

Distributed Creation and Adaptation of Random Scale-Free Overlay Networks

Ingo Scholtes*

*Systemsoftware and Distributed Systems Group

Department of Computer Science

University of Trier

D-54286 Trier, Germany

scholtes@syssoft.uni-trier.de

Abstract—Random scale-free overlay topologies provide a number of properties like for example high resilience against failures of random nodes, small (average) diameter as well as good expansion and congestion characteristics that make them interesting for the use in large-scale distributed systems. A number of these properties have been shown to be influenced by the exponent of their power law degree distribution. In this article, we present a distributed rewiring scheme that is suitable to effectuate random scale-free overlay topologies with an adjustable degree distribution exponent. The scheme uses a biased random walk strategy to sample new endpoints of edges being rewired and relies on an equilibrium model for scale-free networks. The bias of the random walk strategy can be tuned to produce random scale-free networks with arbitrary degree distribution exponents greater than two. We argue that the rewiring strategy can be implemented in a distributed fashion based on a node's information about its immediate neighbors. We present both analytical arguments as well as results that have been obtained in simulations of the proposed protocol.

Keywords-scale-free; overlay networks; adaptation; self-organization; Peer-to-Peer;

During the last decade, the increasing spread and importance of large-scale Peer-to-Peer systems has raised significant research interest in the design and analysis of robust and efficient overlay networks. In this research, structured and unstructured approaches can be distinguished. Mimicking the use of data structures in traditional computing, highly structured overlay topologies facilitate the use of efficient distributed algorithms with deterministic performance. However, the overhead entailed by the construction and maintenance of such deterministically structured topologies questions their usability in large-scale scenarios with dynamic and potentially faulty participants. Constituting a different approach, unstructured overlay networks do not impose constraints about the detailed structure of the emerging network topology. Rather than using costly and potentially complex routines for building and maintaining sophisticated network structures, in such unstructured overlays links can arise in a seemingly random and uncoordinated fashion. They are thus particularly suitable for highly dynamic scenarios in which the operational overhead entailed by structured approaches can possibly dominate a system's overall performance.

While the use of unstructured overlays can reduce construction and maintenance overhead, designing efficient dis-

tributed algorithms with predictable performance is hardly possible when making no assumptions whatsoever about an overlay's structure. Interestingly, based on a stochastic model of the system in question and arguments from random graph theory and complex network science, it is often possible to reason about structural properties of the resulting network topology that hold *almost surely* in the limit of large systems. Similarly, the performance of a number of dynamical processes - many of them relevant to distributed computing systems - has been studied in random network structures. For sufficiently large systems, based on randomized overlay topologies one can thus obtain strong, though probabilistic guarantees about their structure and performance. Considering the classical taxonomy of deterministically structured and completely unstructured overlay networks, this suggests an intermediate class of *probabilistically structured* topologies that promises to combine the benefits of both.

During the last decade, much of the work in the field of random networks has been focused around scale-free networks that are characterized by a power law degree distribution $P(k) \propto k^{-\gamma}$. The fact that networks with such scale-free characteristics seem to emerge naturally in a variety of natural, social and technological contexts has awakened the interest of researchers in disciplines as diverse as mathematics, statistical physics, biology, sociology, and computing. It has since been shown that scale-free networks provide a number of interesting properties like a remarkable robustness against random failures [4], [13], small diameter and average path lengths [8], [11] as well as favorable expansion and congestion properties [16], [33]. Some of these properties make them interesting for the design of large-scale computing systems and - in fact - for certain networked computing systems it has been observed that scale-free structures emerge in a seemingly self-organized way [3], [14], [27], [29].

Based on this observation, during the last couple of years, the performance of distributed algorithms operating in scale-free networks has been studied. For the problems of distributed search [1], [9], [26], information dissemination and entropy reduction protocols [21] as well as synchronization [31], distributed schemes have been derived that

seem to work particularly efficiently in scale-free networks. Considering a scale-free network topology with a degree distribution $P(k) \propto k^{-\gamma}$, it has further been argued that the exponent γ has massive influence on network properties like diameter [11], the vulnerability against targeted attacks as well as the performance of dynamical processes [7]. The reason for this can be found in the fact that the exponent γ determines the skewness of the degree distribution and thus the frequency and magnitude of highly connected hub nodes. For the practical design of scale-free overlay topologies, the degree distribution exponent is thus a critical parameter which largely influences their robustness, the performance of distributed algorithms as well as the distribution of load being imposed on individual machines.

However while these results are clearly important for the design of networked computing systems, many of them have been obtained for *equilibrium statistical ensembles* of random networks, that is for probability spaces in which networks with identical degree distribution are equiprobable. One needs to take care when wanting to apply this rich body of theoretical findings to practical systems. Even though the networks created by a certain process may exhibit a scale-free degree distribution, due to complex correlations that may be introduced by protocols influencing the structure of actual overlay topologies, their properties may very well differ from those of *random scale-free* networks. When wanting to rely on stochastic guarantees that hold for random networks, it is thus crucial to design protocols that recreate the appropriate statistical ensemble for which the guarantees have been obtained in the first place.

In this article, we study the question how random scale-free overlay topologies can be effectuated in practical Peer-to-Peer systems in a way that allows to reasonably apply the rich body of theoretical findings on equilibrium statistical ensembles of uncorrelated scale-free networks. For this, we study a protocol that creates explicitly randomized scale-free topologies with tunable degree distribution exponent. For the networks created by the proposed protocol, all theoretical findings about the properties of random scale-free networks as well as about the relation between the degree distribution exponent and the performance of dynamical processes operating upon them are immediately applicable. We further argue that the proposed protocol is suitable to adapt the degree distribution exponent and thus fine-tune the heterogeneity of overlay connectivity in a distributed and directed fashion while maintaining the overall scale-free characteristics of the topology.

The distributed rewiring mechanism presented in this article, is based on the equilibrium model of uncorrelated scale-free networks that has been considered in [18] as well as [25] and makes use of a biased random walk strategy in order to sample the endpoints of edges being rewired. As we shall see later, the efficiency and thus feasibility of the mechanism is based on the favorable expansion

properties of certain classes of random networks. A detailed description and derivation of the proposed protocol will be presented in section I. Here we further present some analytical arguments on the convergence behavior of the random walk sampling strategy underlying the protocol being presented in the subsequent section II. In section III we present simulation results that have been obtained using an implementation of the proposed rewiring scheme. Having briefly reviewed related work, in section V we conclude the article by summarizing its main contributions and pointing out a number of open issues and threats to validity that could not be addressed so far.

I. CREATING AND ADAPTING SCALE-FREE OVERLAYS

As has been argued above, the exponent γ is a macroscopic, statistical parameter that influences the structural properties of random networks with a power law degree distribution $P(k) \propto k^{-\gamma}$. In the following, we thus intend to derive a distributed protocol that can be used to effectuate network topologies that are randomly drawn from the ensemble of scale-free networks with a particular degree distribution exponent. As initial situation, we assume an arbitrary, connected overlay topology. While for the functioning of the scheme no particular initial state of the overlay is required as long as it is connected, it will later be argued that the initial topology influences the efficiency of the scheme in terms of the number of messages that need to be exchanged. For simplicity, we further assume that each of the n nodes is uniquely identified by a numeric identifier $i \in \{1, \dots, n\}$. However in sufficiently large systems, per-node quantities i that are chosen uniformly at random - and thus not necessarily unique - can be used instead. In order to simplify derivation and analysis, here we further consider a static situation in which no nodes enter or leave the system. Clearly, the main motivation to use a probabilistic overlay topology in the first place is to support highly dynamic systems in which node joins and exits are frequent. In sections II-B and III, we will thus consider dynamic equilibrium situations with fluctuating participants.

