This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2010 proceedings.

Capability and Responsibility Balancing in Online Social Search

Kuang Xu, Victor O.K. Li, Jing Xie, Guang-Hua Yang
Department of Electrical and Electronic Engineering
The University of Hong Kong, Pokfulam, Hong Kong, China

***Abstract*—Online social search (OSS) brings forth a new way to harness the Internet for answers. In this paper, we study the balancing between OSS users' capabilities and responsibilities. Targeting a practical system design, we propose an analytical model that captures the heterogeneity of different referral sessions in OSS, and a distributed socio-aware referral strategy that can achieve the desired balance when the system reaches steady state. We show that configuring the strategy enables the system operator to control the flow of all posed questions in the system. We also discuss the implications of configuring the strategy from a gaming-strategy point of view.**

I. Introduction

Online social search (OSS) brings forth a new way to harness the Internet for answers, utilizing the underlying network structure of online social networks (OSN) to find information [1] [2] via friends, and friends' friends. In a typical OSS system, each user registers his expertise in his profile. A person looking for experts sends the question to his selected contacts. When a user receives a question on which he is an expert, he responds to the questioner. Independently, he can also forward it to his selected contacts. In this way, the question is passed on in the social network. Finally, the questioner may be presented with a great number of potential respondents.

In this paper, we study the balancing between users' capabilities and undertaken responsibilities. Capability refers to expertise, that is, the amount of expertise a user supplies to the OSS system, and responsibility measures the number of questions a node receives. Conceivably, a person may be unwilling to bother himself to answer others' questions on which he is expert though, sometimes, he also uses OSS to find experts on his own questions, thus straining the participating willingness of other OSN users. As opposed to this, our design relies on the principle that more knowledgable (or active) OSN users can provide more help to the online society, that is, the higher the capability, the higher the responsibility.

Targeting a practical system design, we first propose an analytical model that captures the heterogeneity of different referral sessions, relating the question forwarding probability at each node in a referral chain to the expertise of the node who poses this question. This reflects the social connections between the questioner and his reference nodes. Here, we consider a referral session as the process from when a question is injected to the system until it is discarded by an intermediate user in the referral chain or until it exceeds a customized hop-limit. Since in practice an OSN user maintains its local social network, we equip the nodes in our model with the intelligence of awareness, assuming every node is aware of the expertise of its neighbors. We propose a socio-aware referral strategy (SARS), which draws on the classic Metropolis-Hastings algorithm [3]. We prove that our proposed referral strategy can achieve a desired balance between nodes' capabilities and responsibilities when the system reaches the steady state. We analyze in detail the strategy's configuration functions which are accessible to the system operator, enabling the system to control the flow of all posed questions in the system. In particular, the system can assign responsibilities to nodes in any positive relationship with nodes' expertise during a time period of interest, such as proportional, or square proportional. We also evaluate our proposed referral strategy based on the crawled data of a set of real OSNs with various settings, and find that SARS outperforms the traditional random walk strategy.

A salient contribution in this study is that, based on the concept of balance, we substantiate the net benefit of an OSS user by a payoff function, and analyze it from a gaming-strategy point of view. With two concrete examples, we contend that, with the desired balance, inappropriate configuration of the referral strategy may lead to ramifications to the OSS applications.

In contents to follow, we first present our analytical model in Section II, followed by the proposed referral strategy and the analyses in Section III, in which we also give a brief review of the Metropolis-Hastings algorithm we draw on. Then we evaluate the system by simulations in Section IV, and finally, we conclude our study with suggestions for future work in Section V.

II. Model

We consider an OSN as an undirected graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of nodes (OSN users) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges (social ties) in the network. Each edge means one-hop question-forwarding is possible between the corresponding pair of nodes. Let $n = |\mathcal{V}|$ be the number of users in the system. We also denote by $\mathcal{N}_u \subseteq \mathcal{V}$ the set of neighbors of Node u, and $d_u = |\mathcal{N}_u|$ the number of users in this set.

In the system, each user supplies the information of his

* This research is supported in part by the University of Hong Kong Strategic Research Theme of Information Technology.

expertise[1] in his profile. We denote by w_u the expertise of Node u and, in practice, w_u can be published through being embedded in the questions posed by Node u. When a user poses a question to the underlying social network, he hopes that answers could appear from his friends, his friends' friends, and so forth. However, in reality, a user may not be willing to help others forward a question even if he can not respond to the questioner (e.g. not an expert or not interested in answering) since it will cost his effort. For generality, in our model we assume that a node receiving a question will forward it to the next reference with probability q, i.e., discarding the question with probability $1-q$, where q is independent of its expertise. We study the practical case in which q is heterogeneous in different referral sessions, and relate q to Node u who initiates the current referral session, denoting it by q_u.

