Recursive Neighbor Pruning: Decentralized Community Detection with Certified Behavior from Saturated Opinion Dynamics

Anthony Couthures[1], Irinel-Constantin Morărescu[2], Vineeth Satheeskumar Varma[2], Samson Lasaulce[2] and Antoine Girard[3]

***Abstract*— Detecting communities, i.e., groups of densely connected nodes, usually requires a central entity that knows the whole network. We study Recursive Neighbor Pruning (RNP), a fully decentralized alternative inspired by social learning. Each agent holds a private opinion but only shows its neighbors a coarse public action, namely its opinion amplified and clipped to a fixed range. Agents repeatedly move their opinion toward the actions of their neighbors and, once opinions settle, drop the links with neighbors whose action has the opposite sign. Repeating this procedure on the remaining network splits it into communities, without any agent knowing the network or the number of communities. We use tools from nonlinear control to characterize this behavior. We show that splitting requires the amplification to exceed a threshold set by how strongly the network is bottlenecked. We also show that a group in which every member has enough of its links inside the group can no longer be split once its members agree, and that a single run separates such groups exactly when started from suitable opinions. Finally, we bound the number of runs and messages. On two standard families of synthetic networks, RNP is as accurate as the best centralized method as long as each node has fewer than about forty percent of its links outside its community, and it fails close to the value predicted by our analysis. It is more accurate than two decentralized competitors below this level, but exchanges many more messages.**

I. Introduction

Many networks are organized into communities, i.e., groups of nodes with many links among themselves and few links to the rest of the network. Communities shape how such networks function and how information spreads in them [1], [2]. Standard detection methods, based on spectral clustering [3], [4] or on modularity maximization [5], [6], require a single entity that knows every link. In many networked systems this is impossible, because the network is too large or because nodes cannot or will not share their connections. This calls for algorithms in which communities emerge from local interactions between neighbors, and whose behavior can be guaranteed.

Social learning offers a template. People do not know the whole social network, nor what their acquaintances privately think, yet they infer which group they belong to by observing whether their neighbors' visible choices agree with theirs. Opinion dynamics models [7]–[9] formalize this process, but in their basic form, where each agent averages its neighbors' opinions, all agents of a connected network end up agreeing [10], which erases any trace of communities. Persistent disagreement requires an additional mechanism, such as ignoring neighbors whose opinions are too different [11], [12], antagonistic links [13], or coarse communication [14]–[16]. In the last case, each agent communicates an amplified and clipped version of its opinion, like a hesitant person who votes firmly. When the amplification is strong enough, the neutral agreement becomes unstable and groups of agents polarize toward opposite extremes [17]–[19].

Recursive Neighbor Pruning (RNP) turns this polarization into a community detection algorithm. All agents start from random opinions and run the dynamics until opinions settle; each agent then drops its links to neighbors whose public action has the opposite sign; the procedure is repeated on the remaining network until no more links are dropped. The communities are the groups of nodes that remain connected. Agents never share their private opinion exactly, nor any community label, and they need no knowledge of the network or of the number of communities.

Related work. Three lines of work are closest to ours. Morărescu and Girard [20] detect communities as groups of agents that agree quickly, by cutting links whose disagreement exceeds a shrinking tolerance while the dynamics evolve. RNP instead keeps the network fixed during each run and exploits the several stable polarized states created by clipping. Xing et al. [21] show that the steady states of nonlinear opinion dynamics reveal two planted communities, which are then read by a central observer. RNP is fully decentralized and handles an unknown number of communities. Finally, existing decentralized methods either let nodes broadcast community labels (label propagation [22], [23]) or compute spectral quantities by repeated local averaging [24], [25]; we compare RNP with both. This paper builds on [26] but we propose the following specific novelties.

Contributions. (i) For the step-by-step dynamics actually run by the agents, we show that no link can be cut unless the amplification exceeds a threshold determined by the most pronounced bottleneck of the network. We also show that splitting starts along the direction used by spectral clustering. (ii) We give an explicit condition on the fraction of a node's links that leave its group. Under this condition, a group that has agreed can never be split. We also prove that a single run separates such groups exactly from an explicit set of starting opinions. (iii) We bound the number of runs and messages of RNP. (iv) We compare RNP with centralized and decentralized methods on two families of synthetic

[1]UCLouvain, ICTEAM Institute, Louvain-la-Neuve, 1348, Belgium.
[2]Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France.
[3]Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, F-91190, Gif-sur-Yvette, France.

An extended version of this work, which also covers two variants of RNP, has been submitted to Physical Review E. The document has been checked with Claude Opus 5.5 in order to improve to overall presentation

benchmark networks. The numerical simulations validate our theoretical results.

Notation. $\mathbb{R}$ is the set of real numbers. Vectors and matrices are in bold, and $\boldsymbol{x}^\top$ is the transpose of $\boldsymbol{x}$. $\mathbf{1}$ is the vector with all entries equal to 1, $\mathbf{0}$ is the zero vector and $\boldsymbol{I}$ the identity matrix; their dimension is the number N of agents defined in Section II. $\mathrm{diag}(\boldsymbol{x})$ is the diagonal matrix with diagonal $\boldsymbol{x}$, $\mathrm{span}(\boldsymbol{v}) = \{\alpha \boldsymbol{v} \mid \alpha \in \mathbb{R}\}$, $|\mathcal{S}|$ is the number of elements of a set $\mathcal{S}$, $\mathrm{Fix}(f)$ is the set of fixed points of a map f, $\lceil \cdot \rceil$ rounds up to an integer, $\mathcal{U}(a, b)$ is the uniform distribution on $[a, b]$, and $\mathrm{sat}(y) = \min(1, \max(-1, y))$.

