

From local averaging to emergent global behaviors: The fundamental role of network interconnections

Original

From local averaging to emergent global behaviors: The fundamental role of network interconnections / Como, Giacomo; Fagnani, Fabio. - In: SYSTEMS & CONTROL LETTERS. - ISSN 0167-6911. - 95:(2016), pp. 70-76.
[10.1016/j.sysconle.2016.02.003]

Availability:

This version is available at: 11583/2647789 since: 2016-09-09T16:25:33Z

Publisher:

Elsevier

Published

DOI:10.1016/j.sysconle.2016.02.003

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

From local averaging to emergent global behaviors: the fundamental role of network interconnections

Giacomo Como and Fabio Fagnani

G. Como is with the Department of Automatic Control, Lund University, BOX118, SE-22100 Lund, Sweden

giacomo.como@control.lth.se.

F. Fagnani is with the Lagrange Department of Mathematical Sciences, Politecnico di Torino, Corso Stati Uniti 24, 10129, Torino, Italy, fabio.fagnani@polito.it.

Abstract

Distributed averaging is one of the simplest and most widely studied network dynamics. Its applications range from cooperative inference in sensor networks, to robot formation, to opinion dynamics. A number of fundamental results and examples scattered through the literature are gathered here and some original approaches and generalizations are presented, emphasizing the deep interplay between the network interconnection structure and the emergent global behavior.

Keywords: Network dynamics, interconnections, emergent behaviors, distributed averaging, consensus, polarization, electrical networks

1. Introduction

One of the core concepts in the behavioral approach to systems and control developed by Jan Willems in the '80s is that of interconnection [1]. Encompassing the traditional notion of feedback interconnection on which classical input/output control theory is based, the behavioral approach allows for defining interconnections of systems at a more primitive level, as intersections of solution sets of the evolution equations, without the need for specific flow diagrams. As Jan used to repeat, what is an input and what is an output is a matter of the application. This idea of going beyond the classical input/output formalism proved fruitful in applications, e.g., in coding theory, where Willems' study of minimal state space realizations [2] laid the foundations of trellis representations which are the basic tool for the design of efficient decoding algorithms.

More recently, the study of network dynamics is showing deep cultural analogies with the *ansatz* of the behavioral approach. Network dynamics entail a large number of (relatively) simple systems

coupled together along the architecture of a graph. The overall dynamical system can thus be seen as the interconnection of these atomic devices. It does not make much sense to classify *a priori* interconnection signals as input or outputs, rather they are variables coupling the systems, possibly sensor measurements, state positions, epidemic states, and it is often impossible to say who is influenced by whom. The emergence of global behaviors such as synchronization, information fusion, polarization, and diffusion is one of the distinctive features of these complex interconnected systems. Such global behaviors can in fact be seen as the result of the local interactions and of the interconnection graph structure.

This paper focuses on a particularly simple and well studied class of network dynamics: distributed averaging systems. [3, 4] These are linear network dynamics exhibiting many interesting collective behaviors, such as synchronization and transition phenomena. Their applications range from inferential sensor network algorithms [5], to network vehicle formation [6], to models for opinion dynamics [7]. Most of the behavioral approach developed by Jan was in fact focused on linear systems: he used to say that linear systems are sufficiently rich from a theoretical viewpoint and yet containing a huge va-

^{*}This work is dedicated to the memory of Jan C. Willems.

^{**}The first author is a member of the excellence centers LCCC and ELLIIT. His research has been supported by the Swedish Research Council through a junior research grant.

riety of applications. Keeping models as simple as possible was central in Jan’s approach to science.

Using classical results from the Perron-Frobenius theory of non-negative matrices, we first present an asymptotic analysis of the linear averaging dynamics on arbitrary interconnection graphs. As expected, the graph topology plays a crucial role in shaping the emergent global behavior. While it is well known that all states reach an asymptotic consensus on connected graphs, Theorem 2 of Section 2 analyzes the case of a general graph and shows that the asymptotic state of every agent in the network turns out to be a convex combination of the consensus reached by the sink connected components (i.e., components with no outgoing links). The weights of such convex combination have several useful interpretations. They can be seen as hitting probabilities of the dual Markov chain generated by the same averaging matrix or, when the graph is undirected, as voltages of an electrical circuit with suitable boundary conditions on the nodes belonging to the sink components, as explained in Section 3. While analogous electrical interpretations are well known in Markov chain theory [8, 9], they have received relatively minor attention in the distributed averaging literature, with a few exceptions. In particular, to our knowledge, Theorem 3 has not appeared elsewhere in this generality.

A relevant case in the applications is when the sink components all consist of single nodes—called stubborn nodes—that never change their state, e.g., playing the role of opinion leaders in social networks, or anchor nodes in robot formation control. The final part of the paper is dedicated to a deeper understanding of how the asymptotic state is distributed within the network in the presence of such stubborn nodes. It turns out that—depending on the stubborn nodes’ centrality and the graph connectivity—quite different phenomena can emerge ranging from polarized to homogeneous equilibrium configurations. [10] In the polarized case, nodes tend to cluster in subfamilies and converge to values very close to that of a particular stubborn agent, whereas in the homogeneous regime most of the nodes tend to get close to a consensus on a value which is a convex combination of the stubborn node values. In Section 4, we present these phenomena through an example where the transition between the two regimes can be analyzed in detail. We then recall more general results appeared in the literature.

We gather here some notational conventions.

The transpose of A is denoted by A' ; $\mathbf{1}$ is the all-1 vector; $\mathbf{1}_{\mathcal{A}}$ is the vector with all entries equal to 0 except for those whose label is in \mathcal{A} that are equal to 1. The asymptotic notation $a \ll b$ and $a \sim b$ means $\lim a/b = 0$ and $\lim a/b = 1$, respectively.

2. Averaging dynamics on general graphs

Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a directed weighted graph representing the network, where $\mathcal{V} = \{1, \dots, n\}$ is the set of nodes, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of links, and $W \in \mathbb{R}^{n \times n}$ is a matrix of nonnegative link weights such that $W_{ij} > 0$ if and only if $(i, j) \in \mathcal{E}$, with positive diagonal elements of W corresponding to self-loops. We refer to the graph \mathcal{G} as: *connected* if W is irreducible;¹ *undirected* if W is symmetric; *balanced* if $W\mathbf{1} = W'\mathbf{1}$; *unweighted* if $W_{ij} \in \{0, 1\}$ for all $i, j \in \mathcal{V}$. We denote the out-degree vector by $w = W\mathbf{1}$ and assume² that $w_i > 0$ for all nodes i . We then introduce the matrices

$$D = \text{diag}(w), \quad P = D^{-1}W, \quad L = D - W. \quad (1)$$

Observe that the matrices P and $-L$ are respectively row-stochastic and Metzler. Also, $P'w = w$ and $L'\mathbf{1} = 0$ if and only if \mathcal{G} is balanced. Moreover, \mathcal{G} being undirected is equivalent to the detailed balance $w_i P_{ij} = w_j P_{ji}$ for $i, j \in \mathcal{V}$, a property that is referred to as reversibility of P (with respect to w). The matrix L is known as the graph *Laplacian*.

