

**A Generative Model for Dynamic Contextual
Friendship Networks**

Alice Zheng Anna Goldenberg

July 2006
CMU-ML-06-107



A Generative Model for Dynamic Contextual Friendship Networks

Alice X. Zheng* Anna Goldenberg†

July 2006
CMU-ML-06-107

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

*Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA

†Machine Learning Department, Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

Taking inspiration from real-life friendship formation patterns, we propose a new generative model of evolving social networks. Each person in the network has a distribution over social interaction spheres, which we term “contexts.” The model allows for birth and death of links and addition of new people. Model parameters are learned via Gibbs sampling, and results are demonstrated on real social networks. We study the robustness of our model by examining statistical properties of simulated networks, and compare against well-known properties of real social networks.

Keywords: social networks, statistical modeling, generative models, network evolution

1 Introduction

For decades, social scientists have been fascinated with the study of interpersonal relationship networks. Researchers in physics, statistics, and computer science have developed a parallel interest in similar networks such as the World Wide Web, the Internet, and biochemical networks in the cell. The field has taken on new significance in the public consciousness with the appearance of large on-line communities. This is driving the need for models that are capable of encapsulating dynamic social relations.

Two major schools of thought exist in social network modeling. One approach models the network through regression on sufficient statistics of graph *motifs* such as dyads and triads [12]. These models are descriptive in nature and are often degenerate [4]. A radically different approach comes from the random graph community, where generative models are designed to mimic large-scale average behaviors such as the degree distribution [1, 10]. While random graphs are generated dynamically, their links cannot be modified once established. This contradicts real-life behavior of social links.

This paper marks what we believe to be the first step towards a fully generative statistical model of weighted dynamic social networks. We strive for a model that is complex enough to incorporate fundamental properties of social relations, yet simple enough for feasible parameter learning. We focus on the evolution of interpersonal relationships over time, explicitly modeling the birth and gradual decay of links. Our model generates realistic networks and provides a natural interpretation of the underlying social dynamics. In comparison to previous work on latent space models [5, 11], our model goes one step further by learning a parametrized generating function of the latent space.

Let us start with a motivating example. Imagine that Andy moves to a new town. He may find new collaborators at work, make friends at parties, or meet fellow gym-goers while exercising. In general, Andy lives in a number of different spheres of interaction or *contexts*. He may find himself repeatedly meeting certain people in different contexts at different times, consequently developing stronger bonds with them; acquaintances he never meets again may quickly fade away. Andy’s new friends may also introduce him to their friends (a well-known transitive phenomenon called *triadic closures* in social sciences [14]).

With this example in mind, we present our model in Section 2. We show how to learn the parameters of our model using Gibbs sampling in Section 3. Experimental results are discussed in Section 4 and Section 5 contains a brief survey of related work.

2 Dynamic Contextual Friendship Model (DCFM)

2.1 Notation

DCFM allows the addition of new people into the network at each time step. Let T denote the total number of time steps and N_t the number of people at time t . $N = N_T$ denotes the final total number of people. Let M_t denote the number of new people added to the network at time t , so that $N_t = N_{t-1} + M_t$.

Links between people are weighted. Let $\{W^1, \dots, W^T\}$ be a sequence of weight matrices,

where $W^t \in \mathbb{Z}_+^{N_t \times N_t}$ represents the pairwise link weights at time t . We assume that W^t is symmetric, though it can be easily generalized to the directed case.

The intuition behind our model is that friendships are formed in *contexts*. There are a fixed number of contexts in the world, K , such as work, gym, restaurant, grocery store, etc. Each person has a distribution over these contexts, which can be interpreted as the average percentage of time that he spends in each context.

2.2 The Generative Process

At time t , the N_t people in the network each selects his current context R_i^t from a multinomial distribution with parameter θ_i , where θ_i has a Dirichlet prior distribution:

$$\vec{\theta}_i \sim \text{Dir}(\vec{\alpha}), \quad \forall i = 1, \dots, N, \quad (1)$$

$$R_i^t | \theta_i \sim \text{Mult}(\theta_i), \quad \forall t = 1, \dots, T, i = 1, \dots, N_t. \quad (2)$$

