Arbiter: Improved Smart City Operations through Decentralized Autonomous Organization

Francis Mendoza and Hans Walter Behrens
Arizona State University
Tempe, AZ, USA
{fmendoz7, hwb}@asu.edu

***Abstract*—Smart cities have emerged as one of the most promising applications of cyber-physical systems (CPS), carrying the potential to serve the various interests of the public and private sectors at large. However, contemporary smart city infrastructure commonly uses heavily-centralized network architectures, reducing failure resilience and application flexibility. This centralization also imposes high barriers to entry for public access, limiting usage and oversight opportunities. To address these limitations, we describe Arbiter, a novel fog- and edge-based communication architecture based on the concept of a Decentralized Autonomous Organization (DAO). Arbiter aims to improve the socioeconomic equity of the local citizenry by (1) acting as a management layer for citywide CPS assets, (2) providing a compliance layer for managing human capital, and (3) offering a data protection layer to ensure that citizens retain full control of their personal data. We then analyze in detail the technical, socioeconomic, and ethical implications of Arbiter, and contextualize its role in the modern smart city.**

***Index Terms*—cyber-physical systems, edge computing, Internet of Things, decentralized control**

I. Introduction

Smart cities have long been a topic of interest from both public and private sector entities alike. By collecting, collating, and analyzing enormous quantities of data from pervasive cyber-physical systems (CPS) deployed throughout the municipal area, smart cities offer an enormous opportunity for optimization by enabling reactive, real-time responses to dynamic and evolving stakeholder needs [1]. As these connected assets become more prevalent, smart city capabilities will continue to grow over time.

However, even modern industrialized societies have barely begun to tap into these potential opportunities. A lack of sufficient investment, combined with a dearth of rigorously-tested engineering standards for communication and security [2], have prevented nascent smart cities from reaching maturity. Given the current lack of mass adoption and mature use-cases, research and investment have been primarily led by commercial interests, although academic and public-sector interest continues to grow [2]. Consequently, the developmental trajectory of smart cities currently revolves around the private sector, with a focus on leveraging this influx of data to promote and develop emerging business models at the expense of the general public welfare or individual privacy [3]. Furthermore, current approaches rely primarily on centralized, cloud-based architectures operated by a small number of large enterprise providers [4]. While scalable, such approaches concentrate the control of – and the benefits derived from – public sector data in the hands of a few private sector corporations.

This paradigm of enterprise-driven smart city development contains intrinsic flaws, as it does not accommodate the interests of public sector stakeholders and everyday citizens [2] in the policies, procedures, and constraints imposed on the cyber-physical systems and autonomous networks [5] that underpin smart city operations. While residents may be included in an interactive user experience feedback loop to inform the development of various products and services, their role is cast as passive consumers rather than proactive citizens. Beyond their commercial applications, smart cities also have enormous potential to promote and encourage socioeconomic equity within the public sector. For example, systems could increase resource allocation fairness of public human and technical resources [6], reduce bias in public safety enforcement, or actively ensure compliance by policymakers in abiding by the terms of their office and fulfilling their stated duties.

In order to develop mature smart cities, several new technologies must be effectively implemented to support equitable and decentralized paradigms. Blockchains and distributed ledger technologies can offer a reliable compliance layer [7], while fog- and edge-based computing can leverage ubiquitous but underused computing resources to provide low-latency processing. In conjunction, these technologies form the foundation for decentralized autonomous organizations, or DAOs, which serve as an important building block for public-interest smart cities.

To that end, this work makes the following contributions:

1) We describe a network architecture for linking CPS devices, for cooperative computation and communication.
2) We propose Arbiter, a decentralized autonomous organization template that makes use of this cooperative architecture to provide management, compliance, and privacy features for smart city applications.
3) We assess the social and ethical implications of the proposed technology and discuss its potential impacts.

The remainder of the paper is organized as follows. In Section 2, we provide background. Section 3 introduces the proposed CPS architecture. Section 4 describes the Arbiter DAO. Section 5 provides a technical discussion, Section 6 evaluates socioeconomic impacts, and Section 7 discusses ethical considerations. We conclude in Section 8.

