Citation for published version:

Stauffer, AO & Barbosa, VC 2007, 'Probabilistic heuristics for disseminating information in networks', *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, pp. 425-435. <https://doi.org/10.1109/TNET.2007.892877>

DOI:

[10.1109/TNET.2007.892877](https://doi.org/10.1109/TNET.2007.892877)

Publication date:

2007

Document Version

Peer reviewed version

[Link to publication](#)

© 2007 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Probabilistic Heuristics for Disseminating Information in Networks

Alexandre O. Stauffer, Valmir C. Barbosa, *Member, IEEE*

Abstract

We study the problem of disseminating a piece of information through all the nodes of a network, given that it is known originally only to a single node. In the absence of any structural knowledge on the network other than the nodes' neighborhoods, this problem is traditionally solved by flooding all the network's edges. We analyze a recently introduced probabilistic algorithm for flooding and give an alternative probabilistic heuristic that can lead to some cost-effective improvements, like better trade-offs between the message and time complexities involved. We analyze the two algorithms both mathematically and by means of simulations, always within a random-graph framework and considering relevant node-degree distributions.

Index Terms

Random networks, probabilistic flooding, heuristic flooding.

I. INTRODUCTION

A network can be viewed as an undirected graph $G = (N_G, E_G)$ with $n = |N_G|$ nodes and $m = |E_G|$ edges, in which the existence of an edge (u, v) represents the possibility of bidirectional communication between nodes u and v . We consider the problem of disseminating a piece of information, referred to as I , to all the nodes of the network, given that initially only one node, called the originator, has it.

A traditional algorithm to disseminate information in networks when nodes know their immediate neighborhoods and nothing else is what we call uninformed flooding [1], in allusion to the fact that its actions do in no way take into account the structure of the network or any of its properties. In this algorithm, the originator starts by sending I to its neighbors; when receiving I for the first time, each of the other nodes forwards it to its own neighbors. This algorithm has message complexity of $\Theta(m)$, as exactly two messages are transmitted on each edge, and worst-case time complexity of $O(n)$, the latter related to the customary message-passing causal chains of distributed computing under full asynchronism [2]. Also, when edges can be assumed to have approximately equal delays associated with them, the algorithm's average waiting time, which is the average time for a node to receive I , is given by the average distance from the originator in the network, that is, the average number of edges on the shortest paths between the originator and each of the other nodes.

Uninformed flooding can then be very costly with respect to how many copies of I need to be sent overall, and is also inherently taxing on a node's local bandwidth, since it prescribes that a copy of I is to be sent to every one of the node's neighbors. Our interest in this paper is to study the trade-off between the message complexity and the average waiting time when I is no longer forwarded deterministically to all of a node's neighbors, but rather is sent as the result of probabilistic decisions. Also, by deciding probabilistically whether to forward I on an edge or not, we automatically obtain a handle on how to impact a node's bandwidth less heavily as well. Our point of departure has been the recent introduction of probabilistic flooding [3], which prescribes for each node that it forward I to each of its neighbors with fixed probability p . Probabilistic flooding is then also "uninformed," constituting essentially a simple stochastic generalization of the aforementioned (deterministic) uninformed flooding. We introduce in

Research for this paper has been partially supported by CNPq, CAPES, and a FAPERJ BBP grant.

The authors are with the Systems Engineering and Computer Science Program of the Federal University of Rio de Janeiro, Brazil. E-mail addresses: astauffer@pobox.com (A. O. Stauffer), valmir@cos.ufrj.br (V. C. Barbosa).

Section II the alternative that we call heuristic flooding, which is also probabilistic in nature but, unlike probabilistic flooding, takes the network's structure into account insofar as it can be inferred from the nodes' immediate neighborhoods.

Naturally, flooding the network with copies of I that are propagated probabilistically does not ensure that all nodes will eventually receive a copy. There do exist applications, however, that may benefit from the aforementioned trade-off even in the absence of such delivery guarantees. Important examples here come from the area of peer-to-peer computing, in which a central problem is to locate an object that has been replicated throughout the network. Even though systems in this area are often highly organized to facilitate locating objects in response to queries (cf., e.g., [4], [5] and their references), there are many cases in which the objects' or the queries' nature makes such a degree of organization unattainable [3], [6], [7]. In these cases, a typical solution has been to resort to some sort of flooding, that is, to interpret I as a query and the flooding itself as a search. More specifically, the customary approach has been to do flooding while preventing copies of I from penetrating the network too deeply (that is, a node forwards I upon receiving it for the first time only if it is sufficiently close to the originator according to some criterion that depends on the algorithm at hand—cf. [6], [7] and the references therein for mechanisms such as the expanding ring and Gnutella algorithms). Incorporating probabilistic decisions into the traditional algorithm for flooding is then seen to have the potential to contribute significantly further than to the mere broadcasting of information.

We ultimately base our analyses of both probabilistic flooding and heuristic flooding on the formalism laid down in [8], whose details are briefly reviewed and extended in Section III. Our analyses are also based on what we call a flooding digraph, which is a random directed subgraph of the network. In Section IV, we introduce this subgraph and use it to obtain our analytical results.

All our analyses are carried out in the framework of random graphs [9]. Within this framework, we give special attention to random graphs whose node degrees are either Poisson-distributed or distributed according to a power law. The former arise in the classic context of Erdős and Rényi [10], while the latter have recently been found to be potentially representative of networks like the Internet and the WWW [11], [12], [9] (even though the whole issue is still open to debate [13]). Simulation results are given in Section V for these models of random graphs, revealing very good agreement with the analytical results.

A comparative study of the two probabilistic approaches is given in Section VI and focuses on how they behave with respect to the number of messages sent and the average waiting time. In both cases we find that heuristic flooding tends to require fewer messages to be sent, however at the expense of an expanded average waiting time. While in the case of Poisson-distributed node degrees what is gained occurs in similar proportion to what is lost, under power-law-distributed node degrees the trade-off is tipped differently: for the power laws that are thought to approximate the Internet topology, heuristic flooding requires only about 60% of the messages required by probabilistic flooding while taking roughly only 20% longer to complete.

Concluding remarks are given in Section VII.

II. HEURISTIC FLOODING

In probabilistic flooding, each node decides whether to send I to each of its neighbors based on a fixed probability parameter p . As a consequence, a node with a small degree has a smaller probability of receiving the information than a node with a large degree. Given a node u and one of its neighbors v , if a is the degree of u and b the degree of v , our new heuristic flooding is based on a function $h(a, b)$ representing the probability that node u sends I to node v . Intuitively, $h(a, b)$ is expected to have a larger value when at least one of a and b is small than when they are both large, thus attempting to compensate for the inherent drawback of probabilistic flooding that we just mentioned and also to reflect the understanding that a small a may signify the existence of insufficient alternative routes for v to ultimately receive I from u even when b is large.

