Taming the Airwaves: Resilient Spectrum Sharing with BFT

Xiaoxuan Qin, Prashant Krishnamurthy, Amy Babay
Department of Informatics and Networked Systems, University of Pittsburgh

***Abstract*—Dynamic spectrum sharing between incumbents and new entrants increasingly requires near-real-time, partially decentralized coordination to balance critical and lower-priority uses. This makes it important to reason about the robustness of spectrum sharing systems to malicious behavior.**

While the dependability literature includes extensive work on robust coordination under malicious behavior, including Byzantine Agreement, Byzantine Fault Tolerant State Machine Replication, and fault-tolerant sensor averaging protocols, effectively applying these techniques requires a detailed understanding of application requirements. At the same time, application design decisions can have important fault tolerance implications. In this paper, we analyze the requirements of dynamic spectrum sharing, highlighting areas where Byzantine Fault Tolerance techniques can be useful and identifying open research questions.

I. Introduction

Dynamic spectrum sharing based on near-real-time coordination is increasingly needed to address growing radio spectrum demands [1]. The objective of sharing is to protect incumbent users from interference while allowing new entrants to opportunistically exploit idle spectrum across time and space. Critical use cases (typically federal incumbents) and lower-priority users must coexist. An example is the Citizens Broadband Radio Service (CBRS), where US Navy radars are protected while private 5G networks are deployed alongside them. Coordination systems like the CBRS Spectrum Access System (SAS) are deployed by multiple administrators [2] in a decentralized manner for fault tolerance, with no obvious single centralized authority. Any such architecture must also be robust against security breaches and malicious participants [1]. Today, sharing schemes like CBRS remain largely static over long time scales. A resilient architecture could *fundamentally disrupt both how spectrum is allocated and what applications it enables*, by enabling robust sharing of radio spectrum at far more granular time scales.

We examine dynamic spectrum sharing as an emerging disruptive application suited to Byzantine Fault Tolerant (BFT) techniques, while identifying mismatches between spectrum sharing requirements and standard approaches like BFT State Machine Replication (SMR). We propose a research agenda covering the types of misbehavior feasibly tolerable in this setting and the design of BFT spectrum management solutions.

II. Related Work

A. Spectrum Sharing

Spectrum has been largely allocated for "uses" such as data transfer (both terrestrial and satellite, fixed and mobile), radiolocation (for example GPS), and sensing (both active and passive). Such uses are separated in space and frequency to avoid interference. Increasing demand for spectrum in these and other uses has made it necessary to share spectrum both in *space* and *time* in innovative ways. Typically, such sharing involves coordination in space, on various time scales, between incumbents (which we refer to as Type 1 users) and new entrants (which we refer to as Type 2 users); see Figure 1. One example is CBRS, which has three tiers: the incumbent navy radars, licensed allocation, and unlicensed use with registration. Recently, coordination of spectrum between different assignments is being considered in the 37 GHz bands [3]. Decentralized coordination using automated frequency coordination (AFC) [4] is required in the 6 GHz bands (with WiFi) when access points use standard power, so that interference to incumbent microwave links is limited. As more previously allocated and assigned spectrum is opened up by regulatory authorities for sharing, coordination between incumbents and new entrants becomes inevitable.

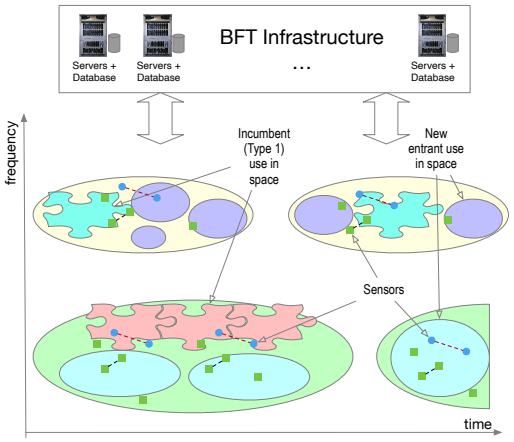

Fig. 1: Big picture view of a BFT spectrum sharing architecture. The x-axis shows time, the y-axis different frequencies. Incumbent users have unpredictable behavior (shown as irregular shapes) while new entrants are controlled by the BFT system with the assistance of sensors that detect incumbent transmissions and may be deployed by different organizations (e.g. blue circles vs green squares).

