

Reconstruction of Graph Signals through Percolation from Seeding Nodes

Santiago Segarra, *Student Member, IEEE*, Antonio G. Marques, *Senior Member, IEEE*, Geert Leus, *Fellow, IEEE*, and Alejandro Ribeiro, *Member, IEEE*

Abstract—New schemes to recover signals defined in the nodes of a graph are proposed. Our focus is on reconstructing bandlimited graph signals, which are signals that admit a sparse representation in a frequency domain related to the structure of the graph. Most existing formulations focus on estimating an *unknown* graph signal by *observing* its value on a subset of nodes. By contrast, in this paper, we study the problem of inducing a *known* graph signal using as input a graph signal that is nonzero only for a small subset of nodes (seeding nodes). The sparse signal is then percolated (interpolated) across the graph using a graph filter. Graph filters are a generalization of classical time-invariant systems and represent linear transformations that can be implemented distributedly across the nodes of the graph. Alternatively, one can interpret graph signals as network states and study graph-signal reconstruction as a network-control problem where the target class of states is represented by bandlimited signals. Three setups are investigated. In the first one, a single simultaneous injection takes place on several nodes in the graph. This setting resembles that of reconstruction after sampling, but in this paper the interpolator is constrained to be a graph filter. In the second one, successive value injections take place on a single node. The third one is a generalization where multiple nodes inject multiple signal values. For noiseless settings, conditions under which perfect reconstruction is feasible are given, and the corresponding schemes to recover the desired signal are specified. Scenarios leading to imperfect reconstruction, either due to insufficient or noisy signal value injections, are also analyzed. Moreover, connections with classical interpolation in the time domain are discussed. Specifically, for time-varying signals, where the ideal interpolator after uniform sampling is a (low-pass) filter, our proposed approach and the reconstruction of a sampled signal coincide. Nevertheless, for general graph signals, we show that these two approaches differ. The last part of the paper presents numerical experiments that illustrate the results developed through synthetic graph signals and two real-world signal reconstruction problems: influencing opinions in a social network and inducing a desired brain state in humans.

Index Terms—Graph signal processing, Signal reconstruction, Interpolation, Percolation, Graph-shift operator, Bandlimited graph signals

I. INTRODUCTION

Sampling and reconstruction of bandlimited signals are cornerstone problems in classical signal processing. The emergence of new fields of knowledge such as network science and big data is generating a pressing need to extend the

results existing for classical time-varying signals to signals defined on graphs [3]–[5]. This not only entails modifying the existing algorithms, but also gaining intuition on the concepts that are preserved and lost when a signal is defined not in the time grid, but in a more general graph domain. In the context of reconstruction of graph signals, two different approaches have been developed. A first approach is related to the interpolation of bandlimited signals, which consists in inferring unobserved values by leveraging the fact that the signal lives in a low-dimensional space [6]–[10]. Although most interpolation approaches are centralized, iterative [11] and distributed [12], [13] interpolation schemes have also been developed. A different approach towards graph signal reconstruction is graph regularization [14], [15] where a notion of smoothness is assumed on the signal and the unobserved values are estimated based on this notion. Both approaches coincide in that they estimate a graph signal from the observation of a subset of the signal values. By contrast, the approach in this paper is to preserve the two-step methodology used when recovering bandlimited time-varying signals, which consists in the generation of a sparse signal followed by the application of a low-pass filter to reconstruct the missing entries, and extend it to the more general graph domain.

To be more specific, we study the reconstruction of bandlimited graph signals through the application of low-pass *graph filters* to sparse *seeding signals*. Graph filters are the generalization of classical time-invariant systems when the signals are defined on a general graph as opposed to the classical time domain [4]. Seeding signals are graph signals attaining nonzero values on a subset of the nodes in the graph, called seeding nodes. Alternatively, one can view graph signals as network states and, thus, reinterpret the seeding values as local control inputs designed to steer the network to a target state associated with a bandlimited signal. Further connections between our proposed approach and the study of controllability for complex networks [16]–[22] are detailed in Section II-B. To describe our contribution more precisely, let \mathbf{y} stand for the graph signal that we want to induce in the graph. Our goal is to design a graph filter \mathbf{H} and a sparse signal \mathbf{x} such that \mathbf{y} can be obtained upon applying \mathbf{H} to \mathbf{x} . The design is accomplished in two steps. In the first step, we design the filter \mathbf{H} leveraging the bandlimitedness of \mathbf{y} , to eliminate the frequencies not present in \mathbf{y} . Then, we use the \mathbf{H} designed in the first step and the specific value of \mathbf{y} to design the signal \mathbf{x} . The challenge is that \mathbf{x} cannot be chosen freely but is rather the output of a *seeding phase* where only a few seeding nodes inject values. The seeding phase requires more elaboration than its counterpart for time-varying signals, not

Work in this paper is supported by Spanish MINECO grant No TEC2013-41604-R and USA NSF CCF-1217963. S. Segarra and A. Ribeiro are with the Dept. of Electrical and Systems Eng., Univ. of Pennsylvania. A. G. Marques is with the Dept. of Signal Theory and Comms., King Juan Carlos Univ. G. Leus is with the Dept. of Electrical Eng., Mathematics and Computer Science, Delft Univ. of Technology. Emails: ssegarra@seas.upenn.edu, antonio.garcia.marques@urjc.es, g.j.t.leus@tudelft.nl, and aribeiro@seas.upenn.edu. Parts of this paper will be presented at EUSIPCO 2015 [1] and GLOBALSIP 2015 [2].

only because graph signals are less regular, but also because it will be shown that in general the seeding values cannot coincide with those of the signal to recover. For a rigorous problem definition see Section II-B and Figure 1 in Section III. Since graph filters act on graph signals through the successive application of local operators, the output of a graph filter can be viewed as the outcome of a diffusion or *percolation* process. Applications include the generation of an opinion profile in a social network [23] by influencing a few agents (Section VII-B) and the synthesis of brain signals [24] by exciting a few neural regions (Section VII-C). Other potential applications for signal reconstruction via local interactions include molecular communications in nanonetworks [25], [26], wireless sensor networks [27], and control of epidemics [28].

The setting considered here differs from the reconstruction of graph signals after sampling [7]–[9]. In the latter, the signal to be recovered is unknown but has some parsimonious (bandlimited) representation, and the goal is to estimate the unknown signal by observing a few values of it. The interpolation operator to obtain the estimated signal from the samples is not constrained to any particular class and, in general, cannot be applied via local interactions. By contrast, in our setting, the signal to reconstruct is given – either pre-specified or estimated via the aforementioned sampling procedure – and we want to induce this signal from just a few seeding values, which in general are *not* samples of the signal to reconstruct. However, we restrict the interpolation operator to the class of graph filters, since we want the reconstruction process to be implemented locally. Notice that this dichotomy – our approach and the reconstruction after sampling – is a feature that arises in graph signal processing but is not present for classical time-varying signals. In classical signal processing, the optimal interpolator of a uniformly sampled signal is a (low-pass) filter, hence, the sampled values coincide with the seeding values and both mentioned settings collapse into one. See Section IV-B for related details.

With these considerations in mind, the paper proposes three different reconstruction schemes, each of them associated with a different seeding phase. In Section III, the seeding phase consists of a unique seeding signal with several nonzero values, which coincides with the intermediate signal \mathbf{x} . By contrast, in Section IV the seeding phase consists of several seeding signals injected by a single seeding node. At each instant, the signal is percolated (diffused) within *one-hop* neighborhoods. The support of \mathbf{x} depends on the duration of the seeding phase and the connectivity of the seeding node. Finally, in Section V we consider a more general scheme which merges the two earlier approaches. In this scheme, the seeding phase consists of several time instants and, in each of them, multiple nodes are allowed to inject a signal. The schemes will be referred to as multiple node-single time (MN-ST), single node-multiple time (SN-MT) and multiple node-multiple time (MN-MT) seeding, respectively. For the three of them, we state conditions on the underlying graph and the seeding nodes to guarantee perfect reconstruction of any bandlimited signal. We also show that, in general, if the interpolator takes the form of a graph filter, the seeding values cannot coincide with those of the signal to interpolate. Furthermore, we discuss how additional seeding values can be used to reduce the complexity of the graph

filter needed for perfect recovery and draw connections with classical interpolation of time-varying signals. In Section VI we study the reconstruction performance in imperfect settings, either because the seeding values are insufficient in number or corrupted by noise. Numerical experiments in Section VII illustrate the reconstruction of signals in noiseless and noisy scenarios using both synthetic and real-world graphs.¹

II. BANDLIMITED GRAPH SIGNALS AND GRAPH FILTERS

Let \mathcal{G} denote a directed graph with a set of N nodes or vertices \mathcal{N} and a set of links \mathcal{E} , such that if node i is connected to j , then $(i, j) \in \mathcal{E}$. The (incoming) neighborhood of i is defined as the set of nodes $\mathcal{N}_i = \{j \mid (j, i) \in \mathcal{E}\}$ connected to i . For any given graph we define the adjacency matrix \mathbf{A} as a sparse $N \times N$ matrix with nonzero elements A_{ji} if and only if $(i, j) \in \mathcal{E}$. The value of A_{ji} captures the strength of the connection from i to j . The focus of this paper is not on analyzing \mathcal{G} , but a graph signal defined on the set of nodes \mathcal{N} . Formally, such a signal can be represented as a vector $\mathbf{x} = [x_1, \dots, x_N]^T \in \mathbb{R}^N$ where the i -th component represents the value of the signal at node i or, alternatively, as a function $f: \mathcal{N} \rightarrow \mathbb{R}$, defined on the vertices of the graph.

The graph \mathcal{G} is endowed with a *graph-shift operator* \mathbf{S} [4], [29]. The shift \mathbf{S} is a $N \times N$ matrix whose entry S_{ji} can be nonzero only if $i = j$ or if $(i, j) \in \mathcal{E}$. The sparsity pattern of the matrix \mathbf{S} captures the local structure of \mathcal{G} , but we make no specific assumptions on the values of the nonzero entries of \mathbf{S} . Choices for \mathbf{S} are the adjacency matrix of the graph [4], [29], its Laplacian [3], and their respective generalizations [30]. The intuition behind \mathbf{S} is to represent a linear transformation that can be computed locally at the nodes of the graph. More rigorously, if \mathbf{y} is defined as $\mathbf{y} = \mathbf{S}\mathbf{x}$, then node i can compute y_i provided that it has access to the value of x_j at $j \in \mathcal{N}_i$. We assume henceforth that \mathbf{S} is diagonalizable, so that there exists a $N \times N$ matrix \mathbf{V} and a $N \times N$ diagonal matrix $\mathbf{\Lambda}$ that can be used to decompose \mathbf{S} as $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$. In particular, \mathbf{S} is diagonalizable when it is normal, i.e., it satisfies $\mathbf{S}\mathbf{S}^H = \mathbf{S}^H\mathbf{S}$ where \mathbf{S}^H denotes the conjugate transpose of \mathbf{S} . In that case, we have that \mathbf{V} is unitary, which implies $\mathbf{V}^{-1} = \mathbf{V}^H$, and leads to the decomposition $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^H$.

We are interested in cases where the graph-shift operator \mathbf{S} plays a role in explaining the graph signal \mathbf{x} . More specifically, cases where \mathbf{x} can be expressed as a linear combination of a *subset* of the columns of $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_N]$, or, equivalently, where the vector $\hat{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}$ is sparse [9]. In this context, vectors \mathbf{v}_i are interpreted as the graph frequency basis, \hat{x}_i as the corresponding signal frequency coefficients, and \mathbf{x} as a bandlimited graph signal. We assume that the set of active frequencies are known and, without loss of generality, that those are the first K ones associated with the eigenvalues of largest magnitude [9], [31]. Under this assumption, if we denote by $\hat{\mathbf{x}}_K := [\hat{x}_1, \dots, \hat{x}_K]^T$ a $K \times 1$ vector collecting the

¹**Notation:** \mathbf{e}_i is the i th $N \times 1$ canonical basis vector (all entries of \mathbf{e}_i are zero except the i th one, which is one); $\mathbf{E}_K := [\mathbf{e}_1, \dots, \mathbf{e}_K]$ is a tall matrix collecting the K first canonical basis vectors while $\mathbf{E}_K := [\mathbf{e}_{K+1}, \dots, \mathbf{e}_N]$ collects the last $N - K$ canonical basis vectors; $\mathbf{0}$ and $\mathbf{1}$ are, respectively, the all-zero and all-one matrices (when not clear from the context, a subscript indicating the dimensions will be used).

coefficients corresponding to those frequencies, it holds that \mathbf{x} is a K -bandlimited signal if

$$\hat{\mathbf{x}} = [\hat{\mathbf{x}}_K^T, 0, \dots, 0]^T, \quad \mathbf{x} = \mathbf{V}\hat{\mathbf{x}} := \mathbf{V}_K \hat{\mathbf{x}}_K, \quad (1)$$

where we have defined the tall matrix $\mathbf{V}_K := [\mathbf{v}_1, \dots, \mathbf{v}_K]$ containing the first K eigenvectors of the shift operator \mathbf{S} .

A. Graph filters

Graph filters $\mathbf{H} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ are linear graph-signal operators of the form $\mathbf{H} := \sum_{l=0}^{L-1} h_l \mathbf{S}^l$; i.e., polynomials (of degree $L-1$) of the graph-shift operator [4]. A particularity of graph filters is that they can be implemented locally, e.g., with $L-1$ exchanges of information among neighbors. This is true because the application of \mathbf{S} on a signal \mathbf{x} can be computed through local interactions.

The graph filter \mathbf{H} can also be written as $\mathbf{H} = \mathbf{V}(\sum_{l=0}^{L-1} h_l \Lambda^l) \mathbf{V}^{-1}$. The diagonal matrix $\hat{\mathbf{H}} := \sum_{l=0}^{L-1} h_l \Lambda^l$ can then be viewed as the frequency response of \mathbf{H} and it can be alternatively written as $\hat{\mathbf{H}} = \text{diag}(\hat{\mathbf{h}})$, where vector $\hat{\mathbf{h}}$ is a vector that contains the frequency response of the filter. Let λ_i denote the i -th eigenvalue of \mathbf{S} and define the $N \times L$ Vandermonde matrix

$$\Psi := \begin{pmatrix} 1 & \lambda_1 & \dots & \lambda_1^{L-1} \\ \vdots & \vdots & & \vdots \\ 1 & \lambda_N & \dots & \lambda_N^{L-1} \end{pmatrix}. \quad (2)$$

Upon defining the vector containing the coefficients of the filter as $\mathbf{h} := [h_0, \dots, h_{L-1}]^T$, it holds that $\hat{\mathbf{h}} = \Psi \mathbf{h}$ and therefore

$$\mathbf{H} = \sum_{l=0}^{L-1} h_l \mathbf{S}^l = \mathbf{V} \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} = \mathbf{V} \text{diag}(\hat{\mathbf{h}}) \mathbf{V}^{-1}. \quad (3)$$

This implies that if \mathbf{y} is defined as $\mathbf{y} = \mathbf{H}\mathbf{x}$, its frequency representation $\hat{\mathbf{y}}$ satisfies

$$\hat{\mathbf{y}} = \text{diag}(\Psi \mathbf{h}) \hat{\mathbf{x}}. \quad (4)$$

Within this context, a *low-pass* graph filter of bandwidth K is one where the frequency response $\hat{\mathbf{h}} := \Psi \mathbf{h}$ is given by

$$\hat{\mathbf{h}} = [\hat{\mathbf{h}}_K^T, 0, \dots, 0]^T, \quad (5)$$

where $\hat{\mathbf{h}}_K$ contains the frequency response for the first K frequencies. Notice that when the low-pass filter in (5) is applied to an arbitrary signal \mathbf{x} , the output signal is K -bandlimited as described in (1). An alternative expression to define a graph filter is [32]

$$\mathbf{H} = a_0 \prod_{l=1}^{L-1} (\mathbf{S} - a_l \mathbf{I}), \quad (6)$$

which also gives rise to a polynomial on \mathbf{S} of degree $L-1$. A specific advantage of the representation in (6) is that it provides a straightforward way to design low-pass filters via successive annihilation of graph frequencies. In particular, if we fix $a_l = \lambda_k$ for some eigenvalue λ_k of \mathbf{S} then the filter \mathbf{H} will eliminate the frequency basis \mathbf{v}_k , i.e., the eigenvector associated with λ_k . For future reference, we denote by D the number of distinct eigenvalues in $\{\lambda_k\}_{k=K+1}^N$.