Ideally, an accurate heuristic function would require some knowledge about the topology of the network beyond the nodes' immediate neighborhoods; realistically, however, and especially so in unstructured

scenarios, a node can only be assumed to have information about its own neighbors. For this reason, heuristic flooding is defined to operate under the assumption by each node u that a neighboring node v can only receive the information from $\min\{a, b\}$ of v 's own neighbors, for each one with equal probability $h(a, b)$, where a and b are the degrees of u and v , respectively. In other words, there are two sides to our assumption. One side is that, if u decides not to forward I to v , then v can still receive I from $\min\{a, b\} - 1$ of its neighbors, if any, in each case with the same probability $h(a, b)$. The other side of the assumption is that the other $b - \min\{a, b\}$ neighbors of v , if any, are unable to receive I if v does not send it to them.

If we let α stand for the desired probability that a node receives the information being broadcast, then, from the perspective of node u upon deciding whether to send the information to its neighbor v , the probability that v does not receive the information (that is, $1 - \alpha$) can be expressed as

$$1 - \alpha = [1 - \alpha + \alpha(1 - h(a, b))]^{\min\{a, b\}}. \quad (1)$$

This expression indicates that v does not receive the information if and only if, for each of the $\min\{a, b\}$ neighbors that the assumption says could send the information to it, either that neighbor has not itself received the information (with probability $1 - \alpha$) or it has but decided not to forward it to v (with probability $\alpha(1 - h(a, b))$). Hence,

$$h(a, b) = \frac{1 - (1 - \alpha)^{1/\min\{a, b\}}}{\alpha}. \quad (2)$$

This heuristic function has the property that, when both a and b are large, $h(a, b)$ is small, which corresponds to the intuition that there may exist other paths along which v can receive I , so it may not be essential that u sends it to v . On the other hand, when at least one of a and b is small, then $h(a, b)$ is large. In particular, when at least one of a and b is equal to 1, that is, $\min\{a, b\} = 1$, then $h(a, b)$ is also 1. We assume $h(a, b) = 0$ if $a = 0$ or $b = 0$.

III. THE UNDERLYING FORMALISM

Henceforth, G is viewed as a random graph whose nodes, for the sake of mathematical tractability, all have degrees that are independent of one another and distributed identically to a random variable K_G . The nodes of G are assumed to be interconnected at random given their degrees, so the degrees of any two adjacent nodes remain independent. The results from [8] reviewed in Section III-A, and also their extensions in Section III-B, hold in the limiting case of a formally infinite number of nodes.

A. Basic results

Denoting by $P_G(a)$ the probability that a randomly chosen node of G has degree a , with $a \geq 0$, or equivalently the probability that K_G equals a , the average degree of G , denoted by Z_G , is the expectation of K_G , that is,

$$Z_G = \sum_{a=0}^{n-1} a P_G(a). \quad (3)$$

More generally, and for $s \geq 1$, the s th moment of K_G , $\langle K_G^s \rangle$, is

$$\langle K_G^s \rangle = \sum_{a=0}^{n-1} a^s P_G(a), \quad (4)$$

where we employ the usual angle-bracket notation to indicate expectation.

We know from [14], [15], [8] that a criterion exists according to which a random graph has almost surely a size- $\Theta(n)$ connected component, referred to as the giant connected component (GCC). When this

is the case, the other components of the graph are small if compared to its number of nodes, that is, the fraction of nodes inside each of them is approximately 0. The criterion in the case of G is that

$$\frac{\langle K_G^2 \rangle}{Z_G} > 2. \quad (5)$$

If this inequality holds, G is said to be above the phase transition. If it does not, then all the components of G are small and a GCC does not exist. In this case, G is said to be below the phase transition.

Let us now consider a node at which we arrive by following a randomly chosen edge of G . The probability that such a node has other $b - 1$ edges incident to it, with $b \geq 1$, is the expected fraction of edges incident to degree- b nodes, which is given by

$$\frac{bP_G(b)}{\sum_{a=0}^{n-1} aP_G(a)} = \frac{bP_G(b)}{Z_G}. \quad (6)$$

Let also two nodes be called r -neighbors of each other, for $r \geq 1$, if the distance between them in G is r . The case of $r = 1$ is simply the case of neighbors in G . Given a neighbor of a randomly chosen node u , the expected number of that neighbor's other neighbors (i.e., excluding u) is $\sum_{b=1}^{n-1} (b-1)bP_G(b)/Z_G$, so the expected number of 2-neighbors of u , which we denote by $Z_G^{(2)}$, is given by

$$Z_G^{(2)} = \sum_{a=0}^{n-1} aP_G(a) \sum_{b=1}^{n-1} (b-1) \frac{bP_G(b)}{Z_G} = Z_G \sum_{b=1}^{n-1} (b-1) \frac{bP_G(b)}{Z_G}. \quad (7)$$

In general, we denote a node's expected number of r -neighbors by $Z_G^{(r)}$. Clearly, $Z_G^{(1)} = Z_G$, and (7) can be generalized to yield

$$Z_G^{(r)} = Z_G^{(r-1)} \sum_{b=1}^{n-1} (b-1) \frac{bP_G(b)}{Z_G} = Z_G \left(\sum_{b=1}^{n-1} (b-1) \frac{bP_G(b)}{Z_G} \right)^{r-1} = Z_G \left(\frac{Z_G^{(2)}}{Z_G} \right)^{r-1}. \quad (8)$$

Note that here we have repeatedly taken into account, as we consider nodes' neighbors that are progressively farther from the initial, randomly chosen node u , that the probability that any two nodes involved in the process are in fact the same node is 0 in the limit as $n \rightarrow \infty$.

