HOW A LAYERED REM ARCHITECTURE BRINGS COGNITION TO TODAY'S MOBILE NETWORKS

JAAP VAN DE BEEK AND TAO CAI, HUAWEI TECHNOLOGIES
SEBASTIEN GRIMOUD AND BERNA SAYRAC, ORANGE LABS
PETRI MÄHÖNEN, JAD NASREDDINE, AND JANNE RIIHIJÄRVI, RWTH AACHEN UNIVERSITY

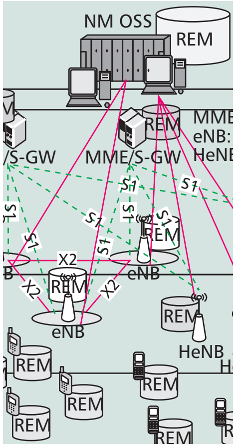

The authors present a layered radio environment map architecture along with its applications to the self-organizing network functionalities of heterogeneous LTE radio access networks comprising macrocells and femtocells.

ABSTRACT

In this article, we present a layered radio environment map architecture along with its applications to the self-organizing network functionalities of heterogeneous LTE radio access networks comprising macrocells and femtocells. In this architecture, the functional blocks reappear with different spatial and temporal granularity at different architectural layers. Although the radio environment map is one of the key promising technologies to enable future cognitive radio networks, it can be already applied to provide limited cognitive capabilities to today's commercial networks too. We explain why, and show how, this architecture can support today's LTE self-organizing network functions like automatic neighbor relation and minimization of drive tests, and also allow the smooth introduction of new radio access technologies through refarming. We also demonstrate some of the quantitative benefits adopting radio environment map technologies can bring using the minimization of drive tests as an example.

INTRODUCTION

Rapidly increasing traffic volumes are causing a significant increase in the complexity of cellular networks. While new technologies are constantly being deployed to keep up with these capacity demands, operators also tend to increase the spectral efficiency requirements of their existing networks in order to delay expensive network transitions and maximally benefit from their investments. Management of the increasingly heterogeneous cellular networks that result from this approach is extremely challenging, and traditional planning tools and radio resource management (RRM) techniques are no longer sufficient. Furthermore, runtime operation of these networks should become as autonomous as possible in order to minimize the operational costs.

The cognitive radio (CR) paradigm was introduced by Mitola [1] to provide context-sensitive optimization, learning, and adaptation capabilities for radio devices and networks. Although a majority of the early research in the CR domain was done in the context of dynamic spectrum access (DSA), the underlying larger goal is still valid and strongly pursued. Recently the research community has also recognized that many enabling technologies considered for future DSA and CR systems can also be applied to present-day networks to bootstrap CR approaches much faster than expected. In this article we show how one such enabler, the radio environment map (REM) [2], can add value to cellular mobile network architectures and offer solutions to the challenges outlined above.

A REM can be thought of as a knowledge base used to dynamically store information related to the radio environment of wireless systems. This information can be represented either by raw radio field measurements or, more efficiently, as the result of modeling processes such as statistical behavioral descriptions. REMs are currently being studied and specified in the European Telecommunications Standards Institute's (ETSI's) emerging standard on reconfigurable radio systems [3]. In contrast to the static databases used in third-generation (3G) and Long Term Evolution (LTE) systems, REMs provide a wireless network with a comprehensive and up-to-date representation of the radio environment, including dynamic knowledge on propagation environment, which can be used to optimize radio resources. This reduces the need for operator's drive tests and measurement campaigns, and hence reduces its operational expense (OPEX).

While much of the recent work on REMs is targeted to envisaged future cognitive radio networks [3, 4], we present here an extensible REM architecture that adds value to conventional existing cellular networks. This architecture, which potentially improves system coverage and capacity of today's cellular networks, has been developed and implemented as one of the core results of the EU-funded FARAMIR project [5, 6]. Today's cellular networks rely on only a small amount of *static* information stored in databasea like the visited location registry (VLR) and home location registry (HLR), which normally include security and location information that

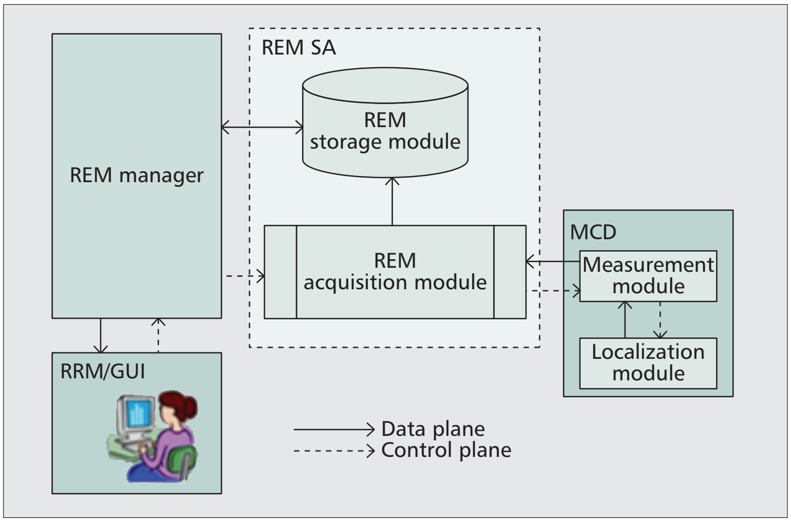

Figure 1. *The functional REM architecture.*

can be used for identification, paging, and handover. In addition, the existing planning tools are applied offline and rely on collected information through expensive measurement campaigns. By enabling the use of a dynamic database such as REM, we enable more autonomous and efficient network operation, more cost-effective network management, along with a more flexible evolution of future network deployments. We give examples of these in coming sections.