Remark 1 (Discrete-time signals) To establish connections with classical time-varying signals, we define the directed cycle graph \mathcal{G}_{dc} , with node set $\mathcal{N} = \{1, 2, \dots, N\}$ and edge set $\mathcal{E}_{dc} = \{(i, \text{mod}_N(i) + 1)\}_{i=1}^N$, where $\text{mod}_N(i)$ denotes the remainder obtained after dividing i by N . Its adjacency and Laplacian matrices are denoted, respectively, as \mathbf{A}_{dc} and $\mathbf{L}_{dc} := \mathbf{I} - \mathbf{A}_{dc}$. Discrete-time periodic signals can be thought as graph signals on the directed cycle \mathcal{G}_{dc} . Setting the shift operator either to $\mathbf{S} = \mathbf{A}_{dc}$ or $\mathbf{S} = \mathbf{L}_{dc}$ gives rise to the Fourier basis \mathbf{F} . More formally, the right eigenvectors of \mathbf{S} satisfy $\mathbf{v} = \mathbf{F}$, with $F_{ij} := \exp(+j2\pi(i-1)(j-1)/N)/\sqrt{N}$ where $j := \sqrt{-1}$. Selecting $\mathbf{S} = \mathbf{A}_{dc}$ has the additional advantage of satisfying $\Lambda_{ii} = \exp(-j2\pi(i-1)/N)$, i.e., the eigenvalues of the shift operator correspond to the classical discrete frequencies. Interpretations for the eigenvalues of \mathbf{L}_{dc} also exist [3]. The frequency representation $\hat{\mathbf{x}}$ of a graph signal \mathbf{x} is given by $\hat{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}$ whereas the frequency response of a filter with coefficients \mathbf{h} is given by $\hat{\mathbf{h}} = \Psi \mathbf{h}$. For general graphs, matrices \mathbf{V}^{-1} and Ψ need not be related. However, for the case of \mathcal{G}_{dc} , if $\mathbf{S}_{dc} = \mathbf{A}_{dc}$, then $\Psi = \sqrt{N}\mathbf{F}^H$ and $\mathbf{V}^{-1} = \mathbf{F}^H$. Thus, the Fourier transforms for signals and filter coefficients are equivalent up to a constant for time-varying signals but this is not true for general graph signals.

B. Signal reconstruction using graph filters

Our objective is to reconstruct a given K -bandlimited signal \mathbf{y} by applying a graph filter \mathbf{H} to a signal \mathbf{x} , where \mathbf{x} is the result of a seeding procedure. More specifically, the reconstruction scheme proceeds in two phases:

- *Seeding phase.* The input to this phase is a set of τ graph signals $\{\mathbf{s}^{(t)}\}_{t=0}^{\tau-1}$, denominated seeding signals. These signals percolate through the graph following the dynamics

$$\mathbf{x}^{(t)} = \mathbf{S}\mathbf{x}^{(t-1)} + \mathbf{s}^{(t)}, \quad \mathbf{x}^{(-1)} = \mathbf{0}. \quad (7)$$

The output of this phase is set as $\mathbf{x} := \mathbf{x}^{(\tau-1)}$.

- *Filtering phase.* The graph signal \mathbf{x} is used as input to a low-pass graph filter \mathbf{H} , generating the output $\mathbf{z} := \mathbf{H}\mathbf{x}$.

The purpose of the seeding phase, which has duration τ , is to inject into the graph the information needed to interpolate the signal \mathbf{y} . The filtering phase further propagates the information available from the seeding phase while annihilating the frequencies with indices $k > K$ that are present in \mathbf{x} but not in \mathbf{y} . This phase has duration $L-1$, which is the order of the filter \mathbf{H} .

The goal of this paper is to design $\{\mathbf{s}^{(t)}\}_{t=0}^{\tau-1}$ and \mathbf{H} such that $\mathbf{z} = \mathbf{y}$. We assume that \mathbf{S} is given, thus, the design of \mathbf{H} consists in determining the filter coefficients \mathbf{h} . In Sections III, IV, and V we present this design for three different seeding schemes, where we impose additional restrictions on the structure and the number of seeding signals.

Remark 2 (Network controllability) The dynamics of the seeding phase in (7) is reminiscent of a control setting where \mathbf{S} represents a discrete time, linear, and time-invariant network dynamics and $\mathbf{s}^{(t)}$ stands for the control input at time t . Hence, inducing a particular graph signal is related to the problem

of controllability of complex networks, which has attracted considerable attention in the past years. The controllability focus has recently shifted from qualitative claims [17] to quantitative analysis that takes into account the energy needed to steer the network to a desired state [18]–[21]. Though related to network controllability, the framework proposed here differs in a few points from the existing works in the area, as detailed next. Firstly, although variants exist, the term controllability usually refers to the ability to drive the network to an arbitrary state whereas we are interested in driving the network to a particular class of (bandlimited) graph signals. Secondly, control approaches are typically concerned with the number of control nodes as well as with the energy needed to drive the network to the desired state based on the controllability Gramian – measured in terms of, e.g., its smallest eigenvalue [18], [19] or its trace [20] –, whereas here the concern is with the number of seeding *values* needed to be injected at the seeding (control) nodes. Equivalently, here we are interested in the ℓ_0 norm of the controller needed. Lastly, even though the seeding phase in (7) is similar to a classical control setting, the filtering phase is not. During this phase, no additional seeding values are injected and one only needs to slightly control the dynamics \mathbf{S} at iteration l via the scalar parameter a_l [cf. (6)]. Consequently, our framework can be seen as one which combines external control input during the seeding phase and internal control of the dynamics during the filtering phase.

Remark 3 (Classical recovery) In classical discrete-time signal processing, recovery of bandlimited signals is a two-step process. Firstly, a sparse regular signal whose nonzero values coincide with those of the signal to recover is generated. Secondly, the (zero) values not specified in the sparse signal are extrapolated from the nonzero ones using a low-pass filter. Our approach in this paper is to preserve this two-step methodology and use it to recover bandlimited graph signals. This provides a way to regenerate a desired signal in a graph – either estimated from samples or otherwise – by acting on a subset of (seeding) nodes. As it will be shown in the ensuing sections, for signals defined on a general graph, the nonzero values of the sparse signal in the first step will not coincide with those of the signal to recover. This deviates from the classical concept of interpolation, which assumes that the nonzero values are the same than those of the signal to reconstruct. The practical advantage of studying recovery schemes that use graph filters is twofold. First, they can be implemented distributedly, using only local exchanges among neighbors. Second, since *graph filters* can be used to *model* diffusion processes (e.g. the spread of an opinion in a social network), our results can be used to reconstruct signals in network applications that implement linear diffusion dynamics.

III. MULTIPLE NODE - SINGLE TIME SEEDING

In multiple node - single time (MN-ST) seeding we consider the particular case where there is only $\tau = 1$ seeding signal \mathbf{s} so that $\mathbf{x} = \mathbf{s}$ [cf. (7)]. Denoting by P the amount of nonzero values in \mathbf{s} , we interpret MN-ST seeding as having P seeding nodes that inject a *single* value, while the remaining $N - P$

nodes keep silent; see left and center panels in Figure 1. Define the signal injected by node i as s_i and assume, without loss of generality, that the seeding nodes are the P first ones. We therefore define the $P \times 1$ and $N \times 1$ seeding vectors as

$$\mathbf{s}_P = [s_1, \dots, s_P]^T, \quad (8)$$

$$\mathbf{s} = [s_1, \dots, s_P, 0, \dots, 0]^T. \quad (9)$$

Then, given a bandlimited signal $\mathbf{y} = \mathbf{V}_K \hat{\mathbf{y}}_K$ [cf. (1)], our goal is to design \mathbf{H} and \mathbf{s} such that

$$\mathbf{y} = \mathbf{H}\mathbf{s}, \quad (10)$$

where \mathbf{H} has the *particular structure of a graph filter* (cf. Section II-A). Exploiting the fact that \mathbf{y} is bandlimited, it is reasonable to write (10) in the frequency domain. To do this, both sides of (10) are left multiplied by \mathbf{V}^{-1} , which yields

$$\hat{\mathbf{y}} = \mathbf{V}^{-1} \mathbf{H}\mathbf{s} = \mathbf{V}^{-1} \mathbf{V} \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{s} = \text{diag}(\Psi \mathbf{h}) \hat{\mathbf{s}}, \quad (11)$$

where we used (3) for the second equality. Utilizing the fact that the seeding signal \mathbf{s} is sparse [cf. (9)] we may write its frequency representation as

$$\hat{\mathbf{s}} = \mathbf{V}^{-1} \mathbf{s} = \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P, \quad (12)$$

where, we recall, $\mathbf{E}_P := [\mathbf{e}_1, \dots, \mathbf{e}_P]$ is a tall matrix collecting the P first canonical basis vectors of size $N \times 1$. By substituting (12) into (11), our goal of designing \mathbf{H} and \mathbf{s} such that $\mathbf{y} = \mathbf{H}\mathbf{s}$ can be reformulated as designing \mathbf{h} and \mathbf{s}_P such that

$$\hat{\mathbf{y}} = \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P, \quad (13)$$

which is a bilinear system of N equations and $L+P$ variables. Leveraging the sparsity of $\hat{\mathbf{y}}$ [cf. (1)], the system of N equations in (13) can be split into two

$$\hat{\mathbf{y}}_K = \mathbf{E}_K^T \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P, \quad (14)$$

$$\mathbf{0}_{N-K} = \bar{\mathbf{E}}_K^T \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P, \quad (15)$$

where, we recall, $\bar{\mathbf{E}}_K := [\mathbf{e}_{K+1}, \dots, \mathbf{e}_N]$ collects the last $N - K$ canonical basis vectors and $\mathbf{0}_{N-K}$ denotes the $(N - K) \times 1$ vector of all zeros. Note that the conditions in (15) are the same for any K -bandlimited signal. On the other hand, the conditions in (14) depend on the specific signal to be interpolated. A natural approach is to use the filter coefficients \mathbf{h} – which are related to the global behavior of the graph – to guarantee that (15) holds, while using the seeding signal \mathbf{s}_P to satisfy (14) and, hence, to guarantee that the output of the interpolation is \mathbf{y} . In this way, the filter coefficients \mathbf{h} to be designed do not depend on the particular signal to reconstruct.

The conditions under which the mentioned approach is guaranteed to find a feasible solution are given in the form of two propositions. Ensuing discussions describe the actual procedure to interpolate the signal.

Proposition 1 *If $L > D$ (cf. Section II-A), there exist infinitely many nonzero $L \times 1$ vectors \mathbf{h}^* such that, after setting $\mathbf{h} = \mathbf{h}^*$, (15) is satisfied for any \mathbf{V}^{-1} and \mathbf{s}_P .*

Proof: Since (15) has to hold for any seeding signal \mathbf{s}_P , we need $\bar{\mathbf{E}}_K^T \Psi \mathbf{h} = \mathbf{0}$. This requires \mathbf{h} to belong to the kernel of the $(N - K) \times L$ matrix $\bar{\mathbf{E}}_K^T \Psi$. Since $\bar{\mathbf{E}}_K^T \Psi$ is a Vandermonde matrix, its number of linearly independent rows is equal to the

number of distinct eigenvalues in $\{\lambda_k\}_{k=K+1}^N$, which is D . Thus, the existence of a solution $\mathbf{h}^* \neq \mathbf{0}$ requires $L > D$. ■

For $L > D$, i.e., when the number of filter coefficients exceeds that of distinct eigenvalues in the inactive frequencies, $\bar{\mathbf{E}}_K^T \Psi$ is rank deficient and the dimension of its kernel space is $L - D$. Hence, setting \mathbf{h} to any nonzero element of the kernel space will satisfy (15). In what follows, we will assume that $L = D + 1$ and set the coefficients \mathbf{h}^* that solve (15) as the unit vector spanning the unidimensional kernel space of $\bar{\mathbf{E}}_K^T \Psi$. For the case where all the eigenvalues $\{\lambda_k\}_{k=K+1}^N$ are distinct, this implies that $L = N - K + 1$.

Once the coefficients of the filter are designed, the next step is to find the optimum seeding signal. With $\hat{\mathbf{h}}_K^* := \bar{\mathbf{E}}_K^T \Psi \mathbf{h}^*$ denoting the frequency response of the low-pass filter in the active frequencies, substituting $\mathbf{h} = \mathbf{h}^*$ into (14) yields

$$\hat{\mathbf{y}}_K = \text{diag}(\hat{\mathbf{h}}_K^*) \mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P. \quad (16)$$

For which the following result holds.

Proposition 2 *The system of K equations in (16) is guaranteed to have a solution with respect to \mathbf{s}_P if the following two conditions hold:*

- i) $\lambda_{k_1} \neq \lambda_{k_2}$ for all $(\lambda_{k_1}, \lambda_{k_2})$ such that $k_1 \leq K$ and $k_2 > K$,
- ii) $\text{rank}(\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P) \geq K$.

