

Big RF for Spectrum Sharing Applications

James Neel

Cognitive Radio Technologies, LLC
Lynchburg, VA, USA
james.neel@crtwireless.com

Shaswar Baban

Kings College London
London, UK
Shaswar.baban@kcl.ac.uk

Neal Mellen

Wireless Spectrum Management, LLC
Tempe, AZ, USA
neal@wsmgmnt.com

Ihsan Akbar

Harris Corporation
Lynchburg, VA, USA
iakbar@harris.com

Charles Sheehe

NASA
Cleveland, OH, USA
Charles.J.Sheehe@NASA.gov

Peter Cook

Peter G. Cook Consultancy
Mesa, AZ, USA
pgcook@pgcook.com

Abstract— In the 3550 MHz band, the FCC proposed the use of Spectrum Access Systems (SAS) to manage interference between commercial and federal systems and between secondary systems. To guide the interference management, the SAS will have to gather significant amounts of real-time data across these networks to detect and analyze potential interference issues as well as possibly determine the presence of mobile incumbent systems and to make appropriate exclusion zone adjustments.

Turning this large set of highly dynamic data into actionable information will require the use of sophisticated analysis techniques, such as those proposed in Big Data. More generally, Big Data techniques are critical components of an increasing number of wireless applications, from network management and planning to radio resource optimization.

This paper reviews emerging commercial applications of Big Data to RF problems ("Big RF"), describes available tools and architectural implications for leveraging Big RF in cognitive radio applications, and places Big RF in the broader context of the evolution of wireless networks.

I. INTRODUCTION

Across the technology space, researchers are turning to Big Data to help solve some of the most pressing problems. Big Data is being used to help cure cancer [1], to reduce traffic jams [2] and to gain an edge in financial transactions [3]. The Cognitive Radio Work Group (CRWG) in the Wireless Innovation Forum (WInnF) believe that Big Data tools and techniques can similarly be applied to address wireless issues and to help facilitate the design and deployment of cognitive radio applications. The CRWG has adopted the term "Big RF" to refer to the class of applications and tools that arise from bringing Big Data to bear on wireless problems.

A. Big Data

Big Data is a collection of tools and techniques that help analysts pull out actionable information - trends, predictions, characteristics, markers - from a sea of data. Commonly, large complex data sets are distilled into easily understood graphics that aid the decision processes of C-level executives. Big Data tends to differ from more traditional analytics by the scale and

nature of the data and is generally distinguished by the following "Four Vs" of Big Data.[4]

- **Volume** - the scale of stored data elements is too large for traditional processing to handle
- **Velocity** - new (streaming) data is arriving at a rate that on its own would be problematic
- **Variety** - insights are gained by combining data from multiple different sources with different syntaxes and semantics
- **Veracity** - decisions must be made under uncertainty due to physical limitations in data collection or flaws or variances in processes for maintaining data quality.

B. Relationship Between Big RF and Cognitive Radio

Many different wireless applications quickly come to mind that have these same characteristics - realizing the national spectrum dashboard [5], the potential management and optimization of millions of spectrum sharing devices, detecting and mitigating flaws in large-scale network deployments. While one can conceive of this information being presented to wireless executives, the decision processes of cognitive radios and cognitive networks could also be improved by having access to information generated from Big Data (or Big RF) analyses.

While many traditional, simpler, reactive cognitive radio algorithms would see little value in such a capability (e.g., choose the channel with the least interference, access the TVWS spectrum allowed at the specified location), *context-aware* cognitive radio (CACR) applications generally require insights beyond the capabilities of a single radio (or subnet). Rather than solely reacting to instantaneous performance measures, CACR applications attempt to optimize based on a broader understanding - what is the user trying to do, what is the mission of the network, what social information is available? [6]

More generally, many potential insights that a cognitive radio could leverage require processing capabilities and access to data far beyond the capacity of a single device. Similarly,

many insights that a single cognitive radio might find desirable will also be useful to other cognitive radios, e.g., the locations of known coverage holes, and could be used for better network management. Further, some insights have more applicability at a network level (e.g., cognitive networks), but extracting the appropriate actionable information from the sea of data is a non-trivial process.

