

# Efficient Network Generation under General Preferential Attachment

James Atwood  
School of Computer Science  
University of Massachusetts  
Amherst  
Amherst, MA, 01003  
jatwood@cs.umass.edu

Bruno Ribeiro  
School of Computer Science  
Carnegie Mellon University  
Pittsburgh, PA, 15213  
ribeiro@cs.cmu.edu

Don Towsley  
School of Computer Science  
University of Massachusetts  
Amherst  
Amherst, MA, 01003  
towsley@cs.umass.edu

## ABSTRACT

Preferential attachment (PA) models of network structure are widely used due to their explanatory power and conceptual simplicity. PA models are able to account for the scale-free degree distribution observed in many real-world large networks through the remarkably simple mechanism of sequentially introducing nodes which attach preferentially to nodes with high degree. The ability to efficiently generate instances from PA models is a key asset in understanding both the models themselves and the real networks that they represent. Surprisingly, little attention has been paid to the problem of efficient instance generation. In this paper, we show that the complexity of generating network instances from a PA model depends on the preference function of the model, provide efficient data structures that work under any preference function, and present empirical results from an implementation based on these structures. We demonstrate that, by indexing growing networks with a simple augmented heap, we can implement a network generator which scales many orders of magnitude beyond existing capabilities ( $10^6 - 10^8$  nodes). We show the utility of an efficient and general PA network generator by investigating the consequences of varying the preference functions of an existing model. We also provide “quicknet”, a freely-available open-source implementation of the methods described in this work.

## 1. INTRODUCTION

There is a clear need for scalable network generators, as the ability to efficiently generate instances from models of network structure is central to understanding both the models and the real networks that they represent. Ideally, researchers should be able to generate networks on the same scale as real communication and social networks, many of which have millions to billions of nodes. Furthermore, network generation is the primary tool both for empirically validating the theoretical behavior of models of network structure and for investigating behaviors that are not captured

by theoretical results. The generation of very large networks is of particular importance for these tasks because theoretically derived behavior is often asymptotic.

However, the generation of large networks is difficult because of its high complexity. In the case of preferential attachment (PA), arguably the most widely used generative model of network evolution, a non-local distribution over node degrees must be both sampled from and updated at each timestep. If we naively index this distribution, we will need to update every node at every timestep, which implies that simulating a network will have complexity of at least  $O(|V|^2)$ .

PA models are of particular interest because they account for the scale-free distribution of degree that is observed in many large networks [15]. For instance, scale-free degree distributions have been observed in the World Wide Web [5, 6, 7], the Internet [9, 12, 20], and telephone call graphs [1, 2]. Scale-free degree distributions have also been observed in bibliographic networks [11] and social networks [18].

Preferential attachment models generate networks by sequentially introducing nodes which prefer to attach to nodes with high degree. While many extensions to this model class exist, all members share the same basic form: At each timestep, sample a node from the network with probability proportional to its degree; introduce a new node to the network; and add an edge from the new node to the sampled node. This behavior has important implications for implementation. First, PA is inherently sequential, because the next action taken depends on the state of the network, and the state of the network changes at each timestep. This implies that the algorithm is not easily parallelized. Second, network nodes must be indexed such that they can be efficiently sampled by degree, and, because we are introducing a new node at each timestep, the index must also support efficient insertion. Furthermore, at each timestep, the degrees of existing nodes can only be incremented, so the index can take advantage of the constraint that unnormalized node probabilities will only increase. Third, the relevant distribution over nodes is non-local, in that the introduction of a new node and edge affects the probability of every node in the network through the normalization factor.

Note that work in this area has typically focused on the asymptotic regime. A model is defined, and a limiting degree distribution (as  $|V|$  approaches infinity) is obtained analytically. Less effort has focused simulating the models to some finite horizon. In the following sections, we provide a robust framework for generating networks via PA. This

