

On the Use of Eigenvector Centrality for Cooperative Streaming

Luca Baldesi*, Leonardo Maccari*, Renato Lo Cigno*

*Department of Information Engineering and Computer Science, University of Trento, Italy

{luca.baldesi, leonardo.maccari, renato.locigno}@disi.unitn.it

Abstract—The timely and efficient cooperative distribution of a streamlined content in a communication network is a key feature for many applications and services. One of the unsolved problems is the assignment of transmission rates to nodes given the constraints imposed by the topology, so that all nodes receive the stream with the minimal global use of resources. This paper addresses the problem exploiting the notion of eigenvector centrality. It shows that the problem can be solved efficiently in a distributed way if every node is aware of the full network topology and that in certain cases only local information on the network graph is sufficient.

Index Terms—Live Streaming, Rate Allocation, Cooperative Streaming, Video Distribution

I. INTRODUCTION

Distributed cooperative delivery of contents in networks is essential for several applications, in particular for peer-to-peer live video streaming. Live streaming tolerates some packet loss but puts a strong focus on timely delivery and on a balanced distribution of delays among the peers [1], [2], [3].

Consider a generic network represented by a graph $G(\mathcal{V}, \mathcal{E})$; $v_i \in \mathcal{V}$ is the generic node, and $e_{ij} \in \mathcal{E}$ means there is a connection between v_i and v_j , thus the neighborhood of v_i is defined as $\mathcal{N}_i = \{v_j \in \mathcal{V} : e_{ij} \in \mathcal{E}\}$ (in the rest of the paper, we use capital letters for matrices and vectors, small case letters for scalars, and calligraphic letters for sets). G can be the graph describing a peer-to-peer network, a virtual network in a data-center, a physical infrastructure like a wireless mesh network. The streaming source is a network node that generates a packetized content with a rate of one packet every τ s and sends one copy of each packet to a set of nodes chosen from its neighborhood.

Each node v_i stores the received packets in a temporary buffer, and every τ_i s it picks another node v_j among its neighbors and a packet stored in the buffer and forward the latter to the former with a unicast transmission. This way packets percolate from the source to all the nodes in the network. We assume that neighbor nodes share the composition of their buffer one another, so that when a node receives a new packet it knows for which of its neighbors this packet is *useful* (the neighbor does not own it). This assumption is realistic because the buffer maps can be easily piggybacked on content. The same assumption is at the base of all the works we compare our results with [2], [4].

This work has been partially funded by the European Commission, H2020-ICT-2015 Programme, Grant Number 688768 'netCommons' (Network Infrastructure as Commons)

$G(\mathcal{V}, \mathcal{E})$	Network graph
v_i	i th node in \mathcal{V}
\mathbf{A}	Normalized adjacency matrix of $G(\mathcal{V}, \mathcal{E})$
τ	Packets generation interval
τ_i	Packets transmission interval of v_i
θ_i	Number of transmissions of v_i per τ ($\frac{\tau_i}{\tau}$)
Θ	Vector of all θ_i
\mathbf{X}	Normalized dominant eigenvector of \mathbf{A}
$\mathbf{I}_{\mathbf{X}}$	Diagonal matrix of \mathbf{X}
$\mathbf{1}$	One vector
\mathbf{A}^*	Reception-equal probability transition matrix for $G(\mathcal{V}, \mathcal{E})$
Θ^*	Reception-equal column of node transmission rates per τ

TABLE I: Nomenclature

The problem we tackle is twofold: i) define the values of τ_i per each node; ii) define a strategy for the choice of the target neighbor. The constraints we pose are that at steady state every node receives the entire content, and at the same time the overall resources used to achieve this goal are minimal.

We show that this can be obtained exploiting the *eigenvector centrality* [5] of the normalized adjacency matrix describing G . We implement the solution extending SSSim [6]¹ and we measure the performance improvement in terms of reduced packet loss and delivery delay compared to other solutions currently in use. Finally we highlight that the eigenvector centrality can be efficiently estimated with distributed algorithms, so that the implementation of our solution in real systems is feasible and can scale up to systems made of hundreds of nodes.

