ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) | 978-1-7281-7605-5/20/$31.00 ©2021 IEEE | DOI: 10.1109/ICASSP39728.2021.9413689

PERIODIC SIGNAL DENOISING: AN ANALYSIS-SYNTHESIS FRAMEWORK BASED ON RAMANUJAN FILTER BANKS AND DICTIONARIES

Pranav Kulkarni and P. P. Vaidyanathan

Dept. of Electrical Engineering, MC 136-93
California Institute of Technology, Pasadena, CA 91125, USA
E-mail: pkulkarn@caltech.edu, ppvnath@systems.caltech.edu

ABSTRACT

Ramanujan filter banks (RFB) have in the past been used to identify periodicities in data. These are analysis filter banks with no synthesis counterpart for perfect reconstruction of the original signal, so they have not been useful for denoising periodic signals. This paper proposes to use a hybrid analysis-synthesis framework for denoising discrete-time periodic signals. The synthesis occurs via a pruned dictionary designed based on the output energies of the RFB analysis filters. A unique property of the framework is that the denoised output signal is guaranteed to be periodic unlike any of the other methods. For a large range of input noise levels, the proposed approach achieves a stable and high SNR gain outperforming many traditional denoising techniques.

***Index Terms*—** Periodicity, denoising, Ramanujan filter banks, pruned synthesis dictionary, hidden periods.

1. INTRODUCTION

Denoising a measurement to get a better estimate of the true signal is important in signal and image processing and has been researched widely [1–6]. Some of the traditional transform domain denoising methods include Discrete Fourier Transform (DFT) denoising and wavelet denoising. DFT denoising works well for harmonic signals when the number of available samples is large. The transform domain methods usually incorporate the a-priori knowledge about the underlying signal and involve thresholding in the transform domain followed by inverse transformation. For example, the hard and soft thresholding methods of wavelet denoising [2] for natural images are based on the fact that most of the energy for natural images is concentrated in only a few dominant wavelet coefficients.

Instead of using square matrices, overcomplete dictionaries can be used for signal representation. If the dictionary is chosen appropriately, signal denoising can be framed as a sparse vector recovery problem. A denoising method based on OMP (Orthogonal Matching Pursuit) is one such method [3]. Further, to avoid explicit handcrafting of the dictionaries that lead to sparse coefficients for the desired class of signals, the dictionaries can either be learned beforehand from clean data or directly from the patches of the noisy signal to be denoised. A denoising method based on the well-known K-SVD [7] has demonstrated good results for images [4].

Many real-world signals like ECG and gravitational wave signals can be closely approximated by periodic signals [8,9]. Discrete-time periodic sequences also occur naturally in proteins and DNA [10,11]. Accurately estimating the component periods and denoising such signals is an important task when the received signal is noisy. Although there exist many denoising methods, we will demonstrate in the paper that these methods usually do not respect the periodicities in the signal and produce a non-periodic signal as denoised output. Only a few methods available in the literature are particularly designed for denoising periodic signals. Non-harmonic analysis (NHA) [5] is one such method that poses denoising as a mean squared error minimization problem. Thus developing periodicity-aware denoising methods is essential.

Scope and Outline. In this paper we propose a novel analysis-synthesis framework for denoising discrete-time periodic signals based on Ramanujan analysis filter banks and synthesis dictionaries. The method produces a truly periodic signal as the denoised output. To our knowledge, none of the existing methods have this property. Our method outperforms many traditional denoising techniques in terms of SNR gain achieved. In Sec. 2, we review Ramanujan filter banks and Ramanujan dictionaries. Then in Sec. 3, we describe our new analysis-synthesis framework and different versions of the associated denoising method in detail. We compare our method with some of the well-known denoising algorithms in Sec. 4. Sec. 5 presents conclusions and future directions.

Notations. A discrete-time signal $x(n)$ is said to be periodic with period P if P is the smallest positive integer such that $x(n) = x(n+P)$ for all n. The notation (k, q) represents the gcd of the integers k and q. So $(k, q) = 1$ means that k and q are coprime. $\text{lcm}\{q_1, q_2, \cdots, q_K\}$ denotes the least common multiple of the integers $q_1, q_2, \cdots, q_K$. The notation $q_i|q$ means that q_i is a divisor of q. Finally, $\phi(q)$ is the Euler totient (number of integers k in $1 \le k \le q$ satisfying $(k, q) = 1$), $\Phi(q) = \sum_{k=1}^{q} \phi(k)$, and $W_q = e^{-j2\pi/q}$.

