CIRCULANT STRUCTURES AND GRAPH SIGNAL PROCESSING

Venkatesan. N. Ekambaram, Giulia. C. Fanti, Babak Ayazifar, Kannan Ramchandran

Department of EECS, UC Berkeley
{venkyne,gfanti,ayazifar,kannanr}@eecs.berkeley.edu

ABSTRACT

Linear shift-invariant processing of graph signals rests on circulant graphs and filters. The spatial features of circulant structures also permit shift-varying operations such as sampling. Their spectral features—as described by their Graph Fourier Transform profiles—enable novel multiscale signal processing systems and methods. To extend the reach of circulant structures, we present a method to decompose an arbitrary graph or filter into a combination of circulant structures. Our decomposition is analogous to resolving a linear time-varying system into a bank of linear time-invariant systems. As an application, we perform multiscale decomposition on temperature data spanning the continental United States.

Index Terms— Circulant graph, graph signal, graph Fourier transform, graph downsampling, circulant decomposition.

1. INTRODUCTION

Numerous modern applications present graph-structured data. Examples include social media (e.g., Facebook and Twitter [1]), wireless sensors (e.g., temperature measurements [2]), power networks [3], computer graphics [4, 5], and finite-element meshes [6]. Researchers in machine learning and computer graphics have spent considerable effort developing techniques to tackle data from large-scale networks. The signal processing community, too, has witnessed a surge of interest in a unified theory for the analysis and processing of *graph signals*—signals measured at, or otherwise defined on, the vertices of a graph. Traditional signal processing deals with signals measured over regular or well-ordered structures, such as time lines or Cartesian grids. Graph signals, however, present a substantially more nuanced paradigm.

Motivated by the impact of Fourier and wavelet analysis in classical signal processing, researchers have generalized some of these tools to the more abstract, irregular domains of graphs and manifolds. Questions of interest include: What functions defined on graphs can serve as Fourier bases? Which graph signals have content at high or low frequencies? What interesting properties, if any, does a graph Fourier decomposition possess with respect to signal analysis or filter design? Is there a fundamental tradeoff, such as an uncertainty principle, governing graph signals?

To tackle such questions, some researchers have focused on the spectral properties of the Laplacian matrix of a graph [5, 7–9]. For example, they have defined a *Graph Fourier Transform* [4, 10–12] based on the eigenvectors of the Laplacian matrix. Earlier work in computer graphics used mesh processing to analyze 3-D structures, and exploited Laplacian matrix properties to define graph-based filters for smoothing [4–6]. Recent effort has also been directed at basic operations on graph signals—such as shifting, sampling, and filtering—and understanding the properties of the GFT in relation to these. Researchers have defined them in a variety of ways and derived corresponding properties [13, 14]. Some authors have developed sampling theorems for special classes of graphs, such as bipartite graphs [12]. There are also results for sampling on general graphs [10, 15]. Other authors have derived uncertainty principles that describe the tradeoff between a graph signal's spectral and spatial spreads [16]. Multiscale graph signal processing research has focused on multilevel bases analogous to wavelets [11, 17–20].

Recent progress notwithstanding, much work remains to be done in the quest for a unified—let alone comprehensive—signal processing theory for graphs. A major challenge in defining even fundamental operations, such as signal shifting, is a general graph's irregular structure. A typical graph looks different from different nodes, so it is unclear how to shift or sample a signal on it. Authors have either resorted to unconventional definitions for these operations [13] or defined them in domains different from those of the signals [14]. Each approach has its own justifications, advantages, and disadvantages.

Our approach is motivated by linear time-invariant (LTI) filtering in classical signal processing. We begin with networks that have *circulant* structures, because their graphs are amenable to *linear shift-invariant (LSI)* operations.[1] We show that even certain shift-varying operations, such as sampling, are natural on such graphs.

A corpus of research literature exploits the representation of *linear time-varying (LTV)* systems as a bank of LTI systems [22–26]. We outline a method to decompose a general graph into a bank of circulant graphs. We do the same for filters. A linear filter defined on an N-vertex graph is a square $N \times N$ matrix. If the filter matrix is circulant for some node ordering, it can perform *linear shift-invariant (LSI)* processing. A matrix $\mathbf{F}$ is circulant if each of its rows is circularly shifted to the right by one element relative to its preceding row—that is, $F((i)_N, (j)_N) = F((i+1)_N, (j+1)_N)$, where $()_N$ denotes $\mod N$. A general filter is *not* circulant, so it can perform only *linear shift-varying (LSV)* processing. However, we can decompose any LSV filter into a bank of LSI filters. To demonstrate the utility of our techniques, we perform a multiscale decomposition on a temperature dataset extending over the United States.