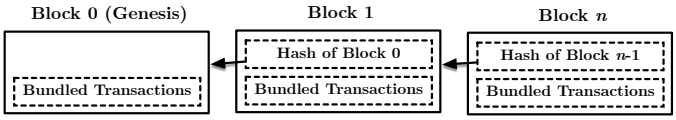

Fig. 1. A traditional blockchain composed of n blocks. Each block contains a set of transactions and a hash of the previous block. Blocks are linked iteratively; falsifying a block requires retroactively changing all prior blocks.

II. Background

A. Blockchains and Distributed Ledger Technology

Blockchains are a relatively recent development in decentralized data storage and operate by shifting trust away from a centralized authority to a set of commonly agreed-upon protocol standards [7]. A subset of distributed ledger technology, blockchains' primary value lies in the enforcement of data integrity, authenticity, and nonrepudiation even when shared between mutually-distrustful parties. Several variants exist, which are associated with various tradeoffs based on the design and application criteria involved. In general, all participants (or *peers*) in the network maintain their own copy of the ledger, which stores the transaction history of the entire network since its initial creation. These transactions are combined into larger collections, called *blocks*, which are protected from tampering through the use of cryptographic hash functions (*hashing*). The iterative process of hashing blocks and appending them to the ledger allows honest peers to detect malicious or improper transactions by cross-referencing new blocks with their local copy of the historical record. Once a sufficient number of peers agree a block is valid, it is added to the ledger – this is referred to as reaching *consensus* on a block [7].

Conceptually, blockchains fall into two categories: (1) public blockchains, which anyone may join or observe, and (2) private blockchains, in which peers must be vetted and approved before joining. The well-known Bitcoin network [8] exemplifies a public blockchain, operating on the assumption that all peers hold equal weight when coming to consensus in the network. Because all nodes have an equal say, this conceptually resembles an electronic version of direct democracy and parallels the equality exemplified by the democratic participation of citizens in civic life. Similarly, peers' anonymity is preserved from oversight, preventing misuse of this data by powerful third parties. Permissioned blockchains, in contrast, require the identities of participating peers to manage the access control layer which governs the list of valid participants [7]. This allows for better network security, faster performance, and more fine-grained control, but sacrifices privacy. This paradigm maps effectively to granting access to sensitive assets, such as to subject matter experts or authorized officials, for whom accountability is desirable.

One final aspect of blockchains is the concept of a smart contract – this describes a piece of computer code which has been published to the ledger, and which waits for predefined conditions to occur before automatically executing permutations to the ledger. These smart contracts can be used to encode and express application logic for the services which must be provided to relevant stakeholders.

B. Decentralized Autonomous Organizations

Decentralized autonomous organizations (or DAOs) are collections of smart contracts that collectively describe a management model for a decentralized organization, instantiated on a public blockchain. A DAO observes its surroundings, and reacts or takes actions according to the behaviors encoded in its contracts. Thus, DAOs can serve as effectively functioning companies, organizations, or governments [9] despite their decentralized and automated nature. Since all rules and governance are encoded on the ledger, revisions are possible using a consensus procedure similar to that used for transactions.

Upon creation, a DAO is initially funded with a valuable utility or currency to use as an enforcement mechanism for its own rules. Thus, third-party peers may enter as stakeholders into a DAO in exchange for benefits such as voting rights, control over particular assets, or limited governance rights, for example. After funding is complete, the DAO becomes active, and operates according to its specifications to dispense rewards or mete out punishment based on its founding principles and input from its stakeholders, ultimately operating independently from its creators. All actions undertaken by a DAO are recorded to its ledger, which makes them both immutable and publicly auditable, available for any peer in the network (not just stakeholders) to review. While many existing DAOs are based on cryptocurrency-based ledgers, new types of DAOs have been proposed [9], [10] to use cyber-physical systems (CPS) and their ledger-based "digital twins" [11] to interface with and control real-world assets.