B. BFT Spectrum Sensing

Byzantine behavior has been investigated in spectrum sensing for cognitive radio (surveyed in [5]). In this case, the typical model is that cognitive radios (CRs) report sensing data to cooperatively determine whether a primary (Type 1) user is

active, and each CR has some probability of being Byzantine. Existing work develops data fusion schemes that (with high probability) accurately detect primary user presence/absence in this environment. Classical fault-tolerant sensor averaging [6] and its privacy-preserving variants [7] tackle a similar problem in a threshold-based rather than probabilistic model, where a fixed number of sensors may be Byzantine.

While sensing is an important component of dynamic spectrum management, this line of work does not provide solutions for how to coordinate spectrum allocation. One exception is [8], which considers joint sensing and allocation, but it requires a specific deployment model where Type 2 users are responsible for sensing and can be penalized for misreporting.

C. Blockchains for Spectrum Management

Prior work has identified blockchains a tool for spectrum management [9], which has led to a large number of blockchain-based spectrum management systems being proposed (e.g. [10]–[16], surveyed in [17], [18]).

While blockchains often (but not always) use Byzantine Fault Tolerant consensus to provide robustness to malicious behavior, understanding the guarantees of the overall system requires a clear statement of the threat model and properties needed from the blockchain infrastructure, which is often lacking or left implicit in this work. Moreover, this work typically focuses only on coordination and does not consider the problem of spectrum sensing. In this paper, we aim to lay a foundation for the development of comprehensive spectrum sharing systems that provide resilient solutions for both sensing and coordination with well defined guarantees.

III. Defining the Dynamic Spectrum Sharing Problem

At a high level, the goal of dynamic spectrum sharing is to maximize spectrum use while protecting incumbent (Type 1) users from interference. To increase spectrum utilization, new entrants (Type 2 users) may transmit in idle spectrum slices that are not currently occupied by a Type 1 user. To avoid interference to Type 1 users, Type 2 users must evacuate the slice (stop transmitting) when a Type 1 transmission begins.

To be more precise, we can define the application-oriented correctness criteria as follows:

Definition 1 (Safety-T1-Protection)**.** *If a Type 1 user u transmits with range r in a spectrum slice S from time T_1 to T_2, no other user u' with transmission range r' located within distance $r + r'$ of u transmits in S from time $T_1 + \epsilon$ to T_2.*

Definition 2 (Safety-T2-Protection)**.** *If a Type 2 user u with transmission range r has been allocated a spectrum slice S from time T_1 to T_2, no other Type 2 user u' with transmission range r' located within distance $r + r'$ of u transmits in S from time T_1 to T_2.*

Definition 3 (Liveness)**.** *If a Type 2 user u with transmission range r submits a request to access spectrum slice S from time T_1 to T_2, either u's request is granted, or some other user u' with transmission range r' located within distance $r + r'$ of u transmits in slice S at some time in the interval $[T_1, T_2]$.*

Note that our definition of Safety-T1-Protection allows for a small amount of time ϵ in which a Type 1 user may experience interference, representing the time between a Type 1 transmission starting and any interfering Type 2 users evacuating the slice. We further note that Safety-T1-Protection and Safety-T2-Protection can trivially be satisfied by simply not granting any T2 requests: the current situation, where allocations are largely static, satisfies *only* Safety-T1-Protection. Other wireless-specific issues are discussed in Section V.

From a system design perspective, there are three core challenges in providing these properties: (1) detecting Type 1 transmissions, (2) coordinating spectrum access, and (3) enforcing allocation decisions. Understanding what types of solutions are feasible for these challenges strongly depends on the specific properties of the application.

IV. Mapping the Solution Space

A. Detecting Type 1 Transmissions

To avoid interference with Type 1 users, any system for coordinating spectrum access must first be able to *detect* transmission by Type 1 users. There are two primary approaches to detecting Type 1 transmissions: *sensing* and *reporting*.

The Environmental Sensing Capability (ESC) in CBRS is an example of a sensing approach, where sensors are deployed in specific areas (particularly along the coast) to detect incumbent US Navy transmissions. Sensors report detected transmissions to the SAS coordinating spectrum access. However, deploying and managing sensor infrastructure can become expensive. The costs depend on the spectrum band and incumbent characteristics discussed in Section V, as well as the fault/threat model assumed for sensors. For example, if incumbents are located in a specific limited geographic area (like Navy users in CBRS), deploying sensors to cover that area is much easier than if incumbent users are distributed throughout a larger area where Type 2 users are also operating. Similarly, the number of sensors needed to cover a geographic area also depends on propagation characteristics and the sensitivity of the sensors.