2. GRAPH FOURIER TRANSFORM

Throughout, we consider only an undirected connected graph $G = (V, E)$ having no self-loops or multiple edges (i.e., the graph is simple), where $V = \{0, 1, \cdots, N-1\}$ is the set of N vertices (nodes) and E is the set of edges. Let $\mathbf{A}$ be the adjacency matrix for the graph, where $A(i,j) > 0$ if nodes i and j are connected by an edge; otherwise, $A(i,j) = 0$. A graph signal $\boldsymbol{x} =$

This research is supported in part by AFOSR grant FA9550-10-1-0567.

[1] Existing work discusses the shift-invariant nature of circulant graphs, but the authors do not explore the properties in substantive depth [21].

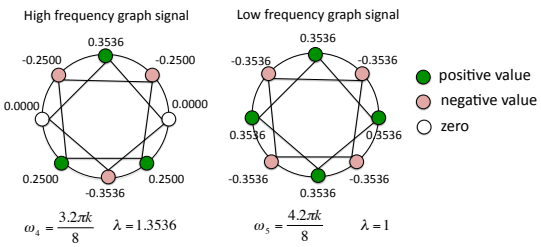

Fig. 1. Example for the permutation. A higher frequency DFT vector is a lower frequency vector on the graph. Note that only the real part of the eigenvectors is shown here.

$[x(0)\ x(1)\ \cdots\ x(N-1)]^{\mathsf{T}}$ is a real-valued function defined on the vertices of the graph; the symbol T denotes matrix or vector transposition. Let $\mathbf{D}$ be the diagonal matrix of vertex degrees, where the diagonal entry $D(i,i) = \sum_{j=0}^{N-1} A(i,j)$. The *Laplacian matrix* of the graph is defined as $\mathbf{L} = \mathbf{D} - \mathbf{A}$. It is a positive semi-definite matrix and has the spectral decomposition $\mathbf{L} = \mathbf{U}\mathbf{\Lambda}\mathbf{U}^{\mathsf{H}}$, where $\mathbf{\Lambda}$ is a diagonal matrix of non-negative real eigenvalues $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_{N-1}$; $\mathbf{U} = [\boldsymbol{u}_0 \mid \boldsymbol{u}_2 \mid \cdots \mid \boldsymbol{u}_{N-1}]$ is the corresponding matrix of eigenvectors $\boldsymbol{u}_k$, which form an orthonormal basis for $\mathbb{R}^N$; and H denotes Hermitian transpose. The *Graph Fourier Transform (GFT)* $\boldsymbol{X}^G$ of the signal $\boldsymbol{x}$ is obtained from the decomposition of $\boldsymbol{x}$ with respect to the orthonormal basis specified by the columns of $\mathbf{U}$ (see, e.g., [5]); that is, $\boldsymbol{x} = \mathbf{U}\boldsymbol{X}^G = \sum_{k=0}^{N-1} X^G(k)\,\boldsymbol{u}_k$, where $X^G(k) = \langle \boldsymbol{x}, \boldsymbol{u}_k \rangle = \boldsymbol{u}_k^{\mathsf{H}}\boldsymbol{x}$. Clearly, $\boldsymbol{X}^G = \mathbf{U}^{\mathsf{H}}\boldsymbol{x}$.

The definition is supported by nodal domain theorems that justify high-frequency and low-frequency interpretations of Laplacian matrix eigenvectors [27]. In general, a graph eigenvector corresponding to a larger eigenvalue is associated with a higher frequency (see Fig. 1). Qualitatively, a high-frequency graph signal varies significantly over neighboring nodes. Once a node ordering is fixed at the outset, we can map each element of $\mathbf{u}_k$ to a corresponding node on the graph and thus view each $\mathbf{u}_k$ as a graph signal in its own right. The numbers of zero crossings of the eigenvectors increase monotonically with increasing eigenvalues. For a ring graph, the eigenvectors are the Discrete Fourier Transform (DFT) vectors, which provide the spectral representation of periodic discrete-time signals in the classical setting. Sinusoids are eigenfunctions of the double-differential operator in the time domain. Similarly, the Laplacian matrix is analogous to the double-differencing operator on a graph. Other analogies to the time domain, too, motivate the choice of the Laplacian matrix eigenvectors as a useful orthogonal basis [16].

