

# Dynamic Stochastic Blockmodels for Time-Evolving Social Networks

Kevin S. Xu, *Member, IEEE*, and Alfred O. Hero, III, *Fellow, IEEE*

**Abstract**—Significant efforts have gone into the development of statistical models for analyzing data in the form of networks, such as social networks. Most existing work has focused on modeling static networks, which represent either a single time snapshot or an aggregate view over time. There has been recent interest in statistical modeling of *dynamic networks*, which are observed at multiple points in time and offer a richer representation of many complex phenomena. In this paper, we present a state-space model for dynamic networks that extends the well-known *stochastic blockmodel* for static networks to the dynamic setting. We fit the model in a near-optimal manner using an extended Kalman filter (EKF) augmented with a local search. We demonstrate that the EKF-based algorithm performs competitively with a state-of-the-art algorithm based on Markov chain Monte Carlo sampling but is significantly less computationally demanding.

**Index Terms**—State-space social network model, dynamic network, on-line estimation, extended Kalman filter.

## I. INTRODUCTION

THE study of networks has emerged as a topic of great interest in recent years. Many complex physical, biological, and social phenomena ranging from protein-protein interactions to the formation of social acquaintances can be naturally represented by networks. Analysis of data in the form of networks has recently captured the attention of the signal processing community [2]–[6]. To date, much research has focused on static networks, which either represent a single time snapshot of the phenomenon of interest or an aggregate view over time. Accordingly, many statistical models for networks have been developed for static networks; see [7] for a survey of the literature. However, most complex phenomena, including social behavior, are time-varying, which has led researchers in recent years to examine *dynamic, time-evolving networks*. Previous studies have typically examined aspects of dynamic networks related to their growth over time, including densification [8],

[9] and shrinking diameters [9], and to their structural changes, including the temporal evolution of communities in the network [10]–[12].

In this paper, we consider dynamic networks represented by a sequence of snapshots of the network at discrete time steps. Both nodes and edges of the network could be *added or removed* over time. In a dynamic social network, for example, the nodes could correspond to people, and the edges could correspond to interactions between people. The presence of an edge between two nodes  $i$  and  $j$  at time step  $t$  would then indicate that an interaction between  $i$  and  $j$  occurred during the time window represented by time step  $t$ .

We characterize such dynamic networks using a set of unobserved *time-varying states* from which the observed snapshots are derived. We utilize a state-space model for dynamic networks first proposed in [1] that combines two types of statistical models: a static model for the individual snapshots and a temporal model for the evolution of the states. The network snapshots are modeled using the *stochastic blockmodel* (SBM) [13], a simple parametric model commonly used in the analysis of static social networks. The state evolution is modeled by a stochastic dynamic system.

Using a Gaussian approximation, which becomes increasingly accurate as the SBM block sizes increase, we employ a *near-optimal* procedure for fitting the proposed model in the *on-line* setting where only past and present network snapshots are available. The inference procedure consists of an extended Kalman filter (EKF) [14] augmented with a local search strategy. The proposed algorithm is considerably faster than a state-of-the-art algorithm that uses Markov chain Monte Carlo (MCMC) sampling, yet our experiments show that the proposed algorithm has comparable accuracy to the more computationally demanding MCMC-based algorithm in recovering the true states. We apply the proposed algorithm to analyze the Enron network [15], [16], a dynamic social network of email communication, and reveal several interesting trends that cannot be identified by aggregate statistics such as the total number of emails sent at each time step.

We first proposed the main elements of our model, including the Gaussian approximation and EKF algorithm for tracking dynamic SBMs, in [1]. This paper refines and extends the analysis of the model in several ways. First, a detailed development of the approximate inference procedure is given, including a study of the approximation accuracy. Second, a more extensive performance analysis is presented, for both simulated and real data, that establishes the advantages of the proposed dynamic SBM method relative to other methods.

Manuscript received September 15, 2013; revised December 19, 2013; accepted February 25, 2014. Date of publication March 06, 2014; date of current version July 16, 2014. This work was supported in part by Army Research Office grant W911NF-12-1-0443. The work of K. S. Xu was supported in part by an award from the Natural Sciences and Engineering Research Council of Canada. Preliminary parts of this work appeared in the conference publication [1]. The guest editor coordinating the review of this manuscript and approving it for publication was Prof. Vikram Krishnamurthy.

K. S. Xu is with the Technicolor Palo Alto Research Center, Palo Alto, CA 94301 USA (e-mail: kevinxu@outlook.com).

A. O. Hero III is with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109-2122 USA (e-mail: hero@umich.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSTSP.2014.2310294

## II. BACKGROUND

### A. Notation

Time steps are denoted by superscripts, while matrix or vector indices are denoted by subscripts. We represent a dynamic network by a time-indexed sequence of graphs, with  $W^t = [w_{ij}^t]$  denoting the adjacency matrix of the graph observed at time step  $t$ . We define  $w_{ij}^t = 1$  if there is an edge from node  $i$  to node  $j$  at time  $t$ , and  $w_{ij}^t = 0$  otherwise. We assume that the graphs are directed, i.e.,  $w_{ij}^t \neq w_{ji}^t$  in general, and that there are no self-edges, i.e.,  $w_{ii}^t = 0$ .  $W^{(t)}$  denotes the set of all snapshots up to time  $t$ ,  $\{W^t, W^{t-1}, \dots, W^1\}$ .  $|V^t|$  and  $|E^t|$  denote the number of observed nodes and edges, respectively, at time  $t$ . We write  $X^{t_1|t_2}$ ,  $t_1 \geq t_2$  to denote a quantity  $X$  at time  $t_1$  computed using only observations from time  $t_2$  and earlier. The notation  $i \in a$  indicates that node  $i$  is a member of class  $a$ .  $|a|$  denotes the number of nodes in class  $a$ . The classes of all nodes at time  $t$  is given by a vector  $\mathbf{c}^t$  with  $c_i^t = a$  if  $i \in a$  at time  $t$ . We denote the submatrix of  $W^t$  corresponding to the relations between nodes in class  $a$  and nodes in class  $b$  by  $W_{[a][b]}^t$ . Finally, we denote the vectorized equivalent of a matrix  $X$ , i.e. the vector obtained by simply stacking columns of  $X$  on top of one another, by  $\mathbf{x}$ . Doubly-indexed subscripts such as  $x_{ij}$  denote entries of matrix  $X$ , while singly-indexed subscripts such as  $x_i$  denote entries of the vectorized equivalent  $\mathbf{x}$ .

### B. Static Stochastic Blockmodels

We present a brief summary of the *static stochastic blockmodel* (SSBM) [13], which we use as the static model for individual network snapshots. Consider a snapshot at an arbitrary time step  $t$ . An SSBM is parameterized by a  $k \times k$  matrix  $\Theta^t = [\theta_{ab}^t]$ , where  $\theta_{ab}^t$  denotes the *probability of forming an edge* between a node in class  $a$  and a node in class  $b$ , and  $k$  denotes the number of classes. The SSBM decomposes the adjacency matrix into  $k^2$  blocks, where each block is associated with relations between nodes in classes  $a$  and  $b$ . Each block  $(a, b)$  corresponds to a submatrix  $W_{[a][b]}^t$  of the adjacency matrix  $W^t$ . Thus, given the *class membership vector*  $\mathbf{c}^t$ , each entry of  $W^t$  is an independent realization of a Bernoulli random variable with a block-dependent parameter; that is,  $w_{ij}^t \sim \text{Bernoulli}(\theta_{c_i^t c_j^t}^t)$ .

