Unleashing the Power of Multi-Server PIR for Enabling Private Access to Spectrum Databases

Mohamed Grissa, Bechir Hamdaoui, and Attila A. Yavuz

Abstract

The emergence of internet of things (IoT) is driving a rapid growth in the use of wireless devices around the world. In addition, 5G, through its support for large numbers of devices, is expected to unleash the full potential of IoT globally. The growing number of wireless devices and the massive traffic stemming from the emergence of these technologies will result in a dramatic increase of the demand for spectrum resources. Dynamic spectrum sharing, enabled through the cognitive radio network technology, emerges as a key solution for coping with these rising spectrum demands. One important approach that is currently being adopted as a potential solution for promoting dynamic spectrum sharing is the deployment and reliance on white space geo-location spectrum databases for locating spectrum availability in the TV white space. Despite the great benefits these databases offer in terms of their ability to help locating spectrum opportunities for secondary usage, they suffer from location privacy issues, as users need to reveal their location in the process of querying these databases for spectrum availability. Knowing that their whereabouts may be exposed, users can be discouraged from querying the databases, thereby hindering the adoption and deployment of this technology in future generation networks. In this article, we focus on the location privacy problem in database-driven dynamic spectrum access. Specifically, we present and compare key approaches that aim to protect the location information of secondary users in database-driven spectrum sharing and discuss some key research challenges that remain unaddressed.

I. Introduction

Geo-location white space databases have risen to become the de facto approach, as recently promoted by FCC (federal communications commission), to enable SUs to identify idle spectrum bands (also known as spectrum holes or white spaces). In this approach, SUs query a certified geo-location spectrum database by including their location information, and the database is expected to return a list of available channels in SUs' vicinity along with other transmission parameters.

In the TV band, ten entities, including Google and Microsoft, were designated by FCC as TV spectrum database administrators that are required to abide by PAWS (protocol to access white space) standard [1] which sets guidelines and operational requirements for all the entities in TV white space database-driven CRNs. As stipulated by PAWS, SUs may be served by several spectrum databases that are required to synchronize their records and provide exactly the same spectrum availability information, in any region, in response to SUs' queries. Simply put, SUs can access the same copy of the spectrum database through multiple, distinct TV white space administrators.

There also exists standards now in the 802 family, developed to regulate the utilization of specific spectrum bands, mainly in the TV band space, using database-driven CRNs. IEEE 802.22, for instance, aims to harness database-driven CRNs to

enable spectrum sharing for wide-area regional networks in the TV bands. We expect also that database-driven CRNs will be deployed in 5G to meet its tremendous data rate and spectrum requirements.

A. Location Privacy Issue in Database-Driven CRNs

Despite their great benefits in improving spectrum utilization efficiency, database-driven CRNs face serious privacy challenges especially in terms of protecting SUs' location privacy. This is mainly due to the fact that SUs have to share their location with the database to have accurate spectrum availability information. Most users will refuse to expose their location as their whereabouts could be used for malicious purposes and eventually to gain more sensitive information especially in the presence of malicious service providers. In fact, the location information can easily reveal an individual's shopping patterns, religious beliefs, or even health condition, especially when combined with publicly available data. Because of this, we envision that the public's acceptance of this paradigm will greatly depend on the security guarantees that CRNs will offer regarding these privacy risks. It is, therefore, paramount to design privacy-preserving mechanisms that protect SUs' location privacy while allowing them to harness the benefits of database-driven CRNs.

Such privacy-preserving mechanisms would have great benefits for both SUs and PUs. Indeed, they would promote dynamic spectrum sharing by encouraging SUs to trust the spectrum databases which would increase spectrum utilization efficiency and enable SUs to harness more frequency resources. This would also lead to the PUs having less harmful interference from SUs concerned about their location privacy and not willing to rely on geo-location spectrum databases to check for spectrum opportunities. However, providing location privacy guarantees to SUs cannot be achieved without a cost. In fact, this will introduce additional computation, communications, and storage overheads, which can lead to additional delay experienced by SUs when querying the spectrum database, causing them to obtain an outdated spectrum availability information in the extreme case. This, in turn, may lead to the use of occupied primary channels and causing harmful interference to PUs. Therefore, privacy-preserving mechanisms for database-driven CRNs must be designed carefully so as to guarantee SUs' privacy without compromising databases' service quality.

