Scaling Network Information Services to Support HetNets and Dynamic Spectrum Access

Esa Piri and Henning Schulzrinne

Abstract: **Wireless network information services allow end systems to discover heterogeneous networks and spectrum available for secondary use at or near their current location, helping them to cope with increasing traffic and finite spectrum resources. We propose a unified architecture that allows end systems to find nearby base stations that are using either licensed, shared or unlicensed spectrum across multiple network operators. Our study evaluates the performance and scalability of spatial databases storing base station coverage area geometries. The measurement results indicate that the current spatial databases perform well even when the number of coverage areas is very large. A single logical spatial database would likely be able to satisfy the query load for a large national cellular network. We also observe that coarse geographic divisions can significantly improve query performance.**

Index Terms: **Heterogeneous networks, information services, spatial databases, spectrum sharing.**

I. INTRODUCTION

Increasing traffic demands and user mobility pose challenges to current cellular networks. It is widely anticipated that data traffic will continue to grow substantially [1], [2]. However, the radio spectrum reserved for mobile networks is limited and strictly regulated, and insufficient spectrum availability already restricts deployments of wireless networks.

Both researchers and standardization organizations have advocated the use of handovers between heterogeneous networks ("HetNets") to better use available spectrum and wired infrastructure resources. IEEE has specified the 802.21 Media Independent Handover (MIH) Services standard [3] for facilitating heterogeneous handovers. 3GPP has published the Access Network Discovery and Selection Function (ANDSF) [4] to allow cellular network users to seamlessly switch to a non-3GPP network in range that satisfies the required service level. Both standards specify an information service, namely, to provide access to relatively static information about networks, base stations (BSs) and capabilities the network supports. Overall, this information provides a good basis for heterogeneous mobility management, but the information services can also help use cases beyond handovers. To improve radio spectrum efficiency, spectrum sharing has been proposed and promoted [5], [6], where legacy users share spectrum bands with secondary users, but it has not yet been widely implemented.

Currently, wireless end systems typically listen for beacons and network announcements in each band, and then associate with one or possibly more suitable BSs, i.e., operated by the user's cellular carrier or a roaming partner. As the number of cellular bands and IEEE 802.11 access points increases and as multiple network operators share their BS infrastructure, this scanning model is no longer efficient. The increasing number of bands and BS choices is fortunately accompanied by the ability of end systems to know their precise geographic location, often within a few meters. Thus, we evaluate a model where mobile devices query databases to find suitable additional bands and BSs that cover their current location, or are close by. Close-by BSs are useful as future handover targets or may afford opportunities to defer transmission of delay-tolerant data until the mobile node reaches less congested or cheaper wireless coverage. All of these applications can be abstracted to the same geographic query model, namely finding a set of BSs and their parameters based on the current location and other attributes, such as current price or roaming access. As we point out later, the number of such coverage regions can be very large, and devices may need to query the database frequently as they move. Thus, our goal is to determine whether such databases can scale to a national wireless network, and how to improve scalability.

The rest of this paper is organized as follows. In Section II we relate our study to other information service and spatial database performance studies. We present the coverage area information service with use cases in Section III. Section IV presents our results and Section V concludes this paper.

II. RELATED WORK

Information services for providing a variety of information about networks have not been thoroughly studied. To the best of our knowledge, [7] is the only available study that exploits an information service enhanced with an information element about cell coverage areas. The paper [7] studies how the cell coverage area information can help cell selection in a dense HetNet environment, whereas this paper presents the coverage area information service in more details, considering its architecture and use cases, and evaluating its scalability to support a large number of cells. The studies of [7], briefly introduced later, are based on the results of this paper.

Despite the scarce number of information service studies, there are a few studies focusing on the IEEE 802.21 information service, such as [8]. They present and evaluate a hierarchical IEEE 802.21 architecture for the Media Independent Information Service (MIIS). As its basic hierarchical structure, the architecture is similar to the one presented in this study as an example architecture for the mobile network operator (MNO) maintained coverage area enhanced information services. However, MIIS, as specified in the standard, provides only relatively static information about networks and their BSs, whereas the coverage

E. Piri is with VTT Technical Research Centre of Finland (e-mail: esa.piri@vtt.fi).

H. Schulzrinne is with the Department of Computer Science, Columbia University, New York, NY USA (email: hgs@cs.columbia.edu).

area information often requires frequent database updates, posing also challenges for the service architecture. The expected query load and the database update overhead argues for a more centralized database solution at the MNO's core network level, where a single or several consistent database servers can handle the queries received without forwarding to another database server in the operator's core network and without excessive forwarding of queries to regional databases that reside closer to the access networks. Thus, the hierarchical architecture of [8] is not well-suited for managed dynamic spectrum access and sharing across MNOs.