SSBMs are used in two settings: the *a priori* blockmodeling setting, where class memberships are known or assumed, and the objective is to estimate the matrix of edge probabilities  $\Theta^t$ , and the *a posteriori* blockmodeling setting, where the objective is to simultaneously estimate  $\Theta^t$  and the class membership vector  $\mathbf{c}^t$ . Since each entry of  $W^t$  is independent, the likelihood for the parameters  $\Phi^t$  of the SSBM is given by

$$f(W^t; \Phi^t) = \prod_{i \neq j} \left( \theta_{c_i^t c_j^t}^t \right)^{w_{ij}^t} \left( 1 - \theta_{c_i^t c_j^t}^t \right)^{1 - w_{ij}^t}. \quad (1)$$

The likelihood (1) can be rewritten as

$$f(W^t; \Phi^t) = \exp \left\{ \sum_{a=1}^k \sum_{b=1}^k [m_{ab}^t \log(\theta_{ab}^t) + (n_{ab}^t - m_{ab}^t) \log(1 - \theta_{ab}^t)] \right\}, \quad (2)$$

where  $m_{ab}^t = \sum_{i \in a} \sum_{j \in b} w_{ij}^t$  denotes the number of *observed edges* in block  $(a, b)$ , and

$$n_{ab}^t = \begin{cases} |a||b| & a \neq b \\ |a|(|a| - 1) & a = b \end{cases} \quad (3)$$

denotes the number of *possible edges* in block  $(a, b)$  [17]. The parameters are given by  $\Phi^t = \Theta^t$  in the *a priori* setting, and  $\Phi^t = \{\Theta^t, \mathbf{c}^t\}$  in the *a posteriori* setting. In the *a priori* setting, a sufficient statistic for estimating  $\Theta^t$  is the matrix  $Y^t$  of *block densities* corresponding to ratios of observed edges relative to possible edges within each block, which has entries  $y_{ab}^t = m_{ab}^t / n_{ab}^t$ . The matrix  $Y^t$  is also the maximum-likelihood estimate of  $\Theta^t$  [17].

Parameter estimation in the *a posteriori* setting is more involved, and many methods have been proposed, including Gibbs sampling [18], label-switching [17], [19], and spectral clustering [20], [21]. The label-switching methods use a heuristic for solving the combinatorial optimization problem of maximizing the likelihood (2) over the set of possible class memberships, which is too large for an exhaustive search to be tractable. The spectral clustering methods utilize the eigenvectors of the adjacency matrix  $W^t$  or a similar matrix to estimate the class memberships.

### C. Related Work

Several statistical models for dynamic networks have previously been proposed for modeling and tracking dynamic networks [7]. Guo *et al.* [22] proposed a temporal extension of the exponential random graph model (ERGM) called the hidden temporal ERGM. Sarkar and Moore [23] proposed a temporal extension of the latent space network model and developed an algorithm to compute point estimates of node positions over time using conjugate gradient optimization initialized from a multidimensional scaling solution. In [24], Sarkar *et al.* proposed a Gaussian approximation that allowed for approximate inference on the dynamic latent space model using Kalman filtering. The approach of [24] is similar in flavor to the approach we employ in this paper; however, our approach involves a different static model, namely the stochastic blockmodel, for the network snapshots and uses this model to develop an extended Kalman filter (EKF) to track the model parameters.

Hoff [25] proposed a dynamic latent factor model analogous to an eigenvalue decomposition with time-invariant eigenvectors and time-varying eigenvalues. The model is applicable to many types of data in the form of multi-way arrays, including dynamic social networks, and is fit using MCMC sampling. In [26], Lee and Priebe proposed a latent process model for attributed (multi-relational) dynamic networks using random dot product spaces. The authors fit mathematically tractable first- and second-order approximations of the random dot process model, for which individual network snapshots are drawn from attributed versions of the Erdős-Rényi and latent space models, respectively. Perry and Wolfe [27] proposed a point process model for dynamic networks of directed interactions and a partial likelihood inference procedure to fit their model. The authors model interactions using a multivariate counting process

that accounts for effects including homophily. Their model operates in continuous time, unlike the proposed model in this paper, which operates on discrete-time snapshots.

More closely related to the state-space dynamic network model we consider in this paper are several temporal extensions of stochastic blockmodels (SBMs). Xing *et al.* [28] and Ho *et al.* [29] proposed temporal extensions of a mixed-membership version of the SBM using linear state-space models for the real-valued class memberships. In [30], Yang *et al.* proposed a temporal extension of the SBM that is similar to our proposed model. The main difference is that the authors explicitly modeled nodes changing between classes over time by using a transition matrix that specifies the probability that a node in class  $i$  at time  $t$  switches to class  $j$  at time  $t+1$  for all  $i, j, t$ . The authors fit the model using a combination of Gibbs sampling and simulated annealing, which they refer to as *probabilistic simulated annealing* (PSA). We use the performance of the PSA algorithm as a baseline for comparison with the less computationally demanding EKF-based approximate inference procedure we utilize in this paper.

### III. DYNAMIC STOCHASTIC BLOCKMODELS

We now present a state-space model for dynamic networks that accomplishes a temporal extension of the SSBM. First we review the dynamic SBM and approximate inference procedure for both a priori and a posteriori blockmodeling proposed in [1]. Next we present an analysis of the time complexity of the inference procedure and discuss hyperparameter estimation. Finally we investigate the validity and accuracy of two key approximations used in the inference procedure.

#### A. A Priori Blockmodels

In the a priori SSBM setting,  $Y^t$  is a sufficient statistic for estimating  $\Theta^t$  as discussed in Section II-B. The entries of  $W_{[a][b]}^t$  are independent and identically distributed (iid) Bernoulli ( $\theta_{ab}^t$ ). Thus the sample mean  $y_{ab}^t$  follows a re-scaled binomial distribution. For large block size,  $y_{ab}^t$  is approximately Gaussian by the Central Limit Theorem with mean  $\theta_{ab}^t$  and variance

$$(\sigma_{ab}^t)^2 = \theta_{ab}^t (1 - \theta_{ab}^t) / n_{ab}^t, \quad (4)$$

where  $n_{ab}^t$  was defined in (3). We assume that  $y_{ab}^t$  is indeed Gaussian for all  $(a, b)$  and posit the linear observation model  $Y^t = \Theta^t + Z^t$ , where  $Z^t$  is a zero-mean independent Gaussian noise matrix with variance  $(\sigma_{ab}^t)^2$  for the  $(a, b)$ th entry.

