# Distributed Ranking in Networks with Limited Memory and Communication

Kyomin Jung
KAIST
Daejeon, Korea
kyomin@kaist.edu

Bo Young Kim
KAIST
Daejeon, Korea
combicola@kaist.ac.kr

Milan Vojnović
Microsoft Research
Cambridge, UK
milanv@microsoft.com

***Abstract*—We study a basic information ranking problem in networks where each node holds an individual preference over a set of items and the goal for each node is to identify a sorted list of items with the largest aggregate preference. We would like to achieve this with a fully decentralized algorithm that uses a limited per-node memory and limited pair-wise communications. We show how this problem can be reduced to a plurality selection problem where the goal for each node is to identify an item with the largest aggregate ranking score, and show that solving the reduced problem solves the original ranking problem with high probability. Then we introduce a simple and natural plurality selection algorithm for the selection over $m > 1$ items that uses only $\log_2(m) + 1$ bits of per-node memory and per pair-wise communication. We prove correctness of the algorithm with high probability as the number of nodes grows large for the case when each node communicates with any other node, and establish tight convergence time bounds.**

**The information ranking problem studied in this paper is a basic ranking problem that arises in various applications such as sorting elements in distributed computing systems, parallel databases, and may as well serve as a model of decentralized inference and opinion formation in distributed environments.**

## I. Introduction

In this paper we consider information ranking in networks where each node has an individual preference over a set of items and the goal for each node is to identify a sorted list of items with the largest aggregate preference (*a top rank list*). This is to be achieved in a decentralized fashion based on limited information exchanged at pair-wise node communications and limited amount of states stored by individual nodes. Such decentralized network information aggregation problems recently attracted quite some research interest, e.g. quantized distributed averaging [10], composite hypotheses testing [2] and consensus problem [2], [6], and are of interest in applications in the context of distributed computing systems, sensor systems, and online social networks.

Specifically, we consider a *ranking problem* over a set of $m > 1$ items where each node has a ranking score associated with each item and then each item is associated with *aggregate ranking score* defined as the sum of ranking scores of this item over all nodes in the network. The goal for each node is to identify a set of $l$ items with the largest aggregate ranking scores and sort them in decreasing order with respect to their aggregate ranking scores. We observe that for the special case of lists of size $l = 1$, the problem corresponds to a *plurality selection* where the goal for each node is to identify an item with the largest aggregate ranking score. Previously-studied consensus problem [19], [6], [21] is a special case of plurality selection where the selection is over a set of two items. Notice that the given setting admits various interpretations. For example, in a distributed storage system, an item may refer to a key value of a dataset where preference of a node for an item may reflect the number of records of given key value stored in a local storage associated with this node. In an online service system, an item may correspond to an opinion or a state of nature and the node may refer to a user that has a preference over the set of opinions or states of nature. Note also that the plurality selection problem may be interpreted as identifying a mode of an empirical distribution that is defined by the frequency counts of items that are partitioned across nodes in a distributed network system. Our work could be seen as studying the following two broadly defined questions: *Q1) How one could best design a fully decentralized algorithm for the information ranking problem?* and *Q2) What convergence time bounds could be achieved for the plurality selection problem and the ranking problem in distributed environments under given node memory and communication constraints?*

In this paper we show how the ranking problem over a set of items can be reduced to a plurality selection problem over a set ordered lists of items, and introduce and study a simple and natural decentralized plurality selection algorithm. First, for the ranking problem we propose a simple and fully decentralized reduction algorithm where each node chooses an ordered list of $l$ items called a *local rank list* in a way that solving the reduced plurality selection problem over the local rank lists solves the original ranking problem with high probability. For the plurality selection algorithm, we introduce and study a simple and natural algorithm that uses a simple automata with only $2m$ states of per-node memory and per pair-wise communication for plurality selection over $m \geq 2$ items. The state of a node indicates the identity of an item that the node would select as a plurality winner according to the current node's belief and one extra bit that may be interpreted as the node's confidence in its current belief. The state held by a node can be encoded with $\log_2(m)+1$ bits. For the case of majority selection over two items, this algorithm includes the ternary protocol studied in [19] as a special case. We provide analysis of the correctness and efficiency of the plurality selection algorithm for the case of complete graphs where each node communicates with any other node at a given rate. The assumption of the complete graphs may be enforced by the system design in a similar manner as found in the context of peer-to-peer networks where peering relations between nodes are pseudo-random and may be justified in general systems where node interactions are not restricted by a specific network topology. This assumption was admitted in many prior works, e.g. [17], [19].