II. Problem Setting and Model

A. Graph and decentralized constraints

We consider N agents interacting over an undirected, connected, simple graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, $\mathcal{V} = \{1, \dots, N\}$, with adjacency matrix $\boldsymbol{A}$, neighborhoods $\mathcal{N}_i = \{j \in \mathcal{V} \mid (i, j) \in \mathcal{E}\}$, degrees $d_i = |\mathcal{N}_i| \geq 1$, and $\boldsymbol{D} = \mathrm{diag}(d_1, \dots, d_N)$. The random-walk matrix $\boldsymbol{P} = \boldsymbol{D}^{-1}\boldsymbol{A}$ is self-adjoint for the inner product $\langle \boldsymbol{x}, \boldsymbol{y} \rangle_{\boldsymbol{D}} = \boldsymbol{x}^\top \boldsymbol{D} \boldsymbol{y}$, with induced norm $\|\boldsymbol{x}\|_{\boldsymbol{D}} = \langle \boldsymbol{x}, \boldsymbol{x} \rangle_{\boldsymbol{D}}^{1/2}$; vectors orthogonal for this inner product are called $\boldsymbol{D}$-orthogonal. Hence the eigenvalues of $\boldsymbol{P}$ are real, $-1 \leq \lambda_1 \leq \cdots \leq \lambda_{N-1} < \lambda_N = 1$, and the associated eigenvectors $\boldsymbol{v}_1, \dots, \boldsymbol{v}_N$ can be chosen $\boldsymbol{D}$-orthogonal, with $\boldsymbol{v}_N = \mathbf{1}$ since $\boldsymbol{P}\mathbf{1} = \mathbf{1}$. The eigenvector $\boldsymbol{v}_{N-1}$ is the Fiedler vector of the random-walk Laplacian $\boldsymbol{I} - \boldsymbol{P}$. We write $\rho(\mathcal{G}) = \max(\lambda_{N-1}, -\lambda_1)$ and $K^*(\mathcal{G}) = 1/\lambda_{N-1}$ with convention $K^*(\mathcal{G}) = +\infty$ if $\lambda_{N-1} \leq 0$.

We seek a partition of $\mathcal{V}$ into communities under three constraints: *locality* (agent i only interacts with $\mathcal{N}_i$ and knows no global quantity), *limited disclosure* (agents exchange a single scalar per step, and never labels, degrees or neighbor lists), and *uniformity* (identical rules, no leader or central aggregation).

B. Nonlinear social learning dynamics

At each time step $t = 0, 1, 2, \dots$, each agent i holds a private opinion $x_i(t) \in [-1, 1]$ and broadcasts the public action $s(x_i(t))$, where $s : [-1, 1] \to [-1, 1]$ is the action function specified in Assumption 1 below. It moves toward the average action of its neighbors:

$$x_i(t+1) = x_i(t) + h\Big(\frac{1}{d_i} \sum_{j \in \mathcal{N}_i} s(x_j(t)) - x_i(t)\Big), \qquad (1)$$

with step size $h \in (0, 1]$. With $\boldsymbol{x} = (x_1, \dots, x_N)^\top$ and $\boldsymbol{s}(\boldsymbol{x}) = (s(x_1), \dots, s(x_N))^\top$, this reads

$$\boldsymbol{x}(t+1) = (1-h)\boldsymbol{x}(t) + h\boldsymbol{P}\boldsymbol{s}(\boldsymbol{x}(t)).$$

This is the Euler discretization of $\dot{\boldsymbol{x}} = \boldsymbol{P}\boldsymbol{s}(\boldsymbol{x}) - \boldsymbol{x}$ studied in [19], [26]. As $x_i(t+1)$ is a convex combination of $x_i(t)$ and values of s, $\mathcal{X} = [-1, 1]^N$ is forward invariant.

Assumption 1. *The function $s : [-1, 1] \to [-1, 1]$ satisfies $0 \leq (s(a) - s(b))/(a - b) \leq K$ for all $a \neq b$, is differentiable at 0 with $s'(0) = K > 1$, and $\mathrm{Fix}(s) = \{-1, 0, 1\}$.*

The natural example, used in our experiments, is $s_K(x) = \min(\max(Kx, -1), 1) := \mathrm{sat}(Kx)$, which is odd. Monotonicity reflects cooperative behavior and makes (1) a monotone system [27]. Note that a neighbor observing $s_K(x_i)$ can reconstruct x_i exactly when $|x_i| < 1/K$, but only its sign once $|x_i| \geq 1/K$. This is not a formal privacy guarantee: the sign is always revealed, and repeated observations may reveal more. Assuming $K > 1$ makes consensus at $\mathbf{0}$ repelling and $\pm\mathbf{1}$ attracting, creating a symmetry breaking.

Algorithm 1: Recursive Neighbor Pruning (RNP)

Input : $\mathcal{G} = (\mathcal{V}, \mathcal{E})$; s, h, $T_{\max}$, N_{retry}
Output: Pruned graph, communities as connected components

$n \leftarrow 0$
while $n < N_{\mathrm{retry}}$ **do**
  foreach $i \in \mathcal{V}$ **do** $x_i \leftarrow \mathcal{U}(-1, 1)$
  for $t = 1$ **to** $T_{\max}$ **do**
    foreach $i \in \mathcal{V}$ **do**
      $x_i \leftarrow x_i + h\big(\frac{1}{d_i} \sum_{j \in \mathcal{N}_i} s(x_j) - x_i\big)$
  $\mathcal{E}_{\mathrm{cut}} \leftarrow \{(i, j) \in \mathcal{E} \mid s(x_i)s(x_j) < 0\}$
  if $\mathcal{E}_{\mathrm{cut}} \neq \emptyset$ **then** $\mathcal{E} \leftarrow \mathcal{E} \setminus \mathcal{E}_{\mathrm{cut}}$; update $\mathcal{N}_i, d_i$; $n \leftarrow 0$
  else $n \leftarrow n + 1$

III. The Recursive Neighbor Pruning Algorithm

A *round* of RNP (Algorithm 1) consists of:
1. independent draws $x_i(0) \sim \mathcal{U}(-1, 1)$
2. $T_{\max}$ steps of dynamics (1) on the current graph
3. removal of every edge (i, j) with $s(x_i)s(x_j) < 0$, a local action since agent i knows its action and observes that of j.