In the dynamic setting where past snapshots are available, the observations would be given by the set  $Y^{(t)}$ . The set  $\Theta^{(t)}$  can then be viewed as the states of a dynamic system that generates the noisy observation sequence. We complete the model by specifying a model for the state evolution over time. Since  $\theta_{ab}^t$  is a probability and must be bounded between 0 and 1, we instead work with the matrix  $\Psi^t = [\psi_{ab}^t] \in \mathbb{R}^{k \times k}$  where  $\psi_{ab}^t = \log(\theta_{ab}^t) - \log(1 - \theta_{ab}^t)$ , the logit of  $\theta_{ab}^t$ . A simple model for the state evolution is given by the linear dynamic system

$$\psi^t = F^t \psi^{t-1} + \mathbf{v}^t, \quad (5)$$

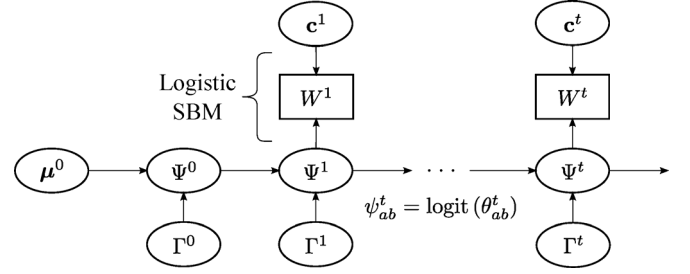


Fig. 1. Graphical representation of the dynamic SBM. The rectangular boxes denote observed quantities, and the ovals denote unobserved quantities. The logistic SBM refers to applying the logistic function (7) to each entry of  $\Psi^t$  to obtain  $\Theta^t$  then generating  $W^t$  using  $\Theta^t$  and  $\mathbf{c}^t$ .

where  $F^t$  is the state transition model applied to the previous state,  $\psi^t$  is the vector representation of the matrix  $\Psi^t$ , and  $\mathbf{v}^t$  is a random vector of zero-mean Gaussian entries, commonly referred to as process noise, with covariance matrix  $\Gamma^t$ . The entries of the process noise vector  $\mathbf{v}^t$  are not assumed to be independent (unlike the entries of  $Z^t$ , which are independent by construction) to allow for states to evolve in a correlated manner. The state transition matrix  $F^t$  may either be known, as in the case of structural time series models [31], or can be estimated using methods for system identification [32]. For the simplest model where  $F^t = I$  for all  $t$ , the state evolution follows a multivariate Gaussian random walk.

The observation model can be written in terms of  $\psi^t$  as<sup>1</sup>

$$\mathbf{y}^t = h(\psi^t) + \mathbf{z}^t, \quad (6)$$

where the function  $h: \mathbb{R}^{k^2} \rightarrow \mathbb{R}^{k^2}$  is defined by

$$h_i(\mathbf{x}) = 1 / (1 + e^{-x_i}), \quad (7)$$

i.e., the logistic function applied to each entry of  $\mathbf{x}$ . We denote the covariance matrix of  $\mathbf{z}^t$  by  $\Sigma^t$ , which is a diagonal matrix<sup>2</sup> with entries given by  $(\sigma_{ab}^t)^2$ . A graphical representation of the dynamic network model is shown in Fig. 1.

The inference procedure for the model is *on-line*, i.e., the state estimate at time  $t$  is formed using only observations from time  $t$  and earlier. We assume the initial state is Gaussian distributed, i.e.,  $\psi^0 \sim \mathcal{N}(\mu^0, \Gamma^0)$ , and that  $\{\psi^0, \mathbf{v}^1, \dots, \mathbf{v}^t, \mathbf{z}^1, \dots, \mathbf{z}^t\}$  are mutually independent. If (6) was linear in  $\psi^t$ , then the optimal estimate of  $\psi^t$  in terms of minimizing mean-squared error would be given by the Kalman filter [14]. Due to the non-linearity, we apply the extended Kalman filter (EKF), which linearizes the dynamics about the predicted state and provides a *near-optimal* estimate of  $\psi^t$ . The EKF equations for the model specified by (5), (6) are as follows. The predicted state estimate is

$$\hat{\psi}^{t|t-1} = F^t \hat{\psi}^{t-1|t-1}, \quad (8)$$

<sup>1</sup>Note that we have converted the block densities  $Y^t$  and observation noise  $Z^t$  to their respective vector representations  $\mathbf{y}^t$  and  $\mathbf{z}^t$ .

<sup>2</sup>The indices  $(a, b)$  for  $(\sigma_{ab}^t)^2$  are converted into a single index  $i$  corresponding to the vector representation  $\mathbf{z}^t$ .

and the predicted covariance estimate is

$$R^{t|t-1} = F^t R^{t-1|t-1} (F^t)^T + \Gamma^t. \quad (9)$$

Let  $H^t$  denote the Jacobian of  $h$  evaluated at the predicted state estimate  $\hat{\psi}^{t|t-1}$ .  $H^t$  is a diagonal matrix with  $(i, i)$ th entry given by

$$h'_i(\hat{\psi}_i^{t|t-1}) = \frac{\exp(-\hat{\psi}_i^{t|t-1})}{[1 + \exp(-\hat{\psi}_i^{t|t-1})]^2}. \quad (10)$$

The near-optimal (when the estimation errors are small enough to make the EKF linearization accurate) Kalman gain is given by

$$K^t = R^{t|t-1} (H^t)^T [H^t R^{t|t-1} (H^t)^T + \Sigma^t]^{-1}, \quad (11)$$

from which the updated state estimate

$$\hat{\psi}^{t|t} = \hat{\psi}^{t|t-1} + K^t [y^t - h(\hat{\psi}^{t|t-1})] \quad (12)$$

and the updated covariance estimate

$$R^{t|t} = (I - K^t H^t) R^{t|t-1} \quad (13)$$

are obtained [14]. The updated state estimate  $\hat{\psi}^{t|t}$  provides a near-optimal fit to the model at time  $t$  given the observed sequence  $W^{(t)}$ .

### B. A Posteriori Blockmodels

In many applications, the class memberships  $\mathbf{c}^t$  are not known a priori and must be estimated along with  $\Psi^t$ . This can be done using label-switching methods as in [17], [19], but rather than maximizing the likelihood (2), we maximize the posterior state density given the entire sequence of observations  $W^{(t)}$  up to time  $t$  to account for the prior information. This is done by alternating between label-switching and applying the EKF to arrive at a maximum a posteriori probability (MAP) estimate of  $\mathbf{c}^t$ .

The posterior state density is given by

$$f(\psi^t | W^{(t)}) \propto f(W^t | \psi^t, W^{(t-1)}) f(\psi^t | W^{(t-1)}). \quad (14)$$

By the conditional independence of current and past observations given the current state,  $W^{(t-1)}$  drops out of the first multiplicative factor on the right side of (14). This factor can thus be obtained simply by substituting  $h(\Psi^t)$  for  $\Theta^t$  in (2). We approximate the second term in (14) with  $f(\psi^t | y^{(t-1)})$  using the estimated class memberships at all previous time steps. By applying the Kalman filter to the linearized temporal model [14],  $f(\psi^t | y^{(t-1)}) \sim \mathcal{N}(\hat{\psi}^{t|t-1}, R^{t|t-1})$ . Thus the logarithm of the posterior density is given by