One of the most popular network dynamics can be seen as the interconnection of local averaging systems, i.e., multi-input/single-state dynamics placed at the nodes $i \in \mathcal{V}$ and governed by the linear updates $x_i(t+1) = \alpha x_i(t) + (1-\alpha) \sum_j P_{ij} u_j(t)$. Here, $\alpha \in [0, 1]$ is an inertia parameter. By putting $u_j(t) = x_j(t)$ one obtains the interconnected system

$$x_i(t+1) = \alpha x_i(t) + (1-\alpha) \sum_j P_{ij} x_j(t), \quad (2)$$

for $i \in \mathcal{V}$. In (2), the sum index j runs in principle over the whole node set \mathcal{V} , but is in fact restricted to the out-neighborhood $\mathcal{N}_i := \{j : W_{ij} > 0\}$ of node i in \mathcal{G} . By assembling all the node states in

¹Note that this convention deviates from the one adopted by some authors who refer to \mathcal{G} as *strongly* connected if W is irreducible and simply connected if $W + W'$ is irreducible.

²This assumption implies no loss of generality since one can add a self-loop with $W_{ii} > 0$ to nodes i with $w_i = 0$ without modifying connectivity and other properties of \mathcal{G} .

a column vector $x(t) \in \mathbb{R}^n$, (2) can be compactly rewritten as

$$x(t+1) = P_\alpha x(t), \quad (3)$$

where $P_\alpha = \alpha I + (1 - \alpha)P$. Hence, the state vector $x(t)$ of the *distributed averaging* dynamics (3) evolves as $x(t) = P_\alpha^t x(0)$, so that its asymptotic behavior is dictated by the eigen-structure of P_α . Being a stochastic matrix, P is non-expansive in the $\|\cdot\|_\infty$ norm, so that its spectrum is contained in the unitary disk centered in 0. Hence, for $0 \leq \alpha \leq 1$, the matrix P_α has 1 as eigenvalue (corresponding to right eigenvector $\mathbf{1}$) and its whole spectrum is contained in the closed disk of diameter coinciding with the segment joining the points $-1 + 2\alpha$ and 1 in the complex plane. Finer properties of the spectrum of P_α are closely related to the geometrical properties of the graph \mathcal{G} as summarized below.

First we consider the case when the graph \mathcal{G} is connected. In this case, it is a standard result of the Perron-Frobenius theory that P_α^t converges to a matrix $\mathbf{1}\pi'$ where π can be uniquely characterized as the left eigenvector $\pi' = \pi'P$ such that $\mathbf{1}'\pi = 1$. Connectivity of \mathcal{G} implies that all the entries of π —which is referred to as the *centrality* vector—are strictly positive. For a balanced graph, π is proportional to the degree vector, namely, $\pi = w/(\mathbf{1}'w)$. For general, unbalanced, connected graphs such simple expression does not hold true, while one can express the entries π_i in terms of infinite sums. For $\alpha \in [0, 1)$, let the *mixing time* of P_α be

$$\tau_\alpha := \inf \left\{ t \geq 0 : \max_{i \in \mathcal{V}} \sum_j |(P_\alpha^t)_{ij} - \pi_j| \leq \frac{1}{2e} \right\}.$$

The mixing time is a popular index to study the speed of convergence of P_α^t . In certain cases it can be estimated from knowledge of the second largest eigenvalue of P_α or coupling techniques. E.g., for the unweighted d -dimensional toroidal grid, one has $\tau_\alpha \sim C_d n^{2/d}$ where C_d is a constant depending on the dimension d but not on the graph size n . For general large-scale graphs whose spectrum analysis is unfeasible and for which no effective couplings are known, it proves more convenient to relate the mixing time to the graph *conductance*

$$\Phi := \min_{\emptyset \neq \mathcal{U} \subsetneq \mathcal{V}} \frac{\sum_{i \in \mathcal{U}} \sum_{j \in \mathcal{V} \setminus \mathcal{U}} \pi_i P_{ij}}{\sum_{i \in \mathcal{U}} \pi_i \cdot \sum_{j \in \mathcal{V} \setminus \mathcal{U}} \pi_j},$$

that is a measure of the lack of bottlenecks in the graph. Results in [11, Section 4.3] imply that

$$\frac{1 - 2/e}{\Phi} \leq \tau_{1/2} \leq \frac{1}{\Phi^2} \log \frac{e^2}{\pi_*}, \quad (4)$$

where $\pi_* = \min_{i \in \mathcal{V}} \pi_i$. By combining the bounds above with estimates of the conductance, it can be shown, e.g., that the Erdos-Renyi random graphs in the connected regime³ exhibit, with probability 1, mixing times of the order of $\log n$. Such graphs are thus mixing faster than the d -dimensional tori.

These results imply that the state $x(t)$ of the averaging dynamics (3) on a connected graph converges to a consensus on a value $\bar{x} = \pi'x(0)$ that is the average of the nodes' initial values weighted by their centralities. The speed of this convergence is captured by the mixing time as in the following.

Proposition 1. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a connected graph. Then, for every $\alpha \in (0, 1)$, the distributed averaging dynamics (3) satisfy*

$$\|x(t) - \mathbf{1}\bar{x}\|_\infty \leq \|x(0) - \mathbf{1}\bar{x}\|_\infty \exp(-[t/\tau_\alpha]), \quad (5)$$

for all $x(0) \in \mathbb{R}^n$, where

$$\bar{x} = \pi'x(0), \quad (6)$$

so that, in particular, $\lim_{t \rightarrow +\infty} x(t) = \mathbf{1}\bar{x}$.