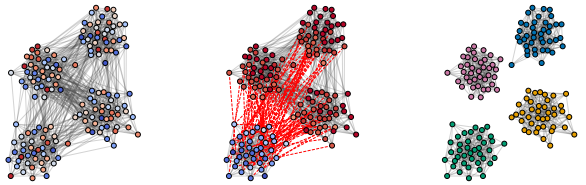
Then, if an edge is cut, a new round starts on the new graph; otherwise the round is repeated, up to N_{retry} consecutive unsuccessful rounds. The communities are the connected components of the final graph.

In Fig. 1 we illustrate the implementation of this algorithm using a Stochastic Block Model (SBM) with four blocks of 40 nodes, in which two nodes are linked with probability $p_{\mathrm{in}} = 0.3$ if they belong to the same block and $p_{\mathrm{out}} \leq 0.05$ otherwise.

The initialization is random rather than spectral. A Fiedler initialization would require global information and restrict rounds to bisections, whereas random initial conditions sample the multistable attractor landscape and often cut the graph into more than two components in one round. Since the dynamics on distinct components are decoupled, the recursion refines the partition component by component, without any agent knowing the number of communities.

IV. Theoretical Guarantees

For $\mathcal{W} \subseteq \mathcal{V}$ and $i \in \mathcal{W}$, let $d_i^E = |\mathcal{N}_i \setminus \mathcal{W}|$ be the number of neighbors of i outside $\mathcal{W}$ and $\xi_i^{\mathcal{W}} = d_i^E/d_i$ its *external degree ratio*.



(a) Initialization. (b) Edge removal. (c) Final partition.

Fig. 1: RNP on an SBM with $p_{\text{in}} = 0.3$ and $p_{\text{out}} \leq 0.05$; $s = s_K$, $K = 5$, $h = 0.1$. (b) At steady state, edges joining agents with opposite actions (red) are cut. (c) Recursion yields four communities.

A. Spectral threshold for the discrete-time dynamics

Let $\mathbf{1}^{\perp_D} = \{\boldsymbol{x} \in \mathbb{R}^N \mid \langle \mathbf{1}, \boldsymbol{x} \rangle_{\boldsymbol{D}} = 0\}$ be the subspace of vectors D-orthogonal to the consensus direction $\mathbf{1}$. This subspace is spanned by $\boldsymbol{v}_1, \dots, \boldsymbol{v}_{N-1}$ and invariant under $\boldsymbol{P}$. For $\boldsymbol{x} \in \mathbb{R}^N$, let $\bar{x} = \mathbf{1}^\top \boldsymbol{D}\boldsymbol{x}/W$ be its degree-weighted average, where $W = \sum_i d_i$, and $\Pi\boldsymbol{x} = \boldsymbol{x} - \bar{x}\mathbf{1}$ its D-orthogonal projection onto $\mathbf{1}^{\perp_D}$, i.e., its disagreement component. One has $\Pi\boldsymbol{P} = \boldsymbol{P}\Pi$ and, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^N$,

$$\langle \Pi\boldsymbol{x}, \Pi\boldsymbol{y} \rangle_{\boldsymbol{D}} = \frac{1}{2W} \sum_{i,j=1}^{N} d_i d_j (x_i - x_j)(y_i - y_j). \tag{2}$$

Proposition 1. *Let Assumption 1 hold and $h \in (0, 1]$.*

i) If $K\lambda_{N-1} < 1$, every equilibrium of (1) is of the form $\alpha\mathbf{1}$ with $\alpha \in \{-1, 0, 1\}$.

ii) If $K\rho(\mathcal{G}) < 1$, then for all $t \geq 0$, $\|\Pi\boldsymbol{x}(t)\|_{\boldsymbol{D}} \leq \big(1 - h(1 - K\rho(\mathcal{G}))\big)^t \|\Pi\boldsymbol{x}(0)\|_{\boldsymbol{D}}$.

iii) If $K\lambda_{N-1} < 1$, the consensus set $\operatorname{span}(\mathbf{1})$ is globally attractive for $\dot{\boldsymbol{x}} = \boldsymbol{P}\boldsymbol{s}(\boldsymbol{x}) - \boldsymbol{x}$.

Proof. (i) First, let us note that equilibria satisfy $\boldsymbol{x} = \boldsymbol{P}\boldsymbol{s}(\boldsymbol{x})$. To simplify the writing we will replace $\boldsymbol{s}(\boldsymbol{x})$ by $\boldsymbol{s}$ when no confusion arises. Since $\Pi\boldsymbol{P} = \boldsymbol{P}\Pi$, equilibria satisfy $\Pi\boldsymbol{x} = \boldsymbol{P}\Pi\boldsymbol{s}$. Furthermore, the Rayleigh quotient of $\boldsymbol{P}$ on $\mathbf{1}^{\perp_D}$ gives $\langle \Pi\boldsymbol{s}, \Pi\boldsymbol{x} \rangle_{\boldsymbol{D}} \leq \lambda_{N-1} \|\Pi\boldsymbol{s}\|_{\boldsymbol{D}}^2$. By Assumption 1 one has

$$(s(a) - s(b))(a - b) \geq \frac{1}{K}(s(a) - s(b))^2,$$

so (2) yields $\langle \Pi\boldsymbol{s}, \Pi\boldsymbol{x} \rangle_{\boldsymbol{D}} \geq \|\Pi\boldsymbol{s}\|_{\boldsymbol{D}}^2 / K$. Hence,

$$\left(\frac{1}{K} - \lambda_{N-1}\right) \|\Pi\boldsymbol{s}\|_{\boldsymbol{D}}^2 \leq 0.$$

Therefore, if $K\lambda_{N-1} < 1$ one gets that $\|\Pi\boldsymbol{s}\|_{\boldsymbol{D}}^2 \leq 0$ yielding $\boldsymbol{s} = \alpha\mathbf{1}$ for some $\alpha \in \mathbb{R}$, $\boldsymbol{x} = \boldsymbol{P}\boldsymbol{s} = \alpha\mathbf{1}$ and $\alpha = s(\alpha)$. Recalling that Assumption 1 imposes $\operatorname{Fix}(s) = \{-1, 0, 1\}$, i.e. $\alpha \in \{-1, 0, 1\}$, completes the proof.

