MIT Open Access Articles

A cochlear heterodyning architecture for an RF fovea

Massachusetts Institute of Technology

A Cochlear Heterodyning Architecture for an RF Fovea

Soumyajit Mandal and Rahul Sarpeshkar

Massachusetts Institute of Technology, Cambridge, MA 02139

Email: rahuls@mit.edu

Abstract— We describe a novel wireless receiver architecture that is a broadband generalization of narrowband heterodyning systems commonly used in radio. It can be constructed with cochlea-like traveling-wave structures. We show how this architecture exploits the efficiency of cochlear spectrum analysis to perform parallel, multi-scale analysis of wideband signals. We discuss analogies between spectrum analysis in our architecture and the process of successive-subranging analog-to-digital conversion. When combined with our prior work on an RF cochlea, such architectures may be useful in cognitive radios for creating "RF foveas" that select the narrowband components present within wideband, but spectrally sparse signals.

I. Introduction

The commercially interesting range of frequencies for wireless applications extends over approximately three orders of magnitude from 100MHz to 100GHz. Nevertheless, this natural resource is becoming increasingly crowded. At the same time, customers want higher quality of service, particularly increased reliability and higher data rates. In order to satisfy these conflicting goals, radios must utilize the RF spectrum in a more efficient manner. The *cognitive radio* paradigm has been proposed as a solution to this important problem [1]–[3].

Cognitive radios add intelligence to traditional wireless communication systems. They are designed to sense their local RF environment, learn from it, and dynamically adapt their operating parameters, such as transmit power, carrier frequency and modulation strategy, appropriately. In particular, they actively search for spectrum holes, i.e., locations in the time-frequency plane that are suitable for communication, and operate there. As a result, they must be built upon dynamically reconfigurable hardware platforms, such as software-defined receivers and transmitters [4], [5]. Unused VHF and UHF TV channels (54-806 MHz in the United States) currently provide most of the spectrum holes that cognitive radios seek to exploit [6]. However, in the future a wider range of frequencies will need to be monitored for potentially usable holes.

We have recently demonstrated a single-chip RF spectrum analyzer based upon a model of the biological cochlea, or inner ear, that we termed an *RF cochlea* [7]. The RF cochlea performs spectrum analysis over a broad range of RF frequencies (600 MHz to 8 GHz in the implementation in [7]) with low-power, wide dynamic range, low latency, and efficient hardware usage as discussed in [7]. In this paper we describe a novel architecture based upon the RF cochlea that can benefit software-defined receivers by providing real-time estimates of interesting portions of the local RF spectrum with scalably high precision and low hardware and power costs. The architecture is based upon a novel analogy between analog-to-digital conversion and frequency estimation.

II. ADCs and Frequency Estimators

We define *frequency estimators* as systems that estimate the instantaneous frequencies present in a signal. They separate an incoming signal that is spread over a bandwidth B into N outputs, or frequency bins. *Spectrum analyzers* are frequency estimators that can, in addition to estimating the frequencies present, also determine their amplitude, i.e., the power spectral density, and/or phase. Spectrum analyzers form a subset of frequency estimators, and may also be viewed as a parallel, broadband generalization of narrowband radios that estimate amplitude or phase of a narrowband signal centered around a single carrier frequency [8].

We have found interesting analogies between frequency estimators and analog-to-digital converters (ADCs). The problems that frequency estimators and ADCs solve are similar, but in different signal domains. ADCs find the value of an unknown analog amplitude (voltage or current), while frequency estimators do the same with an analog frequency variable. As a result, analogous algorithms work in the two cases. However, a signal can have many instantaneous frequencies of interest, but only one instantaneous amplitude. A more accurate analogy is therefore to view frequency estimators as the frequency-domain analogs of multiple parallel ADCs. Each ADC corresponds to estimating the frequency value (and possibly the power content) of one output bin.

The basic operation of subtraction in the amplitude domain is extensively used by ADCs and has a natural frequency-domain analog through multiplication (mixing or heterodyning). A low-pass or band-pass filter at the output of the mixer can be used to select the difference frequency component, while rejecting the unwanted sum frequency component, as in direct conversion receivers or traditional super-heterodyne receivers, respectively. In this paper, we shall assume that all signals are real. However, the architecture that we describe can be extended to the quadrature signal domain as well.