The number of all possible pairwise meetings at time t is $\text{DYAD}^t = \{(i, j) \mid 1 \leq i \leq N_t, i < j \leq N_t\}$. For each pair of people i and j who are in the same context at time t (i.e., $R_i^t = R_j^t$), we sample a Bernoulli random variable F_{ij}^t with parameter $\beta_i \beta_j$. If $F_{ij}^t = 1$, then i and j meets at time t . The parameter β_i may be interpreted as a measurement of friendliness and is a beta-distributed random variable (making it possible for people to have different levels of friendliness):

$$\beta_i \sim \text{Beta}(a, b), \quad \forall i = 1, \dots, N, \quad (3)$$

$$F_{ij}^t | R_i^t, R_j^t \sim \begin{cases} \text{Ber}(\beta_i \beta_j) & \text{if } R_i^t = R_j^t \\ I_0 & \text{o.w.} \end{cases} \quad \forall (i, j) \in \text{DYAD}^t, \quad (4)$$

where I_0 is the indicator function for $F_{ij}^t = 0$.

In addition, the newcomers at time t have the opportunity to form triadic closures with existing people. The probability that a newcomer j is introduced to existing person i is proportional to the weight of the links between i and the people whom j meets in his context. Let $\text{TRIAD}^t = \{(i, j) \mid 1 \leq i \leq N_{t-1}, 1 \leq j \leq N_t\}$ denote the pairs of possible triadic closures. For all $(i, j) \in \text{TRIAD}^t$, we have:

$$G_{ij}^t | W^{t-1}, F_{\cdot j}^t, R_{\cdot}^t \sim \begin{cases} \text{Ber}(\mu_{ij}^t) & \text{if } R_i \neq R_j \\ I_0 & \text{o.w.,} \end{cases} \quad (5)$$

where $\mu_{ij}^t := \sum_{\ell=1}^{N_t} W_{i\ell}^{t-1} F_{\ell j}^t / \sum_{\ell=1}^{N_t} W_{i\ell}^{t-1}$.

Connection weight updates are Poisson distributed. Our choice of a discrete distribution allows for sparse weight matrices, which are often observed in the real world. Pairwise connection weights may drop to zero if the pair have not interacted for a while (though nothing prevents the connection from reappearing in the future). If i and j meets ($F_{ij}^t = 1$ or $G_{ij}^t = 1$), then W_{ij}^t has a Poisson distribution with mean equal to a multiple (γ_h) of their old connection strength. γ_h signifies the rate of weight increase as a result of the ‘‘effectiveness’’ of a meeting; if $\gamma_h > 1$, then the weight will in general increase. (The weight may also decrease under the Poisson distribution, a consequence




Figure 1: Graphical representation of one time step of the generative model. R^t is a N_t -dimensional vector indicating each person's context at time t . F^t is a $N_t \times N_t$ matrix indicating pairwise dyadic meetings. G^t is a $N_{t-1} \times M_t$ matrix that indicate triadic closure for newcomers at time t . W^t is the matrix of observed connection weights at time t . θ , β , γ_h , and γ_ℓ are parameters of the model (hyperparameters are not shown).

perhaps of unhappy meetings.) If i and j do not meet, their mean weight will decrease with rate $\gamma_\ell < 1$. Thus

$$W_{ij}^t \mid W_{ij}^{t-1}, F_{ij}^t, G_{ij}^t, \gamma_h, \gamma_\ell \sim \begin{cases} \text{Poi}(\gamma_h(W_{ij}^{t-1} + \epsilon)) & \text{if } F_{ij}^t = 1 \text{ or } G_{ij}^t = 1 \\ \text{Poi}(\gamma_\ell W_{ij}^{t-1}) & \text{o.w.} \end{cases} \quad (6)$$

where $W_{ij}^{t-1} = 0$ by default for $(i, j) \notin \text{TRIAD}^t$, and ϵ is a small positive constant that lifts the Poisson mean away from zero. As W_{ij}^{t-1} becomes large, γ_h and γ_ℓ control the increase and decrease rates, and the effect of ϵ diminishes. γ_h and γ_ℓ have conjugate Gamma priors:

$$\gamma_h \sim \text{Gamma}(c_h, d_h), \quad (7)$$

$$\gamma_\ell \sim \text{Gamma}(c_\ell, d_\ell). \quad (8)$$

Figure 1 contains a graphical representation of our model. The complete joint probability is:

$$P(\vec{\theta}, \vec{\beta}, \gamma_h, \gamma_\ell, W^{1:T}, R^{1:T}, F^{1:T}, G^{1:T}) = P(\vec{\theta})P(\vec{\beta})P(\gamma_h)P(\gamma_\ell) \prod_t P(R^t \mid \vec{\theta})P(F^t \mid R^t, \vec{\beta})P(G^t \mid R^t, F^t, W^{t-1})P(W^t \mid G^t, F^t, W^{t-1}) \quad (9)$$