2. REVIEW OF RAMANUJAN FILTER BANKS AND RAMANUJAN DICTIONARIES

Ramanujan sums are particularly useful in identifying component periods of a discrete-time periodic signal [12–14]. For an integer $q > 0$, the q-th Ramanujan sum is defined as [15]:

$$c_q(n) = \sum_{\substack{k=1 \\ (k,q)=1}}^{q} e^{j2\pi kn/q} = \sum_{\substack{k=1 \\ (k,q)=1}}^{q} W_q^{-kn} = \sum_{\substack{k=1 \\ (k,q)=1}}^{q} W_q^{kn} \tag{1}$$

It is well-known that the q-th Ramanujan sum is integer valued for all n and is periodic with period q. A review of Ramanujan sums and their properties can be found in [12]. The q-th Ramanujan sum lies in the q-th Ramanujan subspace defined as [12] $S_q = span\{W_q^{kn}; (k, q) = 1\}$. S_q is a $\phi(q)$-dimensional space containing signals $x(n)$ with period q that have non-zero DFT value only at the frequency indices k that are coprime to q. S_q also admits an integer basis, namely [12] $S_q = span\{c_q(n-l);\ 0 \le l \le \phi(q) - 1\}$.

A signal $x \in \mathbb{R}^N$ which is expected to have periodic components only upto P_{max} can be represented in terms of the

This work was supported in parts by the ONR Grant N00014-18-1-2390, the NSF Grant CCF-1712633, and the California Institute of Technology.

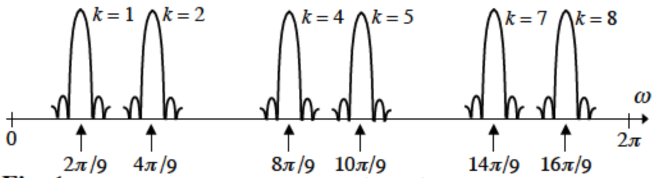

Fig. 1. Frequency response magnitude of the FIR Ramanujan filter $C_9^{(l)}(e^{j\omega})$, when l is finite [20].

dictionary $\boldsymbol{D}$ as [16]

$$\boldsymbol{x} = \boldsymbol{D}\boldsymbol{b}, \text{ where } \boldsymbol{D} = [\boldsymbol{H}_1 \ \ \boldsymbol{H}_2 \ \ \cdots \ \ \boldsymbol{H}_{P_{max}}]. \quad (2)$$

Here, $\boldsymbol{H}_q$ is a $N \times \phi(q)$ matrix whose columns are the q-th Ramanujan sum and its $\phi(q)$ shifted versions spanning the q-th Ramanujan subspace S_q. The columns are extended periodically to obtain N rows. The size of $\boldsymbol{D}$ is thus $N \times \Phi(P_{max})$. The number of samples N is usually three or four times the maximum component period, whereas $\Phi(P_{max})$ grows as $O(P_{max}^2)$ [16]. Thus the matrix $\boldsymbol{D}$ is usually fat. Period identification is posed as a sparse vector recovery problem in this setting, and the component periods can be found by identifying the set of subspaces that correspond to non-zero entries in the coefficient vector $\boldsymbol{b}$. Sub-matrices $\boldsymbol{H}_q$ can in fact be replaced by any other basis of S_q, and such dictionaries are called nested periodic dictionaries [17]. For example, the Farey dictionary [16] uses the coprime frequency index columns from the $q \times q$ DFT matrices to span S_q, and the natural periodic dictionary uses first $\phi(q)$ columns of the $q \times q$ identity matrix to span S_q.