In this article we show how a high-level functional REM architecture maps onto contemporary cellular system architectures. We envisage a layered architecture where different (partial) instances of the functional REM reappear in different network components. REM layers (REM instances of the same scale and with similar scope and information characteristics) are connected hierarchically.[1]

We address the Third Generation Partnership Project (3GPP) LTE network, and the added value brought by giving the REM functional blocks a context in the architecture of this cellular system. In particular, we describe and explore FARAMIR's layered REM architecture for HetNets, 3GPP's heterogeneous LTE networks [8] comprising macro- and femtocells. We study the particular application of the REM architecture to various LTE functions: automatic neighbor relation (ANR) [9, 10], minimization of drive tests (MDT) [11–13], and the introduction of new radio access technologies (RATs) through refarming. We also discuss the effectiveness of deploying these applications through only one enabler, the layered REM. For the MDT case we also demonstrate quantitatively some of the potential benefits that can be gained by adopting REMs as part of the solution.

FARAMIR'S LAYERED REM ARCHITECTURE

Figure 1 shows the high-level functional REM architecture we explore in this article. It is composed of measurement-capable devices (MCDs), a REM data storage and acquisition (REM SA) unit, and a REM manager. A set of interfaces connects the REM to different applications (RRMs or user interfaces that monitor system performance and track error causes).

MCDs are network entities that acquire the measurement information from the environment. The REM SA unit has two functions. Its data acquisition module accounts for all communication with the various MCDs. It sends measurement instructions to MCDs, collects measurement reports, and stores these in the storage module. The latter is essentially a database where, besides the raw reported data, processed data are also stored, in the form of maps and quantitative learned models. Finally, the REM manager, being the brain of the architecture, generates and maintains the REMs. It decides which measurements are performed, by which nodes, and when. It sends measurement instructions to the REM SA and processes raw data to generate REMs.

Let us now turn our attention to how this functional REM architecture can be of advantage in a contemporary cellular system architecture. We start by analyzing two important performance aspects: signaling overhead and information reliability. They are important because they directly relate to key performance metrics such as network capacity and energy efficiency. Based on this analysis, we then present architectural principles.

The creation, modification, and dissemination of information associated with REMs is inevitably associated with signaling overhead. While in some situations this overhead is small (e.g., a REM user may be collocated with the node hosting the relevant REM instance), other cases may come with a substantial overhead that may lead to performance degradation (e.g., the amount of relevant data is large or the connection between the communicating nodes has limited capacity). Therefore, a network operator must balance signaling overhead with the performance improvement that comes with exploiting REM information. This is one of the main challenges that has made the deployment of distributed databases for environmental information less attractive to operators. By designing an architecture that carefully considers the type and amount of information exchanged between the different layers, signaling overhead can be significantly reduced. Hence, one of the most important aspects of a layered REM architecture is the time of validity of REM information. Raw measurement data of rapidly changing radio environmental characteristics typically need to be made available through local REM instances, if at all. Statistical characteristics of such data, on the other hand, have longer times of validity and therefore can easily be shared through various REMs.

Signaling overhead can be reduced by a suitable choice of dissemination strategy. A pull-based on-demand model is most appropriate when information is needed only rarely. For example, the positions of base stations or their frequency allocations typically do not change rapidly and can be requested periodically or when activated. A push-based proactive model, on the other hand, may result in better performance for commonly needed and dynamically changing data. Interference statistics and sub-

---

[1] *Parts of this article are based on our earlier conference paper [7], which has been substantially extended and revised. In particular, we have formalized the REM design with three leading architectural principles and added results of a quantitative study into the value of deploying REMs to reduce the need for drive tests in a network.*

channel allocation in scenarios where femtocells share frequency channels with macrocells will, for example, have to be pushed by the network in each frame.

Besides low signaling overhead, information reliability of the REM is important for RRM techniques too. The required accuracy depends on the objectives of the RRM technique and their timescale. For instance, fast power control requires high precision, while dynamic spectrum allocation performs well with approximate or statistical information.

The reliability also depends on the aggregation of information in space and time. For instance, statistical information about shadowing coefficients can only be accurate if the region from which the data is collected is homogeneous and the measurement time is higher than the coherence time in this region. A choice of large regions for such data aggregation to reduce signaling overhead would come at the expense of the accuracy of the model or statistical characterization. Similarly, local caching of results obtained from REMs can be very effective in reducing overhead but only yields relevant results when the validity of the cached data is assured.

