When the Hammer Meets the Nail: Multi-Server PIR for Database-Driven CRN with Location Privacy Assurance

Mohamed Grissa, Attila A. Yavuz, and Bechir Hamdaoui
Oregon State University, grissam,attila.yavuz,hamdaoui@oregonstate.edu

***Abstract*—We show that it is possible to achieve information theoretic location privacy for secondary users (SUs) in database-driven cognitive radio networks (CRNs) with an end-to-end delay less than a second, which is significantly better than that of the existing alternatives offering only a computational privacy. This is achieved based on a keen observation that, by the requirement of Federal Communications Commission (FCC), all certified spectrum databases synchronize their records. Hence, the same copy of spectrum database is available through multiple (distinct) providers. We harness the synergy between multi-server private information retrieval (PIR) and database-driven CRN architecture to offer an optimal level of privacy with high efficiency by exploiting this observation. We demonstrated, analytically and experimentally with deployments on actual cloud systems that, our adaptations of multi-server PIR outperform that of the (currently) fastest single-server PIR by a magnitude of times with information theoretic security, collusion resiliency and fault-tolerance features. Our analysis indicates that multi-server PIR is an ideal cryptographic tool to provide location privacy in database-driven CRNs, in which the requirement of replicated databases is a natural part of the system architecture, and therefore SUs can enjoy all advantages of multi-server PIR without any additional architectural and deployment costs.**

***Keywords*—Database-driven cognitive radio networks, location privacy, dynamic spectrum access, private information retrieval.**

I. Introduction

The rapid growth of connected wireless devices has dramatically increased the demand for wireless spectrum and led to a serious shortage in spectrum resources. Cognitive radio networks (CRNs) [1] have emerged as a promising technology for solving this shortage problem by enabling dynamic spectrum access (DSA), which improves the spectrum utilization efficiency by allowing unlicensed/secondary users (SUs) to exploit unused spectrum bands (aka spectrum holes or white spaces) of licensed/primary users (PUs).

Currently, two approaches are being adopted to identify these white spaces: spectrum sensing and geolocation spectrum databases. In the spectrum sensing-based approach, SUs need to sense the PU channel to determine whether the channel is available for opportunistic use. The spectrum database-based approach, on the other hand, does not require that SUs perform sensing to check for spectrum availability. It instead requires that SUs query a database (DB) to learn about spectrum opportunities in their vicinity. This approach, already promoted and adopted by the Federal Communications Commission (FCC), was introduced as a way to overcome the technical hurdles faced by the spectrum sensing-based approaches, thereby enhancing the efficiency of spectrum utilization, improving the accuracy of available spectrum identification, and reducing the complexity of terminal devices [2]. Moreover, it pushes the responsibility and complexity of complying with spectrum policies to DB and eases the adoption of policy changes by limiting updates to just a handful number of databases, as opposed to updating large numbers of devices [3].

FCC has designated nine entities (e.g. Google [4], iconectiv [5], and Microsoft [6]) as TV bands device database administrators which are required to follow the guidelines provided by PAWS (Protocol to Access White Space) standard [3]. PAWS sets guidelines and operational requirements for both the spectrum database and the SUs querying it. These include: SUs need to be equipped with geo-location capabilities, SUs must query DB with their specific location to check channel availability before starting their transmissions, DB must register SUs and manage their access to the spectrum, DB must respond to SUs' queries with the list of available channels in their vicinity along with the appropriate transmission parameters. As specified by PAWS standard, SUs may be served by several spectrum databases and are required to register to one or more of these databases prior to querying them for spectrum availability. The spectrum databases are reachable via the Internet, and SUs querying these databases are expected to have some form of Internet connectivity [7].

A. Location Privacy Issues in Database-Driven CRNs

Despite their effectiveness in improving spectrum utilization efficiency, database-driven CRNs suffer from serious security and privacy threats. Since they could be seen as a variant of of *location based service (LBS)*, the disclosure of location information of SUs represents the main threat to SUs when it comes to obtaining spectrum availability from DBs. This is simply because *SUs have to share their locations with DBs to obtain spectrum availability information in their vicinity.* The fine-grained location, when combined with publicly available information, can easily reveal other personal information about an individual including his/her behavior, health condition, personal habits or even beliefs. For instance, an adversary can learn some information about the health condition of a user by observing that the user regularly goes to a hospital for example. The frequency and duration of these visits can even reveal the seriousness of a user illness and even the type of illness if the location corresponds to that of a specialty clinic. The adversary could even sell this information to pharmaceutical advertisers without the user's consent.

Being aware of such potential privacy threats, SUs may

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

978-1-5386-0683-4/17/$31.00

refuse to rely on DB for spectrum availability information, which may present a serious barrier to the adoption of database-based CRNs, and to the public acceptance and promotion of the dynamic spectrum sharing paradigm. Therefore, *there is a critical need for developing techniques to protect the location privacy of SUs while allowing them to harness the benefits of the CRN paradigm without disrupting the functionalities that these techniques are designed for to promote dynamic spectrum sharing*.

B. Research Gap and Objectives

Despite the importance of the location privacy issue in CRNs, only recently has it started to gain interest from the research community [8]. Some works focus on addressing this issue in the context of collaborative spectrum sensing [9]–[13]; others address it in the context of dynamic spectrum auction [14]. Protecting SUs' location privacy in database-driven CRNs is a more challenging task, merely because SUs are required, by protocol design, to provide their physical location to DB to learn about spectrum opportunities in their vicinities. The existing location privacy preservation techniques for database-driven CRN (e.g., [2], [15]–[19]) generally rely on three main lines of privacy preserving technologies, (i) *k-anonymity* [20], (ii) *differential privacy* [21] and (iii) single-server *Private Information Retrieval (PIR)* [22]. However, the direct adaptation of *k-anonymity* based techniques have been shown to yield either insecure or extremely costly results [23]. The solutions adapting *differential privacy* (e.g., [18]) not only incur a non-negligible overhead, but also introduce a noise over the queries, and therefore they may negatively impact the accuracy of spectrum availability information.

Among these alternatives, single-server PIR seems to be the most popular alternative in the context of CRNs. PIR technology is a suitable choice for database-driven CRNs, as it permits privacy preserving queries on a public database, and therefore can enable a SU to retrieve spectrum availability information from the database without leaking his/her location information. However, single-server PIR protocols rely on highly costly partial homomorphic encryption schemes, which need to be executed over the entire database for each query. Indeed, as we also demonstrated with our experiments in Section IV, the execution of a single query even with some of the most efficient single-server PIR schemes [24] takes approximately 20 seconds with a 80 $Mbps$/ 30$Mbps$ bandwidth on a moderate size database (e.g., 10^6 entries). An end-to-end delay with the orders of 20 seconds might be undesirable for spectrum sensing needs of SUs in real-life applications. Also, some of the state-of-the-art efficient computational PIR schemes [25] that are used in the context of CRNs have been shown to be broken [24]. We provide a discussion about the existing privacy enhancing techniques and their potential adaptations to database-driven CRN settings in Sections IV and V.

There is a significant need for practical location privacy preservation approaches for database-driven CRNs that can meet the efficiency and functionality requirements of SUs.