From a fault tolerance perspective, the required sensing density also depends on the fault/threat model: if we assume a number of sensors may be subject to Byzantine faults and invent false transmissions, then we need to require multiple sensors to corroborate reports of the same transmission in order to accept it as valid, which requires that each point in the geographic space be redundantly covered to ensure enough reports can be collected. Moreover, since Byzantine sensors may also fail to report transmissions, we need to ensure that the sensing density is high enough that even if some sensors fail to report, enough redundant reports can be collected to determine that a transmission is active. Thus, one open research question is: **How can we optimize sensor deployments in order to provide the necessary level of protection for Type 1 transmissions while maximizing availability for Type 2 use and minimizing sensing cost?**

An alternative to sensing is to require Type 1 users to report their transmissions. While this eliminates the need for potentially expensive sensing infrastructure, its feasibility depends on the characteristics of incumbent applications. Incumbents may be unwilling to report their transmissions due to national security (in the case of the Navy) or commercial interests. Thus, solutions that allow for reporting transmissions in an obfuscated manner that supports well defined confidentiality guarantees for users are likely to be needed to make such reporting practical. An open question is: **How can we design reporting subsystems that provide enough information about Type 1 users' transmissions to support effective interference protection and substantial Type 2 use, while meeting Type 1 users' privacy requirements?**

B. Coordinating Access

Beyond detecting Type 1 transmissions, dynamic spectrum sharing requires a system to coordinate access, where Type 2 users can request and receive authorization (spectrum grants) to use idle spectrum. This coordination service is essentially a database tracking ongoing transmissions based on reports of Type 1 activity (either via sensors or Type 1 users themselves) and requests from Type 2 users. Such coordination systems are being implemented, often in a partially decentralized manner where multiple companies are responsible for running the service (e.g. CBRS Spectrum Access System). In this environment, where servers may be deployed by competing commercial entities, we would like to build a system that is robust to incorrect behavior of a subset of coordination servers, due to security breaches resulting in compromises, out-of-spec implementations, or accidental or malicious misconfiguration.

While the dependability literature provides many solutions for robust coordination under malicious behavior, significant work is needed to assess their applicability in this context and adapt them to specific application needs. For example, BFT SMR offers strong consistency guarantees in the presence of malicious behavior and seems like a promising approach to ensure servers make consistent decisions about spectrum grant requests (see Figure 1). Assuming a sensor-based detection approach, sensors and Type 2 users act as clients, and servers implement a database replicated using BFT SMR. At a high-level, this can work as follows: Servers accept *requests* from Type 2 users and *reports* from sensors and order these requests and reports using a BFT SMR protocol. Servers execute requests and reports sequentially in the order determined by the BFT SMR protocol. Upon executing a Type 2 request, a server checks whether it interferes with any ongoing transmission; if so, it is denied, otherwise it is granted. Upon executing a sensor report, a server checks whether the reported transmission interferes with any ongoing Type 2 transmission; if so, the interfering Type 2 transmission is revoked.

While this seems straightforward to implement, in fact, this client-server interaction pattern does not match the typical assumptions of BFT SMR. The standard assumption for BFT SMR is that each client submits requests to the system and receives replies to its own requests. Moreover, basic BFT SMR protocols often assume each client is limited to a single outstanding request at a time (or do not guarantee per-client FIFO ordering if a client issues multiple simultaneous requests). In the proposed spectrum sharing implementation, sensors submit reports but do not need to receive replies. Sensors should also never be blocked from submitting new reports, since the system's primary concern is the current state of ongoing transmissions, not the complete history of transmission start and end events. In this sense, sensor reports have "overtaken-by-event" semantics: if a sensor is temporarily disconnected such that its reports are not executed, upon re-connection, we want to execute reports of *ongoing* detected transmissions as quickly as possible, rather than reliably executing all reported events during the disconnection.