We can then obtain an approximation for the expected path length of G by summing the values of $Z_G^{(r)}$ from $r = 1$ up until the sum becomes equal to $n - 1$. When this happens, the current value of r can be taken as L_G , the desired expected path length. Thus, $\sum_{r=1}^{L_G} Z_G^{(r)} = n - 1$, which yields

$$L_G = \frac{\ln \left[\left(\frac{n-1}{Z_G} \right) \left(\frac{Z_G^{(2)}}{Z_G} - 1 \right) + 1 \right]}{\ln(Z_G^{(2)}/Z_G)}. \quad (9)$$

We henceforth assume that G is above the phase transition, that is, G almost surely has a GCC. Given two adjacent nodes u and v , let the reach of u through v be the set of nodes reachable by a path starting at u whose first edge is (u, v) . A randomly chosen node is outside the GCC if and only if all of its neighbors are also outside the GCC, that is, it has a small, size- $o(n)$ reach through each of its neighbors. Denoting by q the probability that this happens for each of those neighbors, we can express q as

$$q = \sum_{a=1}^{n-1} q^{a-1} \frac{aP_G(a)}{Z_G}, \quad (10)$$

that is, a node has a small reach through one of its neighbors if and only if that neighbor itself has a small reach through each of its other neighbors. As a result, the probability that a randomly chosen node is outside the GCC is $\sum_{a=0}^{n-1} q^a P_G(a)$, so the fraction of nodes inside the GCC of G , denoted by θ_G , is given by

$$\theta_G = 1 - \sum_{a=0}^{n-1} q^a P_G(a). \quad (11)$$

B. Extensions within the GCC

One further assumption that we make in this paper is that the originator, the node that initially has the information and starts the dissemination, is inside the GCC. In this case our analyses must be based on a degree distribution that is conditioned upon this membership in the GCC. The conditional probability that a node inside the GCC has degree a , denoted by $P_G(a \mid \text{GCC})$, can be written using Bayes' rule as $P_G(a \mid \text{GCC}) = P_G(\text{GCC} \mid a)P_G(a)/P_G(\text{GCC})$. Here $P_G(\text{GCC})$ is the probability that a randomly chosen node is inside the GCC, that is, θ_G . $P_G(\text{GCC} \mid a)$, in turn, is the probability that a degree- a node is inside the GCC. Such a node is outside the GCC if and only if it has a small reach through each of its neighbors, which occurs with probability q for each neighbor. Then a degree- a node is inside the GCC with probability $P_G(\text{GCC} \mid a) = 1 - q^a$. We then have

$$P_G(a \mid \text{GCC}) = \frac{(1 - q^a) P_G(a)}{\theta_G}. \quad (12)$$

We can now follow our earlier steps and calculate the expected number of neighbors and 2-neighbors of a randomly chosen node inside the GCC, which we refer to as Z_{GCC} and $Z_{\text{GCC}}^{(2)}$, respectively. We obtain

$$Z_{\text{GCC}} = \sum_{a=0}^{n-1} a P_G(a \mid \text{GCC}) = \sum_{a=0}^{n-1} a \frac{(1 - q^a) P_G(a)}{\theta_G} \quad (13)$$

and

$$Z_{\text{GCC}}^{(2)} = Z_{\text{GCC}} \sum_{a=1}^{n-1} (a-1) \frac{a P_G(a)}{Z_G}. \quad (14)$$

We can also retrace the steps that led us to (9) and obtain L_{GCC} , the expected path length inside the GCC:

$$L_{\text{GCC}} = \frac{\ln \left[\left(\frac{n\theta_G - 1}{Z_{\text{GCC}}} \right) \left(\frac{Z_{\text{GCC}}^{(2)}}{Z_{\text{GCC}}} - 1 \right) + 1 \right]}{\ln(Z_{\text{GCC}}^{(2)}/Z_{\text{GCC}})}. \quad (15)$$

IV. MATHEMATICAL ANALYSIS

We start by considering a generic probabilistic algorithm for disseminating a piece of information I on the network represented by G . In this algorithm, when a node of degree a receives I for the first time, it forwards I to each of its neighbors with probability $f(a, b)$, where b is the degree of the prospective recipient. The forwarding process is then suggestive of the random directed subgraph $F = (N_F, E_F)$ of G that we henceforth call the flooding digraph. This digraph has the same nodes as G (that is, $N_F = N_G$) and its edges are such that, for $(u, v) \in E_G$, the directed edge $(u \rightarrow v)$ exists in E_F with probability $f(a, b)$, given that a and b are the degrees of nodes u and v in G , respectively, and similarly for the edge $(v \rightarrow u)$.

Clearly, this generic algorithm can stand for both probabilistic flooding and heuristic flooding. In the former case, $f(a, b) = p$ regardless of a or b ; in the latter, $f(a, b)$ is the heuristic function $h(a, b)$ given by (2). As probabilistic flooding is a special case of heuristic flooding, for the remainder of this section we concentrate solely on the latter and use the $h(a, b)$ of (2) instead of $f(a, b)$.

A. Random digraphs

Let us first review some of the properties of random digraphs in general and also how they apply to the case of F .

For u a node of a digraph, its in-neighbors are those nodes from which an edge exists directed toward u ; its out-neighbors are those nodes toward which an edge exists directed from u . Moreover, a node is an r -in-neighbor of u if the directed distance from it to u is r . Similarly, a node is an r -out-neighbor of

Fig. 1. The structure of the GWCC of F and its relation to G and to the GCC of G .

u if the directed distance from u to it is r . When a directed path exists from node u to node v , we say that v is reachable from u (equivalently, u reaches v).

A connected component of the undirected graph that underlies a digraph (i.e., the graph that we obtain when edge directions are ignored) is called a weakly connected component of the digraph. Thus a digraph has a size- $\Theta(n)$ weakly connected component, referred to as the giant weakly connected component (GWCC), if and only if the undirected graph that underlies it has a GCC.

The GWCC of a digraph has four distinguished types of sub-digraphs. In the case of F , they are as illustrated in Figure 1 and such that:

- The giant strongly connected component (GSCC) is the largest sub-digraph of the GWCC that is maximal with respect to the property that any of its nodes is reachable from any other.
- The giant in-component (GIN), comprising a fraction θ_F^{in} of the nodes of G (or a fraction $\theta_F^{\text{in}}/\theta_G$ of the nodes of the GCC of G), contains all the nodes that can reach the GSCC (including, by definition, those of the GSCC).
- The giant out-component (GOUT), comprising a fraction θ_F^{out} of the nodes of G (or a fraction $\theta_F^{\text{out}}/\theta_G$ of the nodes of the GCC of G), contains all the nodes that are reachable from the GSCC (including, by definition, those of the GSCC).
- Each of the so-called tendrils consists of some of the remaining nodes [16].