B. Related Work

The location privacy issue in database-driven CRNs has only recently started to gain some interest from the research community despite its importance. Broadly speaking, there are mainly two lines of privacy-preserving technologies that were adopted by existing works in this context: (i) *k-anonymity* [2] and (ii) *single-server private information retrieval (PIR)* [3]. In k-anonymity-based approaches, an SU includes $k-1$ additional locations in its query to make its location indistinguishable from these additional locations. Zhang et al. [4] use this concept in their privacy-preserving approach, by generating $k-1$ properly selected dummy locations. Their protocol requires a tradeoff between some spectrum utility and high location privacy level, meaning that achieving the latter results in a decrease in the former. Clearly, the privacy level in such approaches depends on the value of k. Small values of k may compromise privacy while large values can incur high overheads.

Single-server PIR-based approaches [5], on the other hand, which seem to have attracted more focus, are a special type of PIR technology, which essentially allows a user to retrieve a record from a database while hiding the identity of the record

from the database. They are designed to ensure privacy against a computationally bounded server that has to solve a certain computationally hard problem (e.g. discrete logarithm), depending on the underlying cryptosystem, to learn the identity of the record of interest. A PIR-based approach in the context of CRNs aims to hide the record that an SU is interested in from the spectrum database, as this record is typically associated with the user's location as depicted in Figure 1, which gives a simplified high-level overview of how single-server PIR-based approaches work. An SU needs to provide the index i of the record that it is interested in, usually through an inverted-index like mechanism. Generally, using this index of interest, SU constructs a vector whose size is equal to the number of elements in the database. All entries of this vector are equal to the identity element of the plaintext space except for the i^{th} entry, which contains a non-identity element. In Figure 1, the identity element with respect to addition is 0. The elements of this vector are encrypted, typically using an additive homomorphic cryptosystem. This vector of ciphertexts is then multiplied with records of the database and the results are aggregated together thanks to the homomorphic properties of the ciphertexts. The aggregated ciphertext is then returned to SU, who can decrypt it and retrieve the record of interest without revealing the location information. The only information that the spectrum database can learn is the vector of ciphertexts sent by SU.

Troja et al. [6] leverage PIR protocols to preserve the location privacy of mobile SUs. They make an SU send a series of several PIR queries to the database to learn spectrum availability in and around its current location. SUs gradually build a trajectory-specific spectrum knowledge cache as they move and can privately share this knowledge with other SUs within their communication range with whom they form groups and interact in a peer-to-peer manner.

Gao et al. [7] also adapt PIR in their protocol, which enables an SU to hide its location coordinates within other locations and uses some irreversible mathematical transformation to blind this information in such a way that only the SU is able to revert it, before sending it to the database. The database multiplies this blinded query with the spectrum availability matrix and returns the result to SU, which will be able to recover the record of interest by relying on the secure parameters used to transform the original query. Despite the wide-spread usage of single-server PIR protocols in the context of database-driven CRNs, these protocols incur very large overhead as they involve highly costly cryptographic operations, such as modular multiplication over large moduli, that need to be executed over the entire database for every query coming from SUs [5], [8].

Some existing approaches [9] have also adopted a new concept termed ϵ-geoindistinguishability [10], which is derived from the well-known concept of *differential privacy* to suit LBS (location-based services) applications. This is because differential privacy cannot be used as it is to provide location privacy for database-driven CRNs since it is designed for statistical databases (databases that contain private user data). More specifically, differential privacy aims to allow queries to access statistical/aggregate information (e.g., average, sum, etc.) about the data held in the entire database, but without revealing any information about the individual data records themselves. It formalizes the idea that a query should not disclose whether a user's record is present in a database, nor does it disclose any information about the record itself. The basic idea behind ϵ-geoindistinguishability consists of adding controlled random noise to SUs' locations to obfuscate them. The problem with this approach is that it may compromise the accuracy of the spectrum availability information. Moreover, the ϵ-geoindistinguishability-based approach protects an SU's location only within a radius r with a privacy level that depends on

r which results in lower privacy guarantees than PIR-based approaches in general. Note that the work in [9] targets a different goal, protecting bilateral location privacy of both PUs and SUs.