The dictionary method, however, is not as effective when the data to be analyzed is streaming and has a periodicity structure that evolves over time. For this, Ramanujan filter banks (RFB) are more suited [18–20]. To detect component periods up to P_{max}, RFB has FIR filters corresponding to each period from 1 to P_{max}. The q-th Ramanujan filter $C_q^{(l)}(z) = \sum_{n=0}^{lq-1} c_q(n) z^{-n}$ has the first l periods of the q-th Ramanujan sum as its filter coefficients, where l is a fixed repetition number for all the filters in RFB. Thus, the q-th Ramanujan filter has ql possible non-zero filter coefficients.

Fig. 1 qualitatively shows the frequency response of the 9^{th} Ramanujan filter. The q-th Ramanujan filter with a finite l has $\phi(q)$ passbands around the coprime frequency locations. As l increases, the passbands become narrower and the frequency response approaches a set of Dirac-delta functions, like the ideal Ramanujan filter $C_q(e^{j\omega})$. Based on the outputs of the different filters, the component periods can be estimated owing to the following theorem proved in [19]:

Theorem 1. *The lcm property of Ramanujan filter banks*: Let $x(n)$ be a period-P input signal with $1 \leq P \leq P_{max}$. Let nonzero outputs be produced by the subset of ideal Ramanujan filters $C_{q_i}(e^{j\omega})$ corresponding to periods $q_1, q_2, \cdots, q_K$. Then the period P is given by $P = \text{lcm}\{q_1, q_2, \cdots, q_K\}$. $\diamondsuit$

Since the filter outputs are convolutions of $x(n)$ with $c_q(n)$, the outputs are approximately signals in the Ramanujan subspace S_q. Using a sliding window average, the energies of the filter outputs can be plotted as a function of time for a streaming signal to get a time-period plane plot [18, 19] similar to time-frequency plots. RFB has applications in finding periodicities in proteins [10] and DNAs [11].

3. HYBRID ANALYSIS-SYNTHESIS FRAMEWORK

The Ramanujan filter bank is an analysis bank of FIR filters that correspond to different periods. It is natural to look for the synthesis counterpart of RFB. However, note that the RFB is not like a traditional filter bank where the prototype filter is a lowpass filter and the other filters in the filter bank are frequency shifted versions of the prototype filter. The Ramanujan filters have multiple passbands located at coprime frequencies. Decimating the outputs of these filters will inevitably lead to aliasing in a rather complex manner in the frequency domain. In absence of any such decimation, the RFB outputs have a lot of redundant information, and it is not easy to combine outputs of the filters to perfectly reconstruct the original signal. Thus, the RFB is not like traditional analysis filter banks for which there exist synthesis counterparts with perfect reconstruction property [21].

3.1. Using dictionary for synthesis

To overcome this problem, we propose a novel analysis-synthesis framework where analysis is done by the RFB and the synthesis is done via a dictionary generated based on the outputs of the RFB. Consider a noiseless signal $\boldsymbol{x} \in \mathbb{R}^N$ having a non-changing periodicity structure, meaning the periodic components do not change over time. When the signal is passed through an ideal Ramanujan filter bank (which has $C_q(e^{j\omega})$ instead of $C_q^{(l)}(z)$) containing filters up to P_{max}, let the filters with indices $r_1, r_2, \cdots, r_m$ produce non-zero outputs. As long as $\boldsymbol{x}$ does not have any component period larger than P_{max}, from Theorem 1 we can conclude that the signal has non-zero frequency components possibly only at $2\pi k_{ij}/r_i$, where $(k_{ij}, r_i) = 1$, $i = 1, \ldots, m$. This means that the signal can be written as a linear combination of the signals from $S_{r_1}, S_{r_2} \ldots S_{r_m}$. As H_q defined in Eq. (2) spans S_q, we must have

$$\boldsymbol{x} = \widehat{\boldsymbol{D}}\boldsymbol{b}, \text{ where } \widehat{\boldsymbol{D}} = [\boldsymbol{H}_{\boldsymbol{r_1}} \ \ \boldsymbol{H}_{\boldsymbol{r_2}} \ \ \cdots \ \ \boldsymbol{H}_{\boldsymbol{r_m}}] \quad (3)$$