Definition 1. *The **forwarding probability** q_u is an increasing function of the expertise of Node u, i.e., $\frac{\partial q_u}{\partial w_u} > 0$.*

The above definition captures the social connections between the questioner and his reference nodes in a referral session. The scenario for this is that people will likely be more willing to help a person who could potentially be more helpful to them (i.e. more knowledgeable). In addition, we allow a questioner to receive multiple answers in a referral session. Obviously, the number of returned answers is closely related to the number of references a question is forwarded to, which is determined by the forwarding probability q. A higher q indicates a higher visibility of the question, which is more beneficial to the questioner. Thus the above definition says a node with lower w is less able to approach multiple experts on its question. This discourages nodes from publishing underestimated expertise in the system.

Definition 2. *The **responsibility** Node u takes, denoted by r_u, is defined as the number of questions Node u receives during the time period of interest.*

The responsibility directly represents the workload a node undertakes. We hope that a more knowledgable node can make more contributions to the online society by answering more questions from others, i.e., in our context, receiving more questions.

We consider a practical system design that imposes a limit T on the number of hops a question can be forwarded. At each hop of a referral session, the node forwards the question with probability q to a randomly selected neighboring reference. We can obtain S, the expected number of references a question is directed to,

$$\begin{aligned} S &= \sum_{j=1}^{T-1} j \cdot q^{j-1}(1-q) + T \cdot q^{T-1} \\ &= \frac{1-q^{T-1}}{1-q} + q^{T-1}. \end{aligned}$$

We assume n is large, and denote by e_i the *expert density* [4], that is, the probability that a node is an expert on Question i. Thus the expected number of answers obtained by the questioner posing Question i is

$$\begin{aligned} A_i &= e_i \cdot S \\ &= e_i \cdot [\frac{1-q^{T-1}}{1-q} + q^{T-1}] \end{aligned} \tag{1}$$

Thus, the benefit of registering more expertise can be observed quantitatively, by verifying

$$\begin{aligned} \frac{\partial A_i}{\partial w_u} &= \frac{\partial A_i}{\partial q_u} \cdot \frac{\partial q_u}{\partial w_u} \\ &> 0. \end{aligned}$$

From the viewpoint of a system operator, it is rational to have a node with higher expertise undertake higher responsibility, such that all nodes contribute to the online society in a balanced manner. In the following, we propose a distributed referral strategy that can match the responsibility to the expertise.

III. Matching Responsibility to Capacity

Based on the above model, we study how to achieve the balance between nodes' expertise and responsibilities. Since in reality a node maintains its local social network, it is reasonable to assume every node is aware of the expertise of its neighbors. In this section, we propose a socio-aware referral strategy (SARS) based on this social context. SARS is distributed, and draws on the Metropolis-Hastings algorithm [3] [5], which is a standard approach to guide the configuration of biased random walk such that it uniquely converges to an arbitrary distribution. For ease of presentation, in the following we first give a brief technical review of the Metropolis-Hastings algorithm, followed by our proposed referral strategy. Then we analytically study how to configure the strategy in practice. We also discuss the implications from a gaming-strategy point of view.

A. Metropolis-Hastings algorithm

A random walk on graph $\mathcal{G}$ starts at a node v_0, and π_0 is an arbitrary initial distribution. If the random walk is at node u_t at time Step t, it moves to its randomly selected neighboring node u_{t+1} at Step $t+1$ with a certain probability distribution. Let π_t be the distribution of node u_t such that $\pi_t(v) = p(u_t = v)$, for each $v \in \mathcal{V}$. The transition matrix of the random walk is denoted by P, of which an entry P_{uv}, $u, v \in \mathcal{V}$, is the probability that the random walk moves from Node u to its neighbor v in one step. Thus, $\pi_{t+1} = \pi_t \cdot P$.