In order to derive a distributed scheme that can be used to influence the structure of scale-free overlays, we first need a model that is capable of generating uncorrelated network topologies with tunable degree distribution exponents. Here we use a simple equilibrium model for scale-free networks with a fixed number of nodes that has been introduced in [18] and analyzed in [25]. In this model it is assumed that each node $i \in \{1, \dots, n\}$ is assigned a weight

$$w_i = i^{-\alpha}$$

for some parameter α in the range $(0, 1)$. It is then assumed that m edges are created between pairs of nodes (i, j) chosen with probabilities p_i and p_j that are given by the normalized weights

$$p_i = \frac{w_i}{\sum_{k=1}^n w_k}. \quad (1)$$

As has been argued in [25], this simple model produces uncorrelated random scale-free networks with a degree distribution

$$P(k) \propto k^{-(1+\frac{1}{\alpha})}$$

Hence, for $\alpha \rightarrow 0$, the model yields a scale-free network with the exponent $\gamma \rightarrow \infty$, while for $\alpha \rightarrow 1$ the exponent γ converges to two. Hence, it provides a parameter that can be adjusted to effectuate random scale-free networks with an arbitrary degree distribution exponent γ in the range $(2, \infty)$.

In order to apply this simple model in practical networked systems, a distributed scheme is required that creates edges between two nodes i and j in an overlay network with probability $p_i p_j$. For this, we assume that we start with a random, connected overlay topology consisting of n nodes and m edges. In practice, this initial topology may emerge by means of an arbitrary bootstrapping method that connects joining nodes to existing participants either deterministically or at random. We can then view the above model as a rewiring scheme that gradually replaces existing edges so that edges between node pairs emerge with the desired node-dependent probabilities. For this, a node initiating the rewiring of an edge must be able to sample two new endpoints for the edge being rewired according to the probability measure given in equation 1. While one can imagine different mechanisms by which this can be achieved [23], a simple and well-established method to sample nodes in Peer-to-Peer systems is by means of random walks [17], [35]. For this we consider that nodes wishing to rewire an edge sample two new endpoints by means of two independent random walks through the current network topology. For a classical, unbiased random walk, the probability $\pi_i(t)$ to find the walker at an arbitrary time t at node i converges to

$$\pi_i(t) \rightarrow \frac{d_i}{N \cdot \bar{d}} \quad (t \rightarrow \infty)$$

where \bar{d} is the average node degree of the network. In order to sample nodes with the probabilities given in equation 1 we need to introduce a random walk bias that influences the transition probabilities accordingly. Considering a random walk in a connected overlay topology $G(V, E)$ as Markov chain with state space V and stationary distribution π , the random walk bias can be configured by means of a Metropolis-Hastings chain in such a way that a desired stationary distribution π holds [5], [20], [28]. In general, this can be achieved by introducing a bias as shown in the following transition matrix T :

$$T_{ij} = \begin{cases} \frac{1}{d_i} \min \left\{ \frac{\pi_j d_i}{\pi_i d_j}, 1 \right\} & (i, j) \in E, i \neq j \\ 1 - \frac{1}{d_i} \sum_{(k,i) \in E} P_{ik} & i = j \\ 0 & (i, j) \notin E \end{cases} \quad (2)$$

Here d_i denotes the current degree of node $i \in V$ and an entry T_{ij} gives the probability that a random walk residing at node i moves to node j . The fact that this transition matrix has stationary distribution π follows from the reversibility of the underlying Markov chain, as well as from its irreducibility (assuming a connected network topology) and aperiodicity (self-loops are possible). Under these restrictions, the Markov chain convergence theorem ensures that the probability $\pi_i(l)$ to find a random walker that has been started in an arbitrary node resides at node i after l steps converges to π as l goes to infinity.

From this, one can easily configure a random walk bias that results in a stationary distribution suitable to sample nodes in a way that - after rewiring - a scale-free network with degree distribution exponent γ emerges. From the probability p_i in equation 1 and the fact that it gives rise to a scale-free network with degree distribution exponent $1 + \frac{1}{\alpha}$, we obtain the desired stationary distribution

$$\pi_i^\gamma = \frac{i^{-\frac{1}{\gamma-1}}}{\sum_{k=1}^n k^{-\frac{1}{\gamma-1}}} \quad (3)$$

which, with equation 2 and $\frac{\pi_j}{\pi_i} = \left(\frac{i}{j}\right)^{\frac{1}{\gamma-1}}$, yields the following transition matrix P :

$$P_{ij} = \begin{cases} \frac{1}{d_i} \min \left\{ \left(\frac{i}{j}\right)^{\frac{1}{\gamma-1}} \frac{d_i}{d_j}, 1 \right\} & (i, j) \in E, i \neq j \\ 1 - \frac{1}{d_i} \sum_{(k,i) \in E} P_{ik} & i = j \\ 0 & (i, j) \notin E \end{cases} \quad (4)$$

Thus, a random walk with the above bias can be used to sample endpoints of edges and thus perform rewiring operations that effectuate random scale-free network topologies with a particular degree distribution exponent.

A. Bounding the Random Walk Length

The goal of this article is to study the practical applicability of the above strategy in a distributed rewiring scheme. Hence, an important question that needs to be answered is how many steps a random walk with the above bias needs to take before the probability $\pi_i(l)$ to find it in a node i after l steps is sufficiently close to the desired stationary limit π_i . In the rewiring protocol presented in the following section, this translates to the number of messages that need to be exchanged for a single rewiring operation. To assess this convergence behavior analytically, one first needs to give a formal definition of when two probability distributions π and π' shall be considered *sufficiently close*. For this we use

the usual definition of the *total variation distance* D which - for two probability measures π and π' and a finite state space V - is defined as follows:

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

The configuration of the random walk bias according to equation 4 and the Markov convergence theorem ensure that $D(\pi(l), \pi) \rightarrow 0$ for $l \rightarrow \infty$. For an arbitrarily chosen total variation distance $\epsilon > 0$ we can then assess the number of steps l our random walk needs to take until $D(\pi(l), \pi) \leq \epsilon$. In order to bound the minimally required number of steps l , the arguments put forth in [32] can be used. Here it is argued that an upper bound for l is given by

$$l \leq \frac{\ln\left(\frac{1}{\pi_s \epsilon}\right)}{1 - |\lambda_2(P)|}$$

where π is the stationary distribution of the Markov chain, $\lambda_2(P)$ is the second smallest eigenvalue of the transition matrix P and s is the initial state. Thus, finding an upper bound for the number of random walk steps requires to find a lower bound for the second smallest eigenvalue $\lambda_2(P)$ of the transition matrix. Unfortunately, obtaining good bounds for the eigenvalues of stochastic matrices is a non-trivial task. Nevertheless, based on the canonical path approach introduced in [12], [32], analytical arguments concerning the convergence behavior of random walks with a Zipf stationary distribution have been put forth in [35], [36]. In the following we briefly repeat these arguments for the particular random walk strategy considered in this article. In [36] it has been argued that, if the stationary distribution π is highly skewed, a lower bound for the eigenvalue gap $1 - |\lambda_2(P)|$ is given by

$$1 - |\lambda_2(P)| \geq \frac{\pi_{min}}{D \cdot d_{max}}$$

Here D denotes the diameter of the network topology upon which the random walk operates, π_{min} is the minimum probability ascribed to any vertex by the stationary distribution and d_{max} is the maximum degree of any vertex in the network. Thus, for the special case of Zipf stationary distributions, an asymptotic upper bound for the random walk length l required to achieve a total variation distance smaller than ϵ is given as [32], [36]:

$$l \leq \ln\left(\frac{1}{\pi_s \epsilon}\right) \cdot \frac{D \cdot d_{max}}{\pi_{min}} \quad (5)$$

For a random walk strategy configured to eventually effectuate a degree distribution exponent γ and thus stationary distribution π^γ , for the inverse stationary probability of the starting node s , the following bound holds:

$$\frac{1}{\pi_s^\gamma} = s^{\frac{1}{\gamma-1}} \cdot \sum_{k=1}^n k^{\frac{-1}{\gamma-1}} \leq s^{\frac{1}{\gamma-1}} \cdot \sum_{k=1}^n 1 = n \cdot s^{\frac{1}{\gamma-1}}$$

While this holds for arbitrary $\gamma \in [2, \infty)$ and starting nodes s , for the special case of node n we can give a better bound by observing that - due to the increasing skewness - node n is ascribed minimal probability for $\gamma = 2$, that is for $\gamma \in [2, \infty)$

$$\pi_{min}^\gamma \geq \pi_{min}^{\gamma=2}$$

holds. With this, we can bound the inverse minimal probability by considering the logarithmic growth of the harmonic series, so that

$$\frac{1}{\pi_{min}^\gamma} \leq \frac{1}{\pi_{min}^{\gamma=2}} = n \cdot \sum_{k=1}^n \frac{1}{k} = n \cdot H_n = n \cdot (\ln(n) + \tau + r_n)$$

where τ denotes the Euler-Mascheroni constant and $r_n \rightarrow 0$ in the limit of large n . Assuming an initial scale-free topology with n nodes and degree distribution exponent γ_i allows to asymptotically bound diameter and maximum degree as $O(\ln(n))$ and $O(n^{\frac{1}{\gamma_i}})$ respectively [36]. Thus, for large n and a random walk started in node s , an asymptotic upper bound for the minimal length l to achieve total variation distance smaller than ϵ can be given as follows:

$$l = O\left(\ln\left(\frac{n \cdot s^{\frac{1}{\gamma-1}}}{\epsilon}\right) \cdot \ln(n)^2 \cdot n^{1+\frac{1}{\gamma_i}}\right) \quad (6)$$

This theoretic bound scales worse than linear with the network size n . However, as has previously been observed for example in [36], the underlying bounding technique is not necessarily tight, that is the actual convergence behavior of a random walk can be considerably better. Since at present obtaining tight upper bounds for the convergence of Markov chains in complex network topologies is an open research issue, in section III we present simulations that have been performed to derive practicable random walk lengths empirically. As will be argued later, the results of these simulations suggest that the adaptation scheme presented in this article can be practically implemented with reasonable random walk lengths. Although these results suggest that the analytical bounds shown above are not tight and thus uninformative with respect to the performance of the scheme in practice, they can nevertheless be used to study by which parameters the convergence behavior of a random walk is influenced. From equation 6 one can for example infer that the upper bound for the minimal random walk length will generally be higher when wanting to effectuate highly skewed scale-free networks with exponents close to two.

II. PROTOCOL DEFINITION

The arguments laid out in the previous section suggest a rewiring protocol that consists of the following three basic operations: (1) In periodic intervals, a node a selects an edge to a random neighbor b that has not yet been rewired. (2) A random walk with the bias presented in equation 4 is started to sample two nodes x and y with probabilities proportional to π_x and π_y respectively. (3) The edge (v, w) is replaced by the edge (x, y) and the latter is marked as having resulted from a rewiring operation. After all m edges of the overlay have been rewired, a scale-free overlay is obtained whose exponent depends on the particular choice of the random walk bias defined in equation 4. In the algorithms 1 - 4, we give a detailed algorithmic description of the protocol. In these algorithms, d_v denotes the degree of node v , i_v is the ID of node v and $self$ denotes the node at which the code is executed. We further assume that nodes have information about the IDs and the degrees of their nearest neighbors.

A. Rewiring Procedure

The detailed algorithm of the main program loop that is responsible for initiating random walks is shown in algorithm 1. Rewiring operations are initiated by nodes in regular intervals only for those edges that have not yet been rewired. By this means, at most m rewirings are performed where m is the number of edges in the initial random network topology. We further assume, that the timing of these intervals is controlled by a *delay* value, nodes sleeping for *delay* milliseconds between two rewiring cycles. While this is not necessary for the functioning of the protocol, setting an appropriate *delay* value allows to adjust the number of rewiring operations and thus message transfers taking place within a certain time interval. When a node with an unmarked edge wakes up, a rewiring operation is initiated. In order to prevent both endpoints of an edge to initiate rewiring operations for the same edge, rewirings are only started by the node with higher degree or - if the degrees are equal - by the node with the smaller ID. As we shall see later in section III, the particular choice of letting a rewiring be initiated by the better connected endpoint is important since it can improve the performance of the scheme. To find the endpoints of a new edge by which the previously unmarked edge shall be replaced, a node initiates a biased random walk through the overlay (lines 6 – 11). In order to retain connectedness and prevent nodes from being isolated we further assume that only edges from nodes with degree greater than 1 are rewired.

When a node v receives a random walk message, it needs to ensure that the message is forwarded with the bias given in equation 4. In algorithm 2, this is done in lines 14 – 21. Comparing the algorithm with the stochastic matrix P defined in equation 4, here we select a neighbor uniformly at random and draw a random value uniformly in the interval $[0, 1]$ that indicates whether the random walk

transitions along this edge or whether it stays in the current node. One can imagine different schemes by which the two endpoints v and w of the new edge (v, w) are sampled. The node initiating the rewiring could for example start one random walk for each endpoint of the new edge, collect the target nodes of both walks and connect them to each other. In order to simplify the implementation, in algorithms 1 and 2 we propose to sample both endpoints of the new edge in a single random walk of length $2l$, assuming that after l steps, the node at which the random walk currently resides is stored in the field *target* of the message being forwarded. By this means, all information related to a rewiring operation is stored in a single random walk message. Hence the node at which the random walk arrives after $2l$ steps has all information necessary to initiate the rewiring operation. For this, it creates a connection to the *target* node stored in the message while initiating the deletion of the edge between node a that has started the random walk and its neighbor b . As can be seen in algorithm 4 a disconnection requires - apart from removing the edge - no further action at the side of the node from which the edge is removed. As shown in algorithm 3, both endpoints of the newly created edge mutually mark each other in order to prevent it from being rewired again in future invocations of the protocol. We emphasize that this is to prevent unnecessary rewiring operations and thus message exchanges rather than being required for the functioning of the protocol.