A simplified view of the n-th stage of a generic successive-subranging ADC is shown in Fig. 1(a). The input V_n is limited to the range $\pm V_{FS}/2^n$, where V_{FS} is the initial, or full-scale range. The output of the comparator is $b_n \in [-1, 1]$, the next output bit. It is converted into an analog signal by a one-bit digital-analog converter (DAC) whose output voltage is

$b_n V_{FS}/2^{n+1}$. The DAC output is subtracted from the input V_n to extract a residue V_{n+1} that spans half the range of V_n and is the input to the next conversion stage. In general the one-bit quantizer and DAC can be replaced by multiple-bit versions, and the stages pipelined or reused during the conversion cycle.

We can begin to map the ADC of Fig. 1(a) to a frequency estimator in Fig. 1(b) by replacing the full-scale voltage range with an initial frequency range B and by replacing the voltage comparator with a frequency comparator. However, frequency estimation requires estimation of bandpass input signals, analogous to an ADC that requires signals within a certain positive voltage range to be digitized. Hence, at each recursive conversion stage in a frequency estimator, the offset or 'common-mode' frequency used to compare the input frequency against and the bandwidth of the conversion decrease by a factor of 2. Thus, in Fig. 1(b), the DAC of Fig. 1(a) is replaced with a digitally-controlled oscillator (DCO) whose output frequency is $\omega_{DCO} = \left(\omega_{REF} + b_n B \right)/2^{n+1}$ and the subtractor of Fig. 1(a) is replaced by a mixer and band-pass filter as shown.
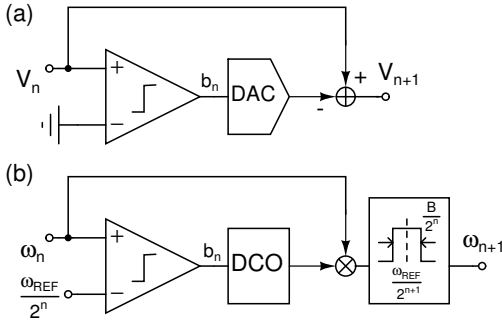

Fig. 1. A single successive-subranging stage for (a) an analog-to-digital converter and (b) a frequency estimator.

The frequency comparator can be built by converting the input and reference frequencies into voltages, and then using a voltage comparator. Frequency-to-voltage converters can be created using tuned circuits (discriminators) or phase-locked loops. The converters should be matched; otherwise the comparator will have a static offset. In addition, the functional relationship between input frequency and output voltage should be monotonic (it need not be linear).

III. THE RF FOVEA

The RF spectrum is sparse over large bandwidths, i.e., dominated by a limited number of relatively narrowband signals which cumulatively occupy only a small fraction of the total bandwidth. Some of these signals are of interest to the user, while others, referred to as interferers, are not. In practical situations interferers can be much larger than desirable signals. A system capable of simultaneously observing several narrow, arbitrarily placed frequency bands can therefore capture essentially all significant information present within the entire RF spectrum. An architecture that uses successive subranging in the frequency domain to solve this problem is shown in

Fig. 2. This *cascaded super-heterodyne architecture* decides where observed regions of the spectrum should be located by using filter banks to decompose spectra into multiple channels.
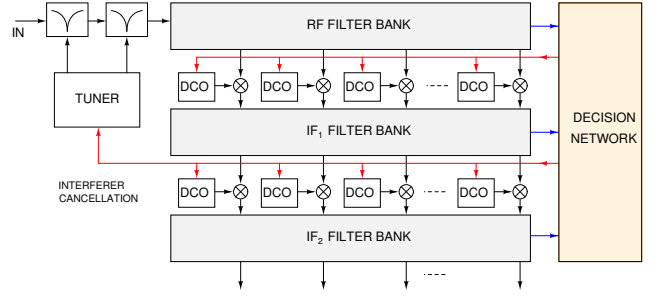

Fig. 2. A successive-subranging architecture for hierarchical analysis of spectrally-sparse broadband signals. Interfering tones can be rapidly estimated to a scalably high degree of precision and then canceled out using tunable notch filters.