Theorem 1 [5]. *Let π be a desired probability distribution. For each neighbor v of Node u, let*

$$P_{uv} = \begin{cases} \frac{1}{d_u} & \text{if } \frac{\pi(u)}{d_u} \le \frac{\pi(v)}{d_v}, \\ \frac{1}{d_v} \cdot \frac{\pi(v)}{\pi(u)} & \text{if } \frac{\pi(u)}{d_u} > \frac{\pi(v)}{d_v}, \end{cases}$$

and $P_{uu} = 1 - \sum_{v \in \mathcal{N}_u} P_{uv}$. Then π is a converged probability distribution of a random walk with the transition matrix P.

Assuming that π is the desired distribution, and π_t is the probability distribution of node visitation by the random walk

[1]In this paper, a user's expertise is measured in terms of the number of topics on which he is expert.

at Step t, the extent to which the convergence is achieved at Step t is measured by,

$$||\pi - \pi_t|| = \frac{1}{2} \sum_{v \in \mathcal{V}} \pi_t(v) - \pi(v),$$

where the factor $\frac{1}{2}$ is used to ensure that $||\pi - \pi_t||$ never exceeds 1. Obviously, $||\pi - \pi_t|| = 0$ represents complete convergence. The convergence time, defined as

$$\tau(\epsilon) = min\{t : \forall t' \geq t, ||\pi - \pi'_t|| \leq \epsilon\},$$

measures the time π_t takes to converge to π. Fast convergence means that $||\pi - \pi_t||$ decreases quickly as t grows. The convergence time of a random walk is bounded as follows:

Theorem 2 [6]. *Let $\pi_{min} = min_{\pi_i > 0} \pi_i$. Then, $\tau(\epsilon) \leq \frac{1}{1-\lambda_2} log((\pi_{min}\epsilon)^{-1})$, λ_2 is the second largest eigenvalue of the transition matrix P.*

The Metropolis-Hastings algorithm provides a convergence-guaranteed mechanism to configure the biased random walk. In the next section, we will show that our proposed referral strategy leverages the core of this algorithm, and additionally introduces the forwarding probability q (see Definition 1).

B. Socio-aware referral strategy

Intuitively, more knowledgable nodes could contribute more to OSS. We seek to develop a referral strategy, such that when the systems reaches the steady state, we can achieve the principle – the higher the capability, the higher the responsibility. Different from the study in [7] which relies on a centralized expertise finder, we study the distributed referral strategy which only requires the information of a node's local social network.

If a node u that receives Question i has expertise in i, it responds to the node that poses i with an answer. Since the system allows a questioner to receive multiple answers, independently, u decides whether to forward i to a next hop node according to the forwarding probability q and, if so, forward i to a randomly selected neighbor according to Theorem 1. If not, the current referral session ends. Each referral session is also initiated with a hop-limit, and considered finished if this hop limit is exceeded.

Lemma 1 *When the system reaches the steady state, the probability of Node $v \in \mathcal{V}$ being visited by all posed questions in the system is $\pi(v)$.*

By "reaches steady state", we mean the random walk converges. Lemma 1 directly follows Theorem 1, which ensures that each node is visited with the desired probability after the convergence. We also note that since the node receiving a question forwards it to a next hop node with a probability q, the transition matrix $P^* = q \cdot P$. Fig. 1 illustrates the transition probabilities between two neighboring nodes in our system. We assume $\frac{\pi(u)}{d_u} \leq \frac{\pi(v)}{d_v}$, thus $P^*_{uv} = \frac{q}{d_u}$, and $P^*_{vu} = \frac{q}{d_u} \cdot \frac{\pi(u)}{\pi(v)}$. We can verify the flow balance condition that $\pi(u)P^*_{uv} = \pi(v)P^*_{vu}$ in this situation. Considering our purpose is to study the balance between nodes' expertise capacities and responsibilities, following Lemma 1 we draw the following conclusion.

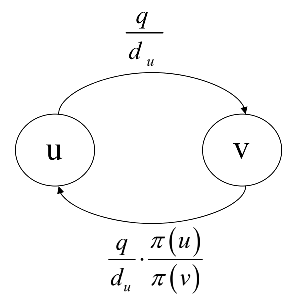

Fig. 1. Example of transition between u and v ($\frac{\pi(u)}{d_u} \leq \frac{\pi(v)}{d_v}$).

Proposition 1 *When the system reaches the steady state, the responsibility that Node $v \in \mathcal{V}$ undertakes in the system is proportional to $\pi(v)$.*

It is easy to see that the above proposition holds independent of the imbalance among nodes in terms of the number of questions they pose.

In practice, any distributed referral strategy would incur communication overhead. Due to space limitations, we do not present the analysis of the communication cost of SARS, and leave the detailed study as future work.

