

LINEAR NETWORK OPERATORS USING NODE-VARIANT GRAPH FILTERS

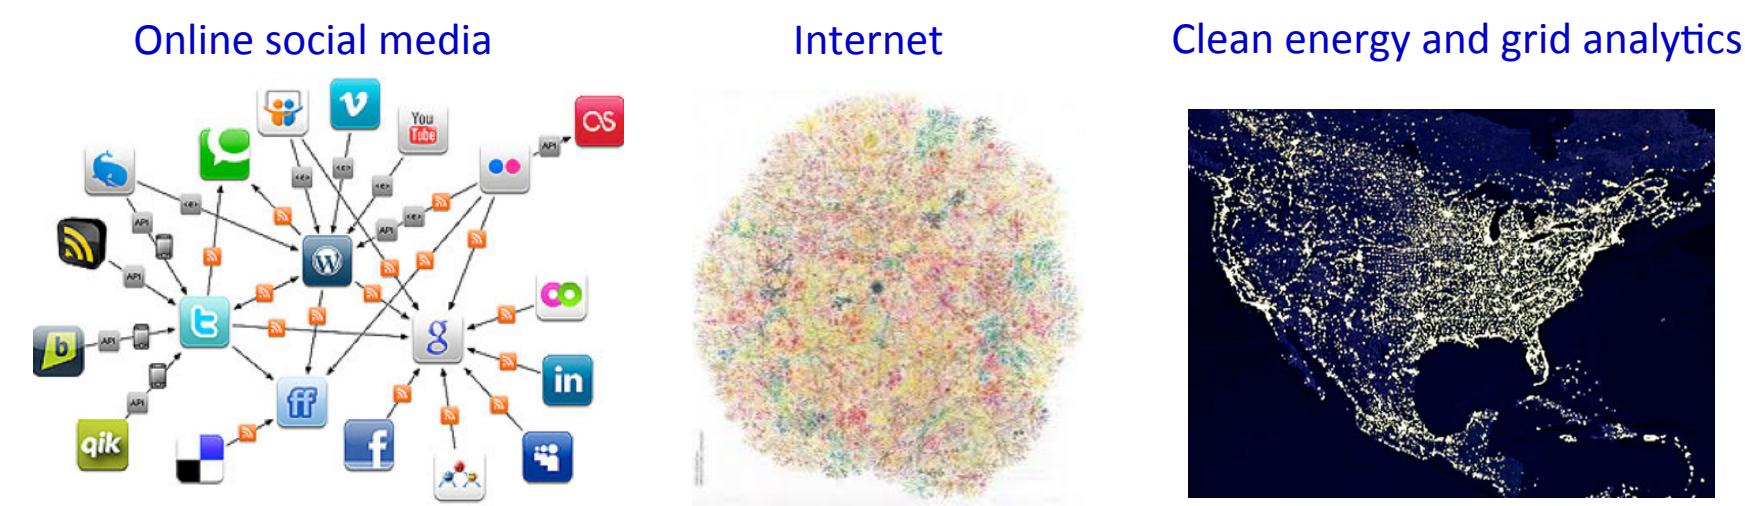
Santiago Segarra*, Antonio G. Marques†, and Alejandro Ribeiro* E-mail: ssegarra@seas.upenn.edu