REM-ARCHITECTURAL PRINCIPLES

The above considerations form the motivation for three REM-architectural principles: layeredness, subsidiarity, and proportionality.

LAYEREDNESS

Different functional REM components may reside in any of the network's hierarchical layers, in a standalone database, an operator's management system, an eNB/HeNB, an HeNB gateway, or, if necessary, in terminals. One of the key insights behind the work presented here is that for many application scenarios the REM system architecture requires a layered structure, characterized by instances of the same functional architecture reappearing at different hierarchical levels. REM SA instances may, for example, appear in user terminals or base stations, obviously storing very different types of information. A layered architecture is needed because some network nodes have limited computation power or memory. Alternatively, some data is needed only in certain network elements, and dissemination to other elements would be a waste of capacity.

SUBSIDIARITY

REM information is stored and handled in the smallest node, lowest in hierarchy, or least centralized possible. As one example, on a nationwide scale three layers can be identified, each layer having its specific spatial-temporal granularity, as shown in Table 1. A highest layer, national level, contains a cross-operator REM-database and manager. A second national/metropolitan layer contains an intra-operator REM, while the lowest layer contains a local neighborhood-level REM. As a second example, coverage maps of eNBs are stored in the particular eNB where the underlying measurements and power configuration can also be found. This enables proper coordination between a femtocell and a macrocell with minimum signaling overhead and delay, and maximum information efficiency. REMs of these eNBs will also have information about other macro-/femtocells in the vicinity, information not necessarily stored in femtocell REMs, and can be used by the eNodeB for efficient resource management.

	National REM	Operator national REM	Operator local REM
Node locations	X	X	X
Node resource capacities	X	X	X
List of eNBs/HeNBs		X	X
General user data/price plan		X	X
Location-based user service preference		X	X
User behavior/traffic statistics		X	X
Local interference statistics			X

Table 1. *Three-layered REM implementation in macro- and femtocells.*

PROPORTIONALITY

The REM update rate must be in proportion to its level in the architectural hierarchy. For fast updating rates (milliseconds, seconds), the REM SA and REM manager must typically be deployed in the BS because of the direct connection and low latency in the connection with the MCDs. For slow updating rates, on the other hand (hours, days, weeks), the data storage module of the REM SA and the REM manager are deployed in the operator's management system, the data collection module being deployed in the base station.

Figure 2 shows how the REM architecture can be mapped onto the basic 3GPP architecture of a heterogeneous LTE cellular network that encompasses both macrocells and femtocells. Femtocells are mainly envisioned for private use where coordination with the network's macrocells is essential in order to control interference levels and maintain the required QoS in the macrocells. In this heterogeneous LTE network, a layered REM plays a key role in the appropriate spectrum sharing between macro- and femtocells or other typically envisioned RRM use cases (mobility management, interference management, admission control, and offloading).

In the following sections we describe detailed system architectural considerations, and the functional components and their placement for three use cases and applications.

APPLICATION TO LTE'S AUTOMATIC NEIGHBOR RELATION

Accurate identification of a BS's neighbors is vital for RRM procedures like mobility and interference management to work properly. This includes intrafrequency, interfrequency, as well

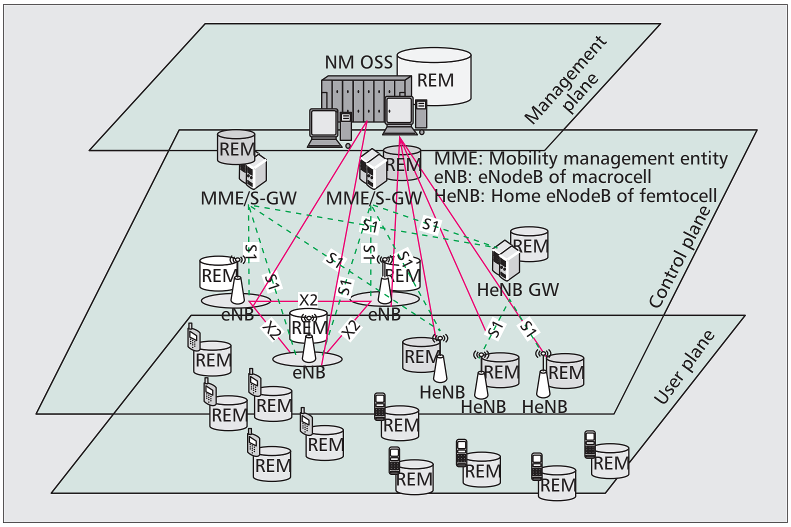

Figure 2. *Layered REM on a heterogeneous LTE architecture that includes macrocells and femtocells. Different REM colors represent different layers with different characteristics.*

as inter-RAT neighboring BSs (eNBs). Up to release 8 of 3GPP's LTE specifications, operators manage these neighboring relations manually. As of release 8, however, an automatic neighbor relation (ANR) functionality is part of the LTE standard. As one of the first functionalities associated with self-organizing networks (SONs) [9, 10], its purpose is to relieve the operator of the burden of manually managing these neighbor relations.