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

```

Price( $n, \lambda$ ):
   $G \leftarrow G_0$ 
  for  $i \leftarrow 1$  to  $n - |G_0|$  do
    existing_node  $\leftarrow$  sample_in_degree( $G, \lambda$ )
    new_node = add_node( $G$ )
    add_edge( $G, \text{new\_node}, \text{existing\_node}$ )
  end for
  return  $G$ 

Krapivsky( $n, p, \lambda, \mu$ ):
   $G \leftarrow G_0$ 
  while  $|V| < n$  do
     $u \leftarrow$  uniform draw
    if  $u < p$  then
      existing_node  $\leftarrow$  sample_in_degree( $G, \lambda$ )
      new_node = add_node( $G$ )
      add_edge( $G, \text{new\_node}, \text{existing\_node}$ )
    else
      existing_node_tail  $\leftarrow$  sample_in_degree( $G, \lambda$ )
      existing_node_head  $\leftarrow$  sample_out_degree( $G, \mu$ )
      add_edge( $G, \text{existing\_node\_tail}, \text{existing\_node\_head}$ )
    end if
  end while
  return  $G$ 

```

**Figure 1: Generating a network with  $n$  nodes under Price and Krapivsky’s models.**  $G_0$  is some small seed network.  $\lambda$  and  $\mu$  are scalars which give the fitness of nodes for incoming and outgoing edges, respectively.

framework is quite efficient and easily scales to many millions of nodes on commodity hardware. We also provide “quicknet”, a freely-available open-source C implementation of the framework<sup>1</sup>.

The remainder of the paper is structured as follows. In Section 2, we present an analysis of the complexity of generating networks from PA models. Section 3 describes candidate methods for efficiently implementing preferential attachment generators and presents results from a simulator which implements them. Section 4 describes several applications of a PA network generator which scales to many millions of nodes. We describe related work in Section 5 and present conclusions and future work in Sections 6 and 7, respectively.

## 2. COMPLEXITY

In this section, we describe the two PA models considered in this work. We then provide a more formal definition of a PA model, which is followed by an analysis of the complexity of generating networks from PA models.

### 2.1 Description of Considered Models

#### 2.1.1 Price’s Model

Figure 1 describes Price’s algorithm. Briefly, at each timestep, a node is sampled from the network with probability proportional to its in-degree; a new node is introduced to the network; and a directed edge is added from the new node to the sampled node. Notice that a node is added at each timestep, so that the generation of a network with  $|V|$  nodes takes  $|V|$  steps.

<sup>1</sup><https://github.com/hackscience/quicknet>

#### 2.1.2 Krapivsky’s Model

Figure 1 also describes the algorithm of Krapivsky et al. (which, for brevity, we will refer to as the “Krapivsky model”). At each step, the algorithm of Price’s model is followed with probability  $p$ , and a “preferential edge step” is taken with probability  $1 - p$ . During a preferential edge step two nodes,  $n_o$  and  $n_i$ , are sampled from the network by out- and in-degree, respectively, and an edge is added from  $n_o$  to  $n_i$ . Note that a node is no longer added at every step; rather, a node is added at a given step with probability  $p$ . This implies that the number of iterations required to generate a network with  $|V|$  nodes is a random variable with expected value  $|V|/p$ .  $|V|/p$  is  $\Theta(|V|) \forall p$ , so asymptotically this is no different than Price’s model. More generally, the number of iterations required to generate a network with  $|V|$  nodes via a PA model is  $\Theta(|V|)$ .

## 2.2 Definitions

Let  $G_t = (V_t, E_t)$  be the network that results from  $t$  iterations of a PA simulation.  $V_t$  is the set vertices (or nodes) within the network and  $E_t$  is the set of edges between elements of  $V_t$ . Let  $T(G_t)$  be the worst-case time complexity of generating  $G_t$ ; that is, the worst-case time complexity of a preferential attachment simulation of  $t$  iterations.

Recall that the number of iterations required to generate a network with  $|V|$  nodes via PA is  $\Theta(|V|)$ . Accordingly, we will omit  $t$  and frame our discussion of complexity  $T(G)$  in terms of  $|V|$ .

Let  $f(a, \lambda) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$  be a function which maps an attribute  $a \in \mathbb{R}$  of  $v \in V$  and  $v$ ’s fitness  $\lambda \in \mathbb{R}$  to a preference mass  $\mu \in \mathbb{R}^+$  for attribute  $a$  of  $v$ . The “preference mass” of an attribute  $a$  of node  $v$  is a non-negative value that is proportional to the probability of selecting  $v$  by  $a$  under the PA model. We will refer to  $f$  as a preference function of the PA model. Note that, in this work, we restrict our attention to the degree attribute  $d \in \mathbb{N}$ , which implies that  $f$  is defined over the natural numbers:  $f(d, \lambda) : \mathbb{N} \times \mathbb{R} \rightarrow \mathbb{R}^+$ . The restriction is purely elective; any real-valued attribute of nodes could be specified.