C. Analysis

To achieve the balance between nodes' capabilities and responsibilities, we can compare the workload taken by a node and the benefit it gains by registering its expertise. Let the cost of receiving/handling a question be α, and the benefit of obtaining an answer be β. We further denote by C the number of forwardings caused by all posed questions in the system during the time period of interest, and c_u, the number of questions posed by Node u[2]. Thus the net benefit Node u obtains is

$$O_u = \beta \cdot \sum_{i=1}^{c_u} A_{u_i} - \alpha C \cdot \pi(u), \tag{2}$$

where A is given in Eqn. (1). By adjusting the probability distribution $\pi(u)$, the system can achieve the desired balance. In the following analysis, we set

$$\pi(u) = \frac{\mathcal{F}(w_u)}{\sum_{v \in \mathcal{V}} \mathcal{F}(w_v)}, \tag{3}$$

where $\mathcal{F}$ is an arbitrary increasing function, such as

$$\mathcal{F}(x) = x \quad \text{or} \quad \mathcal{F}(x) = x^2.$$

Obviously, the user registered expertise w, is the only information required for the system to achieve the desired performance. Thus for a node in the system, its benefit of registering expertise, which could be embodied by obtaining answers when it poses questions, is balanced by the cost of undertaking the responsibility of answering questions.

1) Effect of $\mathcal{F}$: The system can configure the function $\mathcal{F}$ to adjust the balance of each user.

Consider two users u and v with expertise w_u and w_v, respectively. We assume u is more knowledgable than v, and

[2]We consider the c_u questions posed by Node u are different questions.

let $w_u = \rho \cdot w_v$, where $\rho > 1$. We study the implication of configuring $\mathcal{F}(x)$. Let $\mathcal{F}(x) = x^m$, $m > 0$. We can compare the responsibilities the two users undertake. By Definition 2 and Eqn. (3), we obtain

$$\frac{r_u}{r_v} = \frac{C\pi(u)}{C\pi(v)} = \rho^m.$$

Note that ρ times difference in the capacities of two users leads to a ρ^m times difference in the responsibilities undertaken by them. This implies that a function $\mathcal{F}(x)$ with a fast increasing speed (e.g. a function with a large m here) can have a considerable proportion of questions directed to users with higher expertise. This enables the system to control the flow of posed questions by configuring the function $\mathcal{F}(x)$.

2) Game among users: By setting π and $\mathcal{F}$, the system can achieve the desired balance. However, the goal of a practical OSS system is to benefit its users as well as the system as a whole. As a user, the benefit is to obtain satisfactory service by participating; as the system operator, the benefit is to maintain and draw as many active users as possible so as to maximize the revenue. Ensuring all users receive satisfactory service in a balanced manner is an important consideration. However, in order to achieve balance, inappropriate configuration of the system may bring losses to the service provider. In the following, we analyze the system from a gaming-strategy point of view, and discuss the implication to applications.

In the above analysis, we assume each user honestly registers his expertise. However, in reality, an all-round expert is likely to be unwilling to register all he knows [8], and publishes an underestimated expertise in the system. We can model this scenario as a strategic game and assume that a user's objective is to maximize its net benefit in Eqn. (2), which is considered as a node's payoff function. The action of a user is to register its expertise by w. Denote by w^* the real capacity of a user's expertise, we have $0 < w \leq w^*$.

For clarity, we substantiate the forwarding probability q with a simple form $q_u = \rho \cdot w_u$, $\rho > 0$ (by Definition 1), and configure $\mathcal{F}(x) = x$. Combining Eqns. (1), (2), and (3), we can obtain the Nash equilibrium/equilibria by solving the following system of equations

$$\frac{\partial O_u}{\partial w_u} = \varphi_u \rho \frac{1 + (T-2)(\rho w_u)^{T-1} - (T-1)(\rho w_u)^{T-2}}{(1-\rho w_u)^2} - \alpha C \frac{\sum_{v \in \mathcal{V} \setminus u} w_v}{\sum_{v \in \mathcal{V}} w_v} = 0$$

for all $u \in \mathcal{V}$, where $\varphi_u = \beta \cdot \sum_{i=1}^{c_u} e_{u_i}$.