### A. Related Work

Our work relates to that on information aggregation in networks where a variety of problems were studied [3], [10], [4], [2], [16], [9], [11], [12], [1], [5], [21]. Our work is also related to the selection and sorting problems considered in the context of distributed systems, e.g. [14], [17], and some early work on hypothesis testing with limited memory in the context of information theory [7], [13], where we consider a network setting unlike to previous work, as well as that of interacting particle systems such as the standard voter model [15], [18].

The work that is closely related to ours is that on the consensus problem [19] which is equivalent to the majority selection over a set of two items. In [19], the authors studied a ternary algorithm which uses three states for per-node memory and per pair-wise communication. It was shown that for the case of the complete graph of $n$ nodes, the algorithm fails to correctly identify the majority item with probability of error that diminishes to zero exponentially fast with $n$ and that the convergence time is $\Theta(\log(n))$. Our plurality selection algorithm encompasses that in [19] as a special case. A quaternary algorithm for two items under a two-way communication model was proposed in [2] which guarantees correct selection with probability 1 for any finite connected graph and whose expected convergence time was showed to be $O(\frac{1}{\beta(n)}\log(n))$, where $\beta(n) > 0$ is a lower bound on the absolute eigenvalues of some matrices that characterize the rates of node pair-wise interactions [6]; for the case of the complete graph of $n$ nodes, $\beta(n)$ is independent of $n$ and is equal to the fractional margin by which the majority item is preferred over the other item. Our work considers more general problem of plurality selection for any number of items compared with all the aforementioned prior works.

### B. Summary of our Main Contributions

Our main contributions can be summarized as follows:

1) Distributed ranking: we propose a simple and decentralized randomized algorithm for distributed ranking which reduces the problem of ranking to a plurality selection problem over the local rank lists (Section III). We prove that the probability that the algorithm does not output a correct answer diminishes exponentially fast in the number of nodes.
2) Plurality selection: we introduce a simple and natural plurality selection algorithm for the general case of $m > 1$ items. The algorithm uses only $2m$ states of per-node memory, and per pair-wise node communication (Section IV). We show that under the limit dynamics of many nodes, the algorithm converges to a correct answer and show that the $\delta$-convergence time, defined as the earliest time at which for given $\delta \in (0,1)$, all but at most $\delta$ portion of nodes is in a correct state, is $\Theta(m(\log(1/\delta) + \log(1/\gamma)))$, where $\gamma > 0$ is the gap between the aggregate ranking score of a plurality item and the largest ranking score of a non-plurality item. Note that if $\delta = 1/n$, which may be interpreted as requiring all but one node to be correct, the convergence time is $\Theta(m\log(n))$, for every fixed gap $\gamma$. Therefore, we have that the limit dynamics suggests a convergence time that is logarithmic in the number of nodes $n$. Finally we establish a convergence result that relates the original stochastic system with the limit dynamics, and show that in our original stochastic system, at a $\delta$-convergence time with $\delta = 1/n^\alpha$ for sufficiently small $\alpha > 0$, $1 - O(n^{-\alpha})$ fraction of nodes are in a correct state with probability $1 - O(1/n)$.

Due to space limitations, we omitted all the proofs and simulation results, which are available in the online companion of this paper [8].