C. Our Observation and Contribution

The objective of this paper is to develop efficient techniques for database-driven CRNs that preserve the location privacy of SUs during their process of acquiring spectrum availability information. Specifically, we will aim for the following design objectives: (i) (*location privacy*) Preserve the location privacy of SUs while allowing them to receive spectrum availability information; (ii) *(efficiency and practicality)* Incur minimum computation, communication and storage overhead. The cryptographic delay must be minimum to permit fast spectrum availability decision for the SUs, and storage/processing cost must be low to enable practical deployments. (iii) (*fault-tolerance and robustness)* Mitigate the effects of system failures or misbehaving entities (e.g., colluding databases). *It is a very challenging task to meet all of these seemingly conflicting design goals simultaneously.*

The main idea behind our proposed approaches is to harness special properties and characteristics of the database-driven CRN systems to employ private query techniques that can overcome the significant performance, robustness and privacy limitations of the state-of-the-art techniques. Specifically, our proposed approach is based on the following observation:

Observation: *FCC requires that all of its certified databases synchronize their records obtained through registration procedures with one another [26], [27] and need to be consistent across the other databases by providing exactly the same spectrum availability information, in any region, in response to SUs' queries [28]. That is, the same copy of spectrum database is available and accessible to the SUs via multiple (distinct) spectrum database administrators/providers. Is it possible exploit this observation to achieve efficiency location preservation techniques for database-driven CRN?*

In practice, as stated in PAWS standard [3], SUs have the option to register to multiple spectrum databases belonging to multiple service providers. Currently, many companies (e.g. Google [4], iconectiv [5], etc) have obtained authorization from FCC to operate geo-location spectrum databases upon successfully complying to regulatory requirements. Several other companies are still underway to acquire this authorization [29]. Thus, it is more natural and realistic to take this fact into consideration when designing privacy preserving protocols for database-based CRNs. Based on this observation, our main contribution is as follows:

TABLE I: Performance Comparison

Scheme	Comm.	Delay			Privacy
		DB	SU	end-to-end	
LP-Chor	753 KB	0.48 s	0.008 s	0.62 s	$(\ell-1)$-*private*
LP-Goldberg	6000 KB	1.21 s	0.32 s	1.78 s	t-*private* ℓ-*comp.-private*
PriSpectrum [2]	512.8 KB	21 s	0.084 s	24.2	*underlying PIR broken*
Troja et al [17]	8.4 KB	11760 s	5.62 s	11766 s	*computationally-private*
Troja et al [16]	12120 KB	11760 s	48 s	11820 s	*computationally-private*
XPIR [24]	4321 KB	17.66 s	0.34 s	20.53 s	*computationally-private*

Parameters: $n = 560\ MB,\ b = 560\ B,\ r = 10^6,\ \ell = 6,\ w = 8,\ k = 6$

Our Contribution: *To the best of our knowledge, we are the first to exploit the observation that multiple copies of spectrum DBs are available by nature in database-driven CRNs, and therefore it is possible to harness multi-server PIR techniques [22], [30] that offer information-theoretic privacy with substantial efficiency advantages over single-server PIR. We show, analytically and experimentally with deployments on cloud systems, that our adaptation of multi-server PIR techniques significantly outperforms the state-of-the-art location privacy preservation methods as demon-*

strated in Table I and detailed in Section IV. Moreover, our adaptations achieve the information theoretical privacy while existing alternatives offer only computational privacy. This feature provides an assurance against even post-quantum adversaries [31] and can avoid recent attacks on computational PIR [24].

Notice that, multi-server PIR techniques require the availability of multiple (synchronized) replicas of the database. Therefore, despite their high efficiency and security, they received a little attention from the practitioners. For instance, in traditional data outsourcing settings (e.g., private cloud storage), the application requires a client to outsource only a single copy of its database. The distribution and maintenance of multiple copies of the database across different service providers brings additional architectural and deployment costs, which might not be economically attractive for the client.

In this paper, we showcased one of the first natural use-cases of multi-server PIR, in which the multiple copies of synchronized databases are already available by the original design of application (i.e., spectrum availability information in multi-database CRNs), and therefore multi-server PIR does not introduce any extra overhead on top of the application. Exploiting this synergy between multi-database CRN and multi-server PIR permitted us to provide informational theoretical location privacy for SUs with a significantly better efficiency compared to existing single-server PIR approaches.

Desirable Properties: We outline the desirable properties of our approaches below.

- *Computational efficiency:* The adapted approaches are much more efficient than existing location privacy preserving schemes. For instance, as shown in Table I, $LP\text{-}Chor$ and $LP\text{-}Goldberg$ are more than 3 orders of magnitudes faster than the schemes proposed by Troja et al. [16], [17], and 10 times faster than XPIR [24] and $PriSpectrum$ [2].
- *Information Theoretical Privacy Guarantees:* They can achieve information-theoretic privacy which is the optimal privacy level that could be reached as opposed to computational privacy guarantees offered by existing approaches. In fact some of these approaches are prone to recent attacks on computational-PIR protocols [24] and are not secure against post-quantum adversaries [31].
- *Low communication overhead:* Both approaches provide a reasonable communication overhead that is a middle ground between the fastest computational PIR [24] and the most communication efficient computational PIR [32].
- *Fault-Tolerance and Robustness*: Our proposed approaches are resilient to the issues that are associated with multi-server architectures: failures, byzantine behavior, and collusion. Both $LP\text{-}Chor$ and $LP\text{-}Goldberg$ can handle collusion of multiple DBs. In addition, $LP\text{-}Goldberg$ can also handle faulty and byzantine DBs.
- *Experimental evaluation on actual cloud platforms*: We deploy our proposed approaches on a real cloud platform, GENI [33], to show their feasibility. In our experiment, we create multiple geographically distributed VMs each playing the role of a DB. A laptop plays the role of a SU that queries DBs, i.e. VM s. Our experiments confirm the superior computational advantages of the adaption of multi-server PIR over the existing alternatives.

II. Preliminaries and Models

A. Notation and Building Blocks

We summarize our notations in Table II. Our adaptations of multi-server PIR rely on the following building blocks.

TABLE II: Notations

DB	Spectrum database
SU	Secondary user
CRN	Cognitive radio network
ℓ	Number of spectrum databases
$\boldsymbol{D}$	Matrix modeling the content of DB
r	Number of records in $\boldsymbol{D}$
n	Size of the database in bits
b	Size of one record of the database in bits
w	Size of one word of the database in bits
s	Number of words per block
β	Index of the record sought by SU
t	Privacy level (tolerated number of colluding DBs)
k	Number of responding DBs
ϑ	Number of byzantine DBs

Private Information Retrieval (PIR): PIR allows a user to retrieve a data item of its choice from a database, while preventing the server owning the database from gaining information on the identity of the item being retrieved [34]. One trivial solution to this problem is to make the server send an entire copy of the database to the querying user. Obviously, this is a very inefficient solution to the PIR problem as its communication complexity may be prohibitively large. However, it is considered as the only protocol that can provide information-theoretic privacy, i.e. perfect privacy, to the user's query in single-server setting. There are two main classes of PIR protocols according to their privacy level: information-theoretic PIR ($itPIR$) and computational PIR ($cPIR$).

- *Information-theoretic or multi-server PIR:* It guarantees information-theoretic privacy to the user, i.e. privacy against computationally unbounded servers. This could be achieved efficiently only if the database is replicated at $k \geq 2$ non-communicating servers [22], [30]. The main idea behind these protocols consists on decomposing each user's query into several sub-queries to prevent leaking any information about the user's intent.
- *Computational or single-server PIR:* It guarantees privacy against computationally bounded server(s). In other words, a server cannot get any information about the identity of the item retrieved by the user unless it solves a certain computationally hard problem (e.g. prime factorization of large numbers), which is common in modern cryptography. Thus, they offer weaker privacy than their $itPIR$ counterparts [25], [35].

