

Feature Interview with Ranveer Chandra: Author of ‘White Space Networking with Wi-Fi like Connectivity’, Best SIGCOMM Paper 2009

Joseph Camp

Department of Electrical Engineering, Southern Methodist University (SMU), Dallas, TX
camp@smu.edu

ABSTRACT

Ranveer Chandra is a Senior Researcher in the Mobility & Networking Research Group at Microsoft Research. His research is focused on mobile devices, with particular emphasis on wireless communications and energy efficiency. Ranveer is leading the white space networking project at Microsoft Research. He was invited to the FCC to present his work, and spectrum regulators from India, China, Brazil, Singapore and US (including the FCC chairman) have visited the Microsoft campus to see his deployment of the world’s first urban white space network. The following interview captures the essence of his work on white spaces by focusing on his work published in ACM SIGCOMM 2009, which received the Best Paper Award.¹

1. INTERVIEW

1. *When did you start working on this project?*

We started working on the KNOWS (Kognitive Networking over White Spaces) project in 2005. We first built a prototype radio, and submitted it to the FCC. We published an invited paper on this design in IEEE LANMAN 2007 [1]. We kept working on the concept of variable channel widths, which is important for obtaining good performance over the fragmented white space spectrum. This was the SIGCOMM 2008 paper [2]. All of these were building blocks towards the SIGCOMM 2009 paper [3] where we took all the concepts and showed how you could build a system that did Wi-Fi (i.e., what minimal changes are needed to make Wi-Fi work) over the white space frequencies. We called our system WhiteFi, which got a lot of publicity/press.

2. *Where has the work gone? For example, what advances have been made since the paper’s writing?*

The SIGCOMM work has gone in two directions: towards productization and policy impact. The IEEE 802.11af standard picked up some of the key concepts from this paper. Microsoft filed to be a white space database provider, and right now 3 different database providers are certified by the FCC. The other direction in which this work has gone is to influence the regulatory policy on spectrum sharing. There was widespread skepticism back then on whether this type of white space networking was even feasible. The work we did was the first of its kind to show that you could build high-throughput, long-range data networks over TV White

Spaces. Both that you could do it, and how you could do it. This paper was submitted in January 2009, which overlapped with the first FCC ruling. We were passing our learnings on to the FCC before the paper was even published.

Since the work has been published, we have done additional research on different aspects of white space networking with a focus on making them real. In particular we have published papers on building white space networks without sensing (IEEE DySPAN 2011 [4]), deploying a real white space network on campus (MC2R [5]), coexisting with wireless MICs (CoNEXT 2011 [6]), coexistence among white space networks (CoNext 2012 [7]), and recently, on indoor white spaces in big cities published in ACM MobiCom 2013 [8].

The two most significant milestones were the database and the white space deployment. The first FCC white space ruling required devices to both sense the channel and query the database (2008 ruling). When we talked to chip companies, we learned that that an accurate sensing component was expensive, and prone to errors because sensing is very hard to achieve and very complicated. Thus, in follow-up work, we built a system that did not require sensing published in IEEE DySPAN in 2011 [4]. Instead, it would only query the geolocation database. The key question was how to make it conservative so that it does not interfere at all with the primary users – with the TV broadcasts or the wireless microphones. If you built such a system that was conservative, then how many white spaces do you lose, how many TV channels are you not able to use to form this kind of White Space network? We built that system which required this database that we built, a web service [9]. In parallel, we also built a campus WhiteFi testbed where we enabled a unique scenario over the TV white spaces. Microsoft Redmond is a big campus (1 mile x 1 mile), there are lots of buildings, and a shuttle that takes people from one building to another. When people sit inside these shuttles, they do not have Internet connectivity. You could connect using the cellular network, but you could not connect directly to corporate resources. So, the particular scenario we enabled was inside the shuttle, we would do WiFi, but from the shuttle to the buildings, we would hop over the White Spaces. The White Spaces were well suited for this application because they have much longer range than the higher frequencies used by Wi-Fi. This was a unique application that showcased the propagation benefits of white space frequencies as well.