It is worth noting that PIR, on the other hand, is designed for databases that contain data records that are not private (not owned by some specific users), in that data can be accessed by any (legitimate) users as it is the case for database-driven CRNs. PIR aims to prevent the database owner from learning the identity of the record that is being queried by the user.

The common feature among these approaches is that they only consider one spectrum database, whereas in the case of the CRN paradigm, there exists, by design requirement, multiple databases that all concurrently serve SUs. The ability to query multiple databases at no database duplication cost is very unique and inherent to the database-driven CRN paradigm, and can thus be leveraged to design highly efficient privacy-preservation mechanisms. The focus of this paper is on filling this research gap by presenting a new class of approaches that leverage multiple spectrum databases to protect SUs' location privacy. To the best of our knowledge, we are the first to harness the natural existence of multiple spectrum databases to employ multi-server PIR technology [11]. Thanks to this, our approach departs from existing protocols by offering information-theoretic privacy as opposed to existing approaches which can only offer computational privacy at best. In addition, unlike the ϵ-geoindistinguishability-based technique, our multi-server-based approach does not introduce noise, but instead relies on secret sharing to divide the query among multiple databases. Moreover, our adaptation of multi-server PIR brings great performance benefits, making our approach significantly outperform the state-of-the-art location privacy-preservation methods.

Next, we first begin by presenting our proposed approach to address the location privacy issue in database-driven CRNs. Then, we discuss the different challenges that we envision to be the most prominent to database-driven CRNs.

II. Multi-Server PIR for Location Privacy Preservation in Database-driven CRNs

As mentioned earlier, multiple TV spectrum band databases exist by system design requirement, each operated by a different service provider. This FCC requirement, though set for market competitiveness and service reliability reasons, creates the opportunity to make use of multi-server PIR approaches (also known as *information-theoretic PIR*) for developing efficient privacy-preserving schemes. Multi-server PIR protocols guarantee optimal privacy against computationally unbounded servers, qualified as information-theoretic privacy [12]. However, they require duplication of the database at least two non-colluding servers to achieve this privacy level. To prevent leaking any information about the identity of the record of interest to the user, information-theoretic PIR protocols decompose the user's query into several subqueries instead of encrypting it as it is, as in the case of single-server PIR, resulting in much better performances in terms of both computation and communication. Each of these subqueries is processed by a different database and the results are then combined by SU as illustrated in Figure 2. On top of providing stronger privacy guarantees when compared to their single-server PIR counterparts, multi-server PIR protocols are also more efficient (in terms of communications/computation overheads), when the database duplication requirement comes at no cost as in the case of CRNs. Table I summarizes the differentiating aspects of both single-server and multi-server PIRs.

The fact that database-driven CRNs meet the database duplication requirement at no cost allows us to harness the benefits of multi-server PIR to achieve information-theoretic privacy without compromising protocol efficiency. We summarize the

major components of our multi-server PIR-based scheme in Figure 2 and present its steps in the next section.

A. Multi-Server PIR-Based Location Privacy Preservation

In this section, we use the first and simplest multi-server PIR protocol [12] to illustrate the concepts and benefits that multi-server PIR brings to database-driven CRNs. Let us consider a CRN system with ℓ synchronized databases operated by different service providers but all share the same content, modeled as an n-bit binary matrix, $\boldsymbol{D}$, and having r records each of size b bits. Each record of the database is indexed using a unique combination of location information, a channel number, and a timestamp. We assume that SUs and the databases agree on an inverted index technique to enable SUs to associate their location information and channel of interest at a specific time to an index k corresponding to the record of interest in the database. This index is used by SUs to privately query the databases about spectrum opportunities in their vicinity.

Based on the obtained index, SU constructs a standard basis vector $\boldsymbol{e}_k$ of size r, whose elements are equal to 0 except the k^{th} element which is equal to 1. The product of $\boldsymbol{e}_k$ with $\boldsymbol{D}$ yields the k^{th} record of the database. Since k is associated with its location information, SU cannot simply send $\boldsymbol{e}_k$ to the databases. Instead, SU decomposes this vector into several subqueries, with each being sent to a different database to prevent leaking any information about it. For this, SU picks $\ell - 1$ r-bit binary strings $\boldsymbol{\rho}_1, \cdots, \boldsymbol{\rho}_{\ell-1}$ uniformly at random, computes $\boldsymbol{\rho}_\ell$ by XORing the previously generated binary strings with $\boldsymbol{e}_k$, and distributes these ℓ binary strings among the ℓ databases as illustrated in Figure 3.