*Electrical & Systems Engineering, University of Pennsylvania


†Signal Theory & Communications, King Juan Carlos University



Graph signal processing - 101



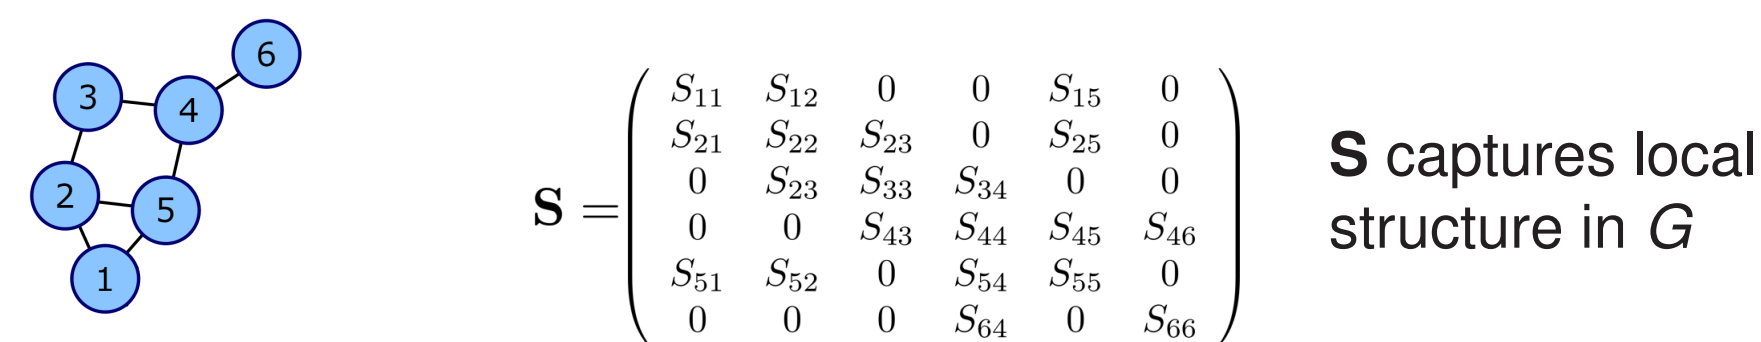
- **Desiderata:** Process, analyze and learn from **network data**
- **Network as graph** $G = (\mathcal{V}, \mathcal{E})$: encode pairwise relationships
- Interest here not in G itself, but in **data** associated with **nodes** in \mathcal{V}
 - ⇒ The object of study is a **graph signal**
- **Ex:** Opinion profile, buffer congestion levels, **neural activity**, epidemic



- Graph SP: extending classical SP results to graph signals
 - ⇒ **Our view:** GSP is well-suited to study network processes
- Filtering, smoothing, prediction, signal synthesis, compression

Graph signals and graph-shift operator

- **Graph signals** are mappings $\mathbf{x} : \mathcal{V} \rightarrow \mathbb{R}$
 - ⇒ May be represented as a vector $\mathbf{x} \in \mathbb{R}^N$ (with $|\mathcal{V}| = N$)
 - ⇒ To understand and analyze \mathbf{x} , useful to account for G 's structure
- Graph G is endowed with a **graph-shift operator** \mathbf{S}
 - ⇒ Matrix $\mathbf{S} \in \mathbb{R}^{N \times N}$ satisfying: $S_{ij} = 0$ for $i \neq j$ and $(i, j) \notin \mathcal{E}$