Shamir Secret Sharing: This is a concept introduced by Shamir et al. [36] to allow a secret holder to divide its secret $\mathcal{S}$ into ℓ shares $\mathcal{S}_1, \cdots, \mathcal{S}_\ell$ and distribute these shares to ℓ parties. In (t, ℓ)-Shamir secret sharing, where $t < \ell$, if t or fewer combine their shares, they learn no information about $\mathcal{S}$. However, if more than t come together, they can easily recover $\mathcal{S}$. Given a secret $\mathcal{S}$ chosen arbitrarily form a finite field, the (t, ℓ)-Shamir secret sharing scheme works as follows: the secret holder chooses ℓ arbitrary non-zero distinct elements $\alpha_1, \cdots, \alpha_\ell \in \mathbb{F}$. Then, it selects t elements $\sigma_1, \cdots, \sigma_t \in \mathbb{F}$ uniformly at random. Finally, the secret holder constructs the polynomial $f(x) = \sigma_0 + \sigma_1 x + \sigma_2 x^2 + \cdots + \sigma_t x^t$, where

$\sigma_0 = \mathcal{S}$. The ℓ shares $\mathcal{S}_1, \cdots, \mathcal{S}_\ell$, that are given to each party, are $(\alpha_1, f(\alpha_1)), \cdots, (\alpha_\ell, f(\alpha_\ell))$. Any $t+1$ or more parties can recover the polynomial f using Lagrange interpolation and thus they can reconstruct the secret $\mathcal{S} = f(0)$. However, t or less parties can learn nothing about $\mathcal{S}$. In other words, if $t+1$ shares of $\mathcal{S}$ are available then $\mathcal{S}$ can be easily recovered.

B. System Model and Security Definitions

We consider a database-driven CRN that contains ℓ DBs, where $\ell \geq 2$, and a SU registered to these DBs to learn spectrum availability information in its vicinity. We assume that these DBs share the same content and that they are synchronized as mandated by PAWS standard [3]. We also assume that DBs may collude in order to infer SU's location. In the following, we present our security definitions.

Definition 1. Byzantine *DB*: *This is a faulty DB that runs but produces incorrect answers, possibly chosen maliciously or computed in error. This might be due to a corrupted or obsolete copy of the database caused by a synchronization problem with the other DBs.*

Definition 2. t-private *PIR*: *The privacy of the query is information-theoretically protected, even if up to t of the ℓ DBs collude, where $t < \ell$.*

Definition 3. ϑ-Byzantine-robust *PIR*: *Even if ϑ of the responding DBs are* Byzantine*, SU can reconstruct the correct database item, and determine which of the DBs provided incorrect response.*

Definition 4. k-out-of-ℓ *PIR*: *SU can reconstruct the correct record if it receives at least k-out-of-ℓ responses, $2 \leq k \leq \ell$.*

Definition 5. Robust *PIR*: *It can deal with DBs that do not respond to SU's queries and allows SU to reconstruct the correct output of the queries in this situation.*

III. Proposed Approaches

In the proposed approaches, we tailor multi-server PIR to the context of multi-DB CRNs. We start by illustrating the structure of the spectrum database that we consider. Then, we give two approaches, each adapts a multi-server PIR protocol with different security and performance properties. We model the content of each DB as an $r \times s$ matrix $\boldsymbol{D}$ of size n bits, where s is the number of words of size w in each record/block of the database and r is the number of records in the database, i.e. $r = n/b$, where $b = s \times w$ is the block size in bits. The k^{th} row of $\boldsymbol{D}$ is the k^{th} record of the database.

$$\boldsymbol{D} = \begin{bmatrix} w_{11} & w_{12} & \dots & w_{1s} \\ w_{21} & w_{22} & \dots & w_{2s} \\ \vdots & \vdots & \ddots & \vdots \\ w_{r1} & w_{r2} & \dots & w_{rs} \end{bmatrix}$$

We further assume that each row of the database corresponds to a unique combination of the tuple (l_x, l_y, C, ts), where l_x and l_y represent one location's latitude and longitude, respectively, C is a channel number, and ts is a time-stamp. We also assume that SUs can associate their location information with the index β of the corresponding record of interest in the database using some inverted index technique that is agreed upon with DBs. An SU that wishes to retrieve record $\boldsymbol{D}_\beta$ without any privacy consideration can simply send to DB a row vector e_β consisting of all zeros except at position β where it has the value 1. Upon receiving e_β, DB multiplies it with $\boldsymbol{D}$ and sends record $\boldsymbol{D}_\beta$ back to SU as we illustrate below:

$$[0 \quad \dots \quad 0 \quad 1 \quad 0 \quad \dots \quad 0] \begin{bmatrix} w_{11} & w_{12} & \dots & w_{1s} \\ w_{21} & w_{22} & \dots & w_{2s} \\ \vdots & \vdots & \ddots & \vdots \\ w_{r1} & w_{r2} & \dots & w_{rs} \end{bmatrix}$$
$$= [w_{\beta 1} \quad w_{\beta 2} \quad \dots \quad w_{\beta s}]$$

This trivial approach makes it easy for DBs to learn SU's location from the vector e_β as $\boldsymbol{D}$ is indexed based on location. In the following we present two approaches that try to hide the content of e_β from DBs, and thus preserve SU's location privacy. The approaches present a tradeoff between efficiency, and some additional security features.

A. Location Privacy with Chor (LP-Chor)

Our first approach, termed LP-$Chor$, harnesses the simple and efficient $itPIR$ protocol proposed by Chor et al. [22]. We describe the different steps of LP-$Chor$ in Algorithm 1 and highlight these steps in Fig. 1. Elements of $\boldsymbol{D}$ in this scheme belong to $GF(2)$, i.e. $w = 1$ bit and $b = s$.

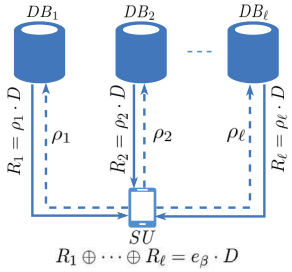

Fig. 1: Main steps of LP-$Chor$ Algorithm

Algorithm 1 $\boldsymbol{D}_\beta \leftarrow LP\text{-}Chor(\ell,\ r,\ b)$

SU
1: $\beta \leftarrow InvIndex(l_x, l_y, C, ts)$
2: Sets standard basis vector $\boldsymbol{e}_\beta \leftarrow \overrightarrow{1}_\beta \in \mathbb{Z}^r$
3: Generates $\boldsymbol{\rho}_1, \cdots, \boldsymbol{\rho}_{\ell-1} \in_R GF(2)^r$
4: $\boldsymbol{\rho}_\ell \leftarrow \boldsymbol{\rho}_1 \oplus \cdots \oplus \boldsymbol{e}_\beta$
5: Sends $\boldsymbol{\rho}_i$ to DB_i, for $1 \leq i \leq \ell$