Algorithm 1 Main Loop

```

1: loop
2:   Sleep(delay)
3:   if neighbors.Count > marked.Count then
4:      $n = \text{RandomUnmarkedNeighbor}()$ 
5:     if  $d_n > 1$  and  $d_{self} > 1$  and  $(d_{self} > d_n$  or  $(d_{self} = d_n$ 
6:       or  $i_{self} < i_n))$  then
7:         {Initiate random walk}
8:          $msg.Hops \leftarrow 0$ 
9:          $msg.a \leftarrow self$ 
10:         $msg.b \leftarrow n$ 
11:         $msg.target \leftarrow null$ 
12:        Send({walk, msg},  $n$ )
13:      end if
14:    end loop

```

Concluding the description of the proposed protocol, we consider the size and number of messages that need to be sent across the network. Sampling the two endpoints of the new edge requires at most $2l$ messages¹, where l is the number of steps taken by a single random walk to sample a node with a probability sufficiently close to the stationary distribution π . Once both endpoints of the new edge have been found, the rewiring requires two messages to disconnect nodes a and b and one message to connect to

¹At most $2l$ since self-loops are allowed to ensure aperiodicity of the underlying Markov chain. While a self-loop is a considered a *step* of the random walk, it does not entail a message exchange.

Algorithm 2 Node receives $\{walk, msg\}$

```
1:  $msg.Hops \leftarrow msg.Hops + 1$ 
2: if  $msg.Hops = l$  then
3:   {Store Endpoint}
4:    $msg.target \leftarrow self$ 
5: else if  $msg.Hops = 2l$  then
6:   {Rewire}
7:   if  $\neg neighbors.Contains(msg.target) \ \&\& \ msg.target \neq self$ 
   then
8:      $Send(\{disconnect, msg.a\}, msg.b)$ 
9:      $Send(\{disconnect, msg.b\}, msg.a)$ 
10:     $Send(\{connect, self\}, msg.target)$ 
11:     $Send(\{connect, msg.target\}, self)$ 
12:   end if
13: else
14:    $n \leftarrow self.RandomNeighbor$ 
15:   if  $random.Next() \leq \frac{d_{self}}{d_n} \left( \frac{i_{self}}{i_n} \right)^{\frac{1}{\gamma_a - 1}}$  then
16:     {Forward Random Walk}
17:      $Send(\{walk, msg\}, n)$ 
18:   else
19:     {Self-Loop}
20:      $Send(\{walk, msg\}, self)$ 
21:   end if
22: end if
```

Algorithm 3 Node receives $\{connect, y\}$

```
1:  $neighbors.Add(y)$ 
2:  $marked.Add(y)$ 
```

the node *target* that has been stored in the random walk message.

Since the IDs of the initial node, its neighbors and the intermediate *target*, as well as the current hop count need to be stored in the random walk message, the required number of bits for a message is logarithmic in the number n of nodes in the system. Thus, the number of bits that need to be transferred per rewiring operation is $O(l \cdot \log(n))$. Since exactly one rewiring operation is executed for each of the m edges in the overlay topology, the total number of bits that need to be transferred in order to create a scale-free topology with the desired exponent is $O(m \cdot l \cdot \log(n))$. We further require to store one additional bit per edge, indicating whether an edge has previously been rewired or not.

B. Join and Leave Procedure in Dynamic Equilibrium States

Above, we have considered a simple distributed protocol for the gradual rewiring of existing edges by means of the random walk sampling mechanism discussed theoretically in section I. So far, we have not discussed the impact of nodes dynamically joining and leaving the system. In the following, we present a simple protocol extension that is suitable to handle the dynamics of nodes in dynamic equilibrium states, that is in networks of roughly fixed size with nodes joining and leaving the system at roughly balanced rates.

In this case, a simple join procedure as shown in algorithms 5 and 6 is sufficient. Upon initialization, a node

Algorithm 4 Node receives $\{disconnect, b\}$

```
1:  $neighbors.Remove(b)$ 
```

joining the system assigns a node ID that is drawn uniformly at random from the range $[1, n]$ with n being the current size of the network. It then creates k overlay links to arbitrary bootstrap nodes. We further assume that these newly created edges are *unmarked* and thus subject to future rewiring operations according to the protocol presented above. Nodes leaving the system do not require any particular handling, except for the removal of stale links as shown in algorithm 7.

Algorithm 5 Node joins the system

```
1: {Draw random ID}
2:  $i_{self} \leftarrow Random(1, n)$ 
3: for  $i \leftarrow 0; i < k; i \leftarrow i + 1$  do
4:    $x \leftarrow ArbitraryBootstrapNode()$ 
5:   {Create link to bootstrap node}
6:    $Send(\{join, self\}, x)$ 
7:    $neighbors.Add(x)$ 
8: end for
```

Algorithm 6 Node receives $\{join, w\}$

```
1:  $neighbors.Add(w)$ 
```

Algorithm 7 Node finds that neighbor w left the system

```
1:  $neighbors.Remove(w)$ 
2: if  $marked.Contains(w)$  then
3:    $marked.Remove(w)$ 
4: end if
```

A simple yet instructive way of looking at the interplay between join and rewiring operations is in terms of how they change the disorder or entropy present in the overlay topology. The dynamics of nodes results in a gradual increase of disorder since links are created to arbitrary (bootstrap) nodes. Hence the probabilistic structure of the topology gradually fades and its entropy increases. In contrary, the rewiring procedure removes unmarked “high entropy” links and replaces them by connections according to a statistical ensemble of random scale-free networks. In this respect, the stochastic rewiring maintains the topology’s probabilistic structure and thus decreases the network’s entropy. In order to maintain a random scale-free network topology, both the loss of structure that is due to node fluctuations as well as the maintenance of probabilistic structure in terms of rewiring operations need to occur at balanced rates.

III. EVALUATION

Having given a description of the rewiring protocol as well as analytical arguments about its convergence behavior, in this section we present simulation results that have been

obtained using an implementation of the proposed scheme. This evaluation is split up in two parts. In a first step, we seek to establish, by simulation, a practicable lower bound for the minimally required random walk length l . We further study the influence of the initiating node's degree on the convergence time of a random walk. Based on these results, in a second step we then simulate the rewiring protocol and study its influence on a network's degree distribution. Apart from considering static topologies, here we also present preliminary simulation results that have been obtained for network topologies in dynamic equilibrium states with fluctuating nodes.