2.1. The Case for Circulant Graphs

LTI filter design has had an enormous impact on signal processing. Accordingly, we wish to develop linear-shift-invariant (LSI) filters on graphs. In particular, if we want to shift a signal value to an adjacent node, shift-invariance requires the signal on the graph to "look the same" after shifting, which is not true of a general graph. *Circulant* graphs allow for this. A graph G is circulant if there exists some ordering of nodes for which the adjacency matrix (or equivalently, the Laplacian matrix) of the graph is circulant. An alternative definition of a circulant graph arises from the way it is constructed. A graph G is circulant with respect to a generating set $S = \{s_1, s_2, ..., s_M\}$, where $0 < s_k \leq N-1$ whenever there is an edge between the node pair $(i, (i+s_k)_N)$, for every $s_k \in S$. An image of size $r \times r$ can be thought of as a circulant graph with set $S = \{1, r\}$. To simplify our analysis, we consider graphs of size $N = 2^n$. The following lemma offers important analytical simplifications.

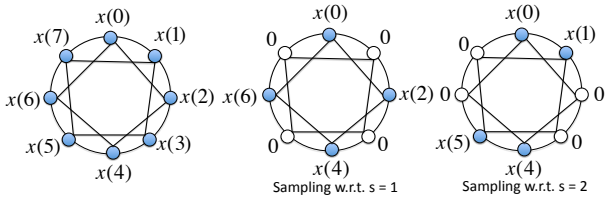

Fig. 2. Sampling with respect to different elements in S.

Lemma 1. *For any connected circulant graph of* 2^n *nodes,* $1 \in S$.

The GFT of a signal defined on a circulant graph has a close relationship with the signal's *Discrete Fourier Transform (DFT)*. It is well known that the columns of the DFT matrix are eigenvectors of any circulant matrix of the same size. The DFT coefficients of the first column of the circulant matrix are the eigenvalues. There is, however, a difference in the ordering of the Fourier basis. In particular, the k^{th} column of the DFT matrix, which corresponds to the frequency $\omega_k = 2\pi k/N$, does not necessarily correspond to $\mathbf{u}_k$, the k^{th} basis vector associated with λ_k. If $\boldsymbol{X}^G$ is the GFT, and $\boldsymbol{X}^F$ the DFT, of a graph signal $\boldsymbol{x}$, then there is a permutation σ that maps the components of $\boldsymbol{X}^F$ to those of $\boldsymbol{X}^G$, i.e., $X^G(k) = X^F(\sigma(k))$.

The permutation is completely determined by taking the DFT of the first column of the Laplacian matrix and reordering the coefficients. Due to the even multiplicities of the eigenvalues (other than $\lambda_0 = 0$ and $\lambda_{N/2}$), we could also have a set of real eigenvectors. In this case, we must modify the permutation relationship between the GFT and the DFT. For simplicity, we work with the DFT vectors, though our results can be extended to any choice of eigenvectors. We first define the notions of shifting and sampling on circulant graphs and then describe a circulant decompositions of general graphs.

2.1.1. Shifting

In the time domain, a shift by one unit involves moving every signal value to an adjacent time sample. Analogously, we define a shift by one unit, as moving the signal value on each node to a neighboring node. In general, each node has multiple neighbors determined by the set S. So we define shifting with respect to each element of S. For example, a shift by one unit with respect to $s = 3$, is a shift of the signal value from node v to node $(v+3)_N$, for every $v \in V$. Each shift can be defined in terms of the matrix $\mathbf{P} = \begin{bmatrix} \mathbf{0}_{N-1}^{\mathsf{T}} & 1 \\ \mathbf{I}_{N-1} & \mathbf{0}_{N-1} \end{bmatrix}$, where $\mathbf{0}_{N-1}$ is a column vector of $N-1$ zeros, and $\mathbf{I}_{N-1}$ the identity matrix of size $N-1$. For a connected circulant graph of size N, a shift by one unit defined with respect to $s = 1$ in S corresponds to $\mathbf{P}\boldsymbol{x}$. In general, a shift by k units with respect to s in S corresponds to $\mathbf{P}^{ks}\boldsymbol{x}$. Note that $\mathbf{P^k} = \mathbf{P}^{\mathbf{k} \bmod \mathbf{N}}$ for any integer k.