- **Ex:** Adjacency \mathbf{A} , Degree \mathbf{D} and Laplacian \mathbf{L}

Locality of S and frequency-domain representation

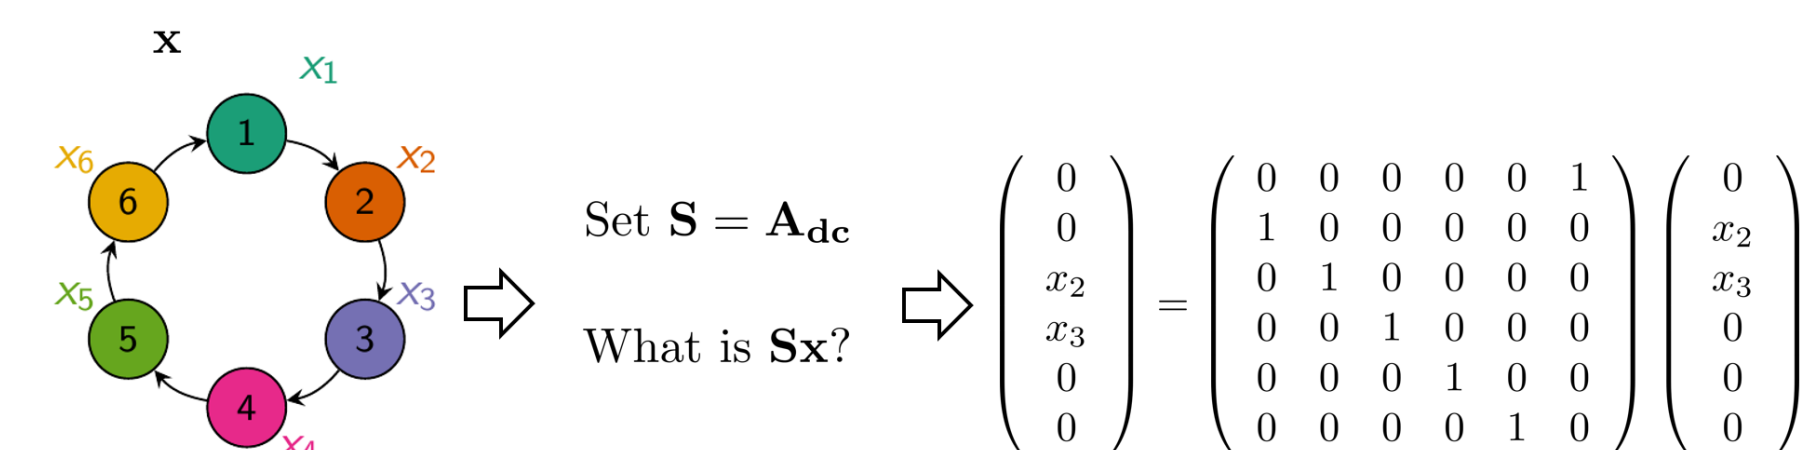
- \mathbf{S} is a **local linear** operator \Rightarrow If $\mathbf{y} = \mathbf{S}\mathbf{x}$, $y_i = \sum_{j \in \mathcal{N}_i^+} S_{ij}x_j \Rightarrow$ 1-hop info
- Spectrum of \mathbf{S} useful to analyze \mathbf{x}
 - ⇒ Consider **diagonalizable** shifts $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$
- Leverage \mathbf{S} to define graph Fourier transform (GFT) and iGFT

$$\tilde{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}, \quad \mathbf{x} = \mathbf{V}\tilde{\mathbf{x}}$$

- **Key message:** the two basic elements of GSP are \mathbf{x} and \mathbf{S}

Connection with discrete-time signal processing

- **Q:** Why is \mathbf{S} called shift? **A:** Resemblance to time shifts



- When $\mathbf{S} = \mathbf{A}_{dc}$, we have that \mathbf{V}^{-1} is the DFT matrix
- **Point of contact** between GSP and DSP
- When particularizing GSP to the **directed cycle**
 - ⇒ We recover known result from DSP

Linear (shift-invariant) graph filter

- A **graph filter** $H : \mathbb{R}^N \rightarrow \mathbb{R}^N$ is a **map** between **graph signals**
 - ⇒ Focus on linear filters $\Rightarrow N \times N$ matrix
- Filter \mathbf{H} is a polynomial in \mathbf{S} of degree $L - 1$, with coeff. $\mathbf{h} = [h_0, \dots, h_{L-1}]^T$

$$\mathbf{H} := h_0\mathbf{S}^0 + h_1\mathbf{S}^1 + \dots + h_{L-1}\mathbf{S}^{L-1} = \sum_{l=0}^{L-1} h_l\mathbf{S}^l$$
- **Key properties:** shift invariance and distributed implementation
 - ⇒ Satisfies $\mathbf{H}(\mathbf{S}\mathbf{x}) = \mathbf{S}(\mathbf{H}\mathbf{x})$, only L -hop information to form $\mathbf{y} = \mathbf{H}\mathbf{x}$

Frequency response of a graph filter

- Using $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$, filter is $\mathbf{H} = \sum_{l=0}^{L-1} h_l\mathbf{S}^l = \mathbf{V} \left(\sum_{l=0}^{L-1} h_l\mathbf{\Lambda}^l \right) \mathbf{V}^{-1}$
- Since $\mathbf{\Lambda}^l$ are diagonal, the GFT-iGFT can be used to write $\mathbf{y} = \mathbf{H}\mathbf{x}$ as

$$\tilde{\mathbf{y}} = \text{diag}(\tilde{\mathbf{h}})\tilde{\mathbf{x}}$$
 - ⇒ Output at frequency k depends only on input at frequency k
- **Frequency response** of \mathbf{H} is $\tilde{\mathbf{h}} = \Psi\mathbf{h}$, where $\Psi_{ij} = \lambda_j^{i-1}$ (Vandermonde)
- Note that GFTs for signals ($\tilde{\mathbf{x}} = \mathbf{V}^{-1}\mathbf{x}$) and filters ($\tilde{\mathbf{h}} = \Psi\mathbf{h}$) are different
 - ⇒ If $\mathbf{S} = \mathbf{A}_{dc}$ (periodic signal), both Ψ and \mathbf{V}^{-1} are equal to the DFT