3 Learning Parameters via Gibbs Sampling

Our model utilizes $O(N)$ parameters to represent the distribution of a sequence of T integer-valued weight matrices each of size $O(N^2)$. (Note that the number of parameters does not depend on the number of time steps.) There are also a number of hidden variables that indicate the underlying pairwise interaction states: $\{R^t, F^t, G^t\}_{t=1\dots T}$. We use Gibbs sampling to sample from the posterior distribution of these random variables given observed weight matrices $\{W^1, \dots, W^T\}$.¹

3.1 Posterior Distributions of Parameters

$$\vec{\theta}_i \mid \dots \sim \text{Dir}(\vec{\alpha} + \vec{\alpha}'_i), \quad (10)$$

$$P(\beta_i \mid \dots) \propto \beta_i^{A_i + a - 1} (1 - \beta_i)^{b - 1} \prod_{j \neq i} (1 - \beta_i \beta_j)^{B_{ij}}, \quad (11)$$

$$\gamma_h \mid \dots \sim \text{Gamma}(c_h + w_h, (v_h + 1/d_h)^{-1}), \quad (12)$$

$$\gamma_\ell \mid \dots \sim \text{Gamma}(c_\ell + w_\ell, (v_\ell + 1/d_\ell)^{-1}). \quad (13)$$

We use \dots as a shorthand for “all other variables in the model.” In Equation 10, $\alpha'_{ik} := \sum_{t=1}^T I_{(R_i=k)}$ is the total number of times person i is seen in context k . In Equation 11, $A_i := |\{(j, t) \mid R_j^t = R_i^t \text{ and } F_{ij}^t = 1\}|$ is the total number of dyadic meetings between i and any other person, and $B_{ij} := |\{t \mid R_i^t = R_j^t \text{ and } F_{ij}^t = 0\}|$ is the total number of times i has “missed” an opportunity for a dyadic meeting. Let $H := \{(i, j, t) \mid F_{ij}^t = 1 \text{ or } G_{ij} = 1\}$ represent the union of the set of dyadic and triadic meetings, and $\mathcal{L} := \{(i, j, t) \mid (i, j) \in \text{DYAD}^t \text{ and } F_{ij}^t = 0\}$ the set of missed dyadic meeting opportunities. $w_h := \sum_{(i,j,t) \in \mathcal{H}} W_{ij}^t$ is the sum of updated weights after the meetings, and $v_h := \sum_{(i,j,t) \in \mathcal{H}} (W_{ij}^{t-1} + \epsilon)$ is the sum of the original weights plus a fixed constant. $w_\ell := \sum_{(i,j,t) \in \mathcal{L}} W_{ij}^t$ is the sum of weights after the missed meetings, and $v_\ell := \sum_{(i,j,t) \in \mathcal{L}} W_{ij}^{t-1}$ is the sum of original weights. (Here we use zero as the default value for W_{ij}^{t-1} if j is not yet present in the network at time $t - 1$.)

Due to coupling from the pairwise interaction terms $\beta_i \beta_j$, the posterior probability distribution of β_i cannot be written in a closed form. However, since β_i lies in the range $[0, 1]$, one can perform coarse-scale numerical integration and sample from interpolated histograms. Alternatively, one can design Metropolis-Hasting updates for β_i , which has the advantage of maintaining a proper Markov chain.