$$p^t = \log f(W^t | \psi^t) + \log f(\psi^t | y^{(t-1)}), \quad (15)$$

- 1:  $\hat{\mathbf{c}}^t \leftarrow \hat{\mathbf{c}}^{t-1}$  {Initialize class memberships}
- 2: Compute block densities  $Y^t$  using  $W^t$  and  $\hat{\mathbf{c}}^t$
- 3: Compute  $\hat{\psi}^{t|t}$  using EKF equations (8)–(13)
- 4: Compute log-posterior  $p^t$  by substituting  $\hat{\psi}^{t|t}$  for  $\psi^t$  in (15)
- 5: **while**  $iter \leq max\_iter$  **do** {Local search algorithm}
- 6:  $\bar{p}^t \leftarrow -\infty$  {Log-posterior of best neighboring solution up to a constant}
- 7:  $\bar{\mathbf{c}}^t \leftarrow \hat{\mathbf{c}}^t$  {Solution currently being visited}
- 8: **for**  $i = 1$  **to**  $|V^t|$  **do** {Visit all neighboring solutions}
- 9: **for**  $j = 1$  **to**  $k$  such that  $\bar{c}_i^t \neq j$  **do**
- 10:  $\bar{c}_i^t \leftarrow j$  {Change class of a single node}
- 11: Compute block densities  $\tilde{Y}^t$  using  $W^t$  and  $\bar{\mathbf{c}}^t$
- 12: Compute  $\tilde{\psi}^{t|t}$  using EKF equations (8)–(13)
- 13: Compute log-posterior  $\tilde{p}^t$  by substituting  $\tilde{\psi}^{t|t}$  for  $\psi^t$  in (15)
- 14: **if**  $\tilde{p}^t > \bar{p}^t$  **then** {Visited solution is better than best neighboring solution}
- 15:  $[\bar{p}^t, \bar{\psi}^{t|t}, \bar{\mathbf{c}}^t] \leftarrow [\tilde{p}^t, \tilde{\psi}^{t|t}, \bar{\mathbf{c}}^t]$
- 16:  $\bar{c}_i^t \leftarrow \bar{c}_i^t$  {Reset class membership of current node}
- 17: **if**  $\bar{p}^t > p^t$  **then** {Best neighboring solution is better than current best solution}
- 18:  $[p^t, \hat{\psi}^{t|t}, \hat{\mathbf{c}}^t] \leftarrow [\bar{p}^t, \bar{\psi}^{t|t}, \bar{\mathbf{c}}^t]$
- 19: **else** {Reached local maximum}
- 20: **break**
- 21: **return**  $[\hat{\psi}^{t|t}, \hat{\mathbf{c}}^t]$

Fig. 2. A posteriori blockmodel inference procedure at time  $t$  using the EKF.

where

$$\begin{aligned} \log f(W^t | \psi^t) &= \sum_{a=1}^k \sum_{b=1}^k \{m_{ab}^t \log [h(\psi_{ab}^t)] \\ &\quad + (n_{ab}^t - m_{ab}^t) \log [1 - h(\psi_{ab}^t)]\}, \\ \log f(\psi^t | y^{(t-1)}) &= -\frac{1}{2} (\psi^t - \hat{\psi}^{t|t-1})^T \\ &\quad \times (R^{t|t-1})^{-1} (\psi^t - \hat{\psi}^{t|t-1}) + C, \end{aligned}$$

and  $C$  is a constant term independent of  $\psi^t$  that can be ignored<sup>3</sup>.

We use the log-posterior (15) as the objective function for label-switching to obtain an a posteriori fit to the dynamic SBM. We find that a simple local search (hill climbing) algorithm [33] initialized using the estimated class memberships at the previous time step suffices, because only a small fraction of nodes change classes between time steps in most applications. Pseudocode for the a posteriori inference procedure is shown in Fig. 2. At the initial time step, we employ the spectral clustering algorithm of Sussman *et al.* [21] for the SSBM to generate an initial estimate  $\hat{\mathbf{c}}^0$  of the class memberships at  $t = 1$ . The spectral clustering algorithm is given in Fig. 3. Using the spectral clustering solution as the initialization prevents the local search from getting stuck in a poor local maximum at the initial time step.

<sup>3</sup>At the initial time step,  $\hat{\psi}^{1|0} = \mu^0$  and  $R^{1|0} = F^1 \Gamma^0 (F^1)^T + \Gamma^1$ .

- 1: Compute singular value decomposition of  $W^1$
- 2:  $\tilde{\Sigma} \leftarrow$  diagonal matrix of  $k$  largest singular values
- 3:  $(\tilde{U}, \tilde{V}) \leftarrow$  left and right singular vectors for  $\tilde{\Sigma}$
- 4:  $\tilde{Z} \leftarrow [\tilde{U}\tilde{\Sigma}^{1/2}, \tilde{V}\tilde{\Sigma}^{1/2}]$  {concatenate scaled left and right singular vectors}
- 5:  $\hat{\mathbf{c}}^0 \leftarrow$  k-means clustering on rows of  $\tilde{Z}$
- 6: **return**  $\hat{\mathbf{c}}^0$

Fig. 3. SSBM spectral clustering initialization to obtain initial estimate  $\hat{\mathbf{c}}^0$  of class memberships at  $t = 1$ .

### C. Time Complexity

We begin with an analysis of the time complexity of the inference procedure for a priori blockmodeling at each time step. Calculating the matrix of block densities  $Y^t$  involves summing over all edges present at time  $t$ , which has  $O(|E^t|)$  time complexity. Application of the EKF requires only matrix-vector multiplications and a matrix inversion (to calculate the Kalman gain). The size of both the observation and state vectors in the EKF is  $k^2$ , so the time complexity of the EKF is dominated by the  $O(k^6)$  complexity of the matrix inversion. Hence the overall time complexity at time  $t$  is  $O(|E^t| + k^6)$ .

A posteriori blockmodeling involves performing a local search at each time step in addition to applying the EKF. At each iteration of the local search, all  $|V^t|(k-1)$  neighboring class assignments are visited. For each class assignment, we compute the EKF estimate  $\hat{\psi}^{t|t}$  and substitute it into the log-posterior (15). As previously mentioned, computing the EKF estimate is dominated by the  $O(k^6)$  complexity of an inverting a  $k^2 \times k^2$  matrix. Evaluating the log-posterior (15) also requires inversion of a  $k^2 \times k^2$  matrix. The matrix inversions are independent of the class assignments so they only need to be performed once at each time step rather than at each iteration of the local search. Thus the time complexity at each local search iteration is reduced to  $O(k^4)$  for the matrix-vector multiplications. The overall time complexity at time  $t$  then becomes  $O(|E^t| + k^6 + |V^t|lk^5)$ , where  $l$  denotes the number of local search iterations. We note that the local search algorithm is easily parallelized; specifically, each visit to a neighboring class assignment can be executed on a separate core or processor, which can significantly reduce the run-time of a posteriori inference.

### D. Estimation of Hyperparameters

The EKF-based inference procedure requires four hyperparameters to be specified:

- 1) the mean  $\mu^0$  of the initial state  $\psi^0$ ,
- 2) the covariance matrix  $\Gamma^0$  of the initial state  $\psi^0$ ,
- 3) the covariance matrix  $\Sigma^t$  of the observation noise  $\mathbf{z}^t$ , and
- 4) the covariance matrix  $\Gamma^t$  of the process (state evolution) noise  $\mathbf{v}^t$ .

The first two hyperparameters relate to the initial state. In the absence of prior information about the network, specifically the matrix  $\Theta^0$  of probabilities of forming edges, we employ a diffuse prior [31]; that is, we let the variances of the initial states approach  $\infty$ . This can be implemented by simply taking  $\hat{\psi}^{1|1} = g(\mathbf{y}^1)$  and  $R^{1|1} = G^1 \Sigma^1 (G^1)^T$ , where  $g_i(\mathbf{x}) = h_i^{-1}(x) =$

$\log(x_i) - \log(1 - x_i)$  is the logit of the  $i$ th entry of  $\mathbf{x}$ , and  $G^1$  is the Jacobian of  $g$  evaluated at  $\mathbf{y}^1$ , which is a diagonal matrix with entries given by  $g'_i(\mathbf{y}^1) = 1/y_i^1 + 1/(1 - y_i^1)$ . Thus the initial state mean and covariance are given by the transformed initial observation and its covariance.

