propose to add an abstraction layer to current smart contracts which will define the elasticity policies for a particular contract. Therefore, executing a so called *elastic smart contract* will transparently consider elasticity aspects on top of the actual functionality provided by the contract. Furthermore, stakeholders should consider executions costs for contracts (e.g. gas for Ethereum smart contracts) as well as infrastructural costs of the blockchains to plan the actual architecture of chains in levels; to this end, a decentralized market of agents [9] would allow the dynamic reconfiguration of the ecosystem taking cost information into account.

4.3 Visionary Use Case

To exemplify the applicability of the proposal we outline a supporting architecture grounded on the current capabilities of blockchain technological state of the art. In such a context, in the current evolution state of the technology towards richer ecosystems, we expect continuous improvements and revisions of the conceptual frameworks presented. In this use case, (Figure 4 shows a fragment of the envisioned blockchain) we can conceptualize an architecture of different virtual chains (composed of “virtual” blocks) that coexist in the same blockchain ecosystem (composed of “grounded” blocks) with smart contract capabilities (such as Ethereum). In such a framework, each grounded block would be a container for multiple virtual blocks that correspond to the different levels and contain either data or contracts related to that level. Specifically, in Figure 4 we exemplify a fragment of the blockchain (Blocks i to i+7)) including two levels (note that in a real scenario there potentially exist a higher number of levels): inside Level 1 we can identify information generated by two agents (s_1 and s_2) and one elastic analytical smart contract (EAC$_1$) in charge of creating derived data from the activity in the level. Next, in Level 2, we can see two kinds of elastic glue contract

---

[1] https://ipfs.io/

(EGC): on the one hand, EGC_1 aggregates the information from Level 1 and incorporates the aggregation as new data in Level 2; on the other hand, EGC_2 (implemented as an oracle) imports data off-chain to the Level 2. Finally, we can see EAC_2 analyze the data of the level to create a new kind of information.

In this context, we can identify different examples of multiple interleaved analytics that can be mapped to the abstract blockchain fragment presented in Figure 4: from low-level analytics regulating small physical spaces that mainly involve sensor data to high-level analytics involving other kind of data sources such as human actor decisions or off-chain census data. For the sake of clarity, we propose a simple example that would correspond with two low levels of analytics representing an adaptable urban lighting system:

Level 1 (street section) would represent a section of a street composed by a number of sensors and lights; concretely in the chain fragment depicted, agents s1, s2 could represent two presence sensors for a given road section that introduce their observations as data in the chain with different time resolution. The analytics contract EAC_1 would periodically perform an analysis over the sensors data to calculate a presence prediction ($\boldsymbol{\alpha}$) in the section; this analytical information would be used to actuate into adaptable street lights in the street section that switch on in the presence of cars, so they dynamically adapt their switch-off latency to the actual prediction.

Level 2 (street) the glue contract EGC_1 could aggregate the presence prediction of different sections calculated in Level 1 in order to create an estimation of the traffic flow in the street ($\boldsymbol{\beta}$); in this level the glue contract EGC_2 could include weather forecast as off-chain data ($\boldsymbol{\gamma}$) so the analytics contract EAC_2 could calculate an estimation of the congestion risk ($\boldsymbol{\delta}$) in order to optimize the traffic lights rules for the given the street.

Furthermore, in a potential superior *Level N* we could leverage advanced use cases such as a new generation contract for waste management service that regulates the actual resource assignment algorithm based on the data harvested by the sensors; this could be implemented by a combination of elastic smart contracts using the analytics gathered and calculating the actual bills automatically having a total transparency and non-tamper management procedure.

Examples of the three elasticity dimensions emerge from our use case: (i) *resources* range from the information providers that can correspond with things (e.g., sensors), software (e.g., government information systems) or people (e.g., an approval from a stakeholder); (ii) depending on the type of resource, a taxonomy of *quality* aspects can be defined (such as resolution data in sensors, availability of the government information system or readiness of the stakeholder); (iii) finally, *costs* involved in the process can also be structured in terms of the resource type (e.g. energy cost of the sensor, infrastructure cost of the information system, or personnel costs). All these concerns would be taken into account to create the elasticity policies for each elastic contract; as an example in the use case, EAC_1 would have a policy to select the number of sensors (resources) filtered by a particular data frequency (quality) and constrained by a maximum number of gas used in the execution of the analytics (cost).