(ii) It is straightforward that

$$\Pi\boldsymbol{x}(t+1) = (1 - h)\Pi\boldsymbol{x}(t) + h\boldsymbol{P}\Pi\boldsymbol{s}.$$

The $\boldsymbol{D}$-norm of $\boldsymbol{P}$ on $\mathbf{1}^{\perp_D}$ is $\rho(\mathcal{G})$, and (2) with $|s(x_i) - s(x_j)| \leq K|x_i - x_j|$ gives $\|\Pi\boldsymbol{s}\|_{\boldsymbol{D}} \leq K\|\Pi\boldsymbol{x}\|_{\boldsymbol{D}}$. Replacing this in the equality above yields the inequality.

(iii) This result is already stated in [26]. □

Items (i) and (iii) involve λ_{N-1} only; item (ii) also requires $K|\lambda_1| < 1$, which only matters for graphs with a strong bipartite component. With $s(x) = x$ ($K = 1$) item (i) holds on every connected graph: linear averaging cannot reveal communities. Graphs without bottleneck, such as complete graphs ($\lambda_{N-1} \leq 0$), cannot be split for any K. Item (iii) matters because item (i) alone does not rule out a cut. A nonlinear system can have all its equilibria at consensus and still have trajectories that never reach them, for example oscillations. Its limitation is that RNP runs the discrete-time dynamics, where item (iii) does not apply as is. The gap only matters for nearly bipartite graphs with a large step h; for small h the discrete dynamics follows the continuous one.

Proposition 2 (Unstable directions)**.** *Under the assumptions of Proposition 1, the linearization of (1) at $\mathbf{0}$ is*

$$\boldsymbol{J} = (1 - h)\boldsymbol{I} + hK\boldsymbol{P},$$

with eigenvalues $\nu_k = 1 + h(K\lambda_k - 1)$ and eigenvectors $\boldsymbol{v}_k$, $k = 1, \dots, N$. The origin is unstable along $\mathbf{1}$ ($\nu_N > 1$). Moreover, the invariant subspace $\mathbf{1}^{\perp_D}$ contains an unstable direction if and only if $K\lambda_{N-1} > 1$, the dominant one being the Fiedler vector $\boldsymbol{v}_{N-1}$; and $\nu_N > \nu_{N-1}$.

The proof follows from the spectral decomposition of $\boldsymbol{P}$. For $K > K^*(\mathcal{G})$, disagreements are amplified along the Fiedler vector, whose sign pattern is the spectral cut [4]. However, the consensus mode grows faster, so a round may end in $\pm\mathbf{1}$ without any cut: this is why RNP needs retries.

B. Robustness and confinement of dense cores

Let $\mathcal{W} \subseteq \mathcal{V}$, $\sigma \in \{-1, 1\}$, and define $L_\sigma \geq 0$ the Lipschitz constant of s on an interval $\mathcal{I}_\sigma$ which defines a neighborhood of σ (note that for a saturation function s_K one has $L_\sigma = 0$ on $\{x \mid \sigma x \geq 1/K\}$). Let us also introduce

- $\eta = \min_{i \in \mathcal{W}} \left(1 - L_\sigma(1 - \xi_i^{\mathcal{W}})\right)$, $\gamma = \max_{i \in \mathcal{W}} \frac{\xi_i^{\mathcal{W}}}{1 - L_\sigma(1 - \xi_i^{\mathcal{W}})}$
- distance of the group to the polarized state σ: $\delta(t) = \max_{i \in \mathcal{W}} |x_i(t) - \sigma|$
- disturbance from external neighbors: $u(t) = \max_{i \in \mathcal{W}, j \in \mathcal{N}_i \setminus \mathcal{W}} |s(x_j(t)) - \sigma|$

The following input-to-state stability (ISS) property holds (see [26]): if $\eta > 0$ and $x_i(\tau) \in \mathcal{I}_\sigma$ for all $i \in \mathcal{W}$, $\tau \leq t$, then

$$\delta(t) \leq (1 - h\eta)^t \delta(0) + \gamma \max_{0 \leq \tau < t} u(\tau). \tag{3}$$

The gain γ is smaller for more internally connected agents, but (3) does not prevent a sign change. The next result does.

Proposition 3 (Confinement)**.** *Let $s = s_K$, $h \in (0, 1]$, and $\mathcal{W} \subseteq \mathcal{V}$ satisfy*

$$\xi_i^{\mathcal{W}} \leq \frac{K - 1}{2K} \qquad \forall i \in \mathcal{W}. \tag{4}$$

Then $\mathcal{H}^+(\mathcal{W}) = \{\boldsymbol{x} \in \mathcal{X} \mid x_i \geq 1/K,\ \forall i \in \mathcal{W}\}$ and $\mathcal{H}^-(\mathcal{W}) = \{\boldsymbol{x} \in \mathcal{X} \mid x_i \leq -1/K,\ \forall i \in \mathcal{W}\}$ are forward invariant. Moreover, without information on the external neighbors (4) cannot be weakened.