II. PROBLEM FORMULATION

We call a *distribution process* the couple (Θ, \mathbf{A}) defined as follows. Given the packet generation interval τ , we call θ_i the ratio between τ and the packet transmission interval τ_i of v_i : $\theta_i = \frac{\tau}{\tau_i}$, and Θ the column vector of all θ_i . Given $G(\mathcal{V}, \mathcal{E})$, the *probability transition matrix* $\mathbf{A} \in [0, 1]^{|\mathcal{V}| \times |\mathcal{V}|}$ is a column-stochastic adjacency matrix in which each element a_{ij} is 0 if $v_i \notin \mathcal{N}_j$, else a_{ij} expresses the probability that node v_j selects the neighbor v_i as the target for the transmission of a packet following any suitable discrete probability distribution. Every τ s, each node v_i receives on average a number of packets $\sum_j a_{ij} \theta_j$, where $a_{ij} \theta_j$ is the average number of packets received by v_i from v_j every τ s. Hence, we introduce the column vector $\Phi = \mathbf{A} \Theta$, whose elements ϕ_i express the average number of packets that v_i receives per τ s. We consider

¹The extended simulator is available at <https://ans.disi.unitn.it/redmine/projects/sssimsim>

only strongly connected graphs (there is a path from v_i to $v_j \forall (i, j)$), otherwise some node can be isolated and have $\phi_i = 0$ by construction. Note that v_i can send on average more than one packet from its buffer per τ , so the transmission rate of v_i can be higher than the rate of the source.

We say that a distribution process is *complete* if every node receives all the packets of the stream: since the original stream is composed of one packet every τ s, v_i should receive on average at least one packet per each τ interval, thus in a complete distribution $\phi_i \geq 1 \forall i$. If $|\mathcal{V}|$ packets must be received every τ s, a necessary condition for completeness is that at least $|\mathcal{V}|$ packets are sent by the nodes. We say that a distribution process is *sustainable* if:

$$\frac{1}{|\mathcal{V}|} \sum_{i=1}^{|\mathcal{V}|} \theta_i \geq 1 \quad (1)$$

If the equality holds in Eq. (1), then the distribution process is sustainable and *minimal*, meaning that it sends the minimum number of packets in the network. However, there is no guarantee that a sustainable process is also complete, but only that the overall amount of packets disseminated is large enough to potentially deliver all the content to all the nodes. The completeness of the distribution process depends also on how nodes send information one another, that is, on the choice of the transition matrix.

To clarify this key concept we introduce a widely-used strategy for the definition of \mathbf{A} and a simple example. Let \mathbf{N} be the column vector whose elements $n_i = |\mathcal{N}_i|$ are the neighborhood sizes of node i , $a_{ij} = \frac{1}{n_j}$ for every non-zero element. We call a matrix with this property “column-uniform”. Let also $\theta_i = 1 \forall i$ ($\Theta = \mathbf{1}$) so that every node sends packets with the same rate, choosing the target neighbors with uniform probability. This choice for the transition matrix, coupled with a Latest-Useful packet choice (i.e., the sender transmits the newest useful packet) is called *high-bandwidth peer first* (HPF) in [2]. A decentralized HPF content distribution has been widely studied as it can be used on a vast range of application contexts and we pick it as our base comparison.

Now consider the simple example given by the network in Fig. 1. The source sends one copy of each packet to some other node, and by definition it does not receive any packet, so it has no incoming links. Eq. (2) reports the corresponding \mathbf{A} and Φ under the assumption that the adopted strategy is HPF and that $\theta_i = 1 \forall i$ (the distribution is minimal). Since the source does not receive packets, we remove it from \mathbf{A} , so the graph remains strongly connected. The distribution is not complete: some of the values of Φ are lower than 1, and some nodes fail to receive the entire stream.

$$\mathbf{A} = \begin{bmatrix} 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & \frac{1}{3} \\ 0 & 0 & 0 & \frac{1}{2} & 0 & \frac{1}{3} \\ 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{3} \\ \frac{1}{2} & 0 & 0 & \frac{1}{2} & 1 & 0 \end{bmatrix}; \Phi = \begin{bmatrix} 0.5 \\ 1.333 \\ 0.833 \\ 0.5 \\ 0.83 \\ 2 \end{bmatrix} \quad (2)$$

To achieve a complete distribution one could use a scalar multiplier α to increase the transmission rates, so that the input vector becomes $\alpha\Theta$ with α large enough in order to have




Fig. 1: Sample network of six nodes, source depicted in grey.