Proof : Condition i) is required to guarantee that all the elements of $\hat{\mathbf{h}}_K^*$ are nonzero. We prove this by contradiction. Recall that the following facts hold true (cf. Proposition 1): a) $\bar{\mathbf{E}}_K^T \Psi \mathbf{h}^* = \mathbf{0}_{N-K}$; b) $\mathbf{h}^* \neq \mathbf{0}_L$ and c) rank of $\bar{\mathbf{E}}_K^T \Psi$ is $L - 1$. Assume, without loss of generality, that the element of $\hat{\mathbf{h}}_K^*$ that is zero is the K -th one. Then, we can use a) to write $\bar{\mathbf{E}}_{K-1}^T \Psi \mathbf{h}^* = \mathbf{0}_{N-K+1}$. Condition i) and fact c) guarantee that $\bar{\mathbf{E}}_{K-1}^T \Psi$ has rank L ; then, satisfying $\bar{\mathbf{E}}_{K-1}^T \Psi \mathbf{h}^* = \mathbf{0}_{N-K+1}$ requires $\mathbf{h}^* = \mathbf{0}_L$, which contradicts b). Hence, all the elements of $\hat{\mathbf{h}}_K^*$ are nonzero. This guarantees that $\text{diag}(\hat{\mathbf{h}}_K^*)$ is invertible, so that (16) can be written as

$$\text{diag}(\hat{\mathbf{h}}_K^*)^{-1} \hat{\mathbf{y}}_K = (\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P) \mathbf{s}_P, \quad (17)$$

where $\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P$ is a $K \times P$ submatrix of \mathbf{V}^{-1} . To guarantee that the system of equations in (17) has at least one solution, we need condition ii) to hold. ■

Different from the time domain, where all the eigenvalues of \mathbf{S} (frequencies) are distinct, in the more general graph domain there can be graph topologies that give rise to \mathbf{S} with repeated eigenvalues. Condition i) is required because a graph filter \mathbf{H} always produces the same frequency response if the corresponding eigenvalues are the same. Therefore, it is not possible for \mathbf{H} to eliminate one of the frequencies without eliminating the other. An alternative to bypass this problem is discussed in Section III-A. Condition ii) requires the rank of the $K \times P$ matrix $\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P$ being at least K . At the very least, this requires P , the number of seeding nodes, to be equal to K , the number of frequencies present in \mathbf{y} . For the particular case of $P = K$, if the conditions in the proposition are satisfied, the \mathbf{s}_P that recovers \mathbf{y} is

$$\mathbf{s}_P = (\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_P)^{-1} \text{diag}(\hat{\mathbf{h}}_K^*)^{-1} \hat{\mathbf{y}}_K. \quad (18)$$

However, if condition ii) is not satisfied, there may be cases where setting $P = K$ fails. To see why this is true, notice that $[\mathbf{V}^{-1}]_{k,p}$ can be viewed as how strongly node p expresses frequency k . Suppose for example that there exists a k such that $[\mathbf{V}^{-1}]_{k,p} = 0$ for all nodes $p = 1, \dots, P$, then it is not possible to reconstruct a signal \mathbf{y} with $\hat{\mathbf{y}}_k \neq 0$ using that set of nodes. This problem is also present when sampling graph signals by observing the value of the signal in a subset of nodes [7]. Recall that from a control perspective, \mathbf{s}_P in (18) can be interpreted as a controller that acts in a single time instant. The control input \mathbf{s}_P is designed such that, after the application of the filtering phase, the network is steered to the desired state represented by the bandlimited signal \mathbf{y} . Note finally that Proposition 2 states conditions for *perfect* reconstruction in a noiseless setting. In noisy scenarios, the specific set of nodes selected to inject the seeding signal has an impact on the reconstruction error. This is analyzed in Section VI.

Once the mathematical description of the MN-ST has been presented, it is instructive to revisit the differences between our approach and that of sampling and subsequent reconstruction of bandlimited graph signal. More specifically, in the context of sampling of graph signals, we define \mathbf{C} as a binary selection matrix (K rows out of the $N \times N$ identity matrix) that identifies the set of K nodes where the signal is observed or sampled. The sampled signal is then $\bar{\mathbf{x}} = \mathbf{C}\mathbf{x}$, which can be used to recover the K -bandlimited signal \mathbf{x} as $\mathbf{x} = \mathbf{V}_K(\mathbf{C}\mathbf{V}_K)^{-1}\bar{\mathbf{x}}$, assuming that $(\mathbf{C}\mathbf{V}_K)$ is full rank [7], [8]. The previous expression reveals that interpolation amounts to multiplying $\bar{\mathbf{x}}$ by the $N \times K$ matrix $\mathbf{V}_K(\mathbf{C}\mathbf{V}_K)^{-1}$, which does not correspond to a graph filter; hence, it cannot be implemented distributedly using local percolation dynamics, as we require in our context. Nevertheless, the sampling schemes studied in [7]–[10] can be used as a preprocessing step for our algorithms in order to estimate the desired signal \mathbf{y} from a few observations. Formally, if one wants $\mathbf{y} = \mathbf{x}$, then the input to our algorithm can be found as $\hat{\mathbf{y}}_K = (\mathbf{C}\mathbf{V}_K)^{-1}\bar{\mathbf{x}}$ in the frequency domain or $\mathbf{y} = \mathbf{V}_K(\mathbf{C}\mathbf{V}_K)^{-1}\bar{\mathbf{x}}$ in the graph domain.

Remark 4 (Computational complexity) Given \mathbf{S} and \mathbf{y} , we have outlined a procedure to compute the optimal graph filter \mathbf{h}^* and seeding signal \mathbf{s}_P . To do this, first we need to have access to the spectrum of \mathbf{S} to obtain \mathbf{V}^{-1} and Ψ , which requires, in general, in the order of $O(N^3)$ computations. Designing the filter response \mathbf{h}^* amounts to finding a vector in the null space of $\bar{\mathbf{E}}_K^T \Psi$, which can be done, e.g., via a QR decomposition [33] and requires $O((N - K)^3)$ computations. Finally, when $P = K$ obtaining the seeding values in \mathbf{s}_P using as input the target signal \mathbf{y} [cf. (18)] requires either $O(K^3)$ operations, if the frequency coefficients $\hat{\mathbf{y}}_K$ are known, or $O(K^3 + NK)$, if only the signal \mathbf{y} is available. Note however that the computations to find the spectrum of \mathbf{S} , the filter coefficients \mathbf{h}^* and the mapping from \mathbf{y} to \mathbf{s}_P depend only on \mathbf{S} and not on \mathbf{y} . As a result, in scenarios where we want to induce multiple signals \mathbf{y} on the same graph, most of the computations can be reused and the cost of computing a new set of seeding signals \mathbf{s}_p is reduced to $O(K^2)$ or $O(NK)$ depending on whether we are given $\hat{\mathbf{y}}_K$ or \mathbf{y} as input [cf. (18)].

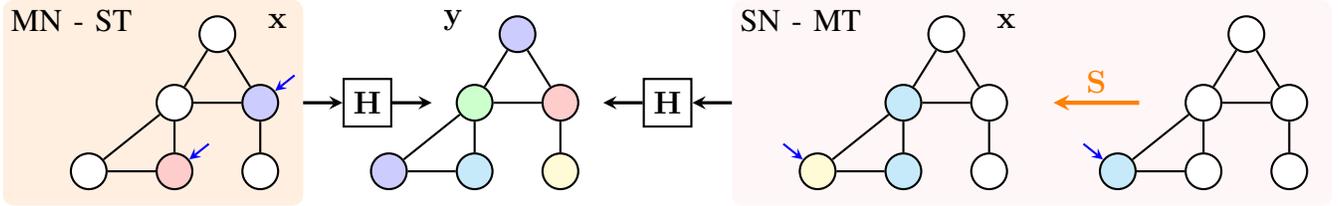


Fig. 1: Two different schemes to reconstruct signal y . On the left, MN-ST seeding injects $P = 2$ values simultaneously (blue arrows), after which the low-pass filter \mathbf{H} is applied. On the right, SN-MT seeding first injects a single value that percolates to the two neighboring nodes. After a second value injection at the same node, filter \mathbf{H} completes the reconstruction.

A. Filter degree reduction in MN-ST seeding

The MN-ST reconstruction scheme requires a low-pass filter \mathbf{H} of degree D , equal to the number of distinct eigenvalues in the inactive frequencies, which grows with the size of the graph. Since the degree of \mathbf{H} corresponds to the number of local interactions needed to implement the filter, the communication overhead can be a problem for large graphs. In this context, we look for solutions that reduce the degree of \mathbf{H} by increasing the number of seeding nodes P . This can be done by splitting the system of equations in (13) as [cf. (14)-(15)]

$$[\hat{\mathbf{y}}_K^T, \mathbf{0}_{P-K}^T]^T = \mathbf{E}_P^T \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P, \quad (19)$$

$$\mathbf{0}_{N-P} = \bar{\mathbf{E}}_P^T \text{diag}(\Psi \mathbf{h}) \mathbf{V}^{-1} \mathbf{E}_P \mathbf{s}_P. \quad (20)$$

The filter coefficients must be obtained now to annihilate the $N - P$ frequencies in (20) and the P seeding nodes must inject a signal whose spectrum, after being filtered by \mathbf{H} , matches that of the desired signal (19). Notice that (19)-(20) can also be used when (14)-(15) fail due to a violation of condition i) in Proposition 2. More specifically, for every frequency index $k_2 > K$ with the same eigenvalue as a frequency index $k_1 \leq K$ we can induce a zero frequency coefficient in the reconstructed signal via the seeding values [cf. (19)] instead of through the low-pass filter [cf. (20)] and, hence, drop condition i) as a requirement for recovery. Further notice that for (20) to hold for any \mathbf{s}_P , the degree of the filter needs to be at least equal to the number of distinct eigenvalues in $\{\lambda_k\}_{k=P+1}^N$ (cf. Proposition 1). In the extreme case of $P = N$, i.e., when every node has an associated seeding value, the trivial solution $\mathbf{h} = [1, 0, \dots, 0]^T$ (0-order filter) and $\mathbf{s}_P = \mathbf{y}$ satisfies (19)-(20).

B. Relation to classical interpolation

In the classical time domain, sinc (low-pass) interpolation of a bandlimited signal leads to perfect reconstruction. If the sampling is performed at the minimum possible rate, the bandwidth of the low-pass filter has to be exactly the same than that of \mathbf{y} . By contrast, if the signal is oversampled, the bandwidth can be larger. Equivalently, if more samples than the minimum required number are available, then the low-pass filter does not have to cancel all the frequencies not present in \mathbf{y} . The analysis in Section III-A reveals that this is also the case when signals are defined in more general graph domains.

The main differences between the MN-ST reconstruction scheme and classical time interpolation come from the fact that the basis \mathbf{V} of a general graph shift \mathbf{S} is not as structured as the Fourier basis \mathbf{F} . A difference of particular relevance is

that, for general graphs, the seeding values \mathbf{s}_P do not coincide with the values of the desired signal \mathbf{y} . This contrasts with the classical interpolation of uniformly sampled time-varying signals, where \mathbf{s}_P is a subset of the signal \mathbf{y} . In fact, it can be rigorously shown that requiring such a condition for general graphs would lead to an infeasible interpolation. To be concrete, suppose that $P = K$, so that $\mathbf{s} = [\mathbf{s}_K^T, \mathbf{0}^T]^T$ [cf. (9)], and that \mathbf{s}_K is equal to the first entries of \mathbf{y} . We can then leverage the fact that \mathbf{s} and $\hat{\mathbf{y}}$ are sparse to write

$$\mathbf{s}_K = \mathbf{E}_K^T \mathbf{y} = \mathbf{E}_K^T \mathbf{V} \hat{\mathbf{y}} = \mathbf{E}_K^T \mathbf{V} \mathbf{E}_K \hat{\mathbf{y}}_K. \quad (21)$$

Secondly, we write the goal of $\mathbf{y} = \mathbf{H}\mathbf{s}$ into the frequency domain as $\hat{\mathbf{y}} = \text{diag}(\hat{\mathbf{h}}) \mathbf{V}^{-1} \mathbf{s}$ and use again the sparsity of \mathbf{s} and $\hat{\mathbf{y}}$ to write

$$\begin{aligned} \hat{\mathbf{y}}_K &= \mathbf{E}_K^T \text{diag}(\hat{\mathbf{h}}) \mathbf{V}^{-1} \mathbf{E}_K \mathbf{s}_K = \mathbf{E}_K^T \text{diag}(\hat{\mathbf{h}}) \mathbf{E}_K \mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{s}_K \\ &= \text{diag}(\hat{\mathbf{h}}_K) \mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{s}_K, \end{aligned} \quad (22)$$

where $\hat{\mathbf{h}}_K := \mathbf{E}_K^T \hat{\mathbf{h}}$ contains the first K components of $\hat{\mathbf{h}}$. Substituting (21) into (22) yields

$$\hat{\mathbf{y}}_K = \text{diag}(\hat{\mathbf{h}}_K) \mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{E}_K^T \mathbf{V} \mathbf{E}_K \hat{\mathbf{y}}_K. \quad (23)$$

Since (23) must hold for all $\hat{\mathbf{y}}_K$, it can only be satisfied if $\text{diag}(\hat{\mathbf{h}}_K) \mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{E}_K^T \mathbf{V} \mathbf{E}_K = \mathbf{I}$. This requires matrix $(\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{E}_K^T \mathbf{V} \mathbf{E}_K)$ to be diagonal. While this is true when $K = N$, it is not true for a general K . However, in the time domain where $\mathbf{V} = \mathbf{F}$ (see Remark 1), for some cases the multiplication of submatrices of \mathbf{F} is guaranteed to be diagonal. For example, if the K seeding nodes are chosen uniformly (equally) spaced, then $(\mathbf{E}_K^T \mathbf{V}^{-1} \mathbf{E}_K \mathbf{E}_K^T \mathbf{V} \mathbf{E}_K) = K/N \mathbf{I}$. This implies not only that (23) is satisfied, but also that all the entries in $\hat{\mathbf{h}}_K$ must be set to N/K . In other words, the optimal low-pass interpolator after uniform sampling in the time domain has the same response for all the active frequencies, as known from classical signal processing.

IV. SINGLE NODE - MULTIPLE TIME SEEDING

In single node - multiple time (SN-MT) seeding, we consider the particular case where all the $\tau = P$ seeding values are injected at a *single* node; see right and center panels in Figure 1. To be more specific, assume without loss of generality that the first node is the one injecting the seeding values, so that the seeding signal $\mathbf{s}^{(t)}$ at time t is of the form $\mathbf{s}^{(t)} = [s^{(t)}, 0, \dots, 0]^T$. Then, define $\mathbf{s}_P := [s^{(P-1)}, \dots, s^{(0)}]^T$ to be a $P \times 1$ vector grouping the seeding values. We present the relation between the seeding values \mathbf{s}_P and the output of the seeding phase \mathbf{x} in the following lemma.

Lemma 1 *The frequency representation of the intermediate signal \mathbf{x} in SN-MT seeding is given by*

$$\widehat{\mathbf{x}} = \text{diag}(\widehat{\mathbf{e}}_1)\Psi\mathbf{s}_P, \quad (24)$$

where $\widehat{\mathbf{e}}_1 := \mathbf{V}^{-1}\mathbf{e}_1$ is the frequency representation of the first canonical basis vector.

Proof: Since \mathbf{x} is obtained after P injections of seeding values following the dynamics in (7), it holds that

$$\mathbf{x} = \mathbf{x}^{(P-1)} = \sum_{l=0}^{P-1} \mathbf{S}^l \mathbf{s}^{(P-1-l)} = \sum_{l=0}^{P-1} \mathbf{S}^l s^{(P-1-l)} \mathbf{e}_1. \quad (25)$$

Equation (25) relates the signal \mathbf{x} to the successive inputs $s^{(\cdot)}$ of the seeding node and can be interpreted as the application of the graph filter

$$\bar{\mathbf{H}} := \sum_{l=0}^{P-1} s^{(P-1-l)} \mathbf{S}^l \quad (26)$$

of degree $P-1$ to the canonical basis vector \mathbf{e}_1 . Building on this interpretation, we may use (4) to write

$$\widehat{\mathbf{x}} = \text{diag}(\Psi\mathbf{s}_P)\widehat{\mathbf{e}}_1, \quad (27)$$

and, by exploiting the fact that for generic vectors \mathbf{a} and \mathbf{b} it holds that $\text{diag}(\mathbf{a})\mathbf{b} = \text{diag}(\mathbf{b})\mathbf{a}$, the lemma follows. ■

The proof of Lemma 1 exploits the reinterpretation of the seeding phase as the application of a graph filter $\bar{\mathbf{H}}$, whose coefficients are given by the seeding values, to the canonical basis vector \mathbf{e}_1 . Equation (24) reveals how $\widehat{\mathbf{x}}$ depends on the structure of the graph Ψ and the seeding values \mathbf{s}_P , as well as on the particular node chosen to inject the values via $\widehat{\mathbf{e}}_1$, whose elements represent how strongly the node expresses each of the graph frequencies.