Here, the size of $\widehat{\boldsymbol{D}}$ is $N \times \sum_{i=1}^{m} \phi(r_i)$ and $\boldsymbol{b}$ is a $\sum_{i=1}^{m} \phi(r_i)$ dimensional coefficient vector. This operation of retaining only a few subspaces from $\boldsymbol{D}$ will be called **dictionary pruning**. When $\boldsymbol{x}$ has only a few component periods in the range $1 \leq P \leq P_{max}$, the matrix $\widehat{\boldsymbol{D}}$ has significantly lesser number of columns than $\Phi(P_{max})$, and might even be a tall matrix instead of a dictionary. Thus, if the filters are indeed ideal, we can conclude that the signal $\boldsymbol{x}$ lies in the column space of matrix $\widehat{\boldsymbol{D}}$ from Eq. (3). Note that given the RFB outputs alone, there is no way to recover the input $\boldsymbol{x}$, that is, there is no known synthesis bank for perfect reconstruction. This is also not the goal of this paper. The goal is to get an estimate of the noise-free version $\boldsymbol{x}$ of a signal from the noisy version $\boldsymbol{y}$. In the following subsection, we develop a denoising framework to achieve this by combining the information at RFB output with the noisy input signal $\boldsymbol{y}$ to extract a cleaner version $\widehat{\boldsymbol{x}}$ of the input. This is a unique aspect of our method.

3.2. Denoising framework

Consider a noisy signal $\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{e}$ as an input to the RFB, where $\boldsymbol{x}$ is a noiseless signal having no periodic component with a period larger than P_{max}, and $\boldsymbol{e}$ is a random noise vector. Consider a practical RFB with filters up to P_{max} with finite repetitions l. If the l is sufficiently large, it is reasonable to assume that the sidelobes of the filters are sufficiently suppressed in the passbands of other Ramanujan filters. Under moderate noise conditions, the output energies of the filters corresponding to certain indices $r_1, r_2 \ldots r_m$ are much larger than the output energies of the other filters. Thus, if we appropriately threshold the output energies to attenuate the noise components, we retain only the subspace indices $r_1, r_2 \ldots r_m$, which are the filter index values that would have produced non-zero outputs if the noiseless signal $\boldsymbol{x}$ was passed through the RFB instead of $\boldsymbol{y}$. Once these contributing

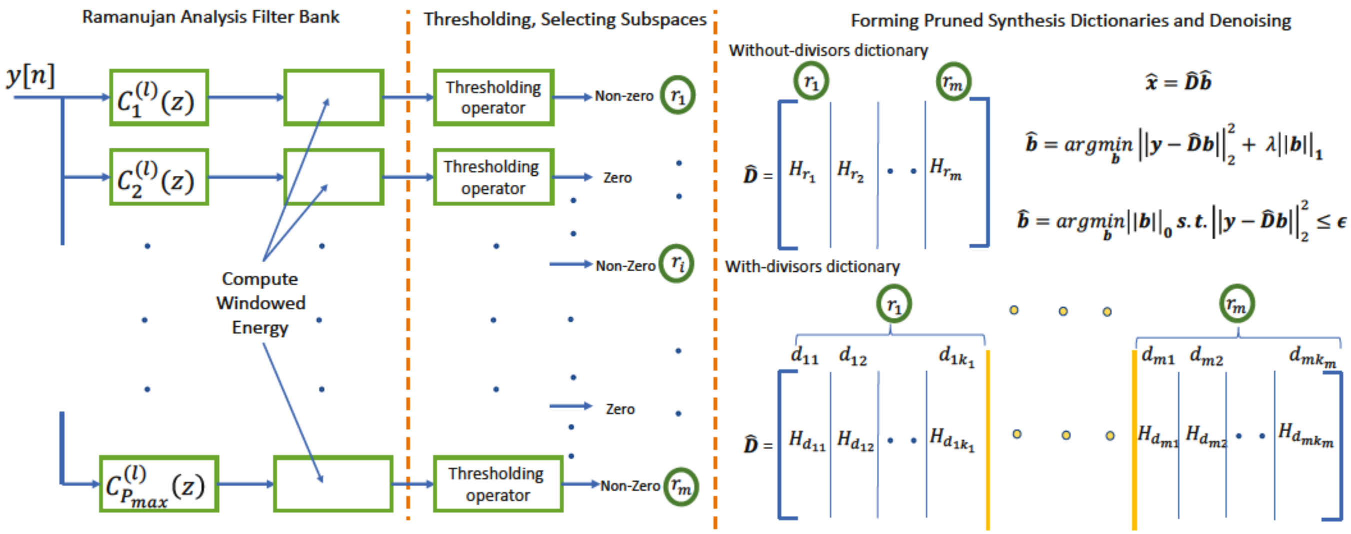

Fig. 2. Denoising framework using Ramanujan analysis filter bank and adaptive synthesis dictionary. See Sec. 3 for details.

subspace indices for x are known, we form a dictionary $\widehat{D}$ like in Eq. (3).