Each *DB*$_i$
6: Receives $\boldsymbol{\rho}_i = \boldsymbol{\rho}_{i1} \cdots \boldsymbol{\rho}_{ir} \in \{0,1\}^r$
7: $\boldsymbol{R} \leftarrow \bigoplus_{\substack{1 \leq j \leq r \\ \boldsymbol{\rho}_{ij}=1}} \boldsymbol{D}_j$, $\boldsymbol{D}_j$ is the j^{th} block of $\boldsymbol{D}$
8: Sends $\boldsymbol{R}_i$ to SU

SU
9: Receives $\boldsymbol{R}_1, \cdots, \boldsymbol{R}_\ell$
10: $\boldsymbol{D}_\beta \leftarrow \boldsymbol{R}_1 \oplus \cdots \oplus \boldsymbol{R}_\ell$

In LP-$Chor$, SU starts by invoking the inverted index subroutine $InvIndex(l_x, l_y, C, ts)$ which takes as input the coordinates of the user, its channel of interest, and a time-stamp and returns a value β. This value corresponds to the index of the record $\boldsymbol{D}_\beta$ of $\boldsymbol{D}$ that SU is interested in. SU then constructs $\boldsymbol{e}_\beta$, which is a standard basis vector $\overrightarrow{1}_\beta \in \mathbb{Z}^r$ having 0 everywhere except at position β which has the value

1 as we discussed previously. SU also picks $\ell-1$ r-bit binary strings $\boldsymbol{\rho}_1, \cdots, \boldsymbol{\rho}_{\ell-1}$ uniformly at random from $GF(2)^r$, and computes $\boldsymbol{\rho}_\ell = \boldsymbol{\rho}_1 \oplus \cdots \oplus \boldsymbol{e}_\beta$. Finally, SU sends $\boldsymbol{\rho}_i$ to DB_i, for $1 \leq i \leq \ell$. Upon receiving the bit-string $\boldsymbol{\rho}_i = \boldsymbol{\rho}_{i1} \oplus \cdots \boldsymbol{\rho}_{ir}$ of length r, DB_i computes $\boldsymbol{R}_i = \boldsymbol{\rho}_i \cdot \boldsymbol{D}$, which could be seen also as the XOR of those blocks $\boldsymbol{D}_j$ in $\boldsymbol{D}$ for which the j^{th} bit of $\boldsymbol{\rho}_i$ is 1, then sends $\boldsymbol{R}_i$ back to SU. SU receives $\boldsymbol{R}_i$s from DB_is, $1 \leq i \leq \ell$, and computes $\boldsymbol{R}_1 \oplus \cdots \oplus \boldsymbol{R}_\ell = (\boldsymbol{\rho}_1 \oplus \cdots \oplus \boldsymbol{\rho}_\ell) \cdot \boldsymbol{D} = \boldsymbol{e}_\beta \cdot \boldsymbol{D}$, which is the β^{th} block of the database that SU is interested in, from which it can retrieve the spectrum availability information.

LP-Chor is very efficient thanks to its reliance on simple XOR operations only as we discuss in Section IV. It is also $(\ell-1)$-*private*, by Definition 2, as collusion of up to $\ell-1$ DBs cannot enable them to learn $\boldsymbol{e}_\beta$, and consequently its location. In fact, only if ℓ DBs collude, then they will be able to learn $\boldsymbol{e}_\beta$ by simply XORing their $\{\boldsymbol{\rho}_i\}_{i=1}^{\ell}$. However this approach suffers from two main drawbacks. First, it is not *robust* since even if one DB fails to respond, SU will not be able to recover $\boldsymbol{D}_\beta$. Second, it is not *byzantine robust*; if one or more DBs return a wrong response, SU will reconstruct a wrong block and also will not be able to recognize which DB misbehaved so as not to rely on it for future queries. In Section III-B we discuss a second approach that improves on these two aspects but with some additional overhead.

B. Location Privacy with Goldberg (LP-Goldberg)

Our second approach, termed *LP-Goldberg*, is based on Goldberg's *itPIR* protocol [30] which uses Shamir secret sharing to hide $\boldsymbol{e}_\beta$, i.e. SU's query. It is a modification of Chor's scheme [22] to achieve both *robustness* and *byzantine robustness*. Rather than working over $GF(2)$ (binary arithmetic), this scheme works over a larger field $\mathbb{F}$, where each element can represent w bits. The database $\boldsymbol{D} = (w_{jk}) \in \mathbb{F}^{r \times s}$ in this scheme, is an $r \times s$ matrix of elements of $\mathbb{F} = GF(2^w)$. Each row represents one block of size b bits, consisting of s words of w bits each. Again, $\boldsymbol{D}$ is replicated among ℓ databases DB_i. We summarize the main steps of *LP-Goldberg* protocol in Algorithm 2 and illustrate them in Fig. 2.

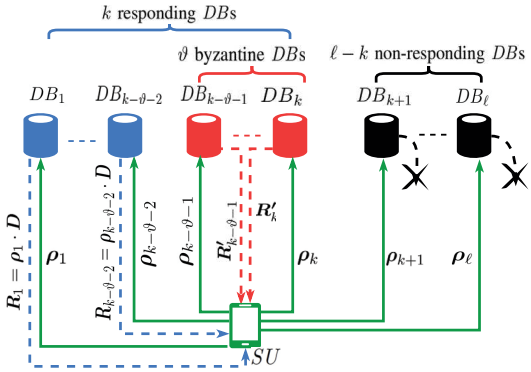

Fig. 2: Illustration of *LP-Goldberg*

To determine the index β of the record that corresponds to its location, SU starts by invoking the subroutine $InvIndex(l_x, l_y, C, ts)$ then constructs the standard basis vector $\boldsymbol{e}_\beta \in \mathbb{F}^r$ as explained earlier. SU then uses (ℓ, t)-Shamir secret sharing to divide the vector $\boldsymbol{e}_\beta$ into ℓ independent shares $(\alpha_1, , \boldsymbol{\rho}_1) \cdots, (\alpha_\ell, \boldsymbol{\rho}_\ell)$ to ensure a *t-private PIR* protocol as in Definition 2. That is, SU chooses ℓ distinct non-zero elements $\alpha_i \in \mathbb{F}^*$ and creates r random degree-t polynomials $f_1, \cdots, f_r$ satisfying $f_j(0) = \boldsymbol{e}_\beta[j]$. SU then sends to each DB_i its share corresponding to the vector $\boldsymbol{\rho}_i = \langle f_1(\alpha_i), \cdots, f_r(\alpha_i) \rangle$. Each DB_i then computes the product $\boldsymbol{R}_i = \boldsymbol{\rho}_i \cdot \boldsymbol{D} = \langle \sum_j f_j(\alpha_i)\boldsymbol{w}_{j1}, \cdots, \sum_j f_j(\alpha_i)\boldsymbol{w}_{js} \rangle \in \mathbb{F}^s$ and sends $\boldsymbol{R}_i$ to SU.