Proof. Let us consider $\boldsymbol{x}(t) \in \mathcal{H}^+(\mathcal{W})$, $i \in \mathcal{W}$, and let $\bar{s}_i = \sum_{j\in\mathcal{N}_i} s_K(x_j(t))/d_i$ be the average action of its neighbors. Internal neighbors broadcast 1 and external ones at least -1, so $\bar{s}_i \geq 1 - 2\xi_i^{\mathcal{W}} \geq 1/K$, and

$$x_i(t+1) = (1-h)x_i(t) + h\bar{s}_i \geq \min(x_i(t), \bar{s}_i) \geq 1/K.$$

If (4) fails for i and all its external neighbors broadcast -1, then $\bar{s}_i < 1/K$ and, for $h = 1$, x_i leaves $\mathcal{H}^+(\mathcal{W})$. A similar reasoning can be done for $\mathcal{H}^-(\mathcal{W})$. □

This closed form specializes the general force balance of [26]. Since $(K-1)/2K < 1/2$, no saturated dynamics can confine a group whose members have more external than internal neighbors.

C. Exact separation in one round

Let us notice that Proposition 3 provides a condition for the preservation of a given community. The next result treats all groups simultaneously, provides a sufficient value for $T_{\max}$ and also gives a sufficient condition to reach polarization not just to stay there.

Theorem 1 (One-round separation)**.** *Let $s = s_K$, $h \in (0,1]$, and $\{\mathcal{V}_1, \dots, \mathcal{V}_m\}$ be a partition of $\mathcal{V}$ into m groups such that $\xi_i^{\mathcal{V}_c} \leq \varepsilon < (K-1)/2K$ for all $c \in \{1,\dots,m\}$ and $i \in \mathcal{V}_c$. Let $m^* = \varepsilon/((1-\varepsilon)K-1) \in (0, 1/K)$ and $\sigma = (\sigma_1, \dots, \sigma_m) \in \{-1,1\}^m$.*

If $\sigma_c x_i(0) > m^*$ for all $i \in \mathcal{V}_c$ and all $c \in \{1,\dots,m\}$, then $\sigma_c x_i(t) \geq 1/K$ for all $i \in \mathcal{V}_c$ and all $t \geq t^*$, where

$$t^* = \Big\lceil \frac{1/K - m_0}{h\big((1-\varepsilon)K-1\big)(m_0 - m^*)} \Big\rceil,$$

with $m_0 = \min_{c,\, i\in\mathcal{V}_c} \sigma_c x_i(0)$. Consequently, if $T_{\max} \geq t^*$, the round cuts exactly the edges between groups with opposite signs σ_c, and no other edge.

Proof. Let $m(t) = \min_{c,\, i\in\mathcal{V}_c} \sigma_c x_i(t)$. Since s_K is odd and non-decreasing, for $i \in \mathcal{V}_c$ the internal neighbors satisfy $\sigma_c s_K(x_j) \geq s_K(m)$ and the external ones $\sigma_c s_K(x_j) \geq -1$. So, with $\xi_i = \xi_i^{\mathcal{V}_c}$ and $m = m(t)$ one gets

$$\sigma_c x_i(t+1) \geq (1-h)m + h[(1-\xi_i)s_K(m) - \xi_i].$$

The bracket is non-increasing in ξ_i because $s_K(m) \geq -1$, hence $m(t+1) \geq g(m(t))$ with

$$g(m) = (1-h)m + h[(1-\varepsilon)s_K(m) - \varepsilon],$$

and g is non-decreasing. On $[m^*, 1/K]$ one has

$$g(m) - m = h((1-\varepsilon)K - 1)(m - m^*) \geq 0,$$

so the iterates of g from m_0 increase by at least $h((1-\varepsilon)K-1)(m_0-m^*)$ per step until they exceed $1/K$, which happens before t^*. For $m \geq 1/K$, $g(m) \geq \min(m, 1-2\varepsilon) \geq 1/K$, since $1 - 2\varepsilon > 1/K$. Thus $m(t) \geq 1/K$ for $t \geq t^*$: every agent of $\mathcal{V}_c$ broadcasts σ_c, and an edge is cut if and only if its endpoints lie in groups with opposite signs. □

Theorem 1 is a deterministic recovery guarantee for one round. It thus certifies *which* cut a round produces from a given set of initial conditions, not the typical behavior. Section V shows that the actual basin is much larger. This is mainly due to the fact that the proposed uniform initialization covers a small set for large N.

D. Consequences for RNP

Corollary 1 (Certificates for RNP)**.** *Let Assumption 1 hold and $h \in (0,1]$.*

i) Cost: *RNP stops after at most $(|\mathcal{E}|+1)N_{\text{retry}}$ rounds, and sends at most $2T_{\max}|\mathcal{E}|$ scalar messages per round. At least $2T_{\max}N_{\text{retry}}$ messages per edge are spent after the last cut.*

ii) Core preservation: *Let $s = s_K$ and $\mathcal{W}$ satisfy* (4)*. If the state enters $\mathcal{H}^+(\mathcal{W})$ or $\mathcal{H}^-(\mathcal{W})$ during a round, no edge inside $\mathcal{W}$ is cut in that round, and* (4) *keeps holding in all later graphs where $\mathcal{W}$ is intact.*

iii) Indivisibility: *If a component $\mathcal{G}_c$ of the current graph satisfies $K\rho(\mathcal{G}_c) < 1$, every round on $\mathcal{G}_c$ converges to the consensus set at the rate of Proposition 1; if the limit is $\pm\mathbf{1}$, no edge of $\mathcal{G}_c$ is cut.*

Proof. (i) Every round removes an edge or increments a counter that is reset only after a removal. (ii) Proposition 3, and pruning only decreases d_i^E. (iii) Proposition 1-(ii), applied to $\mathcal{G}_c$; components evolve independently. □

Items (ii)–(iii) and Theorem 1 suggest choosing K in the window $K^*(\mathcal{G}) < K < \min_c K^*(\mathcal{G}_c)$, where $\mathcal{G}_c$ is the subgraph induced by the c-th community to be found (with $\rho(\mathcal{G}_c)$ in place of λ_{N-1} for $\mathcal{G}_c$ in discrete time): the graph can then be split, while isolated communities are certified indivisible. This window competes with the confinement threshold $(K-1)/2K$, which increases with K.