Corujo et al. present an open source IEEE 802.21 implementation in [9] and they observe that in a Proxy Mobile IPv6 (PMIPv6) based handover scenario most of the information exchanged during the overall handover procedure is related to information service queries. They propose using filtering mechanisms to decrease this overhead. The cell coverage area database proposed in this paper improves the determination of the cells within and nearby the mobile device and their operational range. This information would enable better directing of the information queries to the most likely handover targets, decreasing the overall information amount returned. In [10], the processing performance of query handling in an MIIS server between the supported message encodings of IEEE 802.21, namely Type-Length-Value (TLV) and RDF/XML, is compared. The author observes that TLV encoding is recommended to be used with resource limited information servers. Both encoding methods specified in IEEE 802.21 are also suitable for delivering coverage area related information. The 3GPP ANDSF and IEEE 802.21 are compared for inter-system handovers in [11]. The information services in both standards are observed as similar, but the 3GPP ANDSF allows MNOs to deliver inter-system mobility policies for the mobile devices to better select the access points supported and recommended by the MNO to end-users. These policies could, for example, determine the usage of differently sized cells for different mobility types and speeds, as introduced later in this paper.

The scalability of spatial databases has not been evaluated in depth, either. Performances of various spatial databases have been studied by Ray et al. in [12] and Zhou et al. in [13]. The main contribution of both papers is to introduce a spatial database benchmark for performance evaluation. They measure different spatial operations for various geometries, considering points, lines and polygons. However, the studies do not evaluate the database performances in terms of database size, but focus more on comparing the performances of different databases in scenarios emulating various applications. Both papers measure the performance of MySQL and PostgreSQL with PostGIS, which are also used in this study. In [12], MySQL was observed to perform better than PostgreSQL in the "PointWithinArea" spatial join queries, resulting in on average two times lower query times for the set of database scenarios measured. The query time results reported in [13] show the opposite, PostgreSQL performs slightly better in the same spatial join query. However, the paper [13] does not present the database versions used in the experiments and the introduction of measurement scenarios is also insufficient to allow making any comparison between the results of these studies. The "PointWithinArea"

type of query is used in this study to find polygon areas covering the inquirer's location. However, we use a newer MySQL version than those used in [12] and [13], which more fully supports the OpenGIS geometry published by the Open Geospatial Consortium, yielding substantially better MySQL performance.

BS coverage areas can form very complex polygons to be stored in the database. Due to the characteristics of radio propagation, the obtainable coverage areas always vary more or less. Thus, the polygons stored in the database can be somewhat simplified. Song et al. have presented an algorithm for polygon simplification in [14], reducing the total number of vertices in the polygons.

III. INFORMATION SERVICE WITH COVERAGE AREAS

In the network design phase, the coverage area of each deployed BS is estimated in detail, first through simulations and then the accuracy is improved with field measurements. Even complex coverage areas can be depicted as polygon geometries and stored in geospatial databases. Fig. 1 illustrates the overview of the network information service, where the server maintains a database storing BS coverage polygons. The inquirer sends its current location to the server along with the information about the parameters of interest, obtaining in return BSs within range and nearby.

Due to the varying sensitivity of mobile devices, the coverage area stored in the information services could comprise an area with relatively strong signal strength, e.g., -80 dBm, or possibly several with different strength levels. Setting the coverage area signal strength limit relatively high improves the reliability of the information. For example, temporary obstructions, varying fast fading, and different antenna configurations in the mobile stations vary the actual coverage area at different times and for different devices. However, whether a single coverage area with a certain signal strength threshold per cell can provide sufficient accuracy for different mobile devices and use cases is an interesting future work item. For mobile handsets, the coverage areas can likely be considered as rather similar. Bigger portable devices with high gain antennas obtain larger coverage areas, but their usage of the information service may not be as dynamic as that of smaller handsets. Effectively, defining more than one coverage area per cell for different kind of end systems should not cause problems with respect to the storage needed for the database and also the query times remain similar, as the number of cells per database table does not change. A number of other useful network and BS parameters relevant for network information services are defined in, for instance, IEEE 802.21 and ANDSF, mostly related to frequency band, location, network address, and capability information.

A. Service Architecture

The number of polygons in the database depends on the deployment architecture. For mobility, MNOs may maintain their own database. However, MNOs should collaborate to allow inter-MNO queries and mobility. In addition, virtual operators known as MVNOs may in the future leverage the resources of multiple facilities-based MNOs.
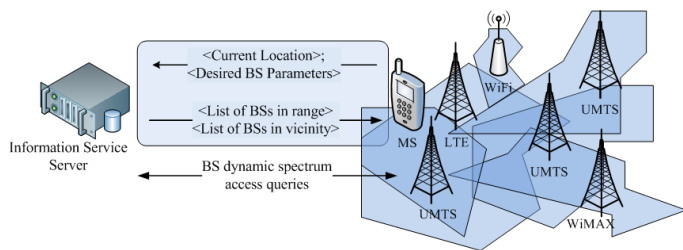
Fig. 1. Overview of coverage area information service.