We discussed in the previous subsection that it is not straightforward to compute b just from the RFB output values. However, for the denoising application, the denoised signal must have some resemblance to the input noisy signal y depending on the noise level. Thus the denoising problem can be solved by first solving the optimization problem

$$\widehat{b} = \arg\min_{b} ||y - \widehat{D}b||_2^2 + \lambda||b||_1 \tag{4}$$

to find $\widehat{b}$ and then taking the denoised signal to be $\widehat{x} = \widehat{D}\widehat{b}$. Here λ is a regularization hyperparameter that can either be chosen heuristically or based on the expected noise variance. This l_1 regularized least squares problem can be solved by convex optimization algorithms such as [22]. Since the denoised output $\widehat{x}$ lies in the column space of $\widehat{D}$, it is guaranteed to be a periodic signal. Instead of solving Eq. (4), we can also estimate $\widehat{b}$ by solving

$$\widehat{b} = \arg\min_{b} ||b||_0 \; s.t. \; ||y - \widehat{D}b||_2^2 \leq \epsilon \tag{5}$$

where ϵ is again a noise-dependent hyperparameter. This l_0 norm minimization problem can be solved approximately by sparse coding algorithms such as OMP [3]. In practice, we noticed that both versions of estimating $\widehat{b}$ (i.e. Eq. (4) and Eq. (5)) lead to very similar SNR gains.

3.3. Further discussion

Note that the denoising effect here is because of two operations. The outputs of RFB are thresholded to retain only the subspaces corresponding $r_1, \cdots, r_m$ in $\widehat{D}$ that have energy above a threshold. This thresholding attenuates the noise components. Another possible contributor to the denoising effect is Eq. (4) or Eq. (5), where some of the less significant subspaces among the retained subspaces $S_{r_1}, \cdots, S_{r_m}$ are eliminated from the representation of $\widehat{x}$ due to the sparsity promoting optimization problems with l_1 and l_0 norms.

The RFB output thresholding is done heuristically based on the expected noise variance. However, note that some of the low energy harmonics of a component signal may have a lower energy output at the corresponding filter than the noise energy at other filters. For example, for a period 6 signal having both period 2 and period 3 components, the retained subspaces may only be S_2 and S_6 if the period 3 component is of low energy and gets thresholded. In order to mitigate this unwanted dropping of low energy harmonics, we can carefully balance the two denoising contributors. One way is to set the RFB output threshold conservatively so as to not miss some of the low energy harmonics and hope that the noise subspaces retained at this stage will be dropped by the optimization algorithm in the next stage. Another way is to set RFB thresholds slightly higher so as to completely eliminate all the noise subspaces. To form a synthesis dictionary in this case, along with the subspaces corresponding to the indices $r_1, r_2, \cdots, r_m$, we also add subspaces corresponding to the **divisors** of $r_1, r_2, \cdots, r_m$. Let r_i have k_i divisors $d_{i1}, d_{i2}, \cdots, d_{ik_i}$. Thus, the synthesis dictionary is given by

$$\widehat{D} = [\underbrace{H_{d_{11}}, \cdots, H_{d_{1k_1}}}_{\text{divisors of } r_1} \; \cdots \; \underbrace{H_{d_{m1}}, \cdots, H_{d_{mk_m}}}_{\text{divisors of } r_m}] \tag{6}$$

This way, the lost subspace S_3 in the above example will be reintroduced, since 3 is a divisor of 6. See Fig. 2 for the complete denoising framework using Ramanujan analysis filter bank and pruned synthesis dictionary.