## II. Problem Formulation

In this section, we introduce notations and formulate the problem. We consider a network represented by a graph $G = (V, E)$ where the set of vertices $V = [n] = \{1, 2, \ldots, n\}$ corresponds to a set of $n > 1$ nodes and the edges corresponds to links between nodes. For each node $i \in V$, a node $j \in V$ is said to be a neighbor of node $i$ if $(i, j) \in E$. For our analysis of plurality selection algorithm, we will assume that the graph is complete, i.e. $(i, j) \in E$, for every pair of nodes $i, j \in V$, $i \neq j$. Each node is assumed to communicate with each of its neighbors at some time instances. Specifically, we admit the standard asynchronous communication model (e.g. [3], [2], [19], [6]) where each edge between a pair of nodes $(i, j) \in E$ is activated at instances of a Poisson process of rate $\lambda_{i,j} \geq 0$. In an instance of such a communication, we will call $i$ an *observer* node and $j$ *a contacted* node. We denote with $I = [m] = \{1, 2, \ldots, m\}$ the set of $m \geq 2$ items. A node $i$'s *state* denotes some bits of information about the preference of node $i$ over items. Initially, node $i$'s state is the original preference of node $i$. The state of $i$ changes by the algorithm as the communication proceeds. At each communication instance, the observer node changes its state based on the state of the contacted node and its own. As a special case, we consider the complete graph and assume that contact rates by individual nodes are identical (without loss of generality assumed to be equal to 1), thus $\lambda_{i,j} = 1/(n-1)$, for every $i, j \in V$, $i \neq j$.

The preference of each node $i \in V$ over items is described by a vector of ranking scores $\vec{v}_i = (v_{i,1}, v_{i,2}, \ldots, v_{i,m})$ such that $v_{i,j} \geq 0$ quantifies the preference of node $i$ for item $j$, and we assume the preference vector is normalized such that $\sum_{j=1}^{m} v_{i,j} = 1$, for every $i \in V$. In particular, if each node prefers exactly one item, we have that for every $i \in V$,

$v_{i,j} = 1$ for some $j \in I$ and $v_{i,h} = 0$, otherwise. In this case, we call $j$ the *selected item* of node $i$. We let $v_j$ denote the *aggregate ranking score* of item $j \in I$ over all nodes, defined by $v_j = \frac{1}{n}\sum_{i=1}^{n} v_{i,j}$. Without loss of generality, we assume that items are enumerated in decreasing order of aggregate ranking scores, so that $v_1 \geq v_2 \geq \cdots \geq v_m$.

A *top rank list* is defined as an ordered list of $l$ items ($1 \leq l \leq m$) sorted in decreasing order of their aggregate ranking scores, such that each item from this list has an aggregate ranking score that is at least as large as the aggregate ranking score of any item that is not in the list. We denote with $k$ the number of items with the largest aggregate ranking score, i.e. $v_1 = \cdots = v_k > v_{k+1} \geq \cdots \geq v_m$. An item $j$ is said to be a *plurality item* if $1 \leq j \leq k$ and is said to be a *non-plurality item* otherwise.

## III. Randomized Distributed Ranking

We present a fully decentralized randomized ranking algorithm that is defined as follows. Given the rank list size $l (1 \leq l \leq m)$, let $p$ be a probability distribution over the set $[l]$ that is assumed to be strictly decreasing.[1] In the following algorithm we consider $\vec{v}_i$ for each node $i$ as a probability distribution over items.

---

**Randomized Distributed Ranking**

1) Each node $i$ creates a local rank list of $l$ distinct items as follows.
2) Node $i$ picks his *selected* item $s_i$ by sampling from the distribution $\vec{v}_i$ and then assigns a set of $l$ distinct items to his local rank list as follows:
   - Item $s_i$ is assigned rank $r$ where $r$ is a random sample from the distribution $p$.
   - Other ranks are assigned items by uniform random sampling without replacement from the set of items $I \setminus \{s_i\}$.
3) Run a plurality selection algorithm over the set of local rank lists of individual nodes.

---

Notice that each node constructs a ranking of $l$ items using only the knowledge of identities of items in the set $I$ and does not use any global knowledge about preferences over items by other nodes. The set of all local rank lists of size $l$ contains $m(m-1)\cdots(m-l+1) = O(m^l)$ elements. The local rank list by a node is constructed by a random assignment that biases the selected item of the node to be ranked higher than other items, which allows for the inference of a top rank list with high probability as we show in theorem 1 below.