Proof As a consequence of the non-expansivity of P' in l_1 -distance we get that

$$\sum_j |(P_\alpha^t)_{ij} - \pi_j| \leq \exp(-[t/\tau_\alpha]), \quad i \in \mathcal{V}.$$

(Cf., e.g., [9, Eqn. (4.34)].) Then, for all $i \in \mathcal{V}$,

$$\begin{aligned} |x_i(t) - \bar{x}| &= \left| \sum_j ((P_\alpha^t)_{ij} - \pi_j)(x_j(0) - \bar{x}) \right| \\ &\leq \left| \sum_j (P_\alpha^t)_{ij} - \pi_j \right| \cdot |x_j(0) - \bar{x}| \\ &\leq \exp(-[t/\tau_\alpha]) \cdot \|x(0) - \bar{x}\|_\infty, \end{aligned}$$

which gives the claim. \blacksquare

We now move on to discussing the asymptotic behavior of the distributed averaging dynamics (3) in arbitrary—not necessarily connected—graphs. The node set \mathcal{V} can always be uniquely partitioned as $\mathcal{V} = \mathcal{V}_1 \cup \dots \cup \mathcal{V}_c$ where, for $1 \leq k \leq c$, the subgraph $\mathcal{G}_k = (\mathcal{V}_k, \mathcal{E} \cap (\mathcal{V}_k \times \mathcal{V}_k), W|_{\mathcal{V}_k \times \mathcal{V}_k})$ is connected and maximal with respect to this property. Such subgraphs are called *connected components*. For any two of them, \mathcal{G}_h and \mathcal{G}_k , we write $\mathcal{G}_h \geq \mathcal{G}_k$ if there is a path in \mathcal{G} connecting some node in \mathcal{V}_k to some node in \mathcal{V}_h . By construction the relation is transitive and such that $\mathcal{G}_h \geq \mathcal{G}_k$ and $\mathcal{G}_k \geq \mathcal{G}_h$ if

³They are constructed by considering n nodes randomly putting a link between any pair of them independently with probability $p = c \log n/n$ for $c > 1$.

and only if $\mathcal{G}_k = \mathcal{G}_h$. Connected components which are maximal with respect to the partial ordering \geq are called *sink components*: any path starting in a sink component will never leave it. The following result characterizes the asymptotic behavior of the averaging dynamics (3) on a general graph.

Theorem 2. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a graph with sink components $\mathcal{G}_k = (\mathcal{S}_k, \mathcal{E} \cap (\mathcal{S}_k \times \mathcal{S}_k), W|_{\mathcal{S}_k \times \mathcal{S}_k})$, for $k = 1, \dots, s$. Let $\mathcal{S} := \bigcup_{1 \leq k \leq s} \mathcal{S}_k$ and $\mathcal{S}_{-k} := \mathcal{S} \setminus \mathcal{S}_k$. Then:*

- (i) *There exists a unique nonnegative $H \in \mathbb{R}^{n \times s}$ such that $H\mathbf{1} = \mathbf{1}$ and, for $i \in \mathcal{S}$, $1 \leq k \leq s$,*

$$LH = 0, \quad H_{ik} = \begin{cases} 1 & i \in \mathcal{S}_k \\ 0 & i \in \mathcal{S}_{-k}; \end{cases} \quad (7)$$

- (ii) *For every $\alpha \in (0, 1)$, the distributed averaging dynamics (3) satisfies*

$$\lim_{t \rightarrow +\infty} x(t) = H\bar{x}, \quad (8)$$

for all $x(0) \in \mathbb{R}^n$, where $\bar{x} \in \mathbb{R}^s$ has entries

$$\bar{x}_k = \sum_{i \in \mathcal{S}_k} \pi_i^{(k)} x_i(0), \quad k = 1, \dots, s, \quad (9)$$

and $\pi^{(k)} \in \mathbb{R}^{\mathcal{S}_k}$ is the centrality vector of \mathcal{G}_k .

Proof Upon reordering nodes in such a way that sink components come last, P_α takes the form

$$P_\alpha = \begin{pmatrix} Q & R^{(1)} & \dots & \dots & R^{(s)} \\ 0 & P^{(1)} & 0 & \dots & 0 \\ 0 & 0 & P^{(2)} & & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & P^{(s)} \end{pmatrix}.$$

Splitting $x(t) = (y(t), x^{(1)}(t), \dots, x^{(s)}(t))$ accordingly, with $y(t) \in \mathbb{R}^{\mathcal{R}}$ where $\mathcal{R} := \mathcal{V} \setminus \mathcal{S}$, and $x^{(k)}(t) \in \mathbb{R}^{\mathcal{S}_k}$ for $k = 1, \dots, s$, recursion (3) reads

$$\begin{aligned} y(t+1) &= Qy(t) + \sum_{1 \leq k \leq s} R^{(k)} x^{(k)}(t) \\ x^{(k)}(t+1) &= P^{(k)} x^{(k)}(t), \quad 1 \leq k \leq s. \end{aligned} \quad (10)$$

First, notice that the evolution of the state on the nodes of the sink components can be studied using Proposition 1. In particular, for $k = 1, \dots, s$, we have that $x^{(k)}(t) \xrightarrow{t \rightarrow \infty} \mathbf{1}\bar{x}_k$, where \bar{x}_k is as in (9) and $\pi^{(k)} \in \mathbb{R}^{\mathcal{S}_k}$ is the centrality vector of \mathcal{G}_k .

On the other hand, Q is nonnegative, and the Perron-Frobenius theory implies that its spectral

radius ρ is an eigenvalue with associated nonnegative left eigenvector $z \in \mathbb{R}^{\mathcal{R}}$. Let $\mathcal{J} \subseteq \mathcal{R}$ be the support of z . Since every node in \mathcal{R} is connected to some node in \mathcal{S} , $\min_{j \in \mathcal{J}} \sum_{i \in \mathcal{J}} Q_{ji} < 1$ (otherwise there would be no links from \mathcal{J} to $\mathcal{V} \setminus \mathcal{J}$). Hence,

$$\rho \sum_{i \in \mathcal{J}} z_i = \sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{R}} Q_{ji} z_j = \sum_{i \in \mathcal{J}} \sum_{j \in \mathcal{J}} Q_{ji} z_j < \sum_{j \in \mathcal{J}} z_j,$$

i.e., $\rho < 1$. Thus, Q is stable and $I - Q$ is invertible. On the other hand, $\lim_{t \rightarrow \infty} x^{(k)}(t) = \mathbf{1}\bar{x}_k$ for $1 \leq k \leq s$ by Proposition 1. It follows that the first line of (10) is a stable LTI system with converging input, so that its state is necessarily converging to

$$\lim_{t \rightarrow +\infty} y(t) = (I - Q)^{-1} \sum_{1 \leq k \leq s} R^{(k)} \mathbf{1}\bar{x}_k.$$

This yields (8), with $H \in \mathbb{R}^{n \times s}$ defined by

$$H_{ik} = \begin{cases} ((I - Q)^{-1} R^{(k)} \mathbf{1})_i & i \in \mathcal{R} \\ 1 & i \in \mathcal{S}_k \\ 0 & i \in \mathcal{S}_{-k}. \end{cases} \quad (11)$$

Note that H is nonnegative since both $R^{(k)}$, for $1 \leq k \leq s$, and $(I - Q)^{-1} = \sum_{l \geq 0} Q^l$ are. Moreover, $P_\alpha \mathbf{1} = \mathbf{1}$ implies $\sum_{1 \leq k \leq s} R^{(k)} \mathbf{1} = (I - Q) \mathbf{1}$, so that

$$\sum_{1 \leq k \leq s} H_{ik} = ((I - Q)^{-1} \sum_{1 \leq k \leq s} R^{(k)} \mathbf{1})_i = 1, \quad i \in \mathcal{R}.$$