4. EXPERIMENTAL RESULTS

In this section, we compare our denoising framework with some of the well-known denoising methods. We consider a random unit-norm period 12 signal of length 100 and add white Gaussian noise so that the SNR is 0 dB. For comparison, we use the SNR gain metric that is defined as the increase in the SNR of the denoised signal compared to that of the noisy signal. Fig. 3 shows the signals denoised by different methods. For denoising with the proposed framework, we set the RFB threshold to 0.4 times the maximum energy output of the filter bank, $P_{max} = 40$, $l = 10$, and $\lambda = 0.01$. Fig. 3(c) shows the signal denoised by the proposed method using a 'with-divisors' synthesis dictionary (Eq. (6)), whereas Fig. 3(d) shows the signal denoised by the 'without-divisors' synthesis dictionary (Eq. (3)). The reconstruction is done by solving Eq. (4). Notice that the 'without-divisors' dictionary loses out on a lower energy harmonic, but the 'with-divisors' dictionary picks up the lost harmonic and provides a much better SNR gain. Although we have shown results with the l_1 method (Eq. (4)), the l_0 method (Eq. (5)) gives very similar SNR gains and denoised outputs with pruned dictionaries. In order to gauge whether the dictionary pruning really helps, we also denoised using the full Ramanujan dictionary as in

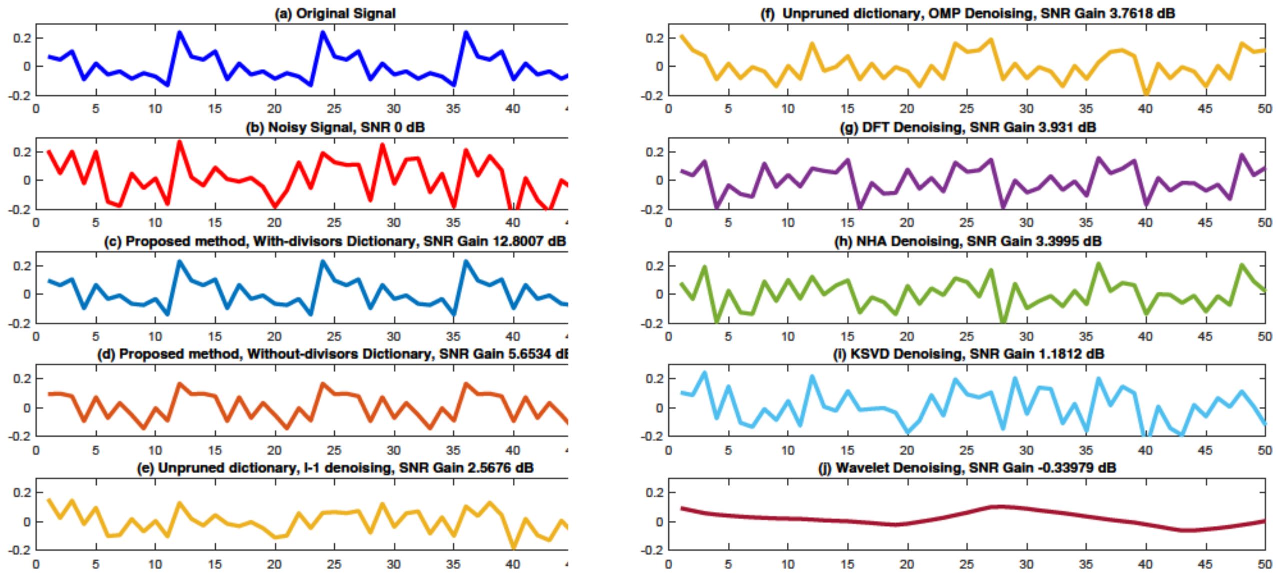

Fig. 3. Comparison of different denoising methods for a period 12 signal. See Sec. 4 for details and discussion.

Eq. (2) but the results from Fig. 3(e,f) make it clear that dictionary pruning is a crucial step without which the denoising performance is not as good.