A. Minimum Random Walk Length

While theoretic asymptotic upper bounds for the required number of steps l of the random walk have been presented in section I-A, here we empirically study the convergence behavior for a number of random walk lengths. By this we intend to derive a practicable random walk length that represents a reasonable trade-off between the imposed number of messages and the resulting total variation distance. We further intend to investigate how the minimum random walk length changes as the network sizes is varied. The following results have been obtained as follows. In each simulation run a number R of random walks was started from a randomly chosen node in a random network topology. For each simulation of a random walk of length l , a hit counter was increased in the node at which the random walk resided in the l -th step. When R random walks had been simulated, the total variation distance was computed based on the observed hit frequencies and the stationary distribution expected for the chosen random walk bias. Depending on network size and the minimum probability π_{min} of the stationary distribution, the number of random walk iterations R was chosen in a range between 10^6 and 10^8 . In particular, it was chosen such that nodes with minimum stationary probability π_{min} were expected to be hit reasonably often to argue about the total variation distance. The above procedure was then repeated for different random network realizations and starting nodes. Finally the minimum, maximum and average total variation distance of a simulation run was computed and the procedure was repeated for different network sizes, random walk biases and random walk lengths.

Figure 1a shows the random walk length l minimally required for the average total variation distance to fall below $\epsilon = 0.05$ in scale-free Barabási-Albert (BA) networks randomly generated by the preferential attachment scheme presented in [6]. Results are shown for different network sizes and for random walks configured to effectuate three different exponents 2.1, 2.5 and 3.5. Rather than the linear scaling behavior suggested by the theoretical upper bound presented in section I-A, the observed required length l rather scales in a sub-linear fashion. The observation that the actual convergence behavior is significantly better than

the theoretical bound that can be obtained by a canonical path approach is consistent with the observations presented in [34] and indicates that the rewiring scheme can be implemented efficiently in practice. Further simulation results that have been obtained for Erdős/Rényi (ER) random graphs indicate that the number of steps required to achieve a total variation distance smaller than 0.05 are in the same range as those for random power law graphs. Informally, this observed fast convergence can be attributed to the good expansion properties and the small diameter of both classical random graphs and random scale-free networks.

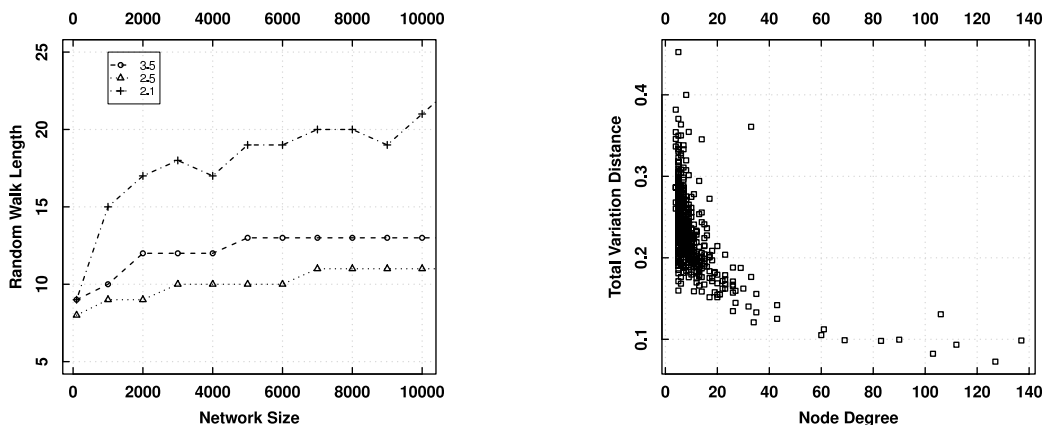
In networks with highly heterogeneous connectivity, a further interesting question is how the choice of the starting node of a random walk influences the total variation distance that can be achieved by a fixed random walk length. To investigate this, a number of random scale-free BA networks was created and a large number of random walks was started from each node of the network². The frequency with which nodes were target of random walks was recorded and the resulting total variation distance to the expected stationary distribution was computed for each starting node individually. Figure 1b shows the relation between the degree of initial nodes and the total variation distance that was achieved in a representative simulation in random Barabási-Albert networks with 1000 nodes, $l = 5$ and a random walk bias to effectuate $\gamma = 3$. These results suggest that random walks provide on average better convergence behavior when being started in highly connected nodes. In fact this is a rather intuitive result since a random walk starting at a high degree node can potentially reach a large number of nodes even in a single step. In the protocol presented in II, this observation justifies the choice that rewiring operations for an edge are initiated by the node with higher degree.

B. Degree Distribution

We now turn to the question of how the rewiring protocol described in section II influences the degree distribution of a network topology. All results presented in the following figures have been obtained for networks consisting of 5000 nodes and roughly 25000 edges. Initial topologies upon which the protocol was started were created using the Barabási-Albert preferential attachment model as well as the Erdős/Rényi model for classical random graphs. For the following measurements we used a random walk length that was long enough to achieve an average total variation distance smaller than $\epsilon = 0.05$. Based on the results presented in the previous section a random walk length $l = 20$ was chosen.

In each simulation run, the protocol presented in algorithm II was applied by all nodes in a network topology with initially unmarked edges until all edges had been rewired

²Here *large* again means sufficiently large to reasonably compute the total variation distance.



(a) Minimum random walk length l required to achieve $D(\pi(l), \pi) \leq 0.05$ in Barabási-Albert networks with random walk biases configured to effectuate exponents 2.1, 2.5 and 3.5 (Lines are drawn to guide the eye)

(b) Correlation between degree of starting node and average $D(\pi(l), \pi)$ in 1000 node Barabási-Albert networks with $\gamma = 3$ and $l = 5$

Figure 1. Results of random walk convergence experiments

(and thus marked). The *delay* interval between individual rewiring iterations was chosen such that - on average - a single rewiring took place per time unit. However, this choice has been made merely for illustration purposes since it allows to trace the evolution of a network topology as links are rewired progressively. Clearly, performing more than one per rewiring at a time would be more appropriate in actual applications of the protocol. When using a *delay* value so that one rewiring is expected to take place per unit of simulated time, for the chosen network size an adaptation cycle is expected to be completed after roughly 25000 time steps. The degree distribution of the network topology was computed each 200 time units and a fit to the current degree distribution exponent was performed. For this, an *R* implementation of the maximum likelihood power law fit procedure described in [10] was used. This procedure yields the fitted degree distribution exponent γ_f that holds with maximum likelihood, the minimum network degree d_{min} above which the fit holds, as well as the Kolmogorov-Smirnov (KS) statistic D . In general, better fits result in smaller values of D , thus allowing to evaluate whether the “power law nature” of the degree distribution is strengthened or fades away under the application of the rewiring scheme. All results are averages of at least 5 independent applications of our protocol on randomly chosen network realizations of identical size. Simulation code, data analysis scripts, datasets, simulation videos as well as some further graphical representations of results that could not be included in this article are available on the author’s website³.