The CRWG previously examined how Big Data (Big RF) techniques can enable CACR and other CR applications [6][7] and developed the block diagram for a Big RF system shown in Figure 1. Leveraging a wide variety of wireless and non-wireless data sources, the system performs a mix of streaming and batch analyses to provide insights to humans, other systems, and cognitive radios subscribed to the service.

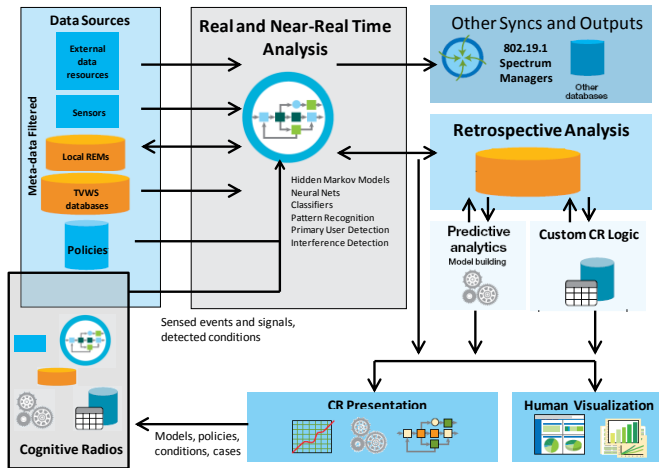


Figure 1: Major Big RF Processes. [7]

Note that unlike the traditional view of a Big Data system [4], in a Big RF-enabled CR application, the decision maker for the system (the CR or network) is also a potential information source for the system. While some Big RF applications may not exhibit this characteristic (e.g., a regulator viewing the analyses from the Spectrum Dashboard), the existence of this feedback loop has two important implications for the deployment of Big RF CR applications.

First, because a CR is frequently responding to very-short-lived phenomena (ms to seconds), the time available for analysis may be very short. Thus the design of a Big RF system for CR applications should incorporate techniques for reducing latency. Second, the nonlinearities in the analytics (e.g., the various machine learning techniques) combined with the feedback loop of CR to analysis to CR could produce unexpected behaviors. Accordingly, designers of CR-enabled Big RF systems should study the stability of the system or incorporate mechanisms to limit the available actions of the system (e.g., policy constraints).

C. Related Efforts and Past Work in CRWG

Several different efforts have examined the collection and synthesis of RF domain information including the Radio Environment Map at Virginia Tech (VT) [8], the Orange Radio Environment Map [9], Faramir (<http://www.ict-faramir.eu/>), and DARPA (Radio Map) [10].

Over the last couple of years, the CRWG has undertaken a project to examine the role of context in cognitive radio and how the explosive growth of information systems could be leveraged to develop new CR applications to improve wireless network performance. Previous publications developed a general model of how information is communicated [11], how informational contextual elements are communicated [12], techniques to enable contextually aware CR applications [6], and described applications to homeland security [7].

D. Document Organization

The remainder of this document is organized as follows. Section 2 discusses potential applications of Big RF to enhance the management and performance of commercial wireless networks. Section 3 describes techniques and tools for implementing Big RF and Section 4 considers architectural implications of Big RF. Section 5 reviews the trends that are pushing the wireless industry towards Big RF.

II. BIG DATA AND BIG RF TO ENHANCE COMMERCIAL WIRELESS NETWORK MANAGEMENT

Big Data can provide actionable insights into large complex data sets, such as is gathered during the operation of a wireless network. This section describes how Big Data / Big RF techniques are being used to improve existing networks, could how the performance of near-term spectrum-sharing networks could be improved, and how networks in the 5-10 year horizon could incorporate these techniques.

A. Improving Management of Existing Networks

Several companies have turned to sophisticated data analysis to improve existing network operations. First, T-Mobile has published that they use Big Data techniques to identify their most "important" users and then adapt their network operations to ensure that these users experience the best possible network quality. [13]. Second, companies such as Intucell (since acquired by Cisco) have shown that by collecting significant amounts of data on call quality and performance that network problems such as coverage holes and dropped calls can be proactively and automatically identified and addressed thereby improving network performance for everyone. [14]

B. Enabling Active Commercial Spectrum Sharing

Spurred on by the publication of the PCAST Report, the United States Government led by the Federal Communications Commission and the National Telecommunications and Information Administration have begun opening up bands once reserved principally for federal government use for commercial operations. As typified by the rules proposed for adoption in 3550-3700 MHz band [15], there are a couple critical differences between these bands and the earlier TV White Space efforts [16] that make this band better suited for Big Data-based management.