V. Numerical Performance

We use two benchmarks. The main one is the Artificial Benchmark for Community Detection (ABCD) [28], whose communities, unlike those of the Lancichinetti–Fortunato–Radicchi (LFR) benchmark [29], remain well connected as mixing increases. Its mixing parameter ξ is the expected external degree ratio, i.e., the mean over nodes i of $\xi_i^{\mathcal{V}_c}$, where $\mathcal{V}_c$ is the community of i. We generate graphs with $N = 1000$, degree exponent 2.5, degrees in $[10, 100]$, community size exponent 1.5 and sizes in $[20, 250]$, using the Python implementation `abcd-graph`. Unless stated otherwise, RNP uses s_K with $K = 4$, $h = 0.5$, $T_{\max} = 300$, $N_{\text{retry}} = 50$. The second benchmark is LFR, generated with the `networkx` implementation and the same degree and community size parameters, with 8 graphs per nominal mixing $\mu \in [0.05, 0.65]$.

Baselines are Louvain [5] (centralized reference), asynchronous label propagation [22] and a distributed spectral bisection (DSB) in the spirit of [24], [25]. In DSB, each component runs 30 steps of the lazy walk $\boldsymbol{y} \leftarrow (\boldsymbol{y} + \boldsymbol{P}\boldsymbol{y})/2$ from a random vector; the local differences $y_i(t) - y_i(t+1)$ cancel the consensus mode, and their signs approximate the Fiedler cut, which is applied as in RNP. A split is kept only if it increases the modularity; this global test is granted to DSB

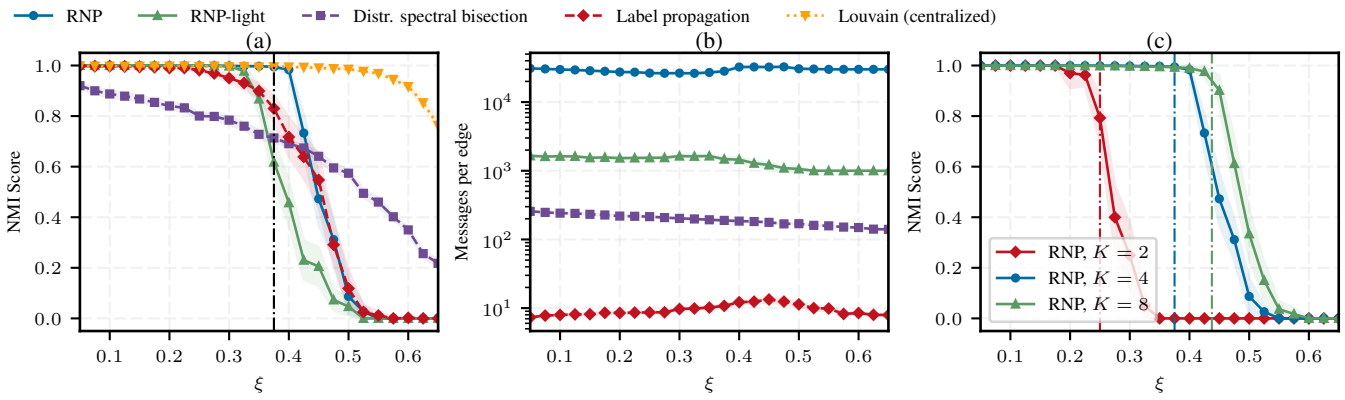


Fig. 2: ABCD graphs ($N = 1000$, 20 graphs per ξ, 5 runs of each randomized method per graph; lines are means over the graphs, shaded areas 95% confidence intervals). (a) NMI of RNP ($K = 4$, $T_{\max} = 300$, $N_{\text{retry}} = 50$), RNP-light ($T_{\max} = 50$, $N_{\text{retry}} = 10$) and baselines; the vertical line is $(K-1)/2K = 0.375$. (b) Scalar messages per edge. (c) NMI of RNP for $K \in \{2, 4, 8\}$; vertical lines are the confinement thresholds $(K-1)/2K \in \{0.25, 0.375, 0.44\}$.

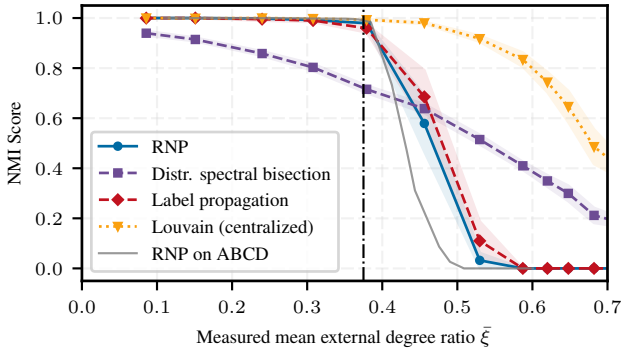


Fig. 3: Second benchmark: LFR graphs ($N = 1000$, 20 graphs per μ, 5 runs of each randomized method per graph), plotted against the measured mean external degree ratio $\bar{\xi}$. The gray line is RNP on ABCD (Fig. 2a), also plotted against its measured $\bar{\xi}$; the vertical line is $(K-1)/(2K) = 0.375$.

for free, which favors it. We report the normalized mutual information (NMI) between the detected and true partitions (1 for a perfect match, 0 for independent partitions) and the number of scalar messages per edge.