We consider a simple scenario of only two users (i.e., $|\mathcal{V}| = 2$). Thus the payoff functions are calculated by

$$O_1(w_1, w_2) = \frac{\varphi_1}{1 - \rho w_1} - \frac{\alpha C w_1}{w_1 + w_2}, \quad (4)$$

$$O_2(w_1, w_2) = \frac{\varphi_2}{1 - \rho w_2} - \frac{\alpha C w_2}{w_1 + w_2}. \quad (5)$$

Without loss of generality, we further simplify the scenario by classifying a node's behavior into only two actions, registering an underestimated expertise capacity $w' < w^*$ and registering its real capacity w^*. Then one can analyze the system based on the above game. Since our purpose here is to contend that appropriate settings of π and $\mathcal{F}$ are important for the system in addition to achieving a desired balance, we do not investigate the game in detail. Instead, two examples are used in the following to show the implications to the OSS applications.

u1 \ u2 (w1 \ w2)	8	9
8	(0, 0)	(0.29, 4.71)
9	(4.71, 0.29)	(5, 5)

(a) Example 1

u1 \ u2 (w1 \ w2)	8	9
2	(-0.75, -3)	(-0.57, 1.82)
3	(-1.30, -2.73)	(-1.07, 7.5)

(b) Example 2

Fig. 2. Gaming examples

In the first example, we assume the action set for both players is $\{8, 9\}$. Fig. 2(a) depicts their payoffs according to (4) and (5), where $\varphi_1 = \varphi_2 = 1$, $\rho = 0.1$, and $\alpha C = 10$. We can see the action profile $(9, 9)$ constitutes the unique Nash equilibrium, which is also optimal for both players. This is desired for the system since each user will actively participating by registering its real expertise while receiving satisfactory service. Applying the same setting, Fig. 2(b) depicts another example, in which the action set for User 1 is $\{2, 3\}$ and that for User 2 is $\{8, 9\}$. The action profile $(2, 9)$ constitutes the Nash equilibrium, in which User 1 unfairly maximizes the payoff by not contributing all of its expertise while User 2 receives non-optimal payoff by contributing all its effort. Thus User 1 may get a "free ride" and User 2 is frustrated and leaves the system. This brings loss to the system. Note that the aforementioned balance holds in both examples. Therefore, the implication is that with the desired balance, inappropriate configuration of the referral strategy may lead to ramifications to the OSS applications.

IV. Evaluation

In this section, we study empirically the performance of our proposed referral strategy. In particular, we investigate the following performance metric.

Balance Index during a time period of interest in an OSS system is defined as

$$BI = \sum_{u \in \mathcal{V}} \left| \pi_u - \frac{r_u}{C} \right|,$$

where C records the total number of forwardings caused by all posed questions in the system during the time period of interest, π_u is set by Eqn. (3), and r_u is give by Definition 2. We study the convergence rate of BI when C increases under various settings.

We utilize the connectivity data of a set of OSNs, namely Orkut, LiveJournal, collected by Mislove et al. [9]. Orkut

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE Globecom 2010 proceedings.

OSN	Orkut	LiveJournal
Number of nodes	3,072,441	5,284,457
Estimated crawled fraction	11.3%	95.4%
Number of links	223,534,301	77,402,652
Av. no. of friends per node	106.1	16.97
Fraction of symmetric links	100.0%	73.5%

TABLE I
STATISTICS OF THE OSN DATASETS [9]

is a website of explicitly defined social network to help people meet new friends and maintain existing relationships. LiveJournal is an online social network of bloggers. The major statistics of these datasets are summarized in Table I. We believe it is more realistic to evaluate the system on these real social network data[3]. Since the networks are too big for evaluation, we sample several different portions from each network with Snowball sampling [10]. We set the size of the sampled networks to 1×10^4. We also set $T = 25$, and the forwarding probability in a referral session $q = 0.9 + 0.1w$, where the expertise w of each node is randomly selected in the range of $[1, 10]$. During the simulation, the nodes posing questions are randomly selected from the population.

We first study the effectiveness of our proposed SARS. Here, we set $\mathcal{F}(x) = x$, and compare the performances achieved by SARS and the simple random walk (SRW) referral strategy. SRW is popularly studied and utilized in the literature [11] [12]. In SRW, if a node decides to forward a question, it simply forwards the question to a uniformly selected neighboring node. Fig. 3(a) shows the simulation result. We can clearly observe that for both social network datasets, SARS outperforms SRW (i.e. the BI achieved by SARS converges much faster than that achieved by SRW). We note that when C is smaller than 1×10^4, the BI of SARS is slightly larger than that of SRW. This is because the system does not reach the steady state (note that the size of the network is set to 1×10^4). We also observe that no clear distinction exists in the resulting performances from Orkut and LiveJournal while the two datasets have different connectivities. By investigating the sampled networks, we find that both networks have a large average degree. This leads us to an empirical conclusion: the system performance in terms of the balance index is independent of its underlying network connectivity when the system has a large number of users.