3.2 Posterior Distributions of Hidden Variables

The variables F_{ij}^t and G_{ij} are conditionally dependent given the observed weight matrices. If a pairwise connection W_{ij} increases from zero to a positive value at time t , then i and j must either have a dyadic or a triadic meeting. On the other hand, dyadic meetings are possible only when i and j are in the same context, and triadic meetings when they are in different contexts. Hence F_{ij}^t

¹Learning results do not seem to be sensitive to the values of hyperparameters. In our experiments, we set the hyperparameters $(\vec{\alpha}_i, a, b, c_h, d_h, c_\ell, d_\ell)$ to reasonable fixed values based on simulations of the model.

and G_{ij}^t may never both be 1. In order to ensure consistency, F_{ij}^t and G_{ij}^t must be updated together. For $(i, j) \in \text{TRIAD}^t$,

$$\begin{aligned}
P(F_{ij}^t = 1, G_{ij} = 0 \mid \dots) &\propto I_{(R_i^t=R_j^t)}(\beta_i\beta_j)\text{Poi}(W_{ij}^t; \gamma_h\epsilon), \\
P(F_{ij}^t = 0, G_{ij} = 1 \mid \dots) &\propto I_{(R_i^t \neq R_j^t)}\mu_{ij}\text{Poi}(W_{ij}^t; \gamma_h\epsilon), \\
P(F_{ij}^t = 0, G_{ij} = 0 \mid \dots) &\propto \left[I_{(R_i^t=R_j^t)}(1 - \beta_i\beta_j) + I_{(R_i^t \neq R_j^t)}(1 - \mu_{ij}) \right] I_{(W_{ij}^t=0)}.
\end{aligned} \tag{14}$$

For $(i, j) \in \text{DYAD}^t \setminus \text{TRIAD}^t$,

$$\begin{aligned}
P(F_{ij}^t = 1 \mid \dots) &\propto I_{(R_i^t=R_j^t)}(\beta_i\beta_j)\text{Poi}(W_{ij}^t; \gamma_h(W_{ij}^{t-1} + \epsilon)), \\
P(F_{ij}^t = 0 \mid \dots) &\propto (I_{(R_i^t=R_j^t)}(1 - \beta_i\beta_j) + I_{(R_i^t \neq R_j^t)})\text{Poi}(W_{ij}^t; \gamma_h W_{ij}^{t-1}).
\end{aligned} \tag{15}$$

There are also consistency constraints for R^t . For example, if $F_{ij}^t = F_{jk}^t = 1$, then i, j , and k must all lie within the same context. If $G_{kl} = 1$ in addition, then l must belong to a different context from i, j , and k . The F variables propagate transitivity constraints, whereas G propagates exclusion constraints.

To update R^t , we first find connected components within F^t . Let p denote the number of components and I the index set for the nodes in the i -th component. We update each R_I^t as a block. Imagine an auxiliary graph where nodes represent these connected components and edges represent exclusion constraints specified by G , i.e., I is connected to J if $G_{ij} = 1$ for some $i \in I$ and $j \in J$. Finding a consistent setting for R^t is equivalent to finding a feasible K -coloring of the auxiliary graph, where K is the total number of contexts. We sample R_I^t sequentially according to an arbitrary ordering of the components. Let $\pi(I)$ denote the set of components that are updated before I . The posterior probabilities are:

$$P(R_I^t = k \mid R_{\pi(I)}^t, G) \propto \begin{cases} 0 & \text{if } G_{IJ} = 1 \text{ and } R_J^t = k \text{ for some } J \in \pi(I) \\ \prod_{i \in I} \theta_{ik} & \text{o.w.} \end{cases} \tag{16}$$

These sequential updates correspond to a greedy K -coloring algorithm; they are approximate Gibbs sampling steps in the sense that they do not condition on the entire set of connected components.

4 Experiments

In this section we would like to address three points. First, we show that our Gibbs sampler is able to recover the true parameters of simulated networks. Second, we present DCFM learning results on a real co-authorship network. Finally, we give a brief overview of the range of the behaviors that DCFM can simulate, and show that the model captures well-known properties of social networks such as power law distribution of node degrees.

4.1 Parameter Learning Results

As a sanity check for our parameter learning algorithm, we apply it to a network generated from the model. Overall we find that the learning procedure quickly converges to a stable set of parameter

values. The hyperparameters are set to be those used in the simulation, however we find that Gibbs is robust to changes in hyperparameters. Below are convergence plots for γ_h , γ_ℓ and β parameters for a small dataset of 78 people in 10 contexts where links form and dissolve over 84 time steps. The number of Gibbs iterations is 10,000.

Figure 2 contains a scatter plot of the friendliness β parameters (mean of the posterior vs. true values). Figure 3 contains the convergence plot and the posterior distribution of γ_h and γ_ℓ . Note that, due to noise in the sampling process, the γ_h values oscillate around the median of 1.90 (true value being 2) and γ_ℓ values have median 0.96 (true value being 1). We observe similar convergence trends for θ_i .