A PA model has one or more preference functions. Price’s model, for example, has a single linear preference function. Krapivsky’s model has two: one for in-degree and another for out-degree. A “linear preferential attachment model” only admits linear preference functions of the form  $f(d, \lambda) = c_1 d + \lambda$ , a “quadratic preferential attachment model” only admits quadratic preference functions of the form  $f(d, \lambda) = c_2 d^2 + c_1 d + \lambda$ , and so on.

### 2.3 Derivation of Generation Complexity

We obtain a trivial lower bound on  $T(G)$  by noting that, in order to generate  $G$ , we must at the very least output  $|V|$  nodes, so

$$T(G) = \Omega(|V|)$$

A discussion of the upper bound follows. Recall that the salient problem in generating networks from a PA model is indexing the network’s nodes in such a way that sampling, insertion, and incrementation can be accomplished efficiently. Tonelli et al. [19] provide a very clever method for accomplishing all three tasks in constant time, provided that the preference function is linear and the fitness is both uniform across all nodes and constant. Given constant insertion and sampling, the generation of a network with  $|V|$

```

Sample(tree):
  sampled_node ← NULL
  u ← uniform sample
  if tree.root != NULL then
    sampled_node ← SampleItem(tree, tree.root, 0., u)
  end if
  return sampled_node
SampleItem(item, η, u):
  if item.left != NULL then
    if u < (η + item.left.subtree_mass) / tree.total_mass
    then
      return SampleItem(tree, item.left, η, u)
    end if
    η ← η + item.left.subtree_mass
  end if
  η ← η + item.node_mass
  if u < observed_mass / tree.total_mass then
    return item.node
  end if
  if item.right != NULL then
    return SampleItem(tree, item.right, η, u)
  end if

```

**Figure 2: General algorithm to sample from the augmented tree structure.**  $\eta$  is the mass observed thus far and  $u$  is a sample from the standard uniform distribution.

nodes takes  $O(|V|)$  time. Considering that the lower bound is  $\Omega(|V|)$ , we have the asymptotically tight bound of

$$T(G) = \Theta(|V|)$$

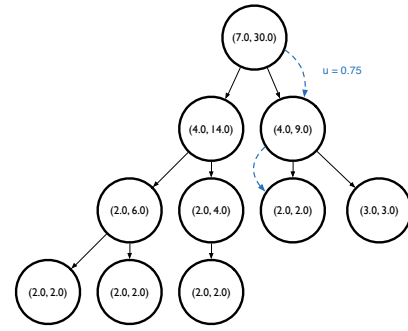
However, this method becomes very inefficient when extended to nonlinear preferential attachment. We can improve performance by shifting to data structures which provide  $O(\log|V|)$  insertion, sampling, and incrementation, giving an overall complexity of

$$T(G) = O(|V|\log|V|)$$

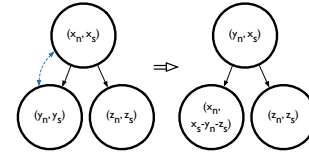
We accomplish this with a set of augmented tree structures. Each tree supports a preference function of the model by indexing the preference mass assigned to each node in the network by that preference function. Each item in the tree indexes a node in the network. The tree items are annotated with the preference mass of the network node under the preference function, and the subtree mass, which is the total preference mass of the subtree which has the item as root; see Figure 3. Note that we refer to “items” in the tree rather than the more typical “nodes”; this is to avoid confusion between elements of the tree and elements of the network. We annotate each tree with the sum of the preference masses of all items within the tree, that is, the normalization factor or partition function. We can sample from such a structure by recursively comparing the properly normalized subtree mass of a given item and its children to a uniform random draw; see Figures 2 and 3.

Note that, at each iteration of a standard PA simulation, we must sample a node, update that node’s mass, and insert a new node. In what follows we show that each of these steps can be accomplished in asymptotically logarithmic time.

### 3. IMPLEMENTATION



**Figure 3: An example of an augmented tree structure.** Each node is annotated with  $(\mu_n, \mu_s)$ , where  $\mu_n$  the node mass and  $\mu_s$  is the subtree mass. The preference function associated with this tree is  $f(d) = d + 2.0$ . The sample path through the tree structure is illustrated for  $u = 0.75$ .