Motivation and problem formulation

- Design **graph filters** to implement a given **linear transformation**
 - ⇒ Implementation is **distributed** by construction
 - ⇒ Conditions for **perfect** and **approximate** implementation
 - ⇒ Leverage results from classical time-invariant systems
- Given a **linear transformation** \mathbf{B} , find the **filter coefficients** \mathbf{h} s. t.

$$\mathbf{B} = \sum_{l=0}^{L-1} h_l\mathbf{S}^l$$

- Graph-shift operator \mathbf{S} is **given**
 - ⇒ Well-suited for cases where \mathbf{S} is a **network process**
 - ⇒ E.g., diffusion in a social network
 - ⇒ Agents exchange information and weigh info observed
 - ⇒ Choosing $\mathbf{h} \Rightarrow$ **fixing the weights**

Conditions for perfect implementation

- Notation: a) D # of distinct $\{\lambda_k\}_{k=1}^N$; b) $\{\gamma_k\}_{k=1}^N$ eigenvalues of \mathbf{B}


Perfect implementation of linear graph operators

The **linear transformation** \mathbf{B} can be implemented using a **graph filter** \mathbf{H} if the three following conditions hold true:

- Matrices \mathbf{B} and \mathbf{S} are simultaneously diagonalizable.
- For all (k_1, k_2) such that $\lambda_{k_1} = \lambda_{k_2}$, it holds that $\gamma_{k_1} = \gamma_{k_2}$.
- The degree $L - 1$ of \mathbf{H} satisfies $L \geq D$.

- i) \Rightarrow frequency basis of \mathbf{B} and \mathbf{S} the same \Rightarrow necessary
- ii) \Rightarrow two equal freqs. in \mathbf{S} must be equal in $\mathbf{B} \Rightarrow$ necessary

- **Restrictive** conditions but **not impossible**
 - ⇒ **Consensus** $\mathbf{B}_{\text{con}} = \mathbf{1}\mathbf{1}^T$ favors i) and
 - ii) because it is **rank-one**




Node-variant graph filters: definition

- A generalization of graph filters:

$$\mathbf{H}_{\text{nv}} := \sum_{l=0}^{L-1} \text{diag}(\mathbf{h}^{(l)})\mathbf{S}^l$$

- ⇒ When $\mathbf{h}^{(l)} = h_l\mathbf{1} \Rightarrow$ regular (node-invariant) filter



- In general, when \mathbf{H}_{nv} is applied to a signal \mathbf{x}
 - ⇒ **Each node** applies **different weights** to the shifted signals $\mathbf{S}^l\mathbf{x}$
 - ⇒ More flexible and still distributed, not shift-invariant

Node-variant graph filters: frequency response

- Collect the coefficients of node i in \mathbf{h}_i , such that $[\mathbf{h}_i]_l = [\mathbf{h}^{(l)}]_i$
- Focus on the filter output at node i , $\mathbf{e}_i^T\mathbf{H}_{\text{nv}}\mathbf{x}$, and define $\mathbf{u}_i := \mathbf{V}^T\mathbf{e}_i$

$$\eta_i^T = \mathbf{e}_i^T\mathbf{H}_{\text{nv}} = \sum_{l=0}^{L-1} [\mathbf{h}_i]_l \mathbf{e}_i^T\mathbf{V}\mathbf{\Lambda}^l\mathbf{V}^{-1} = \mathbf{u}_i^T \text{diag}(\Psi\mathbf{h}_i)\mathbf{V}^{-1}$$

- The output of the filter at node i , $\eta_i^T\mathbf{x}$ is the **inner product** of
 - ⇒ $\mathbf{V}^{-1}\mathbf{x} \Rightarrow$ the **frequency representation** of the input, and
 - ⇒ $\mathbf{u}_i \Rightarrow$ how **strongly the frequencies are expressed** by node i
 - ⇒ **Modulated** by $\Psi\mathbf{h}_i \Rightarrow$ **Frequency response** associated to i