The third hyperparameter  $\Sigma^t$  denotes the covariance matrix of the observation noise. In many applications of state-space models, it is assumed to be time-invariant and estimated jointly with  $\Gamma^t$ . In the dynamic SBM setting, however,  $\Sigma^t$  is necessarily time-varying because it is related to the current state  $\psi^t$  through (4) and the logistic function  $h$  (7). Taking advantage of this relationship, we use a plug-in estimator for  $\Sigma^t$  by substituting  $h(\hat{\psi}^{t|t-1})$  for  $\Theta^t$  in (4).

The final hyperparameter  $\Gamma^t$  denotes the covariance matrix of the process noise  $\mathbf{v}^t$ . Unlike  $\Sigma^t$ , we assume  $\Gamma^t$  to be time-invariant. Furthermore, it is not necessarily diagonal because states could evolve in a correlated manner. For example, if  $\psi_{ab}^t$  increases from time  $t-1$  to time  $t$ , it may be a result of some action by nodes in class  $a$ , which could also affect the other entries in row  $a$  of  $\Psi^t$ . Although  $\Gamma^t$  is not necessarily diagonal, it is desirable to impose some structure on  $\Gamma^t$  so that one does not have to estimate all  $O(k^4)$  covariances individually. Accordingly we assume the structure of  $\Gamma^t$  is such that

$$\gamma_{ij}^t = \begin{cases} s_{diag}, & i = j \\ s_{nb}, & i, j \text{ are neighboring cells in } \Psi^t \\ 0, & \text{otherwise} \end{cases}, \quad (16)$$

where  $i, j$  being neighboring cells means that the matrix indices  $(a_i, b_i)$  corresponding to  $i$  in  $\Psi^t$  are in the same row or column as matrix indices  $(a_j, b_j)$ . This choice for  $\Gamma^t$  exploits the fact that the state  $\Psi^t$  is actually a matrix that has been flattened into a vector  $\psi^t$ .

The objective is then to estimate  $s_{diag}$  and  $s_{nb}$ . Many cost functions and algorithms have been proposed for learning hyperparameters in non-linear dynamic systems; see [34] for a survey of approaches. The typical approach involves iteratively optimizing a cost function, such as the prediction error or likelihood, over the hyperparameters and states. Wan and Nelson [35] present a dual EKF approach for optimizing prediction error separately over the states and hyperparameters for non-linear state evolution models. Ghahramani and Roweis [36] use an expectation-maximization (EM) algorithm to maximize the likelihood of the observation sequence  $W^{(t)}$ . For the experiments in this paper, we assume that the state transition matrix  $F^t = I$  and choose hyperparameters to minimize the mean-squared prediction error.

### E. Approximation Accuracy

As we noted in Section III.A, the proposed inference procedure makes use of two approximations. The first approximation involves modeling the block densities  $\mathbf{y}^t$  by Gaussian random variables. It is simply the Gaussian approximation to the binomial distribution and is valid provided that the blocks are sufficiently large and dense. A rule of thumb often presented in introductory statistics textbooks is that the Gaussian approximation to a binomial( $n, p$ ) distribution is a reasonable approximation provided  $np(1-p) > 5$ . In the context of the pro-

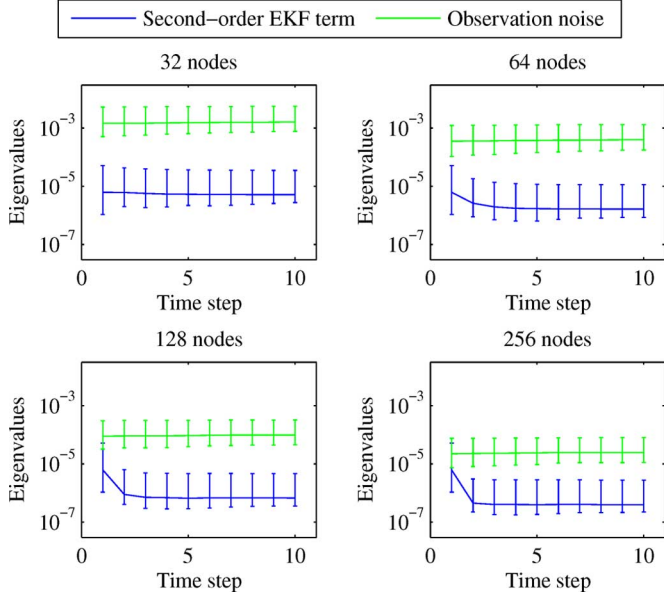


Fig. 4. Eigenvalues of second-order EKF term compared to the observation noise variances  $(\sigma_{ab}^t)^2$  over time, averaged over 50 simulation runs. The line denotes the median of the eigenvalues, while the error bars denote the minimum and maximum. The eigenvalues of the second-order EKF term are generally much smaller than the observation noise variances, suggesting that the linear approximation in the EKF should be very accurate.

posed dynamic SBM, the rule of thumb would correspond to  $n_{ab}^t \theta_{ab}^t (1 - \theta_{ab}^t) > 5$  for all  $a, b, t$ . Recall from (3) that  $n_{ab}^t$  varies as the product of the number of nodes in classes  $a$  and  $b$ . Thus the Gaussian approximation is reasonable even for small values of  $\theta_{ab}^t$ ; for example, 16 nodes in both classes  $a$  and  $b$  is sufficient for  $\theta_{ab}^t = 0.02$ .

The second approximation involves the linearization of the system dynamics in the EKF. The state estimate from the EKF is only approximately optimal due to the linearization, which is a first-order Taylor approximation about the predicted state. Other approximately optimal filters have been proposed for non-linear systems, including the second-order compensated EKF, the unscented Kalman filter, and the particle filter. These filters are often found to perform better than the EKF for complex non-linear models at the cost of higher computational complexity [37]. We argue that the EKF is sufficient for the system model posed by (5) and (6). The state transition model (5) is linear so the only non-linearity is due to the logistic functions  $h_i, i = 1, \dots, k^2$  in the observation model. The EKF should work well when the bias and variance of the second-order term in the Taylor approximation is negligible compared to the observation noise variance [37], i.e.,

$$\frac{1}{4} \left[ \text{tr} \left( h_i'' \left( \hat{\psi}^{t|t-1} \right) R^{t|t-1} \right)^T \text{tr} \left( h_j'' \left( \hat{\psi}^{t|t-1} \right) R^{t|t-1} \right) \right]_{ij} + \frac{1}{2} \left[ \text{tr} \left( h_i'' \left( \hat{\psi}^{t|t-1} \right) R^{t|t-1} h_j'' \left( \hat{\psi}^{t|t-1} \right) R^{t|t-1} \right) \right]_{ij} \ll \Sigma^t,$$

where  $h_i''(\cdot)$  denotes the Hessian matrix of  $h_i$ ,  $[x_{ij}]_{ij}$  denotes a matrix  $X$  with  $(i, j)$ th entry given by  $x_{ij}$ , and  $A \ll B$  denotes that the eigenvalues of  $A$  are much smaller than those of  $B$ .




Fig. 5. MSE comparison for SSBM, EKF, and two particle filters (PFs) with 10 000 particles for 50 simulation runs. PF-G uses the approximate Gaussian observation model, and PF-B uses the actual re-scaled binomial observation model. Edges of the boxes denote 25th and 75th percentiles, and the red (center) lines denote medians. Red (plus) markers denote outliers. The MSE of the EKF is comparable to those of the PFs, confirming the near-optimality of the EKF estimate.