$\phi_i \geq 1 \forall i$. This solution is easy to achieve because it does not touch \mathbf{A} but it is not minimal anymore. A better approach is to modify Θ and \mathbf{A} to guarantee that the distribution is both complete and minimal.

We can now formally state the goal of the paper: given $G(\mathcal{V}, \mathcal{E})$ and τ , find a couple (Θ^*, \mathbf{A}^*) that satisfies:

$$\mathbf{1} = \Phi = \mathbf{A}^* \Theta^* \quad (3)$$

$$|\Theta^*| = |\mathbf{1}| \quad (4)$$

$$\mathbf{1}^T \mathbf{A}^* = \mathbf{1}^T \quad (5)$$

$$a_{ij} = 0 \iff a_{ij}^* = 0 \quad (6)$$

Eq. (3) states that the distribution is complete, for this reason we call this approach a *reception-equal* strategy; Eq. (4) states that the distribution is sustainable and minimal; Eq. (5) states that \mathbf{A}^* must be stochastic; and finally Eq. (6) states that we do not add or drop links present in the original graph, as we are working on a given topology.

We prove that (Θ^*, \mathbf{A}^*) exists and can be computed by each node and we present numerical results that show the effective gains of this methodology.

III. THE RECEPTION-EQUAL DISTRIBUTION PROCESS

Theorem 1. *Let $G(\mathcal{V}, \mathcal{E})$ be a strongly connected network graph and let \mathbf{A} be its associated column-stochastic transmission matrix. Then it is possible to find \mathbf{A}^* and Θ^* such that the conditions given by Eqs. (3) to (6) hold.*

Proof. An adjacency matrix of a strongly connected graph is irreducible and each column has all non negative entries with at least one non-zero element. Since \mathbf{A} is also column stochastic, we can apply the Perron-Frobenius theorem and state that the largest eigenvalue of \mathbf{A} is 1, and it exists a corresponding eigenvector \mathbf{X} with all strictly positive entries that sum to 1:

$$\mathbf{A}\mathbf{X} = \mathbf{X}; \mathbf{X}^T \mathbf{1} = 1 \quad (7)$$

Given a generic array \mathbf{Y} , let \mathbf{I}_Y be the diagonal matrix with the same size of \mathbf{Y} with the elements of \mathbf{Y} on the diagonal: $\mathbf{Y} = \mathbf{I}_Y \mathbf{1}$. Then we have:

$$\mathbf{A}\mathbf{I}_X \mathbf{1} = \mathbf{X} \quad (8)$$

and, since the elements of \mathbf{X} are strictly positive:

$$\mathbf{I}_X^{-1} \mathbf{A}\mathbf{I}_X \mathbf{1} = \mathbf{I}_X^{-1} \mathbf{X} = \mathbf{1} \quad (9)$$

If we choose $\Theta^* = (\mathbf{1}^T \mathbf{I}_X^{-1} \mathbf{A}\mathbf{I}_X)^T$ and $\mathbf{A}^* = \mathbf{I}_X^{-1} \mathbf{A}\mathbf{I}_X \mathbf{I}_{\Theta^*}^{-1}$, then, Eq. (3) is satisfied by:

$$\mathbf{A}^* \Theta^* = \mathbf{I}_X^{-1} \mathbf{A}\mathbf{I}_X \mathbf{I}_{\Theta^*}^{-1} \Theta^* = \mathbf{I}_X^{-1} \mathbf{A}\mathbf{I}_X \mathbf{1} = \mathbf{I}_X^{-1} \mathbf{A}\mathbf{X} = \mathbf{I}_X^{-1} \mathbf{X} = \mathbf{1} \quad (10)$$

Eq. (4) is satisfied by:

$$|\Theta^*| = \Theta^{*T} \mathbf{1} = \mathbf{1}^T \mathbf{I}_X^{-1} \mathbf{A} \mathbf{I}_X \mathbf{1} = \mathbf{1}^T \mathbf{I}_X^{-1} \mathbf{A} \mathbf{X} = \mathbf{1}^T \mathbf{I}_X^{-1} \mathbf{X} = \mathbf{1}^T \mathbf{1} = |\mathbf{1}| \quad (11)$$

Eq. (5) is satisfied by:

$$\mathbf{1}^T \mathbf{A}^* = \mathbf{1}^T \mathbf{I}_X^{-1} \mathbf{A} \mathbf{I}_X \mathbf{I}_{\Theta^*}^{-1} = \Theta^{*T} \mathbf{I}_{\Theta^*}^{-1} = \mathbf{1}^T \quad (12)$$