In the current specifications [9], user equipment (UE) terminals measure the received signal reference power from intrafrequency, interfrequency, and/or inter-RAT neighbors, and inform its serving eNB or HeNB of those neighboring cells from which a strong reference signal is received. Based on these UE measurement reports, the BS extracts its own local neighbor relations table (NRT) containing the identities of its neighbors. Finally, this local neighboring information is communicated to the operation and maintenance center (OMC) where monitoring and higher-level management of neighboring relations are carried out by the operator for possible local conflict resolution and compliance with global coherence in the whole radio access network. Thus, the automatically generated neighbor relations can be managed by the operator. The addition or deletion of NRTs is possible from the operator's management system, and the attributes of the NRTs can be changed too. In [10] a field study in a commercially deployed LTE network was reported illustrating the feasibility of automatically configuring LTE neighbor relations based on this bottom-up signaling of measurement reports.

In short, the current ANR functionality is specified through a bottom-up dissemination approach, and it is in this aspect where a layered REM architecture can add value. By placing REM coverage maps with different temporal and spatial characteristics at different hierarchical layers, it is possible to obtain and maintain more precise, up-to-date, and reliable neighboring information. Specifically, downlink terminal coverage measurements reported to the BSs are used to construct a low-level, local, and quasi-real-time REM coverage map. Using extrapolated geolocalized coverage information stored in these REM coverage maps, the BS can verify the coherence between this information and the information of the neighboring BSs already stored higher in the radio access network.

Using the coverage measurements reported by the BS to higher architectural layers, global REM coverage maps can be constructed, typically at the OMC. Equipped with global geolocalized coverage information, the OMC does not need to resolve conflicting neighboring information coming from the eNBs, HeNBs, or HeNB gateways. It will have more precise and realistic information about coverage, and thus on the neighboring relations. These relations are then disseminated through a top-down approach.

This layered REM architecture has implications for the mapping of the REM's functional architecture components onto the entities of the radio access network (Table 2). First, the UE measurements are satisfied by existing UE capabilities: user terminals are the sole embodiments of the MCD. In some cases the HeNB can also include instances of the MCD, since there is significantly higher probability that the femtocell does not have any active users than a macrocell. Second, the REM SA (both the acquisition and storage components) that contains the coverage map composed of raw and processed measurements has a two- or three-layered structure. Both femtocell and macrocell networks embody REM coverage maps at the BS (low-level) and at the OMC (high-level). In dense femtocell networks an additional middle-level map may appear at the HeNB gateway. Finally, the REM manager appears at each layer, each manager managing the coverage maps at its respective level.

APPLICATION TO LTE'S MINIMIZATION OF DRIVE TESTS

Network operators today predominantly conduct drive tests to collect performance metrics that guide network deployment and operation. Orange Romania, for example, schedules drive tests once a year in the top 60 cities and the main roads, and twice a year in the top 25 cities. When coverage holes are detected, new sites may be deployed, power configurations may be optimized, or antenna tilts and azimuth may be changed to improve the service quality. In general, drive tests are used to identify coverage problems, to help improving user experience, and to optimize network capacity.

As they are costly, and inevitably also cause the emission of CO_2, it is desirable to develop automated solutions to reduce the number of drive tests. Such minimization of drive tests (MDT) is currently being studied and specified by 3GPP for release 10 of the LTE standard [11–13]. The main functionality of MDT is to enable terminals to carry out various radio network measurements including cell identity information, received power, and quality information. In addition, perceptual evaluation of speech

	ANR	MDT	new RATs
OMC	REM manager & REM-SA	REM manager & REM-SA	REM manager & REM-SA
HeNB gateway	REM manager & REM-SA	REM manager & REM-SA	REM manager & REM-SA
eNodeB (BTS)	REM manager & REM-SA	MCD & REM manager & REM-SA	REM-SA
UE	MCD	MCD & REM-SA	MCD

Table 2. *REM functionality in various network layers for the three use case examples.*

In our simulations, the locations with received signal power values that are below a coverage threshold are considered as being in outage (or out of coverage). Those points are used for evaluating the performance of outage predictions in an MDT scenario.

quality for voice services and the throughput for FTP traffic can be measured at the application layer. All these types of information are tagged with available time/location information and then aggregated into an MDT measurement report. Recent studies focus on tailoring the measurement reports with respect to signaling overhead and their capability to support coverage hole detection, handover parameterization, and interference problems in the network [11].

These reports can be handled in two ways. Measurements can be reported to the network immediately when the terminal is in connected mode. Alternatively, measurements can be performed in idle mode, stored/logged in the memory of a terminal and retrieved by the network at a later time through signaling [11]. Naturally, the operator will aim to facilitate the user by ensuring that temporary storage of MDT measurements and their reporting signaling will consume a minimum of data storage capacity and transmission power. As an example, the current 3GPP MDT feature specifies a minimum buffer size requirement to be 64 kbytes [14]. Anyway, memory requirements will increase and become a burden for the terminal.