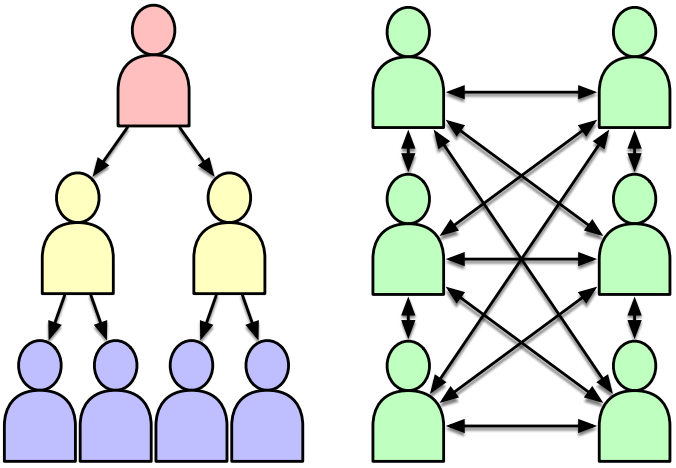

Fig. 2. A hierarchical organization compared to a decentralized organization. Peers communicate directly as required to complete their tasks, rather than relying on top-down centralization to distribute responsibilities.

C. Fog and Edge Computation

Contemporary "cloud computing" describes powerful, Internet-accessible servers that are widely accessible from anywhere despite being typically located inside only a few datacenters. In contrast, fog and edge computing leverage

devices located near to users, even in the same building, to perform the same tasks while dramatically reducing the distance communication must travel. The Internet of Things, which describes network-connected, low-power cyber-physical devices, is projected to grow to 20 billion by the end of 2020 [12]. The scale of these devices is rapidly outpacing the cloud computing infrastructure on which they often rely [13]. At the current rate, existing infrastructure may be overwhelmed by requests from IoT devices, failing to supply sufficient computing, storage, or bandwidth to meet demand.

Although traditional scaling techniques have for the past decades relied on improving the speed of computing cores, Moore's Law is now showing diminishing returns [4]. However, cost per core continues to decrease, encouraging their addition in new infrastructure and scaling out, rather than up. This has led to a large number of computationally-capable devices entering deployment and engendered the conditions for a ubiquitous, low-latency compute infrastructure.

Edge and fog networks take advantage of this abundant but latent source of computing resources [4]; however, they are not designed to completely replace cloud computing. Rather, they enable load-shifting away from overburdened centralized servers. Edge and fog (collectively, *edgefog*) networks are distinct yet complementary concepts. Edge computing refers to computation done directly by low-powered devices furthest from the centralized server (they are "at the edge"), and are highly heterogeneous and extremely numerous. Due to their limited individual capability, results are relayed "upward" towards more-central nodes to allow for aggregation and post-processing. Fog networks typically refer to aggregate collections of edge computing, serving as "localized clouds" to handle many processing tasks locally, while still able to relay some tasks back towards central cloud servers.

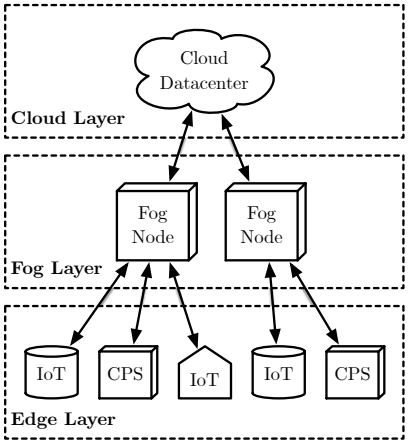

Fig. 3. A simplified network architecture illustrating relationships between the core cloud layer, a localized fog layer, and heterogeneous devices located at the edge. Message traffic "ascends" layers only when necessary.

III. Arbiter Network Architecture

We first describe the CPS layer underpinning our approach. Although the Arbiter has partial operational utility without CPS integration, such as being able to manage and allocate capital investment according to initiatives approves by the public and subject matter experts, it becomes far more useful when paired with public sector smart assets drawn from an existing CPS infrastructure.

A. Edge Network Layer

Infrastructure and enterprise IoT devices reside on the edge layer, comprising the mechanisms for enabling automated public sector services. Potential applications include autonomous transport, smart utilities, intelligent supply chains for example [2]. Existing legacy infrastructure can be retrofitted using "smart kits" to improve utility and reduce costs by obviating the need for replacement. Industrial IoT products typically specify a lifecycle of at least 10 years, but will also require maintenance [1]. Consequently, predictive maintenance based on observed data improves deployment reliability while also stimulating the local economy through the creation of demand for skilled technicians. As automation increases, reskilling programs can preserve career advancement for positions that would otherwise be at risk. Network extensions and improvements are governed by data-driven decision making, under the supervision of local citizenry.

