

D-BigBand: Sensing GHz-Wide Non-Sparse Spectrum on Commodity Radios

Lixin Shi, Haitham Hassanieh, and Dina Katab
Massachusetts Institute of Technology
{lixin, haitham, dina}@csail.mit.edu

ABSTRACT

This paper presents D-BigBand, a system that senses a GHz-wide spectrum in real time using ADCs sampling at only tens of MS/s speed. It is an advanced version of our previous work, BigBand, which senses GHz-wide sparse spectrum using commodity radios. However, unlike BigBand that requires the spectrum to be sparse, D-BigBand does not assume sparsity utilization of the spectrum. This is particularly important for the spectrum under 2GHz which is typically not sparse, crowded with different wireless technologies like TV broadcasting, mobile, AM radios, etc. The key idea of D-BigBand is to run sparse recovery on the changes of the spectrum: since only a small fraction of the spectrum is likely to change its occupancy over short intervals of a few milliseconds, the changes of the spectrum is sparse and we can apply sparse recovery to it. Our evaluation shows that D-BigBand works even if 95% of the spectrum is occupied.

Keywords Spectrum Sensing; Sparse Fourier Transform; Wireless; ADC; Software Radios

1. INTRODUCTION

Spectrum sensing has been a recurring topic in the past two decades, not only in the research community [1], but also in the government [4] and the industry [9]. Specifically, the recent concerns of spectrum crisis [3] and the plan to open up GHz-wide of spectrum by the FCC [11] has driven the community to look at real-time wideband spectrum sensing at the bandwidth of multi-GHz [7].

Realtime GHz signal sensing, however, is challenging. GHz-wide bandwidth requires high-speed analog-to-digital converters (ADCs), which are costly and power hungry, and have a low bit resolution [10]. Instead, typical spectrum sensing platforms like Microsoft Observatory [9] sequentially scan the spectrum; they hop from one band to the next, sensing only tens of MHz at any point in time. As a result, each band is monitored only occasionally, making it easy to miss short lived signals (e.g., radar).

Our previous work, BigBand [7], solves this problem by assuming and leveraging the sparsity in the spectrum. BigBand captures GHz of spectrum in real time but uses only a few ADCs that each

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

53'14, September 7, 2014, Maui, Hawaii, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3073-2/14/09 ...\$15.00.

<http://dx.doi.org/10.1145/2645884.2645887>.

samples the signal at tens of MS/s. To achieve this goal, it builds on advances in the area of sparse Fourier transform [6, 5], which permit signals whose frequency domain representation is sparse to be recovered using only a small subset of their samples — i.e., BigBand can recover GHz of spectrum without sampling it at the Nyquist rate.

In this paper, we extend BigBand to D-BigBand, which senses non-sparse spectrums when the spectrum utilization is dense. According to multiple spectrum reports [8], the spectrum under 2GHz, is especially crowded with various wireless services like TV stations, LTEs, AM radios, etc. Sensing and monitoring the spectrum utilization for these bands is equally important as for sensing spectrum at higher frequencies which tends to be sparse, but sparsity-based techniques such as BigBand is likely to fail at spectrum under 2GHz because of lack of sparsity. For example, The ADC speed required by BigBand is proportional to the spectrum sparsity, which means under 2GHz, it falls back to use GHz ADCs in order to sense these non-sparse spectrum.

The key idea behind our system, D-BigBand, is that even if the spectrum itself is densely occupied, only a small fraction of the spectrum is likely to change its occupancy over short intervals of time (e.g., a few milliseconds). In other words, the difference of spectrum occupancy within a small time window is still sparse, and we can use sparse Fourier transform, the same technique that BigBand is based on, to track the sparse changes in the spectrum.¹ D-BigBand builds on this basic idea to sense densely occupied spectrum using sub-Nyquist sampling. We also evaluate our design empirically showing that it can detect frequency bands that change occupancy even when the spectrum is 95% occupied.

2. PRIMER ON BIGBAND

Since D-BigBand leverages sparse Fourier transform, which is the same sparsity recovery principle as in BigBand, we provide a brief primer on how BigBand recovers sparse spectrum. For more details, we encourage the readers to refer to BigBand [7]. At a high level, the sparse recovery consists of two steps:

- **Bucketization:** BigBand bucketizes the wide frequency spectrum into different buckets, so each of the bucket contains only one or a very small number of non-zero frequencies. The way of bucketization in BigBand is using low-speed ADCs: 1) Each ADC samples at sub-Nyquist rate, which creates aliasing in the frequency domain; 2) aliasing adds up frequencies which effectively creates buckets; 3) the frequencies that are aliased together hash into the same bucket.
- **Recovery:** BigBand recovers the frequency positions and their values by bucketizing the same spectrum multiple times so that

¹The above gives the intuition. However, technically, we compute changes in the signal power, not the actual signal.