To verify Eq. (6) let us explicit each element of Θ^* and \mathbf{A}^* :

$$\theta_j^* = \sum_{k=1}^{|\mathcal{V}|} \frac{a_{kj} x_j}{x_k}; \quad a_{ij}^* = a_{ij} \frac{x_j}{x_i \theta_j^*} \quad (13)$$

since elements of \mathbf{X} are positive and at least one element per column in \mathbf{A} is larger than zero, then $\forall j \theta_j^* > 0$. Therefore, $\forall i, j a_{ij}^* \geq 0$ and $a_{ij}^* = 0 \iff a_{ij} = 0$ which satisfies Eq. (6). \square

Eq. (13) defines the new matrix \mathbf{A}^* and the vector Θ^* that produce the steady-state reception-equal minimal distribution. Every node v_j that knows the full matrix \mathbf{A} is able to compute \mathbf{X} and thus θ_j^* and the column $a_{\cdot,j}^*$. This way v_j can choose the value of τ_j and the values of the probabilities to send a packet to each neighbor that guarantee a reception-equal behavior.

In some cases \mathbf{A} is not known to all the nodes, or it is too large to compute \mathbf{X} , however there are two common situations in which the explicit computation of \mathbf{X} is not required. The first case follows from this proposition.

Proposition 1. *Let $G(\mathcal{V}, \mathcal{E})$ be an unweighted, non directed graph and \mathbf{A} its stochastic, column-uniform associated transmission matrix. Then $\mathbf{I}_X^{-1} \mathbf{A} \mathbf{I}_X = \mathbf{A}^T$.*

Proof. Let z_{ij} a binary variable indicating if $a_{ij} = 0$, since the graph is non directed, $z_{ij} = z_{ji}$. Each element of \mathbf{A} is defined by $a_{ij} = \frac{z_{ij}}{n_j}$. Let \mathbf{N} be the column made of all the values of n_i , then \mathbf{N} is an eigenvector of \mathbf{A} , $\mathbf{A} \mathbf{N} = \mathbf{N}$.

For the uniqueness statement of the Perron-Frobenius theorem $\exists k \in \mathbb{R} : \mathbf{N} = k \mathbf{X}$ hence, the element (i, j) of $\mathbf{I}_X^{-1} \mathbf{A} \mathbf{I}_X$ is given by:

$$\{\mathbf{I}_X^{-1} \mathbf{A} \mathbf{I}_X\}_{ij} = a_{ij} \frac{x_j}{x_i} = \frac{z_{ij} x_j}{n_j x_i} = \frac{z_{ij} x_j}{k x_j x_i} = \quad (14)$$

$$\frac{z_{ij}}{k x_i} = \frac{z_{ji}}{k x_i} = \frac{z_{ji}}{n_i} = a_{ji} \quad (15)$$

\square

Proposition 1 says that given a stochastic matrix \mathbf{A} with symmetric zero-elements, there exist a similarity relationship between \mathbf{A} and \mathbf{A}^T and the corresponding change of base matrix is \mathbf{I}_X . In this case \mathbf{I}_X is said to be a symmetrizer of \mathbf{A} . While there is a large body of work about symmetrizer matrices, to the best of our knowledge this simple and elegant result was never reported in the networking literature.

Corollary 1. *From Theorem 1 and Proposition 1, it follows that if $G(\mathcal{V}, \mathcal{E})$ is an unweighted, undirected graph and \mathbf{A} its associated stochastic column-uniform transmission matrix then $\Theta^* = \mathbf{A} \mathbf{1}$ and $\mathbf{A}^* = \mathbf{A}^T \mathbf{I}_{\Theta^*}^{-1}$.*

One fundamental consequence of Corollary 1 is that computation of \mathbf{X} is no longer needed and that the whole \mathbf{A} needs not

to be known to every node: v_i needs to know only the i th row and column of \mathbf{A} . The i th column is known by construction, since v_i knows all its neighbors, the i th row can be computed if v_i knows also $n_j \forall v_j \in \mathcal{N}_i$ (the size of the neighborhood of all its neighbors, which can easily be piggybacked in packets containing buffer maps or data packets). Then v_i will then be able to compute the transmission rate θ_i^* and the i th column of the probability transition matrix.