Some DBs may fail to respond to SU's query and only k-out-of-ℓ send their responses to SU. SU collects k responses from the k responding DBs and tries to recover the record at index β from the $\boldsymbol{R}_i$s by using the $EasyRecover()$ subroutine from [30] which uses Lagrange interpolation to recover $\boldsymbol{D}_\beta$ from the secret shares $(\alpha_1, \boldsymbol{R}_1), \cdots, (\alpha_k, \boldsymbol{R}_k)$. This is possible thanks to the use of (ℓ, t)-Shamir secret sharing as long as $k > t$ and these k DBs are honest. In fact, by the linearity property of Shamir secret sharing, since $\{(\alpha_i, \boldsymbol{\rho}_i)\}_{i=1}^{\ell}$ is a set of (ℓ, t)-Shamir secret shares of $\boldsymbol{e}_\beta$, then $\{(\alpha_i, \boldsymbol{R}_i)\}_{i=1}^{\ell}$ will be also a set of (ℓ, t)-Shamir secret shares of $\boldsymbol{e}_\beta \cdot \boldsymbol{D}$, which is the β^{th} block of the database. Thus, it is possible for SU to reconstruct $\boldsymbol{D}_\beta$ using Lagrange interpolation as explained in Section II, by relying only on the k responses which makes *LP-Goldberg* robust by Definition 5. Also, the $EasyRecover()$ can detect the DBs that responded honestly, thus those that are byzantine as well, which should discourage DBs from misbehaving. More details about this subroutine could be found in [30].

Algorithm 2 $\boldsymbol{D}_\beta \leftarrow$ *LP-Goldberg*(ℓ, r, b, t, w)

SU
1: $\beta \leftarrow InvIndex(l_x, l_y, C, ts)$
2: Sets standard basis vector $\boldsymbol{e}_\beta \leftarrow \overrightarrow{1}_\beta \in \mathbb{Z}^r$
3: Chooses ℓ distinct $\alpha_1, \cdots, \alpha_\ell \in \mathbb{F}^*$
4: Creates r random degree-t polynomials $f_1, \cdots, f_r \in_R \mathbb{F}[x]$ s.t. $f_j(0) = \boldsymbol{e}_\beta[j]$' $\forall j \in [1, \cdots, r]$
5: $\boldsymbol{\rho}_i \leftarrow \langle f_1(\alpha_i), \cdots, f_r(\alpha_i) \rangle, \forall i \in [1, \cdots, \ell]$
6: Sends $\boldsymbol{\rho}_i$ to DB_i, $\forall i \in [1, \cdots, \ell]$

Each honest *DB*$_i$
7: Receives $\boldsymbol{\rho}_i$
8: $\boldsymbol{R}_i \leftarrow \boldsymbol{\rho}_i \cdot \boldsymbol{D} = \langle \sum_j f_j(\alpha_i)\boldsymbol{w}_{j1}, \cdots, \sum_j f_j(\alpha_i)\boldsymbol{w}_{js} \rangle$
9: Sends $\boldsymbol{R}_i$ to SU

SU
10: Receives $\boldsymbol{R}_1, \cdots, \boldsymbol{R}_k$
11: **if** $k > t$ **then**
12: $\quad \boldsymbol{D}_\beta \leftarrow EasyRecover(t, w, [\alpha_1, ..., \alpha_k], [\boldsymbol{R}_1, \cdots, \boldsymbol{R}_k])$
13: **else if** Recovery fails **and** $\vartheta < k - \lfloor \sqrt{kt} \rfloor$ **then**
14: $\quad S_q \leftarrow \langle \boldsymbol{R}_1[q], \cdots, \boldsymbol{R}_k[q] \rangle, \forall q \in [1, s]$
15: $\quad \boldsymbol{D}_\beta \leftarrow HardRecover(t, w, [\alpha_1, ..., \alpha_k], [S_1, \cdots, S_s])$

Moreover, ϑ DBs among the k responding ones may even be *byzantine*, as in Definition 1, and produce incorrect response. In that case, it would be impossible for SU to simply rely on Lagrange interpolation to recover the correct responses. Since Shamir secret sharing is based on polynomial interpolation, the problem of recovering the response in the case of *byzantine* failures corresponds to noisy polynomial reconstruction, which is exactly the problem of decoding Reed-Solomon codes [37]. Thus, SU would rather rely on error correction codes and more precisely on the Guruswami-Sudan list decoding [38] algorithm which can correct $\vartheta < k - \lfloor \sqrt{kt} \rfloor$ incorrect responses. In fact, the

vector $\langle \boldsymbol{R}_1[q], \boldsymbol{R}_2[q], \cdots, \boldsymbol{R}_\ell[q] \rangle$ is a Reed-Solomon codeword encoding the polynomial $g_q = \sum_j f_j \boldsymbol{w}_{jq}$, and the client wishes to compute $g_q(0)$ for each $1 \leq q \leq s$ to recover all the s words forming the record $\boldsymbol{D}_\beta = \langle g_1(0), \cdots, g_s(0) \rangle$. This is done through the $HardRecover()$ subroutine from [30]. This makes $LP\text{-}Goldberg$ also ϑ-*Byzantine-robust*, by Definition 3, and solves the robustness issues that $LP\text{-}Chor$ suffers from, however, this comes at the cost of an additional overhead as we discuss in Section IV.

Corollary 1. *$LP\text{-}Chor$ and $LP\text{-}Goldberg$ directly inherit the security properties of Chor's [22] PIR and Goldberg's [30] PIR respectively.*

IV. Evaluation and Analysis

A. Analytical Comparison

We start by studying $LP\text{-}Chor$ and $LP\text{-}Goldberg$'s performance analytically and we compare them to existing approaches. For $LP\text{-}Goldberg$, we choose $w = 8$ to simplify the cost of computations as in [37]; since in $GF(2^8)$, additions are XOR operations on bytes and multiplications are lookup operations into a 64 KB table [37]. We summarize the system communication complexity and the computation incurred by both DB and SU and we illustrate the difference in architecture and privacy level of the different approaches in Table III. As we mentioned earlier, existing research focuses on the single DB setting. We compare $LP\text{-}Chor$ and $LP\text{-}Goldberg$ to these approaches despite the difference of architecture to show the great benefits that multi-server PIR brings in terms of performance and privacy as we discuss next. We briefly discuss these approaches in the following.

Gao et al. [2] propose a PIR-based approach, termed $PriSpectrum$, that relies on the PIR scheme of Trostle et al. [25] to defend against the new attack that they identify. This new attack exploits spectrum utilization pattern to localize SUs. Troja et al. [16], [17] propose two other PIR-based approaches that try to minimize the number of PIR queries by either allowing SUs to share their availability information with other SUs [16] or by exploiting trajectory information to make SUs retrieve information for their current and future positions in the same query [17].

Despite their merit in providing location privacy to SUs these PIR-based approaches incur high overhead especially in terms of computation. This is due to the fact that they rely on $cPIR$ protocols to provide location privacy to SUs, which are known to suffer from expensive computational cost. In fact, answering an SU's query through a $cPIR$ protocol, requires DB to process all of its records, otherwise DB would learn that SU is not interested in them and would then learn partial information about the record $\boldsymbol{D}_\beta$, and consequently SU's location. This makes the computational cost of most $cPIR$ based location preserving schemes linear on the database size from DB side as we illustrate in Table III. Now this is not exclusive to $cPIR$ protocols as even $itPIR$ protocols may require processing all the records to guarantee privacy, however, the main difference with $cPIR$ protocols is that the latter have a very large cost per bit in the database, usually involving expensive group operations like multiplication modulo a large modulus [24] as opposed to multi-server $itPIR$ protocols. This could be seen clearly in Table III as both $LP\text{-}Chor$ and $LP\text{-}Goldberg$ require DB to perform a very efficient XOR operation per bit of the database. The same applies to the overhead incurred by SU which only performs XOR operations in both $LP\text{-}Chor$ and $LP\text{-}Goldberg$, while performing expensive modular multiplications and even exponentiations over large primes in the $cPIR$-based approaches.