B. Number of nodes reached

We now consider the expected fraction of nodes of G that receive I in heuristic flooding. Recall that, by assumption, G almost surely has a GCC and the originator is one of the nodes of the GCC. Two cases must be considered. The first case corresponds to the originator being outside the GIN, belonging therefore either to the portion of the GOUT that does not intersect the GSCC or to a tendril. This case happens with probability $1 - \theta_F^{\text{in}}/\theta_G$ and leads to a negligibly small number of nodes reached by the flooding. The second case is the case in which the originator is inside the GIN. It occurs with probability $\theta_F^{\text{in}}/\theta_G$ and the flooding necessarily reaches all the nodes of the GOUT. Nodes outside the GOUT may also be reached but contribute negligibly, since they belong either to the portion of the GIN that does not intersect the GSCC or to a tendril. So, in essence, the probability that a node is reached, denoted by P_n , is

$$P_n = \frac{\theta_F^{\text{in}}\theta_F^{\text{out}}}{\theta_G^2}. \quad (16)$$

In order to calculate θ_F^{in} and θ_F^{out} , let us first define the ancestry of a node as the set of nodes from which it can be reached in F , and the descent of a node as the set of nodes reachable from it in F . Then a randomly chosen node is inside the GIN if and only if it has a large, size- $\Theta(n)$ descent. Now consider

a randomly chosen node u having degree a in G . We say that a neighbor v of u in G is a dead end with respect to u if either $(u \rightarrow v)$ is not an edge of F or it is but v has a small descent. If b is the degree of v in G , we denote by q_b^{out} the conditional probability that v has a small descent, given that it has a directed edge in F incoming from u . The probability that a degree- b node is a dead end with respect to a degree- a neighbor in G , which we denote by $w_{a,b}^{\text{out}}$, is then

$$w_{a,b}^{\text{out}} = 1 - h(a, b) + h(a, b)q_b^{\text{out}}. \quad (17)$$

But the probability that a node's neighbor in G has degree b is given by (6), so the probability that a given neighbor of a degree- a node is a dead end is $\sum_{b=1}^{n-1} w_{a,b}^{\text{out}} b P_G(b) / Z_G$. We also know that such a degree- b node has a small descent if and only if each of its $b-1$ other neighbors in G is itself a dead end with respect to it. Thus, q_b^{out} is such that

$$q_b^{\text{out}} = \left(\sum_{c=1}^{n-1} w_{b,c}^{\text{out}} \frac{c P_G(c)}{Z_G} \right)^{b-1}. \quad (18)$$

It now suffices to recognize that a randomly chosen node is outside the GIN if and only if all of its neighbors in G are dead ends with respect to it. The fraction of nodes inside the GIN is then given by

$$\theta_F^{\text{in}} = 1 - \sum_{a=0}^{n-1} \left(\sum_{b=1}^{n-1} w_{a,b}^{\text{out}} \frac{b P_G(b)}{Z_G} \right)^a P_G(a). \quad (19)$$

The equations in (17) and (18) lead to a nonlinear system of $n-1$ equations (letting $b = 1, \dots, n-1$ in (18)) on $n-1$ variables (q_1^{out} through q_{n-1}^{out}). A solution of this system within $[0, 1]^{n-1}$ can be used, via (17), to calculate θ_F^{in} in (19).

The calculation of θ_F^{out} follows a completely analogous development, since a randomly chosen node is inside the GOUT if and only if it has a large, size- $\Theta(n)$ ancestry. This leads to

$$w_{a,b}^{\text{in}} = 1 - h(b, a) + h(b, a)q_b^{\text{in}}, \quad (20)$$

where

$$q_b^{\text{in}} = \left(\sum_{c=1}^{n-1} w_{b,c}^{\text{in}} \frac{c P_G(c)}{Z_G} \right)^{b-1} \quad (21)$$

is the conditional probability that a degree- b node having a degree- a neighbor in G has a small ancestry, given that in F it has an edge outgoing to that degree- a neighbor. The counterpart of (19) is then

$$\theta_F^{\text{out}} = 1 - \sum_{a=0}^{n-1} \left(\sum_{b=1}^{n-1} w_{a,b}^{\text{in}} \frac{b P_G(b)}{Z_G} \right)^a P_G(a). \quad (22)$$

C. Number of messages sent

Let P_m be the ratio of the expected number of messages sent by heuristic flooding to the expected number of messages sent by uninformed flooding. Given that the originator of both algorithms is inside the GCC, the expected number of messages sent by uninformed flooding is $Z_{\text{GCC}} n \theta_G$. Letting Z_r be the expected out-degree in F of the nodes reached by heuristic flooding, the expected number of messages in this case is $Z_r n \theta_G P_n$. We then have

$$P_m = \frac{Z_r P_n}{Z_{\text{GCC}}}. \quad (23)$$

The value of Z_r can be approximated by the average out-degree of the nodes inside the GOUT, which we denote by Z_{GOUT}^- . Letting $P_F^-(i | \text{GOUT})$ be the conditional probability that a randomly chosen node has out-degree i in F , given that it is inside the GOUT, Z_{GOUT}^- can be expressed as

$$Z_{\text{GOUT}}^- = \sum_{i=0}^{n-1} i P_F^-(i | \text{GOUT}). \quad (24)$$

If we let $P_F^-(i | a, \text{GOUT})$ be the conditional probability that a node has out-degree i in F , given that it has degree a in G and that it is inside the GOUT, then $P_F^-(i | \text{GOUT})$ can be written as

$$P_F^-(i | \text{GOUT}) = \sum_{a=i}^{n-1} P_F^-(i | a, \text{GOUT}) P_G(a | \text{GOUT}). \quad (25)$$

Now recall that every node's degree in G is by assumption independent and identically distributed with respect to all others'. Thus, letting h_a^- be the expectation of the heuristic function $h(a, b)$ as b varies, that is,

$$h_a^- = \sum_{b=1}^{n-1} h(a, b) \frac{b P_G(b)}{Z_G}, \quad (26)$$

the conditional probability that a degree- a node has out-degree i in F , which we denote by $P_F^-(i | a)$, is given by the binomial distribution:

$$P_F^-(i | a) = \binom{a}{i} (h_a^-)^i (1 - h_a^-)^{a-i}. \quad (27)$$

But given a node inside the GOUT, (6) still gives the probability that one of its neighbor's degree is b in G , albeit only approximately. Therefore, a node with degree a in G that is inside the GOUT has out-degree i in F with probability given by (27), that is, $P_F^-(i | a, \text{GOUT}) = P_F^-(i | a)$, which leads to

$$P_F^-(i | \text{GOUT}) = \sum_{a=i}^{n-1} P_F^-(i | a) P_G(a | \text{GOUT}). \quad (28)$$

We can then re-write (24) as

$$\begin{aligned} Z_{\text{GOUT}}^- &= \sum_{i=0}^{n-1} i \sum_{a=i}^{n-1} \binom{a}{i} (h_a^-)^i (1 - h_a^-)^{a-i} P_G(a | \text{GOUT}) \\ &= \sum_{a=0}^{n-1} P_G(a | \text{GOUT}) \sum_{i=0}^a i \binom{a}{i} (h_a^-)^i (1 - h_a^-)^{a-i}, \end{aligned} \quad (29)$$

where the rightmost summation gives the expected out-degree in F of a node whose degree in G is a , i.e., $a h_a^-$. Therefore,

$$Z_{\text{GOUT}}^- = \sum_{a=0}^{n-1} P_G(a | \text{GOUT}) a h_a^-. \quad (30)$$