A. Results

Accuracy. RNP achieves NMI ≥ 0.99 and finds the correct communities (relative error below 0.01) for $\xi \leq 0.4$, similar to Louvain (Fig. 2a). In this regime it outperforms label propagation and DSB (NMI 0.72 and 0.69 at $\xi = 0.4$, respectively), and ≤ 0.92 even for well-separated graphs). Beyond $\xi \approx 0.40$, the NMI of RNP drops (about 0.47 at $\xi = 0.45$, 0.09 at 0.5, 0 from 0.55) whereas Louvain stays above 0.96 up to $\xi = 0.55$; DSB degrades more slowly and is the most accurate decentralized method for $\xi \geq 0.45$. When it fails, RNP returns a single community (relative error close to -1): it *fails safely* rather than producing a noisy partition.

Communication cost. RNP needs about 3×10^4 messages per edge (Fig. 2b), two orders of magnitude more than DSB (140 to 260) and three more than label propagation (about 10). This number of message is due to the certification of indivisibility: by Corollary 1, the final N_{retry} unsuccessful rounds alone cost $2T_{\max}N_{\text{retry}} = 3 \times 10^4$ messages per edge. RNP-light reduces it twentyfold with same accuracy for $\xi \leq 0.3$, but start degrading earlier (0.46 at $\xi = 0.4$). This comes from the parameterN_{retry} that trades messages for accuracy, and not from $T_{\max}$. Since rounds start from independent initial conditions, a component that a round cuts with probability p is wrongly declared indivisible with probability $(1-p)^{N_{\text{retry}}}$, and p is small near the breakdown. With $N_{\text{retry}} = 10$ and $T_{\max} = 300$, RNP returns the same partitions for $\xi \leq 0.275$ with three times fewer messages, but its NMI drops quickly as ξ increases.

Breakdown follows the confinement threshold. Fig. 2c varies K. The breakdown point moves with K as predicted by Proposition 3: RNP keeps an NMI of at least 0.97 up to $\xi = 0.2$ for $K = 2$ (threshold 0.25), up to 0.4 for $K = 4$ (threshold 0.375), and up to 0.425 for $K = 8$ (threshold 0.44). On these graphs, $K^*(\mathcal{G})$ grows from 1.1 to 1.8 and $\min_c K^*(\mathcal{G}_c)$, over the subgraphs $\mathcal{G}_c$ induced by the true communities, decreases from 2.4 to 1.76 as ξ goes from 0.1 to 0.5. Hence $K = 4$ and $K = 8$ lie outside the indivisibility window, yet no over-segmentation is observed (relative error at most 0.06). The window is therefore conservative in practice, and confinement is the binding constraint: the best choice is a large gain.

Second benchmark. The LFR generator benchmark is mixing more the communities with its nominal parameter μ (e.g., a measured mean external ratio $\bar{\xi} = 0.38$ for $\mu = 0.25$ and 0.46 for $\mu = 0.3$). Simplify exposition, Fig. 3 uses $\bar{\xi}$ as x-axis. RNP is again exact up to $\bar{\xi} \approx 0.38$ (NMI ≥ 0.97), drops to 0.58 at $\bar{\xi} = 0.46$ and to 0.03 beyond 0.53, i.e., it breaks down at the same external ratio as on ABCD, close to the confinement threshold. Below the breakdown, RNP is slightly more accurate than label propagation and clearly more accurate than DSB; at the breakdown, label propagation degrades slightly later, and DSB is the most accurate decentralized method beyond it. Louvain remains

above 0.9 up to $\bar{\xi} \approx 0.53$. On these graphs, $K^*(\mathcal{G})$ exceeds $\min_c K^*(\mathcal{G}_c)$ from $\mu \approx 0.33$ on, and 6% of the communities are disconnected at $\mu = 0.4$, which confirms that LFR communities lose their internal cohesion as mixing increases.

VI. Limitations and Discussion

We can observe the following limitation for RNP:

Irreversible cuts. A cut is decided from a single trajectory and never undone. Proposition 3 protects a core only once it has polarized, and Theorem 1 covers an exponentially small set of initial conditions; frontier nodes are then assigned to one side or isolated depending on the initialization.

Cost of stopping. Because the consensus mode dominates (Proposition 2), indivisibility is only certified statistically, after N_{retry} unsuccessful rounds. The retry counter also requires a component-wide agreement on whether a cut occurred; running a fixed number of rounds on a common clock is a fully decentralized alternative.

Gain selection. The relevant thresholds depend on $K^*(\mathcal{G})$, $K^*(\mathcal{G}_c)$ and the external ratios, which agents cannot compute. Our results support a large initial gain, but no local rule adapting K is available yet.

Disclosure and scope. Saturation limits but does not prevent information leakage, and no formal privacy guarantee is claimed. Proposition 1 requires undirected graphs, and item (ii) the stronger condition $K\rho(\mathcal{G}) < 1$.

These limitations motivate two variants studied in the extended version: cutting edges during the transient with a decaying confidence bound, in the spirit of [20], and replacing hard cuts by edge reliability scores over many rounds.

VII. Conclusion

We analyzed Recursive Neighbor Pruning, a decentralized community detection algorithm in which agents only exchange saturated actions and cut ties with neighbors displaying opposite actions. For the discrete-time dynamics run by the agents, we proved a spectral threshold for splitting, identified the Fiedler direction of symmetry breaking, gave a closed-form confinement condition $\xi_i^{\mathcal{W}} \leq (K-1)/2K$ on the external degree ratios and an exact one-round separation result, and bounded the cost of the recursion. On ABCD and LFR benchmarks, RNP matches Louvain up to a mean external ratio of about 0.4, its breakdown follows the confinement threshold on both benchmarks and for three gains, and it fails safely beyond. Its main weaknesses are its communication cost, dominated by stopping, and the lack of guarantees on the typical basin of the separating rounds, which future work will address together with local gain adaptation.

References

[1] S. Fortunato, "Community detection in graphs," *Physics Reports*, vol. 486, no. 3, pp. 75–174, Feb. 2010.

[2] J. Li, S. Lai *et al.*, "A comprehensive review of community detection in graphs," *Neurocomputing*, vol. 600, p. 128169, Oct. 2024.