In terms of communication overhead, the proposed approaches incur a cost that is linear in the number of records r and their size b. As an optimal choice of these parameters is usually $r = b = \sqrt{n}$ [22], [24], [30], [37] then this cost could be seen as $\mathcal{O}(\sqrt{nw})$ to retrieve a record of size $\sqrt{nw}$ bits, which is a reasonable cost for an information theoretic privacy.

Moreover, as illustrated in Table III, existent approaches fail to provide information theoretic privacy as the underlying security relies on computational PIR schemes. The only approaches that provide information theoretic location privacy are $LP\text{-}Chor$ and $LP\text{-}Goldberg$ which are $(\ell-1)$-*private* and t-*private*, respectively, by Definition 2. It is worth mentioning that $PriSpectrum$ [2] relies on the well-known $cPIR$ of Trostle et al. [25] representing the state-of-the-art in efficient $cPIR$. However, this $cPIR$ scheme has been broken [24], [39]. Since the security of $PriSpectrum$ follows that of Trostle et al. [25] broken $cPIR$, then $PriSpectrum$ fails to provide the privacy objective that it was designed for. However, we include it in our performance analysis for completeness.

B. Experimental Evaluation

We further evaluate the performance of the proposed schemes experimentally to confirm the analytical observations.

Hardware setting and configuration. We have deployed the proposed approaches on GENI [33] cloud platform using the percy++ library [40]. We have created 6 virtual machines (VMs), each playing the role of a DB and they all share the same copy of $\boldsymbol{D}$. We deploy these GENI VMs in different locations in the US to count for the network delay and make our experiment closer to the real case scenario where spectrum service providers are located in different locations. These VMs are running Ubuntu 14.04, each having 8 GB of RAM, 15 GB SSD, and 4 vCPUs, Intel Xeon X5650 2.67 GHz or Intel Xeon E5-2450 2.10 GHz. To assess the SU overhead we use a Lenovo Yoga 3 Pro laptop with 8 GB RAM running Ubuntu 16.10 with an Intel Core m Processor 5Y70 CPU 1.10 GHz. The client laptop communicates with the remote VMs through ssh tunnels. We are also aware of the advances in $cPIR$ technology, and more precisely the fastest $cPIR$ protocol in the literature which is proposed by Aguilar et al. [24]. We include this protocol in our experiment to illustrate how multi-server PIR performs against the best known $cPIR$ scheme if it is to be deployed in CRNs. We use the available implementation of this protocol provided in [41] and we deploy its server component on a remote GENI VM while the client component is deployed on the Lenovo Yoga 3 Pro laptop.

Dataset. Spectrum service providers (e.g. Google, Microsoft, etc) offer graphical web interfaces and APIs to interact with their databases allowing to retrieve basic spectrum availability information for a user-specified location. Access to full data from real spectrum databases was not possible, thus, we generated random data for our experiment. The generated data consists of a matrix that models the content of the database, $\boldsymbol{D}$, with a fixed block size $b = 560$ kB while varying the number of records r. The value of b is estimated based on

TABLE III: Comparison with existent schemes

Scheme	Communication	Computation		Setting	Privacy
		DB	*SU*		
LP-Chor	$(r+b)\cdot\ell$	$nt_{\oplus}$	$(r+b)\cdot((\ell-1)\cdot t_{\oplus})$	ℓ *DBs*	$(\ell-1)$-*private*
LP-Goldberg	$r\cdot w\cdot\ell+k\cdot b$	$(n/w)\cdot t_{\oplus}$	$\ell\cdot(\ell-1)\cdot rt_{\oplus}+3\ell\cdot(\ell+1)t_{\oplus}$	ℓ *DBs*	t-*private* ℓ-*comp.-private*
PriSpectrum [2]	$(2\sqrt{r}+3)\cdot\lceil\log p\rceil$	$\mathcal{O}(r)\cdot Mulp$	$4\sqrt{r}\cdot Mulp$	1 *DB*	*underlying PIR broken*
Troja et al [17]	$12\delta\cdot b$	$\mathcal{O}(n)\cdot Mulp$	$4\sqrt{n}\cdot Mulp$	1 *DB*	*computationally-private*
Troja et al [16]	$n_g\cdot\pi\cdot\log_2 q+(2\sqrt{n}+3)\cdot\lceil\log p\rceil$	$\mathcal{O}(n)\cdot Mulp$	$n_g\cdot\pi\cdot(2Expp+Mulp)+4\sqrt{n}\cdot Mulp$	1 *DB*	*computationally-private*
XPIR [24]	$d\cdot(r/\alpha)^{1/d}\cdot\mathcal{C}+\lambda\cdot F^d\cdot b$	$2d\cdot(r/\alpha)\cdot(b/\ell_0)\cdot Mulp$	$d\cdot(r/\alpha)^{1/d}\cdot Enc+d\cdot\alpha\cdot b/\ell_0\cdot Dec$	1 *DB*	*computationally-private*

Variables: $t_{\oplus}$ is the execution time of one XOR operation. p is a large prime, and $Mulp$ and $Expp$ are the execution time of performing one modular multiplication, and one modular exponentiation respectively. π denotes the number of bits that an SU shares with other SUs in [16], n_g is the number of SUs within a same group in [16]. δ is the number of DB segments in [17]. d is the recursion level, α is the aggregation level, $\mathcal{C}$ is the Ring-LWE ciphertext size, λ is the number of elements returned by DB, F is the expansion factor of the Ring-LWE cryptosystem, ℓ_0 is the number of bits absorbed in a cyphertext, all are used in [24]. (Enc, Dec) are respectively the encryption and decryption cost for Ring-LWE cryptosystem used in [24].

the public raw data provided by FCC [42] on a daily basis and which service providers use to populate their spectrum databases.

Results and Comparison. We first measure the query end-to-end delay of the proposed approaches and plot the results in Fig. 3. We also include the delay introduced by the existing schemes based on our estimation of the operations included in Table III. The end-to-end delay that we measure takes into consideration the time needed by SU to generate the query, the network delay, the time needed by DB to process the query, and finally the time needed by SU to extract the β^{th} record of the database. We consider two different internet speed configurations in our experiment. We first rely on a high-speed internet connection of $80Mbps$ on the download and $30Mbps$ on the upload for all compared approaches. Then we use a low-speed internet connection of $1Mbps$ on the upload and download to assess the impact of the bandwidth on *LP-Chor* and *LP-Goldberg*, and also on XPIR as well.

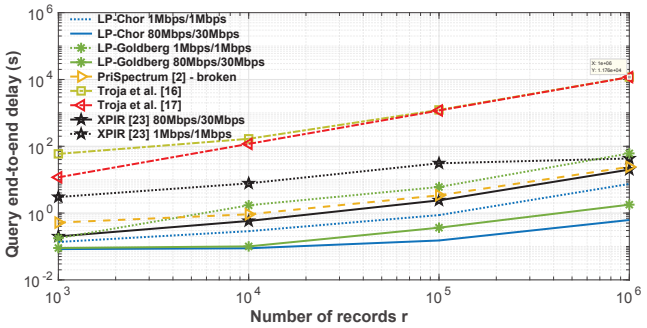

Fig. 3: Query RTT of the different PIR-based approaches

Fig. 3 shows that the proposed schemes perform much better than the existing approaches in terms of delay even with low-speed internet connection. They also perform better than the fastest existing *cPIR* protocol XPIR. This shows the benefit of relying on multi-server *itPIR* in multi-*DB CRN*s. Also, and as expected, *LP-Chor* scheme performs better than *LP-Goldberg* thanks to its simplicity. As we will see later, *LP-Goldberg* also incurs larger communication overhead than *LP-Chor* as well. This could be acceptable knowing that *LP-Goldberg* can handle collusion of up-to ℓ *DB*s, and is robust in the case of $(\ell-k)$ non-responding *DB*s, and ϑ byzantine *DB*s, as opposed to *LP-Chor*. This means that *LP-Goldberg* could be more suitable to real world scenario as failures and byzantine behaviors are common in reality. Fig. 3 also shows that the network bandwidth has a significant impact on the end-to-end latency. This is due to the relatively large amount of data that needs to be exchanged during the execution of these protocols which requires higher internet speeds.