3. *What hardware did you use for outside the shuttle on the Microsoft campus deployment?*

In this particular deployment, we made it very flexible. There

¹This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The authors take full responsibility for this article’s technical content. Comments can be posted through CCR Online.

were antennas on the rooftop and antennas on the shuttle. We had a plug-and-play system depending on what we needed. We had a WiFi-based system using the KNOWS platform and a WiMax-based system that used WiMax over the TV frequencies. We also tried devices from Adaptrum (a Silicon Valley based company) and 6Harmonics, which had proprietary protocols. We tried out many devices in a plug and play manner to see which fit the requirement. Now we are trying with Sora with the same setup to see how it works.

4. *What was the time frame for the outside testbed and all the visitors to the testbed?*

The testbed was follow-on work to the SIGCOMM 2009 paper [3]. We had the first setup ready in October 2009. Subsequently, we had visitors from all over the world come to see the deployment including the FCC Chairman Julius Genachowski (September 2010). We had the Telecom Regulatory Authority of India Chairman visit as well. It is one of the agencies in India that looks at spectrum regulations. We also had regulators from Brazil, Singapore, China, and other countries. Essentially, it was the first time people could see a real live white spaces network in action. Also in October 2010, the FCC passed a new ruling where it was no longer required to both sense the spectrum and consult with a database to use the spectrum. I don't think that the FCC chairman made up his mind because of the testbed, but our demo definitely played a role in showing such a system could work.

5. *Where has the work not gone? For example, what unforeseen challenges came up that prevented certain aspects of the proposed WhiteFi scheme?*

I wouldn't say the work hasn't gone there, but is going slowly. The first ASIC for doing white spaces was announced just a few months back by Neul, a Cambridge (UK) based company. Also, the standardization is taking time. They are moving fast, but still the IEEE 802.11af standard is not yet final. So, there are things which are happening, but they are happening at a slower pace. Another concept that we had used in the SIGCOMM paper was spectrum sensing. However, accurate, low-cost spectrum sensing is still an open research problem and not ready for primetime because of the error and cost. In follow-up work we moved away from having sensing as a requirement in the Senseless paper [4]. However, in our more recent work in MobiCom 2013 [8], we are revisiting the concept of spectrum sensing since, if done right, it can potentially unlock additional unused spectrum for white space networking.

6. *Why do you think it received the SIGCOMM best paper award?*

I think it received the award because of two reasons. First, it was solving a very current, difficult problem. Its impact was at the cutting edge of different fields: research, industry, and policy. It was the first time any research had shown how to build a network over the white spaces. Second, the learnings from this paper were fundamental for a larger problem space – that of forming networks over any dynamic spectrum access network. This is valid even today and might be a panacea for using other (non-TV) spectrum for dynamic spectrum access.

7. *How has the use of spectrum changed since the paper's writing?*

With a proliferation of mobile devices over the last few years there is a need for additional spectrum to service their data

demand. That said, this artificial spectrum scarcity is primarily created by existing regulatory policy. That is, there is enough spectrum, but not the way it is currently allocated. Since we wrote this paper, and in fact even before that, the FCC made the White Spaces ruling for the first time that allowed dynamic spectrum access for opportunistic data communications. Since the ruling, the entire White Spaces concept has taken off with multiple startups that are building these radios and databases for White Space spectrum. There are 9 authorized database providers, 3 have already been certified by the FCC: SpectrumBridge, Google and Telecordia (subsidiary of Ericsson). There is a lot of momentum that is happening around the TV White Spaces, and people are looking for ways to use it – that's all happened since the paper was written and now it's more mainstream for people to talk about spectrum sharing, not only in the context of the TV White Spaces, but in other parts of the spectrum as well. More recently, T-Mobile, for example, is trying to show that it can use some part of the spectrum based on spectrum sharing for its mobile network. More recently, the FCC issued another NPRM (Notice of Proposal for Rule Making) for the 3.5 GHz band. There again, part of this spectrum right now is allocated for radars and satellites. So, the FCC has asked for comments on how to reuse that spectrum for data communications. There were lots of people talking about a multi-tiered approach in which if the spectrum is not being used, how it can be used by other devices for data communication. Again, these are using the concepts that were proposed in the 2009 paper [3].