Once it receives a bitstring from SU, DB_i multiplies it with $\boldsymbol{D}$ and sends the result $\boldsymbol{R}_i$ to SU. SU collects all results from the ℓ databases and XOR them, yielding exactly the k^{th} record that SU queried thanks to the properties of XOR.

Thanks to its reliance on simple XOR operations only, this approach is very efficient in terms of computation compared to existing single-server-based approaches. We have validated this experimentally using GENI [13] cloud platform by deploying 6 virtual machines sharing the same content, emulating spectrum databases as in a real system. These virtual machines are deployed in different GENI aggregates located in Stanford, New York, Chicago, Kentucky, Utah, and Washington to count for the networking delay. As access to real spectrum database data was not possible, we generated for our experiment a random data matrix for modeling the database's content, with a fixed block size of 560 kB and a variable number of records. The block size is estimated based on the public raw data provided by FCC [14] and used by service providers to populate their spectrum databases. It takes into consideration SU's query generation time, network delay, database's query processing time, and SU's reply decoding time. We use an Internet connection of $80\,Mbps$ on the download and $30\,Mbps$ on the upload in our experiment. A laptop plays the role of a SU by querying the virtual machines representing the databases. We have deployed the multi-server PIR approach discussed in this paper along with the best existing single-server PIR protocol [8], [11], which was not used in the context of CRNs, in our testbed. We also compare our location privacy-preserving approach to the state-of-the-art protocols [6], [7], which are based on single-server PIR. Note that in our comparison we consider only PIR-based approaches since they offer stronger privacy guarantees than their *k-anonymity and ϵ-geoindistinguashibility-based approaches* and also because some of these works aims at a different privacy objective. We illustrate the measured performances of the different protocols in terms of query end-to-end delay in Figure 4, and communication overhead in Figure 5; more detailed results can be

found in [11]. The end-to-end delay measures the total delay from when SU generates its query until it receives all responses and extracts the record. As shown in Figure 4, our multi-server PIR-based approach outperforms existing approaches in terms of total delay, especially as the number of records in the database increases. This further proves the benefit of using multi-server information-theoretic PIR in database-driven CRNs over classical approaches. Note also that most of this overhead is incurred by the spectrum databases, leaving only a small number of XOR operations to be performed by SU, which makes the approach more suitable for devices with constrained resources like IoTs.

As explained earlier, each query sent to each of the ℓ databases is of size r bits. As each database performs XOR operations of some of its rows to generate a response, the size of each response is similar to that of one row of the database, that is b bits. Therefore, the total communications overhead incurred every time an SU queries the databases would be $(r + b) \cdot \ell$, which is still better than state-of-the-art techniques; this has also been confirmed via our experimental results as depicted in Figure 5.

In addition, our approach offers information-theoretic privacy as opposed to the weaker computational privacy offered by existing single-server-based approaches. It can also handle collusion of up to $\ell - 1$ databases without revealing any information about $\boldsymbol{e}_k$ of SU (i.e. its location). Only if ℓ databases collude that it will be possible to retrieve the location of SU by XORing the bitstrings that they received from SU. This is unlikely to happen in a real-life scenario, where service providers operating these spectrum databases are competing companies that have no or little interest in colluding.

B. Robustness to Database Failures

Despite its great benefits in terms of efficiency and privacy, the presented multi-server PIR approach is not robust to database failures, since if one (or more) of the ℓ databases fail to respond, SU will not be able to recover the record. The approach is not robust against *Byzantine* failures either, because if one (or more) database returns an erroneous result (maliciously or unintentionally), SU will reconstruct a wrong result, leading to inaccurate spectrum availability information, which can trigger harmful interference to incumbent users. Also, SU will not be able to identify which one of the databases misbehaved so as to not rely on it for future queries. We have addressed these issues in the work presented in [11], which proposes an improved version of the multi-server PIR approach discussed above that still provides information-theoretic privacy while ensuring robustness against these aforementioned database failures.

III. Open Research Challenges

Despite the efforts made so far to protect SUs' location privacy in database-driven CRNs, there remains challenges that need to be addressed, which we highlight next.