Let $x(S)$ be the random variable indicating the fraction of nodes that selects a local rank list $S$ under the above randomized ranking procedure, and let $I_l$ denote the set of all local rank lists of size $l$. Define $x^* = \max_{S \in I_l} \mathbb{E}(x(S))$ and let $T_l = \{S \in I_l : \mathbb{E}(x(S)) = x^*\}$ and $B_l = I_l \setminus T_l$. We say that the aggregate ranking scores $\{v_j\}$ are *$\epsilon$-separated*, if there exists $\epsilon > 0$ such that for every $1 \leq i < m$, either $v_i = v_{i+1}$ or $v_i \geq v_{i+1} + \epsilon$.

***Theorem*** *1:* Suppose that the aggregate ranking scores are $\epsilon$-separated. Then, for every $1 \leq l \leq m$ and a decreasing distribution $p$, there exists $\xi \geq \epsilon[\min_j p(j)]/3$ and a positive constant $C_\xi$ which depends only on $\xi$ such that

$$\begin{aligned} & \mathbb{P}(x(T) \geq x(B) + \xi, \text{ for all } T \in T_l, B \in B_l) \\ \geq \quad & 1 - 2\frac{m!}{(m-l)!}\exp\left(-C_\xi n\right). \end{aligned}$$

The proof of the theorem relies on combining the rearrangement inequality with the concentration of a sum of independent random variables. The theorem holds for $C_\xi = \xi^2/27$. Hence theorem 1 tells us that our random assignment guarantees correct identification of a top rank list with high probability provided that $\epsilon = \Omega\left(\sqrt{\frac{\log(n)}{n}}\right)$.

## IV. A Plurality Selection Algorithm

In this section, we study a simple and natural plurality selection algorithm over a set of $m > 1$ items. We will first present the algorithm and then describe the underlying stochastic system that describes the evolution of the states under the given algorithm and the assumed communication model. We will then consider a limit dynamics that is a system of ordinary differential equations that are justified by appropriate scaling of the original stochastic system. We will first characterize the convergence time under this limit dynamics. Then we will provide a convergence result of the original stochastic system to the limit dynamics that justifies the convergence time characterization derived under the limit dynamics to the original stochastic system.

In our plurality selection algorithm, at each time instance, each node will be in a state $(j, s)$ where $j \in I$ indicates the currently preferred item and $s$ is either 0 (weak) or 1 (strong). The extra bit $s$ can be interpreted as the node's confidence in its current belief. We call a state $(j, 0)$ to be a *weak* state $j$ and $(j, 1)$ a *strong* $j$ state, for every $j \in I$. Hence each node holds a memory of $2m$ states, and at each communication between a pair of nodes, the nodes exchange one of $2m$ states. The following describes our algorithm.

---

**Plurality Selection**

At each communication instance between two nodes:

1) If the observer node is in a strong state $j$ and the contacted node is in a different strong state, then the observer node switches to the weak state $j$.
2) If the observer node is in a weak state $j$, it switches to the state of the contacted node.

---

[1] For example, one may choose $p(j) = M(l + 1 - j)$, $j = 1, 2, \ldots, l$, where $M$ is the normalization constant $M = 2/[l(l + 1)]$.

Notice that this is a rather natural algorithm where the weak state essentially serves to remember the last strong state in which the node was. The state of the node can be encoded by $\log_2(m)+1$ bits of memory per node and each communication of a pair of nodes requires the same bits of transmission.

**State Evolution**. Under the assumed asynchronous communication model, the system state evolves according to a continuous-time Markov process which we introduce in the following. Let $S_i(t)$ and $W_i(t)$ be the number of nodes that are in strong $i$ and weak $i$ state respectively at time $t$. The state evolution $(\vec{S}(t), \vec{W}(t),\ t \geq 0)$ is a continuous-time Markov process specified by the following transition rates:

$$(\vec{S}, \vec{W}) \to \begin{cases} (\vec{S}, \vec{W}) + (-e_i, e_i) & : \quad S_i \frac{\sum_{l \neq i} S_l}{n-1} \\ (\vec{S}, \vec{W}) + (0, -e_i + e_j) & : \quad W_i \frac{W_j}{n-1} \\ (\vec{S}, \vec{W}) + (e_j, -e_i) & : \quad W_i \frac{S_j}{n-1} \end{cases} \tag{1}$$

where $e_i$ is a vector of dimension $m$ whose all elements are equal to 0 but the $i$-th element is equal to 1.