For Type 2 users, revocations as described above represent a *spontaneous reply* [19]. The Upright system suggests that spontaneous replies can be sent over an *unreliable channel* and application-specific logic can handle the case where message loss is detected via a sequence number gap in that channel [19]. However, this still does not provide a complete solution for our application: a Type 2 user who becomes disconnected and unable to receive any revocations must stop transmitting in order to avoid interfering with Type 1 transmissions. Thus, we need a heartbeat mechanism for Type 2 users to detect connection loss and stop transmitting. This also requires a logical timeout mechanism (e.g. [20]) to ensure servers remove disconnected Type 2 users from their database of ongoing transmissions at the same logical point in time. Due to the nature of sensor and Type 2 user interactions with the system, surprisingly, it is not trivial to answer the question: **How can we combine or redesign existing mechanisms to construct a *correct* BFT-SMR-based coordination system for dynamic spectrum sharing?**

In addition to correctness, the *performance* of the coordination system can become important. The latency to process reports of Type 1 transmissions and evict Type 2 users if needed defines the amount of time that a newly detected Type 1 transmission may experience interference. Similarly, high latency to process and grant requests by Type 2 users can result in poor user experience and reduce total spectrum utilization. Due to the fact that interference effects are inherently regional, only affecting users that are sufficiently close to each other, spectrum sharing coordination is an excellent fit for scaling via sharding [21]. Preliminary work proposes a *regional agreement* model, where a BFT-SMR-replicated database of ongoing transmissions is partitioned into shards that each correspond to a discrete geographic region, with transmissions occurring at region boundaries requiring cross-shard transactions [22]. Realizing the performance benefits of sharding requires answering the question: **How can we design concrete protocols and sharding strategies that enable low-latency allocation decisions, while maintaining correctness?** We aim to develop sharding strategies that can partition the relevant geographic space into shards in a manner that balances low latency for operations within a shard (easier with smaller areas) with minimizing the number of *cross-shard*

transactions (easier with larger areas) based on application characteristics (Section V).

Finally, an interesting question is: **Can simpler coordination mechanisms provide the needed application guarantees with lower coordination overhead compared to BFT SMR?** The same preliminary work in [22] suggests a *dynamic agreement model* where the servers that need to agree to grant a given Type 2 request are selected dynamically based on the Type 2 user's location and run a BFT interactive consistency instance [23] to vote on whether the request should be granted based on their local states. However, a complete protocol specification and analysis, taking the application characteristics into account, is needed to understand whether this type of approach will improve performance in practice.

C. Enforcing Allocation Decisions

So far, we have considered the effects of faulty sensors and servers. However, it is also important to consider Byzantine user behavior: even if the coordination system makes optimal allocation decisions, without any enforcement mechanism, a Type 2 user could simply ignore denial or revocation actions and transmit anyway, causing interference. While there has been work on incentive mechanisms in blockchains, such solutions are mainly focused on incentivizing nodes to participate correctly in the consensus protocol (a role fulfilled by the servers in our setting) [24]. Some work has targeted malicious clients in BFT SMR, but it assumes the clients aim to corrupt the replicated state of the system by submitting incorrect updates [25]. In contrast, the enforcement problem for dynamic spectrum sharing is concerned with user's transmission behavior, which is not necessarily visible to the coordination system. Guardians that physically disable transmission have been proposed in real-time networks [26], [27], but it is not clear how to implement this in our wireless setting.

Addressing this likely requires monitoring and reporting Type 2 transmission behavior. Deploying sensors to detect Type 2 transmissions (similar to Section IV-A for Type 1) is possible, but demands an even more extensive, expensive infrastructure. A key question is: **Can we design cost-effective sensing solutions to detect Type 2 user misbehavior?**

V. Application Characteristics

The above discussion simplifies several application-specific challenges that arise in coordinating spectrum access.

A. Spectrum Band Characteristics

Radio propagation characteristics vary with frequency and are site-specific (line-of-sight, indoor vs outdoor propagation). If a sharing system has to coordinate 10 MHz channels in the 3.5 GHz band, the upper 6 GHz band, and the 37 GHz band, for the same transmit power, the *ranges* of transmission and sensing will be different. Environmental effects are also different, as are the effects of antenna characteristics. For simplicity, we have assumed omnidirectional antennas in Sections III and IV-A. With directional antennas, the optimization of sensor deployments will change. A practical answer to the question of optimizing sensor deployments in Section IV-A must account for the specific characteristics of the band(s) involved and the available sensor hardware.