First, the incumbents are not static, requiring more real-time information to be incorporated, perhaps via sensors, to determine spectrum availability. While in the 3550 band, incumbent mobility is effectively limited to the sea (thereby

dramatically limiting the scope of the problem [17]), this is unlikely to be the case when spectrum sharing is employed in the 1755 MHz bands (AWS-3 / SGLS bands) where federal mobile terrestrial and airborne systems currently operate. Second rather than enforcing relatively fixed purely location-based exclusion zones for interference management, the FNPRM proposes that the Spectrum Access Systems (SAS) tasked with managing these systems protect incumbents and licensed users based on actual and predicted aggregate interference levels rather than purely based on device locations. While this has the potential to significantly increase spectral efficiency, it requires the repeated use and updating of detailed propagation models and sensors that must adjust as the devices move.

C. Enabling Future Wireless Networks

Big-Data-like approaches will also likely be needed to enable networks that will be deployed further into the future, even when interference-level based spectrum sharing with mobile incumbents is not necessarily a core function. For instance, the FCC recently issued a Notice of Inquiry requesting input into how to best manage spectrum above 24 GHz and to describe what kinds of applications could be enabled in that band. [18]. Often considered as a possible candidate for 5G networks, these high frequencies employed by these networks tend to have very small ranges with omni-directional antennas and instead will likely be deployed as collections of "hotspots" with pencil-width like beams adaptively providing very high throughputs to very narrow coverage areas. Frequently it will be the case that users could be serviced by many different beams from many different base stations at the same location. Depending on how these ultra-narrow beams are managed, they could be a significant benefit (ala cooperative MIMO) or a significant detriment (interference between beams intended for different users in an exaggerated form of cell-piercing). The real-time coordination across potentially hundreds of such devices supporting thousands of mobile users in complex propagation environments will require significant data processing power.

III. BIG RF DEPLOYMENT CONSIDERATIONS

Commercial communications architectures are built to maximize the profit from their expected use, which may evolve as new uses and technologies are needed. While profits come from the use of the network (e.g., data from cell phone, laptop, Bluetooth headsets, and remote sensing devices like home alarm systems), the networks must also accommodate the overhead associated with flow of management data. A goal of any communications network architecture, then, is to maximize the use of the network while minimizing the amount of management data that the network carries.

A. Satisfying Competing Objectives

One way to satisfy these competing objectives is to push as much of the management processes as possible to the edge of the network. This has a related benefit in reducing latencies thereby enabling the management of real-time phenomena and traffic - important for cognitive radio applications. This highlights an important distinction between traditional Big

Data analytics and Big RF applications - many Big RF applications are applied to the subset of real time information concerning the use and loading of the Radio Frequency (RF) environment to actively manage the interactions between the users and the networks.

To illustrate an application of Big RF to edge management in communications network architectures, consider a hypothetical cellular network depicted in Figure 2. Large networks cover tens of thousands of square miles areas with high population density, such as Cleveland, Ohio, and have users who expect their devices to work. The population of a city moves, performs a variety of activities and congregates in different areas. Management of network events that are cyclical in nature can be analyzed by Big Data Analytics to enhance the network's management and control. Currently, in preparation of these anticipated events, such as large sporting events, service providers deploy numerous temporary base stations to handle the expected sustained surge in demand.

However, there are also unexpected events, like a flash mob, that also lead to significant surges in traffic, but with less time to respond or anticipate the event's onset. In this case, the inspection of information packet headers (e.g., data, text, video) or by advanced context reasoning over the information flowing through the network at a given moment (e.g., tweets, posts), a Big RF Manager could rapidly cache popular information at local edge servers to decrease latency and reduce demand on the network backbone. Or in a disaster-response scenario [19], priority could be given to first responders while context-aware devices in conjunction with real-time Big RF management could negotiate alternatives while maintaining network resource availability, though this would require many changes to existing network management practices.




Figure 2: A Honeycomb cell structure over Cleveland Ohio with a Big RF Manager.¹

¹ Image used with permission of Focal Plane Photography, LLC 4078 Elmwood Rd South Euclid Ohio 44124 Use of base photo: "Downtown Cleveland Ohio vertical aerial" <https://www.focal-plane.com/large-area-mapping.html>.