We note that the state evolution can be equivalently represented by a Markov process $(\vec{S}(t), \vec{U}(t),\ t \geq 0)$ where $U_i(t)$ denotes the number of nodes that are in either strong or weak state $i$ at time $t$, i.e. $U_i(t) = S_i(t) + W_i(t)$, for $i \in I$.

**The Limit O.D.E.**. The Markov process $(\vec{S}(t), \vec{W}(t),\ t \geq 0)$ with transition rates (1) is a density-dependent Markov process whose scaled version $(\vec{s}^n(t), \vec{w}^n(t)) = (\frac{1}{n}\vec{S}(t), \frac{1}{n}\vec{W}(t))$ such that $\lim_{n\to\infty}(\vec{s}^n(0), \vec{w}^n(0)) = (\vec{s}(0), \vec{w}(0))$, for some fixed $(\vec{s}(0), \vec{w}(0))$, uniformly converges over any compact time interval to the solution of the following system of ordinary differential equations:

$$\frac{d}{dt}s_i(t) \quad = \quad (1 - 2s(t) + s_i(t))s_i(t) \tag{2}$$

$$\frac{d}{dt}w_i(t) \quad = \quad s_i(t)(s(t) - s_i(t)) - s(t)w_i(t), \tag{3}$$

where $i = 1, 2, \ldots, m$, $s(t) = \sum_j s_j(t)$ and $t \geq 0$.

We will also use an equivalent alternative representation of (3) by defining $u_i(t) = s_i(t) + w_i(t)$, for $i \in C$ and $t \geq 0$:

$$\frac{d}{dt}u_i(t) \quad = \quad s_i(t) - s(t)u_i(t). \tag{4}$$

### *A. Convergence of the Limit Dynamics*

**The Limit Point**. We first present our main result that establishes that under the limit dynamics, the fraction of nodes in any non-plurality state diminishes to zero with time.

***Theorem** 2:* Suppose $s_1(0) = \cdots s_k(0) > s_{k+1}(0) \geq \cdots \geq s_m(0)$ and $\sum_j s_j(0) = 1$. Then, for every $t \geq 0$,

$$u_1(t) = \cdots = u_k(t) > u_{k+1}(t) \geq \cdots \geq u_m(t), \tag{5}$$

$$s_1(t) = \cdots = s_k(t) > s_{k+1}(t) \geq \cdots \geq s_m(t). \tag{6}$$

Moreover,

$$\lim_{t\to\infty} u_i(t) = \begin{cases} \frac{1}{k}, & i = 1, 2, \ldots, k \\ 0, & \text{otherwise} \end{cases}, \tag{7}$$

and

$$\lim_{t\to\infty} s_i(t) = \begin{cases} \frac{1}{2k-1}, & i = 1, 2, \ldots, k \\ 0, & \text{otherwise} \end{cases}. \tag{8}$$

The theorem says that for every initial value, the system is order preserving with respect to the fractions of nodes in the states. If there is a majority preferred item, i.e. $k = 1$, all nodes converge to the strong state corresponding to the majority item.

**Convergence Time**. We consider the convergence time of the limit dynamics so that only some given fraction $0 < \delta < 1$ of nodes is in a state that corresponds to a non-plurality item.

*Definition 1:* For given $0 < \delta < 1$, $t_\delta \geq 0$ is said to be $\delta$-convergence time if $\sum_{i=1}^k u_i(t_\delta) = 1 - \delta$.

We first characterize the rate of convergence at which the nodes in any non-plurality state depletes, asymptotically for large $t$.

***Lemma** 1:* For every non-plurality item $i = k+1, \ldots, m$, we have

$$\lim_{t\to\infty} \frac{1}{t}\log(s_i(t)) = \lim_{t\to\infty} \frac{1}{t}\log(w_i(t)) = -\frac{1}{2k-1}.$$

The lemma tells us that the fraction of nodes that are in a non-plurality state diminishes to zero exponentially with a rate $1/(2k-1)$. Hence, this shows that the larger the number of plurality items, the slower the rate of convergence. In particular, this rate of convergence is inversely proportional to the number $k$ of plurality items, asymptotically for large $k$.