Perfect reconstruction with node-variant filters

- **Node-variant filters** can implement a larger class of transformations
 - ⇒ Pick $\mathbf{h}^{(l)}$ for $l = 0, \dots, L - 1$ so that

$$\mathbf{B} = \sum_{l=0}^{L-1} \text{diag}(\mathbf{h}^{(l)})\mathbf{S}^l$$

- Defining $\mathbf{b}_i := \mathbf{B}^T\mathbf{e}_i$ and $\tilde{\mathbf{b}}_i := \mathbf{V}^T\mathbf{b}_i$

Perfect implementation using node-variant filters

The linear transformation \mathbf{B} can be implemented using the node-variant filter \mathbf{H}_{nv} if the three following conditions hold for all i :

- $[\tilde{\mathbf{b}}_i]_k = 0$ for those k such that $[\mathbf{u}_i]_k = 0$.
- For all (k_1, k_2) s. t. $\lambda_{k_1} = \lambda_{k_2}$, it holds $[\tilde{\mathbf{b}}_i]_{k_1}/[\mathbf{u}_i]_{k_1} = [\tilde{\mathbf{b}}_i]_{k_2}/[\mathbf{u}_i]_{k_2}$.
- The degree $L - 1$ of \mathbf{H}_{nv} satisfies $L \geq D$.

- **Less** restrictive that for **node-invariant** filters
 - ⇒ Not surprising since we have **additional freedom** in the design

- Some shift operators \mathbf{S} can implement **any** transformation \mathbf{B}

Corollary

Any linear transformation \mathbf{B} can be implemented using a node-varying filter of degree $N - 1$ if the graph-shift operator $\mathbf{S} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{-1}$ satisfies:

- all the entries of \mathbf{V} are non-zero.
- all the eigenvalues $\{\lambda_k\}_{k=1}^N$ are distinct.

- Under these conditions, unique set of optimal coefficients $\{\mathbf{h}_i^*\}_{i=1}^N$

$$\mathbf{h}_i^* = \Psi^{-1} \text{diag}(\mathbf{u}_i)^{-1} \mathbf{V}^T \mathbf{b}_i$$

Approximate implementation: No prior knowledge

- When **perfect** reconstruction is **infeasible**
 - ⇒ Minimize a **pre-specified error metric**, design \mathbf{H}_{nv} to resemble \mathbf{B}
 - ⇒ We look at minimizing $\|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_F$ and $\|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_2$
- Define $\Phi_i := (\mathbf{V}^{-1})^T \text{diag}(\mathbf{u}_i)\Psi$ and $\tilde{\mathbf{U}} := [\text{diag}(\mathbf{u}_1), \dots, \text{diag}(\mathbf{u}_N)]^T$
- Collect the desired coefficient vectors in $\Gamma = [\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N]$

Approximate implementation with no knowledge of x

The **optimal filter coefficients** defined as $\{\mathbf{h}_{i,F}^*\}_{i=1}^N := \arg\min_{\{\mathbf{h}_i\}_{i=1}^N} \|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_F$ and $\{\mathbf{h}_{i,2}^*\}_{i=1}^N := \arg\min_{\{\mathbf{h}_i\}_{i=1}^N} \|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_2$ are, respectively, given by

$$\mathbf{h}_{i,F}^* = \Phi_i^\dagger \mathbf{b}_i \triangleq (\Phi_i^T \Phi_i)^{-1} \Phi_i^T \mathbf{b}_i,$$

for all i , where \triangleq holds if Φ_i has full column rank, and

$$\begin{aligned} \{\Gamma_2^*, \mathbf{s}^*\} &= \arg\min_{\{\Gamma, \mathbf{s}\}} \mathbf{s} \\ \text{s. to } &\begin{bmatrix} \mathbf{s} \mathbf{I} & (\mathbf{I} \odot \Psi \Gamma)^T \tilde{\mathbf{U}} \mathbf{V}^{-1} - \mathbf{B} \\ ((\mathbf{I} \odot \Psi \Gamma)^T \tilde{\mathbf{U}} \mathbf{V}^{-1} - \mathbf{B})^T & \mathbf{s} \mathbf{I} \end{bmatrix} \succeq 0. \end{aligned}$$