B. Hybrid Architecture

To achieve the sub 1 ms latency, as advocated for in the realm of next generation 5G networks, [22] has suggested that some of the core network and Internet related delays could be reduced by bringing content closer to the consumer. This is usually done through caching, and in this case, through intelligent caching [23]. However, to provide this sort of intelligence reliably and in real-time, tracking and intensive computation of end-users behavior will be required to predict the behavior of masses of people, and maybe machines as well.

Conversely, there are many other applications that are not as sensitive to latency and/or may require visibility into a wider array of information sources. In these cases, caching services at the edge may be less efficient. To handle both cases, a hybrid architecture could be adopted.

An architecture depicting the mix of user devices, base stations, sensors, servers, routers, Big RF and Big Data analysis routines is shown in Figure 3. In this view, the architecture is hierarchical as more advanced analytics for low latency applications (Big RF) or higher-latency applications (Big Data) are implemented at different levels in the design. Similarly, cell-towers continue implementing their existing algorithms while hand-held electronics implement contextually-aware applications. Combined, this architecture leverages the network information and the context of that information to improve the user experience by automatically adjusting radio resources and caching popular data, in a manner similar to how some search engines monitor and adapt to searches in real-time.

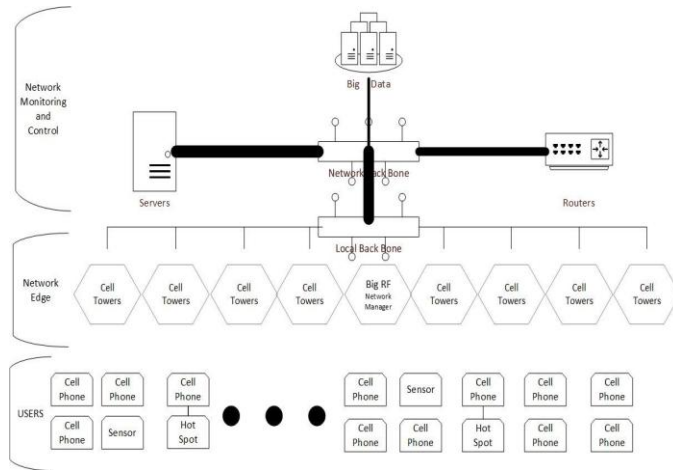


Figure 3: The Relationships Between the Network Elements

IV. BIG DATA TOOLS FOR THE HYBRID ARCHITECTURE

Several tools have been developed, primarily to leverage cluster computing, that could allow portions of Big Data computations to be performed closer to the network edge, thereby opening up the power of Big Data for cognitive radio and enabling the hybrid architecture of Figure 3. This Section examines Apache Spark as an example cluster computing

system among others, e.g., Apache Hadoop or Apache Storm, that could be applied for this purpose. It first briefly explains the Spark ecosystem and architecture and then describes the distributed transformations Spark applies to the input raw data, stage-by-stage, to produce actionable knowledge.

A. Apache Spark Overview

Apache Spark is an open-source, easy to use, fast and general data processing engine, originally developed in the AMPLab at UC Berkeley [20]. Unlike Hadoop's disk-based MapReduce runtime algorithm [21], Spark uses in-memory processing, leading to faster processing time, especially for iterative processes such as machine learning algorithms. In contrast to Storm, Spark also allows for micro-batch and stream cluster computing, bringing together the best of both processing worlds (batch and stream) in one place.

B. Spark Ecosystem

The SPARK Ecosystem consists of the following components.

- **Spark Core Engine:** The kernel of the Spark system containing the main routines to execute the overlaying packages.
- **Spark SQL:** Data can be queried as a distributed dataset (RDD) through Spark SQL, which also facilitates running SQL queries alongside complex analytic algorithms.
- **Spark Streaming:** Enables Spark to run streaming jobs for applications that require real-time response.
- **MLlib:** The Machine Learning libraries (MLlib) package is interoperable with NumPy in Python and could be used to build CR Big Data algorithms by adapting the machine learning processes to wireless domain problems.
- **GraphX:** A graph computing package to build, manipulate and reason about graphs, such as may be useful in reasoning about social networks.
- **YARN (Yet Another Resource Negotiator):** An Apache project within the Hadoop ecosystem used for managing the computing cluster.
- **HDFS (Hadoop Distributed File System):** A distributed and scalable file-system used to store large files.