A. Private collection of SUs' usage data

Most of the existing location privacy preservation efforts for database-driven CRNs have focused on hiding SUs' information when querying the databases. However, after querying the spectrum database, there is an optional but important notification phase in which SU may notify the database about its operational parameters (e.g., band to be used, transmit power level, etc). Conveying such information could be beneficial in helping the database manage the spectrum resources more efficiently.

However, such information could also be easily linked to SUs' location. Now, it is true that this notification phase is optional in the PAWS standard, as it is not concerned, at this point, with the interference and coexistence among SUs. However, some applications may have more stringent requirements about their QoS and the interference among SUs, and may, thereby, require SUs' spectrum usage information to ensure better coexistence between them. Thus, there is a significant need for understanding how and what spectrum usage data should be collected and analyzed to assess spectrum utilization without an infringement on SU's privacy. However, privately updating the database with the usage information could be cumbersome, and standard privacy enhancing techniques may not be used off-the-shelf in this context.

B. Partial replication of database content

The multi-server PIR-based approach presented in this paper is efficient and provides information-theoretic privacy by exploiting the fact that database-driven CRNs involve multiple service providers by design. However, one aspect that needs to be improved is the communication overhead incurred by this approach. One potential solution could be to rely on coding techniques used in distributed storage systems to store the spectrum availability information in a way that each service provider will have a combination of different stripes of the data while ensuring redundancy but also robustness against failures. This way, each spectrum database will have to handle less data which could boost the performance and lower the communication overhead. Privacy implications of such a design must also be studied carefully. Partial replication could be also suitable for database-driven CRNs in 5G, where each cell could be managed by a base station that has its own spectrum database covering only that cell. In this scenario, partial replication-based PIR schemes could be very useful. A global database can even oversee these spectrum databases at the base station level.

C. Coexistence and interference among SUs

The PAWS protocol, in its current version, does not handle interference avoidance and coexistence among SUs. It is, however, expected to evolve and include such attributes in the future. From a privacy-preserving perspective, attributing the task of dealing with these issues to the spectrum database itself could be tricky as they both involve exchanging information that is highly correlated to SUs' location [5]. With this in mind, designing an approach to address contention over spectrum resources and coexistence among SUs by only involving SUs themselves in a distributed way would be attractive from a privacy perspective. One potential way to achieve this goal could be to rely on distributed consensus mechanisms augmented with blockchain technology [15] to enable SUs to share the availability information among each other and privately agree on how to organize their use of the spectrum without having to interact with the service providers. This could be thought of as a hybrid approach in managing the spectrum resources combining the use of spectrum databases to privately learn spectrum availability and the use of a secure distributed mechanism to handle coexistence among SUs. Blockchain can be used as an infrastructure for SUs to coordinate their use of the spectrum and to commit to their intended transmission parameters. Each record in the databases can also include a smart contract, a self-executing script running on top of the blockchain, to enforce permissible transmission parameters and how a certain channel can be shared in a specific location. Permissioned SUs can

even rely on advanced cryptographic techniques to preserve their privacy by anonymously participating in the blockchain while privately proving their membership to the system.

D. Multi-SU coordinated queries

The overhead that a privacy preserving mechanism begets on the spectrum database could be reduced if SUs' queries are minimized. This is motivated by the fact that in some situations, several SUs will seek spectrum availability in the same area, thus, they will send the same queries. This repeated work from the database side could be minimized by exploiting SUs proximity for instance. Indeed, a distributed clustering mechanism could be run among SUs to create groups based on location, and only a group representative can query the spectrum database and share spectrum availability with its group members. Such a mechanism will also minimize SUs' interaction with the database, reducing the exposure of their private information. Another way that is worth investigating to amortize the overhead is to batch SUs' queries, especially if there is a large number of SUs querying the database concurrently at a given time. Though promising and beneficial especially to spectrum databases, batching SUs' queries may not be desirable in applications with stringent latency requirements. Therefore, the optimization technique may depend on the application and its requirements.