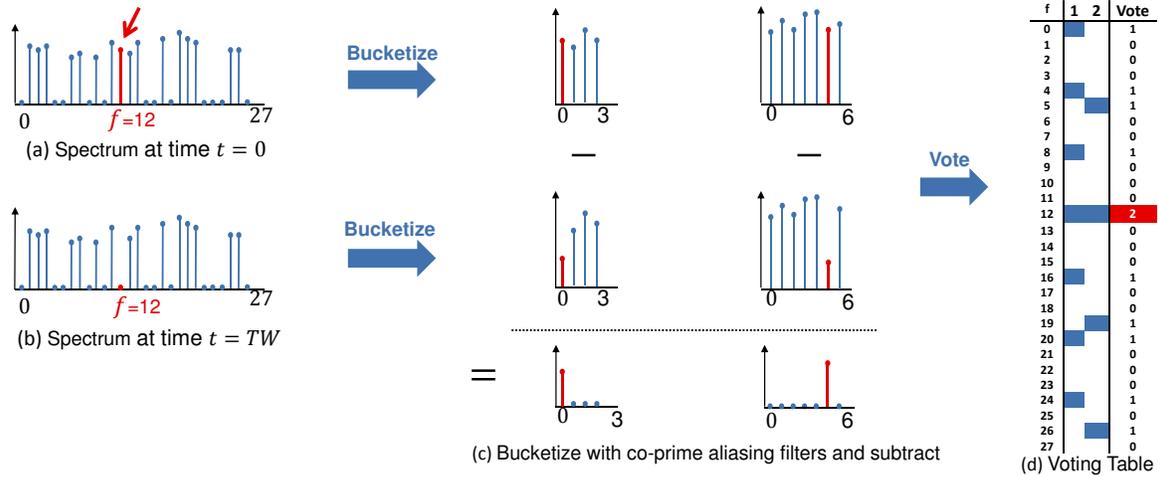


Figure 1—Sensing one change in non-sparse spectrum: (a) $f=12$ is occupied at $t = 0$. (b) $f=12$ is empty at $t = TW$. (c) Bucketize the spectrum at $t = 0$ and $t = TW$ using co-prime aliasing filters and subtract the two bucketizations to discover changing buckets. Changes are sparse. (d) Each co-prime filter votes for the frequencies that hash to a changing bucket. Only $f=12$ gets two votes.

it randomizes the bucketization process. Now the frequencies that hash into the same bucket in one bucketization do not collide again in other bucketizations so that we can decouple them. [7] proves that if the buckets are co-prime (i.e., the subsampling ratios of the ADCs are co-prime), we can recover the spectrum accurately by voting techniques.

3. D-BIGBAND

The key idea of D-BigBand is that although the spectrum might be not sparse, the changes in spectrum usage are sparse: i.e. over short intervals, only few frequencies are freed up or become occupied. We refer to this as *differential sparsity*. Differential sparsity allows D-BigBand to leverage the same sparsity recovery principle as in BigBand, but operates on the differential sparsity instead of the sparsity in the spectrum utilization. To see how this allows D-BigBand to sense a non-sparse spectrum we will start with an example.

3.1 Illustrative Example

In this example, we are going to assume that the state of any frequency can either be occupied or empty. However, if a frequency is occupied, its value does not change over time. We will later explain how to deal with the fact that values of occupied frequencies change over time. Let us consider the case where one frequency $f = 12$ which was occupied becomes unoccupied after time TW as shown in Fig. 1(a,b). Now if we bucketize the spectrum, all buckets will be non-empty and will have collisions. Hence, we cannot directly use the BigBand algorithm. However, since frequency $f = 12$ became empty after time TW , the power in the bucket it hashes to will become lower after time TW . Further, since it is the only frequency that changed state, only the power of that bucket changes. Hence if we subtract the bucketization at time TW from that at time 0, we can find which buckets have frequencies that changed state as shown in Fig. 1(c).