The filter banks in Fig. 2 have band-pass transfer functions with center frequencies and bandwidths that scale arbitrarily with horizontal position i. The outputs at each position are analyzed by a cascade of successive subranging stages, each of which is similar to that shown in Fig. 1(b). We may identify $\omega_{REF}/2^n$ and $B/2^n$ with the local center frequency and bandwidth, respectively. The latter can remain constant with i, or scale with center frequency to generate constant-Q characteristics. Here $Q = \omega_{REF}/B$ is the quality factor. In addition, for a fixed value of i the center frequencies and bandwidths change with vertical position j, usually decreasing by a constant factor α at each row. The bandwidth of the i-th filter in the j-th bank, which we denote by $B_{j,i}$, is given by

$$B_{j,i} = \frac{B_{j-1,i}}{\alpha} = \frac{B_{0,i}}{\alpha^m}. \tag{1}$$

A common value for α is 2, in which case the decision network consists of one-bit quantizers, i.e., frequency comparators that determine which half of the input bandwidth $B_{j-1,i}$ should be included in the output bandwidth $B_{j,i}$. In general we must use quantizers with $\log_2(\alpha)$ bits. Thus, the output of each successive filter bank allows the user to determine the frequencies of N signals of interest in parallel with exponentially-increasing precision, where N is the number of filters in each filter bank. The outputs of the decision network can be viewed as the results of N parallel frequency-to-digital conversions (analogous to analog-to-digital conversions). The architecture is also efficient in terms of hardware requirements because it only needs M filter banks and $M \times N$ mixers to increase precision by a factor of α^M. Finally, it retains amplitude information about each output signal and thus acts as a spectrum analyzer as well as a frequency estimator.

The ultimate precision of successive subranging architectures for frequency estimation is limited by three factors: quantization noise, random errors (caused by device mismatch, thermal noise, and flicker noise), and interference from signals in adjacent channels. The first two factors also affect the performance of ADCs, while the third is unique to frequency

estimation. Such interference is due to overlap between the transfer functions of the initial filter bank, and can be reduced by using higher-order filters with larger roll-off slopes and stopband rejections. The main function of subsequent filter banks is to further reduce the relative amount of interference through additional filtering. They can be omitted, or their order reduced, if interference levels are low enough.

One application of the cascaded super-heterodyne structure is to accurately estimate the frequencies and bandwidths of unwanted interferers. The interferers can then be canceled out, allowing weak signals to be detected in their presence. A simple interference-cancellation strategy that uses tunable notch filters is shown in Fig. 2. This scheme was simulated using MATLAB; the results are shown in Fig. 3. Two tones one octave apart are fed into the structure, with one being 80dB larger in amplitude than the other. The curve labeled "original" shows the results without interferer cancellation: The smaller tone is invisible. The lower curve shows filter bank outputs after the large interfering tone has been estimated and notched out: The small tone is now clearly resolved.
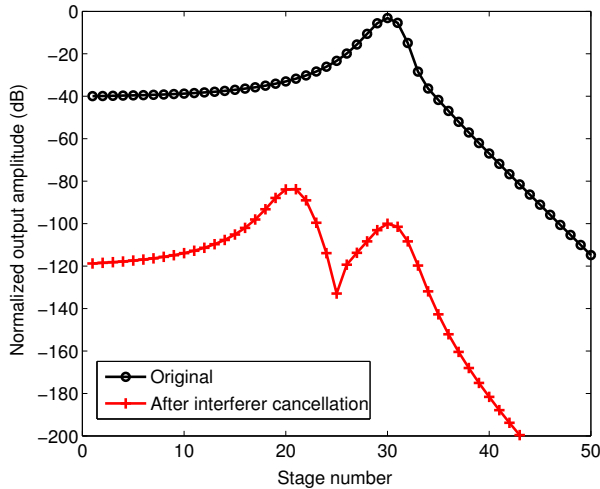

Fig. 3. Simulated performance of the cascaded super-heterodyne architecture. The plots show outputs of the final filter bank before and after a large interfering signal was estimated and cancelled out.

Cochlea-like traveling-wave structures consist of transmission lines or filter cascades with maximal-response frequencies that decrease exponentially with position. They are an efficient way to synthesize high-order, constant-Q filter banks because of extensive hardware reuse: Each output is filtered by multiple stages, each of which simultaneously produces its own output. Unlike other spectrum analysis algorithms, such as the FFT or simple filter banks, each cochlear stage is used multiple times to synthesize output transfer functions with large roll-off slopes, and consequently, high frequency selectivities. As a result, the hardware and power costs of the cochlea are much lower than comparable algorithms [7].

An added advantage of constant-Q frequency decomposition is that it is better matched to the real RF environment.