C. Big RF Deployment with SPARK

To put the Spark cluster computing into next-generation communication systems context, Figure 4 depicts an example Big Data-enabled cellular networking architecture, designed to support the ambitious less-than-1 ms round-trip-time (RTT) latency performance. The envisioned system includes a number of sensors and mobile phones connecting via a Remote Radio Unit (RRU) to the network. Through the application of Big RF analytics, relevant content can be identified and cached on local data stores at the Baseband Unit (BBU), which is a pool of digital processing units each of which were previously located at different basestations. Thereby, each BBU can be considered a cloudlet within the Cloud RAN [24] subsystem. This architecture leverages the network information and the context of that information to improve the user experience by automatically adjusting radio resources and caching popular data, in a manner similar to

how some search engines monitor and adapt to searches in real-time.

The Big RF enabling architecture in Figure 4 has the following three advantages over a more traditional design.

- **Offloading computation to the local BBU:** Information processing on the cached data offloads communication and processing overhead that would otherwise be required if one main cloud computing system was used.
- **Enhanced Security and Privacy:** Localizing information reduces the attractiveness of any one target and reduces the number of potentially interceptable information transfers.
- **Reduced Latency:** For those applications that require only local communications, running algorithms locally on the data cached at the BBU reduces the end-to-end RTT delay.

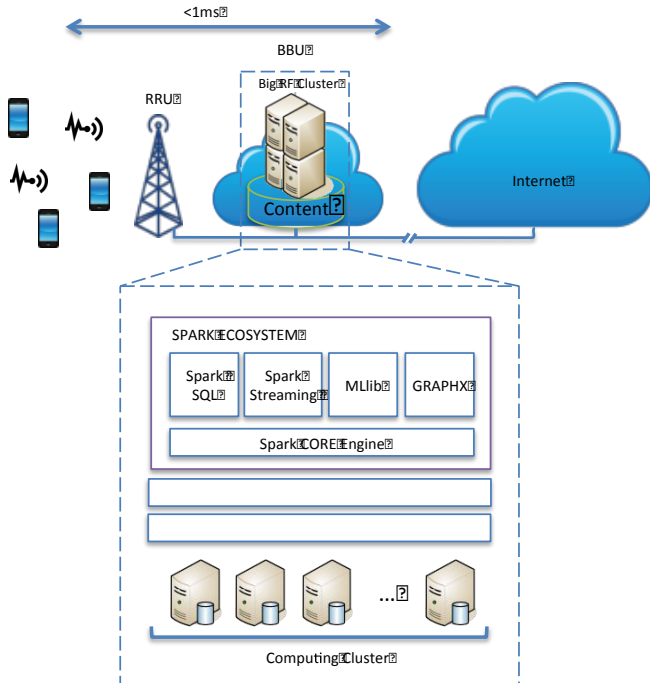


Figure 4: Smart Caching at the Edge based on Big RF analytics. Adapted from [22]

V. TRENDS TOWARDS BIG RF

This Section considers how Big Data and Big RF can be integrated into the operation of wireless systems. To provide context, the Section first looks at the history and evolution of cellular wireless systems. Then the Section presents key enabling trends to deploying Big RF.

A. Evolution of Cellular Systems

Broadly, cellular systems (and communications systems in general) progress through a series of overlapping technology waves (1G, 2G, 3G, 4G, 5G). Each wave progresses through a sequence of stages wherein the technology is researched, developed, and matured with significant gains occurring as we learn how to best optimize these technologies while maintaining compatibility with older technologies still in the field. However, each wave is composed of smaller waves (e.g., 3.75G, GPRS, or EDGE) as sub-technologies follow a similar pattern of R&D, adoption, and deployment. For

instance the cellular generational waves were themselves defined by more specific technology progressions, such as FDMA to TDMA to SC-OFDMA, GMSK to OFDM, circuit-switched to packet switched network core, flip phones to smart phones, single antennas to sectorized antennas to MIMO systems, and self-contained networks to hybrid data offloading networks (e.g., WiFi). These waves of cellular technologies are illustrated in Figure 5. Note that in this model, Big RF applications are beginning to emerge now, but we predict will first see significant deployments as part of 5G.