The next step is to analyze the output of the filtering phase in the frequency domain $\widehat{\mathbf{z}}$. To do this, recall that \mathbf{h}^* denotes the coefficients of a low-pass filter (cf. Section II-A) that eliminates all frequencies with indices $k > K$. Defining $\widehat{\mathbf{h}}^* := \Psi\mathbf{h}^*$, we may analyze the application of the low-pass filter in the frequency domain as

$$\widehat{\mathbf{z}}_K = \mathbf{E}_K^T \text{diag}(\widehat{\mathbf{h}}^*)\widehat{\mathbf{x}} = \mathbf{E}_K^T \text{diag}(\widehat{\mathbf{h}}^*)\mathbf{E}_K \mathbf{E}_K^T \widehat{\mathbf{x}}. \quad (28)$$

Further recalling that $\widehat{\mathbf{h}}_K^* := \mathbf{E}_K^T \widehat{\mathbf{h}}^*$ and substituting (24) into (28), we obtain [cf. (16)]

$$\widehat{\mathbf{y}}_K = \text{diag}(\widehat{\mathbf{h}}_K^*)\mathbf{E}_K^T \text{diag}(\widehat{\mathbf{e}}_1)\Psi\mathbf{s}_P. \quad (29)$$

Expression (29) relates the frequencies present in \mathbf{y} to the seeding values \mathbf{s}_P . Provided that the amount of seeding values is not smaller than the bandwidth of the signal \mathbf{y} , i.e., $K \leq P$, the following proposition states the conditions under which (29) can be solved with respect to \mathbf{s}_P .

Proposition 3 *Let U_1 be the number of values in $\{\{\widehat{\mathbf{e}}_1\}_k\}_{k=1}^K$ that are zero and let D_1 be the number of repeated values in $\{\lambda_k\}_{k=1}^K$. Then, the system of K equations in (29) is guaranteed to have a solution with respect to \mathbf{s}_P if the following two conditions hold:*

- i) $\lambda_{k_1} \neq \lambda_{k_2}$ for all $(\lambda_{k_1}, \lambda_{k_2})$ such that $k_1 \leq K$ and $k_2 > K$,
- ii) $U_1 = 0$ and $D_1 = 0$.

Proof: If we rewrite (29) as

$$\widehat{\mathbf{y}}_K = \left(\text{diag}(\widehat{\mathbf{h}}_K^*) \right) \left(\mathbf{E}_K^T \text{diag}(\widehat{\mathbf{e}}_1)\mathbf{E}_K \right) \left(\mathbf{E}_K^T \Psi \right) \mathbf{s}_P, \quad (30)$$

then it becomes clear that conditions i) and ii) ensure invertibility of the two square matrices, and full row rank of the rectangular matrix $\mathbf{E}_K^T \Psi$. To be specific, condition i) is required to guarantee that all the entries of vector $\widehat{\mathbf{h}}_K^*$ are nonzero and, hence, matrix $\text{diag}(\widehat{\mathbf{h}}_K^*)$ is invertible (cf. proof of Proposition 2). Condition $U_1 = 0$ in ii) ensures that $\mathbf{E}_K^T \text{diag}(\widehat{\mathbf{e}}_1)\mathbf{E}_K$ is invertible since it is a diagonal matrix with no zero elements in its diagonal. Finally, $D_1 = 0$ guarantees that $\mathbf{E}_K^T \Psi$ has rank K whenever $K \leq P$ since it is a row-wise Vandermonde matrix with no repeated rows. ■

Condition i), which is equivalent to that in Proposition 2, guarantees that the low-pass filter with coefficients \mathbf{h}^* does not eliminate any of the frequencies present in \mathbf{y} . Condition ii) states requirements for recovery on both the seeding node and the global structure of the graph. The seeding node is required to be able to act on every active frequency ($U_1 = 0$), while the graph is required to have every active frequency distinguishable from each other ($D_1 = 0$). Condition ii) ensures that the rank of matrix $\mathbf{E}_K^T \text{diag}(\widehat{\mathbf{e}}_1)\Psi$ is equal to K when $P \geq K$, guaranteeing that (29) can be solved with respect to \mathbf{s}_P . For the particular case of $P = K$, i.e. the amount of seeding values equals the bandwidth, the seeding values can be found as

$$\mathbf{s}_P = \left(\mathbf{E}_K^T \text{diag}(\widehat{\mathbf{e}}_1)\Psi \right)^{-1} \text{diag}(\widehat{\mathbf{h}}_K^*)^{-1} \widehat{\mathbf{y}}_K. \quad (31)$$

When comparing the conditions ii) in Propositions 2 and 3, we observe that for MN-ST seeding we should require a rank condition on a submatrix of \mathbf{V}^{-1} . By contrast, for SN-MT seeding, the Vandermonde structure of Ψ allows reformulating the rank condition in terms of the graph related quantities U_1 and D_1 , providing further insight on specifying the situations when recovery is possible. This dual behavior is also present when sampling graph signals. When following “selection sampling” [7], which is the counterpart of MN-ST interpolation, perfect reconstruction depends on the invertibility of a submatrix of \mathbf{V} , whereas when following an “aggregation sampling” scheme [8], which is the counterpart of SN-MT interpolation, the conditions for perfect reconstruction can be written in terms of specific graph related quantities. From a network controllability viewpoint, \mathbf{s}_P in (31) contains the successive control inputs that have to be injected at the single seeding node such that, after the application of the filtering phase, the desired network state is achieved.

Even though Proposition 3 guarantees perfect recovery under SN-MT seeding in a noiseless case, in noisy scenarios the selection of the seeding node is essential to reduce the reconstruction error. This is analyzed in Section VI under a more general seeding scheme.

A. Filter degree reduction in SN-MT seeding

Mimicking the filter degree reduction technique presented in Section III-A, SN-MT seeding can also implement a lower-degree filter if a higher number of seeding values is injected. To achieve this, we need the additional seeding values to generate a signal whose spectrum is zero for the inactive frequencies that are not eliminated by the filter. More specifically, the seeding values \mathbf{s}_P and the filter coefficients \mathbf{h} have to

satisfy [cf. (19)-(20)]

$$[\widehat{\mathbf{y}}_K^T, \mathbf{0}_{P-K}^T]^T = \text{diag}(\mathbf{E}_P^T \Psi \mathbf{h}) \mathbf{E}_P^T \text{diag}(\widehat{\mathbf{e}}_1) \Psi \mathbf{s}_P, \quad (32)$$

$$\mathbf{0}_{N-P} = \text{diag}(\bar{\mathbf{E}}_P^T \Psi \mathbf{h}) \bar{\mathbf{E}}_P^T \text{diag}(\widehat{\mathbf{e}}_1) \Psi \mathbf{s}_P, \quad (33)$$

where $N - P$ is the number of frequency coefficients eliminated by the low-pass filter \mathbf{h} . As done in Section III-A, \mathbf{h} will be designed to solve (33) for any choice of \mathbf{s}_P , while \mathbf{s}_P will be chosen to solve the P equations in (32). A sufficient degree for \mathbf{h} is presented next.

Proposition 4 *Let U_2 be the number of values in $\{\widehat{\mathbf{e}}_1\}_k\}_{k=K+1}^N$ that are zero and D_2 be the number of repeated values in $\{\lambda_k\}_{k \in \mathcal{K}_U}$, where $\mathcal{K}_U := \{k \mid K < k \leq N \text{ and } [\widehat{\mathbf{e}}_1]_k \neq 0\}$. Then, (33) can be solved with respect to \mathbf{h} for any choice of \mathbf{s}_P provided that $L-1 \geq \max(0, N-P-U_2-D_2)$.*

Proof: Notice that in (33) the filter eliminates the last $N - P$ frequencies, however, since the ordering is arbitrary, any subset of $N - P$ frequencies (not containing the K first ones) can be chosen to be annihilated by \mathbf{h} . Thus, our objective is to show that the proposed degree is enough to nullify a particular choice of $N - P$ frequency coefficients. Define as \mathcal{R} the set of indices corresponding to zero elements in $\widehat{\mathbf{e}}_1$ or repeated rows in Ψ . Since condition *ii*) in Proposition 3 must be satisfied – otherwise, perfect recovery would be infeasible –, the cardinality of \mathcal{R} is $U_2 + D_2$ and every index in \mathcal{R} must be greater than K .

First assume that $U_2 + D_2 < N - P$ and pick the $N - P$ frequencies to be eliminated to include the ones in \mathcal{R} . This is equivalent to picking a frequency ordering such that every index in \mathcal{R} is greater than P . Based on \mathcal{R} , define the selection matrices $\mathbf{E}_{\mathcal{R}} := [\mathbf{e}_{k_1}, \mathbf{e}_{k_2}, \dots, \mathbf{e}_{k_{U_2+D_2}}]$ for all $k_i \in \mathcal{R}$ and $\bar{\mathbf{E}}_{\mathcal{R}} := [\mathbf{e}_{k_1}, \mathbf{e}_{k_2}, \dots, \mathbf{e}_{k_{N-P-U_2-D_2}}]$ for all $k_i \in \{P+1, \dots, N\} \setminus \mathcal{R}$ where \setminus represents the set difference operator. Hence, the system of equations in (33) can be split into two

$$\mathbf{0}_{N-P-U_2-D_2} = \text{diag}(\bar{\mathbf{E}}_{\mathcal{R}}^T \Psi \mathbf{h}) \bar{\mathbf{E}}_{\mathcal{R}}^T \text{diag}(\widehat{\mathbf{e}}_1) \Psi \mathbf{s}_P, \quad (34)$$

$$\mathbf{0}_{U_2+D_2} = \text{diag}(\mathbf{E}_{\mathcal{R}}^T \Psi \mathbf{h}) \mathbf{E}_{\mathcal{R}}^T \text{diag}(\widehat{\mathbf{e}}_1) \Psi \mathbf{s}_P. \quad (35)$$

Condition (34) can be guaranteed for any \mathbf{s}_P if $\mathbf{h} = \mathbf{h}^*$, where \mathbf{h}^* are the coefficients of a low-pass filter of degree $L-1 = N - P - U_2 - D_2$, as stated by the proposition. To complete the proof, we need to show that $\mathbf{h} = \mathbf{h}^*$ also guarantees that (35) holds. To see why this is the case, notice that U_2 rows of $\mathbf{E}_{\mathcal{R}}^T \text{diag}(\widehat{\mathbf{e}}_1) \Psi$ are exactly zero, trivially satisfying (35) for any \mathbf{s}_P . Also, each of the remaining D_2 equations in (35) corresponds to a repeated eigenvalue and, thus, can be obtained by multiplying one of the $N - K - U_2 - D_2$ homogenous equations in (34) and (32) by a *scalar*, guaranteeing that \mathbf{h}^* also solves these D_2 equations.

For the case where $U_2 + D_2 \geq N - P$, we pick a frequency ordering such that every index greater than P is contained in \mathcal{R} . Thus, (33) is implied by the homogenous equations in (32), and no filter (degree 0) is needed. ■

Proposition 4 explicitly states that every additional seeding value decreases the required filter degree. However, in contrast to the situation for MN-ST, this reduction of the filter degree

does not entail a reduction in the number of applications of the graph-shift operator, because it requires the length of the seeding phase to be extended. More interestingly, the additional seeding values can be used as a mean to guarantee perfect reconstruction when condition *i*) in Proposition 3 is not satisfied, as explained in Section III-A for MN-ST seeding.

When $P \geq N - U_2 - D_2$ the seeding phase suffices to recover the signal. This can be of interest in scenarios where the graph-shift operator \mathbf{S} describes an intrinsic graph diffusion dynamic and the design of the filter coefficients is not feasible. It is also of interest if \mathbf{y} is not bandlimited. By rendering the filtering phase unnecessary, SN-MT seeding can be used to induce arbitrary (non-bandlimited) signals provided that the amount of seeding values is large enough. See Section VII for further discussions.

B. Relation to classical interpolation

When $\mathbf{S} = \mathbf{A}_{dc}$, applying the l th power of \mathbf{S} to a signal amounts to shifting the signal l time instants. Consequently, the intermediate signal \mathbf{x} obtained after the seeding phase in SN-MT reconstruction coincides with the seeding signal \mathbf{s} in MN-ST reconstruction, provided that the seeding nodes are chosen adjacent to each other. Moreover, for the extreme case of the number of seeding values P being enough to eliminate the filtering phase, which entails $L - 1 = 0$, Proposition 4 requires setting $P = N$, because both U_2 and D_2 are zero if $\mathbf{S} = \mathbf{A}_{dc}$. The design of the $P = N$ seeding values \mathbf{s}_P that guarantee that $\mathbf{x} = \mathbf{y}$ can be carried out trivially by setting $\mathbf{s}_P = [x_1, \dots, x_N] = \mathbf{y}$.

V. MULTIPLE NODE - MULTIPLE TIME SEEDING

In the more general multiple node - multiple time (MN-MT) seeding scheme, we can have several seeding signals ($\tau > 1$) and we do not assume any structure on $\mathbf{s}^{(t)}$, so that any node may inject a seeding value at any given time. We concatenate the τ seeding signals into the $N\tau \times 1$ vector $\underline{\mathbf{s}}$ defined as $\underline{\mathbf{s}} := \text{vec}([\mathbf{s}^{(\tau-1)}, \mathbf{s}^{(\tau-2)}, \dots, \mathbf{s}^{(0)}]^T)$.

Defining the $N \times N^2$ matrix $\Theta := [\text{diag}(\widehat{\mathbf{e}}_1), \dots, \text{diag}(\widehat{\mathbf{e}}_N)]$, we may relate \mathbf{x} to $\underline{\mathbf{s}}$ as stated in the following lemma.

Lemma 2 *The frequency representation of the intermediate signal \mathbf{x} in MN-MT seeding is given by*

$$\widehat{\mathbf{x}} = \Theta (\mathbf{I} \otimes \Psi) \underline{\mathbf{s}}, \quad (36)$$

where \otimes represents the Kronecker product.

Proof: If we denote by $\mathbf{x} := \mathbf{x}^{(\tau-1)}$ the signal obtained after the seeding phase, it holds that [cf. (7)]

$$\mathbf{x} = \sum_{l=0}^{\tau-1} \mathbf{S}^l \mathbf{s}^{(\tau-1-l)} = \sum_{l=0}^{\tau-1} \mathbf{S}^l \left(\sum_{i=1}^N s_i^{(\tau-1-l)} \mathbf{e}_i \right) = \sum_{i=1}^N \mathbf{H}_i \mathbf{e}_i, \quad (37)$$

where the filter \mathbf{H}_i is given by

$$\mathbf{H}_i = \sum_{l=0}^{\tau-1} s_i^{(\tau-1-l)} \mathbf{S}^l. \quad (38)$$

Writing the input-output relationship of those filters in the frequency domain, we have that [cf. (4)]

$$\widehat{\mathbf{x}} = \sum_{i=1}^N \text{diag}(\Psi \mathbf{s}_i) \widehat{\mathbf{e}}_i = \sum_{i=1}^N \text{diag}(\widehat{\mathbf{e}}_i) \Psi \mathbf{s}_i. \quad (39)$$

Recalling the definitions of Θ and $\underline{\mathbf{s}}$, the sum in (39) can be written in matrix form, giving rise to (36). ■

As was the case for Lemma 1 in SN-MT seeding, Lemma 2 leverages the reinterpretation of the seeding phase as the application of a filter – in this case, N different filters, one per node – to the canonical basis vectors [cf. (37)]. The coefficients of the filter associated with the i -th node are given by the values injected by that i -th node [cf. (38)]. Notice that, as expected, (36) reduces to (24) whenever the seeding values are forced to be zero for every seeding node except for the first one.