E. Protecting the privacy of PUs

Database-driven spectrum sharing may pose privacy issues to PUs as well. Seemingly innocuous queries from SUs can reveal a great deal of sensitive information about a PU including its location, antenna parameters, transmitter identity, and times of operation. For instance, a malicious SU can request spectrum availability information beyond what it needs, such as information pertaining to a geographical area that it has no intention to be in. This might not be problematic in commercial systems, such as TV spectrum. However, when it comes to federal government systems, including public safety or military systems, such information may be considered very sensitive. In some countries, collecting and analyzing spectrum usage data from military systems is even regarded as a felony. Military PUs may decline to share their location or capabilities with a spectrum database unless there is a mechanism to provide availability information to SUs without disclosing sensitive information about PUs, which is still an open research area. This has become even more urgent after the calls made by FCC to share the 3.5 GHz band, initially intended for incumbent federal users, including military systems, with non-federal users via a dynamic spectrum access system similar to that of the TV white space.

F. Spectrum access policy enforcement.

While it is paramount to protect SUs' location information, it is also necessary to ensure that these SUs comply with the rules mandated by the spectrum databases. For instance, a SU must not use a channel that is occupied by a PU. SU's transmit power must not exceed the threshold provided by the database, so as to avoid harming PUs or other SUs. A rogue SU that does not obey these rules would be difficult to detect especially if a privacy-preserving mechanism is in place. Therefore, the goal of a privacy-preserving mechanism should be to protect SUs privacy while preventing them from misbehaving.

IV. Conclusion

In this paper, we present existing and new techniques that are designed to protect the location privacy of SUs in database-driven CRNs. We describe techniques that harness the key observation that, by design, database-driven CRNs contain multiple synchronized spectrum databases sharing the same content and operated by different service providers. Based on this observation, we leverage multi-server PIR as a more viable and realistic alternative to efficiently provide information-theoretic location privacy to SUs. We also discuss privacy and security challenges that still need to be addressed to promote the adoption of the CRN technology in future generation networks.

V. Acknowledgment

This work was supported in part by the US National Science Foundation under NSF award CNS-1162296.

References

[1] V. Chen, S. Das, L. Zhu, J. Malyar, and P. McCann, "Protocol to access white-space (paws) databases," Tech. Rep., 2015.

[2] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, vol. 10, no. 05, pp. 557–570, 2002.

[3] E. Kushilevitz and R. Ostrovsky, "Replication is not needed: Single database, computationally-private information retrieval," in *Foundations of Computer Science, 1997. Proceedings., 38th Annual Symposium on*. IEEE, 1997, pp. 364–373.

[4] L. Zhang, C. Fang, Y. Li, H. Zhu, and M. Dong, "Optimal strategies for defending location inference attack in database-driven crns," in *Communications (ICC), 2015 IEEE International Conference on*. IEEE, 2015, pp. 7640–7645.

[5] M. Grissa, B. Hamdaoui, and A. A. Yavuz, "Location privacy in cognitive radio networks: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 3, pp. 1726–1760, 2017.

[6] E. Troja and S. Bakiras, "Leveraging p2p interactions for efficient location privacy in database-driven dynamic spectrum access," in *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, 2014, pp. 453–456.

[7] Z. Gao, H. Zhu, Y. Liu, M. Li, and Z. Cao, "Location privacy in database-driven cognitive radio networks: Attacks and countermeasures," in *INFOCOM, 2013 Proceedings IEEE*. IEEE, 2013, pp. 2751–2759.

[8] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, "Xpir: Private information retrieval for everyone," *Proceedings on Privacy Enhancing Technologies*, vol. 2016, no. 2, pp. 155–174, 2016.

[9] Z. Zhang, H. Zhang, S. He, and P. Cheng, "Achieving bilateral utility maximization and location privacy preservation in database-driven cognitive radio networks," in *MASS*. IEEE, 2015.

[10] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*. ACM, 2013, pp. 901–914.

[11] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "When the hammer meets the nail: Multi-server pir for database-driven CRN with location privacy assurance," in *2017 IEEE Conference on Communications and Network Security (CNS)*, Oct 2017, pp. 1–9.

[12] B. Chor, O. Goldreich, E. Kushilevitz, and M. Sudan, "Private information retrieval," in *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*. IEEE, 1995, pp. 41–50.

[13] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar, "Geni: A federated testbed for innovative network experiments," *Computer Networks*, vol. 61, no. 0, pp. 5 – 23, 2014, special issue on Future Internet Testbeds – Part I.

[14] "Cdbs data," https://transition.fcc.gov/Bureaus/MB/Databases/cdbs/, accessed: 2018-05-17.