Hence, $H\mathbf{1} = \mathbf{1}$. Furthermore, one has that

$$\begin{aligned} LH &= D(I - P)H = \frac{1}{1 - \alpha} D(I - P_\alpha) \\ &= \frac{1}{1 - \alpha} D(I - Q)(I - Q)^{-1} \sum_{1 \leq k \leq s} R^{(k)} \mathbf{1}\mathbf{1}'_{\{k\}} \\ &\quad - \frac{1}{1 - \alpha} D \sum_{1 \leq k \leq s} R^{(k)} \mathbf{1}\mathbf{1}'_{\{k\}} = 0, \end{aligned}$$

proving (7). Uniqueness of the solution of (7) follows from invertibility of the $\mathcal{R} \times \mathcal{R}$ block of L . ■

Theorem 2 states that the nodes belonging to a sink component \mathcal{S}_k asymptotically reach consensus on the value $\bar{x}_k = \sum_{i \in \mathcal{S}_k} \pi_i^{(k)} x_i(0)$. The state of every other node $i \in \mathcal{R}$ converges to a convex combination of the consensus values of the various sink components with weights H_{ik} characterized by (7). The initial states $x_i(0)$ of the nodes $i \in \mathcal{R}$ have thus no influence on the equilibrium state x . Equivalently, the equilibrium state x of the averaging dynamics (3) can be characterized as the solution of

$$Lx = 0, \quad x_j = \bar{x}_k \quad j \in \mathcal{S}_k, \quad 1 \leq k \leq s, \quad (12)$$

which is referred to as the *Laplace equation* on \mathcal{G} with *boundary conditions* on $\mathcal{S} = \mathcal{S}_1 \cup \dots \cup \mathcal{S}_k$.

Two special cases are worth being examined:

- If there is a single sink component, $H = \mathbf{1}$, thus Theorem 2 implies that $P_\alpha^t x(0) \xrightarrow{t \rightarrow \infty} \mathbf{1} \pi' x(0)$. I.e., the system converges to consensus on a convex combination of the initial states of the nodes of the sink component, while all other nodes' initial states do not play any role. This is a generalization of Proposition 1.
- If the sink components are all singletons, i.e., $\mathcal{S}_k = \{v_k\}$ for $1 \leq k \leq s$, then $x(t) \xrightarrow{t \rightarrow \infty} H\bar{x}$, where $\bar{x}_k = x_{v_k}(0)$ for $1 \leq k \leq s$. In this case, sink nodes keep their state constant in time. They are sometimes referred to as *stubborn nodes* and are used to model opinion leaders in social networks [10], or anchor nodes in robotic formation models [12].

An interesting probabilistic interpretation comes from considering a discrete-time Markov chain $X(t)$ with state space \mathcal{V} and transition probability matrix P_α . I.e., $X(t+1)$ is conditionally independent from the past $X(0), \dots, X(t-1)$ given the present $X(t)$, and $\mathbb{P}(X(t+1) = j | X(t) = i) = (P_\alpha)_{ij}$. It is well known that, with probability one, $X(t)$ will enter one of the sink components in finite time and never leave it ever after. For $k = 1, \dots, s$, let A_k be the event that $X(t)$ enters the sink component \mathcal{S}_k before any other. Consider the matrix $M \in \mathbb{R}^{n \times s}$ with entries $M_{ik} = \mathbb{P}(A_k | X(0) = i)$. Then, a simple conditioning argument yields

$$\begin{aligned} M_{ik} &= \mathbb{P}(A_k | X(0) = i) \\ &= \sum_{l \in \mathcal{V}} (P_\alpha)_{il} \mathbb{P}(A_k | X(1) = l) = (P_\alpha M)_{ik} \end{aligned}$$

for all $i \in \mathcal{V}$ and $1 \leq k \leq s$. Hence, M solves (7) and we can deduce that $H = M$. In other terms, the weight H_{ik} that agent i puts on \bar{x}_k in determining its asymptotic state x_i in (3) can be interpreted as the probability that a Markov chain started at node i and moving with transition probability matrix P_α hits the sink component \mathcal{S}_k before any other.

3. The electrical network interpretation

Consider a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ with sink components $\mathcal{S}_1, \dots, \mathcal{S}_s$, where $s \geq 2$. Put $\mathcal{S} := \bigcup_{1 \leq k \leq s} \mathcal{S}_k$, $\mathcal{R} := \mathcal{V} \setminus \mathcal{S}$, and $\mathcal{S}_{-k} := \mathcal{S} \setminus \mathcal{S}_k$ for every $1 \leq k \leq s$.

Then, Theorem 2 guarantees that the state of the averaging dynamics (3) converges to an equilibrium $x = H\bar{x}$, where $H \in \mathbb{R}^{n \times s}$ is the stochastic matrix satisfying (7) and $\bar{x} \in \mathbb{R}^s$ is the vector of the weighted averages of the initial condition in the sink components. In this section, we focus on the special case when the restriction of the graph \mathcal{G} to the node set \mathcal{R} is connected and undirected, i.e., when

$$W_{ij} = W_{ji}, \quad i, j \in \mathcal{R}. \quad (13)$$

We will first interpret the equilibrium states of (3) as voltages in an electrical network associated to \mathcal{G} , then relate them to the effective resistances in the network. Let us start by defining the link flows

$$f_{ij} = W_{ij}(x_i - x_j), \quad i, j \in \mathcal{V}. \quad (14)$$

The key consequence of (13) is that then $x = H\bar{x}$ and (7) imply the following conservation law:

$$\sum_j f_{ij} = 0, \quad \forall i \in \mathcal{R}. \quad (15)$$

Indeed, one can give the following interpretation: the link weights W_{ij} represent conductances and their inverses are resistances; x_i is the voltage in node i ; and f_{ij} is the electrical current flowing from node i to node j . Then, (14) and (15) can be read as the Ohm law and, respectively, the Kirchoff law.