[3] A. Pothen, H. D. Simon, and K.-P. Liou, "Partitioning sparse matrices with eigenvectors of graphs," *SIAM J. Matrix Anal. Appl.*, vol. 11, no. 3, Jul. 1990.

[4] D. A. Spielman and S.-H. Teng, "Spectral partitioning works: Planar graphs and finite element meshes," *Linear Algebra and its Applications*, vol. 421, no. 2, pp. 284–305, Mar. 2007.

[5] V. D. Blondel, J.-L. Guillaume *et al.*, "Fast unfolding of communities in large networks," *J. Stat. Mech.*, vol. 2008, no. 10, p. P10008, Oct. 2008.

[6] V. A. Traag, L. Waltman, and N. J. van Eck, "From Louvain to Leiden: Guaranteeing well-connected communities," *Sci Rep*, vol. 9, no. 1, p. 5233, Mar. 2019.

[7] J. R. P. French, "A formal theory of social power." *Psychological Review*, vol. 63, no. 3, pp. 181–194, 1956.

[8] M. H. Degroot, "Reaching a Consensus," *J. Am. Stat. Assoc.*, vol. 69, no. 345, pp. 118–121, Mar. 1974.

[9] N. E. Friedkin and E. C. Johnsen, "Social influence and opinions," *J. Math. Sociol.*, vol. 15, no. 3-4, pp. 193–206, Jan. 1990.

[10] R. L. Berger, "A Necessary and Sufficient Condition for Reaching a Consensus Using DeGroot's Method," *J. Am. Stat. Assoc.*, vol. 76, no. 374, pp. 415–418, 1981.

[11] G. Deffuant, D. Neau *et al.*, "Mixing beliefs among interacting agents," *Advs. Complex Syst.*, vol. 03, no. 01n04, pp. 87–98, Jan. 2000.

[12] R. Hegselmann and U. Krause, "Opinion dynamics and bounded confidence models, analysis, and simulation," *J. Artif. Soc. Soc. Simul.*, vol. 5, no. 3, 2002.

[13] C. Altafini, "Consensus Problems on Networks With Antagonistic Interactions," *IEEE Trans. Autom. Control*, vol. 58, no. 4, pp. 935–946, Apr. 2013.

[14] A. C. R. Martins, "Continuous opinions and discrete actions in opinion dynamics problems," *Int. J. Mod. Phys. C*, vol. 19, no. 4, pp. 617–625, 2008.

[15] R. Gray, A. Franci *et al.*, "Multiagent Decision-Making Dynamics Inspired by Honeybees," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 2, pp. 793–806, Jun. 2018.

[16] A. Bizyaeva, A. Franci, and N. E. Leonard, "Nonlinear Opinion Dynamics With Tunable Sensitivity," *IEEE Trans. Automat. Contr.*, vol. 68, no. 3, pp. 1415–1430, Mar. 2023.

[17] A. Fontan and C. Altafini, "Multiequilibria Analysis for a Class of Collective Decision-Making Networked Systems," *IEEE Trans. Control Netw. Syst.*, vol. 5, no. 4, pp. 1931–1940, Dec. 2018.

[18] N. E. Leonard, K. Lipsitz *et al.*, "The nonlinear feedback dynamics of asymmetric political polarization," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 118, no. 50, p. e2102149118, Dec. 2021.

[19] A. Couthures, V. S. Varma *et al.*, "Global Synchronization of Multi-Agent Systems With Nonlinear Interactions," *IEEE Control Syst. Lett.*, vol. 9, pp. 354–359, 2025.

[20] I.-C. Morărescu and A. Girard, "Opinion Dynamics With Decaying Confidence: Application to Community Detection in Graphs," *IEEE Trans. Autom. Control*, vol. 56, no. 8, pp. 1862–1873, Aug. 2011.

[21] Y. Xing, A. Bizyaeva, and K. H. Johansson, "Learning Communities from Equilibria of Nonlinear Opinion Dynamics," in *2024 IEEE 63rd Conference on Decision and Control (CDC)*, 2024, pp. 2325–2330.

[22] U. N. Raghavan, R. Albert, and S. Kumara, "Near linear time algorithm to detect community structures in large-scale networks," *Phys. Rev. E*, vol. 76, no. 3, p. 036106, Sep. 2007.

[23] J. Xie, B. K. Szymanski, and X. Liu, "SLPA: Uncovering Overlapping Communities in Social Networks via a Speaker-Listener Interaction Dynamic Process," in *2011 IEEE 11th Int. Conf. Data Min. Workshop*, Dec. 2011, pp. 344–349.

[24] D. Kempe and F. McSherry, "A decentralized algorithm for spectral analysis," *Journal of Computer and System Sciences*, vol. 74, no. 1, pp. 70–83, 2008.

[25] F. Lin and W. W. Cohen, "Power iteration clustering," in *Proc. 27th Int. Conf. Machine Learning (ICML)*, 2010, pp. 655–662.

[26] A. Couthures, G. Bainier *et al.*, "From Consensus to Robust Clustering: Multi-Agent Systems with Nonlinear Interactions," no. arXiv:2511.21228, Nov. 2025.

[27] M. Hirsch and H. Smith, "Chapter 4 Monotone Dynamical Systems," in *Handbook of Differential Equations: Ordinary Differential Equations*. Elsevier, 2006, vol. 2, pp. 239–357.

[28] B. Kamiński, P. Prałat, and F. Théberge, "Artificial Benchmark for Community Detection (ABCD)—Fast random graph model with community structure," *Netw. Sci.*, vol. 9, no. 2, pp. 153–178, Jun. 2021.

[29] A. Lancichinetti, S. Fortunato, and F. Radicchi, "Benchmark graphs for testing community detection algorithms," *Phys. Rev. E*, vol. 78, no. 4, p. 046110, Oct. 2008.