B. Fog Network Layer

The fog network is a natural extension of core cloud networks at a more localized level. By leveraging local processing and forwarding only what is necessary, these fog networks allow cloud computing providers to effectively scale with rapidly increasing CPS-driven demand. In parallel, electronics manufacturers benefit from the maturation of the nascent fog network market, driving new sales. Likewise, the creation of local fog datacenters stimulates the local economy and fosters demand for skilled local labor. Although a mature fog network implementation has yet to appear, current efforts are focused on the sale and deployment of specialized equipment at carefully-chosen locations within the broader IoT network.

C. Cloud Network Layer

While edgefog networks play a critical role in alleviating demand on centralized servers, sensitive data or overly-demanding computations may still need to be relegated to the cloud. In particular, the cloud can play an important role in connecting independent networks between metropolitan areas, allowing for collaboration and cooperation across wide geographic areas despite gaps in CPS infrastructure.

IV. Arbiter DAO

Blockchain technology in the form of a DAO, operating on the hybrid network architecture previously described, forms the core of the Arbiter system. Arbiter acts as:

1) A management layer: To marshal public resources, city-wide CPS assets, human resources, and financial capital towards citizen-driven socioeconomic equity proposals.
2) A compliance layer: To ensure that public servants and enterprises operating within the metropolitan network comply with pre-established rules, and enforcing those

rules through the issuance or seizure of assets and resources.

3) A data protection layer: To prevent the unethical collection and usage of citizen data by unauthorized or malicious parties without their consent, by giving citizens full control over their data and its release.

To accomplish these goals, the Arbiter DAO replies on a smart contract layer to allow governance and behavioral encoding, paired with CPS represented by digital twins in the network. By leveraging the ubiquitous computing capabilities of the edgefog network to protect transmission and storage of data, and to reliably service stakeholders and citizens in real-time, Arbiter encourages improvements in socioeconomic equity through decentralized, collective decision making.

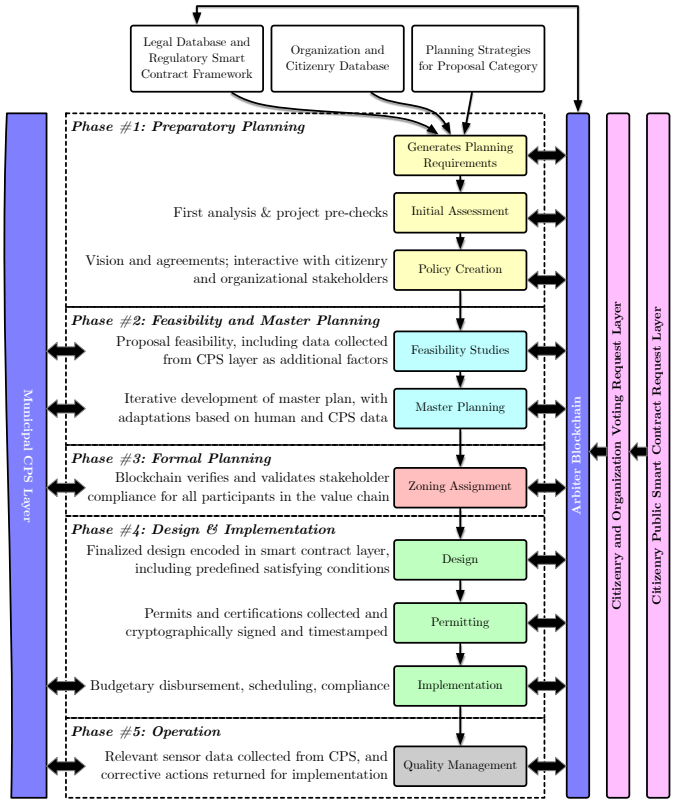

Fig. 4. Using Arbiter to empower the urban planning process. Here, it marshals resources and talent for public works projects as proposed by the citizenry, based on municipal data and CPS infrastructure.