³see <http://syssoft.uni-trier.de/scholtes>

1) *Static Topologies*: In the following, we consider a simple static situation in which no nodes enter or leave the initial network topology. Figures 3a and 3d show the effect of the proposed protocol on the degree distribution of a network that was initially created by the BA, respectively ER model. For BA networks, the average fitted exponent γ_f of the initial topology was on average 2.9, while for ER networks the used fitting procedure yielded 3.5 with an at least 10-fold value of the KS-statistic D . The results suggest that the protocol does lead to an adaptation of the degree distribution exponent of the overlay topology. In particular, the evolution of the Kolmogorov-Smirnov statistic D that is shown in Figure 3b demonstrate that the scale-free characteristic of BA networks is preserved. The increase of the minimum degree above which the fit holds can be explained by the exponent-dependent finite-size effects in scale-free networks. For ER networks, the roughly 10-fold decrease of the KD-statistic D that can be seen in Figure 3e indicates the emergence of scale-free characteristics, that is the power law fit to the degree distribution becomes more reliable. In Figures 3c and 3f, the evolution of the average maximum degree is shown. The results are consistent with the maximum degree expected in networks of the given size and with different degree distribution exponents. In Figure 2, the average fit parameters for the network topology eventually reached after adaptation are shown. The results demonstrate that - as expected from the underlying theoretical model - the protocol can be applied to transform arbitrary initial topologies into random scale-free networks whose degree distribution is described by a power law with an exponent reasonably close to the intended value.

So far, we have only studied simulations using a single

	γ_t	2.1	2.3	2.5	2.7	2.9	3.1	3.3	3.5
BA	γ_f	2.24	2.40	2.60	2.82	2.99	3.24	3.44	3.5
	D	0.012	0.01	0.01	0.01	0.01	0.01	0.012	0.02
	d_{min}	6.6	8.4	8.6	10	10	11.6	13.2	11.6
ER	γ_f	2.252	2.41	2.61	2.80	3.03	3.25	3.45	3.5
	D	0.013	0.01	0.01	0.009	0.009	0.01	0.012	0.02
	d_{min}	7.6	8.8	9.2	9.6	11.4	12.6	12.2	12.4

Figure 2. Average fit parameters after adaptation with targeted exponents $\gamma_t \in [2.1, 3.5]$ for 5000 node Erdős/Rényi (ER) and Barabási/Albert (BA) networks with roughly 25000 edges

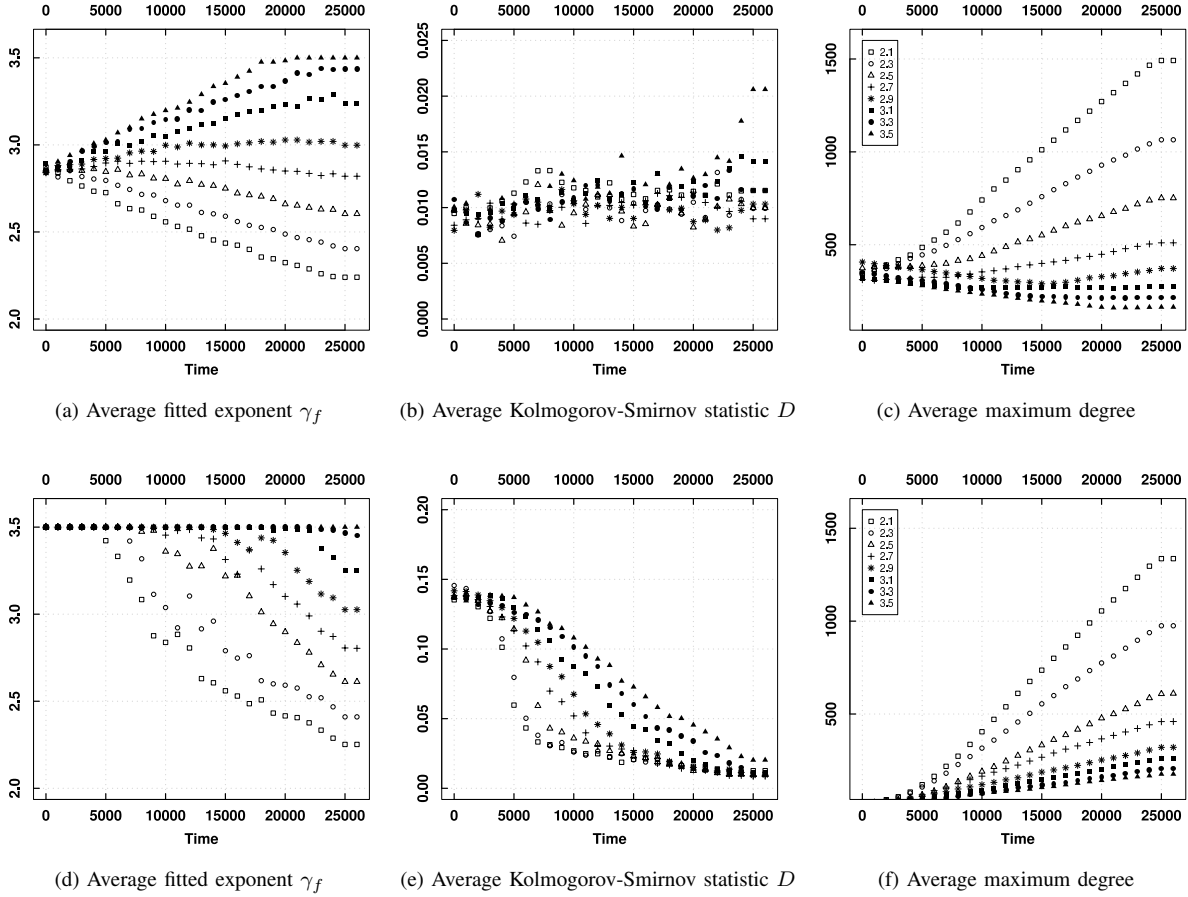


Figure 3. Time Evolution of 5000 node Barabási/Albert (a-c) and Erdős/Rényi (d-f) networks during adaptation runs with $\gamma_t \in [2.1, 3.5]$ and $l = 20$

“cycle” of the proposed adaptation protocol. In Figure 4, results are shown for a simulation in which three adaptation cycles targeting different exponents were subsequently initiated in a BA network with 10^4 nodes and roughly $5 \cdot 10^4$ edges. The chosen random walk length of $l = 22$ was again consistent with the values found in section III-A. In Figure 4a-4c, time steps in which adaptation cycles were started are indicated by vertical lines. The targeted degree distribution exponents were 2.9, 2.1 and 3.5 respectively. Again the results indicate that the proposed scheme achieves the desired adaptation. Furthermore, Figure 4b shows how the “power law nature” of the degree distribution - and thus the scale-free characteristic of the network - temporarily

fades during the adaptation while being restored near the ends of adaptation cycles.