Fig. 2. Example information service architecture.

An example MNO-maintained hierarchical information service architecture is shown in Fig. 2. MNOs may divide their operational geographical area into smaller regions, each including a network information service responsible for maintaining an information database about the BSs within the region. When a mobile device connects to the network, the closest information service can be assigned to it as a home server. Most of the queries are, thus, handled solely by its respective information service server, Region 1 IS in Fig. 2, and the database size in the regional servers does not necessarily grow very large. However, regional information services need coordination for frequent coverage area updates and the inter-MNO usage increases the overall number of BS queries. In order to avoid excessive query forwardings, each MNO has a Root IS server (or possibly several consistent servers) which not only knows the geographical coordinates of each region but also keeps a consistent centralized database about the MNO's BS infrastructure. For example, when the mobile device moves near the border of Region 1, towards Region 2, the information queries regarding the BSs within the neighboring region needs to be forwarded to the root server. Instead of forwarding the query to the Region 2 IS server, the root information service could answer the queries directly. This solution can be seen as a compromise between attempts to minimize the signaling overhead for the network traffic and to optimize the query load. At the same time, root servers can also act as coordinators for the consistency of the regional databases, facilitating the information service maintenance. To further optimize the query load distribution, the home information service of the mobile device can be changed when moving into the region of another server. A mapping between the location of the mobile device and the service boundary is common for many location-based services and there are also protocols for such operations, for example, Location-to-Service (LoST) [15]. LoST addresses also the ability to federate different location databases.

In inter-MNO queries, the MNOs having mutual roaming agreements need only to provide access to the root information service servers.

For dynamic spectrum sharing, the need for inter-MNO consistency and rapid updates argue for a single logical database. To manage dynamic spectrum access and spectrum sharing, a nation-wide database for the spectrum usage and license information, possibly maintained by a national regulatory authority, would ease the management of dynamic spectrum access and spectrum sharing. Moreover, searching for BSs of all MNOs may often require excessive amount of queries, advocating also a single logical database, shown as Inter-MNO IS in Fig. 2.
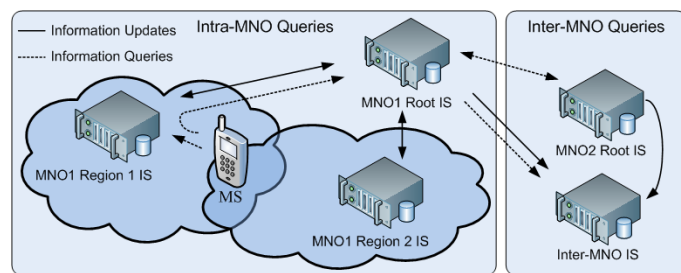
However, rather than specifying a service architecture, the goal of this paper is to study whether such large centralized information service databases based on BS coverage area polygons are feasible.

The database size rapidly grows when the database stores coverage areas from a large area and, especially, from numerous MNOs. For example, according to CTIA's statistics[1], there were over 300,000 commercial wireless cell sites in the USA in December 2012, a 41% increase from 2007. In the near future, the average cell size is expected to decrease through the increasing deployment of femto, micro, and picocells, increasing significantly the total number of cells. When factoring in the access points operating on ISM bands, largely based on the IEEE 802.11 Wireless Local Area Networks (WLAN) technology, the number would dramatically rise. Even if private WLAN networks are excluded, the number of public WLAN hotspots, either free of charge or chargeable, increases the number of cells significantly. The Wireless Broadband Alliance predicts that the number of public hotspots will rise to 5.8 million by 2015 worldwide, up a 350% from 2011 [16]. Traffic offloading has recently attracted the interest of both MNOs and network equipment manufacturers, causing commercial ISM band hotspots to become more popular. For example, 3GPP has specified WLAN interoperability with 3GPP-specified radio and network technologies [17], which allows MNOs to have a consistent way for load balancing through offloading. Thus, the scalability of a database storing BS coverage areas is of crucial importance of this kind of information service.

The database must be dynamic. For instance, the coverage area polygons and frequency bands must be updated frequently, based on end system performance measurements and reflecting infrastructure changes. The coverage areas could be refined, for example, through measurement tools available for end-users' mobile devices measuring the speed of the Internet connection used. Those applications could be developed to also measure the signal strength, such as Portolan [18] and Netradar[2], and to report the results along with the location and the BS ID to the information service database. Naturally, the different mobile devices and the positions in which the devices are held affect the results, but still they could provide valuable data for evaluating the current coverage area in the database and updating it to better describe the prevailing situation. The mobile devices keeping a local information database based on the previous queries should frequently query for an update to the coverage areas.