A. Management Layer

The first core goal of Arbiter lies in effectively marshaling public sector resources, such as smart assets from the citywide CPS infrastructure or available human resources, in satisfying predetermined outcomes as voted on by the citizenry and relevant subject matter expert boards. To highlight the role of Arbiter's management layer, consider an urban planning process as a motivational scenario. In Figure 4, we see a flow diagram of the decision making process extended to include the Arbiter, including its blockchain layer, smart contract layer, the CPS assets it interacts with, and ancillary databases it may need to access to accomplish a public works task – in this case, a new park.

Arbiter takes into account available stablecoin cryptocurrency assets (a digital representation mirroring concrete financial resources), the assets that are needed for the development and operation of the project, requests represented by smart contracts from citizens and relevant stakeholders, and the applicable regulatory framework. Requests to update the properties of a particular object are submitted to the voting layer, where citizens or subject matter experts can determine the relative priorities of the encoded issues, and decide which proposals will receive resources, and how much should be allocated.

This voting layer is enabled by a *federated* blockchain, which combines positive aspects of both public and permissioned blockchains. Namely, it is operated and controlled by an individual city, allowing different municipalities to reflect the needs of their citizens; this necessitates a level of control unique to permissioned blockchain systems. This also permits the credentialing of subject matter experts to form committees that address topics that require specialized knowledge to make valid determinations. However, by preserving data accessibility to residents, the transparency and auditability of public blockchains are preserved.

B. Compliance Layer

Continuing our motivating example, we now address the compliance layer offered by Arbiter. This consists of several steps, which collectively protect the human network of citizens and stakeholders (e.g., enterprises or other organizations from the public and private sector).

The first phase occurs during the smart contract proposal voting layer, by ensuring an accurate accounting of votes on deciding a particular issue. Due to the federated, public/permissioned model adopted by Arbiter, each vote must be associated with particular allocated voting rights. This allocation permits high-fidelity validation of vote integrity, including evaluation of permission to vote on particular issues, rights to propose new projects, or possession of relevant expertise. Implicitly, votes may not be cast without correct permissions to do so, validating the integrity of the vote before it has even begun.

The second phase occurs once voting has concluded, during the execution of the proposals themselves. To participate in a proposal, e.g. to be considered for a construction contract, a stakeholder must have provided digital collateral to secure their participation. Peers or stakeholders who fail to comply with the terms of the agreement, or who attempt to subvert the integrity of the voting process itself, will forfeit their initial stake.

Note that if necessary, stakes could be forfeit for all failing proposals; this would drastically reduce their frequency and increase network scalability, but would also strongly disincentivize risk-taking or minority-favored proposals. Alternatively, a hybrid approach could be used, where low-risk, low-reward proposals would require low or no fees, while significant

allocations of civic capital might require a larger risk to penalize misbehavior. This could be useful, for example, to discourage lowball bids designed to secure a proposal, but which are unrealistic in their scope. By failing to satisfy their initial bid, unscrupulous stakeholders would lose not only the assets associated with the proposal itself but also their initial invested stake as well.

C. Data Protection Layer

Although smart contracts play a role in Arbiter's data protection layer, it benefits from the network architecture as well. By leveraging an edgefog approach, the amount of data sent, and the distance it must be sent, are both reduced. Citizen-users remain in full control of their data due to the peer-to-peer (P2P) architecture used for blockchain transaction propagation, which also applies to smart contract execution [7]. Since the route data will take is nondeterministic, and intermediate nodes are untrusted, intercepting messages becomes extremely difficult or impossible for a malicious actor, much less decrypting protected data.

Smart contracts can further reduce data leakage, enforcing strict data protections, and stringent access permissions without depriving citizens of the utility of public services. By replacing the traditionally-hierarchical structures of municipal services with DAO-driven transaction processing logic, only the minimal subset of data needed to request service must be provided. By automating the provision of services, alternative providers may also offer competitive proposals for citizen-users to choose, giving additional choice to users over how their data is used, and by whom.

V. Technical Discussion

While Arbiter offers promising opportunities for empowering the smart cities of tomorrow, there are additional technical considerations that must be noted prior to deployment. Here, we highlight the advantages and disadvantages of Arbiter.

A. Arbiter Strengths

Given the current state of development of smart cities, biased towards exclusive reliance on centralized cloud networks and the data silos inherent to such architectures, the edgefog network architecture proposed for Arbiter provides superior scalability for the many billions of IoT devices expected to be deployed in the coming years.