**Figure 4: A diagram of the heap exchange process in the augmented heap.** Each node is annotated with  $(\mu_n, \mu_s)$ , where  $\mu_n$  the node mass and  $\mu_s$  is the subtree mass. Note that exchanges maintain the subtree mass invariant.

The tree structure that we described in the previous section can be implemented in a number of different ways that have a generation time of  $O(|V|\log|V|)$ . They differ in their computational time for finite  $|V|$ . In this section, we empirically evaluate a set realizations of the annotated tree structure. Specifically, we investigate a simple binary max-heap where priority is defined by node mass, and a set of binary treaps with various sort and priority keys.

Note that, in the discussion of the heap-based and treap-based implementations of the tree structure, we will often refer to a “sort invariant” and a “heap invariant”. The sort invariant states that, for any three nodes  $Y \leftarrow X \rightarrow Z$  where  $Y$  and  $Z$  are the left and right children of parent  $X$ , respectively, and a key  $k$  (sometimes referred to as the “sort key”) that is associated with each item,  $Y.k \leq X.k \leq Z.k$ . The heap invariant states that for any three nodes  $Y \leftarrow X \rightarrow Z$  (defined in the same fashion) and some priority  $p$  (sometimes referred to as the “priority key”) associated with each item,  $X.p \geq Y.p$  and  $X.p \geq Z.p$ .

We first describe the binary maximum heap. We annotate each item in the heap with a node mass, which is defined by the preference function, and a subtree mass, which is initialized to the node mass. When inserting a new item  $i$ ,  $i$ ’s node mass is added to all traversed items, so that the subtree mass is remains accurate upon insertion. Sampling is accomplished via the algorithm of Figure 2. Node mass may only increase, so we implement an augmented version of increase-key which maintains subtree mass under exchanges; see Figure 4 for a diagram of the exchange op-

Increment(heap,item,mass\_function):

```

new_mass ← mass_function(item.mass + 1)
additional_mass ← new_mass - item.mass
item.mass ← new_mass
item.subtree_mass += additional_mass
while item != heap.root && parent(item).mass <
item.mass do
    parent(item).subtree_mass += additional_mass
    heap_exchange(heap, item, parent(item))
    item ← parent(item);
end while
while item != heap.root do
    parent(item).subtree_mass += additional_mass
    item ← parent(item);
end while

```

**Figure 5: The Increment operation of the heap-based tree structure. The constant-time operation heap\_exchange is demonstrated in Figure 4. The two while loops collectively over an item’s  $O(\log|V|)$  ancestors, so Increment is a  $O(\log|V|)$  operation.**

Insert(heap, item):

```

heap.add(item)
node_mass ← item.node_mass
while item has a parent do
    item ← parent(item)
    item.subtree_mass += node_mass
end while

```

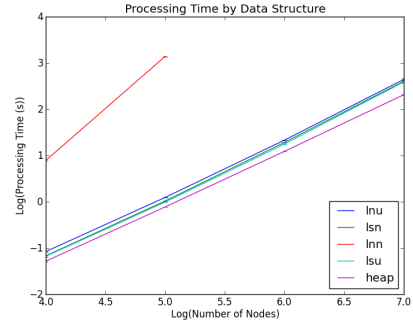
**Figure 6: The Insert operation of the heap-based tree structure. Note that each item has  $O(\log|V|)$  ancestors, so Insert is a  $O(\log|V|)$  operation.**

eration. The increase-key operation supports the Increment operation, which is described below. We set priorities to be equivalent to node masses so that the most probable nodes can be accessed more quickly.

The PA process is supported by the binary maximum heap via the operations Sample, Increment and Insert. As previously mentioned, sampling is performed via the algorithm of Figure 2. Increment increases an item’s mass and then performs heap exchanges to account for any violation of the heap invariant; it is described in Figure 5. Insert adds an item to the index, appropriately updating the subtree masses of any parent items; see Figure 6.

Note that, if we were to annotate each item with a node’s *probability* mass rather than *preference* mass, insertion would be a linear time operation. When a new node is introduced, the probability of every existing node decreases because the normalization factor increases. Thus, upon insertion, every item’s probability mass would need to be updated. There are  $|V|$  items, so insertion becomes a  $\Theta(|V|)$  operation in this situation. Conversely, the preference mass of each node is unaffected by the introduction of a new node. Insertion in this scenario is a  $O(\log|V|)$  operation; see Figure 6.