Now, if we apply Bayes' rule and write $P_G(a | \text{GOUT})$ as

$$P_G(a | \text{GOUT}) = \frac{P_G(\text{GOUT} | a) P_G(a)}{P_G(\text{GOUT})} = \frac{P_G(\text{GOUT} | a) P_G(a)}{\theta_F^{\text{out}}}, \quad (31)$$

and furthermore recognize that a node of degree a in G is outside the GOUT (with probability $1 - P_G(\text{GOUT} | a)$) if and only if each of its neighbors in G is like the degree- b nodes at the end of

Section IV-B, and also that this occurs for each such node with probability $\sum_{b=1}^{n-1} w_{a,b}^{\text{in}} b P_G(b) / Z_G$, then we obtain

$$P_G(\text{GOUT} \mid a) = 1 - \left(\sum_{b=1}^{n-1} w_{a,b}^{\text{in}} \frac{b P_G(b)}{Z_G} \right)^a, \quad (32)$$

culminating in

$$Z_{\text{GOUT}}^- = \sum_{a=0}^{n-1} a h_a^- \left[1 - \left(\sum_{b=1}^{n-1} w_{a,b}^{\text{in}} \frac{b P_G(b)}{Z_G} \right)^a \right] \frac{P_G(a)}{\theta_F^{\text{out}}}. \quad (33)$$

D. Average waiting time

As we mentioned in Section I, assessing a distributed algorithm's time-related complexities in a fully general asynchronous setting requires that message-passing causal chains be considered [2]. In our current context, and taking into account our current knowledge of random graphs and their analysis, obtaining accurate analytical estimates of those complexities for flooding seems infeasible. Our choice is then to settle for a less general form of asynchronism in which the delay associated with the delivery of a message over an edge of G is roughly constant for all of G 's edges. In this case, the waiting time before a node is reached by a flooding can be taken to be proportional to the distance in G from the originator to that node.

In order to compare the expected waiting time of heuristic flooding with that of uninformed flooding, we calculate the ratio P_t of the expected path length from the originator in F to the expected path length from the originator in G . Recall that, even though we employ the common denomination as a path, in the case of F paths are directed, while in G they are undirected.

For uninformed flooding with originator inside the GCC, the expected path length from the originator can be approximated by L_{GCC} as in (15). In order to obtain the corresponding expression for heuristic flooding, we proceed as we did for P_n and P_m and consider only the case in which the originator is inside the GIN, which occurs with probability $\theta_F^{\text{in}} / \theta_G$. We also assume that only the nodes inside the GOUT can be reached by the flooding. Letting L_{GOUT}^- be the expected path length from the originator under these assumptions, P_t is such that

$$P_t = \frac{\theta_F^{\text{in}} L_{\text{GOUT}}^-}{\theta_G L_{\text{GCC}}}. \quad (34)$$

The average out-degree inside the GOUT, Z_{GOUT}^- , is given by (33). In order to obtain the expected number of 2-out-neighbors of a randomly chosen node inside the GOUT, denoted by $Z_{\text{GOUT}}^{-(2)}$, we first consider two adjacent nodes u_1 and u_2 having degrees a_1 and a_2 , respectively, in G . If the directed edge $(u_1 \rightarrow u_2)$ does not exist in F , which occurs with probability $1 - h(a_1, a_2)$, then u_1 has no 2-out-neighbors reachable through u_2 . On the other hand, if the directed edge does exist, and this has probability $h(a_1, a_2)$ of occurring, then the number of 2-out-neighbors that u_1 can reach through u_2 is expected to be $(a_2 - 1) h_{a_2}^-$. So each neighbor of u_1 in G provides an expected number of 2-out-neighbors to it that is given approximately by $\sum_{a_2=1}^{n-1} (a_2 - 1) h_{a_2}^- h(a_1, a_2) a_2 P_G(a_2) / Z_G$. Now let $t_{a_1, a_2}(x)$ be the function

$$t_{a_1, a_2}(x) = x (a_2 - 1) h(a_1, a_2) \frac{a_2 P_G(a_2)}{Z_G}. \quad (35)$$

Then $Z_{\text{GOUT}}^{-(2)}$ is given by

$$Z_{\text{GOUT}}^{-(2)} = \sum_{a_1=0}^{n-1} a_1 P_G(a_1 \mid \text{GOUT}) \sum_{a_2=1}^{n-1} t_{a_1, a_2}(h_{a_2}^-). \quad (36)$$

Proceeding likewise, we see that, for $r > 1$, the expected number of r -out-neighbors of a node inside the GOUT is given by

$$Z_{\text{GOUT}}^{-(r)} = \sum_{a_1=0}^{n-1} a_1 P_G(a_1 \mid \text{GOUT}) \sum_{a_2=1}^{n-1} t_{a_1, a_2} \left(\sum_{a_3=1}^{n-1} t_{a_2, a_3} \left(\cdots \sum_{a_r=1}^{n-1} t_{a_{r-1}, a_r} (h_{a_r}^-) \right) \right). \quad (37)$$

But unlike the $Z_G^{(r)}$ of (8) or the $Z_{\text{GCC}}^{(r)}$ that implicitly led to (15), now the sequence $Z_{\text{GOUT}}^-, Z_{\text{GOUT}}^{-(2)}, \dots$ is not a geometric progression and L_{GOUT}^- cannot be expressed by an equation analogous to (15). We could, however, in principle, use (37) to obtain $\lfloor L_{\text{GOUT}}^- \rfloor$ and $\lceil L_{\text{GOUT}}^- \rceil$ by extending the sum $Z_{\text{GOUT}}^- + Z_{\text{GOUT}}^{-(2)} + \dots$ up until a number greater than $n - 1$ were obtained. The last two values to go into the sum would be $Z_{\text{GOUT}}^{-(\lfloor L_{\text{GOUT}}^- \rfloor)}$ and $Z_{\text{GOUT}}^{-(\lceil L_{\text{GOUT}}^- \rceil)}$. But the time complexity of calculating $Z_{\text{GOUT}}^{-(r)}$ is $O(n^r)$, so this method is not really practical.