Using a DAO to encode automated behaviors and rule enforcement uses well-understood cryptographic principles to ensure safety and security thanks to its blockchain underpinnings. Furthermore, the use of permissioned systems allows for compartmentalization of sensitive data, which may be necessary for some proposals and allows for weighted votes should citizens decide to delegate voting power to subject matter experts on some topics.

B. Arbiter Weaknesses

The use of blockchain, while secure, nevertheless brings several undesirable inefficiencies. For example, each participant is assumed to keep a copy of the transaction ledger, which for an extremely dynamic system like Arbiter will quickly become impractical for the low-powered devices acting as peers. While partially mitigated by permissioned blockchain systems, this challenge highlights a need for long-term storage of voluminous data streams. Since Arbiter specializes in the optimization of available resources, analysis of existing data is a critical aspect, and that data must be stored before it can be assessed. Sharded or geospatially-partitioned approaches could mitigate this drawback by limiting the amount of information stored by each node.

Additionally, the use of edgefog networks and blockchains exposes the network to potential risks from adversarial interference. More computational power can be used to execute so-called "proof of work" algorithms to counteract this behavior, but again, low-powered edge devices may not be able to allocate sufficient resources to effectively participate in this way. Alternative proof mechanisms are in active development to address this concern and make adversarially-resilient systems more computationally tractable.

Finally, DAOs represent extremely complex collections of smart contracts. If flawed contracts are introduced to the system, it can undermine its fundamental operation. Indeed, such a flaw brought down the eponymous DAO [14], the first decentralized autonomous organization ever created. Improvements to simplify the creation and secure the execution of smart contracts are therefore necessary to reinforce the technical foundations of Arbiter.

VI. Socioeconomic Impact

By promoting and enforcing socioeconomic equity through intelligent use of citywide CPS assets and the automation of compliance and management activities at scale, Arbiter has a drastic effect on social and economic factors. It can speed transaction settlement times, prevent fraud, and reduce friction between citizens, stakeholders, and institutions while preserving the broader social contract. In addition, by enabling new autonomous business models and streamlining public services, it allows for synergy with the developing IoT service economy as well. It also promotes increased demand for engineers, technicians, and business leaders for its service and expansion, fostering local job growth. It also opens new fields of work, severely reduces overhead costs, and increases quality of life across the community.

By reallocating decision-making power away from powerful yet corruptible intermediates and towards a collectively-driven set of objectively-enforced rules, Arbiter can further decrease incidences of fraud and corruption. In a virtuous cycle, decreases in the misallocation of resources further build public trust and participation in the system. Thus, collective and collaborative behaviors become more rewarding, and zero-sum competitive behavior is disincentivized, leading to further increases in participant satisfaction.

VII. Ethical Considerations

There are several ethical considerations to consider prior to a prospective large-scale deployment of Arbiter, especially with

regards to the scale of automation, its breadth of integration with existing CPS infrastructure, and the privacy of its data.

Due to its disruption of existing management and allocation paradigms, as well as its control over financial, material, and service-oriented assets, it is reasonable to expect extensive pushback from entrenched interests. We hope that the objective efficiency and inherent fairness of Arbiter will lead to grassroots support, but whether such tactics would be effective against organized lobbying efforts remains to be seen.

Further, given the breadth of Arbiter's integration, its misuse could easily result in dystopian outcomes if sufficiently numerous malicious actors colluded to seize majority control of the network. The extent to which this technology should therefore be embedded into our daily lives, and its appropriateness to the role, has yet to be determined.

Finally, while data privacy issues can be mitigated through alternative P2P messaging infrastructures [15], [16], it is important to differentiate which messages should remain private and which should be made public, and to determine which channels are amendable to decentralization and which should remain centralized.

Beyond these aspects, the fundamental issue of collective decision making lies in ensuring it does not fall prey to the so-called tyranny of the majority, in which a majority of participants enforce their wishes to the detriment of the minority. Especially when rules are enforced dispassionately and objectively via well-defined algorithms, this question becomes even more important, since human compassion manifests only collectively at one remove from the voting members.