Such interpretation has deep implications. First, by simply verifying first-order conditions, one can show that (13) implies that, for $k = 1, \dots, s$, the k -th column of H coincides with the solution of the following quadratic optimization problem

$$\frac{1}{R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}} = \min_{\substack{y \in \mathbb{R}^n : \\ y_i = 1 \quad i \in \mathcal{S}_k \\ y_i = 0 \quad i \in \mathcal{S}_{-k}}} \frac{1}{2} \sum_{i, j \in \mathcal{V}} W_{ij} (y_i - y_j)^2. \quad (16)$$

The quantity that is to be minimized in the right-hand side of (16) represents the energy dissipation in the network when the voltages are y_i . Hence, $H\mathbf{1}_{\{k\}}$ is the vector of voltages with minimal energy dissipation under the constraints that the voltage is 1 in \mathcal{S}_k and 0 in $\mathcal{S} \setminus \mathcal{S}_k$. The inverse $R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}$ of such minimal energy dissipation is known as the *effective resistance* between the node sets \mathcal{S}_k and \mathcal{S}_{-k} . A classical duality result known as Thompson's principle [9, Th. 9.10] states that $R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}$ coincides with the minimal energy dissipation of a

unitary flow from \mathcal{S}_k to \mathcal{S}_{-k}

$$R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}} = \min_{\substack{\theta \in \mathbb{R}^{n \times n} : \\ (\theta \mathbf{1})' \mathbf{1}_{\mathcal{S}_k} = 1 \\ (\theta \mathbf{1})_i = 0 \quad i \in \mathcal{R}}} \frac{1}{2} \sum_{i,j \in \mathcal{V}} \frac{1}{W_{ij}} \theta_{ij}^2, \quad (17)$$

and that the minimum above is achieved by

$$\theta_{ij} = W_{ij}(H_{ik} - H_{jk})R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}. \quad (18)$$

In particular, the fact that θ as defined in (18) is a unitary flow from \mathcal{S}_k to \mathcal{S}_{-k} (i.e., it satisfies the constraints in the righthand side of (17)) implies that its normalized version $R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}^{-1} \theta$ satisfies the following flow conservation equations

$$\begin{aligned} \frac{1}{R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}} &= \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{S}_k} W_{ij}(1 - H_{ik}) && \begin{pmatrix} \text{flow} \\ \text{out of } \mathcal{S}_k \end{pmatrix} \\ &= \sum_{i \in \mathcal{U}} \sum_{j \in \mathcal{W}} W_{ij}(H_{ik} - H_{jk}) && \begin{pmatrix} \text{flow} \\ \text{from } \mathcal{U} \\ \text{to } \mathcal{W} \end{pmatrix} \\ &= \sum_{i \in \mathcal{V}} \sum_{j \in \mathcal{S}_{-k}} W_{ij}H_{ik}, && \begin{pmatrix} \text{flow} \\ \text{into } \mathcal{S}_{-k} \end{pmatrix} \end{aligned} \quad (19)$$

where $\mathcal{V} = \mathcal{U} \cup \mathcal{W}$, $\mathcal{U} \cap \mathcal{W} = \emptyset$ is any cut of \mathcal{G} such that $\mathcal{S}_k \subseteq \mathcal{U}$ and $\mathcal{S}_{-k} \subseteq \mathcal{W}$.

One key advantage of the electrical network interpretation is that the relations (16) and (19) imply simple and powerful rules, e.g., the parallel and series laws [9, pp. 135–136], to compute or estimate the equilibrium state x . Another fundamental property is monotonicity with respect to changes of the network (known as Rayleigh's law): the effective resistance is never increased when new links are added, or when the conductance, i.e., the weight, of some existing links is increased—including when two nodes in \mathcal{R} are glued together, that is equivalent to the addition of an infinite weight link between them. [9, Theorem 9.12] In fact, the equilibrium states x_i can be expressed purely in terms of effective resistances as stated in the following.

Theorem 3. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a graph with sink components $\mathcal{S}_1, \dots, \mathcal{S}_s$, $s \geq 2$, whose restriction to $\mathcal{R} = \mathcal{V} \setminus \mathcal{S}$, where $\mathcal{S} = \bigcup_{1 \leq k \leq s} \mathcal{S}_k$, is connected and undirected (as per (13)). Then, the equilibrium state vector x of (3) satisfies, for $i \in \mathcal{V}$,*

$$x_i = \frac{1}{2} \sum_{1 \leq k \leq s} \bar{x}_k \left(1 + \frac{R_{i \leftrightarrow \mathcal{S}_{-k}} - R_{i \leftrightarrow \mathcal{S}_k}}{R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}} \right), \quad (20)$$

where $R_{i \leftrightarrow \mathcal{S}_k}$ and $R_{i \leftrightarrow \mathcal{S}_{-k}}$ are the effective resistances between node i and \mathcal{S}_k and \mathcal{S}_{-k} , respectively,

in the graph obtained from \mathcal{G} by glueing together all nodes in \mathcal{S}_k and \mathcal{S}_{-k} respectively into single nodes.

Proof Let $\hat{\mathcal{G}} = (\hat{\mathcal{V}}, \hat{\mathcal{E}}, \hat{W})$ be the connected undirected graph obtained from \mathcal{G} by merging all nodes in \mathcal{S}_k and \mathcal{S}_{-k} into single nodes v and \bar{v} , respectively, and making the links incident in v and \bar{v} bidirectional. Let its Laplacian $\hat{L} = \text{diag}(\hat{W}\mathbf{1}) - \hat{W}$ have eigenvalues $0 = \lambda_1 < \lambda_2 \leq \dots \leq \lambda_n$ and corresponding orthonormal base of eigenvectors $\frac{1}{\sqrt{n}}\mathbf{1} = \phi_{(1)}, \phi_{(2)}, \dots, \phi_{(n)}$. Define the Green matrix

$$G = G' = \sum_{2 \leq l \leq n} \frac{1}{\lambda_l} \phi_{(l)} \phi_{(l)}'$$

and observe that $G\mathbf{1} = 0$ and

$$LG = \sum_{2 \leq l \leq n} \frac{1}{\lambda_l} L\phi_{(l)} \phi_{(l)}' = \sum_{2 \leq l \leq n} \phi_{(l)} \phi_{(l)}' = I - \frac{1}{n} \mathbf{1}\mathbf{1}'.$$

From the above and $L\mathbf{1} = 0$ it can be deduced that, for $h, j \in \hat{\mathcal{V}}$, all solutions y of $Ly = (\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}})$ can be written as $y = G(\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}}) + \alpha\mathbf{1}$ for some scalar α . Now, let \hat{R}_{hj} be the effective resistance between h and j in $\hat{\mathcal{G}}$ and $x^{(h,j)}$ be the solution of $(Lx^{(h,j)})_i = 0$ for $i \neq h, j$, $x_h^{(h,j)} = 1$, $x_j^{(h,j)} = 0$. Arguing as in (19) gives $Lx^{(h,j)} = \hat{R}_{hj}^{-1}(\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}})$, from which we deduce that

$$\hat{R}_{hj}x^{(h,j)} = G(\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}}) + \alpha\mathbf{1}, \quad (21)$$

for some scalar α . It follows from (21) that

$$\begin{aligned} \hat{R}_{hj} &= \hat{R}_{hj}(x_h^{(h,j)} - x_j^{(h,j)}) \\ &= (\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}})' G(\mathbf{1}_{\{h\}} - \mathbf{1}_{\{j\}}) \\ &= G_{hh} - G_{hj} - G_{jh} + G_{jj} \\ &= G_{hh} - 2G_{hj} + G_{jj}. \end{aligned} \quad (22)$$