A way out does exist, however, in some special cases. In the case of probabilistic flooding, for example, we have $h(a_1, a_2) = h_a^- = p$ for $1 \leq a_1, a_2, a \leq n - 1$, so $t_{a_1, a_2}(x)$ is independent of a_1 , leading (37) to be simplified as

$$Z_{\text{GOUT}}^{-(r)} = \sum_{a_1=0}^{n-1} a_1 P_G(a_1 \mid \text{GOUT}) \left(\sum_{a_2=1}^{n-1} t_{a_1, a_2}(1) \right)^{r-1} p. \quad (38)$$

In this case, the sequence $Z_{\text{GOUT}}^{-(1)}, Z_{\text{GOUT}}^{-(2)}, \dots$ is indeed a geometric progression, so we can proceed as we did for the undirected case and obtain

$$L_{\text{GOUT}}^- = \frac{\ln \left[\left(\frac{n\theta_F^{\text{out}} - 1}{Z_{\text{GOUT}}^-} \right) \left(\frac{Z_{\text{GOUT}}^{-(2)}}{Z_{\text{GOUT}}^-} - 1 \right) + 1 \right]}{\ln(Z_{\text{GOUT}}^{-(2)} / Z_{\text{GOUT}}^-)}. \quad (39)$$

The case of a Poisson degree distribution is also amenable to further analysis, since we can approximate $Z_{\text{GOUT}}^{-(r)}$ by $Z_{\text{GOUT}}^- \rho^{r-1}$, where ρ is the expected number of out-neighbors that a node chosen by moving along the direction of a randomly chosen edge of F has [17]. Let us consider such an edge of F . The joint probability that the node from which this edge outgoes has degree a_1 in G and that the node to which it incomes has degree a_2 in G is denoted by $P_G(a_1, a_2 \mid e)$, where e represents the event that the edge exists in F . Using Bayes' rule, this probability can be written as

$$P_G(a_1, a_2 \mid e) = \frac{P_G(e \mid a_1, a_2) P_G(a_1, a_2)}{P(e)} = \frac{h(a_1, a_2) a_1 P_G(a_1) a_2 P_G(a_2)}{Z_G \sum_{b=1}^{n-1} h_b^- b P_G(b)}, \quad (40)$$

which leads to

$$\rho = \sum_{a_1=1}^{n-1} \sum_{a_2=1}^{n-1} (a_2 - 1) h_{a_2}^- \frac{h(a_1, a_2) a_1 P_G(a_1) a_2 P_G(a_2)}{Z_G \sum_{b=1}^{n-1} h_b^- b P_G(b)}, \quad (41)$$

and finally to

$$L_{\text{GOUT}}^- = \frac{\ln \left[\left(\frac{n\theta_F^{\text{out}} - 1}{Z_{\text{GOUT}}^-} \right) (\rho - 1) + 1 \right]}{\ln \rho}. \quad (42)$$

V. SIMULATION RESULTS

In our simulations we have used random graphs as models for networks [9]. The network may thus not be connected, so simulations have only been carried out inside the largest connected component of the random graph. Also, in order to ensure that such a component encompasses a large number of nodes, we have restricted ourselves to graphs for which a GCC is almost surely guaranteed to exist. Our interest has been to analyze P_n , P_m , and P_i , and to this end we have simulated probabilistic flooding for $p = 0.30, 0.60, 0.90$ and heuristic flooding for $\alpha = 0.90, 0.95, 0.99$.

Fig. 2. Simulation of probabilistic (a–c) and heuristic (d–f) flooding on random graphs with Poisson-distributed degrees. The plots show P_n (a and d), P_m (b and e), and P_t (c and f) for $p = 0.30, 0.60, 0.90$ and $\alpha = 0.90, 0.95, 0.99$. Solid lines give the analytical predictions of Section IV.

We have considered two random-graph models. In the first model G is generated on n nodes by creating an edge between any two distinct nodes with fixed probability $z/(n-1)$ for suitable z . A node in the resulting graph has degree a with the Poisson probability of mean z [18], that is, $P_G(a) = e^{-z} z^a / a!$. Also, we have $Z_G = z$ and $\langle K_G^2 \rangle = z^2 + z$, which by (5) implies that we need $z > 1$ for the GCC to almost surely exist.

Figure 2 shows the results obtained by simulating probabilistic and heuristic flooding on random graphs with Poisson-distributed degrees for $n = 10000$ and z varying from 1 to 10. For each value of z , the simulation consisted of generating 15 random graphs and, for each one, doing 1000 instances of each type of flooding (uninformed, probabilistic, and heuristic), each from a randomly chosen node inside the GCC. The simulation results are given as averages of P_n , P_m , and P_t over the 15000 samples. Notice first of all that agreement between simulation and analytical results is very good throughout, with a slight deviation only in parts (c) and (f) of the figure. Such deviations are attributed to the various approximations of Section IV-D. Note also that both probabilistic and heuristic flooding do indeed improve on uninformed flooding as far as the number of messages used is concerned ($P_m < 1$), but nearly always at the expense of larger waiting times ($P_t > 1$).

The results for probabilistic flooding (Figure 2(a–c)) show that, given a value of p , there exists a value of z above which nearly all the nodes of the network are reached by the flooding. If we require $P_n \geq 0.99$, then the threshold values for z are $z = 5$ for $p = 0.90$ and $z = 9$ for $p = 0.60$. For $p = 0.30$ the threshold is greater than 10. The plots for P_m show that P_m increases with z roughly until these thresholds are reached, remaining constant and approximately equal to p from there onward. The plots for P_t show a similar but not identical behavior: like P_m , P_t often increases roughly until the thresholds are reached; unlike P_m , P_t starts to decrease slightly as z is further increased onward. All the plots show that P_n , P_m , and P_t have value 0 when z is near 1. This happens because, although the network is above the phase transition that gives rise to the GCC of G with high probability for any $z > 1$, for z sufficiently near 1 the flooding digraph F is still below the phase transition at which the GIN and the GOUT appear (since F is a directed subgraph of G), so they do not yet exist.

In the case of heuristic flooding (Figure 2(d–f)), the results show that P_n flattens out much earlier than probabilistic flooding to a value that depends on α . This value is about 0.99 for $\alpha = 0.99$, 0.94 for $\alpha = 0.95$, and 0.86 for $\alpha = 0.90$. The value of P_m , in turn, increases sharply until approximately the value at which the flattening of P_n occurs. Past this value, P_m is seen to decrease continually, which indicates the desirable property that, given a fixed P_n , a progressively smaller fraction of messages is needed as the average degree increases. In contrast, P_t increases onward through the higher values of z . This occurs because heuristic flooding tends to send messages on relatively longer paths, avoiding as it does the sending of messages between high-degree nodes.

The second random-graph model we have considered is the one in which degrees are distributed according to a power law. In such a random graph, the probability that a node has degree a is $P_G(a) = Ca^{-\tau}$, where C is a normalizing constant and $\tau > 1$ is a parameter. Clearly, in the limit as $n \rightarrow \infty$ we have $C = 1/\zeta(\tau)$, where $\zeta(x)$ is the Riemann zeta function [19], that is, $\zeta(x) = \sum_{y=1}^{\infty} y^{-x}$. We also have $Z_G = \zeta(\tau - 1)/\zeta(\tau)$ and $\langle K_G^2 \rangle = \zeta(\tau - 2)/\zeta(\tau)$, so the criterion of (5) for the appearance of the GCC can be seen to translate into $\zeta(\tau - 2)/\zeta(\tau - 1) > 2$. Solving this inequality for τ numerically reveals that the GCC almost surely exists when $\tau < 3.47$.