2) *Dynamic Equilibrium Topologies:* So far, only results for static topologies have been presented, that is the rewiring protocol executing in different network topologies with no nodes entering or leaving the topology. In the following, we present some preliminary results for dynamic topologies in which nodes join and leave the overlay uniformly at random at balanced rates, thus forming a dynamic equilibrium state for a system with roughly fixed size. In the following experiments, a BA network with 5000 nodes and roughly 25000 edges has been used as initial topology. However results for ER networks have shown to be identical. In each

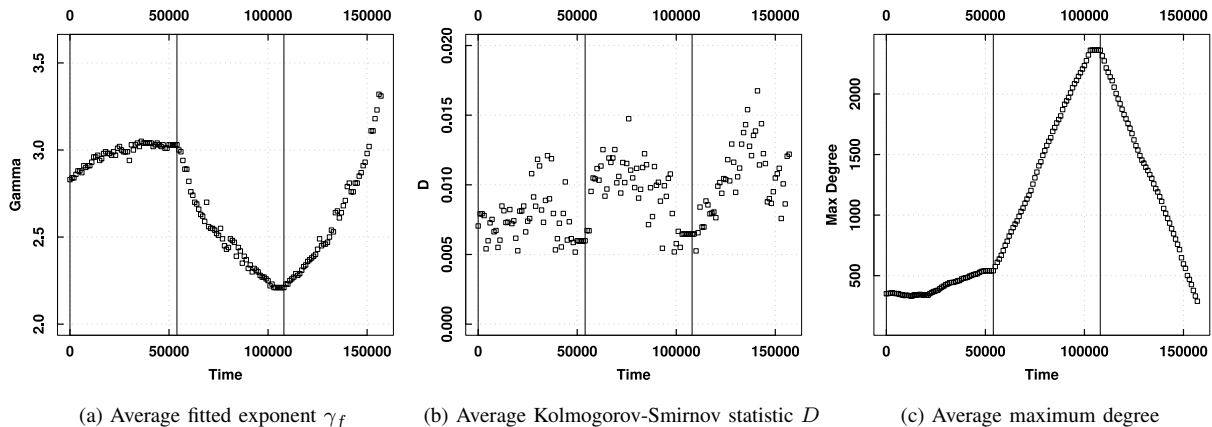


Figure 4. Time Evolution of 10000 node Barabási/Albert network during multiple adaptation cycles with $\gamma_0 = 2.9$, $\gamma_{54000} = 2.1$ and $\gamma_{108000} = 3.5$. Start times of adaptation cycles are indicated by vertical lines.

simulated time step, a single node was removed from while a single other node joined the network. Nodes joining the network were connected to $k = 5$ random bootstrap nodes according to the algorithm depicted in algorithm 5. Stale links left by failing nodes were removed immediately from the network. The delay parameter has been set to a value that was sufficiently small to compensate for the constant creation of unmarked links by joining nodes. Hence, the average number of links being rewired (and thus being marked) within a certain time interval was roughly equal to the average number of random (unmarked) edges created in that same time interval by joining nodes. Figure 5 shows the averaged results of simulations targeting degree distribution exponents $\gamma_t \in [2.1, 3.5]$. In particular, Figure 5a suggests that the degree distribution exponent quickly approaches the targeted value as node fluctuations drive the gradual rewiring of links.

IV. RELATED WORK

During the last couple of years, a number of distributed approaches to the construction, maintenance and adaptation of probabilistically structured overlay topologies have been proposed. Here we briefly summarize a selection of approaches that are related to the present article. The use of random walks for the sampling of random participants in Peer-to-Peer overlays with good expansion (and thus Markov chain convergence) properties has been proposed in [17] and [35]. In particular, in [35] the use of biased random walks for a non-uniform random sampling of Peers is studied and analytical arguments for their convergence behavior are given. As argued in section I, a similar random walk strategy constitutes the foundation for the adaptation scheme presented in this article. We finally emphasize that different approaches to a random sampling in Peer-to-Peer networks have been proposed as well, like for example the gossip-based topology management scheme considered in [22]. To

date, it is however unclear how such alternative sampling mechanisms could be used in our particular scenario.

Considering the problem of creating and adapting overlay networks with scale-free characteristics in a distributed fashion, it has been argued for example in [24] that the degree distribution exponent of scale-free networks can be tuned by adjusting the connection preferences of joining nodes. While this constitutes the basis for an adaptation of growing networks, it remains unclear how the existing theoretical models can be implemented efficiently in practice and in how far theoretical results on random scale-free networks apply to the networks created by this process. Considering practical networked systems, in [19] distributed strategies for the creation of scale-free overlays with connectivity cutoffs based on capacity constraints have been considered. Since the adaptation of the degree distribution exponent does also change the maximum degree in the network, the schemes presented in [19] - although different in nature and intention - can be viewed as being related to the scheme presented in the present article. Finally, the problem of adapting the degree distribution exponent in scale-free overlay networks has been considered in our previous work [30]. However, in contrast to the protocol presented in the present article, no analytical arguments for the functioning of the scheme considered in [30] as well as its precise effects on the degree distribution exponent as well as on other properties of the network could be given. As such, the protocol can be viewed as a companion scheme to the distributed power law monitoring mechanism presented in [30].

V. CONCLUSION

In this article, a simple distributed protocol has been presented that can be used to effectuate random scale-free overlay topologies with tunable degree distribution exponent and thus a tunable heterogeneity of the connectivity distribution. The proposed scheme is based on a rewiring strategy




Figure 5. Time evolution of dynamic 5000 node Barabási/Albert networks with random and uniform churn, $\gamma_t \in [2.1, 3.5]$ and $l = 20$

and the sampling of random nodes by means of a biased random walk. We have shown that the proposed protocol is suitable to transform arbitrary connected topologies into sufficiently random scale-free networks given that the expansion properties of the initial topology provide sufficiently fast convergence of the random walk strategy. In Barabási-Albert and Erdős/Rényi networks, the random walk length required to provide sampling probabilities that are acceptably close to the stationary limit are found to be significantly smaller than suggested by theoretical upper bounds. Based on the observed superior convergence behavior of random walks being started in hub nodes, the performance of the protocol benefits from the fact that rewiring operations are preferentially started by high degree nodes. In a future iteration, it thus seems to be reasonable to choose the length l of each random walk individually based on the degree of the node initiating it. A further potential improvement is the use of two-stage random walks which - in a first stage - preferentially move to hub nodes, and then - in a second stage - switch their bias to sample nodes according to the desired stationary distribution.

Apart from creating randomized scale-free overlay topologies, the proposed protocol can also be used to adapt the degree distribution exponent of scale-free networks in a decentralized and directed fashion. Based on analogies to statistical physics that have been put forth in the study of complex networks [2], [15] and the fact that - in the limit of large systems - a number of important network properties change abruptly when the degree distribution exponent exceeds certain critical points, such a mechanism can be used to actively use *phase transition phenomena* occurring in statistical ensembles of complex networks for an adaptation of network characteristics [30]. Based on the results presented in this article, we thus argue that the proposed protocol is a simple and practicable approach to create and adapt *probabilistically structured overlays* for

large-scale Peer-to-Peer systems.

We conclude this article by summarizing the main threats to validity and open issues that could not be considered so far. An important aspect in any practical application of the proposed scheme is the fact that a sufficiently efficient implementation of the scheme requires to accept moderate total variation distances. While this allows to keep the message overhead in an acceptable range, it also limits the randomness of the resulting network topology. While small total variation distances suggest that the resulting deviation of properties from those of truly random networks are rather moderate, a further investigation of these effects is an open issue. Furthermore, although we have argued that simulations are a reasonable approach to establish empirical bounds on the required random walk length, the range of network sizes and topologies considered so far is fairly limited. A study of further topologies must thus be considered future work.