2.1.2. Sampling

A proper definition of sampling on a circulant graph requires two parameters: a downsampling factor m and an element $s \in S$ with respect to which the downsampling is to be performed. For example, if we want to downsample a graph signal $\boldsymbol{x}$ by the factor 2 with respect to the element $s = 1$, we keep every alternate value of $\boldsymbol{x}$. In general, when we want to downsample by a factor m with respect to an element $s \in S$, we start at Node 0, look at Nodes $s, 2s, ...$ and keep every m^{th} element (see Fig. 2). Then we focus on Node 1, repeating the steps above for Nodes $s+1, 2s+1, ...$ We continue this until we reach Node $gcd(s, N)$. This is equivalent to sampling

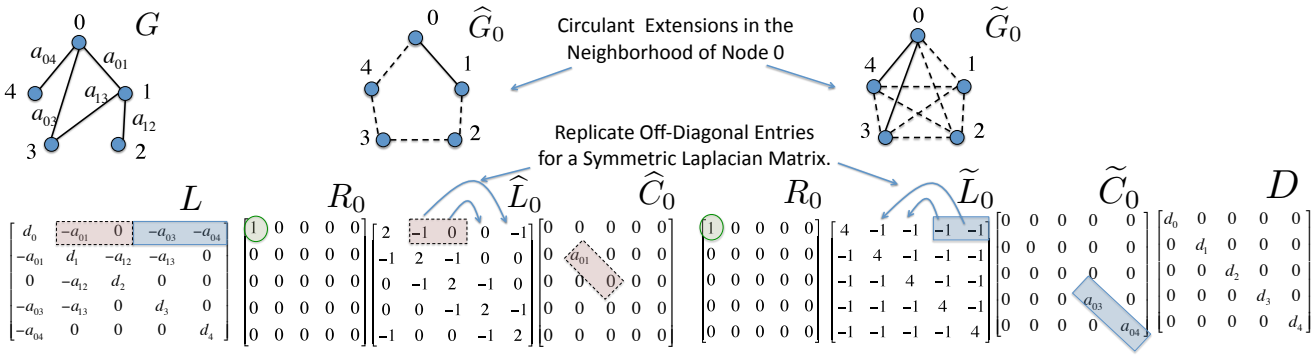

Fig. 3. Example of the circulant decomposition for the given graph G. The example here illustrates the circulant graphs used to reconstruct Row 0 of $\mathbf{L}$. Also shown are the pre- and post-multiplying diagonal matrices that select appropriate row and column entries.

the cosets of the associated Cayley group [28]. In our particular context, given a downsampled graph and its corresponding signal, upsampling involves assigning zero signal values to the nodes that were removed from the original graph. Aliasing expressions can be derived using the properties of the DFT basis and the known permutation. However, we omit their discussion due to space constraints.

2.2. Circulant Decomposition of Non-Circulant Graphs

There are various ways to decompose a non-circulant graph into circulant graphs of identical size; this is because an arbitrary non-circulant matrix can be decomposed into a combination of circulant matrices. It is well known that a circulant matrix $\mathbf{F}$ can be expressed as $\mathbf{F} = \sum_{\ell=0}^{N-1} f(\ell)\,\mathbf{P}^{\ell}$, where $\boldsymbol{f} = [f(0), f(1), ..., f(N-1)]^{\mathsf{T}}$ is the first column of $\mathbf{F}$. Based on a *circulant extension* of the neighborhood of each node, we show that this result can be generalized to the Laplacian matrix of a non-circulant graph.

Fig. 3 depicts a circulant decomposition of the five-node graph G. The key idea is to reconstruct each row of the graph's Laplacian matrix $\mathbf{L}$ using appropriate circulant Laplacian matrices that are pre- and post-multiplied by suitable diagonal matrices. To illustrate the method, we reconstruct row 0 of $\mathbf{L}$. Each of the matrices $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$ in Fig. 3 inherits half of the off-diagonal entries of row 0 of $\mathbf{L}$. We want $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$ to be the Laplacian matrices of undirected circulant graphs, labeled $\widehat{G}_0$ and $\widetilde{G}_0$ in Fig. 3, respectively. The graphs $\widehat{G}_0$ and $\widetilde{G}_0$ can be viewed as circulant extensions of the neighborhood of node 0 of G. Per force, half of the row entries of $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$ completely specify each full matrix. For example, once we fix $\widehat{L}_0(0,1)$ and $\widehat{L}_0(0,2)$, it must be that $\widehat{L}_0(0,3) = \widehat{L}_0(0,2)$ and $\widehat{L}_0(0,4) = \widehat{L}_0(0,1)$. The last two entries of row 0 in $\widetilde{\mathbf{L}}_0$, too, must be mirror images of the corresponding first two off-diagonal entries.