We simulate networks drawn from the dynamic SBM to investigate the contributions of the second-order EKF term. The simulation parameters are chosen based on a synthetic network generator proposed by Newman and Girvan [38]. The network consists of four equally sized classes of nodes. The mean  $\mu^0$  of the initial state  $\psi^0$  is chosen so that  $E[\theta_{aa}^t] = 0.2580$  and  $E[\theta_{ab}^t] = 0.0834$  for  $a, b = 1, 2, 3, 4; a \neq b$ . The initial state covariance  $\Gamma^0$  is chosen to be a scaled identity matrix  $0.04I$ .  $\Theta^t$  evolves according to a Gaussian random walk model on  $\psi^t$ , i.e.,  $F^t = I$  in (5).  $\Gamma^t$  is constructed using (16), with  $s_{diag} = 0.01$  and  $s_{nb} = 0.0025$ . 10 time steps are generated; at each time, we draw a new graph snapshot<sup>4</sup> from the SBM parameterized by  $\Theta^t$  and  $c^t$ .

Since the observation noise variance (4) is inversely proportional to  $n_{ab}^t$ , we can simulate a variety of conditions simply by varying the number of nodes in the network. A comparison of the eigenvalues of the second-order EKF term to the observation noise variances for four choices of the number of nodes is shown in Fig. 4. We find that the contribution of the second-order term is relatively small, suggesting that the first-order Taylor approximation in the EKF is sufficient.

In Fig. 5, we present a comparison of the mean-squared errors (MSEs) for the static SBM (SSBM), the EKF, and two particle filters (PFs), one using the approximate Gaussian distribution for  $y^t$  (PF-G), and one using the actual re-scaled binomial distribution (PF-B). Each data point in the box plot denotes one simulation run, and MSE refers to the mean of the squared tracking error  $\|\hat{\psi}^{t|t} - \psi^t\|_2^2$  over the 10 time steps of the simulation run. The MSEs for the EKF and both PFs are comparable, which confirms that the EKF is sufficient for the proposed model. In addition, there is very little difference in the MSEs for the PF using the approximate Gaussian distribution and the PF using the actual distribution for  $y^t$ , confirming that the Central Limit Theorem approximation is also sufficient.

<sup>4</sup>We draw undirected graph snapshots to match the procedure in [38].






Fig. 6. Comparison of (a) mean-squared tracking error, (b) and class estimation accuracy of a posteriori methods for simulated experiment. The EKF significantly outperforms the SSBM in both tracking and class estimation. The EKF performs slightly better than PSA in tracking and slightly worse in class estimation, but with much less computation time (45 seconds for the EKF compared to 365 seconds for PSA).

#### IV. EXPERIMENTS

##### A. Simulated Stochastic Blockmodels

In this experiment we generate synthetic networks in a manner similar to a simulation experiment in [30]. The network consists of 128 nodes initially split into 4 classes of 32 nodes each. At each time step, 10% of the nodes are randomly selected to leave their class and are randomly assigned to one of the other three classes. The simulated network parameters are chosen to be the same as in the simulated networks described in Section III-E. 10 time steps are generated in each simulation run.

We compare the performance of the EKF-based inference procedure to two baselines. The first is the static stochastic block model (SSBM) fit to each time step individually by spectral clustering [21]. The second baseline is the probabilistic simulated annealing (PSA) algorithm proposed by Yang *et al.* [30], which uses a combination of Gibbs sampling and simulated annealing to perform approximate inference. In the a priori setting, only the EKF and SSBM are applicable, while all three methods are applicable in the a posteriori setting.

1) *Performance Metrics*: The mean-squared errors (MSEs) for the a priori SSBM and EKF are similar to those shown in Fig. 5 for 128 nodes due to the similarity in the experiment setup. The MSEs for the a posteriori methods are shown in Fig. 6(a). The proposed EKF method achieves the lowest MSE in both the a priori and a posteriori settings. The SSBM performs only slightly worse than the EKF in the a priori setting since the observation noise variance is inversely proportional to the square of the number of nodes in each block. However, the SSBM performs extremely poorly in the a posteriori setting due to inaccuracy in the estimation of the true classes.

We evaluate the class estimation accuracy using the adjusted Rand index [39]. An adjusted Rand index of 1 denotes perfect accuracy, and 0 denotes the expected accuracy of a random estimate. The adjusted Rand indices for the a posteriori methods are shown in Fig. 6(b). Both EKF and PSA offer a significant improvement over the SSBM approach. By approximating the posterior distribution over the classes, the PSA method [30] is




Fig. 7. Variation of median adjusted Rand index (over 50 simulation runs) on hyperparameter settings for EKF and PSA in simulated experiment. The proposed EKF method is robust to the choice of hyperparameters, while the PSA method is extremely sensitive to the choice of hyperparameters. (a) EKF: red asterisk denotes hyperparameter estimates by minimizing mean-squared prediction error. (b) PSA: red asterisk denotes hyperparameter estimates by maximizing modularity.  $\hat{\alpha}_{ab} = 1$  for all  $a \neq b$ .

able to achieve slightly higher accuracy in estimating the true classes compared to our EKF approach, which utilizes a MAP estimate of the classes. However, this comes at the expense of computation time. Since the PSA approach uses Gibbs sampling and simulated annealing, it requires significantly more computation time than both the SSBM and EKF approaches (about 6 minutes for PSA compared to under 1 minute for EKF and under 1 second for SSBM on a Linux machine with a quad-core 3.00 GHz Intel Xeon processor).

2) *Hyperparameter Sensitivity*: While PSA is able to slightly outperform our proposed EKF method in class estimation accuracy (at the cost of higher computation time), it is also more sensitive to the choices of hyperparameters. Specifically, Yang *et al.* [30] note that the accuracy in estimating the true classes is sensitive to the choices of the hyperparameters for the conjugate prior of the matrix of edge probabilities  $\Theta^t$ . The conjugate prior for each entry  $\theta_{ab}^t$  in  $\Theta^t$  is Beta distributed with hyperparameters  $\alpha_{ab}, \beta_{ab}$ . In Fig. 7, we plot the variation of adjusted Rand index for different choices of hyperparameters for both the a posteriori EKF and PSA<sup>5</sup>. For the EKF, the hyperparameters are  $s_{diag}, s_{nb}$  as discussed in Section III-D. Note from Fig. 7(a) that the EKF is robust to the choice of hyperparameters, while from Fig. 7(b)

<sup>5</sup>We choose one value of  $\hat{\beta}_{ab}$  for all  $a, b$  and choose  $\hat{\alpha}_{ab} = 1$  for all  $a \neq b$ , identical to [30].