The layered REM concept is likely to play a crucial role and adds value to this functionality by optimizing storage and signaling through environmental awareness and coordination. Increasing the reporting rate or reducing the logging rate can reduce the impact on the terminal memory, in a coordinated way taking into account other aspects such as MDT performance and terminal power consumption. The intelligence embedded in the layered REM allows efficient distribution of MDT data storage at different hierarchical levels of the system architecture. In particular, the terminal memory can be managed as a lowest level of REM storage; note the functional REM distribution across the network layers in Table 2.

The REM manager has a number of other optimization means. First, the optimization can be done dynamically to balance storage capacity available at different network nodes. The REM manager can choose when to retrieve MDT measurement reports from terminals, based on the evaluation of terminal memory, traffic load in the system, as well as energy consumption for reporting. During rush hours, for example, when traffic demand is high, the logged MDT data can be stored at the UE REM level and then be reported in a scheduled manner to alleviate possible uplink signaling congestions. Second, the REM manager can decide to which network node the terminal reports its MDT data. Using environment-aware coordination between macrocells and femtocells, and provided by the layered-REM itself, the terminal can send the logged MDT reports via an HeNB instead of an eNB, regardless of the measurements' origin. This results in reduced uplink transmission power, leading to terminal energy savings and hence to the extension of a terminal's battery life. Third, the REM manager can decide where to store MDT measurements depending on how these measurements are to be used. MDT data related to local coverage holes can be stored in local REMs (eNB, HeNB, and HeNB gateway) for use in localized and fast optimization algorithms. MDT data related to larger area coverage problems, which involve more than one cell or are related to handover, are stored in higher-level REMs at the OSS/MMS level. Various optimization algorithms can then be operated and guided by a higher-level REM manager.

We have studied the performance of the REM-enabled MDT based on realistic simulation data obtained using a real network topology, exact terrain information, and a precise ray-tracing propagation model developed at Orange Labs. Signal power received by a typical mobile device from a BS in an urban environment (Paris region) has been used as the reference real measurement data. The spatial resolution is 5 m, chosen to be in accordance with typical outdoor correlation distance values. Our focus here is on detecting coverage holes, although the techniques themselves have a much wider application scope.

In our simulations, the locations with received signal power values that are below a coverage threshold are considered as being in outage (or out of coverage). Those points are used for evaluating the performance of outage predictions in an MDT scenario. It is assumed that the operator carries out drive test measurements on a randomly selected set of measurement locations over a zone of interest within the coverage area of the BS. The percentage of the measurement locations over the zone of interest is fixed by the operator. REM predictions on received signal power values are calculated over a second set of randomly chosen prediction locations. Similarly, the operators also fix the ratio of the number of prediction locations to the number of measurements. The method used here for REM predictions is one of the state-of-the-art methods from spatial statistics, Bayesian Kriging interpolation.

Example results from our performance evaluation are shown in Fig. 3 where the probabilities