Generation of a random graph with degrees thus distributed can be achieved in two phases [8]. First, the degrees $a_1, a_2, \dots, a_n$ of the n nodes, constituting the graph's so-called degree sequence, are sampled from the power-law distribution and $\sum_{i=1}^n a_i$ labeled balls are put inside an imaginary urn.¹ For $1 \leq i \leq n$, label i is given to exactly a_i balls. In the second phase, a pair of balls (labeled, say, u and v) is picked from the urn at random and an edge (u, v) is added to the graph. This process is repeated until the urn becomes empty. This algorithm clearly generates a multigraph (a graph in which multiple edges and self-loops are allowed to exist), but the resulting graph is consistent with the model used in our analysis, where we assumed edge independence throughout. Furthermore, even though there are other algorithms to generate random graphs with a given degree sequence [20], it is still unclear how to use them within reasonable time bounds. The fact that multiple edges or self-loops may now exist in G has direct impact on how the various flooding algorithms are simulated. Specifically, a node no longer considers whether to forward the information it receives for the first time to each of its neighbors, but rather whether to forward it on each of the edges that are incident to it in G .

Figure 3 shows the results obtained by simulating the three flooding methods on random graphs with power-law-distributed degrees for $n = 10000$ and τ varying from 2 to 3. For each value of τ , the simulation consists of generating 300 random graphs and, for each one, proceeding exactly as indicated for the Poisson case. First notice, again as in the Poisson case, that agreement between simulation and analytical results is very good for P_n and P_m . For P_t , however, the situation is different. Under a power-law distribution for node degrees, expressions like (14) and (36), respectively for $Z_{\text{GCC}}^{(2)}$ and $Z_{\text{GOUT}}^{-(2)}$, are known to be problematic [8], and indeed the calculations do not converge and lead to wrong values for L_{GCC} and L_{GOUT}^- . Even so, it is curious to note that for probabilistic flooding the two errors seem to compensate each other somehow and the prediction for P_t comes close to the simulation results, as shown in part (c) of the figure. The case of heuristic flooding, on the other hand, remains lacking a satisfactory analytical prediction (thence none is shown in part (f) of the figure).

Also noteworthy in Figure 3 is the consistent superiority of probabilistic and heuristic flooding over uninformed flooding regarding the number of messages sent ($P_m < 1$). In addition, and contrasting with the Poisson case of Figure 2, the same holds for probabilistic flooding, and for sufficiently large τ also for heuristic flooding, regarding waiting times ($P_t < 1$). The price for this, of course, is heavy and reflected on the often poor values of P_n , with the exception of probabilistic flooding with $p = 0.90$ and of heuristic flooding, more or less regardless of the value of α , when τ is near 2.

Simulation results for probabilistic flooding (Figure 3(a–c)) show that P_n increases when τ decreases, but the variation is sometimes not really too pronounced. For example, for $p = 0.90$ P_n is about 0.60 when $\tau = 3.0$ and 0.87 when $\tau = 2.0$. Similar observations are true of P_m and P_t . Notice, in addition,

¹If $\sum_{i=1}^n a_i$ turns out to be odd, the degree sampling is repeated until an even sum is obtained.

Fig. 3. Simulation of probabilistic (a–c) and heuristic (d–f) flooding on random graphs with power-law-distributed degrees. The plots show P_n (a and d), P_m (b and e), and P_t (c and f) for $p = 0.30, 0.60, 0.90$ and $\alpha = 0.90, 0.95, 0.99$. Solid lines give the analytical predictions of Section IV.

that unlike the Poisson case there does not seem to exist any interval within the range considered for τ inside which any of the three indicators becomes approximately constant.

In heuristic flooding (Figure 3(d–f)), the value of P_n also increases when τ decreases, but at a faster pace than in the case of probabilistic flooding. The values of P_m and P_t display two different kinds of behavior. While to either side of $\tau \approx 2.4$ the value of P_m decreases as τ is moved farther away, P_t falls relatively sharply to the right of $\tau \approx 2.4$ but rises, albeit slowly and with an inversion with respect to the values of α , to the left. From Figure 3(d), $\tau \approx 2.4$ appears to coincide with a significant change in the second derivative of P_n as a function of τ .

VI. PROBABILISTIC VERSUS HEURISTIC FLOODING

In spite of having already obtained results on both types of flooding, we still cannot compare them to each other directly, as the parameter p of probabilistic flooding has no relation whatsoever to the parameter α of heuristic flooding. In this section, our strategy to make such a comparison possible is to first simulate heuristic flooding with $\alpha = 0.99$ and then use the resulting value of P_n to obtain, via (16), the value of p for which probabilistic flooding is expected to reach the same fraction of the graph's nodes.

Our results under this strategy are shown in Figure 4 for the Poisson case. Notice, first, that the strategy seems indeed to be effective, since the plots for P_n are practically the same for both types of flooding. For z right past 1, all the six plots exhibit an anomaly, which is once again likely due to the fact that, at these values of z , the phase transition that gives rise to the GIN and the GOUT has not yet taken place.

What is most interesting, though, is the evident trade-off between P_m and P_t that can be seen in parts (b) and (c) of the figure. Clearly, under the constraint that P_n is roughly the same for both probabilistic and heuristic flooding, for z larger than approximately 2.5, P_m is significantly larger for probabilistic flooding than it is for heuristic flooding. On the other hand, P_t is significantly larger for heuristic flooding across practically all the spectrum of z values that Figure 4 encompasses.

But notwithstanding this trade-off, Figure 4 also suggests that the two types of flooding become more and more similar to each other as z is increased. In fact, as the nodes' degrees become larger and G more

Fig. 4. Comparison of probabilistic flooding and heuristic flooding in the Poisson case. Plots are shown for P_n (a), P_m (b), and P_t (c).

Fig. 5. Comparison of probabilistic flooding and heuristic flooding for the power-law case. Plots are shown for P_n (a), P_m (b), and P_t (c).

homogeneous (approaching, in the limit, the complete graph on n nodes), the heuristic function of (2) approaches a small constant and heuristic flooding progressively becomes probabilistic flooding.