Comparing the denoised output of the proposed method with the results from Fig. 3 (f-j), we see that our method provides much better results than OMP, DFT, NHA, wavelets, and K-SVD based denoising methods. None of these methods yields a truly periodic signal, which is a unique advantage of our method. For OMP denoising with the full Ramanujan dictionary (Fig. 3(f)), $\epsilon = 1$ was used. We shall note that OMP denoising gives better denoising results if the signal length is very large but P_{max} is kept low, in which case the dictionary is more often tall than fat. Although DFT denoising (Fig. 3(g)) does not perform that well here, it gives good results if the period is a divisor of the signal length. In such cases, the fundamental frequency falls on the DFT grid and hence DFT performs very well.

The NHA based denoising method (Fig. 3(h)) also does not perform very well. The main reason for this is that NHA requires good initialization from DFT to perform well, which the DFT may not provide when the noise level is significant and the period is not a divisor of the signal length. Moreover, since the NHA method requires iteratively adding frequency components one by one, it suffers from error propagation. K-SVD denoising (Fig. 3(i)), which is a data-adaptive technique, does not perform well for this setting. One possible reason is that with the small signal length, K-SVD is not able to learn a good dictionary from the patches of the noisy signal. Note that wavelet denoising (Fig. 3(j)) does not yield acceptable results. Wavelets are designed for good representation of signals with time-localized supports, such as, for example, bump signals [23], whereas the periodic signals are very different than this.

In Fig. 4, we compare the above denoising methods based on the SNR gains offered for different levels of noise in the input. For each value of input noise level, we average the SNR gains obtained for 1000 randomly generated period 12 signals corrupted with noise. The best sets of hyperparameters (such as the RFB threshold, λ, ϵ) are empirically chosen for each method at every input SNR by joint search over appropriate sets of values. Notice that the proposed method that uses the 'with-divisors' dictionary offers the best SNR gain that is stable over a large range of input SNR values, followed next by the 'without-divisors' dictionary method.

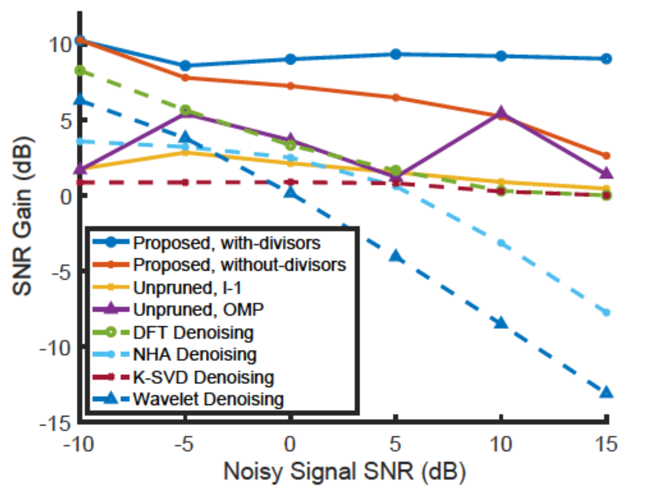

Fig. 4. Comparison of SNR Gains offered by different denoising methods averaged over 1000 Monte-Carlo runs

To compare the denoising performance for signals with multiple periodic components, we also considered signals that are combinations of period 7 and period 12 signals. Since the effective period is 7×12, data with only 100 samples hardly looks periodic. Even in this case, the proposed denoising method demonstrated results superior to the other methods. Details are omitted because of space constraints.

5. CONCLUDING REMARKS

In this paper, we proposed a hybrid analysis filter bank and synthesis dictionary framework for periodic signal denoising that combines the information at RFB output with the noisy input signal to extract a cleaner version of the input. Compared to other well-known methods, our framework offers the best SNR gain that is stable over a large range of input SNR values even when the signal length is not too large. Another unique advantage of our method is that the denoised signal is exactly periodic, unlike the other methods. An interesting task for the future would be to optimize each filter in the analysis bank, based on data, to minimize interference from input components belonging to other Ramanujan subspaces not represented by that filter.

6. REFERENCES

[1] S. Boll, "Suppression of acoustic noise in speech using spectral subtraction," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 27, no. 2, pp. 113–120, 1979.

[2] G. Chen, W. Xie, and Y. Zhao, "Wavelet-based denoising: A brief review," in *2013 Fourth International Conference on Intelligent Control and Information Processing (ICICIP)*, pp. 570–574, 2013.