By applying (21) with $h = v$ and $j = \bar{v}$, and (22) twice, with $h = l$ and $j = v$ first, and then with $h = l$ and $j = \bar{v}$, one gets, for $l \in \hat{\mathcal{V}}$,

$$\begin{aligned} 2\hat{R}_{v\bar{v}}x_l^{(v,\bar{v})} &= 2(G_{lv} - G_{l\bar{v}} + \alpha) \\ &= G_{vv} - \hat{R}_{lv} - G_{v\bar{v}} + \hat{R}_{l\bar{v}} + 2\alpha. \end{aligned} \quad (23)$$

Choosing $l = \bar{v}$ and recalling that $x_{\bar{v}}^{(v,\bar{v})} = 0 = \hat{R}_{v\bar{v}}$ gives $2\alpha = \hat{R}_{v\bar{v}} + G_{v\bar{v}} - G_{vv}$ in (23). Substituting back in (23) and noting that $x_i^{(v,\bar{v})} = H_{ik}$, one gets

$$H_{ik} = \frac{R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}} + R_{i \leftrightarrow \mathcal{S}_{-k}} - R_{i \leftrightarrow \mathcal{S}_k}}{2R_{\mathcal{S}_k \leftrightarrow \mathcal{S}_{-k}}}, \quad (24)$$

for $i \in \mathcal{V}$ and $1 \leq k \leq s$. The claim now follows by substituting (24) into (8). We observe that alternative proofs of (24) have been proposed based on Markov chain arguments (cf. [9, Ex. 10.8]). ■

An insightful special case of Theorem 3 is when \mathcal{G} has two sink components \mathcal{S}^+ and \mathcal{S}^- with values $\bar{x}^+ = 1$ and $\bar{x}^- = 0$, respectively. Then, (20) reads

$$x_i = \frac{1}{2} + \frac{R_{i \leftrightarrow \mathcal{S}^-} - R_{i \leftrightarrow \mathcal{S}^+}}{2R_{\mathcal{S}^+ \leftrightarrow \mathcal{S}^-}}, \quad i \in \mathcal{V}.$$

The sign of the difference between the two effective resistances determines if node i will be more influenced by \mathcal{S}^+ or \mathcal{S}^- . In this sense, formula (20) expresses the bias of a node towards a sink component as determined by the electrical resistance to that sink in comparison to the others. Such electrical equivalence has recently found applications in the design of efficient distributed algorithms for the optimal stubborn node placement problem. [13]

4. Polarization and homogeneous influence

In this section, we consider graphs with $s \geq 2$ singleton sink components $\mathcal{S}_1 = \{v_1\}, \dots, \mathcal{S}_s = \{v_s\}$ with values $\bar{x}_1, \dots, \bar{x}_s \in [0, 1]$. We investigate conditions —on the graph structure and on the size of the sink components— under which most of the entries x_i of the equilibrium state $x = H\bar{x}$ of the averaging dynamics (3) are close to a common value \tilde{x} which is a convex combination of the \bar{x}_j s, or rather they are all close to one of the extreme values $\bar{x}_1, \dots, \bar{x}_s$. In order to formalize these notions, we consider infinite sequences of graphs (typically of increasing size), and briefly refer to them as (large-scale) networks. Following [10], we say that the sink components $\mathcal{S} = \{v_1, \dots, v_s\}$ have *homogeneous influence* on (the rest of) the network if, for all $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \inf_{\bar{x}} \frac{1}{n} |\{i \in \mathcal{V} : |x_i - \tilde{x}| < \varepsilon\}| = 1. \quad (25)$$

On the other hand, we refer to a network as *polarized* if, for all $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \left| \left\{ i \in \mathcal{V} : \min_{1 \leq k \leq s} |x_i - \bar{x}_k| < \varepsilon \right\} \right| = 1. \quad (26)$$

As in the previous section, we confine our discussion to the special case when the restriction of the graph \mathcal{G} to $\mathcal{R} = \mathcal{V} \setminus \mathcal{S}$ is connected and undirected. In fact, for the sake of simplicity, we mostly focus on an even more special graph structure $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ whose node set \mathcal{V} consists of only two stubborn nodes, v_0 and v_1 , and two disjoint sets of regular nodes, \mathcal{U}_0 and \mathcal{U}_1 , such that: the nodes in \mathcal{U}_0 (respectively, \mathcal{U}_1) are all connected by

a weight- γ directed link to v_0 (v_1); the subnetwork obtained by removing the stubborn nodes from \mathcal{G} is undirected and connected; the aggregate weight $\sum_{j \in \mathcal{U}_1} W_{ij}$ (respectively, $\sum_{j \in \mathcal{U}_0} W_{ij}$) of links connecting a node $i \in \mathcal{U}_0$ ($i \in \mathcal{U}_1$) to nodes in \mathcal{U}_1 (\mathcal{U}_0) is a positive constant β_0 (β_1) independent of i . In other words, we consider a network whose weight matrix W has the structure

$$W = \begin{bmatrix} \gamma & 0 & \dots & \dots & 0 & 0 \\ \vdots & & & & & \vdots \\ \gamma & & A & B & & 0 \\ 0 & & & & & \gamma \\ \vdots & & C & D & & \vdots \\ 0 & 0 & \dots & \dots & 0 & \gamma \end{bmatrix} \quad \begin{aligned} A &= A' \\ B &= C' \\ C &= B' \\ D &= D' \\ B\mathbf{1} &= \beta_0\mathbf{1} \\ C\mathbf{1} &= \beta_1\mathbf{1} \end{aligned} \quad (27)$$

Proposition 4. *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a graph with weight matrix W as in (27). Let the stubborn nodes be assigned values $\bar{x}_{v_0} = 0$ and $\bar{x}_{v_1} = 1$. Let $x = H\bar{x}$ be the equilibrium state of the averaging dynamics (3) and $y_h := \frac{1}{n_h} \sum_{i \in \mathcal{U}_h} x_i$ for $h = 0, 1$ be the average states in the two subsets of nodes. Then,*

$$\begin{aligned} |h - y_h| &\leq \left(1 + \frac{n_h}{n_{1-h}} + \frac{\gamma}{\beta_h} \right)^{-1}, \quad h = 0, 1, \\ y_1 - y_0 &\leq \left(1 + \frac{\beta_0}{\gamma} + \frac{\beta_1}{\gamma} \right)^{-1}. \end{aligned}$$

Proof Since the restriction of the graph to \mathcal{R} is undirected, we can use the electrical circuit interpretation of Section 3. In particular, (19) yields

$$\begin{aligned} \frac{1}{R_{v_0 \leftrightarrow v_1}} &= \gamma \sum_{i \in \mathcal{U}_0} x_i = \sum_{i \in \mathcal{U}_0} \sum_{j \in \mathcal{U}_1} W_{ij} (x_j - x_i) \\ &= \gamma \sum_{j \in \mathcal{U}_1} (1 - x_j). \end{aligned} \quad (28)$$

Moreover, we can get a lower bound on the effective resistance by merging, for $h = 0, 1$, all nodes in \mathcal{U}_h into a single node u_h and applying the parallel and series law to the resulting network, thus getting

$$R_{v_0 \leftrightarrow v_1} \geq \frac{1}{\gamma n_0} + \frac{1}{n_0 \beta_0} + \frac{1}{\gamma n_1} = \frac{1}{\gamma n_0} + \frac{1}{n_1 \beta_1} + \frac{1}{\gamma n_1}.$$

Then, the claim follows by substituting the identity (28) in the lefthand side of the above. \blacksquare

It follows from Proposition 4 that, for a network structure as in (27),

- (i) if $\gamma \gg \max\{\beta_0, \beta_1\}$, then $y_0 \rightarrow 0$ and $y_1 \rightarrow 1$;
- (ii) if $\gamma \ll \max\{\beta_0, \beta_1\}$, then $y_1 - y_0 \rightarrow 0$.