The corresponding results for graphs with degrees distributed according to a power law are shown in the plots of Figure 5. Once again, our strategy's effectiveness is corroborated by part (a) of the figure. Also, parts (b) and (c) reveal the same trade-off between P_m and P_t that we observed in the Poisson case, along with a similar tendency, especially in the case of part (b), for the two algorithms to resemble each other as τ is increased toward 3. The justification for part (b) is that, at these values for τ , heuristic flooding has a small value for P_n (cf. part (a)), one that can be achieved by probabilistic flooding with a value of p for which P_m remains relatively low. As for part (c), increasing τ toward its higher values makes the occurrence of nodes of very high degree ever less likely. Consequently, heuristic flooding no longer refrains so strongly from sending the information being broadcast to the high-degree nodes of the network. This makes path lengths tend to lean toward those of probabilistic flooding.

But what is especially noteworthy in the power-law case is that the superiority of heuristic flooding in terms of P_m becomes very pronounced as τ is decreased from roughly $\tau = 2.5$ (and enters the purported Internet interval, by most accounts [9]). This happens because, under a power law, the graph tends to contain a large set of nodes of small degree, many of them of degree 1. In order to reach the same number of nodes as heuristic flooding, probabilistic flooding must run with a value for p that is near 1, which causes an unnecessarily high number of messages to be sent. Heuristic flooding, in turn, is relatively insensitive to the plethora of low-degree nodes, since it attempts to provide each and every node with the same probability of receiving the information.

VII. CONCLUSIONS

Flooding a network probabilistically for information dissemination is an attempt at reducing the heavy communications demand that uninformed flooding incurs in terms of how many messages are needed. The drawback, of course, is that the guarantee of network-wide information delivery may be lost. In this paper we started with the recently introduced probabilistic flooding, in which every node forwards the

information to all of its neighbors with constant probability. Then we introduced an alternative technique, called heuristic flooding, which employs degree-based probabilistic decisions at the nodes, aiming at a pre-defined probability, the same for all nodes, that a node receives the information being disseminated. We have also contributed a detailed mathematical analysis of both probabilistic flooding and heuristic flooding, comprising analytical predictions of the techniques' reachability, their communications requirements, and an indicator related to the time needed for completion of the flooding.

Heuristic flooding is controlled by a parameter that gives the desired probability that a node receives the information under dissemination. While for information broadcasting this parameter is to be set as close to 1 as possible, in retrospect we see that it has a role to play in the peer-to-peer context briefly discussed in Section I as well. Specifically, in such cases the parameter's value should be set as a function of how densely the objects are replicated throughout the network, however unstructured it may be. Even though we clearly see the inverse proportionality that should govern dissemination in this case, further analysis of applications of this specific nature is required to clarify how to set the parameter in each particular case. In any event, though, we see that our new probabilistic heuristic opens up interesting tracks of investigation that may in the future come to impact the design of distributed algorithms for unstructured networks significantly.

Our extensive simulation results have revealed the two techniques' main characteristics on random graphs with Poisson- and power-law-distributed degrees. They have also indicated an excellent agreement between theory and experimentation in most cases. In addition, one final set of simulations designed especially to allow a meaningful comparison between the two flooding techniques demonstrated an interesting trade-off involving them: in general, heuristic flooding outperforms probabilistic flooding in terms of communications requirements, but the opposite holds in terms of the delay required for completion. Curiously, the balance between the two sides of this trade-off is not equally tipped: for example, under a power law with relatively small parameter, the communications-related gain of heuristic flooding over probabilistic flooding significantly surpasses its delay-related loss.

REFERENCES

- [1] A. Segall, "Distributed network protocols," *IEEE Transactions on Information Theory*, vol. IT-29, pp. 23–35, 1983.
- [2] V. C. Barbosa, *An Introduction to Distributed Algorithms*. Cambridge, MA: The MIT Press, 1996.
- [3] F. Banaei-Kashani and C. Shahabi, "Criticality-based analysis and design of unstructured peer-to-peer networks as "complex systems"," in *Proceedings of the Third International Symposium on Cluster Computing and the Grid*, 2003, pp. 351–358.
- [4] I. Stoica, R. Morris, D. Liben-Nowell, D. R. Karger, M. F. Kaashoek, F. Dabek, and H. Balakrishnan, "Chord: a scalable peer-to-peer lookup protocol for internet applications," *IEEE/ACM Transactions on Networking*, vol. 11, pp. 17–32, 2003.
- [5] B. Y. Zhao, L. Huang, J. Stribling, S. C. Rhea, A. D. Joseph, and J. D. Kubiatowicz, "Tapestry: a resilient global-scale overlay for service deployment," *IEEE Journal on Selected Areas in Communications*, vol. 22, pp. 41–53, 2004.
- [6] W. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker, "Search and replication in unstructured peer-to-peer networks," in *Proceedings of the Sixteenth International Conference on Supercomputing*, 2002, pp. 84–95.
- [7] L. A. Adamic, R. M. Lukose, and B. A. Huberman, "Local search in unstructured networks," in *Handbook of Graphs and Networks: From the Genome to the Internet*, S. Bornholdt and H. G. Schuster, Eds. Weinheim, Germany: Wiley-VCH, 2003, pp. 295–317.
- [8] M. E. J. Newman, S. H. Strogatz, and D. J. Watts, "Random graphs with arbitrary degree distributions and their applications," *Physical Review E*, vol. 64, p. 026118, 2001.
- [9] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, pp. 47–97, 2002.
- [10] P. Erdős and A. Rényi, "On random graphs," *Publicationes Mathematicae*, vol. 6, pp. 290–297, 1959.
- [11] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, 1999, pp. 251–262.
- [12] A. Medina, I. Matta, and J. Byers, "On the origin of power laws in Internet topologies," *Computer Communication Review*, vol. 30, no. 2, pp. 18–28, 2000.
- [13] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore, "On the bias of traceroute sampling: or, power-law degree distributions in regular graphs," in *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*, 2005, pp. 694–703.
- [14] M. Molloy and B. Reed, "A critical point for random graphs with a given degree sequence," *Random Structures and Algorithms*, vol. 6, pp. 161–180, 1995.
- [15] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin, "Resilience of the Internet to random breakdowns," *Physical Review Letters*, vol. 85, pp. 4626–4628, 2000.
- [16] S. N. Dorogovtsev, J. F. F. Mendes, and A. N. Samukhin, "Giant strongly connected component of directed networks," *Physical Review E*, vol. 64, p. 025101, 2001.
- [17] R. M. Karp, "The transitive closure of a random digraph," *Random Structures and Algorithms*, vol. 1, pp. 73–93, 1990.

- [18] B. Bollobás, *Random Graphs*, 2nd ed. Cambridge, UK: Cambridge University Press, 2001.
- [19] S. Y. Yan, *Number Theory for Computing*, 2nd ed. Berlin, Germany: Springer-Verlag, 2002.
- [20] C. Gkantsidis, M. Mihail, and E. Zegura, "The Markov chain simulation method for generating connected power law random graphs," in *Proceedings of the Fifth Workshop on Algorithm Engineering and Experiments*, 2003.