While some preliminary results on network topologies with dynamic participants have been presented, a major open issue of this work is the fact that simulations have only been performed for the rather special situation of dynamic equilibrium states with nodes failing uniformly at random. An evaluation of the impact of further realistic (non-uniform) churn models and rates is thus crucial. Considering an application in practical systems, a further potential problem that could not be considered so far is the impact of message losses and how to properly handle them in the protocol.

Finally, the rather simple implementation considered in this article assumes per-node weights that are based on values chosen uniformly at random from a fixed interval. As a result, the number of links that any given node will eventually acquire is determined by mere chance. Clearly, this is not suitable for practical systems and an important next step is to extend the proposed scheme in a way that per-node characteristics like bandwidth capacity, uptime

and stability or processing power are considered in the random assignment of node weights. In the light of these limitations, the scheme proposed in this article must be seen as a mere first step on the road towards stochastic network membership protocols whose self-organization and self-adaptation capabilities can be analyzed exactly in the framework of complex network theory.

ACKNOWLEDGMENT

The author would like to thank Jean Botev and Markus Esch (both University of Luxembourg), Katharina A. Zweig (University of Heidelberg, Germany) as well as the anonymous reviewers for helpful suggestions.

REFERENCES

- [1] L. A. Adamic, R. M. Lukose, A. R. Puniyani, and B. A. Huberman. Search in power-law networks. *Physical review E, Statistical, nonlinear, and soft matter physics*, 64, 2001.
- [2] R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 72(1):48–97, 2002.
- [3] R. Albert, H. Jeong, and A.-L. Barabasi. The diameter of the world wide web. *Nature*, 401:130, 1999.
- [4] R. Albert, H. Jeong, and A.-L. Barabasi. Error and attack tolerance of complex networks. *Nature*, 406:378–382, 2000.
- [5] Y. Azar, A. Z. Broder, A. R. Karlin, N. Linial, and S. Phillips. Biased random walks. In *STOC '92: Proceedings of the twenty-fourth annual ACM symposium on Theory of computing*, pages 1–9, New York, NY, USA, 1992. ACM.
- [6] A.-L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [7] A. Barrat, M. Barthelemy, and A. Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, 2008.
- [8] B. Bollobás and I. Riordan. The diameter of a scale-free random graph. *Combinatorica*, 24:5–34, 2004.
- [9] P. Boykin, V. Roychowdhury, and N. Sarshar. Percolation search in power law networks: making unstructured peer-to-peer networks scalable. *Proceedings of the Fourth International Conference on Peer-to-Peer Computing*, 2004.
- [10] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Review*, 51:661, 2007.
- [11] R. Cohen and S. Havlin. Scale-free networks are ultrasmall. *Phys. Rev. Lett.*, 90(5):058701, Feb 2003.
- [12] P. Diaconis and D. Stroock. Geometric bounds for eigenvalues of markov chains. *Ann. Appl. Prob.*, 1:36–61, 1991.
- [13] S. N. Dorogovtsev and J. F. F. Mendes. Evolution of networks with aging of sites. *Phys. Rev. E*, 62(2):1842–1845, Aug 2000.
- [14] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pages 251–262, New York, NY, USA, 1999. ACM Press.
- [15] I. Farkas, I. Derenyi, G. Palla, and T. Vicsek. Equilibrium statistical mechanics of network structures. *LECT.NOTES PHYS.*, 650:163, 2004.
- [16] C. Gkantsidis, M. Mihail, and A. Saberi. Conductance and congestion in power law graphs. *SIGMETRICS Perform. Eval. Rev.*, 31(1):148–159, 2003.
- [17] C. Gkantsidis, M. Mihail, and A. Saberi. Random walks in peer-to-peer networks: algorithms and evaluation. *Perform. Eval.*, 63(3):241–263, 2006.
- [18] K. I. Goh, B. Kahng, and D. Kim. Universal behavior of load distribution in scale-free networks. *Physical Review Letters*, 87:278701, 2001.
- [19] H. Guclu, D. Kumari, and M. Yuksel. Ad-hoc limited scale-free models for unstructured peer-to-peer networks. *CoRR*, abs/0806.2395, 2008.
- [20] W. K. Hastings. Monte carlo sampling methods using markov chains and their applications. *Biometrika*, 57:97–109, 1970.
- [21] M. Jelasity, A. Montresor, and O. Babaoglu. Gossip-based aggregation in large dynamic networks. *ACM Trans. Comput. Syst.*, 23(3):219–252, 2005.
- [22] M. Jelasity, A. Montresor, and O. Babaoglu. T-man: Gossip-based fast overlay topology construction. *Comput. Netw.*, 53(13):2321–2339, 2009.
- [23] M. Jelasity, S. Voulgaris, R. Guerraoui, A.-M. Kermarrec, and M. van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3):8, 2007.
- [24] P. L. Krapivsky, S. Redner, and F. Leyvraz. Connectivity of Growing Random Networks. *Physical Review Letters*, 85:4629–4632, Nov. 2000.
- [25] D.-S. Lee, K.-I. Goh, B. Kahng, and D. Kim. Evolution of scale-free random graphs: Potts model formulation. *Nuclear Physics B*, 696:351–380, Sept. 2004.
- [26] Q. Lv, P. Cao, E. Cohen, K. Li, and S. Shenker. Search and replication in unstructured peer-to-peer networks. In *Proceedings of the 16th international conference on Supercomputing (ICS 2002)*, pages 84–95, New York, NY, USA, 2002. ACM.
- [27] A. Medina, I. Matta, and J. Byers. On the origin of power laws in internet topologies. *SIGCOMM Comput. Commun. Rev.*, 30(2):18–28, 2000.
- [28] N. Metropolis, A. W. Rosenbluth, A. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [29] M. Ripeanu and I. Foster. *Revised Papers From the First international Workshop on Peer-To-Peer Systems*, chapter Mapping the Gnutella Network: Macroscopic Properties of Large-Scale Peer-to-Peer Systems, pages 85–93. Springer, 2002.

- [30] I. Scholtes, J. Botev, M. Esch, A. Hoehfeld, and H. Schloss. Awareness-driven phase transitions in very large scale distributed systems. In *Proceedings of the Second IEEE International Conferences on Self-Adaptive and Self-Organizing Systems (SaSo), Venice, Italy*. IEEE, October 2008.
- [31] I. Scholtes, J. Botev, M. Esch, and P. Sturm. Epidemic self-synchronization in complex networks of kuramoto oscillators. *Advances in Complex Systems*, 13:33–58, 2010.
- [32] A. Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability and Computing*, 1:351–370, 1992.
- [33] Z. Toroczkai and K. E. Bassler. Jamming is limited in scale-free systems. *Nature*, 428:716, 2004.
- [34] H. Zhang, G. Neglia, D. F. Towsley, and G. L. Presti. Stability and efficiency of unstructured file sharing networks. *IEEE Journal on Selected Areas in Communications*, 26(7):1284–1294, 2008.
- [35] M. Zhong, K. Shen, and J. Seiferas. Non-uniform random membership management in peer-to-peer networks. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 2, pages 1151 – 1161 vol. 2, march 2005.
- [36] M. Zhong, K. Shen, and J. Seiferas. The convergence-guaranteed random walk and its applications in peer-to-peer networks. *IEEE Trans. Comput.*, 57(5):619–633, 2008.