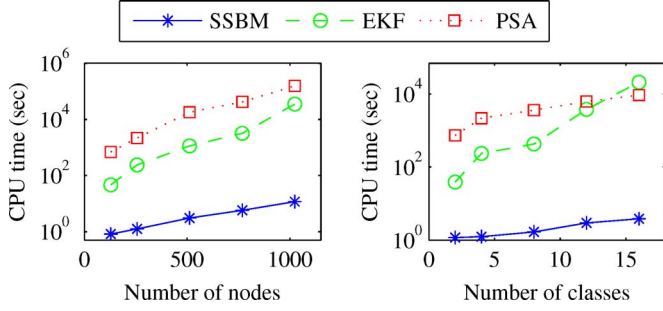


Fig. 8. Scalability of a posteriori methods as number of nodes and number of classes is varied in simulated experiment. The EKF method is faster than PSA (except for a large number of classes), which attains similar levels of accuracy, but slower than the SSBM, which has significantly poorer accuracy.

it can be seen that PSA is extremely sensitive to the choices of  $\hat{\alpha}_{aa}, \hat{\beta}_{ab}$ . In particular, certain choices of  $\hat{\alpha}_{aa}, \hat{\beta}_{ab}$  result in mean adjusted Rand indices close to 0, i.e., barely better than the expected result from randomly assigning nodes to classes. Yang *et al.* [30] recommend choosing values of  $\hat{\alpha}_{aa}, \hat{\beta}_{ab}$  that maximize the modularity criterion [38], which is a measure of the strength of community structure for a given partition (when ground truth is not available). The modularity-based approach is applicable in this experiment because the classes correspond to communities, i.e.,  $\Theta^t$  is diagonally dominant, but is not applicable in the general case where  $\Theta^t$  may not be diagonally dominant. In addition, hyperparameter values that maximize modularity often have extremely poor MSE. In the general case, one might apply a diffuse prior by setting  $\hat{\alpha}_{ab} = \hat{\beta}_{ab} = 1$  for all  $a, b$ . Using this approach, the class estimation accuracy suffers significantly, as shown in Fig. 7(b).

3) *Scalability*: We evaluate the scalability of the EKF-based algorithm and the baselines by varying both the number of nodes and the number of classes in the experiment. Both a priori algorithms require only seconds, so we compare only the a posteriori algorithms. As noted in Section III-C, the time complexity of the a posteriori EKF at each time step is  $O(|E^t| + k^6 + |V^t|lk^5)$ , where  $l$  denotes the number of local search iterations. From [30], the time complexity of PSA at each time step is  $O(|V^t| + |E^t| + k^2 + l(C_1|E^t| + C_2|V^t|))$ , where  $l$  denotes the number of iterations in Gibbs sampling, and  $C_1, C_2$  are constants. The left pane of Fig. 8 shows the variation in computation time as the number of nodes  $|V^t|$  is increased from 128 to 1 024 with the number of classes held constant at 4. The SSBM is fit using spectral clustering [21], which is much faster than the EKF and PSA, but does utilize any temporal model, and suffers from poor accuracy in recovering the true states, as we showed in Fig. 6(a). On the other hand, the EKF and PSA are comparable in accuracy but the EKF is about an order of magnitude faster.

The right pane of Fig. 8 shows the variation in computation time as the number of classes  $k$  is increased from 2 to 16 with the number of nodes held constant at 256. Notice that the EKF is an order of magnitude faster than PSA for  $k \leq 8$ , but slower than PSA once  $k$  reaches 16. This is also an expected result because the time complexity of PSA is quadratic in  $k$ , while the EKF requires  $k^6$ . Yang *et al.* [30] model temporal variation in the class memberships but not in  $\Theta^t$ , unlike the state-space SBM we use;

as a result, PSA has higher tracking error than the EKF as shown in Fig. 6(a). The EKF performs near-optimal state tracking requiring the inversion of a full  $k^2 \times k^2$  covariance matrix as mentioned in Section III-C, which scales poorly in  $k$ . For larger values of  $k$ , one could achieve significant savings in computation time by assuming that the process noise covariance matrix  $\Gamma^t$  is block-diagonal, which would decouple the dynamics. Inference could then be performed by using multiple EKFs with a smaller state space.

### B. MIT Reality Mining

The experiment is conducted on the MIT Reality Mining data set [40]. The data was collected by recording cell phone activity of 94 students and staff at MIT over a year. We construct dynamic networks based on physical proximity, which was measured using scans for nearby Bluetooth devices at 5-minute intervals. We exclude data collected near the beginning and end of the experiment where participation was low. Each time step corresponds to a 1-week interval, resulting in 37 time steps between August 2004 and May 2005.

The affiliation of participants are known, so this data set serves as an excellent real data benchmark for dynamic network analysis. Eagle *et al.* [40] demonstrated that two communities could be found in the time-aggregated network of physical proximity, corresponding to first-year business school students and staff working in the same building. We use these participant affiliations as ground-truth class memberships and compare the class estimation accuracy of the a posteriori dynamic SBM methods. Karrer and Newman [17] showed that class memberships from a posteriori blockmodeling do not agree with community memberships when there is significant degree heterogeneity within communities. To reduce the degree heterogeneity, we connect each participant  $i$  to the 10 other participants who spent the most time in physical proximity of  $i$  during each time step.

A summary of the class estimation performance for the three a posteriori methods is shown in Table I. Both the EKF and PSA, which utilize dynamic models, are more accurate than the SSBM fit using spectral clustering. Notice that the EKF actually has higher class estimation accuracy than the more computationally demanding PSA. We find that this is due to the temporal model of the edge probability matrix  $\Theta^t$  in our proposed dynamic SBM, which PSA does not utilize. Notice from Fig. 9 that PSA does not adapt well to changes in the edge probabilities over time, which degrades its class estimation accuracy compared to our proposed EKF method.

### C. Enron Email Network

We run this experiment on a dynamic social network constructed from the Enron corpus [15], [16], which consists of about 500 000 email messages between 184 Enron employees from 1998 to 2002. We place directed edges between employees  $i$  and  $j$  at time  $t$  if  $i$  sends at least one email to  $j$  during week  $t$ . Each time step corresponds to a 1-week interval. We make no distinction between emails sent “to,” “cc,” or “bcc.” In addition to the email data, the roles of most of the employees within the company (e.g., CEO, president, manager, etc.) are available, which we use as classes for a priori blockmodeling. Employees



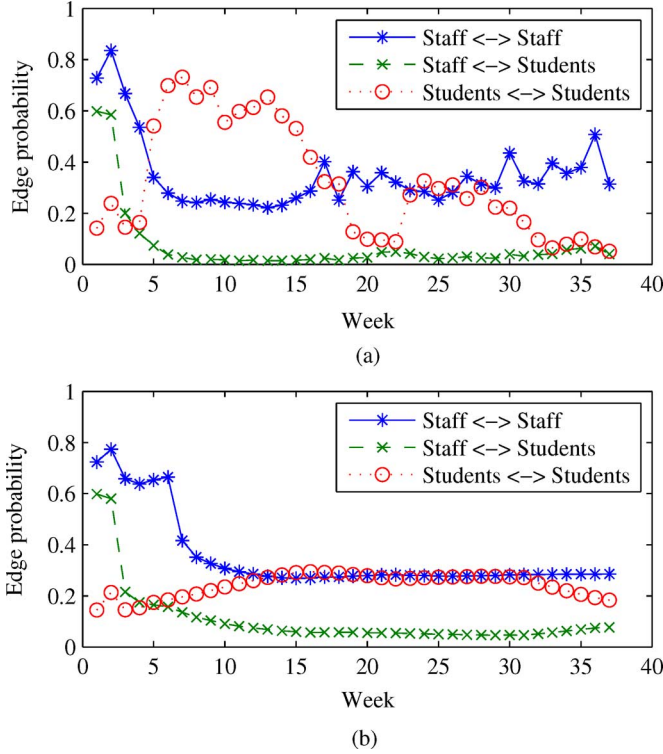


Fig. 9. Estimated edge probabilities for MIT Reality Mining network using (a) EKF and (b) PSA. Temporal dynamics corresponding to fall and winter semesters are visible in the EKF estimates of the edge probabilities between students, but not in the PSA estimates.