- For $\|\cdot\|_F$, optimal coefficients result from a **least-squares** problem
- There is **no closed form** for the $\|\cdot\|_2$ case
 - ⇒ Efficiently obtain it by solving a **semi-definite** program

Approximate implementation: Prior knowledge

- Incorporate **prior knowledge** of the **input signal** \mathbf{x} into the design
- Goal is to minimize a metric of the **error vector** $\mathbf{d} := \mathbf{H}_{\text{nv}}\mathbf{x} - \mathbf{B}\mathbf{x}$
- When \mathbf{x} is zero-mean with covariance \mathbf{R}_x , the error covariance is

$$\mathbf{R}_d := \mathbb{E}[\mathbf{d}\mathbf{d}^T] = (\mathbf{H}_{\text{nv}} - \mathbf{B})\mathbf{R}_x(\mathbf{H}_{\text{nv}} - \mathbf{B})^T$$

- Pick \mathbf{h}_i to minimize some metric of the error covariance \mathbf{R}_d
 - ⇒ $\text{trace}(\mathbf{R}_d) \Rightarrow$ **Mean squared error**
 - ⇒ $\lambda_{\max}(\mathbf{R}_d) \Rightarrow$ **worst-case** error

Approximate implementation with statistical knowledge of x

The **optimal filter coefficients** defined as $\{\mathbf{h}_{i,\text{Tr}}^*\}_{i=1}^N := \arg\min_{\{\mathbf{h}_i\}_{i=1}^N} \text{trace}(\mathbf{R}_d)$ and $\{\mathbf{h}_{i,\lambda}^*\}_{i=1}^N := \arg\min_{\{\mathbf{h}_i\}_{i=1}^N} \lambda_{\max}(\mathbf{R}_d)$ are, respectively, given by

$$\mathbf{h}_{i,\text{Tr}}^* = (\mathbf{R}_x^{1/2} \Phi_i)^{\dagger} \mathbf{R}_x^{1/2} \mathbf{b}_i \triangleq (\Phi_i^T \mathbf{R}_x \Phi_i)^{-1} \Phi_i^T \mathbf{R}_x \mathbf{b}_i,$$

for all i , where \triangleq holds if $\mathbf{R}_x^{1/2} \Phi_i$ has full column rank, and

$$\begin{aligned} \{\Gamma_\lambda^*, \mathbf{s}^*\} &= \arg\min_{\{\Gamma, \mathbf{s}\}} \mathbf{s} \\ \text{s. to } &\begin{bmatrix} \mathbf{s} \mathbf{I} & (\mathbf{I} \odot \Psi \Gamma)^T \tilde{\mathbf{U}} \mathbf{V}^{-1} - \mathbf{B} \\ ((\mathbf{I} \odot \Psi \Gamma)^T \tilde{\mathbf{U}} \mathbf{V}^{-1} - \mathbf{B})^T & \mathbf{s} \mathbf{R}_x^{-1} \end{bmatrix} \succeq 0. \end{aligned}$$

- **No prior knowledge** scenarios equivalent to the cases $\mathbf{R}_x = \mathbf{I}$
 - ⇒ $\|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_F \equiv \text{trace}(\mathbf{R}_d)$ and $\|\mathbf{H}_{\text{nv}} - \mathbf{B}\|_2 \equiv \lambda_{\max}(\mathbf{R}_d)$
- Additional assumptions can be incorporated into the designs

Finite-time consensus

- Local implementation of the consensus operator $\mathbf{B}_{\text{con}} = \mathbf{1}\mathbf{1}^T/N$

Corollary


The operator \mathbf{B}_{con} can be written as a filter $\sum_{l=0}^{N-1} c_l \mathbf{S}^l$ for some \mathbf{S} associated with an undirected graph G if and only if G is **connected**.

- If the **weights** in \mathbf{S} can be **selected** (e.g., choose $\mathbf{S} = \mathbf{L}$)
 - ⇒ consensus is achieved in finite time for connected graphs

- Key: \mathbf{B}_{con} is low-rank (repeated eigenvalues well-suited for approx.)