Figure 2: β parameter scatter plot.

To test the interpretability of the model on real data, we learn the parameters of DCFM on a real-life collaboration network. We have collected interaction information, such as meetings and co-authorships, over 13 years for 120 people connected to our lab. We omit the names of people due to anonymity constraints. More information will be available in the full version of the paper.

Sampling results yield the median values of $\gamma_h = 1$ and $\gamma_\ell = .02$. This shows that lab members either have steady collaboration patterns over time, or have spurious interactions that quickly die off. We also find that the head of the lab, who participates in most but not all of the collaborations, is quite friendly with $\beta = .86$. Interestingly, the next most friendly person with $\beta = .82$ is a student who is not the most prolific but has many co-authors. In the process of parameter learning, we find that our original assumption of 10 contexts is not enough to accommodate all the consistency constraints arising between R , F , and G . Thus we increase the number of contexts to 20. Figure 4 shows the learned context distributions of the above mentioned professor and student. The two are mostly comparable except for contexts 7, 8, and 17. Assuming that the contexts represent topics of study, the student is the most interested in 7 and least in 17, whereas the professor has a rather uniform distribution over all fields, most of all number 8.




Figure 3: The left-hand column shows convergence plots of γ_h and γ_l over Gibbs sampling iterations. The right-hand column contains histogram of the sampled values.




Figure 4: Context distributions of the two most “friendly” people in the co-authorship network.

We are currently in the process of collecting more learning results on larger simulated and real networks.

4.2 Simulations and Reality

We study the range of behaviors of our model under different parameter settings using a set of established metrics. Our goal is to see whether the model can capture behaviors that have been observed in a variety of real networks. We find that DCFM can generate networks whose clustering coefficient, average path length, and effective diameter fall within the range of observed values. [1] and [10] present empirical measurements for a variety of real-life networks and give formal definitions for the metrics used.

Unless otherwise specified, the number of contexts K is set to 10 in all our simulation experiments. The context preference parameter θ_i is drawn from a peaked Dirichlet prior, where $\alpha_{k^*} = 5$ for a randomly selected k^* , and $\alpha_k = 1$ otherwise, so that each person slightly prefers one of the contexts. The friendliness parameter β_i is drawn from a Beta(a, b) distribution, where $a = 1$ and b varies. The weights update rates are $\gamma_h = 2$, $\gamma_\ell = 0.5$, and $\epsilon = 1$. We add one person to the network at every time step, so that $N_t = t$ and $M_t = 1$. All experiments are repeated with 10 trials.

4.2.1 Effects of Friendliness

The parameter β_i determines the “friendliness” of the i -th person and is drawn from a Beta(a, b) distribution. As b increases from 2 to 10, average friendliness decreases from 0.33 to 0.09. We wish to test the effect of b on overall network properties. In order to isolate the effects of friendliness, we fix the context assignments by setting $R_i^t = R_i^1$ for all $t > 1$. In this setting, people do not form triadic closures, and connection weights are updated only through dyadic meetings. As people become less friendly, one as expected observes a corresponding decrease in average node degree (Figure 5, top left). Interestingly, the clustering coefficient goes up as friendliness goes down. This is because low friendliness makes for smaller clusters, and it is easier for smaller clusters to become densely connected. We also observe large variance in average path length and effective diameter at low friendliness levels. This is due to the fact that most clusters now contain one to two people. As small clusters become connected by chance, shortest path lengths varies from trial to trial.

4.2.2 Degree distribution

Under different parameter settings, our model may generate networks with a variety of degree distributions. Lower levels of friendliness typically lead to more power-law-like degree distributions, while higher levels often result in a heavier tail. In Figure 6, we show two degree distribution plots for different friendliness levels. On the left ($b = 3$) the quadratic polynomial is a much better fit to the degree distribution than the linear one. This means that, when people are more friendly, the drop off in the number of people with high node degree is slower than would be expected under the power law. We do observe the power law effect at a lower level of friendliness. In the right-hand




Figure 5: Effects of the friendliness parameter on a network of 200 people with fixed contexts. The x-axes represent different values of b in $\text{Beta}(1, b)$.

side plot, the linear polynomial with coefficient -1.6 gives as good of a fit as a quadratic function. This coefficient value lies well within the normally observed range for real social networks [1].