We use Price’s model as an illustrative example. Recall that, in Price’s model, a new node is introduced at each timestep, and an edge from the new node to an existing node is added preferentially. We first identify an existing node via Sample. We then create a new node and add edge from the



**Figure 7: The empirical run time of the simulator using different index types. The acronyms in the legend indicate the preference function type, the sort key, and the definition of priority for different variants of the treap structure. Error bars, barely visible, indicate the 95% confidence interval.**

new node to the existing node. Increment is called on the existing node to reflect the change in preference mass due to the new incoming edge. Finally, the new node is added to the index via Insert. Sample, Increment and Insert are  $O(\log|V|)$  operations, which implies that a single iteration is  $O(\log|V|)$ , and that a simulation with  $|V|$  iterations is  $O(|V|\log|V|)$ . Generating a network with  $|V|$  nodes takes  $\Theta(|V|)$  iterations, so

$$T(G) = O(|V|\log|V|)$$

The set of treaps was designed to make more efficient use of space. The heap was implemented via a dynamic array, which provides amortized constant-time sampling at the cost of some wasted space. We sought to avoid this wastage by instead using binary treaps, which are an extension of binary trees that maintain a heap invariant over a random priority to ensure balance in expectation [3]. We found that the treap-based implementation consistently underperformed the heap-based implementation, so we devote less space to its description.

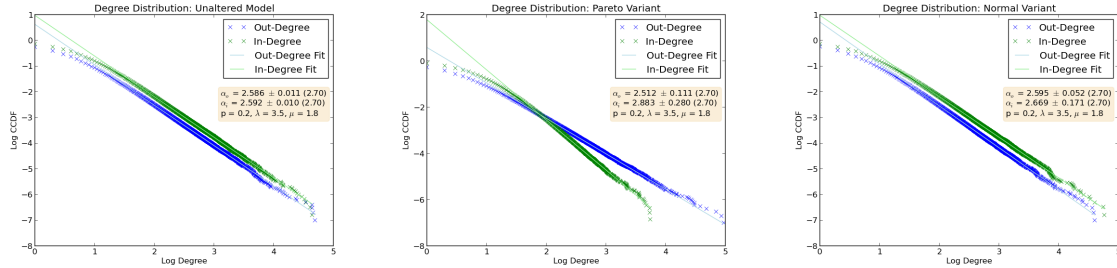
Figure 7 shows the empirical run time of each of these structures as a function of generated network size. All networks were generated from Krapivsky’s model. We found that the binary heap consistently took significantly less time than the treap-based methods to generate networks of several different sizes.

## 4. APPLICATIONS

We validate our generation model by generating sets of networks from the Krapivsky model and comparing the marginal degree distributions inferred from the generated networks with the asymptotic value predicted by the model. We then use the generator to explore some interesting questions. Specifically, we analyze the effect of changing the fitnesses of the Krapivsky model from a constant value to a random variable with various distributions. We also analyze the robustness of Krapivsky’s model to superlinear preference functions.

### 4.1 Validating the Network Generator

We validate our model by comparing the inferred exponents of the marginal distributions of generated networks



**Figure 8: A comparison of the generated and theoretical marginal distributions of in- and out-degree under Krapivsky’s model. The model was parameterized with  $\lambda = 3.5$  and  $\mu = 1.8$  and 10 networks with  $10^7$  nodes were generated. We plot the degree distribution of a single example network. Three variants are investigated: The unaltered model (left), a model with Pareto-distributed fitnesses (center), and a model with normally-distributed fitnesses (right).  $\alpha_o = \mu_o \pm 2\sigma_{SE}(\alpha_o^*)$  specifies the inferred exponent of the marginal out-degree distribution, where  $\mu_o$  is the observed mean of the exponent,  $\sigma_{SE}$  is the standard error, and  $\alpha_o^*$  is the predicted exponent for the unaltered model. The same form holds for  $\alpha_i$ .**

with the known (theoretical, asymptotic) values for the exponents. We generated 10 networks with  $10^7$  nodes each. Figure 8 shows a plot of the base-10 logarithm of both degree and complementary cumulative distribution. The exponents of the marginal degree distributions were inferred via linear regression. We found that, as expected, they both exhibited power-law behavior (evident in the linearity) and that the inferred exponents of the distributions were in relatively good agreement with asymptotic theoretical values. Note that, while networks with  $10^7$  nodes are very large, they are still finite; we believe that this accounts for the small discrepancy between the inferred exponents and the theoretical values.