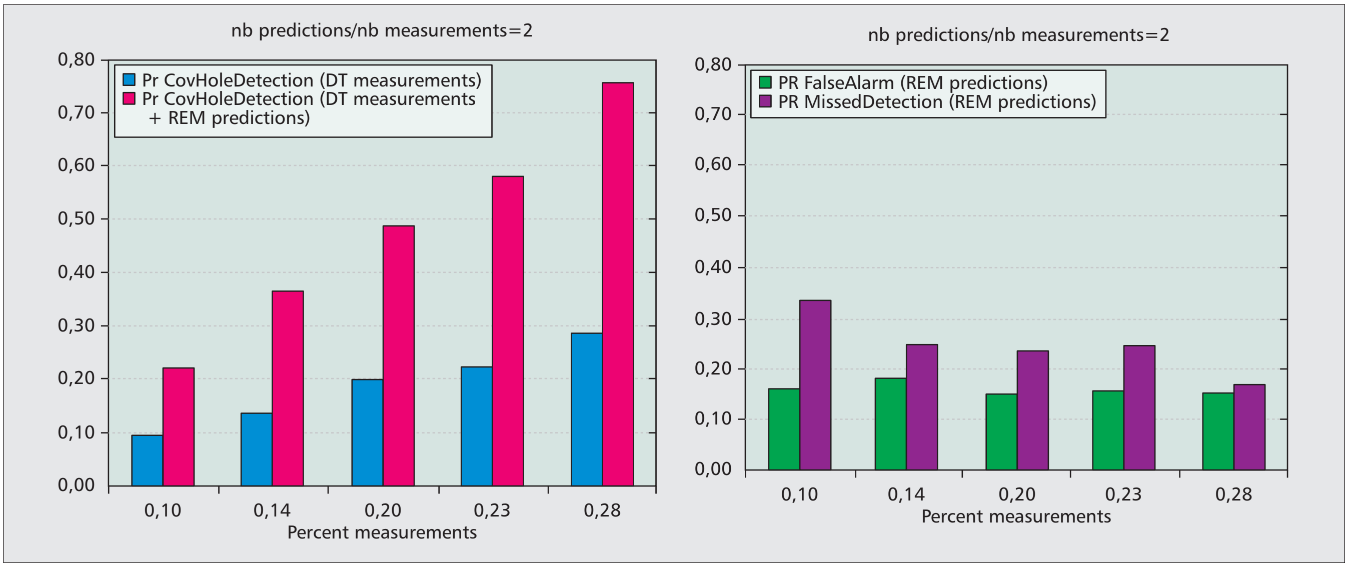

Figure 3. *Comparison of coverage hole detection probability with drive test (DT) measurements and with the combination of DT measurements and REM predictions (left) together with the corresponding false alarm and missed detection probabilities (right) for a prediction-to-measurement ratio of 2.*

of coverage hole detection and false alarm as a function of the percentage of measurement locations are depicted with and without REM. We see that the adoption of REMs significantly increases the probability of detecting coverage holes compared to using raw drive test data only. Even for relatively large measurement counts the latter case results in detection probability of less than 30 percent, which can be increased to 75 percent by incorporating REM techniques. We also see that the false alarm probability remains at a very acceptable level.

We have also studied the required signaling overhead for REM use. The signaling required for building and maintaining REMs in scenarios like the ones discussed here relies mainly on the measurement reporting procedures, with typical update rates of the REM being on the order of weeks or months. Therefore, the cost of REM mainly consists of storage issues rather than signaling overhead on the interfaces. The storage size required by a REM is equal to the product of the size of the relevant 3GPP Measurement Report message [12] by the required number of this message over the geographical zone of interest. The former needs a detailed system architecture design and is a function of the number of reported neighboring cells. The second one depends on the correlation distance, the spatial resolution of the REM (which is linked to the correlation distance), the measurement percentage, and the prediction-to-measurement ratio. We have evaluated these for our performance evaluation scenario, resulting in the estimates for the required REM storage size for 10 percent measurement ratio and a prediction-to-measurement ratio of 1 equaling 834 kb. When the measurement ratio is increased to 20 percent with a prediction-to-measurement ratio of 3, the required REM size increases to 7.5 Mb. Both of these numbers are very manageable for current database and data processing technologies.

APPLICATION TO THE INTRODUCTION OF NEW TECHNOLOGIES THROUGH REFARMING

Operators typically introduce a new RAT in their radio access network through a process commonly known as refarming. A gradual migration, introducing the new RAT in the existing frequency band in a predefined area, is preferred over an abrupt one, for reasons of capital expenditures (CAPEX) associated with the deployment and inertia of the RAT capability evolution of the user terminals in use. Usually, the coexistence of two technologies is handled by leaving enough spectrum (guard bands) or using a less aggressive frequency reuse scheme (guard distances). This inevitably wastes resources.

Layered REMs can be used to automatically find and set relevant parameters (frequency reuse factor, guard band interval, etc.) in these specific areas. In Fig. 4, there is a single access technology (T1) over the entire area except the middle elliptic region where technologies T1 and T2 coexist. A large guard band (usually larger than needed) is cautiously set between these two technologies in order to guarantee no interference. With precise interference information on each location coming from the REM, the guard band can be adjusted in an optimum way to prevent interference and minimize waste of spectrum resources. Geolocation-based measurement is critical for performance optimization across old and new technologies where multimode terminals can act as MCDs. Thus, layered REM can be built with the support of one central entity, controlling entities for old and new technologies (BSCs for GSM, RNCs for UMTS/HSDPA, and eNodeBs for LTE as well as terminals). Here, different instances of REMs are required at different network entities belonging to different technologies. We therefore speak of a horizontal hierarchy.

Table 2 illustrates again how various REM functions are distributed across the various layers of the network. The geo-localized measurements reported by the multimode terminals are collected at the base stations of different technologies where different instances of REM SAs are located. These base stations include BTSs for GSM, macro-/femtocell NodeBs for UMTS/HSPA, and macro-/femtocell eNodeBs for LTE. The REM SAs at the BS level are small-scale high-resolution REMs, with small coverage areas and high temporal and spatial granularities. These low-level maps allow the accurate detection and localization of coverage and interference problems.

The solution of these problems, however, is executed at a higher level in the network architecture. Therefore, the REM manager must be located in the operator's domain (O&M for macro BSs, (H)eNB management system for femtocells) where a more global view that encompasses several technologies is available. Since the problems are solved at the O&M, necessary information must be conveyed to the O&M by the local REM SAs. This information can contain the precise localization of the problem, along with some statistical information about the value of the related KPIs and/or counters.

Large-scale low-resolution maps are also needed at a higher level of the network hierarchy (typically at the O&M) for detecting, localizing, and solving larger coverage and interference problems. In order to avoid communication overhead between BSs and the O&M, the geolocalized measurements must have temporal and spatial granularities that are below certain thresholds. Since the decisions on the solutions to the coverage/interference problems are made at the operator's domain, the GUI and/or RRM entity that accounts for such decisions is found at the operator's domain, collocated with the REM manager, in order to facilitate appropriate information passing.