---

[1] www.ctia.org/your-wireless-life/how-wireless-works/wireless-quick-facts
[2] www.netradar.org

B. Use Case: Dynamic Spectrum Sharing and Access

Due to the limited spectrum allocated for broadband mobile access, the cost of spectrum has soared in some countries. Making new spectrum available for mobile networks is challenging. Often, there is no unallocated spectrum available and reallocating spectrum for broadband use is costly and time consuming. Various regulatory authorities are working on regulations to manage and promote spectrum sharing. For example, the Federal Communications Commission (FCC) has initiatives regarding shared spectrum in the TV white space [19] and the 3.5 GHz band [20]. Moreover, a spectrum sharing demonstration utilizing an industry-driven Authorized Shared Access (ASA) concept is introduced in [21]. The database scalability results presented in this study can be applied to the ASA concept [22], where an ASA repository, indicating the areas allowing spectrum sharing, is based on a geolocation database. In principle, ASA can be perceived as a special case for the utilization of the cell coverage area database, targeted for managing spectrum sharing usage and licensing.

We argue that a centralized database for storing the coverage areas of BSs would ease the management of spectrum sharing and accelerate its deployment. Through the database, geographic exclusion zones for dynamic spectrum access can be found. With additional information related to licensing, possible spectrum allowing secondary usage within range can be discovered. Network entities managing dynamic spectrum access capable BSs can query for spectrum usage in the operational range of their BSs. In order to avoid excessive query forwarding, the spectrum management entities in different locations could be provided either with direct access to the inter-MNO database, or via MNOs central information services. In case of spectrum change in the BS, the new frequency band and possibly an estimated new coverage area can be updated to the database. Overall, a centralized coverage area database would facilitate and improve the authorized controlling of dynamic frequency band allocation for BSs through spectrum loaning or agreeing on shared use of occupied channels as a secondary service.

C. Use Case: Mobility and Network Load Balancing

Current information services have often been developed for the purpose of optimizing seamless mobility and network load balancing. Whether the handover is mobile- or network-initiated, the information services, overall, aim to assist in preparing a handover by finding BSs and providing the necessary information about them for selecting the most suitable target BS. Moreover, they allow adapting the ongoing network services according to the characteristics of the target BS prior to performing the handover. However, the operational ranges of different BSs are unknown or can only be roughly estimated based on the information provided by the current information services, which causes non-optimal handovers. For example, if the mobile device is moving fast, as in the case of vehicular mobility, use of small cells cause excessive number of handovers compared to favoring the use of large macrocells when available. Simulation results with a cell coverage area database presented in [7] show that the number of handovers can be even cut to half when selecting large macrocells located nearby the mobile de-vice instead of the nearest cell. If the traversed route is known beforehand, e.g., from a road navigator, the BSs on the route could even be assessed and possibly selected beforehand. This can potentially decrease the number of handovers even more. As the future network environment will likely be increasingly composed of small cell BSs, the number of BSs nearby a mobile device can be very large, especially in densely populated areas. In order to improve the cell selection, [7] proposes a clustering method for grouping the BS cells found, based on the distance to the BS and the cell area. This enables directing the more detailed assessment of the handover target choices only to those cells that provide sufficiently large coverage for high speed mobility. Moreover, for low speed mobility, small cells may often provide the highest capacity due to a smaller number of users within their coverage area. From the load balancing point of view, also MNOs likely prefer using small cells, such as WLAN access points and picocells, for stationary and pedestrian usage.

Network entities responsible for offloading decisions in networks should be able to reliably find the best potential target access point for the mobile station to be moved to another BS. If all BSs within the range and close-by the mobile station are known, the network is not required to ask the mobile station to scan for target BSs. The handover target can in most cases be assessed based on the information received from the network information service. Moreover, finding close-by BSs and suitable BSs from other MNOs to allow network sharing are not properly handled by current information services. Knowledge of BS coverage areas would confer additional information for mobility decision entities to make more proactive and better decisions.

IV. PERFORMANCE EVALUATION

The main contribution of this paper is to quantify the performance of spatial extensions of two commonly used relational databases, MySQL and PostgreSQL with PostGIS. Spatial extensions enable generating, storing, and analyzing geographical features. A polygon geometry model allows storing coverage areas in databases, where the polygon vertices are given as geographical coordinates, handled in World Geodetic System (WGS) 84 format in our study. For our evaluation, we created uniform polygons, each with a radius of 2000 meters and randomly located across the United States. In reality, the cells are not uniformly distributed as assumed in these measurements. However, rather than defining different cell distribution scenarios specific to certain environments, the random polygon locations are sufficient for our database scalability evaluation. The evaluation results do not depend on the size of the polygon, but our polygons are meant to roughly represent a realistic macrocell BS. Nevertheless, some measurements with polygons with shorter radius lengths, representing pico and microcells, have been carried out for a query time comparison.