The second case in which the computation can be performed without the knowledge of the whole \mathbf{A} is when the graph is stable enough so that the nodes can use a gossiping communication protocol to compute \mathbf{X} in a distributed way. Recently, a distributed eigenvector (or PageRank) centrality computation algorithm based on a gossiping protocol has been proposed [7]. Using such protocol, nodes exchange only local information and converge to the exact computation of their own centrality. Once this value is computed, we can assume the neighbor nodes will exchange their centrality and thus be able to compute Eq. (13).

The reception-equal distribution process also guarantees a natural upper bound to the amount of information a single node can transmit, as the following proposition shows.

Proposition 2. *From Theorem 1, it follows:*

$$\theta_j^* \leq n_j$$

Proof. From Eq. (13) and Eq. (7) we have:

$$\theta_j^* = \sum_{k=1}^{|\mathcal{V}|} \frac{a_{kj} x_j}{x_k} = \sum_{k=1}^{|\mathcal{V}|} \frac{a_{kj} x_j}{\sum_{i=1}^{|\mathcal{V}|} a_{ki} x_i} = \quad (16)$$

$$\sum_{k=1}^{|\mathcal{V}|} \frac{a_{kj} x_j}{a_{kj} x_j + \sum_{i=1, i \neq j}^{|\mathcal{V}|} a_{ki} x_i} \quad (17)$$

Each term in the summation is lower or equal than 1 and it is 0 if and only if $a_{kj} = 0$. Since $a_{kj} = 0 \iff z_{kj} = 0$ then $\frac{a_{kj} x_j}{a_{kj} x_j + \sum_{i=1, i \neq j}^{|\mathcal{V}|} a_{ki} x_i} \leq z_{kj}$ and we have:

$$\theta_j^* \leq \sum_{k=1}^{|\mathcal{V}|} z_{kj} = n_j \quad (18)$$

\square

Proposition 2 states that the information a node has to send is upper bounded by the number of its neighbors: no node i ever sends on average more than n_i packets per τ_i s. This is the mathematical counterpart of the intuitive idea that if a node sends every packet to all its neighbors there is nothing more to send, but its formal proof in a distribution process is critical for real implementations.

IV. NUMERICAL RESULTS

Using the modified version of SSSim we compare the performance of three strategies: reception-equal (R-E for brevity), HPF, and a third strategy introduced by Wu and Li in [4]. Wu and Li assign transmission rates to nodes solving an optimization problem that makes the distribution minimal, but with limited buffer size can still introduce packet loss. This is a computationally intensive centralized solution. This

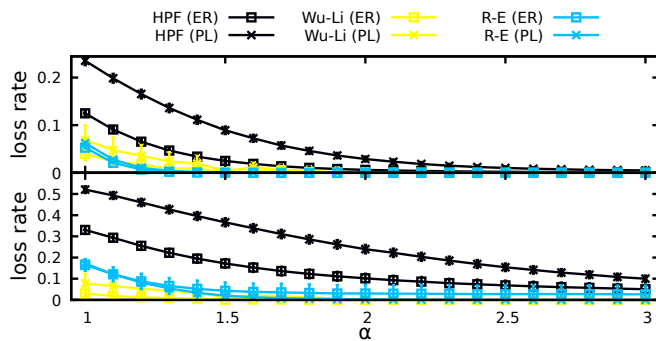


Fig. 2: Loss rate with 99% confidence interval as a function of α with $|\mathcal{V}|=50$; $|\mathcal{E}|=264$ (top) and $|\mathcal{E}|=96$ (bottom)

centralized solution should be interpreted as a lower bound (on a 32 cores server with 64G RAM we were able to find the optimal solution up to 200 nodes).

In theory, an ideal protocol guarantees that at every τ_i , every node has one packet that some neighbor needs [1]. In practice this is not true due to, e.g., limited buffer size or congestion. We set the buffer size to 64 packets in our simulations. It is clear that with minimal resources (when Eq. (4) is satisfied) every lost transmission chance translates into a lost packet for some node. For this reason, and to take into account loss or corruption of packets, real protocols (and our simulator) use a multiplication factor α , so that each node sends $\alpha\theta_i$ packets per τ . We test the minimal distribution ($\alpha = 1$) and also the non-minimal ones ($\alpha > 1$).