Subtracting the bucketizations, allowed us to bucketize the “changes” in the spectrum. However, we still need to estimate which frequency is the one that changed state out of the frequencies that hash to the bucket. To do this, we use an estimation procedure based on voting and co-prime aliasing filters. Both at time 0 and time TW , we perform two bucketizations; one using an aliasing filter with four buckets and another using an aliasing filter with seven buckets as shown in Fig. 1(c). Now every frequency that is hashed

into a bucket that changed gets a vote. However, since the filters are co-prime, frequencies that hash to the same bucket as f in the first filter and get a vote, will hash to a different bucket in the second filter and will not get a second vote. Hence, only frequency $f = 12$ will get two votes which allows us to estimate its position as shown in Fig. 1(d).

The above example gives an intuition of how we can leverage the sparsity of changes in the spectrum to discover which frequencies become occupied and which become empty. However, to be able to generalize the above approach, we need to first address the following issue: Since the values of the occupied frequencies change after a time TW , the values of the buckets will change even if the state of the frequencies that hash to them did not change. Hence, we cannot simply subtract the two bucketizations. However, since FCC typically requires wireless transmissions to be whitened over time, the average power of a bucket will not change if the state of frequencies that hash to it does not change. To estimate the average power over a time window TW , D-BigBand performs the bucketization multiple times and averages the power of the buckets. The longer the time window TW , the better the estimate of the average power of each bucket. However, the longer the time window, the more frequencies change their state. In §4, we show that a time window $TW = 1$ ms allows us to properly detect changes in the buckets.

3.2 D-BigBand Algorithm

D-BigBand’s algorithm works as follows. Over a time window TW , D-BigBand bucketizes the signal multiple times² for each of the co-prime aliasing filters and calculates the average power in the bucket over this time window. It then repeats these bucketizations over the next time window and subtracts the average power of the buckets in the first time window from that in the second time window. After that each filter votes for frequencies that hash to buckets where the power changed. Frequencies that get full votes are picked as the frequencies whose state has changed. Hence, based on our knowledge of the spectrum occupancy during the first time window, we can discover the spectrum occupancy during the second time window.

As with any differential system, we need to initialize the state of spectrum occupancy. However, an interesting property of D-BigBand is that we can initialize the occupancy of each frequency

²The number of times D-BigBand can average is $= TW/T$ where T is the FFT window time.

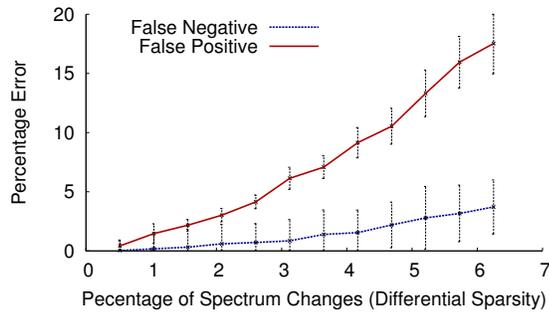


Figure 2—D-BigBand’s effectiveness as a function of Differential Sparsity: For less than 2.5% changes in the spectrum of 1GHz, D-BigBand works reliably well.

in the spectrum to unknown. This is because, when we take the difference in power we can tell whether the frequency became occupied or it became empty. Hence, once the occupancy of a frequency changes, we can tell its current state irrespective of its previous state. This avoids the need for initialization and prevents error propagation.

4. IMPLEMENTATION AND EVALUATION

4.1 Implementing D-BigBand

As a proof of concept, we implement D-BigBand using USRP N210 software radios [2]. Since the USRPs use the same ADCs, it is not possible to have co-prime sub-sampling rates.³ Therefore, to verify that D-BigBand can sense a non-sparse spectrum, we use multiple USRPs sampling adjacent narrowband chunks to capture a full 1 GHz of spectrum. And then, we emulate lower-speed co-prime ADCs by subsampling the 1GHz signal.

However, since our testbed has only 20 USRPs, we divide them into 10 receivers and 10 transmitters and capture 250 MHz at a time. We repeat this 4 times at center frequencies that are 250 MHz apart and stitch them together in the frequency domain to capture the full 1 GHz spectrum. We then perform the inverse FFT to obtain a time signal sampled at 1 GHz. We now subsample this time domain signal using co-prime aliasing filters with the following sampling rates: 1/21, 1/20, 1/23 GHz, and run D-BigBand on these subsampled versions of the signal.

4.2 Evaluation

In this section, we evaluate D-BigBand’s ability to properly sense a non-sparse spectrum.