To analyze the output of the filtering phase \mathbf{z} , recall that $\widehat{\mathbf{h}}_K^* = \mathbf{E}_K^T \Psi \mathbf{h}^*$ denotes the response of a low-pass filter in the active frequencies. Mimicking the procedure in Section IV, we find that the active frequency coefficients in \mathbf{y} can be written in terms of the seeding values $\underline{\mathbf{s}}$ as

$$\widehat{\mathbf{y}}_K = \widehat{\mathbf{z}}_K = \text{diag}(\widehat{\mathbf{h}}_K^*) \mathbf{E}_K^T \Theta (\mathbf{I} \otimes \Psi) \underline{\mathbf{s}}. \quad (40)$$

The system of equations in (40) is underdetermined, since the K values in $\widehat{\mathbf{y}}_K$ can be reconstructed using the $N\tau$ values in $\underline{\mathbf{s}}$. However, our focus is on the case where only $P \ll N\tau$ seeding values are injected during the seeding phase. To this extent, we introduce the $P \times N\tau$ selection matrix \mathbf{C} whose elements are binary $C_{ij} \in \{0, 1\}$ and satisfy $\sum_j C_{ij} = 1$ and $\sum_i C_{ij} \leq 1$ for all i and j , respectively. Since the matrix has exactly one 1 in every row, \mathbf{C} selects P seeding values among the $N\tau$ node-time pairs. If we denote by $\underline{\mathbf{s}}_P := \mathbf{C} \underline{\mathbf{s}}$ the vector containing these P seeding values, (40) can be rewritten as [cf. (16) and (29)]

$$\widehat{\mathbf{y}}_K = \text{diag}(\widehat{\mathbf{h}}_K^*) \mathbf{E}_K^T \Theta (\mathbf{I} \otimes \Psi) \mathbf{C}^T \underline{\mathbf{s}}_P. \quad (41)$$

The conditions under which (41) can be solved with respect to $\underline{\mathbf{s}}_P$ are given in the form of the following proposition.

Proposition 5 *The system of K equations in (41) is guaranteed to have a solution with respect to $\underline{\mathbf{s}}_P$ if the following two conditions hold:*

- i) $\lambda_{k_1} \neq \lambda_{k_2}$ for all $(\lambda_{k_1}, \lambda_{k_2})$ such that $k_1 \leq K$ and $k_2 > K$,*
- ii) $\text{rank}(\mathbf{E}_K^T \Theta (\mathbf{I} \otimes \Psi) \mathbf{C}^T) \geq K$.*

Proof: Condition *i*) is required to guarantee that $\text{diag}(\widehat{\mathbf{h}}_K^*)$ is invertible (cf. proof of Proposition 2). This allows us to rewrite (41) as

$$\text{diag}(\widehat{\mathbf{h}}_K^*)^{-1} \widehat{\mathbf{y}}_K = \mathbf{E}_K^T \Theta (\mathbf{I} \otimes \Psi) \mathbf{C}^T \underline{\mathbf{s}}_P. \quad (42)$$

To guarantee that the system of equations in (42) has at least one solution, we need condition *ii*) to hold. ■

Condition *i*), also present in Propositions 2 and 3, guarantees that the filtering phase does not annihilate any of the frequencies present in \mathbf{y} . Condition *ii*) requires, at the very least, $P \geq K$, i.e. the number of seeding values cannot be smaller than the bandwidth of the desired signal. However,

there may be cases where setting $P = K$ can fail as stated in the discussion ensuing Proposition 2. Moreover, since MN-MT subsumes the particular cases of MN-ST and SN-MT seeding, the result in Proposition 5 generalizes those in Propositions 2 and 3; see Remark 5.

Mimicking the developments in Sections III-A and IV-A, for the general case of MN-MT seeding, additional seeding values can be used to reduce the degree of the low-pass filter. Indeed, for every extra seeding value the degree of the filter needed decreases by one, reducing the communication cost of the reconstruction scheme. Moreover, these extra seeding values can be used either to obtain perfect reconstruction even when condition *i*) in Proposition 5 is violated, as explained in Section III-A, or even to induce a non-bandlimited signal \mathbf{y} by rendering the filtering phase unnecessary.

The selection matrix \mathbf{C} can be designed so that condition *ii*) in Proposition 5 is satisfied, guaranteeing perfect recovery. Furthermore, for the cases in which perfect reconstruction is infeasible due to, e.g., the presence of noise, the choice of \mathbf{C} can be optimized to achieve robust recovery, as analyzed in the following section.

We want to emphasize that there is no particular seeding scheme – MN-ST, SN-MT, or MN-MT – preferred over others in general but, rather, their relevance depends on the application of interest. E.g., for graphs where we can only act on a few nodes (or just one) and the graph-shift \mathbf{S} encodes a network dynamic naturally applied by the medium [25], [34] then SN-MT or MN-MT with few nodes are reasonable approaches. On the other hand, if we can act on a large number of nodes simultaneously, then MN-ST or MN-MT might be preferable to reduce the duration of the seeding phase.

Remark 5 (Special cases of MN-MT) Denoting by $\mathbf{e}_{i,N\tau}$ the i th canonical basis vector of size $N\tau$, define the matrices $\mathbf{C}_{SN} := [\mathbf{e}_{1,N\tau}, \mathbf{e}_{2,N\tau}, \dots, \mathbf{e}_{P,N\tau}]^T$ and $\mathbf{C}_{ST} := [\mathbf{e}_{1,N\tau}, \mathbf{e}_{\tau+1,N\tau}, \dots, \mathbf{e}_{(P-1)\tau+1,N\tau}]^T$. Then, setting $\mathbf{C} = \mathbf{C}_{SN}$ in (41) amounts to selecting the first P elements of $\underline{\mathbf{s}}$ and (41) reduces to (29), so that SN-MT reconstruction is recovered. Similarly, by setting $\mathbf{C} = \mathbf{C}_{ST}$ one selects the elements of $\underline{\mathbf{s}}$ in positions $1, \tau + 1, \dots, (P - 1)\tau + 1$ and (41) reduces to (16), as in MN-ST reconstruction.

VI. IMPERFECT RECONSTRUCTION

We study two settings where perfect reconstruction is infeasible: insufficient number of seeding values (Section VI-A) and additive noise in the injections (Section VI-B). The analysis is focused on the MN-MT seeding scheme, since the results for MN-ST and SN-MT can be obtained by particularizing the value of the selection matrix \mathbf{C} (cf. Remark 5).

A. Insufficient seeding values

When the number of seeding values P is not enough to achieve perfect reconstruction, the goal is to minimize a pre-specified error metric between the reconstructed signal \mathbf{z} and the original K -bandlimited graph signal \mathbf{y} . This situation can arise in practice when, e.g., we want to induce an opinion profile in a social network where only a few agents can be influenced (Section VII-B) or when trying to induce a

brain signal but we have access to only a limited number of brain regions (Section VII-C). Alternatively, settings where the seeding values are insufficient arise, for example, when trying to reconstruct *approximately* K -bandlimited signals. In these cases, perfect reconstruction would require in the order of N seeding values but low reconstruction errors can still be achieved with $P < N$ seeding values by following the strategies described in this section. Three different design scenarios are considered. In the first one, the seeding values \underline{s}_P are designed assuming that both \mathbf{h} and \mathbf{C} are fixed. The second scenario addresses the joint design of \underline{s}_P and \mathbf{h} . In the last one, the joint design of \underline{s}_P and \mathbf{C} is performed.

1) *Designing the seeding values \underline{s}_P* : Assume that condition *i*) in Proposition 5 holds and recall that \mathbf{h}^* stands for the coefficients of a low-pass filter that eliminates all the frequencies $k > K$. Then, the first K frequency coefficients $\hat{\mathbf{z}}_K$ of the reconstructed signal \mathbf{z} are obtained as [cf. (41)]

$$\hat{\mathbf{z}}_K = \text{diag}(\hat{\mathbf{h}}_K^*) \mathbf{E}_K^T \Theta(\mathbf{I} \otimes \Psi) \mathbf{C}^T \underline{s}_P. \quad (43)$$

Since we assume insufficient seeding values, i.e., $P < K$, obtaining $\hat{\mathbf{y}}_K = \hat{\mathbf{z}}_K$ is in general infeasible. A reasonable approach is to design \underline{s}_P to minimize the energy of the reconstruction error. Defining the matrix

$$\Phi_K := \text{diag}(\hat{\mathbf{h}}_K^*) \mathbf{E}_K^T \Theta(\mathbf{I} \otimes \Psi), \quad (44)$$

the optimal seeding values \underline{s}_P^* can be obtained as

$$\underline{s}_P^* := \arg \min_{\underline{s}_P} \|\mathbf{y} - \mathbf{V}_K \Phi_K \mathbf{C}^T \underline{s}_P\|_2^2, \quad (45)$$

where, we recall, $\mathbf{V}_K := \mathbf{V} \mathbf{E}_K$. The minimization problem in (45) has the well-known closed-form solution [35]

$$\underline{s}_P^* = (\mathbf{C} \Phi_K^H \Phi_K \mathbf{C}^T)^{-1} \mathbf{C} \Phi_K^H \mathbf{V}_K^H \mathbf{y}, \quad (46)$$

where we assume that the *fixed* seeding locations \mathbf{C} lead to a matrix $\Phi_K \mathbf{C}^T$ that has full column rank. With $\epsilon := \mathbf{y} - \mathbf{z}$ denoting the reconstruction error, its energy can be written as

$$\|\epsilon\|_2^2 = \mathbf{y}^H \mathbf{V}_K \left(\mathbf{I} - \Phi_K \mathbf{C}^T (\mathbf{C} \Phi_K^H \Phi_K \mathbf{C}^T)^{-1} \mathbf{C} \Phi_K^H \right) \mathbf{V}_K^H \mathbf{y}. \quad (47)$$

Notice that, since \mathbf{h}^* is given, the reconstruction error is zero for the frequency components $k > K$.

2) *Designing the seeding values \underline{s}_P and the filter coefficients \mathbf{h}* : When perfect reconstruction is infeasible, carrying out a separate optimization of \mathbf{h} and \underline{s}_P , where \mathbf{h} is designed to filter the frequencies not present in \mathbf{y} and \underline{s}_P is designed to match the spectrum of \mathbf{y} in the active frequencies, is not jointly optimal. Minimization of the reconstruction error by jointly designing \underline{s}_P and \mathbf{h} is briefly discussed next. Notice that the N frequency coefficients – as opposed to just the first K coefficients – of the reconstructed signal $\hat{\mathbf{z}}$ are [cf. (41)]

$$\hat{\mathbf{z}} = \text{diag}(\Psi \mathbf{h}) \mathbf{E}_K^T \Theta(\mathbf{I} \otimes \Psi) \mathbf{C}^T \underline{s}_P. \quad (48)$$

Hence, if the objective is to minimize $\|\epsilon\|_2^2$, we have that

$$\begin{aligned} \{\underline{s}_P^*, \mathbf{h}^*\} &:= \underset{\{\underline{s}_P, \mathbf{h}\}}{\text{argmin}} \|\mathbf{y} - \mathbf{V} \hat{\mathbf{z}}\|_2^2 \\ &= \underset{\{\underline{s}_P, \mathbf{h}\}}{\text{argmin}} \|\mathbf{y} - \mathbf{V} \text{diag}(\Psi \mathbf{h}) \mathbf{E}_K^T \Theta(\mathbf{I} \otimes \Psi) \mathbf{C}^T \underline{s}_P\|_2^2, \end{aligned} \quad (49)$$

which is a *bilinear* optimization. Bilinear problems are non-convex, but there is a large amount of works dealing with their analysis and efficient solution [36]–[38].

The formulation in (49) considers that \mathbf{h}^* can be chosen as a function of the signal to reconstruct \mathbf{y} . In applications where this is not feasible, formulating the optimal design requires additional assumptions on \mathbf{y} . If the distribution of \mathbf{y} is known, a two-stage stochastic programming approach can be pursued [39]. In the second stage, \mathbf{h} in (49) is considered given and the optimal \underline{s}_P^* is obtained as the minimizer of $\|\epsilon(\mathbf{h}, \mathbf{y}, \underline{s}_P)\|_2^2$, which is a function of \mathbf{h} and \mathbf{y} [cf. (46)]. In the first stage, the solution of the second stage $\underline{s}_P^*(\mathbf{h}, \mathbf{y})$ and the distribution of \mathbf{y} are leveraged to write the expectation of the reconstruction error in (49) as $\bar{\epsilon}(\mathbf{h}) := \mathbb{E}_{\mathbf{y}} [\|\epsilon(\mathbf{h}, \mathbf{y}, \underline{s}_P^*(\mathbf{h}, \mathbf{y}))\|_2^2]$, which only depends on \mathbf{h} . The optimum \mathbf{h}^* is then the minimizer of the expected error $\bar{\epsilon}(\mathbf{h})$. Notice that this two-stage approach is used in Sections III, IV, and V to find conditions for perfect recovery, where bandlimitedness is the prior knowledge of \mathbf{y} .

3) *Designing the seeding values \underline{s}_P and the seeding locations \mathbf{C}* : Suppose now that one can select the specific nodes and time instants where the injections take place. This amounts to choosing the P entries of \mathbf{C} that are nonzero, which is a combinatorial problem. Although for small networks one could try all possible choices of \mathbf{C} and select the one leading to the smallest reconstruction error, for general networks a more scalable approach is required. To formulate the corresponding optimization problem $\underline{s}_P = \mathbf{C} \underline{s}$ is substituted into (45). After that, the product $\mathbf{C}^T \mathbf{C}$ is rewritten as $\text{diag}(\mathbf{c})$ where \mathbf{c} is a binary selection vector of dimension $N\tau \times 1$. Note that having $c_i = 1$ indicates that at time $t = \text{mod}_\tau(N\tau - i)$ the node $(i + t)/\tau$ injects a seeding value. With this notation, the joint design of \underline{s}_P and \mathbf{c} amounts to solving

$$\begin{aligned} \{\underline{s}^*, \mathbf{c}^*\} &:= \underset{\{\underline{s}, \mathbf{c}\}}{\text{argmin}} \|\mathbf{y} - \mathbf{V}_K \Phi_K \text{diag}(\mathbf{c}) \underline{s}\|_2^2 + \gamma \|\mathbf{c}\|_0 \\ \text{s.t.} \quad &\mathbf{c} \in \{0, 1\}^{N\tau}, \end{aligned} \quad (50)$$

where Φ_K is defined in (44). In (50) each seeding location used is penalized with a constant cost γ . By tuning γ , the desired level of sparsity of \mathbf{c} can be achieved. Problem (50) can be further simplified by setting $\mathbf{d} := \text{diag}(\mathbf{c}) \underline{s}$ and requiring sparsity on \mathbf{d}

$$\mathbf{d}^* := \underset{\mathbf{d}}{\text{argmin}} \|\mathbf{y} - \mathbf{V}_K \Phi_K \mathbf{d}\|_2^2 + \gamma \|\mathbf{d}\|_0. \quad (51)$$

Among other advantages, the formulation in (51) is amenable to relaxations that reduce the computational complexity required to find a solution. A straightforward approach is to relax the problem by replacing the 0-norm with the 1-norm to obtain a convex formulation.