CONCLUSIONS

As discussed above a layered radio environment map architecture, with a scope that exhibits subsidiarity and with proportionality of its dynamics, can add value and bring cognition to conventional mobile networks. Such an architecture facilitates radio network optimization in heterogeneous LTE systems comprising macro- and femtocells. In particular, two self-organizing network functionalities in 3GPP's LTE standard, ANR and MDT, can take advantage and gain efficiency, and the introduction of new services through refarming is also facilitated. Applying REM architecture and techniques offers an intriguing possibility to bring cognition into present-day cellular systems and standards without any changes in regulations or radio bearers themselves. Naturally, the developed hierarchical REM system has even wider use in the context of cognitive radio networks. We would like to conclude by noting that the reported above work is not only theoretical and conceptual simulation-based. The FARAMIR project consortium has built the first prototype implementations of the key parts of REM components and has successfully demonstrated its capabilities.

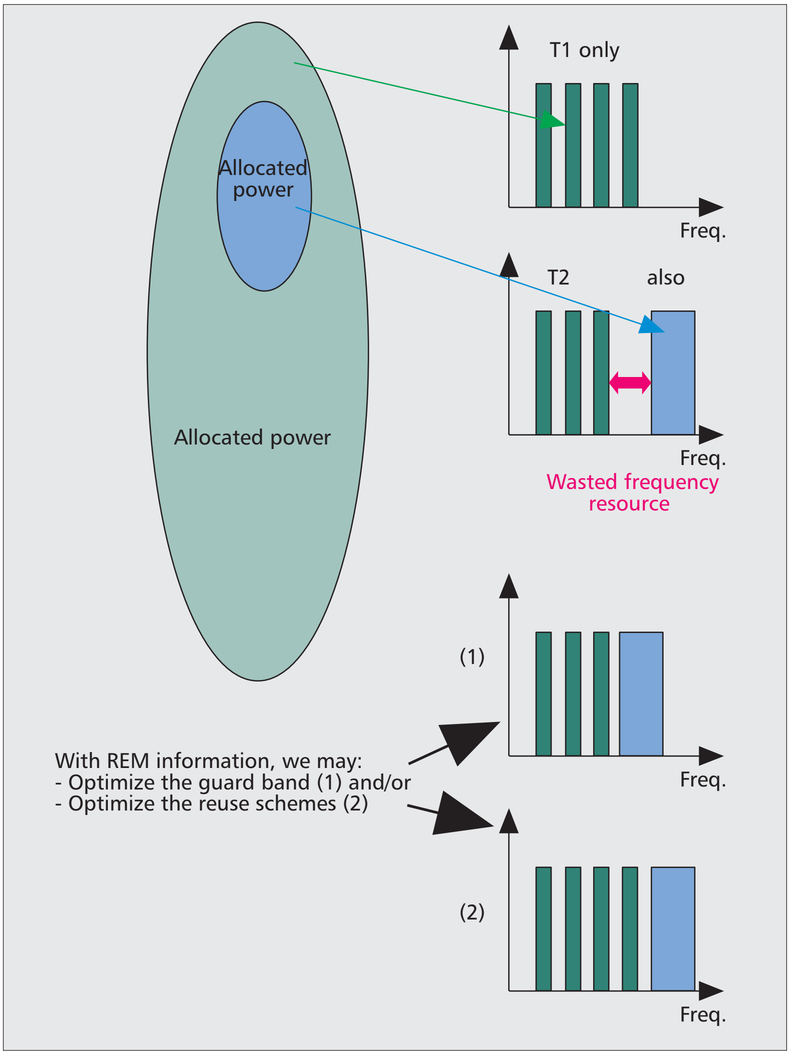

Figure 4. *Optimization of refarming guard bands using REMs.*

ACKNOWLEDGMENTS

We would like to thank our colleagues in the European Union funded FARAMIR project (grant number ICT-248351) for useful discussions and collaboration, and the European Union for partial financial support though the project. One of us (PM) is also acknowledging partial support from German DFG through UMIC research center funding.

REFERENCES

[1] J. Mitola, *Cognitive Radio: An Integrated Agent Architecture for Software Defined Radio*, Ph.D. thesis, KTH Royal Institute of Technology, 2000.
[2] Y. Zhao, B. Le, and J. H. Reed, "Network Support – The Radio Environment Map," *Cognitive Radio Technology*, B. Fette, Ed., Elsevier, 2006.
[3] ETSI-RRS, TR 102 947, "Reconfigurable Radio Systems; Use Cases for Building and Exploitation of Radio Environment Maps for Intra-Operator Scenarios," v 0.0.1, June 2011.
[4] J. Mitola, "Cognitive Radio Architecture Evolution," *Proc. IEEE*, vol. 97, no. 4, Apr. 2009.
[5] http://www.ict-faramir.eu
[6] ICT-248351 FARAMIR, B. Sayrac, Ed., "FARAMIR Document D2.2, Scenario Definitions," v1.01, Aug. 2010.
[7] T. Cai *et al.*, "Design of Layered Radio Environment Maps for RAN Optimization in Heterogeneous LTE Systems," *Proc. IEEE PIMRC 2011*, Toronto, Canada, Sept. 2011.