5. CONCLUSIONS

When facing complex scenarios as those that arise in smart cities, where transparency and accountability of the data and analytics are key goals, blockchains are a natural fit. However as these scenarios are typically composed by a complex ecosystem of IoT sensors, edge devices, fog nodes, and cloud data centers, the application of traditional blockchain technologies poses several challenges concerning elasticity and integration aspects, since the requirements for the analytics to be performed varies dynamically, not only in terms of resources needed, quality and cost aspects, but also in the dimensions of those resources. Thus, in order to support elasticity as well as horizontal and vertical integration, in this paper we introduce the concept of elastic and glue smart contracts.

The evolution of current blockchains towards supporting our envisioned elastic smart contracts needs the introduction of elasticity related information to the contracts logic. We propose an elasticity policy abstraction layer to extend the existing smart contracts introducing rules to account for variations in the three elasticity aspects (resources, quality and cost). Additionally, we characterize the different integration scenarios that can be applied to elastic smart contracts, exemplifying them in the context of smart cities. Our vision is that using approaches such as virtual chains and adapting current elastic services frameworks, we can achieve a greater level of integration inside (and between) the various analytical levels while keeping a flexible reconfiguration of the architecture in case there is a need for vertical or horizontal offloading of computation.

6. ACKNOWLEDGMENTS

This work has been partially supported by the EU Commission (FEDER), Spanish and Andalusian R&D&I programmes under grants TIN2015-70560-R and P12-TIC-1867, and the TU Wien Research Cluster Smart CT.

7. REFERENCES

[1] Gusev, M., Kotoska, M., Jakimovski, B., Dustdar, S., Scekic, O., Rausch, T., Nastic, S., Ristov, S., Fahringer, T. (2018). An Architecture for Streaming IoT Applications: A Deviceless Edge Computing Approach. under submission

[2] Swan, M. 2015. *Blockchain: Blueprint for a New Economy*. O'Reilly Media.

[3] Androulaki, E. et al. 2018. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. *Proceedings of the Thirteenth EuroSys Conference on - EuroSys '18* (New York, New York, USA, 2018), 1–15.

[4] Dorri, A., Kanhere, S.S. and Jurdak, R. 2017. Towards an Optimized BlockChain for IoT. *Proceedings of the Second International Conference on Internet-of-Things Design and Implementation - IoTDI '17* (New York, New York, USA, 2017), 173–178.

[5] Ethereum Foundation, 2018. *A Next-Generation Smart Contract and Decentralized Application Platform*. [Online, last accessed: 7-Dec-2018]. Available: https://github.com/ethereum/wiki/wiki/White-Paper

[6] Xu, X., Pautasso, C., Zhu, L., Gramoli, V., Ponomarev, A., Tran, A.B. and Chen, S. 2016. The Blockchain as a Software Connector. *2016 13th Working IEEE/IFIP Conference on Software Architecture (WICSA)* (Apr. 2016), 182–191.

[7] Nelson, J., Ali, M., Shea, R., and Freedman, M. J. 2016. Extending existing blockchains with virtualchain. *Workshop on Distributed Cryptocurrencies and Consensus Ledgers (DCCL'16)*, (Chicago, IL).

[8] Copil G., Moldovan D., Truong H.-L., Dustdar S. (2016). rSYBL: A Framework for Specifying and Controlling Cloud Services Elasticity. *Transactions on Internet Technology*, Volume 16, Issue 3, pp. 18:1 - 18:20

[9] García, J.M., Fernandez, P., Ruiz-Cortés, A., Dustdar, S. and Toro, M. 2017. Edge and Cloud Pricing for the Sharing Economy. *IEEE Internet Computing*. 21, 2 (2017), 78–84.

[10] M. Herlihy. 2018. Atomic cross-chain swaps. arXiv:1801.09515

[11] Hope-Bailie, A. and Thomas, S. 2016. Interledger: Creating a Standard for Payments. *Proceedings of the 25th International Conference Companion on World Wide Web (WWW '16 Companion)*, 281-282.