[15] K. Kotobi and S. G. Bilén, "Blockchain-enabled spectrum access in cognitive radio networks," in *Wireless Telecommunications Symposium (WTS), 2017*. IEEE, 2017, pp. 1–6.

BIOGRAPHIES

MOHAMED GRISSA is currently a Ph.D. candidate in Electrical and Computer Engineering at Oregon State University. He received his diploma of engineering in Telecommunications from the Ecole Superieure des Communications de Tunis (SUP'COM), Tunisia, in 2011, and his M.S. in ECE from Oregon State University in 2015. His research interests include privacy and security in wireless networks, cognitive radio networks, eHealth systems, and Blockchain technology.

BECHIR HAMDAOUI is a Professor in the School of EECS at Oregon State University. His research interest is on the general areas of networking systems and wireless communication. He won several awards, including the ICC 2017 and IWCMC 2017 Best Paper Awards, the 2016 EECS Outstanding Research Award, and the 2009 NSF CAREER Award. He served as a Distinguished Lecturer for the IEEE Communication Society for 2016 and 2017.

ATTILA ALTAY YAVUZ is an Assistant Professor in the Department of Computer Science and Engineering, University of South Florida (2018). He is a recipient of NSF CAREER Award (2017) and DBSec Best Paper Award. His research on privacy enhancing technologies and intra-vehicular network security are in the process of technology transfer with world-wide deployments. He has authored more than 45 research articles in top conferences and journals along with several patents.

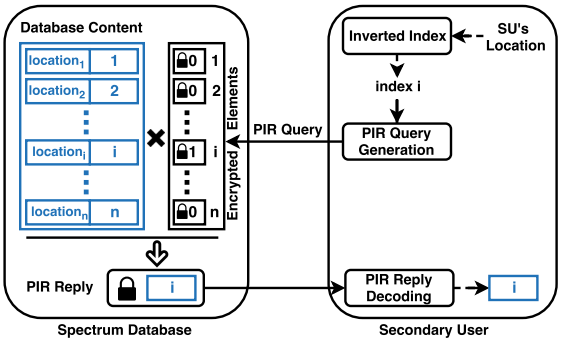

Figure 1: Single-server *PIR*-based location privacy-preserving approach

TABLE I: Comparison between single-server and multi-server PIRs

	Single-server PIR	**Multi-server** PIR
Computation overhead	High	Low
Communication overhead	High	Low
Database replication	Not needed	Required
Privacy level	Not optimal (Computational privacy only)	Optimal (Information-theoretic privacy)