We test the distribution strategies against four sets of network graphs. Each set is composed by ten different graphs with homogeneous type, number of nodes and links. We investigate synthetic graphs (Erdős graphs, labelled ER and Power-Law graphs generated using the Barabási-Albert model, labelled PL) with 50 nodes and 96 or 264 links to investigate different network connectivity properties. For each streaming session, at steady state we measure the average loss at each node, and the average arrival delay of the packets, measured in number of intervals τ from the generation of the packet to its arrival. Results are shown in Figs. 2 and 3; please note that the strategies obtain very similar results for both Erdős and Power-Law graphs, hence, curves may overlap in the figures.

Figure 2 shows the loss rate varying α and the topology type. The upper plot refers to highly connected networks and the lower one to graphs with less links. The improvement provided by R-E over HPF is clear for topology types and link densities. $\alpha = 1.5$ is sufficient for R-E to obtain zero loss. The R-E strategy performs very close to Wu-Li, and in one point even better, due to the fact that Wu-Li overloads some nodes in the shortest paths, and, given a limited buffer, it loses even more packets.

Figure 3 shows the loss rate varying $|\mathcal{V}|$ and keeping $\alpha = 1.5$. Increasing the size of the network quickly impacts the loss rate for HPF, while it does not affect the R-E strategy. Again, the performance with a limited buffer size is even better than Wu-Li, that requires a centralized global optimization and we could not solve for networks larger than 200 nodes.

Figure 4 shows the delay in one particular example, with




Fig. 3: Loss rate with 99% confidence interval as a function of \mathcal{E} , $\alpha = 1.5$, $\frac{|\mathcal{E}|}{|\mathcal{V}|} \simeq 5$.




Fig. 4: Reception delay histogram for an Erdős graph with $|\mathcal{V}|=50$, $|\mathcal{E}|=264$ and $\alpha = 2.5$

$\alpha = 2.5$, the minimum that guarantees a 100% delivery of packets for all strategies. If loss rate is not zero, then the potential delay of lost packets can not be included, and would make the comparison meaningless. R-E outperforms HPF which maintains a longer tail even in this non particularly challenging scenario in which every node uploads 2.5 times the traffic it downloads. The performances of Wu-Li are comparable, if not slightly worse than R-E.

V. CONCLUSIONS

This paper presented a reception-equal rate allocation strategy to guarantee a complete distribution minimal resources. The strategy is based on the analysis of the topology of the network (physical or virtual) that supports the streaming application, exploiting the notion of eigenvector centrality. The theoretical analysis together with the simulation study that show that the scheme can be easily implemented and that it maintains a very large advantage compared to other strategies.

REFERENCES

- [1] J. A. Oliveira, F. Prado, F. M. de Lima, M. Rubinstein, and A. Sztajnberg, "Improving peer neighborhood on P2P video distribution networks using Push/Pull protocol," *Computer Communications*, vol. 61, pp. 17–33, 2015.
- [2] J. Zhang, C. Yang, and X. Zhang, "A high-bandwidth live streaming model in mesh-based Peer-to-Peer networks," *IEEE Communications Letters*, vol. 20, no. 12, pp. 2390–2393, 2016.
- [3] B. Bellalta, E. Belyaev, M. Jonsson, and A. Vinel, "Performance evaluation of IEEE 802.11 p-enabled vehicular video surveillance system," *IEEE Communications Letters*, vol. 18, no. 4, pp. 708–711, 2014.
- [4] C. Wu and B. Li, "Optimal rate allocation in overlay content distribution," in *6th Int. IFIP Networking Conf. (NETWORKING)*, LNCS, vol. 4479, pp. 678–690, Atlanta, GA, May 2007.
- [5] P. Bonacich, "Some unique properties of eigenvector centrality," *Social Networks*, vol. 29, no. 4, pp. 555–564, 2007.
- [6] L. Abeni, C. Kiraly, and R. Lo Cigno, "SSSim: a Simple and Scalable Simulator for P2P Streaming Systems," in *14th IEEE Workshop on Computer Aided Modeling and Design of Comm. Links and Networks (CAMAD)*, pp 1–6, June 2009.
- [7] H. Ishii and R. Tempo, "Distributed randomized algorithms for the PageRank computation," *IEEE Trans. on Automatic Control*, vol. 55, no. 9, pp. 1987–2002, 2010.