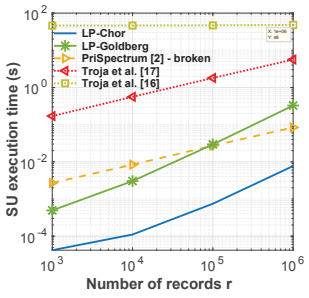

(a) SU Computation Overhead.

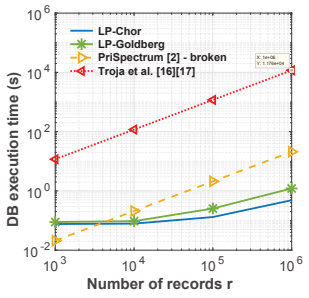

(b) DB Computation Overhead.

Fig. 4: Computation Comparison

We also compare the computational complexity experienced by each SU and DB separately in the different approaches as shown in Table III. We further illustrate this through experimentation and we plot the results in Fig. 4a, which shows that the proposed schemes incur lower overhead on the SU than the existing approaches. The same observation applies to the computation experienced by each DB which again involves only efficient XOR operations in the proposed schemes. We illustrate this in Fig. 4b.

We also study the impact of non-responding *DB*s on the end-to-end delay experienced by the SU in *LP-Goldberg* as illustrated in Fig. 5. This Figure shows that as the number of faulty *DB*s increases, the end-to-end delay decreases since SU needs to process fewer shares to recover the record $\boldsymbol{D}_{\beta}$. As opposed to *LP-Chor*, in *LP-Goldberg*, SU is still able to recover the record β even if only k out-of-ℓ *DB*s respond. Please recall also that our experiment was performed on resource constrained VMs to emulate *DB*s, however in reality, *DB*s should have much more powerful computational resources than of those of the used VMs which will have a tremendous impact on further reducing the overhead of the proposed approaches.

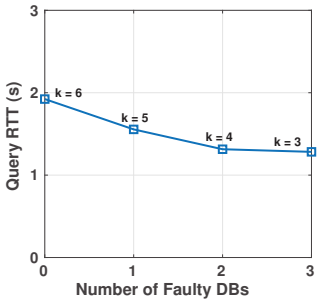

Fig. 5: Impact of the number of faulty DBs on the query RTT.

In terms of communication overhead, most of the approaches, including $LP\text{-}Chor$ and $LP\text{-}Goldberg$, have linear cost in the number of records in the database as shown in Table III. What really makes a difference between these schemes' communication overheads is the associated constant factor which could be very large for some protocols. Based on our experiment and the expressions displayed in Table III, we plot in Fig. 6, the communication overhead that the CRN experiences for each private spectrum availability query issued by SU for the different schemes. The scheme with the lowest communication overhead is that of Troja et al. [17] thanks to the use of Gentry et al. PIR [32] which is the most communication efficient single-server protocol in the literature having a constant communication overhead. However this scheme is computationally expensive just like most of the existing $cPIR$-based approaches as we show in Fig. 3. $LP\text{-}Chor$ is the second best scheme in terms of communication overhead but it also provides information theoretic privacy. As shown in Fig. 6, $LP\text{-}Chor$ incurs much lower communication overhead than $LP\text{-}Goldberg$ thanks to the simplicity of the underlying Chor PIR protocol. However, as we discussed earlier, $LP\text{-}Goldberg$ provides additional security features compared to $LP\text{-}Chor$. XPIR has a relatively high communication overhead especially for smaller database size but its overhead becomes comparable to that of $LP\text{-}Goldberg$ when the database's size gets larger as shown in Fig. 6. This could be a good alternative to the $cPIR$ schemes used in the context of CRNs especially that it introduces much lower latency which is critical in the context of CRNs. Still, the proposed approaches have better performance and also provide information-theoretic privacy to SUs, which shows their practicality in real world.

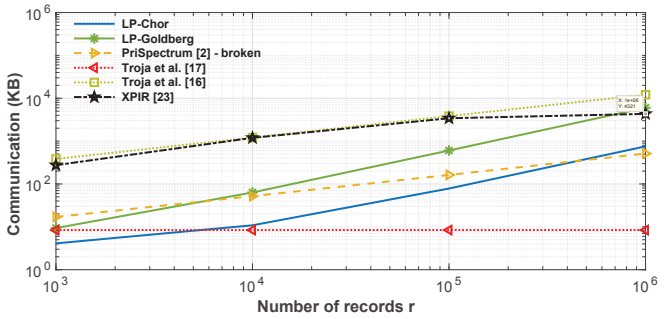

Fig. 6: Comparison of the communication overhead of the different approaches: $b = 560$ B, $k = \ell$, $\vartheta = 0$.

V. Related Work

There are other approaches that address the location privacy issue in database-driven CRNs. However, for the below mentioned reasons we decided not to consider them in our performance analysis. For instance, Zhang et al. [15] rely on the concept of *k-anonymity* to make each SU queries DB by sending a square cloak region that includes its actual location. *k-anonymity* guarantees that SU's location is indistinguishable among a set of k points. This could be achieved through the use of dummy locations by generating $k-1$ properly selected dummy points, and performing k queries to DB, using the real and dummy locations. Their approach relies on a tradeoff between providing high location privacy level and maximizing some utility. This makes it suffer from the fact that achieving a high location privacy level results in a decrease in spectrum utility. However, *k-anonymity*-based approaches cannot achieve high location privacy without incurring substantial communication/computation overhead. Furthermore, it has been shown in a recent study led by Sprint and Technicolor [23] that anonymization based techniques are not efficient in providing location privacy guarantees, and may even leak some location information. Grissa et al [19], [43] propose an information theoretic approach which could be considered as a variant of the trivial PIR solution. They achieve this by using set-membership probabilistic data structures/filters to compress the content of the database and send it to SU which then needs to try several combinations of channels and transmission parameters to check their existence in the data structure. However, LPDB is only suitable for situations where the structure of the database is known to SUs which is not always realistic. Also, LPDB relies on probabilistic data structures which makes it prone to false positives that can lead to erroneous spectrum availability decision and cause interference to PU's transmission. Zhang et al. [18] rely on the *ϵ-geo-indistinguishability* mechanism [44], derived from *differential privacy* to protect bilateral location privacy of both PUs and SUs, which is different from what we try to achieve in this paper. This mechanism helps SUs obfuscate their location, however, it introduces noise to SU's location which may impact the accuracy of the spectrum availability information retrieved.