- We compare the performance of **three methods**: 1) Asymptotic fastest distributed linear averaging (FDLA), 2) **Graph filter** approx., 3) **Node-variant graph filter** approx.

- Compare **worst-case** and **mean error** design (50 nodes)

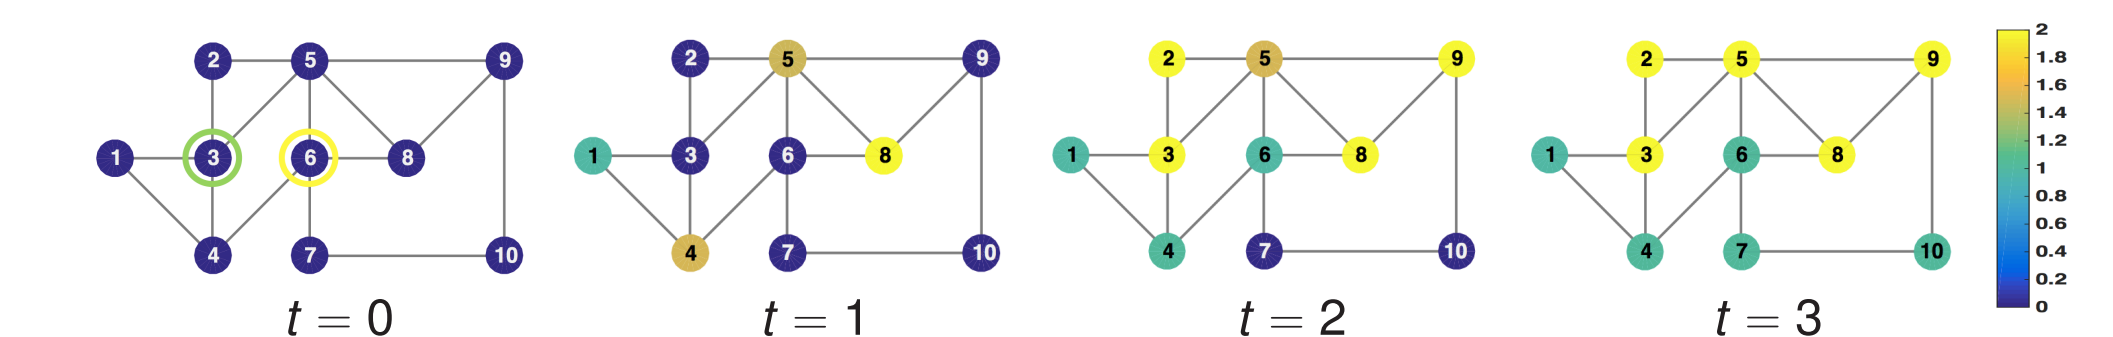


Analog network coding (ANC)

- In **ANC** we focus on the transmission from sources to sinks
 - ⇒ Leverage **graph filters** to design **ANC** schemes
 - ⇒ **Subset of the nodes** in a graph
 - ⇒ Main results **still hold** for the modified framework

- Consider a 10-node undirected graph
- Nodes **3** and **6** are sources $\mathbf{B}_{SR} = \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \end{bmatrix}^T$
- Node **3** transmits to 1, 4, 6, 7, and 10
- Node **6** transmits to the remaining ones
- Denote by g and w the signals injected by nodes **3** and **6**, respectively