Each query was performed with a different random coordinate denoting the location of the inquirer. One query consists of two database queries; the first one searches for polygons covering the inquirer's location and another one searches for polygons that are near, but the inquirer is not within the covered area. We defined the range for close-by BS polygons to be about 350 meters. The vicinity query differs between the databases as MySQL does not support a function directly returning the

polygons within a particular distance. To emulate PostgreSQL's *ST_DWithin*, we used the *ST_Intersects* function with lines of length 350 meters from inquirer.

First we measured how the query time changed as a function of the number of polygons. The query time indicates only the delay of executing the query, not the time taken by the data processing in the server. The queries were repeated 100 times and the average value is shown. Then we measured query load with parallel query processes constantly polling the database over a persistent database connection. We increased the number of querying processes until the system saturated. Each measurement run lasted 30 seconds.

To improve the query performance, we replace a single query to a large nationwide database with a hierarchical query, dividing the United States into either just states, or states and counties. Most mobile nodes will remain within the same state or county for extended periods of time. This regional division substantially decreases the table size and, thus, reduces query delay and increases throughput. For the hierarchical query measurements, we simply add database tables with the state and county polygons to the system.

The measurements were conducted on an Ubuntu Server 12.04 64-bit virtual machine, with 8 GB of memory, four cores and 40% share of the total CPU of a VMware VSphere 5 host. We opted to use a MySQL development version (5.6.6-m9) for better spatial support and the MyISAM indexing engine. The PostGIS extender version of 1.5.3 and the GIST indexing engine were used with the PostgreSQL (v. 9.1.5). We verified the results on another server running on physical hardware and having 32 GB memory, two quad core CPUs, 3.6 GHz each, resulting in 16 threads in total. In addition, to validate the results on a different server machine, we also measured the newer PostGIS version of 2.0.1. Otherwise, the measurement setup and the measurements were similar.

The databases on both machines used the default configuration parameters. Thus, the goal of this study is not to do a performance comparison between the measured databases, but to evaluate how current spatial databases scale with large numbers of polygons. Database queries were carried out with a Perl script using a DBI module as the database interface. The query times were measured using the *gettimeofday()* function.

A. Results: Query Times

Fig. 3 shows the query times with different numbers of 20 edge polygons. The growth trend is observed to be logarithmic. The MySQL total query times stay below 1 millisecond (ms) even with 600,000 polygons. Overall, MySQL performs better than PostgreSQL. The total PostgreSQL query times reach 1.5 ms already with 10,000 polygons, which is almost five times that of MySQL. However, as the number of polygons increases the performance difference narrows. When the PostgreSQL database stores 300,000 polygons or more, the total query times are on average three times higher than those of MySQL with the same numbers of polygons. In the verification tests on the physical server, we observed that PostgreSQL with the newer PostGIS performs rather similarly compared to MySQL. However, we observed some performance differences with the smaller numbers of polygons measured. Unlike on the virtual server, the difference between the total query times with the newer PostGIS and MySQL stays rather stable regardless of the number of polygons in the database. The PostgreSQL total query times are almost three times that of MySQL with each number of polygons in the database. Thus, the newer PostGIS is observed to perform better than the older version with 10,000 and 100,000 polygons in the database when the query times are related to those of MySQL. This can also be observed by comparing the query times between the two servers. The PostgreSQL query times on the physical server are about 30% higher than those shown in Fig. 3 with each number of polygons. With 10,000-300,000 polygons, the MySQL query times on the physical server are over 50% higher than shown in Fig. 3 for the MySQL database, being doubled with 10,000 polygons. The differences in hardware configurations in the measurement servers explain why the query times on the physical server are overall higher. However, the query delays are sufficiently small that they are unlikely to impede network attachment and handoff decisions.

The impact of polygon edges on the total query time was also measured with polygons of 8, 13, and 20 edges. With MySQL, there is no substantial difference observed below 300,000 polygons. With 300,000 polygons and above, the database queries storing 20 edge polygons take on average more time. The difference between 13 edge and 20 edge polygons ranges from 6% to 16% between 300,000 and 1 million (1M) polygons in the database. However, the standard deviation with 20 edge polygons, 0.44 ms with 1M polygons, is not notably different than that with 8 and 13 edge polygons. Each query was executed with a different random inquirer's location and the server machine ran also other native Linux processes that slightly affect the CPU use. These bring minor variation to the results. For example, the median MySQL query time is 8% and 15% lower than shown in Fig. 3 with 300,000 and 1M polygons, respectively. With PostgreSQL, the median values are 2-4% lower than the average. With PostgreSQL, 20 edge polygons increase the total query times by, on average, 15% compared to polygons with 8 and 13 edges. The query times with polygons of 8 edges and 13 edges are similar, until with 1M polygons, the 8 edge polygons are 5% faster.