VIII. Conclusion

We have proposed Arbiter, a novel civic collaboration framework for smart cities that improves socioeconomic equity by leveraging decentralized autonomous organization principles for collective decision making. By integrating with existing cyber-physical systems infrastructure and ubiquitous edgefog networks, it offers management, compliance, and data privacy functionality in addition to managing and optimizing public services. Stakeholders and citizens may offer proposals, vote on issues relevant to them, and allocate public resources fairly, secure in the confidence that their wishes are enforced fairly and transparently.

References

[1] K. Su, J. Li, and H. Fu, "Smart city and the applications," in *2011 International Conference on Electronics, Communications and Control (ICECC)*, Sep. 2011, pp. 1028–1031.

[2] L. Carvalho, "Smart cities from scratch? A socio-technical perspective," *Cambridge Journal of Regions, Economy and Society*, vol. 8, no. 1, pp. 43–60, Mar. 2015.

[3] R. Kitchin, "Getting smarter about smart cities: Improving data privacy and data security," in *Data Protection Unit, Department of the Taoiseach*, Dublin, Ireland, 2016.

[4] T. Taleb, S. Dutta, A. Ksentini, M. Iqbal, and H. Flinck, "Mobile Edge Computing Potential in Making Cities Smarter," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 38–43, Mar. 2017.

[5] J. Jin, J. Gubbi, S. Marusic, and M. Palaniswami, "An Information Framework for Creating a Smart City Through Internet of Things," *IEEE Internet of Things Journal*, vol. 1, no. 2, pp. 112–121, Apr. 2014.

[6] J. Ojasalo and L. Tähtinen, "Integrating Open Innovation Platforms in Public Sector Decision Making: Empirical Results from Smart City Research," *Technology Innovation Management Review*, vol. 6, no. 12, pp. 38–48, 2016.

[7] M. Swan, *Blockchain: Blueprint for a New Economy*. "O'Reilly Media, Inc.", Jan. 2015.

[8] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," Tech. Rep., 2008.

[9] N. Diallo, W. Shi, L. Xu, Z. Gao, L. Chen, Y. Lu, N. Shah, L. Carranco, T.-C. Le, A. B. Surez, and G. Turner, "eGov-DAO: A Better Government using Blockchain based Decentralized Autonomous Organization," in *2018 International Conference on eDemocracy eGovernment (ICEDEG)*, Apr. 2018, pp. 166–171.

[10] K. Christidis and M. Devetsikiotis, "Blockchains and Smart Contracts for the Internet of Things," *IEEE Access*, vol. 4, pp. 2292–2303, 2016.

[11] K. Biswas and V. Muthukkumarasamy, "Securing Smart Cities Using Blockchain Technology," in *2016 IEEE 18th International Conference on High Performance Computing and Communications; IEEE 14th International Conference on Smart City; IEEE 2nd International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, Dec. 2016, pp. 1392–1393.

[12] R. Almadhoun, M. Kadadha, M. Alhemeiri, M. Alshehhi, and K. Salah, "A User Authentication Scheme of IoT Devices using Blockchain-Enabled Fog Nodes," in *2018 IEEE/ACS 15th International Conference on Computer Systems and Applications (AICCSA)*, Oct. 2018, pp. 1–8.

[13] H. El-Sayed, S. Sankar, M. Prasad, D. Puthal, A. Gupta, M. Mohanty, and C.-T. Lin, "Edge of Things: The Big Picture on the Integration of Edge, IoT and the Cloud in a Distributed Computing Environment," *IEEE Access*, vol. 6, pp. 1706–1717, 2018.

[14] M. I. Mehar, C. L. Shier, A. Giambattista, E. Gong, G. Fletcher, R. Sanayhie, H. M. Kim, and M. Laskowski, "Understanding a Revolutionary and Flawed Grand Experiment in Blockchain: The DAO Attack," *Journal of Cases on Information Technology (JCIT)*, vol. 21, no. 1, pp. 19–32, Jan. 2019.

[15] H. W. Behrens and K. S. Candan, "Adversarially-Resistant On-Demand Topic Channels for Wireless Sensor Networks," in *2018 IEEE 37th Symposium on Reliable Distributed Systems (SRDS)*, Oct. 2018, pp. 83–92.

[16] ——, "Pando: Efficient Byzantine-Tolerant Distributed Sensor Fusion using Forest Ensembles," in *ICC 2020 - 2020 IEEE International Conference on Communications (ICC)*, Jun. 2020, pp. 1–6.