The RF spectrum today is divided into licensed and unlicensed bands that roughly follow constant-Q characteristics, i.e., have bandwidths that are approximately proportional to center frequencies. For example, the Q of the unlicensed industrial-scientific-military (ISM) frequency bands allocated internationally only varies by about a factor of ten as the center frequencies vary by three orders of magnitude. It is thus advantageous to replace the filter banks in Fig. 2 with cochlea-like traveling-wave structures. We refer to this special case, which is shown in Fig. 4, as the *RF fovea*, because it uses an RF cochlea to rapidly obtain an overview of the entire spectrum, and then adaptively "locks in" or "foveates" on the largest components in that spectrum with exponentially-increasing precision.

As in Fig. 2, center frequencies and bandwidths in Fig. 4 decrease exponentially with vertical position j. The center frequency of the j-th output from the i-th cochlea is proportional to its bandwidth (constant-Q behavior) and is given by

$$f_{i,j} = f_{0,0}2^{-i/N_{oct}}\alpha^{-j} = f_{0,0}e^{-[i\ln(2)/N_{nat}+j\ln(\alpha)]}, \quad (2)$$
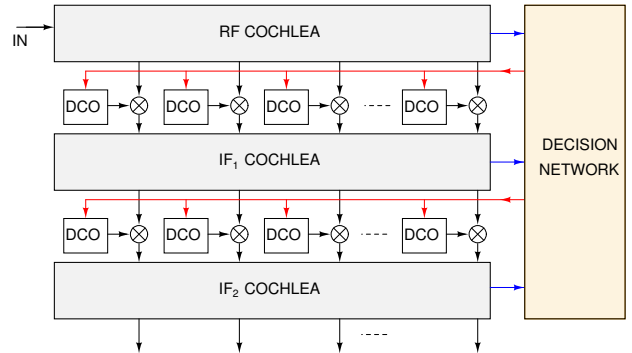

Fig. 4. The RF fovea, an example of the cascaded super-heterodyne architecture that uses cochlear models to obtain frequency selectivity.

where each cochlea is assumed to produce N_{oct} outputs per octave, and α is the constant that represents the exponentially increasing precision as we move from RF to lower-frequency IF stages in Fig. 4. The outputs of each cochlea are downconverted using mixers. Each downconversion step selects $1/\alpha$ of the output bandwidth of the previous cochlear output. For example, if $\alpha = 2$, one-bit decisions determine which half gets selected by changing the DCO frequency f_{ji} fed into the mixer between $f_{ji,0}$ and $f_{ji,1}$, where $1 \le j \le M$ and $1 \le i \le N$ are the cochlear and output indices, respectively. In order to reduce hardware costs we can downconvert only those outputs that appear significant, such as local maxima (peaks). The output of the mixer at location i can be directly fed into the i-th stage of the next cochlea, as shown in Fig. 4. Alternatively, all outputs of the j-th cochlea can be added and fed into the first stage of the $j + 1$-th cochlea.

The total acquisition time of the RF fovea is determined by the sum of the settling times of all the cochleas. It is

dominated by the final cochlea since it analyzes the lowest input frequencies, and is given by

$$T_{acq} \approx \frac{\gamma N_{oct}}{f_{M,N}} \left(\frac{\alpha}{\alpha - 1} \right), \tag{3}$$

where γ is a constant, of order unity, that depends on the cochlear model, $f_{M,N} = f_{0,N}/\alpha^M$ is the lowest frequency analyzed by the final cochlea, there are a total of M cochleas and $f_{0,N}$ is the lowest frequency analyzed by the first (input) cochlea. The hardware complexity of the system is modest: It scales as $O(NM)$, where N is given by

$$N = N_{oct} \left(\log_2(\beta) + 1 \right). \tag{4}$$

This formula assumes (conservatively) that the first N_{oct} stages of each cochlea are necessary for building up the transfer functions and do not produce useful outputs. Also, $\beta = f_{j,N_{oct}}/f_{j,N}$ is defined as the ratio of maximum and minimum frequencies analyzed by each cochlea.

Fig. 5 shows the simulated frequency estimation error of the RF fovea for various values of M. We generated this plot by feeding artificial spectra into the system. Each spectrum consisted of N_f tones with random frequencies and amplitudes, uniformly distributed in log-frequency and log-amplitude space, respectively. We used the unidirectional cochlear model described in previous work [7], and set $N_{oct} = 10$ for the simulations. If simple energy extraction, i.e., no pattern-recognition or phase-based spectral computations, is performed on the cochlear outputs, a single cochlear model can localize the frequency corresponding to a spatial-response peak to a worst-case fractional bandwidth of $\Delta f/f \approx \ln(2)/N_{oct}$ when the output signal-to-noise ratio is 1. If the unknown frequency is uniformly distributed in this interval, the resulting RMS estimation error with this conservative assumption is $(\Delta f/f)/\sqrt{12} = 2.0\%$ for $N_{oct} = 10$, and corresponds to $M = 1$ in Fig. 5. Our simulations confirm that the frequency-estimation error initially decreases with M as expected.