We next show tight bounds on the $\delta$-convergence time that holds for any fixed $m > 1$ and any given number $1 \leq k < m$ of plurality items, and any initial state such that each node is in a strong state and the gap between the initial fraction of nodes that prefer a plurality item and the fraction of nodes that prefer a non-plurality item is at least $0 < \gamma < 1/k$, i.e. $s_1(0) - s_{k+1}(0) > \gamma$. We first present an upper bound.

***Theorem** 3:* For every fixed $m > 1$ and initial state such that $s_{k+1}(0) + \gamma \leq s_1(0)$ for $\gamma > 0$ and $\sum_i s_i(0) = 1$, there exists a constant $C_m > 0$ such that the $\delta$-convergence time $t_\delta$ is such that

$$t_\delta \leq (2m-1)\left[\log\left(\frac{1}{\delta}\right) + \log\left(\frac{1}{\gamma}\right)\right] + C_m.$$

The proof is based on analysis of the system of ordinary differential equations (2)–(3) and, in particular, using comparison with some auxiliary differential systems. The bound of the theorem holds for $C_m = (2m)\log(2m)$. For the special case $m = 2$ and $\delta = 1/n$, note that $t_\delta \leq 3\log(n) + O(1)$ which is exactly of the same order as the bound for binary consensus in [19], and our result is a generalization of that in [19] in revealing how the convergence time depends on the gap $\gamma$. We show that the bound is tight up to a constant factor by exhibiting the following lower bound.

*Theorem 4:* For every even $m > 1$, there exists an initial state with the gap at least $\gamma > 0$ and constant $C_m > 0$ such that the $\delta$-convergence time satisfies, for every sufficiently small $\delta, \gamma > 0$,

$$t_\delta \geq (m-1)\log\left(\frac{1}{\delta}\right) + (2m-1)\log\left(\frac{1}{\gamma}\right) + C_m.$$

The proof is based on considering the following symmetric case: $m$ is even with $k = m/2$ and for $0 < \gamma < 1/k$, the initial state is given as follows

$$s_i(0) = \begin{cases} \frac{1}{m} + \frac{\gamma}{2}, & i = 1, 2, \ldots, \frac{m}{2} \\ \frac{1}{m} - \frac{\gamma}{2}, & i = \frac{m}{2}+1, \ldots, m. \end{cases}$$

### *B. Convergence to the Limit Dynamics*

We consider the convergence of our original stochastic system (1) to the solution of the limit differential system (2)-(3), or (4), as the number of nodes $n$ grows large. Since the stochastic system is a density-dependent Markov process, by Kurtz's convergence theorem [20], we know that the scaled stochastic system $(\frac{1}{n}\vec{S}(t), \frac{1}{n}\vec{W}(t),\ t \geq 0)$ uniformly converges to the solution of the limit differential system (2)-(3) on every compact interval $[0, T]$, where $T > 0$ is fixed. We will extend this result to time interval $[0, T_n]$ where $T_n$ is allowed to grow with $n$ logarithmically fast, i.e. $T_n = O(\log(n))$. This will enable us to relate the convergence time derived for the limit differential system to the convergence of the original stochastic system.

In the following lemma we characterize the convergence of the stochastic system to its limit differential system.

*Lemma 2:* For the Markov process (1), it holds that for every $\epsilon > 0$, $0 < C < 1/4$ and $T_n = C\log(n)$, there exists a constant $C_1 > 0$ such that for sufficiently large $n$,

$$\begin{aligned} & \mathbb{P}\left(\sup_{0\leq t\leq T_n} |(\vec{s}^n(t), \vec{w}^n(t)) - (\vec{s}(t), \vec{w}(t))| \geq \epsilon\right) \\ \leq \quad & C_1 \exp\left(-\frac{\epsilon^2}{(2m)^3 C}\frac{n^{1-4C}}{\log(n)}\right). \end{aligned}$$

The lemma is established by tedious but minor adaptations of the proof of Kurtz's convergence theorem that is available in [20]. Denoting with $u_B^n(t)$ the fraction of nodes that are in a non-plurality state at time $t$, and using lemma 2, we derive the following result that characterizes the fraction of nodes in a non-plurality state of the stochastic system at the $\delta$-convergence time.