VI. Conclusion

In this paper, with the key observation that database-driven CRNs contain multiple synchronized DBs having the same content, we harnessed multi-server PIR techniques to achieve an optimal location privacy for SUs with high efficiency. Our analytical and experimental analysis indicate that our adaptation of multi-server PIR for database-driven CRNs achieve magnitudes of time faster end-to-end delay compared to the fastest state-of-the-art single-server PIR adaptation with an information theoretical privacy guarantee. Specifically, we adapted two multi-server PIR techniques into CRN settings as $LP\text{-}Chor$ and $LP\text{-}Goldberg$. $LP\text{-}Chor$ achieves an end-to-end delay below a second with high collusion resiliency, while $LP\text{-}Goldberg$ offers fault tolerance and byzantine robustness with a significantly higher efficiency compared to single-server PIR based approaches. Given the demonstrated benefits of multi-server PIR approaches without incurring any extra architectural overhead on database-driven CRNs, we hope this work will provide an incentive for the research community to consider this direction when designing location privacy preservation protocols for CRNs.

REFERENCES

[1] J. Mitola and G. Q. Maguire, "Cognitive radio: making software radios more personal," *IEEE personal comm.*, vol. 6, no. 4, pp. 13–18, 1999.

[2] 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*, 2013, pp. 2751–2759.

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

[4] "Google spectrum database," https://www.google.com/get/spectrumdatabase/, accessed: 2017-04-14.

[5] "iconectiv white spaces database," https://spectrum.iconectiv.com/main/home/, accessed: 2017-04-14.

[6] "Microsoft white spaces database," http://whitespaces.microsoftspectrum.com/, accessed: 2017-04-14.

[7] A. Mancuso, S. Probasco, and B. Patil, "Protocol to access white-space (paws) databases: Use cases and requirements," Tech. Rep., 2013.

[8] M. Grissa, B. Hamdaoui, and A. A. Yavuz, "Location privacy in cognitive radio networks: A survey," *IEEE Communications Surveys Tutorials*, vol. PP, no. 99, pp. 1–1, 2017.

[9] S. Li, H. Zhu, Z. Gao, X. Guan, K. Xing, and X. Shen, "Location privacy preservation in collaborative spectrum sensing," in *INFOCOM, 2012 Proceedings IEEE*. IEEE, 2012, pp. 729–737.

[10] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "Lpos: Location privacy for optimal sensing in cognitive radio networks," in *Global Communications Conference (GLOBECOM), 2015 IEEE*. IEEE, 2015.

[11] W. Wang and Q. Zhang, "Privacy-preserving collaborative spectrum sensing with multipleservice providers," *Wireless Communications, IEEE Transactions on*, 2015.

[12] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "An efficient technique for protecting location privacy of cooperative spectrum sensing users," in *INFOCOM WKSHPS*. IEEE, 2016.

[13] ——, "Preserving the location privacy of secondary users in cooperative spectrum sensing," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 2, pp. 418–431, 2017.

[14] S. Liu, H. Zhu, R. Du, C. Chen, and X. Guan, "Location privacy preserving dynamic spectrum auction in cognitive radio network," in *ICDCS*. IEEE, 2013, pp. 256–265.

[15] 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*.

[16] 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.

[17] ——, "Efficient location privacy for moving clients in database-driven dynamic spectrum access," in *ICCCN*. IEEE, 2015.

[18] 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.

[19] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "Location privacy preservation in database-driven wireless cognitive networks through encrypted probabilistic data structures," *IEEE Transactions on Cognitive Communications and Networking*, vol. 3, no. 2, pp. 255–266, 2017.

[20] M. Gruteser and D. Grunwald, "Anonymous usage of location-based services through spatial and temporal cloaking," in *Proceedings of the 1st international conference on Mobile systems, applications and services*. ACM, 2003, pp. 31–42.

[21] C. Dwork, "Differential privacy: A survey of results," in *International Conference on Theory and Applications of Models of Computation*. Springer, 2008, pp. 1–19.

[22] B. Chor, E. Kushilevitz, O. Goldreich, and M. Sudan, "Private information retrieval," *J. ACM*, vol. 45, no. 6, pp. 965–981, Nov. 1998.

[23] H. Zang and J. Bolot, "Anonymization of location data does not work: A large-scale measurement study," in *Proc. of the 17th annual int'l conf. on Mobile computing and networking*. ACM, 2011, pp. 145–156.

[24] C. Aguilar-Melchor, J. Barrier, L. Fousse, and M.-O. Killijian, "Xpir: Private information retrieval for everyone," *Proceedings on Privacy Enhancing Technologies*, vol. 2, pp. 155–174, 2016.

[25] J. Trostle and A. Parrish, "Efficient computationally private information retrieval from anonymity or trapdoor groups," in *International Conference on Information Security*. Springer, 2010, pp. 114–128.

[26] "White space database administrator group database-to-database synchronization interoperability specification," FCC, Tech. Rep., 2012.

[27] F. (2012), "TVWS database system requirements and tests," https://transition.fcc.gov/oet/whitespace/guides/TVWS_Database_Tests4.doc.

[28] R. Ramjee, S. Roy, and K. Chintalapudi, "A critique of fcc's tv white space regulations," *GetMobile: Mobile Computing and Communications*, vol. 20, no. 1, pp. 20–25, 2016.

[29] "White space database administrators guide," https://www.fcc.gov/general/white-space-database-administrators-guide, FCC, accessed: 2017-04-14.

[30] I. Goldberg, "Improving the robustness of private information retrieval," in *Security and Privacy, 2007. IEEE Symp. on*, pp. 131–148.

[31] L. Chen, S. Jordan, Y. Liu, D. Moody, R. Peralta, R. Perlner, and D. Smith-Tone, "Report on post-quantum cryptography. nistir 8105," 2016.

[32] C. Gentry and Z. Ramzan, "Single-database private information retrieval with constant communication rate," *Automata, Languages and Programming*, pp. 103–103, 2005.

[33] 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.

[34] A. Beimel and Y. Ishai, "Information-theoretic private information retrieval: A unified construction," in *International Colloquium on Automata, Languages, and Programming*. Springer, 2001, pp. 912–926.

[35] C. A. Melchor and P. Gaborit, "A fast private information retrieval protocol," in *ISIT 2008*. IEEE, pp. 1848–1852.

[36] A. Shamir, "How to share a secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, 1979.

[37] C. Devet, I. Goldberg, and N. Heninger, "Optimally robust private information retrieval." in *USENIX Security Symp.*, 2012, pp. 269–283.

[38] V. Guruswami and M. Sudan, "Improved decoding of reed-solomon and algebraic-geometric codes," in *Foundations of Computer Science, 1998. Proceedings. 39th Annual Symposium on*. IEEE, 1998, pp. 28–37.

[39] T. Lepoint and M. Tibouchi, "Cryptanalysis of a (somewhat) additively homomorphic encryption scheme used in pir," in *Int'l Conf. on Financial Cryptography and Data Security*. Springer, 2015, pp. 184–193.

[40] "Percy++ library," http://percy.sourceforge.net, accessed: 2017-04-14.

[41] "Xpir implementation," https://github.com/XPIR-team/XPIR, accessed: 2017-04-14.

[42] "Cdbs data," https://transition.fcc.gov/Bureaus/MB/Databases/cdbs/, accessed: 2017-04-20.

[43] M. Grissa, A. A. Yavuz, and B. Hamdaoui, "Cuckoo filter-based location-privacy preservation in database-driven cognitive radio networks," in *Computer Networks and Information Security (WSCNIS), 2015 World Symposium on*. IEEE, 2015, pp. 1–7.

[44] 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.