As in problem (49), the design in (50) and its subsequent simplification in (51) assume that the seeding nodes can be chosen as a function of \mathbf{y} . For applications where this is not convenient, a two-stage stochastic programming approach similar to the one described for (49) can also be used in solving (50). Last but not least, although computationally challenging, a joint optimization of \underline{s}_P , \mathbf{h} and \mathbf{C} can be pursued by combining the approaches in Sections VI-A2 and VI-A3.

B. Noise when injecting the seeding values

The conditions for perfect reconstruction stated in Propositions 2, 3 and 5 require the seeding values to be the exact solution of (16), (29), and (41), respectively. However, in real applications, the injected values can be corrupted with additive noise. This noise can be either attenuated or amplified when the signal percolates through the graph via the successive applications of \mathbf{S} . The goal of this section is to quantify the reconstruction error and to discuss seeding selection schemes tailored to these operating conditions. Their performance will be illustrated via numerical simulations in Section VII.

Let us assume that the injected signal is $\underline{s}_P + \mathbf{w}_P$, where \mathbf{w}_P is a $P \times 1$ noise vector with zero mean and covariance \mathbf{R}_w . The active frequencies of the reconstructed signal can then be written as $\hat{\mathbf{z}}_K = \Phi_K \mathbf{C}^T (\underline{s}_P + \mathbf{w}_P)$ [cf. (41) and (44)]. From this, we may obtain the reconstruction error as

$$\epsilon = \mathbf{V}_K (\hat{\mathbf{z}}_K - \hat{\mathbf{y}}_K) = \mathbf{V}_K \Phi_K \mathbf{C}^T \mathbf{w}_P, \quad (52)$$

with covariance matrix

$$\mathbf{R}_\epsilon = \mathbb{E}(\epsilon \epsilon^H) = \mathbf{V}_K \Phi_K \mathbf{C}^T \mathbf{R}_w \mathbf{C} \Phi_K^H \mathbf{V}_K^H. \quad (53)$$

Ideally, \mathbf{C} should be designed to select the seeding nodes and time instants that minimize the reconstruction error, which can be quantified as a function of \mathbf{R}_ϵ . In what follows, we will focus on minimizing the mean squared error (MSE), which is achieved by minimizing $\text{trace}(\mathbf{R}_\epsilon)$. However, similar approaches can be followed to minimize other commonly used error metrics such as $\lambda_{\max}(\mathbf{R}_\epsilon)$ and $\log(\det(\mathbf{R}_\epsilon))$ [40].

To illustrate the design of \mathbf{C} , two particular scenarios are considered. In the first one, we assume i) that the frequency coefficients $\hat{\mathbf{y}}_K$ are zero mean with covariance $\mathbf{R}_{\hat{\mathbf{y}}} = \mathbf{I}$ and ii) that the noise \mathbf{w}_P is zero mean with covariance $\mathbf{R}_w = \sigma^2 \mathbb{E}(\|\underline{s}_P\|_2^2) \mathbf{I}$. Note that assumption ii) is meaningful if the system operates under a constant signal-to-noise ratio (SNR) regime. In the second scenario, we assume the noise is also uncorrelated but its power is independent of that of the seeding signal, so that $\mathbf{R}_w = \sigma^2 \mathbf{I}$.

For convenience, the optimal seeding strategy for the first scenario is presented in the form of a lemma.

Lemma 3 *Suppose that $\hat{\mathbf{y}}_K$ and \mathbf{w}_P are drawn from zero-mean distributions with covariances $\mathbf{R}_{\hat{\mathbf{y}}} = \mathbf{I}$ and $\mathbf{R}_w = \sigma^2 \mathbb{E}(\|\underline{s}_P\|_2^2) \mathbf{I}$, respectively. Then, the selection \mathbf{c}^* that minimizes the MSE of the reconstruction is given by*

$$\mathbf{c}^* := \underset{\mathbf{c}}{\text{argmin}} \text{trace} \left(\left(\Phi_K \text{diag}(\mathbf{c}) \Phi_K^H \right)^{-1} \right) \text{trace} \left(\Phi_K \text{diag}(\mathbf{c}) \Phi_K^H \right) \quad (54)$$

s.t. $\mathbf{c} \in \{0, 1\}^{N\tau}$, $\|\mathbf{c}\|_0 = P$.

Proof : To prove the lemma, we need to show that the minimization of the objective in (54) is equivalent to the minimization of $\text{trace}(\mathbf{R}_\epsilon)$. By substituting $\mathbf{R}_w = \sigma^2 \mathbb{E}(\|\underline{s}_P\|_2^2) \mathbf{I}$ and $\text{diag}(\mathbf{c}) := \mathbf{C}^T \mathbf{C}$ into (53), it follows that

$$\text{trace}(\mathbf{R}_\epsilon) = \sigma^2 \mathbb{E}(\|\underline{s}_P\|_2^2) \text{trace}(\mathbf{V}_K \Phi_K \text{diag}(\mathbf{c}) \Phi_K^H \mathbf{V}_K^H). \quad (55)$$

Since the trace is invariant to cyclic permutations and $\mathbf{V}_K^H \mathbf{V}_K = \mathbf{I}$, we have that

$$\text{trace}(\mathbf{R}_\epsilon) = \sigma^2 \mathbb{E}(\|\underline{s}_P\|_2^2) \text{trace}(\Phi_K \text{diag}(\mathbf{c}) \Phi_K^H). \quad (56)$$

To find an expression for $\mathbb{E}(\|\underline{s}_P\|_2^2)$, we leverage the fact that $\hat{\mathbf{y}}_K = \Phi_K \mathbf{C}^T \underline{s}_P$ [cf. (41) and (44)] to write $\underline{s}_P = \mathbf{C} \Phi_K^H (\Phi_K \mathbf{C}^T \mathbf{C} \Phi_K^H)^{-1} \hat{\mathbf{y}}_K$ and, consequently, to write $\|\underline{s}_P\|_2^2$ as

$$\|\underline{s}_P\|_2^2 = \underline{s}_P^H \underline{s}_P = \hat{\mathbf{y}}_K^H (\Phi_K \text{diag}(\mathbf{c}) \Phi_K^H)^{-1} \hat{\mathbf{y}}_K. \quad (57)$$

Using the expression for the expected value of a quadratic form, it follows that

$$\mathbb{E}(\|\underline{s}_P\|_2^2) = \text{trace} \left(\left(\Phi_K \text{diag}(\mathbf{c}) \Phi_K^H \right)^{-1} \right). \quad (58)$$

Upon replacing (58) into (56) and recalling that σ^2 does not depend on \mathbf{c} , the expression in (54) follows. ■

The statistical assumption on $\hat{\mathbf{y}}_K$ allows us to design \mathbf{c}^* such that the *expected* performance of the reconstruction scheme is optimized. In this way, the choice of the seeding nodes and instants is independent of the particular signal being reconstructed.

Even though obtaining general relaxations to efficiently approximate the non-convex problem in (54) is out of the scope of the paper, we can gain intuition by specializing (54) for time-varying signals, i.e., by setting $\mathbf{S} = \mathbf{A}_{dc}$. For SN-MT seeding, where designing \mathbf{c} boils down to selecting the seeding node, it can be shown that the objective in (54) does not depend on the particular node chosen. This is as it should be, since in the directed cycle every node is topologically indistinguishable from the others. For MN-ST seeding, the best strategy is to uniformly distribute the seeding nodes, as we formally state next.

Proposition 6 *Suppose that the problem in (54) is particularized for the case of MN-ST seeding of time-varying signals using an ideal low-pass filter. Then, if $K = P = N/\theta$, it holds that the optimal seeding strategy selects the nodes in positions $1, 1 + \theta, \dots, 1 + (K - 1)\theta$.*

Proof: When $\mathbf{S} = \mathbf{A}_{dc}$ we have that: a) $\mathbf{V} = \mathbf{F}$; b) $\hat{\mathbf{h}}_K^* = \alpha \mathbf{1}_K$ for some constant α where $\mathbf{1}_K$ is the $K \times 1$ vector of all ones – since we are considering an ideal low-pass filter –; and c) \mathbf{c} only can take nonzero values in positions $i = 1, 1 + \tau, \dots, 1 + (N - 1)\tau$ (since we are considering MN-ST seeding). Leveraging a), b) and c), problem (54) can be reformulated as

$$\mathbf{c}^* := \underset{\mathbf{c}}{\text{argmin}} \text{trace}(\mathbf{M}^{-1}) \text{trace}(\mathbf{M}) \quad (59)$$

$$\text{s.t. } \mathbf{M} = \mathbf{E}_K^T \mathbf{F}^H \text{diag}(\mathbf{c}) \mathbf{F} \mathbf{E}_K, \quad \mathbf{c} \in \{0, 1\}^N, \quad \|\mathbf{c}\|_0 = K,$$

where \mathbf{c} selects K seeding nodes out of the N possible ones. First, notice that $\text{trace}(\mathbf{M})$ does not depend on the particular choice of \mathbf{c} . To see why this is true, we denote by $\mathcal{I}(\mathbf{c})$ the set containing the indices of the K seeding nodes selected by \mathbf{c} . Then, we can exploit the structure in \mathbf{F} to write

$$\text{trace}(\mathbf{M}) = \sum_{i \in \mathcal{I}(\mathbf{c})} \sum_{j=0}^{K-1} \left| \frac{1}{\sqrt{N}} e^{+j \frac{2\pi}{N} (i-1)(j-1)} \right|^2 = \frac{K^2}{N}, \quad (60)$$

which does not depend on \mathbf{c} . Hence, the optimal \mathbf{c}^* in (59) can be found as the one minimizing $\text{trace}(\mathbf{M}^{-1})$.

If we denote by $\{\gamma_i\}_{i=1}^K$ the K eigenvalues of \mathbf{M} , our goal is then to find the \mathbf{c}^* that minimizes $\sum_i 1/\gamma_i$. Given that all γ_i

$K = 4, P = 4$	MN-ST	SN-MT	MN-MT (2 nodes \times 2 seeds)	
% of recovery	91.3	95.6	93.8	
Min error	.003	.013	.006	
$K = 4, P = 8$	MN-ST	SN-MT	MN-MT (4 \times 2)	MN-MT (2 \times 4)
% of recovery	99.7	95.6	99.5	98.4
Min error	.002	.009	.002	.003

TABLE I: Recovery performance for the three seeding schemes in 10,000 random graphs of size 10 for the cases when the number of seeding values P equals the bandwidth K and when it doubles it.

are nonnegative (\mathbf{M} is positive semi-definite) and (60) implies that $\sum_i \gamma_i = K^2/N$, the minimization is achieved by setting $\gamma_1 = \gamma_2 = \dots = \gamma_K = K/N$. Hence, if we show that uniform seeding leads to $\gamma_i = K/N$ for all i , the proof concludes. To show this, notice that under uniform sampling

$$\text{diag}(\mathbf{c})\mathbf{F}\mathbf{E}_K = \sqrt{\frac{K}{N}}\mathbf{F}^{(K)}, \quad (61)$$

where $\mathbf{F}^{(K)}$ is the Fourier basis of size $K \times K$. Hence, $\mathbf{M} = K/N\mathbf{I}$ [cf. (59)] and every eigenvalue of \mathbf{M} equals K/N . ■

In words, even though for the noiseless case any seeding selection strategy satisfying the conditions in Proposition 2 is equally optimal, uniform seeding in directed cycles is the best MN-ST scheme when noise is present in \underline{s}_P .

A second scenario of interest are setups where the additive noise at different value injections is uncorrelated and of fixed power, i.e., $\mathbf{R}_w = \sigma^2\mathbf{I}$. In this case, (53) can be rewritten as

$$\mathbf{R}_e = \sigma^2\mathbf{V}_K\Phi_K\text{diag}(\mathbf{c})\Phi_K^H\mathbf{V}_K^H. \quad (62)$$

The design of \mathbf{c} that minimizes the MSE of the reconstruction is the solution of the following linear integer program

$$\begin{aligned} \mathbf{c}^* := \underset{\mathbf{c}}{\text{argmin}} \text{trace}(\mathbf{R}_e) &= \underset{\mathbf{c}}{\text{argmin}} \text{trace}(\Phi_K\text{diag}(\mathbf{c})\Phi_K^H) \\ \text{s.t. } \mathbf{c} &\in \{0, 1\}^{N\tau}, \quad \|\mathbf{c}\|_0 = P \end{aligned} \quad (63)$$

which can be approximated by relaxing the binary and 0-norm constraints.

It turns out that the solution of (63) promotes the injection of seeding values at nodes that weakly express the active frequencies, i.e., nodes j such that the values $[\hat{e}_j]_k$ for $k \leq K$ are small. This occurs because the noise power is fixed and those nodes require the injection of seeding signals with high power, leading to a high SNR.

VII. NUMERICAL EXPERIMENTS

We illustrate the reconstruction schemes in noiseless and noisy scenarios using synthetic (Section VII-A) and real-world graphs (Sections VII-B and VII-C).

A. Synthetic graph signals

Figure 2a represents a graph \mathcal{G} with $N = 10$ nodes and adjacency matrix \mathbf{A} generated using an Erdős-Rényi (ER) model with edge probability 0.3 [41]. Define the graph-shift operator $\mathbf{S} = \mathbf{A}$ and let \mathbf{y} be a signal to be recovered. Though seemingly random in the node domain, the structure of \mathbf{y} is highly determined by \mathcal{G} . Indeed, \mathbf{y} has bandwidth $K = 4$, as can be observed from its frequency representation $\hat{\mathbf{y}}$ in Figure 2a.

$K = 4, P = 4$	MN-ST	SN-MT	MN-MT (2 nodes \times 2 seeds)	
% of recovery	100.0	100.0	100.0	
Min error	.004	.125	.011	
$K = 4, P = 8$	MN-ST	SN-MT	MN-MT (4 \times 2)	MN-MT (2 \times 4)
% of recovery	100.0	100.0	100.0	100.0
Min error	.003	.037	.003	.006

TABLE II: Counterpart of Table I based on 1,000 random graphs of size 100.