Each successive row of $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$ is obtained by a right circular shift of its preceding row. The matrix $\mathbf{R}_0$ selects row 0 of $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$. The matrices $\widehat{\mathbf{C}}_0$ and $\widetilde{\mathbf{C}}_0$ select the corresponding entries of row 0 of $\widehat{\mathbf{L}}_0$ and $\widetilde{\mathbf{L}}_0$, respectively—which match the corresponding off-diagonal entries in $\mathbf{L}$. Since $\mathbf{R}_0\,\widehat{\mathbf{L}}_0\widehat{\mathbf{C}}_0 + \mathbf{R}_0\,\widetilde{\mathbf{L}}_0\widetilde{\mathbf{C}}_0$ does *not* reproduce $L(0,0)$, we add $D(0,0)$ to the appropriate entry of $\mathbf{D}$.

Theorem 1. (Circulant Decomposition of Laplacian Matrices) *Consider a graph G whose Laplacian and adjacency matrices are $\mathbf{L}$ and $\mathbf{A}$, respectively. Then $\mathbf{L}$ can be written as*

$$\mathbf{L} \;=\; \mathbf{D} + \overbrace{\sum_{k=0}^{N-1} (\mathbf{R}_k\,\widehat{\mathbf{L}}_k\,\widehat{\mathbf{C}}_k + \mathbf{R}_k\,\widetilde{\mathbf{L}}_k\,\widetilde{\mathbf{C}}_k)}^{=-\mathbf{A}},$$

where $\mathbf{D}$ is the degree matrix of G; each $\mathbf{R}_k$ is a matrix whose only non-zero entry is $R_k(k,k) = 1$; and each $\widehat{\mathbf{C}}_k$ and $\widetilde{\mathbf{C}}_k$ is a diagonal matrix whose non-zero entries we now describe in their most general form. For $i = 1, \ldots, \lfloor N/2 \rfloor$,

$$\widehat{C}_k((k+i)_N, (k+i)_N) = -L(k, (k+i)_N) = A(k, (k+i)_N).$$

For $i = 1 + \lfloor N/2 \rfloor, \ldots, N-1$,

$$\widetilde{C}_k((k+i)_N, (k+i)_N) = -L(k, (k+i)_N) = A(k, (k+i)_N).$$

The Laplacian matrices $\widehat{\mathbf{L}}_k$ and $\widetilde{\mathbf{L}}_k$ represent simple, unweighted, and undirected circulant graphs defined by the sets $\widehat{S}_k$ and $\widetilde{S}_k$, respectively, characterized below:

$$\widehat{S}_k = \Big\{ i \;\Big|\; i \in \{1, \ldots, \lfloor N/2 \rfloor\} \wedge L(k, (k+i)_N) \neq 0 \Big\},$$
$$\widetilde{S}_k = \Big\{ i \;\Big|\; i \in \{1 + \lfloor N/2 \rfloor, \ldots, N-1\} \wedge L(k, (k+i)_N) \neq 0 \Big\}.$$

Proof: Use the construction method as described for Fig. 3. ■

Corollary 1. (Circulant Decomposition of Adjacency Matrices) $\mathbf{A} = \sum_{k=0}^{N-1} (\mathbf{R}_k\,\widehat{\mathbf{A}}_k\,\widehat{\mathbf{C}}_k + \mathbf{R}_k\,\widetilde{\mathbf{A}}_k\,\widetilde{\mathbf{C}}_k)$, *where $\widehat{\mathbf{A}}_k$ and $\widetilde{\mathbf{A}}_k$ are the adjacency matrices of the undirected circulant graphs $\widehat{G}_k$ and $\widetilde{G}_k$, respectively.*

Proof: Note that $\widehat{\mathbf{L}}_k = \widehat{d}_k\mathbf{I} - \widehat{\mathbf{A}}_\mathbf{k}$ and $\widetilde{\mathbf{L}}_k = \widetilde{d}_k\mathbf{I} - \widetilde{\mathbf{A}}_\mathbf{k}$ for some positive values $\widehat{d}_k$ and $\widetilde{d}_k$, and exploit the fact that $\mathbf{R_k}\,\widehat{\mathbf{C}}_\mathbf{k} = \mathbf{R_k}\,\widetilde{\mathbf{C}}_\mathbf{k} = \mathbf{0_{N\times N}}$. ■

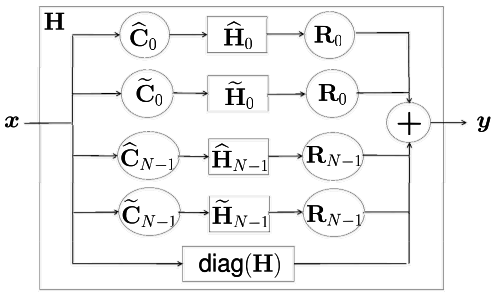

Fig. 4. Decomposition of an LSV filter on a general graph as a bank of LSI filters on individual circulant graphs.