In reality, cell sizes vary greatly. When the radius of polygons is significantly shorter than 2000 meters, the average query times can be even tens of percent shorter than those presented earlier with the high numbers of polygons in the database. With the polygon radius of 500 meters, and 10,000 and 100,000 polygons, the query times are observed to be about 5-10% lower than those shown in Fig. 3 for both MySQL and PostgreSQL. With 300,000 polygons and more, the query times are about 20% shorter, and even more with MySQL, with the 500 meter polygon radii. However, with 1M polygons, the PostgreSQL query times are close to those with the 2000 meter radius. Shortening the radius to 100 meters, the MySQL query times can decrease even over 30% when the database stores 300,000 polygons or more. In PostgreSQL, we did not observe an additional decrease in the total query times. Anyway, the average query times can be even shorter than shown in Fig. 3 when the database stores more small cell coverage areas typical of future mobile networks.

The same number of cells in a significantly smaller measurement area than the area across the United States was not found
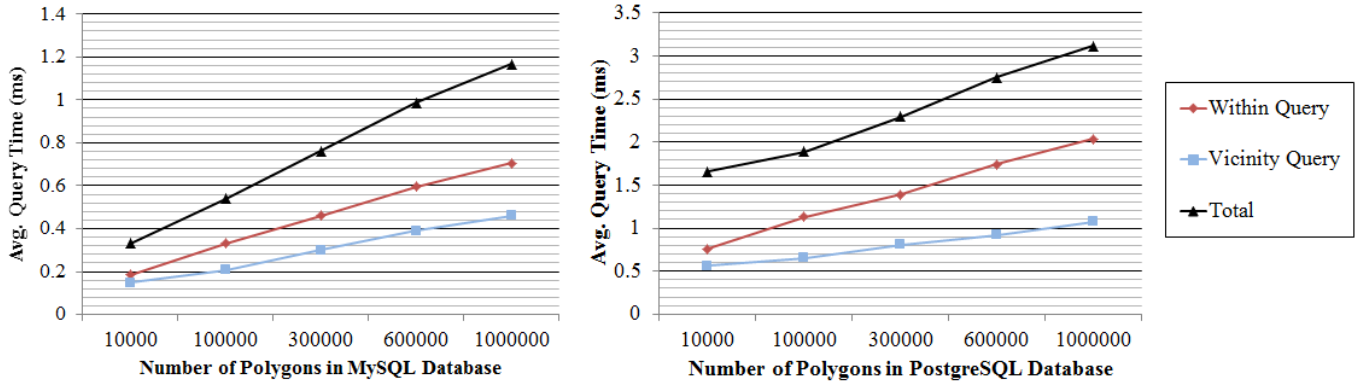
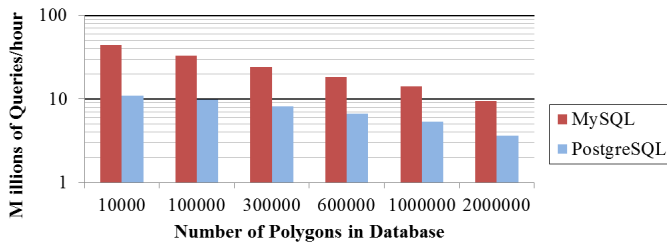Fig. 3. Database query times with polygons of 20 edges.

Fig. 4. Maximum query load.

to decrease the query times much. For example, having 10,000 polygons in a square area with 100 kilometer side lengths, the query times are only a few percent lower. Generating many more polygons to cover such a small area does not reflect the reality anymore.

B. Results: Query Load

Fig. 4 illustrates the query performance in terms of queries per hour. The maximum load was attained with four parallel querying processes constantly stressing the measured database. With 100,000 polygons in the MySQL database, we could sustain a load of 32M queries per hour. Continuing to 300,000, 600,000, and 1M polygons, the query load lowers in steps of about 25%. According to CTIA, there were 326.4M wireless subscriber connections in the United States in December 2012. Assuming they were all connected simultaneously and using one MySQL database storing 300,000 polygons for querying BSs, this would mean that each user can make 74 queries per hour. Of course, the usage should not be restricted only to subscribers as introduced earlier, but networks should be allowed to make queries as well. Also, this calculation neglects important factors such as limitations of network capacity and processing of queries, but, overall, the results indicate that a single logical spatial database can likely satisfy the query load for a large national cellular network. PostgreSQL sustains almost 10M queries per hour with 100,000 polygons. The PostgreSQL query load decrease is gentler than with MySQL, being on average in steps of 18% lower until 1M polygons. 2M polygons drops the performance only by 33% compared to 1M polygons with both databases.