Observe that, since, for all $\varepsilon > 0$ and $h = 0, 1$,

$$\frac{1}{n} |\{i \in \mathcal{U}_h : |x_i - h| \geq \varepsilon\}| \leq \frac{|h - y_h|}{\varepsilon},$$

point (i) above implies that the network is polarized if $\gamma \gg \max\{\beta_0, \beta_1\}$. The intuition behind this result is that, if $\gamma \gg \max\{\beta_0, \beta_1\}$, then, in both \mathcal{U}_0 and \mathcal{U}_1 , the total weight of links towards the stubborn node v_0 (respectively, v_1) is much larger than the total weight of links to the other set of nodes. Then, for $h = 0, 1$, nodes in \mathcal{U}_h get influenced by v_h much more than by the other set of nodes \mathcal{U}_{1-h} .

On the other hand, observe that $y_1 - y_0 \rightarrow 0$ does not imply homogeneous influence. Sufficient conditions for homogeneous influence have been proved in [10] based on finer properties of the graph \mathcal{G} , in particular on its mixing time, as per the following.

Theorem 5 ([10]). *Let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$ be a graph with sink components $\mathcal{S}_1, \dots, \mathcal{S}_s$, let $\mathcal{S} = \bigcup_{1 \leq j \leq s} \mathcal{S}_j$ and $\mathcal{R} = \mathcal{V} \setminus \mathcal{S}$. Let $\tilde{\mathcal{G}} = (\mathcal{V}, \tilde{\mathcal{E}}, \tilde{W})$ be the undirected graph obtained from \mathcal{G} by making all the directed links from some node in \mathcal{R} to some node in \mathcal{S} bidirectional and letting the modified weight matrix \tilde{W} coincide with W in its $\mathcal{R} \times \mathcal{V}$ block, and be such that $\tilde{W}_{ij} = W_{ji}$ for all $i \in \mathcal{S}$ and $j \in \mathcal{R}$, and $\tilde{W}_{ij} = 0$ for all $i, j \in \mathcal{S}$. Assume that $\tilde{\mathcal{G}}$ is connected, and let $\tilde{P} = \text{diag}(\tilde{W}\mathbf{1})^{-1}\tilde{W}$, $\tilde{\pi} = \tilde{P}'\tilde{\pi}$ be its invariant probability vector, and $\tilde{\tau}$ be the mixing time of $\frac{1}{2}(I + \tilde{P})$. Then, the equilibrium state x of (3) satisfies*

$$\frac{1}{n} |\{i \in \mathcal{V} : |x_i - \tilde{x}| \geq \varepsilon\}| \leq \frac{\Delta}{\varepsilon n \tilde{\pi}_*} \psi(\tilde{\tau} \cdot \tilde{\pi}_S),$$

for $\varepsilon > 0$, where $\tilde{x} = \tilde{\pi}'x$, $\Delta = \max_{1 \leq i, j \leq s} \{\bar{x}_i - \bar{x}_j\}$, $\tilde{\pi}_* := \min_{i \in \mathcal{V}} \tilde{\pi}_i$, and $\psi(y) := y \log(e^2/y)$.

Theorem 5 implies that influence is homogenous in networks such that the product $\tilde{\tau} \cdot \tilde{\pi}_S$ of the mixing time and the aggregate centrality of the set of stubborn nodes is vanishing (and $n\tilde{\pi}_*$ is bounded away from 0). Networks such that $\tilde{\tau} \cdot \tilde{\pi}_S \rightarrow 0$ have been referred to as *highly fluid* [10]. Examples of highly fluid networks include d -dimensional tori with $d \geq 3$ when $|\mathcal{S}| \ll n^{1-2/d}$, and expansive networks such as the Erdos-Renyi graph provided that $|\mathcal{S}| \ll n/\log n$.

We conclude this section with an application of Proposition 4 and Theorem 5, highlighting a threshold phenomenon, with a transition from polarization to homogeneous influence. Let $\mathcal{G}_0 = (\mathcal{U}_0, \mathcal{E}_0)$ and $\mathcal{G}_1 = (\mathcal{U}_1, \mathcal{E}_1)$ be two independent and identically distributed Erdos-Renyi random graphs with




Figure 1: A graph \mathcal{G} obtained by matching two independent and identically distributed Erdos-Renyi graphs (black links) with node sets \mathcal{U}_0 and \mathcal{U}_1 by weight- β links (light grey horizontal links) and connecting, for $h = 0, 1$, each node in \mathcal{U}_h with a stubborn node v_h by a weight- γ link (dark grey).

parameters $|\mathcal{U}_0| = |\mathcal{U}_1| = m$ and $p = \omega m^{-1} \log m$, where $\omega > 1$ is a constant independent of m . Hence, distinct pairs of nodes $\{i, j\} \subseteq \mathcal{U}_h$, $h = 0, 1$, are connected by weight-1 undirected links independently with probability p . The scaling $pm/\log m = \omega > 1$ guarantees that \mathcal{G}_0 and \mathcal{G}_1 are connected with high probability as m grows large. [14, Ch. 6] Then, let $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$, where $\mathcal{V} = \{v_0\} \cup \mathcal{U}_0 \cup \mathcal{U}_1 \cup \{v_1\}$ be the graph obtained by interconnecting \mathcal{G}_0 and \mathcal{G}_1 by an arbitrary matching of \mathcal{U}_0 and \mathcal{U}_1 of weight- β links (i.e., every node in \mathcal{U}_h is connected to exactly one node in \mathcal{U}_{1-h} by an undirected weight- β link) and adding a directed weight- γ link from each node in \mathcal{U}_0 to v_0 and from each node in \mathcal{U}_1 to v_1 . (See Figure 1.) Proposition 4 and Theorem 5 imply that

(i) if $\gamma \gg \beta$, then the network is polarized;

(ii) if $\gamma \ll \beta \ll 1$, then influence is homogeneous.