5 Related Work

The principles underlying the mechanisms by which relationships evolve are still not well understood [9]. Current models aim at either describing observed phenomena or predicting future trends. A common approach is to select a set of graph based features, such as degree distribution or the number of dyads and triangles, and create models that mimic observed behavior of the evolution of these features in real life networks. Works of [7, 2, 3] in physics and [13, 6] in social sciences follow this approach. However, under models of average behavior, the actual links between any two given people might not have any meaning. Consequently, these models are often difficult to interpret.

Another approach aims to predict likely future friends and collaborators based on the properties of the network seen so far [10, 9]. These models often have problems of scalability, and cannot encode common network dynamics such as mobility and link modification. Moreover, these models usually do not take into account triadic closure, a phenomenon of great importance in social networks [14, 8].

[11] presents an interesting dynamic social network model. This work builds on [5], which introduces latent positions for each person in order to explain observed links. If two people are close in the latent space, they are likely to have a connection. [5] estimate latent positions in a static data set. [11] adds a dynamic component by allowing the latent positions to be updated based on both their previous positions and on the newly observed interactions, One can imagine




Figure 6: Log-log plot of the degree distributions of a network with 200 people. β_i is drawn from Beta(1, 3) for the plot on the left, and from Beta(1, 8) on the right. Solid lines represent a linear fit and dashed lines quadratic fit to the data.

a generative mechanism that governs such perturbations of latent positions, though authors do not offer one. The model of [11] assumed a fixed number of people in the network.

6 Discussion and Future Work

Researchers have long sought a learnable social network model built upon solid principles from social science. In this paper, we propose a generative model for evolving friendship networks based on the idea of social contexts. Our model adheres to real-life behavior of friendship networks at the cost of increased complexity in the generating process. Despite its structural complexity, the model is relatively parsimonious in the parameters, and parameter learning is possible via Gibbs sampling. The learning algorithm scales on the order of $O(N^2T)$, and we have performed preliminary experiments on networks with as many as 600 people.

Our simulations commence the process of exploring the range behavior of this model, but we have yet to systematically analyze the model for the type of behaviors it can and cannot emulate. Another issue we have not touched upon in this paper is identifiability. We expect to answer these questions and more in future work.

References

- [1] R Albert and A Barabási. Statistical mechanics of social networks. *Rev of Modern Physics*, 74, 2002.

- [2] A Barabasi, H Jeong, Z Neda, E Ravasz, A Schubert, and T Vicsek. Evolution of the social network of scientific collaboration. *Physica A*, 311(3–4):590–614, 2002.
- [3] J Davidsen, J Ebel, and S Bornholdt. Emergence of a small world from local interactions: Modeling acquaintance networks. *Physical Review Letters*, 88, 2002.
- [4] M Handcock. Assessing degeneracy in statistical models of social networks. Working Paper 39, University of Washington, 2003.
- [5] P Hoff, A Raftery, and M Handcock. Latent space approaches to social network analysis. *Journal of the American Statistical Association*, 97:1090–1098, 2002.
- [6] M Huisman and T Snijders. Statistical analysis of longitudinal network data with changing composition. *Sociological Methods and Research*, 32(2):253–287, 2003.
- [7] E Jin, M Girvan, and M Newman. The structure of growing social networks. *Physical Review Letters E*, 64, 2001.
- [8] G Kossinets and D Watts. Empirical analysis of an evolving social network. *Science*, 311(5757):88–90, 2006.
- [9] D. Liben-Nowell and J. Kleinberg. The link prediction problem for social networks. In *Proc. 12th International Conference on Information and Knowledge Management*, 2003.
- [10] M Newman. The structure of scientific collaboration networks. In *Proceedings of the National Academy of Sciences USA*, volume 98, pages 404–409, 2001.
- [11] P Sarkar and A Moore. Dynamic social network analysis using latent space models. *SIGKDD Explorations: Special Edition on Link Mining*, 2005.
- [12] T Snijders, P Pattison, G Robins, and M Handcock. New specifications for exponential random graph models. Submitted for publication, 2004.
- [13] G Van De Bunt, M Van Duijn, and T Snijders. Friendship networks through time: An actor-oriented dynamic statistical network model. *Computation and Mathematical Organization Theory*, 5(2):167–192, 1999.
- [14] S Wasserman and K Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, 1994.