## 4.2 Exploring Extensions to Krapivsky’s Model

### 4.2.1 Pareto Fitness

We use our network generator to investigate the effects of altering the Krapivsky model. Specifically, we generated networks from a variant where the fitnesses assigned to each node were sampled from a Pareto distribution, rather than assigning the same constant value to each node. Results can be seen in Figure 8. The distribution of in-degree fitness is  $\frac{\lambda d_m^\lambda}{d^{\lambda+1}}$  and has expected value  $\frac{\lambda d_m}{\lambda-1}$ . The parameter  $d_m$  is set to be  $(\lambda-1)$  so that the expected value of the distribution simplifies to  $\lambda$ . The same form was used for the out-degree fitness. Note that this variant still exhibits scale-free behavior, that the inferred exponents are in better agreement with the predicted values, and that the variance of the inferred exponents is higher.

### 4.2.2 Normal Fitness

We also simulated a variant of the Krapivsky model where fitnesses were sampled from a truncated normal distribution. Results can be seen in Figure 8. In-degree fitnesses were sampled from  $N(\lambda, (\lambda/4)^2)$  and out-degree fitnesses from  $N(\mu, (\mu/4)^2)$ . The variances were chosen such that the probability of sampling a negative fitness is very small (less than  $10^{-4}$ ); the distributions were truncated so that any negative samples were replaced with zero. Note that scale-free behavior is still observed and that the inferred exponents of the marginal distributions of in and out-degree are in very close agreement with the simulation of the orig-

inal model.

### 4.2.3 Robustness to Superlinear Preference Functions

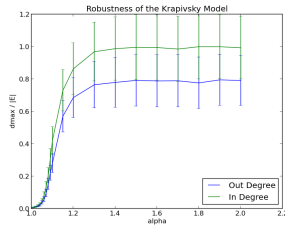
Super-linear preference functions increase the strength of the “rich-get-richer” effect. This can lead to situations where one node quickly overtakes all others and is thus a component of most of the edges in the network. In the extreme case, a star will form; all edges will be connected to the outlier node. We investigate the robustness of Krapivsky’s model to super-linear preference functions by plotting the ratio  $\frac{d_{max}}{|E|}$ , where  $d_{max}$  is the maximum degree, as a function of the preference function exponent  $\alpha$ ; see Figure 9.  $\frac{d_{max}}{|E|}$  will approach 1 as the network approaches a star formation.

There is an interesting side effect to the transition from scale-free to star-structured networks. As the network becomes more star-like, the probability of selecting the most probable node tends to increase. The most probable node always sits at the top of the heap, so it can be accessed in constant time. So, the closer a network’s structure is to a star formation, the higher the probability that an iteration of a PA algorithm will be constant time. For a star structured network in the limit, every iteration will be constant time and the generation of a network with  $|V|$  nodes will be  $\Theta(|V|)$ . This behavior is apparent for finite  $|V|$ ; we have observed that the runtime of the generator tends to decrease as  $\alpha$  increases.

## 5. RELATED WORK

This work is concerned with the problem of efficiently generating networks from PA models. Some examples of PA models include the models of Price (directed networks with scale-free in-degrees) [16], Barabasi and Albert (undirected networks with scale-free degrees) [4], Krapivsky et al. (directed networks with non-independent in and out-degrees which exhibit marginally scale-free behavior) [14], and Capocci et al. (like Krapivsky’s model, but with reciprocity) [8].

There has been some prior work in efficiently generating networks from PA models. Tonelli et al. [19] provide a method for computing an iteration of the linear Yule-Simon cumulative advantage process in constant time. This method can naturally be extended to network generation



**Figure 9: The robustness of the Krapivsky model to superlinear preference functions.**  $\alpha$  indicates the exponent of a preference function of the form  $f(d) = d^\alpha + c$ .  $d_{max}$  is the maximum degree the generated network;  $\frac{d_{max}}{|E|}$  gives the proportion of edges that involve the maximum degree node.  $\frac{d_{max}}{|E|} = 1$  indicates a star formation. Note that there is a phase transition from scale-free to star-structured networks between  $\alpha = 1.0$  and  $\alpha = 1.2$ . 100 networks with  $10^6$  nodes each were generated for each value of  $\alpha$ . Error bars indicate the 95% confidence interval.

through linear PA. However, the extension to nonlinear PA (not shown due to space constraints) is very inefficient in both time and space. Ren and Li [17] describe the simulation of a particular linear PA model, RX, but do not address the general problem of simulating networks from models with general preference functions. Hruz et al. [13] and D’Angelo and Ferretti [10] provide methods for parallelizing the simulation of linear PA, but do not treat the nonlinear case. To the best of our knowledge, our work is the first to address the efficient generation from PA models under possibly nonlinear preference functions.