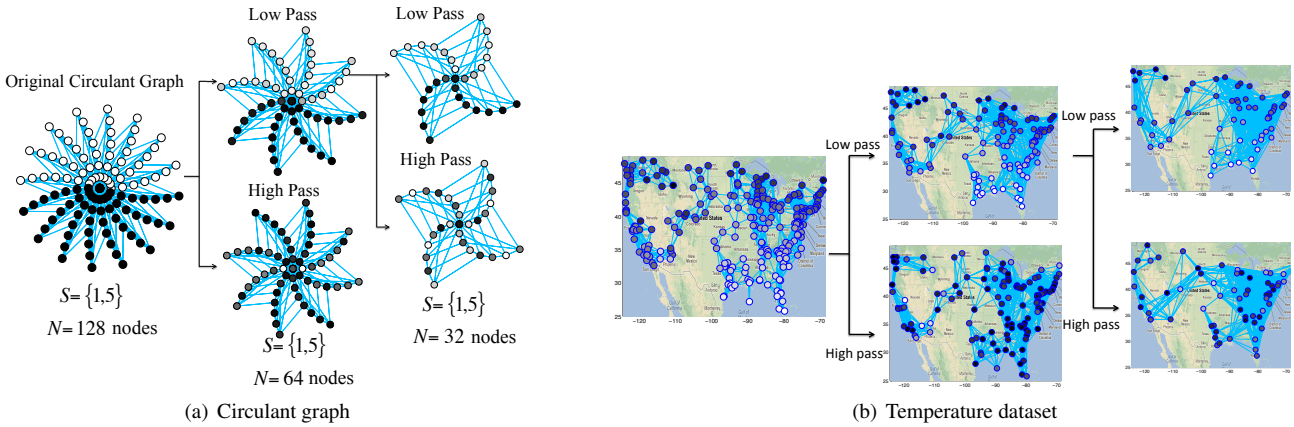

(a) Circulant graph (b) Temperature dataset

Fig. 5. Multiscale decomposition of a synthetic dataset on a circulant graph and a real-world temperature dataset. One can visually observe that the low-pass branches retain most of the content of the original signal and the high-pass branches retain the changes in the signal values.

2.3. Filtering on Graphs

We can now discuss linear filtering operations on graphs. Let $\mathbf{I_A}$ denote the *indicator matrix* for $\mathbf{A}$; that is, $I_{\mathbf{A}}(i,j) = 1$ if $A(i,j) \neq 0$, and $I_{\mathbf{A}}(i,j) = 0$ if $A(i,j) = 0$.

Definition 1: (A k-Hop Linear Filter) We say $\mathbf{H}$ is a 1-hop linear filter on a graph G if $\mathbf{H} = \mathbf{H} \circ (\mathbf{I_A} + \mathbf{I})$, where $\mathbf{I}$ is the identity matrix. Thus we want the non-zero off-diagonal entries of $\mathbf{H}$ coincide with those of the graph's adjacency matrix $\mathbf{A}$. The symbol $\circ$ denotes the standard Hadamard (element-wise) product. It is known that the number of walks of length k from node i to node j is given by the entry (i,j) of $(\mathbf{I_A})^{\mathbf{k}}$ [29]. So, a k-hop linear filter is characterized by the k-hop indicator matrix $\mathbf{I}_{(\mathbf{I_A})^{\mathbf{k}}}$, which is the same as $\mathbf{I}_{\mathbf{A^k}}$ since $\mathbf{A}$ has no negative entries.

Definition 2: (Linear Shift-Invariant Filtering on Graphs) A filter $\mathbf{H}$ is LSI if $\mathbf{P}\,\mathbf{H}\,\boldsymbol{x} = \mathbf{H}\,\mathbf{P}\,\boldsymbol{x}$ for every graph signal $\boldsymbol{x}$.