Indeed, point (i) above follows directly from Proposition 4. In order to verify point (ii), for $h = 0, 1$, consider the network $\tilde{\mathcal{G}}_h$ with regular nodes \mathcal{U}_h and stubborn nodes $\mathcal{S}_h := \{v_h\} \cup \mathcal{U}_{1-h}$ obtained by removing from \mathcal{G} node v_{1-h} along with all the internal links of \mathcal{G}_{1-h} . Let $\tilde{\mathcal{G}}_h$ be the undirected graph obtained from $\tilde{\mathcal{G}}_h$ by making all its links incident to v_h and to any $v \in \mathcal{U}_{1-h}$ bidirectional with weight γ and β , respectively. (See Figure 2.) Let l be the total number of undirected links in the Erdos-Renyi graph \mathcal{G}_h , that is of order $m^2 p = \omega m \log m$ with high probability. Note that, in $\tilde{\mathcal{G}}_h$, the degree of v_h is $m\gamma$ (m weight- γ incident links), the degree of any $i \in \mathcal{U}_{1-h}$ is β (one weight- β link), and the total degree of nodes in \mathcal{U}_h is $m(\gamma + \beta) + 2l$, so that the aggregate centrality of \mathcal{S}_h in $\tilde{\mathcal{G}}_h$ is given by

$$\tilde{\pi}_{\mathcal{S}_h} = (\gamma + \beta)m / (2(\gamma + \beta)m + 2l). \quad (29)$$




Figure 2: The graph $\tilde{\mathcal{G}}_0$ obtained from \mathcal{G} by removing node v_1 along with its incident links and all links connecting pairs of nodes in \mathcal{U}_1 and making links incoming in v_0 bidirectional.

Now, we recall a result in [14, Ch. 6] stating that the connected Erdos-Renyi graph has conductance bounded away from 0 with high probability as the network size grows large. This applies directly to the conductance Φ_h of \mathcal{G}_h , for $h = 0, 1$, while it can be shown to carry over to the conductance $\tilde{\Phi}_h$ of $\tilde{\mathcal{G}}_h$ upon verifying that $\gamma \ll \beta \ll 1$ implies that the centralities and the transition probability between every $i, j \in \mathcal{U}_h$ are of the same order in \mathcal{G}_h and $\tilde{\mathcal{G}}_h$, and that cuts in $\tilde{\mathcal{G}}_h$ separating subsets $\mathcal{U} \subseteq \{v_h\} \cup \mathcal{U}_{1-h}$ have bottleneck ratios bounded away from 0. Hence, $\tilde{\Phi}_h \sim \Phi_h$ with high probability as n grows large, and using the fact that the minimum degree in \mathcal{G}_h is of order $\log m$, and the bound (4), one can show that the mixing time of $\tilde{\mathcal{G}}_h$ is of order $\log m$. By combining this with (29), it follows that, with high probability as m grows large, the network $\hat{\mathcal{G}}_h$, $h = 0, 1$, is highly fluid if $\gamma \ll \beta \ll 1$.

Finally, let x and $\hat{x}^{(h)}$, for $h = 0, 1$, be the solutions of the Laplace equations respectively on \mathcal{G} with boundary conditions $x_{v_0} = 0$, $x_{v_1} = 1$, and on $\hat{\mathcal{G}}_h$ with boundary conditions $\hat{x}_{v_h}^{(h)} = h$ and $\hat{x}_i^{(h)} = x_i$ for all $i \in \mathcal{U}_{1-h}$. Observe that $\hat{x}^{(h)}$ coincides with the restriction of x on $\mathcal{V} \setminus \{v_{1-h}\}$. Then, Theorem 5 implies that influence in $\hat{\mathcal{G}}_h$ is homogeneous, so that all but a vanishing fraction of nodes i in \mathcal{U}_0 and \mathcal{U}_1 have equilibrium state $\hat{x}_i = x_i$ close to y_0 and y_1 , respectively. On the other hand, Proposition 4 implies that, if $\gamma \ll \beta$, then $y_1 - y_0 \rightarrow 0$, so that influence is homogeneous on the whole network \mathcal{G} .

5. Conclusion

Simple and deep at the same time: two features that Jan Willems considered central in science can be well appreciated in the theory of distributed averaging. This paper has presented some fundamen-

tal results for distributed averaging systems in a unified framework, giving a novel coherent perspective to classical material together with new generalizations. The role of the electrical network interpretation in providing insight into the equilibrium analysis has been highlighted and some advanced material on the transition between homogeneous influence and polarization has been presented.

Challenging problems for future research include: more complex heterogeneous networks, robustness to perturbations [15], interconnections of higher order or nonlinear systems (e.g., coupled oscillators).

References

- [1] J. Willems, Control as interconnection, Vol. 202 of LNCIS, Springer, 1995, pp. 261–275.
- [2] J. Willems, Models for dynamics, Dynamics Reported 2 (1989) 171–269.
- [3] J. Tsitsiklis, Problems in decentralized decision making and computation, Ph.D. thesis, Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (1984).
- [4] A. Olshevsky, J. N. Tsitsiklis, Convergence speed in distributed consensus and averaging, SIAM Review 53 (4) (2011) 747–772.
- [5] L. Xiao, S. Boyd, S. Lall, A scheme for robust distributed sensor fusion based on average consensus, in: Symposium on Information Processing of Sensor Networks (IPSN), 2005, pp. 63–70.
- [6] A. Jadbabaie, J. Lin, S. Morse, Coordination of groups of mobile autonomous agents using nearest neighbor rules, IEEE Transactions on Automatic Control 48 (6) (2003) 988–1001.
- [7] M. DeGroot, Reaching a consensus, Journal of the American Statistical Association 69 (345) (1974) 118–121.
- [8] P. Doyle, J. Snell, Random walks and electric networks, Mathematical Association of America, 1984.
- [9] D. Levin, Y. Peres, E. Wilmer, Markov Chains and Mixing Times, American Mathematical Society, 2010.
- [10] D. Acemoglu, G. Como, F. Fagnani, A. Ozdaglar, Opinion fluctuations and disagreement in social networks, Mathematics of Operation Research 38 (1) (2013) 1–27.
- [11] R. Montenegro, P. Tetali, Mathematical aspects of mixing times in Markov chains, Foundations and Trends in Theoretical Computer Science 1 (3) (2006) 237–354.
- [12] M. Ji, G. Ferrari-Trecate, M. Egerstedt, A. Buffa, Containment control in mobile networks, IEEE Transactions on Automatic Control 53 (8) (2008) 1972–1975.
- [13] L. Vassio, F. Fagnani, P. Frasca, A. Ozdaglar, Message passing optimization of harmonic influence centrality, IEEE Transactions on Control of Network Systems 1 (1) (2014) 109–120.
- [14] R. Durrett, Random graph dynamics, Cambridge University Press, 2006.
- [15] G. Como, F. Fagnani, Robustness of large-scale stochastic matrices to localized perturbations, IEEE Transactions on Network Science and Engineering 2 (2) (2015) 53–64.