The transmit power has substantial influence on range r and the interference caused to Type 1 or Type 2 users. Transmit power also has an impact on application needs. A higher transmit power may support higher data rates as the signal to interference plus noise ratio may be higher. The Liveness property in Section III only states that *some* user should be allowed to transmit when there is idle spectrum. But, we would like to *maximize* utilization. If, instead of purely first-come-first-served request processing, servers can batch concurrent requests from Type 2 users and select processing order within a batch based on power, it is unclear which strategy has better utilization: more Type 2 users with lower power, or the opposite. This also affects sensor deployment. A question that arises in spectrum coordination is: **if the transmit power is a variable that servers can regulate, what strategy is optimal in terms of the application correctness properties and spectrum utilization?**

B. Incumbent Characteristics

As shown in Figure 1, incumbent usage may change in space and time (navy radars are typically on the coasts). Further, some incumbent usage may be sensitive or classified. A final complication arises in the case of passive spectrum uses such as radio astronomy and remote sensing (e.g., for weather). In radio astronomy, telescopes and satellites are trying avoidance based coordination, where the telescopes avoid observations when satellites are above and/or satellites steer their transmission beams away for telescopes when they are in their vicinity (e.g., [28]). Avoidance based spectrum coordination may require a hybrid of sensing and incumbent reporting, as well as trustworthy user behavior.

C. Cross-Band Sharing

So far, we have assumed each user transmits in a slice in a single frequency band. However, when sharing across frequency bands and incumbents, we can also ask: **Can multiple bands should be purposed or bundled to suit a Type 2 user's application needs?** What strategies work best with respect to the application correctness criteria in Section III and maximizing utilization, how such requests should be prioritized or ordered, and how to ensure they do not fail remain open questions.

VI. Conclusion

Near real-time spectrum sharing with fault tolerance is an application that raises many open questions, some of which could be addressed through careful application of BFT-SMR protocols. Other questions require deeper analysis and evaluation. We raise some of the open challenges in this paper.

Acknowledgment

This work was funded in part by a subaward (seed grant) and other subawards from SpectrumX: An NSF Spectrum Innovation Center (NSF Award No. 2132700).

REFERENCES

[1] A. Abouzeid *et al.*, "National spectrum strategy research and development plan," *National Science and Technology Council*, October 2024.

[2] WinnForum, "List of SAS administrators," *https://cbrs.wirelessinnovation.org/sas-administrators*, 2024.

[3] The Federal Communications Commission, "In the matter of lower 37 GHz band: Use of spectrum bands above 24 GHz for mobile radio services," *Sixth Report and Order and Notice of Proposed Rulemaking: GN Docket No. 14-177*, April 29, 2025.

[4] (2024, May) Open afc - telecom infra project. [Online]. Available: https://telecominfraproject.com/open-afc/

[5] L. Zhang, G. Ding, Q. Wu, Y. Zou, Z. Han, and J. Wang, "Byzantine attack and defense in cognitive radio networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 3, pp. 1342–1363, 2015.

[6] K. Marzullo, "Tolerating failures of continuous-valued sensors," *ACM Trans. Comput. Syst.*, vol. 8, no. 4, p. 284–304, Nov. 1990. [Online]. Available: https://doi.org/10.1145/128733.128735

[7] C. Jin, C. Yin, M. van Dijk, S. Duan, F. Massacci, M. K. Reiter, and H. Zhang, "Pg: Byzantine fault-tolerant and privacy-preserving sensor fusion with guaranteed output delivery," in *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 3272–3286. [Online]. Available: https://doi.org/10.1145/3658644.3670343

[8] J. Gan, J. Wu, P. Li, Z. Chen, J. Zhang, Z. Chen, J. He, and S. Fan, "Joint spectrum sensing and resource allocation against byzantine attack in overlay cognitive radio networks," *Transactions on Emerging Telecommunications Technologies*, vol. 33, no. 7, p. e4500, 2022.

[9] M. B. H. Weiss, K. Werbach, D. C. Sicker, and C. E. C. Bastidas, "On the application of blockchains to spectrum management," *IEEE Transactions on Cognitive Communications and Networking*, vol. 5, no. 2, pp. 193–205, 2019.

[10] Y. Xiao, S. Shi, W. Lou, C. Wang, X. Li, N. Zhang, Y. T. Hou, and J. H. Reed, "Bd-sas: Enabling dynamic spectrum sharing in low-trust environment," *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 4, pp. 842–856, 2023.

[11] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "Trustsas: A trustworthy spectrum access system for the 3.5 ghz cbrs band," in *IEEE INFOCOM 2019 - IEEE Conference on Computer Communications*, 2019, pp. 1495–1503.