Theorem 2. *If* $\mathbf{H}$ *is a* 1*-hop LSV filter on* G*, then*

$$\mathbf{H} = \mathbf{D_H} + \sum_{k=0}^{N-1} \left[\mathbf{R}_k((\mathbf{I}_{\widehat{\mathbf{A}}_k} + \mathbf{I}) \circ \widehat{\mathbf{H}}_k)\widehat{\mathbf{C}}_k + \mathbf{R}_k((\mathbf{I}_{\widetilde{\mathbf{A}}_k} + \mathbf{I}) \circ \widetilde{\mathbf{H}}_k)\widetilde{\mathbf{C}}_k \right],$$

where $\widehat{\mathbf{H}}_k$ *and* $\widetilde{\mathbf{H}}_k$ *are LSI filters on circulant graphs constructed from* $\mathbf{H}$ *in a manner analogous to how the matrices* $\widehat{\mathbf{L}}_k$ *and* $\widetilde{\mathbf{L}}_k$ *are constructed from* $\mathbf{L}$ *in Theorem 1. The matrices* $\widehat{\mathbf{A}}_k$ *and* $\widetilde{\mathbf{A}}_k$ *are the adjacency matrices corresponding to* $\widehat{\mathbf{L}}_k$ *and* $\widetilde{\mathbf{L}}_k$*, respectively. The matrices* $\mathbf{R_k}$*,* $\widehat{\mathbf{C}}_k$*, and* $\widetilde{\mathbf{C}}_k$ *are as defined in Theorem 1 and* $\mathbf{D_H} = \mathrm{diag}(\mathbf{H})$.

Proof: Decompose $\mathbf{H}$ using the method of Theorem 1. ■

Fig. 4 shows a block diagram depiction of this decomposition.

3. EXPERIMENTS

To demonstrate the utility of this framework, we perform multiscale decompositions of synthetic and real-world graph signals; such decompositions can enable efficient visualization and processing of large datasets. We first consider the circulant graph of Fig. 5(a). The generating set S for the graph contains only odd numbers, and the number N of nodes is a power of 2. We apply ideal high-pass and low-pass filters to achieve a multiscale representation of the signal and the graph (Fig. 5(a)). That the two components at each stage of the decomposition are representative of the low-pass and high-pass contents of the signal is visually apparent. Each downsampled graph is obtained by connecting the retained nodes so that every alternate pair of nodes along the path defined by the element $s \in S$ is connected. In this manner, we can perform multi-level signal decomposition. The high-pass and low-pass branches also allow for perfect reconstruction of graph signals which we do not prove here due to space constraints.

We have also designed critically-sampled perfect-reconstruction filter banks inspired by time-varying lattice filters and Haar filters, which we apply to a real-world temperature dataset. At each stage we retain $N/2$ components at the high-pass and low-pass outputs. From the Federal Climate Complex's weather data we extracted average temperature measurements collected during March 2012 across the continental United States [2]. Based on these measurements we generated a graph of 347 weather stations as nodes. We connected by an edge any node pair whose geographic distance was below the threshold of 400 km. The graph captures temperature correlations across proximate locations.

The temperature signal is plotted on the graph according to intensity—lighter coloration indicating higher temperature. The signal undergoes two rounds of filtering and sampling. Our decomposition uses Haar filters and gives rise to the multiscale representation in Fig. 5(b). Visually, the representation produces results that are to be expected—at least intuitively. However, more work is needed to study the performance of this decomposition technique quantitatively. The framework allows us to obtain an arbitrarily small representation of a graph and its signal while preserving certain spectral properties. This could prove to be a useful tool for analyzing signals defined on large graphs.

4. FUTURE DIRECTIONS

Graph signal processing based on circulant structures presents numerous open questions that inform future research directions. For example, discussion of filters should go beyond low-pass and high-pass varieties. Our vision is to leverage LTI system theory to design efficient filters for graph signals. We have analytical results on aliasing patterns that can arise from the downsampling of graph signals. We also have preliminary results for critically-sampled perfect-reconstruction filter banks. Recently we have begun exploring applications in graph-based semi-supervised learning. We are optimistic that exploiting circulant structures is a promising complement to the existing graph signal processing theories and techniques.

5. REFERENCES

[1] B. Krishnamurthy, P. Gill, and M. Arlitt, "A few chirps about twitter," in *Proceedings of the first workshop on Online social networks*. ACM, 2008, pp. 19–24.

[2] "Gsod : Global surface summary of day," http://www.ncdc.noaa.gov/cgi−bin/res40.pl?page=climvisgsod.html

[3] Å. J. Holmgren, "Using graph models to analyze the vulnerability of electric power networks," *Risk analysis*, vol. 26, no. 4, pp. 955–969, 2006.

[4] G. Taubin, T. Zhang, and G. Golub, "Optimal surface smoothing as filter design," *Computer VisionECCV'96*, pp. 283–292, 1996.

[5] G. Taubin, "A signal processing approach to fair surface design," in *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*. ACM, 1995, pp. 351–358.