[3] T. T. Cai and L. Wang, "Orthogonal matching pursuit for sparse signal recovery with noise," *IEEE Transactions on Information Theory*, vol. 57, no. 7, pp. 4680–4688, 2011.

[4] M. Protter and M. Elad, "Image sequence denoising via sparse and redundant representations," *IEEE Transactions on Image Processing*, vol. 18, no. 1, pp. 27–35, 2009.

[5] T. Yoshizawa, S. Hirobayashi, and T. Misawa, "Noise reduction for periodic signals using high-resolution frequency analysis," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2011, p. 1, Jul. 2011.

[6] P. Chatterjee and P. Milanfar, "Is denoising dead?," *IEEE Transactions on Image Processing*, vol. 19, no. 4, pp. 895–911, 2010.

[7] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4311–4322, 2006.

[8] D. Jia, K. Yanagisawa, Y. Ono, K. Hirobayashi, M. Hasegawa, S. Hirobayashi, H. Tagoshi, T. Narikawa, N. Uchikata, and H. Takahashi, "Multiwindow nonharmonic analysis method for gravitational waves," *IEEE Access*, vol. 6, pp. 48645–48655, 2018.

[9] D. Jia, K. Yanagisawa, M. Hasegawa, S. Hirobayashi, H. Tagoshi, T. Narikawa, N. Uchikata, and H. Takahashi, "Time–frequency-based non-harmonic analysis to reduce line-noise impact for LIGO observation system," *Astronomy and Computing*, vol. 25, pp. 238 – 246, 2018.

[10] S. V. Tenneti and P. P. Vaidyanathan, "Detection of protein repeats using the Ramanujan filter bank," in *2016 50th Asilomar Conference on Signals, Systems and Computers*, pp. 343–348, 2016.

[11] S. V. Tenneti and P. P. Vaidyanathan, "Detecting tandem repeats in dna using Ramanujan filter bank," in *2016 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 21–24, 2016.

[12] P. P. Vaidyanathan, "Ramanujan sums in the context of signal processing — Part I: Fundamentals," *IEEE Transactions on Signal Processing*, vol. 62, pp. 4145–4157, Aug. 2014.

[13] P. P. Vaidyanathan, "Ramanujan sums in the context of signal processing — Part II: FIR representations and applications," *IEEE Transactions on Signal Processing*, vol. 62, pp. 4158–4172, Aug. 2014.

[14] P. P. Vaidyanathan and S. Tenneti, "Srinivasa Ramanujan and signal-processing problems," *Philosophical Transactions of the Royal Society A*, vol. 378, Issue 2163, Dec. 2019.

[15] S. Ramanujan, "On certain trigonometrical sums and their applications in the theory of numbers," *Trans. of the Cambridge Phil. Soc.*, vol. XXII, no. 13, pp. 259–276, 1918.

[16] P. P. Vaidyanathan and P. Pal, "The Farey-dictionary for sparse representation of periodic signals," in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 360–364, 2014.

[17] S. V. Tenneti and P. P. Vaidyanathan, "Nested periodic matrices and dictionaries: New signal representations for period estimation," *IEEE Transactions on Signal Processing*, vol. 63, pp. 3736–3750, Jul. 2015.

[18] S. Tenneti and P. P. Vaidyanathan, "Ramanujan filter banks for estimation and tracking of periodicities," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3851–3855, Apr. 2015.

[19] P. P. Vaidyanathan and S. Tenneti, "Properties of Ramanujan filter banks," in *2015 23rd European Signal Processing Conference (EUSIPCO)*, pp. 2816–2820, Aug. 2015.

[20] P. P. Vaidyanathan and S. Tenneti, "Efficient multiplierless structures for Ramanujan filter banks," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6458–6462, 2017.

[21] P. P. Vaidyanathan, "Multirate systems and filter banks," 1992.

[22] K. Koh, S. Kim, and S. Boyd, "l1ls: Simple matlab solver for l1-regularized least squares problems." `https://web.stanford.edu/~boyd/l1_ls/`, 2008.

[23] D. L. Donoho, I. M. Johnstone, G. Kerkyacharian, and D. Picard, "Wavelet shrinkage: asymptopia?," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 57, no. 2, pp. 301–337, 1995.