However, the error asymptotes to a fixed value as M increases. This asymptotic error increases monotonically with N_f, indicating that the precision of our algorithm is ultimately limited by interference from other parts of the spectrum if they have not been notched out as in Fig. 2. Besides notching, we can also increase N_{nat}, which is proportional to the high-frequency rolloff slopes of the cochlear transfer functions, to reduce such interference. The result is improved precision at the cost of increased hardware and power consumption.

IV. CONCLUSION

We have described an *RF fovea*, a successive-subranging receiver architecture that allows users to focus on selected portions of a broadband input spectrum. The RF fovea architecture is a generalization of heterodyning concepts from the narrow-band domain to the broadband domain. It can be efficiently implemented using traveling-wave cochlear models, creating a *cascaded cochleas* or *cochlear heterodyning* architecture, and may be useful as a dynamic spectrum-sensing front end for
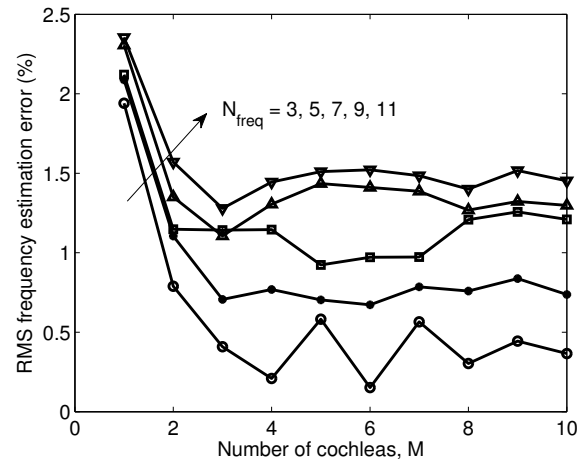

Fig. 5. Frequency estimation error of the RF fovea to inputs consisting of N_f tones with random frequencies and amplitudes as a function of M, the total number of cochleas, obtained from MATLAB simulations.

software-defined and cognitive radio receivers. For example, the dynamic range of such receivers can be increased by detecting strong unwanted signals (interferers) and adaptively removing them, resulting in a form of distributed gain control. Our cochlear models themselves use gain control to increase dynamic range, i.e., they amplify strong signals less than weak signals. The gain-control process causes two-tone suppression: Strong tones reduce the gain of weaker signals at nearby frequencies, thus enhancing spectral contrast [9]. As a result our system can detect interferers in noisy environments more reliably than purely linear schemes.

REFERENCES

[1] J. Mitola III and G. Q. Maguire, Jr., "Cognitive radio: Making software radios more personal," *IEEE Personal Communications*, vol. 6, no. 4, pp. 13–18, 1999.

[2] S. Haykin, "Cognitive radio: Brain-empowered wireless communications," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 201–220, 2005.

[3] S. Haykin, J. H. Reed, G. Y. Li, and M. Shafi, Eds., *Special issue on cognitive radio*, ser. Proceedings of the IEEE, vol. 97, no. 4–5, April–May 2009.

[4] J. Mitola III, "Software radios: Survey, critical evaluation and future directions," *IEEE Aerospace and Electronic Systems Magazine*, vol. 8, no. 4, pp. 25–36, 1993.

[5] ——, "The software radio architecture," *IEEE Communications Magazine*, vol. 33, no. 5, pp. 26–38, 1995.

[6] M. Mishra and A. Sahai, "How much white space is there?" EECS Department, University of California, Berkeley, Tech. Rep., Jan 2009. [Online]. Available: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-3.html

[7] S. Mandal, S. Zhak, and R. Sarpeshkar, "A bio-inspired active radio-frequency silicon cochlea," *IEEE Journal of Solid-State Circuits*, vol. 44, no. 6, pp. 1814–1828, 2009.

[8] ——, "Architectures for universal or software radio," U.S. patent application 11/958,990 filed December 18th 2007. Published US2008/0240301 A1, October 2008.

[9] L. Turicchia and R. Sarpeshkar, "A bio-inspired companding strategy for spectral enhancement," *IEEE Transactions on Speech and Audio Processing*, vol. 13, no. 2, pp. 243–253, 2005.