Experiment 1: Differential Sparsity Range

We start by evaluating how well D-BigBand can sense the spectrum versus how many frequencies in the spectrum are changing their occupancy. We set the sparsity of the spectrum to 50% occupancy on average and we vary the percentage frequencies that change their occupancy over 1 ms between 0.5% up to 6.25% of the total 1 GHz spectrum.

Results 1: We consider two metrics: 1) *False Negatives*: The fraction of occupied frequencies that BigBand incorrectly reports as empty.

³The USRP ADC has a sampling rate of 100 MHz, and the USRP has digital filters but these can only produce sampling rates which are integer dividers of 100 MS/s (i.e. 100/2, 100/3, 100/4, etc.). This means, for 1 GHz bandwidth, it is not possible with USRPs to get even two aliasing filters that sample at $1/p_1$ and $1/p_2$ where p_1 and p_2 are co-prime. Instead, we can of course implement the co-prime aliasing filters using commodity ADCs. However, this would require building a new receiver that uses these ADCs.

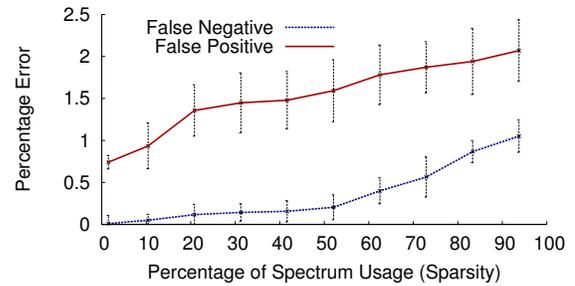


Figure 3—D-BigBand’s effectiveness as a function of Spectrum Sparsity: Over a band of 1 GHz, D-BigBand can reliably detect changes in spectrum occupancy even when the spectrum is 95% occupied, as long as the change in spectrum occupancy is less than 1% every ms.

2) *False Positives*: The fraction of empty frequencies. Fig. 2 shows the number of false positives and false negatives as a function of the percentage of frequencies that change occupancy. The figure shows for less than 2.5% changes in the spectrum (i.e. 25 MHz), false positives are below 5% and false negatives are below 1% and for less than 5% changes in the spectrum (i.e. 60 MHz), false negatives are below 3%. Although, false positives become larger as the percentage of frequencies changing their occupancy every 1 ms becomes larger, the goal in spectrum sharing is typically to find an unoccupied frequency to transmit in so as long as the number of false negative is small, D-BigBand will always allow devices to correctly find some unoccupied frequencies.

Experiment 2: Sparsity Range

Now we want to confirm that D-BigBand can properly sense the spectrum at any sparsity. For this, we fix the number of frequencies that change occupancy every 1 ms to 1% (i.e. 10 MHz) and vary the percentage of total occupied frequencies in the spectrum between 1% to 95%.

Results 2: Fig. 3 shows the percentage of of false positives and false negatives. As the occupancy increases from 1% to 95%, both false negatives and false positives increase by 1% this is because as the spectrum becomes more occupied, more frequencies hash to the same bucket and the variance in estimating the average power per bucket increases.

5. REFERENCES

- [1] D. Chen, S. Yin, Q. Zhang, M. Liu, and S. Li. Mining spectrum usage data: a large-scale spectrum measurement study. In *Mobicom*, 2009.
- [2] Ettus. Inc. USRP. <http://ettus.com>.
- [3] FCC. Mobile broadband: The benefits of additional spectrum.
- [4] FCC, Second Memorandum Opinion & Order 10-174.
- [5] B. Ghazi, H. Hassanieh, P. Indyk, D. Katabi, E. Price, and L. Shi. Sample-optimal average-case sparse fourier transform in two dimensions. In *Allerton*, 2013.
- [6] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Nearly optimal sparse fourier transform. In *STOC*, 2012.
- [7] H. Hassanieh, L. Shi, O. Abari, E. Hamed, and D. Katabi. Ghz-wide sensing and decoding using the sparse fourier transform. In *INFOCOM’14*, 2014.
- [8] M. A. McHenry. NSF spectrum occupancy measurement project summary, 2005.
- [9] Microsoft Spectrum Observatory. <http://spectrum-observatory.cloudapp.net>.
- [10] B. Murmann. A/d converter trends: Power dissipation, scaling and digitally assisted architectures. In *IEEE Custom Integrated Circuits Conf. (CICC)*, 2008.
- [11] PCAST: Realizing the full potential of government held spectrum to spur economic growth, 2012.