*Theorem 5:* Suppose that $t_\delta$ is the $\delta$-convergence time with $\delta = 1/n^\alpha$ for some $0 < \alpha < 1/(8m+2)$. Then,

$$u_B^n(t_\delta) = O(n^{-\alpha}) \text{ with high probability.}$$

## V. Conclusion

There are several interesting directions for future work. Our results may be regarded as a first step towards understanding the fundamental information-theoretic bounds for the underlying ranking problem in network communication systems. Our work focused on establishing upper bounds on the convergence time under given memory and communication constraints. An interesting problem would be to investigate what best lower bounds on the convergence time could be achieved under given memory and communication constraints. It would also be of interest to investigate the information ranking problem for node interaction rates that induce an arbitrary connected graph.

**Acknowledgement.** Kyomin Jung was supported by the Basic Science Research Program through the National Research Foundation of Korea(NRF) funded by the Ministry of Education, Science and Technology(2012032786).

## References

[1] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar. Bayesian learning in social networks. *Review of Economic Studies*, 2010.
[2] F. Benezit, P. Thiran, and M. Vetterli. Interval consensus: from quantized gossip to voting. In *Proc. of ICASSP*, 2009.
[3] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Randomized gossip algorithms. *IEEE Trans. on Information Theory*, 52:2508–2530.
[4] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis. On distributed averaging algorithms and quanitization effects. *IEEE Trans. on Automatic Control*, 54:2506–2517.
[5] F. Chierichetti and J. Kleinberg. Voting with limited information and many alternatives. *Proc. of SODA*, 2012.
[6] M Draief and M.Vojnović. Convergence speed of binary interval consensus. *SIAM Journal on Control and Optimization*, 2012.
[7] M. E. Hellman and T. M. Cover. Learning with finite memory. *The Annals of Mathematical Statistics*, 41(3):765–782, 1970.
[8] K. Jung, B. Y. Kim, and M. Vojnović. Distributed ranking in networks with limited memory and communication. Technical Report MSR-TR-2011-88, Microsoft Research, 2011.
[9] K. Jung, D. Shah, and J. Shin. Distributed averaging via lifted Markov chains. *IEEE Trans. on Information Theory*, 56:634–647.
[10] A. Kashyap, T. Basar, and R. Srikant. Quantized consensus. *Automatica*, 43:1192–1203, 2007.
[11] J. Kleinberg. The small-world phenomenon and distributed search. *SIAM News*, 37, 2004.
[12] J. Kleinberg. Complex networks and decentralized search algorithms. In *Proc. of the International Congress of Mathematicians*, 2006.
[13] J. Koplowitz. Necessary and sufficient memory size for $m$-hypothesis testing. *IEEE Trans. on Information Theory*, IT-21(1):44–46, 1975.
[14] F. Kuhn, T. Locher, and R. Wattenhofer. Tight bounds for distributed selection. In *Proc. of ACM SPAA*, 2007.
[15] T. M. Liggett. *Interacting Particle Systems*. Springer, 2 edition, 2006.
[16] E. Mossel and G. Schoenebeck. Reaching consensus on social networks. *International Journal of Intelligent Systems - Decision Making in Social Networks*, 25, 2010.
[17] B. Patt-Shamir and M. Teplitsky. The round complexity of distributed sorting. In *Proc. of ACM PODC*, 2011.
[18] D. Peleg. *Distributed computing: a locality-sensitive approach*. SIAM, 2000.
[19] E. Perron, D. Vasudevan, and M. Vojnović. Using three states for binary consensus on complete graphs. In *Proc. of IEEE Infocom*, 2009.
[20] A. Shwartz and A. Weiss. *Large Deviations for Performance Analysis - Queues, Communications, and Computing*. Chapman & Hall, 1995.
[21] B. Touri and A. Nedic. On ergodicity, infinite flow, and consensus in random models. *IEEE Transactions on Automatic Control*, 56:1593–1605, 2011.