We found that in the verification measurements the PostgreSQL's performance gap to MySQL was narrowed by 40%. The reason behind this is that PostgreSQL better utilized the 16 threads available in the server, whereas MySQL load saturated with 7-9 parallel processes when storing more than 100,000 polygons. Moreover, as we saw earlier, the newer PostGIS was observed to also perform better than the older version. However, we did not attempt to optimize the database settings. With MySQL, the maximum loads decreased 15% to 30% between 300,000 and 1M polygons compared to Fig. 4. But with 10,000 and 100,000 polygons, the respective maximum loads were interestingly higher. With PostgreSQL, the loads were constantly about 50% higher than shown in Fig. 4. Despite differences in the query loads, the maximum loads on the different servers were observed to decrease in a similar way as the number of polygons in the databases grows.

C. Results: Regional Optimization

As we saw in the previous section, queries over very small databases take about half the time compared to a very large polygon collection. Thus, it is more effective to split the database into regions. Fig. 5 illustrates the results with 300,000 polygons, where the performance of the non-optimized database is compared to databases where the cells are divided into different tables according to states, and states and counties the cells are located. In these measurements, we expect a nationwide information service, where one table stores the state borders and each state has a table storing the coverage areas of the cells within the state. When the database is further split into counties, for each state there is a table storing the county borders and each county has a table storing the cell polygons within it. In the hierarchical information service architecture shown in Fig. 2, the Root IS handles, for example, the county and/or state tables and forwards the queries to respective regional database tables, located in the regional servers but preferably also locally in the Root IS.

In state division, we assume that the BSs are equally distributed among the 50 states of the United States. In county division, we use the New York state as example containing 62 counties and assume that the county of New York includes 1,000 cells. This yields table sizes of 50, 62, and 1,000 polygons for state, county, and cells, respectively. With MySQL, the state division increases the maximum query load performance by 183%. Combined state and county division yields a gain of 195%, although there are two more database queries required in the server than in the non-optimized case. With PostgreSQL
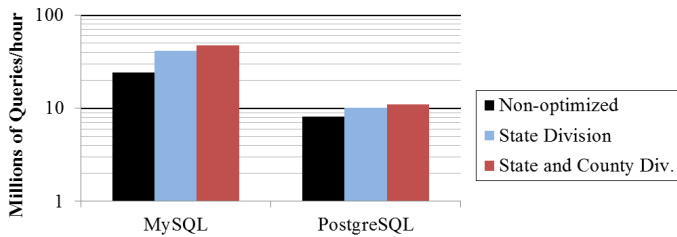
Fig. 5. Comparison of regional optimization models with 300k polygons.

the gains of regional optimization exceed 100% as well, being 130% and 135% with state and state and county divisions, respectively. MySQL outperforms PostgreSQL by over 400%, sustaining the load of 47M queries/hour with the three table optimization scheme.

V. CONCLUSION

Network information services can help operators keep up with increasing mobile traffic by better leveraging heterogeneous networks and shared spectrum. We evaluated whether a unified geographic database can provide fine-grained BS coverage information and found that even though the database sizes grow very large, modern databases perform well, providing query performance in the magnitude of millions of queries per hour even with a million polygons stored in a single database. However, regional optimization was observed to be helpful, more than doubling the query performance in both MySQL and PostgreSQL. Thus, our scalability results for both proved that basing both end-user mobility and shared spectrum access systems on same coverage area enhanced information service is feasible.