We further study the convergence of the balance index with different configurations of the function $\mathcal{F}(x) = x^m, m > 0$. From Fig. 3(b) we can see that the balance index converges under different values of m when the system reaches steady state. In addition, it converges slower under a larger value of m, which matches our analysis (in Section III-C1) that a fast increasing $\mathcal{F}(x)$ enlarges the distinction in responsibilities taken by nodes with different expertise.

[3]Utilizing datasets from actual OSS applications such as Aardvark (which is based on Facebook) would be ideal, but those datasets are unavailable since Facebook prohibits automated crawlers.

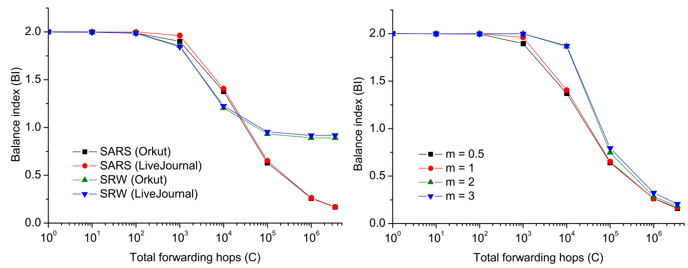

(a) Different strategy (b) Different $\mathcal{F}$ (LiveJournal)

Fig. 3. Performance evaluation result.

V. CONCLUSION

In this paper, we study the balancing between the users' capabilities and responsibilities of OSS. We propose an analytical model that captures the heterogeneity of different referral sessions. With the design principle – the higher the capability, the higher the responsibility, we proposed a distributed socio-aware referral strategy that is proved able to achieve the desired balance when the system reaches steady state. We analyze in detail the strategy's configuration functions that enable the system operator to control the flow of all posed questions in the system. The evaluation based on crawled data of several OSNs validates our analysis. We also discuss the implications of configuring the referral strategy from a gaming-strategy point of view, and show that, with the desired balance, inappropriate configuration of the referral strategy may lead to ramifications to the OSS applications. In the future, we would like to further investigate this interesting issue in detail.

REFERENCES

[1] D. J. Watts, P. S. Dodds, and M. E. J. Newman, "Identity and search in social networks," *Science*, vol. 296, pp. 1302–1305, May 2002.
[2] J. Zhang and M. V. Alstyne, "SWIM: fostering social network based information search," in *Proceedings of ACM CHI '04*, 2004.
[3] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equation of state calculations by fast computing machines," *J. Chemical Physics*, vol. 21, pp. 1087–1092, 1953.
[4] K. Xu, J. Xie, and V. O.K. Li, "Locating experts via online social networks," in *Proceedings of IEEE ICC '10*, 2010.
[5] Y. Azar, A. Broder, A. Karlin, N. Linial, and S. Phillips, "Biased random walks," in *Proceedings of the 24th Annual ACM Symposium on Theory of Computing*, 1992.
[6] P. Diaconis and D. Stroock, "Geometric bounds for eigenvalues of markov chains," *Annuals of Applied Probability*, vol. 1, pp. 36–61, 1991.
[7] J. Zhang, M. S. Ackerman, and L. Adamic, "Expertise networks in online communities: structure and algorithms," in *Proceedings of 16th International Conference on World Wide Web*, 2007.
[8] H. Kautz, B. Selman, and M. Shah, "Referral Web: combining social networks and collaborative filtering," *Communications of the ACM*, vol. 40, no. 3, pp. 63–65, 1997.
[9] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, "Measurement and analysis of online social networks," in *Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement*, 2007.
[10] S. K. Thompson, *Sampling*, John Wiley & Sons Inc., 2 edition, April 2002.
[11] C. Gkantsidis, M. Mihail, and A. Saberi, "Random walks in peer-to-peer networks," in *Proceedings of IEEE INFOCOM '04*, 2004.
[12] L. Lovasz, "Random walks on graphs: a survey," *Combinatorics, Paul Erdos is Eighty*, vol. 2, pp. 1–46, 1993.

978-1-4244-5638-3/10/$26.00 ©2010 IEEE