Efficient spectrum summarization using compressed spectrum scans

Mariya Zheleva[1], Timothy Larock[1], Paul Schmitt[2], and Petko Bogdanov[1]

[1]Department of Computer Science, University at Albany SUNY, {*mzheleva,tlarock,pbogdanov*}@*albany.edu*
[2]Department of Computer Science, Princeton University, *pschmitt@cs.princeton.edu*

***Abstract*—We present AirPress, a spectrum scan compression method that leverages wavelet decomposition for lossy compression of spectrum data and allows up to 64:1 compression ratio of power spectral density traces without adversely impacting the spectrum summarization accuracy. We demonstrate the utility of AirPress on real-world spectrum measurements and show that it enables high-accuracy spectrum summarization of real-world transmitters while reducing the corresponding trace by 94%.**

I. Introduction

Next-generation networks with opportunistic spectrum access require longitudinal spectrum traces with fine granularity in frequency, time and space. This is poised to produce prohibitive amounts of spectrum scans that will pose large bandwidth, storage and processing overhead on spectrum measurement. For example, a one second scan of a 600MHz spectrum band with a USRP sampling at 20Msps amounts to 23GB.

To enable spectrum inventory at scale we propose AirPress, a method that compresses raw spectrum traces and thus enables large-scale spectrum scan collection, storage and processing. AirPress makes use of wavelet decomposition for lossy compression and, depending on the signal complexity, achieves up to 64:1 compression rate while maintaining small error rates. We demonstrate the utility of AirPress by analyzing controlled Wi-Fi and Bluetooth transmissions and a real-world wideband spectrum scan from 400MHz to 1GHz. Our analysis shows that different bands tolerate different compression levels, which creates an opportunity for adaptive compression towards a scalable spectrum inventory.

This paper makes several contributions: (i) we design AirPress, a spectrum compression technique that reduces the volume of spectrum scans by up to 94%, while preserving signal properties; (ii) we demonstrate that AirPress retains accurate spectrum summarization; and (iii) we harness AirPress to map spectrum compressibility across wideband spectrum measurements and show that compressability depends on the spectrum dynamics.

II. AirPress

The key question tackled by AirPress is *Can raw spectrum measurements be compressed in a manner that preserves the underlying spectrum characteristics, while reducing the overhead for storage and analysis?* We propose to address the storage challenge using wavelet-based compression. In AirPress, we adopt a one dimensional wavelet decomposition applied to a scan of signal power $p^t(f)$ over a range of frequencies at a given time instant t, where p is a function over n discrete frequency values. The *Haar* wavelet transform is a good choice for decomposing impulse-like signals such as these in spectrum scans, and hence we use this basis.

The wavelet decomposition of a power scan $p^t(f)$ is a mapping from the n-dimensional original signal $p^t(f)$ to a set of coefficients w^t that correspond to summaries of the signal at different resolutions. The full decomposition w^t has n coefficients and can be used to reconstruct the original signal $p^t(f)$ exactly. Due to the local regularities of the signal many of the coefficients are close to zero and a lossy reconstruction can be obtained by maintaining only a synopsis of the decomposition $\tilde{w}^t$ containing a subset of k coefficients, i.e. $|\tilde{w}^t| = k$. A similar approach has been adopted for approximate query answering in databases [1]. One can also show that for a budget of k coefficients to compute a synopsis, keeping the coefficients of largest absolute value is optimal when minimizing the sum of squared error between the reconstructed and original signal [2].

Our compression approach applies a wavelet transform of the original scan $p^t(f)$ and computes a synopsis $\tilde{w}^t$ of pre-specified size k that can be used to reconstruct an approximation of the original scan $\tilde{p}^t$, and answer various queries regarding occupancy. The savings in storage in our scheme are k/n-fold, i.e. the compression rate is k/n. Of note is that while more aggressive compression leads to drastic reduction in spectrum scan size and smoother signals, it can also eliminate some of the original signal properties. Thus, a tradeoff exists between compression level and the truthfulness of the reconstructed signal. We explore this tradeoff in § III.

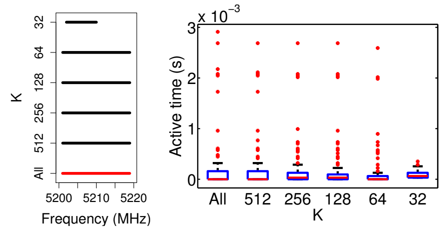

Fig. 1: (left) Bandwidth detection. The bandwidth of the Wi-Fi transmitter is successfully detected as we decrease the number of coefficients to 64. Past 64 coefficient the signal reconstruction deviates significantly from the original signal, which results in false bandwidth detection. (right) Active time detection. The average active time and the active time distribution is successfully detected as we decrease the coefficients to 64. False reconstruction of the signal with 32 coefficients result in inaccurate active time detection.