Figure 5: The Evolution of Digital Cellular Standards

B. Trends Supporting the Development of Big RF

Much research is currently focused on increasing capacity. While PHY technology improvements have a role to play, most capacity gains are foreseen as coming from the following trends: decreased cell size (e.g., the move to femto- and attocells), increasingly larger antenna arrays, gaining access to additional spectrum (e.g., AWS-3), and leveraging heterogeneous and opportunistic spectrum (e.g., WiFi offloading, TVWS, and the 3550 MHz band). The need to support legacy and as yet to be defined standards have in turn led to more frequency agile base stations and user equipment including software defined handsets and base stations, that simplify upgrades and increase flexibility. These trends by themselves, however, will not be sufficient to support future data transmission requirements.

At the same time, patterns of usage continue to evolve as an increasing fraction of wireless transmissions come from relatively static, indoor users, which has led to an increasing use of distributed antenna systems (DAS), which are a step towards the processing offloading envisioned for Big RF. Likewise, the use of opportunistic spectrum with increasingly tighter sharing between primary and secondary users is leading to more sophisticated network management systems that will need to process amounts of data from sensors and user equipment, thereby placing a greater need for Big RF applications. Finally, the increasing support for context-aware applications, from social networking APIs to context-enabling chipsets [25] can provide a rich source of data to enable Big RF applications.

Taking an integrated view of these trends, exponentially increasing number of cells require exponentially escalating

sophistication to manage. Simply performing the management of these future networks will necessitate adopting a Big RF-type solution, and likely one where offloading becomes an integral part of the network, relieving it of data transmission and management.

VI. CONCLUSIONS

The wireless world is evolving to a point where the volume, variety, velocity, and veracity of data needed to enable cognitive radio and other wireless applications requires Big-Data solutions. This will arrive as an evolution of existing wireless technologies as White Space Databases and 802.19.1-style spectrum managers expand their scope to handle more responsibilities and incorporate additional data. Similarly, cellular networks and public safety systems are expanding to leverage the wealth of data to improve every day performance and to better handle unexpected situations such as during disaster response. At the same time, as wireless communications is relied on to enable the control and management of larger-scale systems (e.g., Smart Grid), their RF domain issues will similarly require Big RF solutions.

While these analyses of Big Data tools can be used by engineers or executives to improve network design and operation, we believe the greatest gains will come from making these analyses available to cognitive radios and cognitive networks. However, because of the time-scales CRs typically operate on, many of these applications will also have to contend with tight latency requirements. To address this, this paper has proposed adopting a hybrid distributed-cloud architecture where services and data caches for real-time applications are placed closer to the edge while applications with a broader array of data sources and longer timelines can be deployed in the core of the network. This paper noted that some Big Data tools, such as Spark, appear to be well suited for such an architecture, and described how these tools would fit into the envisioned hybrid architecture.

The realization and implementation of Big RF and CACR will not occur at once; rather it will be the result of a combination of improvements along several different threads of ongoing technology developments. The existing threads of diverse spectrum allocations, smaller cell size, network and handset hardware consolidation will be a contributing factor. The newer threads of Big RF Spectrum Management and network and spectrum virtualization will be the deciding factor on actual implementation.

There are several recent developments that, if made available for rapid analysis and action, could ameliorate the implementation by simplifying the development of Big RF systems. A few examples are receiver harm thresholds, more advanced propagation models, and geo-located measurements from deployed radios. These, and others, could be used to create sophisticated database-driven systems that make more accurate real-time predictions of interference thereby allowing closer spectrum packing. Due to the mobility and scale of such a solution, Big RF would be needed to handle the analysis and data collection.

While offering many opportunities for improved network performance and insights, application enabled through Big RF should be cognizant of the some implementation challenges. The security and backhaul are two limiting factors in most networks and deploying all of the analysis and database resources will be problematic for real-time adaptations. Fortunately, as most data, traffic, and exploitable correlations are, at most, regional phenomena, local reach back may suffice for most Big RF problems. Further gains could be achieved by prioritizing which data is sent back along with information compression, ontologies, and context-aware transmissions to further reduce the bandwidth and security demands.