8. *What were some key observations from the deployment that the FCC chairman or any other visitors able to use in spectrum policy making?*

We demonstrated 3 different concepts. One was one of spectrum efficiency, i.e., we could pack a lot of bits in this spectrum and adjust the channel width based on demand, and you could use it for high speed communications. Second, we talked of spectrum agility. That is, you could move from one part of the spectrum to another. If you were in one part of the network, and a primary user showed up, you could immediately move to another part of the spectrum, without disrupting the ongoing data communication and at the same time, not interfering with the primary user's transmissions. This is one of the core principles behind dynamic spectrum access. The third concept we demonstrated was that of coexistence. This is follow-on work which we published at CoNEXT 2011 [6]. What we showed was how we could transmit data on the same TV channel on which a wireless microphone was operational without affecting the sound recording. This idea has been presented to (and subsequently by) the people who plan and deploy wireless microphones in big events such as the Super Bowl and Grammy's.

9. *You have now launched a similar network in Africa. Can you tell us about this deployment?*

We have been working with the Technology and Policy Group in Microsoft to kickstart various commercial (and non-commercial) international white space deployments, for example, in the UK, Singapore, Kenya, and others. The goal is to help regulators in these countries make use of this unused spectrum to enable applications that would otherwise have not been possible. I will be going to Malawi in a couple of weeks to help the regulators there as well. In Africa this is part

of a larger 4Africa initiative launched by Microsoft, to help extend the reach of Internet and technology to different portions of Africa. In these deployments, we have been helping with (i) building the database for these countries [10], and (ii) helping setup the deployment.

10. *What factors affect which areas are chosen in these deployments?*

We want the reach of this technology to be global. We want people everywhere, including developing countries to benefit from this. The work in Kenya for example is part of a larger 4Africa initiative, and I am going to help the government of Malawi set this up as well. In addition, we recently worked with the government of Singapore to launch a similar deployment. The goal is to share our learnings and experience with the policy and technology in the US to help accelerate the development of policy in these countries.

11. *What hardware and database is being used in these deployments?*

Each of these deployments are using our database, the Microsoft Research database. The first version was built by Rohan, one of our interns, and we have customized that database for each country such as Singapore and Kenya. For hardware, we are working with different startups, e.g. Neul, Adaptrum, in different countries for hardware.

12. *One of the key aspects of this paper was dealing with the network components of frequency agility in white spaces as opposed to some of the other previous works which is really looking at links (non-network aspects). Any other thoughts about networks of nodes and how that has transformed since the writing of this paper?*

One of the main contributions of the paper was, as you mentioned, this networking concept. Before this paper, most papers were about how to build the radios and how does that radio senses the spectrum. We showed three different ways in which white spaces networks are different from traditional networks: spatial variation, temporal variation, and fragmentation. These three components have implications across the board on the networking aspects of discovery, on spectrum allocation, and how you handle disconnections. We proposed ways to handle these factors. At a high level the question we answered was: if I bought an access point and put it in my house, how does the network get setup in a way that gives the best performance (throughput, non-disruptive, etc.)? For example, there are so many channels available, which one should I decide to use? The widest channel might not be the best channel. We proposed a metric to make the decision. There have been a number of papers that have developed on that metric as well. These techniques are fundamental and will impact the deployment of networks using dynamic spectrum access in any band – TV and others.