The first set of experiments illustrates the perfect recovery of \mathbf{y} when $P = 4$ seeding values are injected into \mathcal{G} followed by a low-pass filter of degree $N - P = 6$. The reconstruction is carried out using MN-MT seeding (Section V) where nodes 1 and 2 act as seeding nodes and each of them injects a seeding value for time instants $t \in \{0, 1\}$. After the seeding phase, a filter that successively annihilates the $N - K = 6$ frequencies not active in \mathbf{y} is implemented [cf. (6)]. The evolutions of the reconstructed signal and its frequency representation are depicted in Figures 2b and 2c, respectively. Notice that perfect reconstruction is achieved since the last column in both figures coincide with \mathbf{y} and $\hat{\mathbf{y}}$. Figure 2b illustrates that the reconstructed signal is sparse during the seeding phase, consisting of the first two time instants. More specifically, for $t = 0$ the signal attains nonzero values only for the seeding nodes [cf. (7)] and for $t = 1$ the signal remains zero for every node outside of the one-hop neighborhood of the seeding nodes. During the filtering phase – times $t = 2$ to $t = 7$ – signal values are successively exchanged between neighboring nodes in order to finally recover \mathbf{y} at time $t = 7$. Figure 2c helps to understand the operation of the filtering phase. The signal \mathbf{x} obtained after the seeding phase ($t = 1$) contains every frequency not active in the desired signal \mathbf{y} . Thus, in every successive time instant, one of these frequencies is annihilated. E.g., at time $t = 2$ the frequency with index $i = 5$ is eliminated and at $t = 3$ the frequency $i = 6$ is eliminated. In this way, at time $t = 7$ every frequency not active in \mathbf{y} has been annihilated and perfect recovery is achieved.

To compare the reconstruction performance of MN-ST, SN-MT, and MN-MT seeding, we generate two sets of random Erdős-Rényi (ER) graphs of different sizes: 10,000 graphs with 10 nodes and edge probability 0.3, and 1,000 graphs with 100 nodes and edge probability 0.1. On each graph we define a 4-bandlimited signal and try to recover it through the three seeding schemes presented; see Table I for 10-node graphs and Table II for 100-node graphs. Regarding the amount of seeding values P , we consider cases where P is equal to the minimum necessary for perfect recovery, i.e., $P = K$, and cases where the number of seeding values are more than necessary $P = 2K$. When $P = 4$, we restrict the MN-MT schemes to those consisting of two seeding nodes injecting two values each and when $P = 8$ we consider two variants of MN-MT consisting of two nodes injecting four values each and four nodes injecting two values each.

We first obtain the recovery percentage of the three schemes in noiseless scenarios. More specifically, for a given graph and signal to recover both of size N , we test for perfect recovery for N different choices of seeding nodes. For SN-MT seeding this implies considering all N choices of the single seeding node whereas for the remaining schemes we randomly pick N configurations of seeding nodes. In this way, every entry in

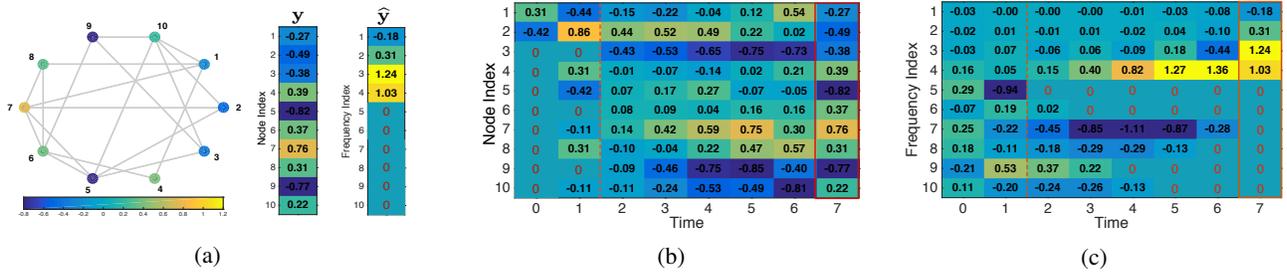


Fig. 2: Perfect recovery of a bandlimited graph signal. (a) The graph \mathcal{G} , the target signal to recover \mathbf{y} and its frequency representation $\hat{\mathbf{y}}$. (b) Evolution of the reconstructed signal. The seeding and filtering phases are separated by a dotted line and the recovered signal is framed in red. (c) Evolution of the frequency components of the reconstructed signal. Successive annihilation during the filtering phase is observed.

Tables I and II is based on 100,000 measurements, obtained from the N seeding choices times the number of graphs considered. E.g., the SN-MT success rate of 95.6% shown in Table I means that in 95,600 experiments the bandlimited signal was recovered perfectly from the seeding values. If we focus on the first row of Table I, the highest recovery percentage of SN-MT suggests that, for the cases where $P = K$, condition *ii*) in Proposition 3 is more commonly satisfied in *random* ER graphs than the respective conditions in Propositions 2 and 5. By contrast, the opposite is true when $P = 2K$. To understand why this is the case, notice that condition *ii*) in Proposition 3 does not depend on the value of P . Intuitively, if the seeding node is incapable of expressing an active frequency $U_1 > 0$ or two active frequencies are indistinguishable $D_1 > 0$, then extra seeding values (larger P) cannot improve recovery. On the other hand, for MN-ST and MN-MT seeding the corresponding condition *ii*) in Propositions 2 and 5 *do* depend on P and, for larger P , it is easier to satisfy the rank condition required. If we now focus on the recovery percentages in Table II we see that perfect recovery is always achieved for all the schemes considered pointing to the fact that in general ER graphs, whenever $K \ll N$, the conditions in Propositions 2, 3, and 5 are satisfied.

We then introduce noise in the injections following the constant SNR model in Section VI-B for $\sigma = 10^{-3}$. Denoting by \mathbf{z} the signal obtained from the reconstruction and by \mathbf{y} the desired signal, we define the reconstruction error as $\epsilon = \|\mathbf{z} - \mathbf{y}\|_2 / \|\mathbf{y}\|_2$. For every given graph and signal \mathbf{y} , we record the minimum ϵ for every choice of seeding nodes within each reconstruction scheme. In Tables I and II we report the median of these values across all the generated graphs in the ‘Min error’ rows. For the cases where $P = K$, observe that MN-ST is an order of magnitude more robust than SN-MT when $N = 10$ and almost two orders when $N = 100$. Moreover, MN-MT seeding schemes achieve intermediate performances in the presence of noise. Finally, we see that increasing the number of seeding values to $P = 8$ entails an increase in robustness for both graph sizes considered and every seeding scheme. It is interesting to notice that for SN-MT where additional seeding values do not improve recovery in noiseless settings, when noise is present these extra values can be used to increase robustness. E.g., for graphs containing 100 nodes, the median of the minimum reconstruction errors in the 1,000 graphs generated was reduced from 0.125 to 0.037 by doubling the number of seeding values.

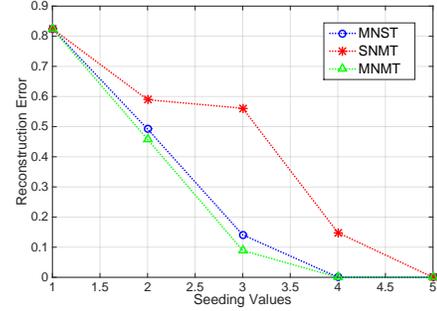


Fig. 3: Reconstruction errors when recovering a signal in a social network with insufficient seeding values.

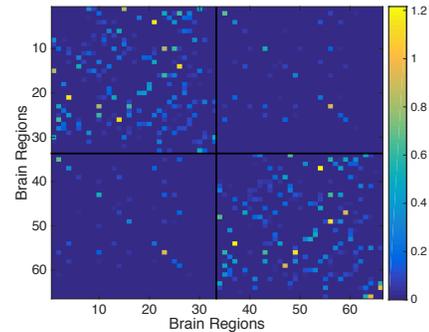


Fig. 4: Heat map of the adjacency matrix \mathbf{A} of brain graph \mathcal{G} .

B. Influencing opinions in social networks

Consider the well-known social network of Zachary’s karate club [42] represented by a graph \mathcal{G} consisting of 34 nodes or members of the club and 78 undirected edges symbolizing friendships among members. Denoting by \mathbf{L} the Laplacian of \mathcal{G} , define the graph shift operator $\mathbf{S} = \mathbf{I} - \alpha\mathbf{L}$ with $\alpha = 1/\lambda_{\max}(\mathbf{L})$, modeling the diffusion of opinions between the members of the club. A signal \mathbf{y} on \mathcal{G} can be interpreted as a unidimensional opinion of each club member regarding a specific topic, and each successive application of \mathbf{S} can be seen as an opinion update influenced by neighboring individuals. Bandlimitedness of \mathbf{y} implies that the opinion discrepancies between neighbors are small. Arbitrary (non-bandlimited) profiles can also be induced when the number of seeding values is sufficiently large. In this context, signal reconstruction can be interpreted as the problem of inducing a desired global opinion profile by influencing the opinion of a subset of members.

We analyze the recovery performance when the number of seeding values is insufficient (Section VI-A). For this, we

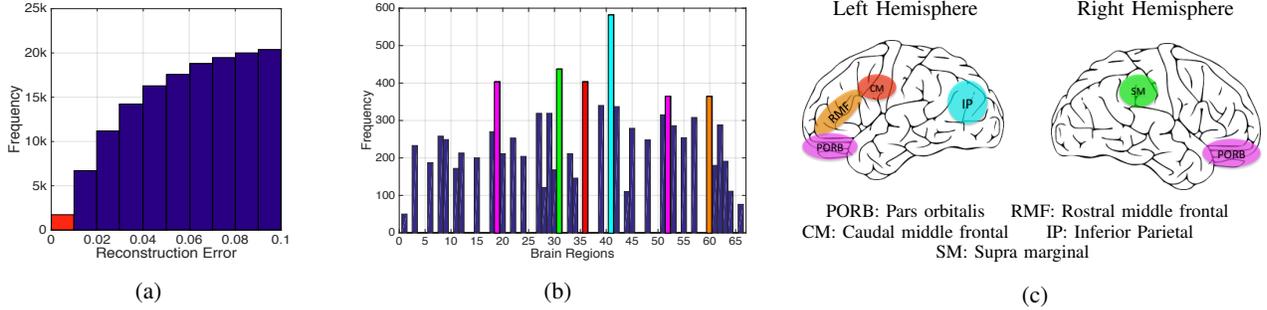


Fig. 6: Inducing a brain state in the presence of noise. (a) Histogram of the reconstruction error for different choices of the six seeding nodes. (b) Frequency of appearance of each brain region among the configurations achieving the lowest reconstruction errors. (c) Anatomical location of the six regions most used in robust seeding configurations.

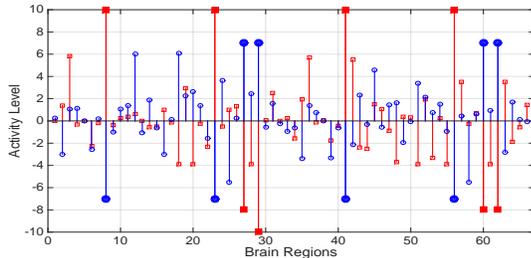


Fig. 5: Initial y_i (red) and target y_t (blue) brain states. High activity is represented by positive activity levels while negative values represent low levels of activity.

generate a signal \mathbf{y} of bandwidth $K = 5$ and try to reconstruct it using P seeding values for $P = 1, \dots, 5$; see Figure 3. For every P , we find the combination of seeding values and locations that minimizes the reconstruction error within each seeding scheme (cf. Section VI-A3). E.g., if $P = 2$ and we are analyzing SN-MT seeding, we consider every individual as possible seeding node and then choose the one achieving the minimum error. In Figure 3 we report the average of these minimum errors across 100 bandlimited signals. As expected, for $P = 1$ the three schemes coincide and for $P = 5$ perfect recovery is achieved for all of them. However, for intermediate values of P , MN-ST presents a considerably lower error than SN-MT. For $P = 3$ this implies that, when trying to induce a global opinion profile, it is more effective to influence the opinion of three individuals once than to influence the opinion of the same individual three times. The fact that MN-MT seeding presents the lowest reconstruction errors is expected since this scheme includes the other two as particular cases. Notice that in this case, as opposed to Tables I and II, we do not restrict MN-MT to any particular combination of seeding values and nodes.

For the above analysis to hold true, we must be able to apply a low-pass filter on the social network as required by the filtering phases of MN-ST, SN-MT, and MN-MT. This can be achieved by assuming that we can modify the rate of exchange of opinions in the network represented by α . Indeed, consider that after the seeding phase, the signal still percolates over the graph – people still communicate their opinions to neighbors – but we can modify the diffusion rate α_l at each discrete time instant. Thus, after $L - 1$ interactions we obtain that

$$\mathbf{z} = \prod_{l=1}^{L-1} (\mathbf{I} - \alpha_l \mathbf{L}) \mathbf{x}, \quad (64)$$

which is equivalent to applying an annihilating filter [cf. (6)]

to \mathbf{x} . Notice that the filter in (64) is a polynomial on \mathbf{L} rather than \mathbf{S} . However, the frequency annihilation procedure is still valid since the eigenvectors – frequency basis – of \mathbf{L} and \mathbf{S} are equal.

C. Inducing a brain state

Upon dividing the human brain into the 66 regions of interest (ROIs) defined in [43], we build a weighted undirected graph \mathcal{G} whose nodes are the ROIs and whose edge weights are given by the density of anatomical connections between regions; see Figure 4. The first 33 ROIs are located on the right hemisphere of the brain while regions 34 to 66 correspond to their left counterparts. From Figure 4 we see that most connections occur within the same cortical hemisphere with few inter hemispheric connections. We define the graph-shift operator $\mathbf{S} = \mathbf{A}$ where \mathbf{A} is the adjacency matrix of \mathcal{G} modeling the electrical connectivity between brain regions. The level of activity of each ROI can be represented by a graph signal \mathbf{y} where larger values represent higher levels of activity. Successive applications of \mathbf{S} on \mathbf{y} model a linear evolution of the brain activity pattern [44]. As a method to inject seeding values to \mathcal{G} , we consider transcranial magnetic stimulation (TMS) [24], a noninvasive method to stimulate ROIs. In this context, reconstructing a brain signal amounts to inducing a specific brain state via TMS. In particular, we consider the problem of driving the brain from a resting state to one associated with high-level cognitive operations. This experiment is included as an illustration of the potential applicability of our framework to neuroscience, rather than as a tested procedure with immediate clinical implementation.

Brain resting states are associated with high activity in the posterior cingulate (PC) and inferior parietal (IP) cortices whereas active states are associated with high activity in the rostral middle frontal (RMF) and superior parietal (SP) cortices [45], [46]. In Figure 5 we present the initial y_i and target y_t signals, where the activity corresponding to the eight regions mentioned – left and right versions of each cortex – is highlighted with larger markers. In order to drive the brain from y_i to y_t we consider a MN-MT seeding scheme with six seeding nodes. Since it is unclear how to implement a low-pass filter in a human brain, we consider that each seeding node injects eleven values, totalizing $P = 66$ seeding values permitting the recovery of the target signal after the seeding phase without the need of a posterior filtering phase. Notice that throughout the paper we assumed the initial signal y_i to

be zero, meaning that there is no signal present on the graph before the reconstruction process. However, our model can accommodate for \mathbf{y}_i different from zero. To see this, if the seeding phase lasts τ instants, then we can design our seeding values to recover the signal $\mathbf{y}_r = \mathbf{y}_t - \mathbf{S}^{\tau-1} \mathbf{y}_i$ in the original formulation so that the negative term cancels the effect of the seeding phase on \mathbf{y}_i and the target signal \mathbf{y}_t is recovered.