[6] H. Zhang, O. Van Kaick, and R. Dyer, "Spectral mesh processing," in *Computer graphics forum*. Wiley Online Library, 2010, vol. 29, pp. 1865–1894.

[7] F. R. K. Chung, "Lectures on spectral graph theory," *CBMS Lectures, Fresno*, 1996.

[8] R. Merris, "Laplacian matrices of graphs: a survey," *Linear algebra and its applications*, vol. 197, pp. 143–176, 1994.

[9] B. Mohar, "The laplacian spectrum of graphs," *Graph theory, combinatorics, and applications*, vol. 2, pp. 871–898, 1991.

[10] I. Z. Pesenson and M. Z. Pesenson, "Sampling, filtering and sparse approximations on combinatorial graphs," *Journal of Fourier Analysis and Applications*, vol. 16, no. 6, pp. 921–942, 2010.

[11] D. K. Hammond, P. Vandergheynst, and R. Gribonval, "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, vol. 30, no. 2, pp. 129–150, 2011.

[12] S. K. Narang and A. Ortega, "Downsampling graphs using spectral theory," in *Acoustics, Speech and Signal Processing (ICASSP), 2011 IEEE International Conference on*. IEEE, 2011, pp. 4208–4211.

[13] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *Signal Processing, IEEE Transactions on*, vol. 61, no. 7, pp. 1644–1656, 2013.

[14] D. I. Shuman, S. K. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "Signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular data domains," *arXiv preprint arXiv:1211.0053*, 2012.

[15] P. E. T. Jorgensen, "A sampling theory for infinite weighted graphs," *Opuscula Mathematica*, vol. 31, no. 2, pp. 209–236, 2011.

[16] A. Agaskar and Y. M. Lu, "An uncertainty principle for functions defined on graphs," in *SPIE Optical Engineering+ Applications*. International Society for Optics and Photonics, 2011, pp. 81380T–81380T.

[17] R. R. Coifman and M. Maggioni, "Diffusion wavelets," *Applied and Computational Harmonic Analysis*, vol. 21, no. 1, pp. 53–94, 2006.

[18] M. Gavish, B. Nadler, and R. R. Coifman, "Multiscale wavelets on trees, graphs and high dimensional data: Theory and applications to semi supervised learning," in *Proc. International Conference on Machine Learning, Haifa, Israel*, 2010.

[19] S. K. Narang and A. Ortega, "Perfect reconstruction two-channel wavelet filter banks for graph structured data," *Signal Processing, IEEE Transactions on*, vol. 60, no. 6, pp. 2786–2799, 2012.

[20] S. K. Narang and A. Ortega, "Local two-channel critically sampled filter-banks on graphs," in *Proc. IEEE International Conference on Image Processing (ICIP)*, 2010, pp. 333–336.

[21] L. J. Grady and J. R. Polimeni, *Discrete Calculus: Applied Analysis on Graphs for Computational Science*, Springer, 2010.

[22] G. Lohar and A. G. Wacker, "Matrix theory approach to the canonical representation of a class of linear discrete-time time-variant systems," *Circuits and Systems, IEEE Transactions on*, vol. 37, no. 2, pp. 303–306, 1990.

[23] G. Lohar, D. P. Mukherjee, and D. Dutta Majumder, "On a decomposition of 2-d circular convolution," *Pattern recognition letters*, vol. 13, no. 10, pp. 701–706, 1992.

[24] C. Herley and M. Vetterli, "Orthogonal time-varying filter banks and wavelet packets," *Signal Processing, IEEE Transactions on*, vol. 42, no. 10, pp. 2650–2663, 1994.

[25] R. Meyer and C. S. Burrus, "A unified analysis of multirate and periodically time-varying digital filters," *Circuits and Systems, IEEE Transactions on*, vol. 22, no. 3, pp. 162–168, 1975.

[26] S. M. Phoong and P. P. Vaidyanathan, "A polyphase approach to time-varying filter banks," in *Acoustics, Speech, and Signal Processing, 1996. ICASSP-96. Conference Proceedings., 1996 IEEE International Conference on*. IEEE, 1996, vol. 3, pp. 1554–1557.

[27] E. BrianDavies, G. M. L. Gladwell, J. Leydold, and P. F. Stadler, "Discrete nodal domain theorems," *Linear Algebra and its Applications*, vol. 336, no. 1, pp. 51–60, 2001.

[28] C. D. Godsil and G. Royle, *Algebraic graph theory*, vol. 8, Springer New York, 2001.

[29] N. Biggs, *Algebraic graph theory*, Cambridge University Press, 1993.