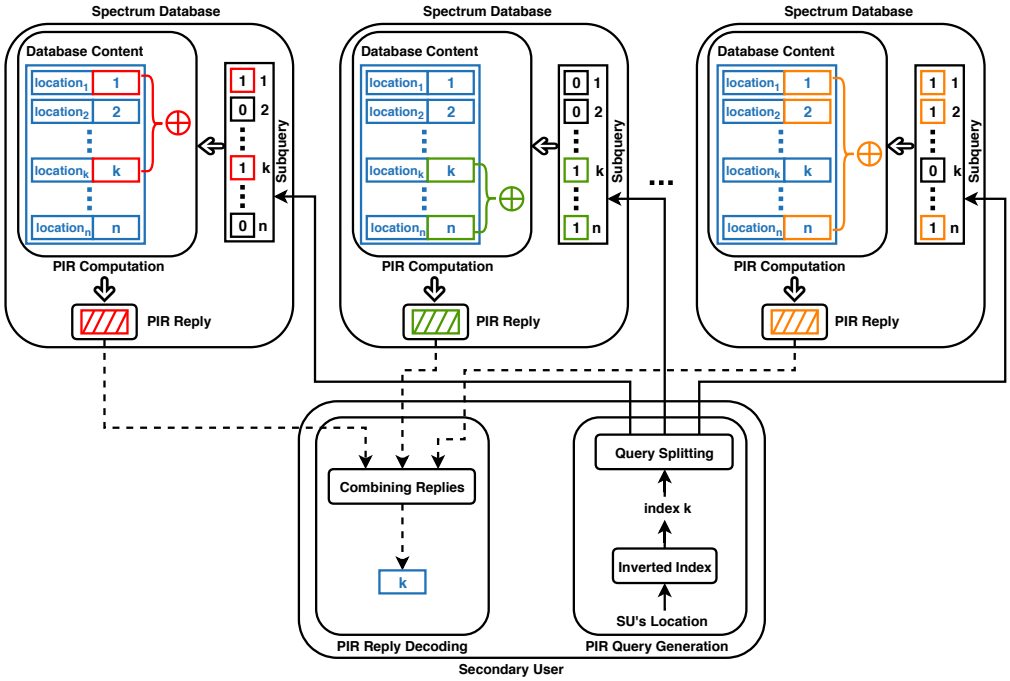

Figure 2: Multi-server *PIR*-based location privacy-preserving approach

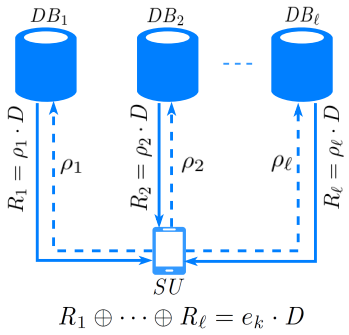

Figure 3: Main steps of multi-server PIR-based location privacy using Chor's PIR

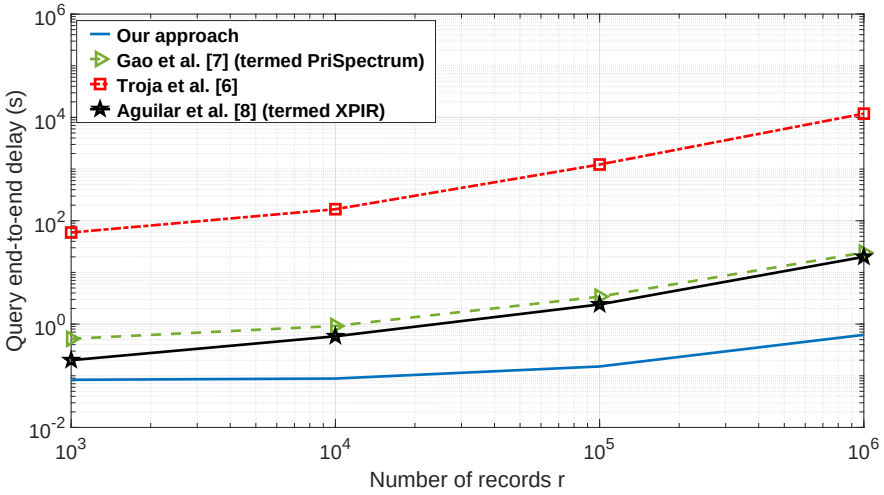

Figure 4: Query end-to-end delays under multi-server PIR and single-server PIR.

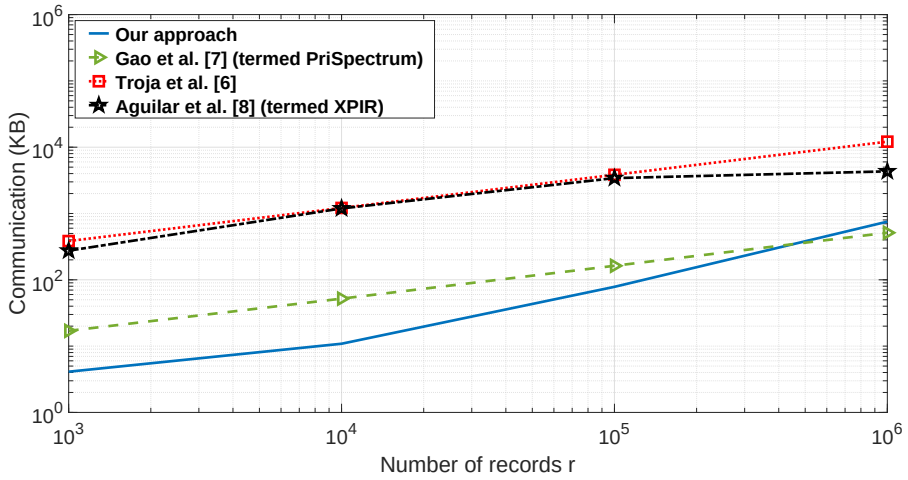

Figure 5: Communication overheads under multi-server PIR and single-server PIR: $b = 560$ B.