We consider noisy injections following the constant SNR model in Section VI-B for $\sigma = 10^{-3}$. Since in this case we have an initial signal present in the graph, we define the reconstruction error as $\epsilon = \|\mathbf{z} - \mathbf{y}_t\|_2 / \|\mathbf{y}_r\|_2$, where \mathbf{z} denotes the reconstructed signal, and compute ϵ for every possible combination of seeding nodes. Given that the seeding values are induced by TMS, we discard as possible seeding nodes the regions inaccessible by TMS like the ones located in the medial cortex and subcortical structures. After discarding inaccessible ROIs, the six seeding nodes can be chosen out of 38 possible ROIs, amounting to 2,760,681 possible configurations. In Figure 6a we present a histogram of the reconstruction error for different seeding configurations where we only show those attaining errors below 0.1. The red bar in this histogram corresponds to the 1,611 configurations that achieve the lowest reconstruction errors. In Figure 6b we present the frequency of appearance of each ROI in these 1,611 robust seeding configurations. The regions with zero appearances correspond to the ROIs inaccessible to TMS, however, among the accessible regions the frequency of appearance is not uniform. For example, the left inferior parietal cortex in position 41 appears 583 times whereas the right bank of the frontal pole in position 1 is only used 50 times. In Figure 6c we depict the 6 regions more commonly used in robust seeding configurations. Notice that both the left and right versions of the Pars Orbitalis are commonly used as seeding nodes, suggesting the importance of this region for robust brain state induction. These robust seeding configurations do not appear to be very sensitive to the noise level σ . For instance, if we increase the noise level to $\sigma = 10^{-2}$, the maximum error achieved among the best 1,611 strategies is 0.108 (as opposed to 0.01 in the current setting), but it is still the case that the inferior parietal (IP) is the most common region in these robust seeding configurations.

VIII. CONCLUSIONS

A novel approach for the recovery of bandlimited graph signals – that admit a sparse representation in the frequency domain – was proposed. The focus was not on estimating an unknown graph signal but rather on inducing a known bandlimited signal through minimal actions on the graph. These actions referred to signal injections at different seeding nodes, which then percolate through the graph via local interactions described by a graph filter. Restrictions on the number of seeding nodes and the amount of injections at each node gave rise to three different reconstruction schemes and their performance in noiseless and noisy settings was analyzed. For the noiseless case, we showed that a K -bandlimited signal can be recovered using K injections followed by a low-pass filter in the (graph) frequency domain. In contrast to classical time-varying signals, it was also shown that if the seeding nodes inject the values of the original signal in those nodes, perfect recovery is not feasible. For scenarios leading to

imperfect reconstruction, we analyzed robust seeding strategies to minimize distortion. Finally, the different reconstruction schemes were illustrated through numerical experiments in both synthetic and real-world graph signals.

Future research directions include consideration of scenarios where \mathbf{S} is not given but, rather, has to be estimated from the observations of graph signals. Moreover, settings where some of the elements of \mathbf{S} can be designed to benefit quick and robust signal reconstruction are also of interest.

REFERENCES

- [1] S. Segarra, A. G. Marques, G. Leus, and A. Ribeiro, "Interpolation of graph signals using shift-invariant graph filters," in *European Signal Process. Conf. (EUSIPCO)*, Nice, France, Aug. 31 - Sept. 4 2015.
- [2] —, "Reconstruction of graph signals: percolation from a single seeding node," in *Global Conf. on Signal and Info. Process. (GlobalSIP)*, Orlando, FL, December 2015 (submitted).
- [3] D. Shuman, S. Narang, P. Frossard, A. Ortega, and P. Vandergheynst, "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains," *IEEE Signal Process. Mag.*, vol. 30, no. 3, pp. 83–98, Mar. 2013.
- [4] A. Sandryhaila and J. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Apr. 2013.
- [5] X. Zhu and M. Rabbat, "Approximating signals supported on graphs," in *IEEE Intl. Conf. Acoust., Speech and Signal Process. (ICASSP)*, March 2012, pp. 3921–3924.
- [6] S. Narang, A. Gadde, and A. Ortega, "Signal processing techniques for interpolation in graph structured data," in *IEEE Intl. Conf. Acoust., Speech and Signal Process. (ICASSP)*, May 2013, pp. 5445–5449.
- [7] S. Chen, R. Varma, A. Sandryhaila, and J. Kovačević, "Discrete signal processing on graphs: Sampling theory," *arXiv preprint arXiv:1503.05432*, 2015.
- [8] A. G. Marques, S. Segarra, G. Leus, and A. Ribeiro, "Sampling of graph signals with successive local aggregations," *arXiv preprint arXiv:1504.04687*, 2015.
- [9] A. Anis, A. Gadde, and A. Ortega, "Towards a sampling theorem for signals on arbitrary graphs," in *IEEE Intl. Conf. Acoust., Speech and Signal Process. (ICASSP)*, May 2014, pp. 3864–3868.
- [10] A. Gadde, A. Anis, and A. Ortega, "Active semi-supervised learning using sampling theory for graph signals," in *Proc. ACM SIGKDD Intl. Conf. Knowl. Disc. Data Min.*, 2014, pp. 492–501.
- [11] S. Narang, A. Gadde, E. Sanou, and A. Ortega, "Localized iterative methods for interpolation in graph structured data," in *Global Conf. on Signal and Info. Process. (GlobalSIP)*, Dec. 2013, pp. 491–494.
- [12] X. Wang, P. Liu, and Y. Gu, "Local-set-based graph signal reconstruction," *IEEE Trans. Signal Process.*, vol. 63, no. 9, pp. 2432–2444, Sept. 2015.
- [13] X. Wang, M. Wang, and Y. Gu, "A distributed tracking algorithm for reconstruction of graph signals," *IEEE J. Sel. Topics Signal Process.*, vol. 9, no. 4, pp. 728–740, June 2015.
- [14] S. Chen, A. Sandryhaila, J. Moura, and J. Kovačević, "Signal recovery on graphs," *arXiv preprint arXiv:1411.7414*, 2014.
- [15] D. Zhou and B. Schölkopf, "A regularization framework for learning from graph data," in *ICML Work. Stat. Relational Learning and its Connections to other Fields*, vol. 15, 2004, pp. 67–68.
- [16] K. Reinschke, *Multivariable Control: A Graph-theoretic Approach*. Springer-Verlag, 1988.
- [17] Y.-Y. Liu, J.-J. Slotine, and A.-L. Barabasi, "Controllability of complex networks," *Nature*, vol. 473, no. 7346, pp. 167–173, 2011.
- [18] F. Pasqualetti, S. Zampieri, and F. Bullo, "Controllability metrics, limitations and algorithms for complex networks," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 1, pp. 40–52, March 2014.
- [19] G. Yan, J. Ren, Y.-C. Lai, C.-H. Lai, and B. Li, "Controlling complex networks: How much energy is needed?" *Phys. Rev. Lett.*, vol. 108, p. 218703, May 2012.
- [20] T. Summers and J. Lygeros, "Optimal sensor and actuator placement in complex dynamical networks," in *World Congr. Intl. Fed. Auto. Control*, vol. 19, no. 1, 2014, pp. 3784–3789.
- [21] J. Sun and A. E. Motter, "Controllability transition and nonlocality in network control," *Phys. Rev. Lett.*, vol. 110, p. 208701, 2013.
- [22] E. D. Sontag, *Mathematical Control Theory: Deterministic Finite Dimensional Systems*. Springer Science & Business Media, 2013, vol. 6.
- [23] D. J. Watts and P. S. Dodds, "Influentials, networks, and public opinion formation," *J. Consumer Research*, vol. 34, no. 4, pp. 441–458, 2007.

- [24] M. Hallett, "Transcranial magnetic stimulation and the human brain," *Nature*, vol. 406, no. 6792, pp. 147–150, 2000.
- [25] T. Nakano, M. Moore, F. Wei, A. Vasilakos, and J. Shuai, "Molecular communication and networking: Opportunities and challenges," *IEEE Trans. Nanobiosci.*, vol. 11, no. 2, pp. 135–148, June 2012.
- [26] M. S. Kuran, H. B. Yilmaz, T. Tugcu, and B. Ozerman, "Energy model for communication via diffusion in nanonetworks," *Nano Commun. Netw.*, vol. 1, no. 2, pp. 86 – 95, 2010.
- [27] C. Intanagonwivat, R. Govindan, and D. Estrin, "Directed diffusion: A scalable and robust communication paradigm for sensor networks," in *Intl. Conf. Mobile Comput. and Netw.* New York, NY, USA: ACM, 2000, pp. 56–67.
- [28] P. Van Mieghem, "The N-intertwined SIS epidemic network model," *Computing*, vol. 93, no. 2-4, pp. 147–169, 2011.
- [29] A. Sandryhaila and J. Moura, "Discrete signal processing on graphs: Frequency analysis," *IEEE Trans. Signal Process.*, vol. 62, no. 12, pp. 3042–3054, June 2014.
- [30] C. Godsil and G. Royle, *Algebraic graph theory*. Springer-Verlag, Graduate Texts in Mathematics, 2001, vol. 207.
- [31] M. Rabbat and V. Gripon, "Towards a spectral characterization of signals supported on small-world networks," in *IEEE Intl. Conf. Acoust., Speech and Signal Process. (ICASSP)*, May 2014, pp. 4793–4797.
- [32] S. Safavi and U. Khan, "Revisiting finite-time distributed algorithms via successive nulling of eigenvalues," *IEEE Signal Process. Lett.*, vol. 22, no. 1, pp. 54–57, Jan. 2015.
- [33] G. H. Golub and C. F. van Loan, *Matrix Computations*. Johns Hopkins University Press, 1983.
- [34] S. Segarra, A. G. Marques, and A. Ribeiro, "Distributed linear network operators using graph filters," *arXiv preprint arXiv:1510.03947*, 2015.
- [35] S. M. Kay, *Fundamentals of Statistical Signal Processing: Estimation Theory*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
- [36] G. McCormick, "Computability of global solutions to factorable non-convex programs: Part I convex underestimating problems," *Math. Program.*, vol. 10, no. 1, pp. 147–175, 1976.
- [37] H. Konno, "A cutting plane algorithm for solving bilinear programs," *Math. Program.*, vol. 11, no. 1, pp. 14–27, 1976.
- [38] S. Choudhary and U. Mitra, "Identifiability scaling laws in bilinear inverse problems," *arXiv preprint arXiv:1402.2637*, 2014.
- [39] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2014, vol. 16.
- [40] F. Pukelsheim, *Optimal Design of Experiments*. SIAM, 1993, vol. 50.
- [41] B. Bollobás, *Random Graphs*. Springer, 1998.
- [42] W. W. Zachary, "An information flow model for conflict and fission in small groups," *J. Anthropological Research*, vol. 33, no. 4, pp. 452–473, 1977.
- [43] P. Hagmann, L. Cammoun, X. Gigandet, R. Meuli, C. J. Honey, V. J. Wedeen, and O. Sporns, "Mapping the structural core of human cerebral cortex," *PLoS Biol.*, vol. 6, no. 7, p. e159, 2008.
- [44] S. Gu, F. Pasqualetti, M. Cieslak, S. T. Grafton, and D. S. Bassett, "Controllability of brain networks," *arXiv preprint arXiv:1406.5197*, 2014.
- [45] M. D. Greicius, B. Krasnow, A. L. Reiss, and V. Menon, "Functional connectivity in the resting brain: A network analysis of the default mode hypothesis," *Proc. Nat. Academy Sci.*, vol. 100, no. 1, pp. 253–258, 2003.
- [46] J. D. Medaglia, M.-E. Lynall, and D. S. Bassett, "Cognitive network neuroscience," *J. Cogn. Neurosci.*, vol. 27, no. 8, pp. 1471–1491, July 2015.



Santiago Segarra (S'12) received the B.Sc. degree in industrial engineering with highest honors (Valedictorian) from the Instituto Tecnológico de Buenos Aires (ITBA), Argentina, in 2011 and the M.Sc. degree in electrical engineering from the University of Pennsylvania, Philadelphia, in 2014. Since 2011, he has been working towards the Ph.D. degree in the Department of Electrical and Systems Engineering at the University of Pennsylvania.

His research interests include network theory, data analysis, machine learning, and graph signal processing. Mr. Segarra received the ITBA's 2011 Best Undergraduate Thesis Award in industrial engineering, the 2011 Outstanding Graduate Award granted by the National Academy of Engineering of Argentina, and the Best Student Paper Award at the 2015 Asilomar Conference.



Antonio G. Marques (SM'13) received the telecommunications engineering degree and the Doctorate degree, both with highest honors, from the Carlos III University of Madrid, Spain, in 2002 and 2007, respectively. In 2007, he became a faculty of the Department of Signal Theory and Communications, King Juan Carlos University, Madrid, Spain, where he currently develops his research and teaching activities as an Associate Professor. From 2005 to 2015, he held different visiting positions at the University of Minnesota, Minneapolis. In 2015 and 2016 he was a visitor scholar at the University of Pennsylvania, Philadelphia.

His research interests lie in the areas of signal processing, networking and communications. His current research focuses on stochastic optimization of wireless and power networks, signal processing for graphs, and nonlinear network optimization. Dr. Marques has served the IEEE in a number of posts (currently, he is an Associate Editor of the Signal Process. Letters) and his work has been awarded in several conferences and workshops.



Geert Leus (F'13) received the MSc and PhD degree in applied sciences from the Katholieke Universiteit Leuven, Belgium, in June 1996 and May 2000, respectively. Currently, Geert Leus is an "Antoni van Leeuwenhoek" Full Professor at the Faculty of Electrical Engineering, Mathematics and Computer Science of the Delft University of Technology, The Netherlands. His research interests are in the area of signal processing for communications. Geert Leus received a 2002 IEEE Signal Processing Society Young Author Best Paper Award and a 2005 IEEE

Signal Processing Society Best Paper Award. He is a Fellow of IEEE and a Fellow of EURASIP. Geert Leus was the Chair of the IEEE Signal Processing for Communications and Networking Technical Committee, and an Associate Editor for the IEEE Transactions on Signal Processing, the IEEE Transactions on Wireless Communications, the IEEE Signal Processing Letters, and the EURASIP Journal on Advances in Signal Processing. Currently, he is a Member-at-Large to the Board of Governors of the IEEE Signal Processing Society and a member of the IEEE Sensor Array and Multichannel Technical Committee. He finally serves as the Editor in Chief of the EURASIP Journal on Advances in Signal Processing.



Alejandro Ribeiro (M'07) received the B.Sc. degree in electrical engineering from the Universidad de la Republica Oriental del Uruguay, Montevideo, in 1998 and the M.Sc. and Ph.D. degree in electrical engineering from the Department of Electrical and Computer Engineering, the University of Minnesota, Minneapolis in 2005 and 2007. From 1998 to 2003, he was a member of the technical staff at Bell-south Montevideo. After his M.Sc. and Ph.D studies, in 2008 he joined the University of Pennsylvania (Penn), Philadelphia, where he is currently the

Rosenbluth Associate Professor at the Department of Electrical and Systems Engineering.

His research interests are in the applications of statistical signal processing to the study of networks and networked phenomena. His current research focuses on wireless networks, network optimization, learning in networks, networked control, robot teams, and structured representations of networked data structures. Dr. Ribeiro received the 2014 O. Hugo Schuck best paper award, the 2012 S. Reid Warren, Jr. Award presented by Penn's undergraduate student body for outstanding teaching, the NSF CAREER Award in 2010, and student paper awards at the 2013 American Control Conference and 2015 Asilomar Conference (as adviser), as well as the 2005 and 2006 International Conferences on Acoustics, Speech and Signal Processing. Dr. Ribeiro is a Fulbright scholar and a Penn Fellow.