## 6. CONCLUSION

We provide an efficient framework for simulating preferential attachment under general preference functions which scales to many millions of nodes. We validate this framework empirically and show applications in the simulation and comparison of large networks.

## 7. FUTURE WORK

We have shown that, for nonlinear preferential attachment, the complexity of generating a network with  $|V|$  nodes is both  $\Omega(|V|)$  and  $O(|V|\log|V|)$ . Future work could provide asymptotically tighter bounds.

## 8. REFERENCES

- [1] W. Aiello, F. Chung, and L. Lü. A random graph model for massive graphs. In *the thirty-second annual ACM symposium*, pages 171–180. ACM Press, 2000.
- [2] W. Aiello, F. Chung, and L. Lu. Random evolution in massive graphs. *Foundations of Computer Science*, 2001.
- [3] C. R. Aragon and R. G. Seidel. Randomized search trees. *Foundations of Computer Science, 1989., 30th Annual Symposium on*, pages 540–545, 1989.
- [4] A. L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, pages 1–4, Oct. 1999.

- [5] A.-L. Barabási, R. Albert, and H. Jeong. Internet: Diameter of the World-Wide Web. *Nature*, 401(6749):130–131, Sept. 1999.
- [6] A.-L. Barabási, R. Albert, and H. Jeong. Scale-free characteristics of random networks: the topology of the world-wide web. *Physica A: Statistical Mechanics and its Applications*, 281(1-4):69–77, June 2000.
- [7] A. Broder, R. Kumar, F. Maghoul, and P. Raghavan. Graph structure in the Web. *Computer*, 2000.
- [8] A. Capocci, V. Servedio, F. Colaiori, L. S. Buriol, and e. al. Preferential attachment in the growth of social networks: The internet encyclopedia Wikipedia. *Physical Review E*, 74(3):036116, Sept. 2006.
- [9] Q. Chen, H. Chang, and R. Govindan. The Origin of Power Laws in Internet Topologies Revisited. . . . *Joint Conference of the . . .*, 2002.
- [10] G. D’Angelo and S. Ferretti. Simulation of scale-free networks. In *2nd International ICST Conference on Simulation Tools and Techniques*. ICST.
- [11] D. J. de Solla Price. Networks of Scientific Papers. *Science*, 1965.
- [12] M. Faloutsos, P. Faloutsos, and C. Faloutsos. *On power-law relationships of the Internet topology*, volume 29. ACM, Oct. 1999.
- [13] T. Hruz, S. Geisseler, and M. Schöngens. Parallelism in simulation and modeling of scale-free complex networks. *Parallel Computing*, 36(8):469–485, Aug. 2010.
- [14] P. Krapivsky, G. Rodgers, and S. Redner. Degree Distributions of Growing Networks. *Physical Review Letters*, 86(23):5401–5404, June 2001.
- [15] M. E. J. Newman. The Structure and Function of Complex Networks. *SIAM Review*, 45(2):167–256, Jan. 2003.
- [16] D. d. S. Price. A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science*, 27(5):292–306, 1976.
- [17] W. Ren and J. Li. A fast algorithm for simulating scale-free networks. pages 264–268, 2009.
- [18] B. Ribeiro, W. Gauvin, B. Liu, and D. Towsley. On MySpace Account Spans and Double Pareto-Like Distribution of Friends. In *INFOCOM IEEE Conference on Computer Communications Workshops*, pages 1–6, 2010.
- [19] R. Tonelli, G. Concas, and M. Locci. Three efficient algorithms for implementing the preferential attachment mechanism in Yule-Simon Stochastic Process. *WSEAS Transactions on Information Science & Applications*, 7, 2010.
- [20] A. Vázquez, R. Pastor-Satorras, and A. Vespignani. Large-scale topological and dynamical properties of the Internet. *Physical Review E*, 65(6):066130, June 2002.