VII. REFERENCES

- [1] M. Helft, "Can Big Data Cure Cancer?", *Fortune*, July 24, 2014. Available online: <http://fortune.com/2014/07/24/can-big-data-cure-cancer/>
- [2] E. Epstein, "Less Traffic on Your Commute? Thank Big Data," *Mashable*, Oct 1, 2014. Available online: <http://mashable.com/2014/10/01/data-traffic-jams/>
- [3] I. Aldridge, "Trends: All Finance Will Soon Be Big Data Finance," *Huffington Post*, Feb 4, 2015. Available online: http://www.huffingtonpost.com/irene-aldrige/trends-all-finance-will-s_b_6613138.html
- [4] IBM Corp, "Big data analytics for communications service providers," white paper, Oct. 2012.
- [5] FCC Spectrum Dashboard, <http://reboot.fcc.gov/spectrumdashboard/>
- [6] J. Neel, P. Cook N. Mellen, I. Akbar, D. Devasirvatham, C. Sheehe, R. Schutz, "The Role of Context in Cognitive Systems," *Journal of Signal Processing Systems*, April 2014.
- [7] J. Neel, S. Baban, N. Mellen, I. Akbar, C. Sheehe, B. Schutz, P. Cook, "Big RF for Homeland Security Applications", *IEEE HST Symposium*, 2015, April 2015.
- [8] Y. Zhao, D. Raymond, C. da Silva, J. Reed, S. Midkiff, "Performance Evaluation of Radio Environment Map-Enabled Cognitive Radios Coexisting with Incumbent Users," *MILCOM 2007*.
- [9] B. Sayrac, et al, "Radio Environmental Maps (REMs): A Cognitive Tool for Environmental Awareness," Talk at CTTC, 2011.
- [10] Advanced RF Mapping (Radio Map) [http://www.darpa.mil/Our_Work/STO/Programs/Advanced_RF_Mapping_\(Radio_Map\).aspx](http://www.darpa.mil/Our_Work/STO/Programs/Advanced_RF_Mapping_(Radio_Map).aspx)
- [11] Wireless Innovation Forum, "IPA – Information Process Architecture Volume I," WINNF-09-P-0020-V1.0.0, Nov 01, 2010.
- [12] Wireless Innovation Forum, "Information Process Architecture Volume 2: Survey of IPA Like Systems," WINNF-09-P-0021, Jun 27 2012.
- [13] <http://www.lightreading.com/analytics/big-data/big-data-saves-t-mobile-big-bucks/d/d-id/709611>
- [14] <http://techcrunch.com/2013/01/23/cisco-buys-israeli-mobile-startup-intucell-for-475-million-cash-and-bessemer-gets-nearly-half/>
- [15] FCC, Further Notice of Proposed Rule Making, Amendment of the Commission's Rules with Regard to Commercial Operations in the 3550-3650 MHz Band, GN Docket No. 12-354.
- [16] <http://www.fcc.gov/encyclopedia/white-space-database-administration>
- [17] Comments of Google, VT, and Federated Wireless to FCC in Docket 12-354, September 3, 2014.
- [18] Federal Communications Commission, "Notice of Inquiry" In the Matter of User of Spectrum Bands Above 24 GHz for Mobile Radio Services (GN Docket No. 14-177)" FCC 14-154, Oct 17, 2014.
- [19] SDR Forum, "Use Cases for Cognitive Applications in Public Safety Communications Systems - Volume 1: Review of the 7 July Bombing of the London Underground," SDRF-07-P-0019-V1.0.0
- [20] M. Zaharia, M. Chowdhury, M. Franklin, S. Shenker, I. Stoica, "Spark: cluster computing with working sets." Proceedings of the 2nd USENIX Conference in Hot Topics in Cloud Computing, 2010.
- [21] J. Dean and S. Ghemawat. "MapReduce: simplified data processing on large clusters." *Communications of the ACM* 51.1 (2008): 107-113.
- [22] GSMA, "Understanding 5G: Perspectives on future technological advancements in mobile", Dec. 2014, Online at: <https://gsmaintelligence.com/files/analysis/?file=141208-5g.pdf>

[23] Bastug, Ejder, Mehdi Bennis, and Mérouane Debbah. "Living on the edge: The role of proactive caching in 5g wireless networks." *Communications Magazine*, IEEE 52.8 (2014): 82-89.

[24] "C-RAN: the road towards green RAN", China Mobile Research Institute, Technical Report, v3.0, Dec 2013.

[25] <https://developer.qualcomm.com/mobile-development/mobile-technologies/context-aware-gimbal>