13. *You talked in the work about how multiple APs would coordinate with one another. Any final thoughts on what has gone on with multi-AP and coordination? The databases solve that to some degree, correct?*

That is still an open question as to how you solve the co-existence problem. The co-existence problem is more severe in the TV white spaces because the signal propagates much better. So, more devices interfere with one another and are likely to trample each other's network. One way in which we

tried to solve it was a modification to 802.11 with adaptive preambles, with lengths to signal different distances to not trample over one another [7]. We called our system Weeble (CoNEXT 2012). Another issue is how to use the database to coordinate different access points where there are different networks by different operators run by different individuals that might have little to no incentive to cooperate. That's what makes this problem very interesting. WiFi is the only mainstream example where you have different devices operated by different entities that coexist. For example, in the Technology For All (TFA) Network in Houston you have been using Wi-Fi and you have been using TV channels, but it is all operated by you. How do you do this with different operators at the same time? What is the right model for sharing? There have been a few solutions using multiple access points with White Spaces, but I feel that the solution has not been nailed.

14. *Do you have any last thoughts or things you would like to add?*

We are currently working on two different directions. First, with white spaces, we are deploying many more trials in different countries. We are particularly interested in co-existence in terms of how to build a larger network. Second, we are working on dynamic spectrum access in non-TV white spaces (e.g., 3.5 GHz or NTIA band). For example, if you were to share the spectrum with radars or satellites, then what should be the spectrum access and sharing policies? The first step to this problem is just trying to observe spectrum occupancy. Through a project called the Microsoft Spectrum Observatory, we are deploying measurement points where we are monitoring the frequency usage from 30 MHz to 6 GHz, constantly scanning this spectrum, and identifying which portions of the spectrum are not actively used. We had an Allerton paper last year where we discussed challenges and characterized how different portions of the spectrum can be used for dynamic spectrum access [11]. We hope to publish more results on this research in the coming year.

2. REFERENCES

- [1] Y. Yuan, P. Bahl, R. Chandra, P. Chou, J. Ferrell, T. Moscibroda, S. Narlanka, and Y. Wu, "Knows: Cognitive radio networks over white spaces," in *IEEE DySPAN*, Apr. 2007.
- [2] R. Chandra, R. Mahajan, T. Moscibroda, R. Raghavendra, and P. Bahl, "A case for adapting channel width in wireless networks," in *ACM SIGCOMM*, Aug. 2008.
- [3] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh, "White space networking with Wi-Fi like connectivity," in *ACM SIGCOMM*, Aug. 2009.
- [4] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl, "SenseLess: A Database-Driven White Spaces Network," in *IEEE DySPAN*, May 2011.
- [5] R. Chandra, T. Moscibroda, P. Bahl, R. Murty, G. Nychis, and X. Wang, "A campus-wide testbed over the tv white spaces," *Mobile Computing and Communications Review*, vol. 15, no. 3, pp. 2–9, July 2011.
- [6] G. Nychis, R. Chandra, T. Moscibroda, I. Tashev, and P. Steenkiste, "Reclaiming the white spaces: Spectrum efficient coexistence with primary users," in *ACM CoNEXT*, Dec. 2011.
- [7] B. Radunovic, R. Chandra, and D. Gunawardena, "Weeble: Enabling low-power nodes to coexist with high-power nodes in white space networks," in *ACM CoNEXT*, Dec. 2012.
- [8] X. Ying, J. Zhang, L. Yan, G. Zhang, M. Chen, and R. Chandra, "Exploring indoor white spaces in metropolises," in *ACM MobiCom*, Oct. 2013.
- [9] "White Space Finder," <http://whitespaces.msresearch.us>, July 2013.
- [10] "Microsoft Research White Space Finder - a Singapore Trail," <http://whitespaces.cloudapp.net>, July 2013.
- [11] A. Chowdhery, R. Chandra, P. Garnett, and P. Mitchell, "Characterizing spectrum goodness for dynamic spectrum access," in *IEEE Allerton*, Oct. 2012.