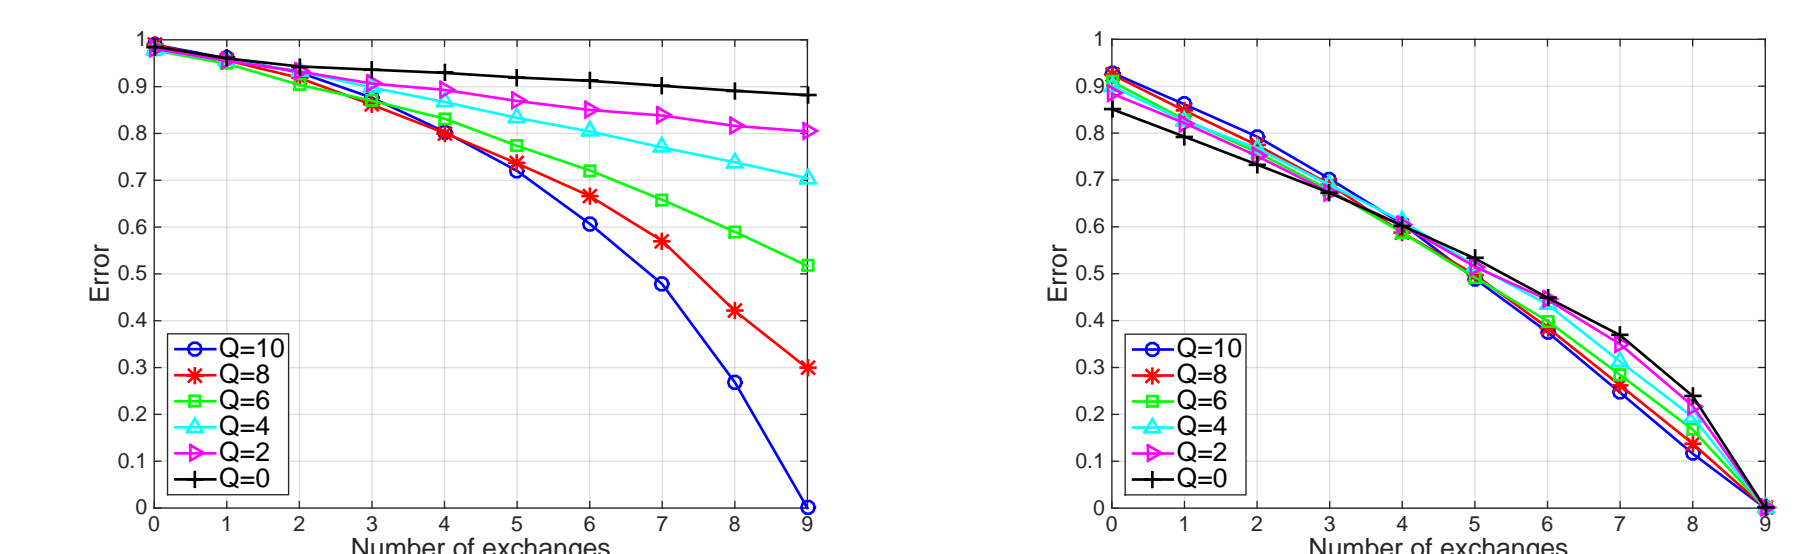
- Node colors represent the **successive best approximations**



node	recover signal	observed signal $\mathbf{z}^{(t)}$	Optimal coeff.
1	g	$t=0$ 0 $t=1$ g $t=2$ $g+w$ $t=3$ $5g+2w$	$[0, 1]$
3	w	$t=0$ 0 $t=1$ 0 $t=2$ $4g+2w$ $t=3$ $4g+3w$	$[-2, 0, 0.5]$
10	g	$t=0$ 0 $t=1$ 0 $t=2$ w $t=3$ $g+2w$	$[0, 0, -2, 1]$

Spectrum mismatch

- Assume that \mathbf{B} and \mathbf{S} share Q eigenvectors
 - ⇒ **Higher Q** , closer to being **simultaneously diagonalizable**
- Generate 1000 **Erdős-Rényi** graphs ($N=10$)
- Plot average errors for **node-invariant** and **node-variant** filters



- Performance of **node-invariant** filter is **sensitive** to changes in Q
- **Node-variant** filters are **robust** to the spectral differences of \mathbf{B} and \mathbf{S}

References

- D. I. Shuman et al., "The emerging field of signal processing on graphs," Signal Process. Mag., 2013.
- A. Sandryhaila and J. Moura, "Discrete signal processing on graphs," IEEE TSP, 2013.
- D. I. Shuman, P. Vandergheynst, and P. Frossard, "Chebyshev polynomial approximation for distributed signal processing," DCOSS, 2011.
- A. Sandryhaila, S. Kar, J. Moura, "Finite-time distributed consensus through graph filters," ICASSP, 2014.
- S. Barbarossa, G. Scutari, and T. Battisti, "Distributed signal subspace projection algorithms with maximum convergence rate for sensor networks with topological constraints," ICASSP, 2009.
- S. Segarra et al., "Distributed linear network operators using graph filters," arXiv, 2015.