The work reported here is not only theoretical and conceptual simulation based. The FARAMIR project consortium has built the first prototype implementations of the key parts of REM components and has successfully demonstrated its capabilities.

[8] 3GPP TS 36.300 v11.1.0, "Evolved Universal Terrestrial Radio Access (E-UTRA) and Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Overall Description; Stage 2."

[9] 3GPP TR 36.902 v9.3.1, "Evolved Universal Terrestrial Radio Access Network (E-UTRAN); Self-Configuring and Self-Optimizing Network (SON) Use Cases and Solutions."

[10] A. Dahlen *et al.*, "Evaluations of LTE Automatic Neighbor Relations," *Proc. 2011 IEEE VTC Spring*, 15–18 May 2011.

[11] J. Puttonen *et al.*, "Coverage Optimization for Minimization of Drive Tests in LTE with Extended RLF Reporting," *Proc. IEEE PIMRC*, 26–30 Sept. 2010.

[12] 3GPP TS 37.320 v10.4.0, "Radio Measurement Collection for Minimization of Drive Tests (MDT); Overall Description; Stage 2."

[13] 3GPP TR 36.805 v9.0.0, "Study on Minimization of Drive-Tests in Next Generation Networks."

[14] 3GPP TS 36.306 v10.5.0, "User Equipment Radio Access Capabilities."

BIOGRAPHIES

JAAP VAN DE BEEK [M] is a principal research engineer with Huawei Technologies Sweden. He has held research positions in the telecommunications industry and in academia since 1994. With some highly cited contributions he was among those pioneering OFDMA as an access scheme for cellular radio in the mid nineties. Later he contributed to the preparation and specification of 3GPP's LTE standard and currently he is engaged in collaborative European research on cognitive radio and future cellular networks. He is an editor of *IEEE Communications Letters* and *IEEE ComSoc Technology News* and a recipient of the IEEE Heinrich Hertz Award.

TAO CAI received his B.Sc. in physics from Fudan University, Shanghai, in 1990 and his Ph.D. in electrical engineering from Beijing University of Posts and Telecoms in 1999. From 1999 to 2001 he worked at Ericsson Research's Beijing branch. From 2001 to 2002 he was a research associate with the Radio Communication Systems Laboratory of Royal Institute of Technology in Sweden. Since August 2002 he has been with Huawei Sweden R&D center as a senior researcher. His research interests include radio resource management, self-organizing networks, and cognitive radio.

SEBASTIEN GRIMOUD received an M.Sc. in 2009 in technologies for broadband communications from University College London, United Kingdom, and an engineering degree from Telecom Bretagne, France. In 2009 he joined France Telecom R&D as a research engineer. Since 2011 he is in charge of cross-border agreement negotiation at ANFR.

PETRI MÄHÖNEN [SM] is currently a full professor and head of the Institute for Networked Systems at RWTH Aachen University. His research interests include cognitive radio systems, self-organized networks, embedded intelligence, and optimization of communications systems. He currently serves as an Associate Editor of *IEEE Transactions on Mobile Communications*.

JAD NASREDDINE [M] is currently working as a senior researcher at the Institute for Networked Systems at RWTH Aachen University. Between 2005 and 2010 he worked as a postdoctoral researcher in different European universities. He received his Ph.D. degree in March 2005 from the University of Rennes I and his engineering diploma in 2001 from the Lebanese University. During his Ph.D., he worked as a research assistant at Telecom Bretagne. He has managed and participated in several European and national projects in France, Spain, and Germany. He has served as chair, TPC member, and reviewer for several conferences and journals. His research interests include cognitive and heterogeneous networks, self-organizing networks, spectrum and radio resource management, femtocells, and radio environment modeling.

JANNE RIIHIJÄRVI [M] is working as a senior research scientist at the Institute for Networked Systems at RWTH Aachen University. Before joining RWTH he worked in a variety of research projects on wireless networks at VTT Electronics and the Centre for Wireless Communications at the University of Oulu. His current research interests are in applications of techniques from spatial statistics and stochastic geometry on characterization of wireless networks, embedded intelligence in general, and design of architectures and protocols for large-scale heterogeneous networks.

BERNA SAYRAC received B.S., M.S., and Ph.D. degrees from the Department of Electrical and Electronics Engineering of Middle East Technical University (METU), Turkey, in 1990, 1992, and 1997, respectively. She worked as an assistant professor at METU between 2000 and 2001, and as a research scientist at Philips Research France between 2001 and 2002. Since 2002, she has worked as a research engineer at Orange Labs. Her research activities also include participating in European Commission funded research projects. Her research interests include self-organizing networks and cognitive radio. She holds several patents and has authored more than 50 peer-reviewed papers in prestigious journals and conferences. She also acts as expert evaluator, reviewer, TPC member, and guest editor for various conferences and journals.