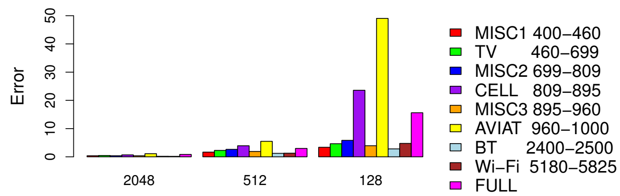

Fig. 2: Compression of the full spectrum (FULL) and individual functional bands using the same coefficient budget. The error of full-band compression is smaller than the cumulative error of split compression.

III. Experimental Evaluation

Effects of spectrum compression on summarization. In this section we demonstrate AirPress's ability to preserve signal quality and enable detailed transmitter identification while performing a 16-fold spectrum scan compression. For this experiment we use a controlled Wi-Fi spectrum scan. In order to summarize the spectrum use we leverage our previous work on spectrum summarization, TxMiner [3].

We focus on detection of transmitter bandwidth and active time. The Wi-Fi transmission in question takes place in channel 40. The transmission utilizes multiple consecutive time chunks of variable duration. As Figure 1(left) shows, we are able to successfully detect the bandwidth as the number of coefficients decreases to 64. At 32 coefficients the signal reconstruction begins to differ drastically in comparison with the original, which leads to false bandwidth detection. Figure 1(right) presents results for active time detection with decreasing coefficients. Each boxplot presents the distribution of active times detected at the corresponding compression rate. As we can see, the mean of the detected active times remains unchanged as the coefficients decrease to 64. We see some of the outliers disappear and be replaced with several smaller outliers. This means that as we denoise the data with increasing compression some of the individual active intervals become more pronounced and are detected as separate intervals. As the number of coefficients reaches 32 the active time distribution changes dramatically, indicating a false detection of transmitter activity.

Wide-band spectrum compression. In order to demonstrate the utility of AirPress in creating a large-scale spectrum inventory we evaluate the compressibility of spectrum bands across a wide frequency range. We collect a spectrum scan from 400MHz to 1GHz with a step of 20MHz and a duration of 1 second. We add to this scan our traces from the controlled Wi-Fi and Bluetooth transmissions. We split the spectrum scan in several sub-bands according to the FCC's spectrum allocation charts as shown in Figure 2. We note that the miscellaneous (MISC) bands host a variety of technologies including air navigation, maritime, amateur radio, Earth exploration satellites, public safety and narrow-band PCS.

For this experiment we select a decreasing number of coefficients for each run. During each run we first compress the entire frequency band with the corresponding number of coefficients. We call this full-scan compression. We then redistribute the total number of coefficients to the individual sub-bands and compress these sub-bands separately. We call this split compression. We redistribute the coefficients based on the sub-band size in frequency; that is larger bands will get a proportionally larger fraction of the coefficients.

Figure 2 presents our results. As expected, for all sub-bands the error increases as we decrease the number of coefficients. Richer sub-bands such as Aviation and Cellular are less compressible when the objective is error minimization. Lastly, we note that the error from full-scan compression is smaller than the cumulative error from split compression. The reason for this is that there are idle fractions in each sub-band, which have to be assigned the same coefficients multiple times in the split compression. At the same time in the full-scan compression these idle fractions get assigned the same coefficients only once, which results in better compressibility and more efficient distribution of coefficients to dynamic bands. Our future analysis will study this hypothesis and design spectrum compression principles that regroup spectrum data based on expected activity as opposed to regulatory allocation.

References

[1] K. Chakrabarti, M. Garofalakis, R. Rastogi, and K. Shim, "Approximate query processing using wavelets," *The VLDB JournalThe International Journal on Very Large Data Bases*, vol. 10, no. 2-3, pp. 199–223, 2001.

[2] E. J. Stollnitz, T. D. DeRose, and D. H. Salesin, *Wavelets for computer graphics: theory and applications*. M. Kaufmann, '96.

[3] M. Zheleva, A. Chowdhery, R. Chandra, A. Kapoor, and P. Garnett, "TxMiner: Identifying Transmitters in Real-World Spectrum Measurements," in *IEEE DySPAN'15*, (Stockholm, Sweden), 29 September - 2 October 2015.