REFERENCES

[1] ITU-R, "Assessment of the Global Mobile Broadband Deployments and Forecasts for International Mobile Telecommunications," Report ITU-R M.2243, Radiocommunication Sector of ITU, 2011.
[2] Cisco, "Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2012-2017." White Paper, 2013.
[3] IEEE, "IEEE Standard for Local and Metropolitan Area Networks. Part 21: Media Independent Handover Services," IEEE Std. 802.21, 2009.
[4] 3GPP, "3GPP System Architecture Evolution (SAE); Security Aspects of non-3GPP Accesses (Release 8)," 3GPP Std. TS 33.402, Rev. 8.8.0, 2012.
[5] PCAST, "Realizing the Full Potential of Government-Held Spectrum to Spur Economic Growth," report, President's Council of Advisors on Science and Technology, July 2012.
[6] European Commission, "Promoting the Shared Use of Radio Spectrum Resources in the Internal Market," COM(2012) 478 final, 2012.
[7] E. Piri, "Cell Coverage Area Information Service to Improve Cell Selection in HetNets," in *Proceedings of the 11th Annual IEEE Consumer Communications & Networking Conference (CCNC)*, (Las Vegas, USA), Jan. 2014.
[8] F. Buiati, L. Villalba, D. Corujo, S. Sargento, and R. Aguiar, "IEEE 802.21 Information Services Deployment for Heterogeneous Mobile Environments," *IET Communications*, vol. 5, no. 18, pp. 2721–2729, 2011.
[9] D. Corujo, C. Guimaraes, B. Santos, and R. Aguiar, "Using an Open-Source IEEE 802.21 Implementation for Network-based Localized Mobility Management," *IEEE Communications Magazine*, vol. 49, no. 9, pp. 114–123, 2011.
[10] A. Cezar, "Practical Approach of Implementing Media Independent Information Service from IEEE 802.21 Standard," in *Proceedings of 1st International Conference on Wireless Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems Technology (Wireless VITAE 2009)*, (Aalborg, Denmark), pp. 92–96, May 2009.
[11] S. Frei, W. Fuhrmann, A. Rinkel, and B. Ghita, "Improvements to Inter System Handover in the EPC Environment," in *Proceedings of 4th IFIP International Conference on New Technologies, Mobility and Security (NTMS)*, (Paris, France), pp. 1–5, Feb. 2011.
[12] S. Ray, B. Simion, and A. D. Brown, "Jackpine: A Benchmark to Evaluate Spatial Database Performance," in *Proceedings of the IEEE International Conference on Data Engineering (ICDE)*, (Hannover, Germany), pp. 1139–1150, Apr. 2011.
[13] Z. Zhou, B. Zhou, W. Li, B. Griklak, C. Caiseda, and Q. Huang, "Evaluating Query Performance on Object-Relational Spatial Databases," in *Procceedings of the 2nd International Conference on Computer Science and Information Technology (ICCSIT)*, (Beijing, China), pp. 489–492, Aug. 2009.
[14] W. Song, J. Lee, and H. Schulzrinne, "Polygon Simplification for Location-Based Services Using Population Density," in *Proceedings of IEEE International Conference on Communications (ICC)*, (Kyoto, Japan), pp. 1–6, June 2011.
[15] T. Hardie and A. Newton and H. Schulzrinne and H. Tschofenig, "LoST: A Location-to-Service Translation Protocol." IETF RFC 5222, 2008.
[16] WBA, "Global Developments in Public Wi-Fi," WBA Industry Report 2011, Wireless Broadband Alliance, 2011.
[17] 3GPP, "3GPP system to Wireless Local Area Network (WLAN) Interworking; System Description (Release 8)," 3GPP Std. TS 23.234, Rev. 8.0.0, 2008.
[18] E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio, "Sensing the Internet Through Crowdsourcing," in *Proceedings of IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, (San Diego, USA), pp. 248–254, Mar. 2013.
[19] FCC, "Second Memorandum Opinion And Order," FCC 10-174, Federal Communications Commission, Sept. 2010. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-174A1.pdf.
[20] FCC, "FCC Proposes to Enable Innovative Small Cell Use of Spectrum in the 3.5 GHz Band," Dec. 2012. Available: http://www.fcc.gov/document/fcc-proposes-innovative-small-cell-use-35-ghz-band.
[21] M. Matinmikko, M. Palola, H. Saarnisaari, M. Heikkilä, J. Prokkola, T. Kippola, T. Hänninen, M. Jokinen, and S. Yrjölä, "Cognitive Radio Trial Environment: First Live Authorized Shared Access-Based Spectrum-Sharing Demonstration," *IEEE Vehicular Technology Magazine*, vol. 8, pp. 30–37, Sept. 2013.
[22] ECC, "Report on ASA concept," FM(12)084 Annex 47, Electronic Communications Committee, 2012.

Esa Piri received his M.Sc. (Tech.) from the University of Oulu, Finland, in 2008. Since then, he has been working as a Research Scientist at VTT Technical Research Centre of Finland. Currently, he is also pursuing a doctoral degree from a topic related to cross-layer issues in heterogeneous network environment in the University of Oulu. His research interests and activities are mainly related to cognitive and seamless networking in heterogeneous networks, network management, network measurements, and testbeds. He has published over 30 journal articles, conference papers and book chapters. Esa has been a member of IEEE since 2008.

Henning Schulzrinne, Levi Professor of Computer Science at Columbia University, received his Ph.D. from the University of Massachusetts in Amherst, Massachusetts. He was an MTS at AT&T Bell Laboratories and an associate department head at GMD-Fokus (Berlin), before joining the Computer Science and Electrical Engineering departments at Columbia University. He served as chair of the Department of Computer Science from 2004 to 2009, as Engineering Fellow at the US Federal Communications Commission (FCC) in 2010 and 2011, and as Chief Technology Officer at the FCC since 2012.

He has published more than 250 journal and conference papers, and more than 70 Internet RFCs. Protocols co-developed by him, such as RTP, RTSP and SIP, are now Internet standards, used by almost all Internet telephony and multimedia applications. His research interests include Internet multimedia systems, ubiquitous computing, and mobile systems.

He is a Fellow of the IEEE, has received the New York City Mayor's Award for Excellence in Science and Technology, the VON Pioneer Award, TCCC service award, the IEEE Region 1 William Terry Award for Lifetime Distinguished Service to IEEE and the UMass Computer Science Outstanding Alumni recognition.