[12] T. Ariyarathna, P. Harankahadeniya, S. Isthikar, N. Pathirana, H. D. Bandara, and A. Madanayake, "Dynamic spectrum access via smart contracts on blockchain," in *2019 IEEE Wireless Communications and Networking Conference (WCNC)*, 2019, pp. 1–6.

[13] Z. Chen, L. Wang, and Y. Zhang, "Blockchain structure electromagnetic spectrum database in distributed cognitive radio monitoring system," *IEEE Transactions on Cognitive Communications and Networking*, vol. 8, no. 4, pp. 1647–1664, 2022.

[14] Z. Cheng, Y. Liang, Y. Zhao, S. Wang, and C. Sun, "A multi-blockchain scheme for distributed spectrum sharing in cbrs system," *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 2, pp. 266–280, 2023.

[15] Z. Li, W. Wang, Q. Wu, and X. Wang, "Multi-operator dynamic spectrum sharing for wireless communications: A consortium blockchain enabled framework," *IEEE Transactions on Cognitive Communications and Networking*, vol. 9, no. 1, pp. 3–15, 2023.

[16] K. Yan, W. Ma, S. Sun, and W. Wang, "Blockchain-based dynamic spectrum sharing for service-centric 6g networks: An evolutionary approach," *IEEE Transactions on Network Science and Engineering*, vol. 13, pp. 485–500, 2026.

[17] L. Perera, P. Ranaweera, S. Kusaladharma, S. Wang, and M. Liyanage, "A survey on blockchain for dynamic spectrum sharing," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 1753–1802, 2024.

[18] S. T. Muntaha, P. I. Lazaridis, M. Hafeez, Q. Z. Ahmed, F. A. Khan, and Z. D. Zaharis, "Blockchain for dynamic spectrum access and network slicing: A review," *IEEE Access*, vol. 11, pp. 17 922–17 944, 2023.

[19] A. Clement, M. Kapritsos, S. Lee, Y. Wang, L. Alvisi, M. Dahlin, and T. Riche, "Upright cluster services," in *Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles*, ser. SOSP '09. New York, NY, USA: Association for Computing Machinery, 2009, p. 277–290. [Online]. Available: https://doi.org/10.1145/1629575.1629602

[20] J. Kirsch, *Intrusion-tolerant replication under attack*. The Johns Hopkins University, 2010.

[21] G. Wang, Z. J. Shi, M. Nixon, and S. Han, "Sok: Sharding on blockchain," in *Proceedings of the 1st ACM Conference on Advances in Financial Technologies*, ser. AFT '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 41–61. [Online]. Available: https://doi.org/10.1145/3318041.3355457

[22] A. Babay, P. Krishnamurthy, I. Murtazashvili, and X. Qin, "Byzantine fault tolerance models for distributed coordination in dynamic spectrum sharing," August 2025, available at SSRN: https://ssrn.com/abstract=5375967.

[23] M. Singhal and N. G. Shivaratri, *Advanced concepts in operating systems*. McGraw-Hill, Inc., 1994.

[24] R. Han, Z. Yan, X. Liang, and L. T. Yang, "How can incentive mechanisms and blockchain benefit with each other? a survey," *ACM Comput. Surv.*, vol. 55, no. 7, Dec. 2022. [Online]. Available: https://doi.org/10.1145/3539604

[25] Y. Amir, C. Danilov, J. Lane, M. Miskin-Amir, and C. Nita-Rotaru, "Enhancing distributed systems with mechanisms to cope with malicious clients," Technical Report CNDS-2005-4, the Distributed Systems and Networks Lab, Johns Hopkins University, Tech. Rep., 2005.

[26] M. I. Alkoudsi, G. Fohler, and M. Völp, "A network-agnostic approach to enforcing collision-free time-triggered communication," in *2023 IEEE 28th Pacific Rim International Symposium on Dependable Computing (PRDC)*, 2023, pp. 71–77.

[27] G. Bauer, H. Kopetz, and W. Steiner, "The central guardian approach to enforce fault isolation in the time-triggered architecture," in *The Sixth International Symposium on Autonomous Decentralized Systems, 2003. ISADS 2003.*, 2003, pp. 37–44.

[28] A. Witze, "Swarms of satellites are harming astronomy. here's how researchers are fighting back," *Nature*, Mar. 2025.