TABLE I  
COMPARISON OF CLASS ESTIMATION PERFORMANCE FOR A POSTERIORI METHODS APPLIED TO THE MIT REALITY MINING NETWORK. THE EKF HAS THE HIGHEST ACCURACY AND IS MUCH FASTER THAN PSA

Method	Mean adjusted Rand index	CPU time (sec)
SSBM	0.457	< 1
EKF	0.558	15
PSA	0.539	244

with unknown roles are placed in an “others” class. We remove the first 56 and last 13 weeks of data, where only a few emails were sent.

1) *Dynamic Link Prediction*: Unlike in the simulated data set, we do not know the ground truth states in this experiment. Thus we turn to the task of dynamic link prediction [41] to provide a basis for comparison. Dynamic link prediction differs from static link prediction [42] because the link predictor must simultaneously predict the new edges that will be added at time  $t + 1$ , as well as the current edges (as of time  $t$ ) that will be removed at time  $t + 1$ , from the observations  $W^{(t)}$ . The latter task is not addressed by most static link prediction methods in the literature.

Since the SBM assumes stochastic equivalence between nodes in the same class, the EKF and PSA methods alone are only good predictors of the block densities  $Y^t$ , not the edges themselves. However, the EKF and PSA methods can be combined with a predictor that operates on individual edges to form a good link predictor. A simple individual-level predictor is the exponentially-weighted moving average (EWMA) [43], [44] given by  $\hat{W}^{t+1} = \lambda \hat{W}^t + (1 - \lambda) W^t$ . Using a convex

TABLE II  
COMPARISON OF DYNAMIC LINK PREDICTION PERFORMANCE FOR THE ENRON NETWORK. THE EKF + EWMA AND PSA + EWMA APPROACHES PERFORM COMPARABLY AND BETTER THAN THE EWMA ALONE

Method	AUC	CPU time (sec)
EWMA	0.913	< 1
A priori EKF + EWMA	0.939	< 1
A posteriori EKF + EWMA	0.941	148
PSA + EWMA	0.942	1,440




Fig. 10. Estimated edge probability matrices for two selected weeks of the Enron email network. Entry  $(i, j)$  denotes the estimated probability of an edge from class  $i$  to class  $j$ . Classes are as follows: (1) directors, (2) CEOs, (3) presidents, (4) vice-presidents, (5) managers, (6) traders, and (7) others. Notice the increase in the probability of edges from CEOs during the week of Skilling’s resignation.

combination of the EKF or PSA and EWMA predictors, we obtain a better link predictor that incorporates both block-level characteristics (through the EKF or PSA) and individual-level characteristics (through the EWMA). We evaluate the performance of the link predictors using the area under the receiver operating characteristic curve (AUC) metric. Since the PSA implementation we used accepts only undirected graphs, we reciprocate all edges to create undirected graphs for the link prediction experiment.

As shown in Table II, the dynamic SBM approaches combined with the EWMA all perform roughly comparably in terms of AUC, and all better than the EWMA alone. Notice that the a priori EKF adds hardly any computation time to the EWMA. If class memberships are not known in advance, one would have to use a posteriori methods. Both a posteriori methods perform roughly equally in terms of AUC, but our proposed EKF method is once again an order of magnitude faster than PSA.

2) *Temporal Dynamics*: Next we investigate the temporal variation of the state estimates from the EKF. Recall that the states  $\Psi^t$  correspond to the logit of the edge probabilities  $\Theta^t$ . We apply the a priori EKF to obtain the state estimates  $\hat{\Psi}^{t|t}$  and their variances (the diagonal of  $R^{t|t}$ ). Applying the logistic function, we can then obtain the estimated edge probabilities  $\hat{\Theta}^{t|t}$  with confidence intervals.

Examining the temporal variation of  $\hat{\Theta}^{t|t}$  reveals some interesting trends. For example, a large increase in the probabilities of edges from CEOs is found at week 89, in which CEO Jeffrey Skilling resigned. Inspection of the content of the emails sent during week 89 confirms Skilling’s resignation to be the cause of the increased probabilities. Fig. 10 shows a comparison of the matrix  $\hat{\Theta}^{t|t}$  during a normal week and during the week Skilling resigned.




Fig. 11. A priori EKF estimated edge probabilities  $\hat{\theta}_{ab}^{t|t}$  (solid lines) with 95% confidence intervals (shaded region) for (a) CEOs to presidents and (b) others to others by week in the Enron network. Estimated edge probabilities from CEOs to presidents peak at times corresponding to three major events (labeled). Edge probabilities between those in other roles increase only after Enron falls under federal investigation.




Fig. 12. Total number of emails sent each week in the Enron data set. Peaks are found around events 2 and 3 but not event 1.

Another interesting trend is highlighted in Fig. 11, where the temporal variation of two selected edge probabilities over the entire data trace with 95% confidence intervals is shown. Probabilities of edges from Enron CEOs to presidents show a steady increase as Enron's financial situation worsens, hinting at more frequent and widespread insider discussions, while probabilities of edges between others (not of one of the six known roles) begin to increase only after Enron falls under federal investigation. Notice also that the estimated edge probabilities from CEOs to presidents peak at three times that align with three major events during the Enron scandal [45]. A plot of the number of emails sent by week (Fig. 12) reveals peaks in email activity around events 2 and 3 but not around event 1 (CEO Skilling's resignation). Specifically, the overall volume of emails did not increase during the week Skilling resigned; only the volume of emails originating from CEOs increased, as we identified from Fig. 10. Indeed the temporal variation of the edge probabilities *between classes*, not the edge probabilities

across the entire network, is what reveals the internal dynamics of this time-evolving social network. Furthermore, the temporal model provides estimates with less uncertainty than one would obtain by fitting a static SBM to each time step, with 95% confidence intervals that are 25% narrower on average.

## V. CONCLUSION

This paper presented a statistical model for time-evolving networks that utilizes a set of unobserved time-varying states to characterize the dynamics of the networks. The model extends the well-known stochastic blockmodel for static networks to the dynamic setting and can be used for either a priori or a posteriori blockmodeling. We utilized a near-optimal on-line inference procedure based on the EKF that is much faster than an existing algorithm based on MCMC sampling yet shows comparable accuracy.

We applied the EKF-based inference procedure to the Enron email network and discovered some interesting trends when we examined the estimated states. One such trend was a steady increase in edge probabilities from Enron CEOs to presidents as Enron's financial situation worsened, while edge probabilities between other employees remained at their baseline levels until Enron fell under federal investigation. Furthermore, examining the temporal variation of edge probabilities between classes revealed a spike in edge probabilities that corresponded to the resignation of Enron CEO Jeffrey Skilling; this spike could not be found by simply examining the number of emails sent by week. The proposed procedure also showed promising results for predicting future email activity. We believe the proposed model and inference procedure can be applied to reveal the internal dynamics of many other time-evolving networks.

## ACKNOWLEDGMENT

We would like to thank Tianbao Yang for providing the source code for the probabilistic simulated annealing algorithm and Kevin Murphy for providing the Kalman filtering toolbox for MATLAB. We are also grateful to Greg Newstadt for his assistance in the implementation of the particle filters. Finally we thank Mark Kliger, Brandon Oselio, and the reviewers for their comments on the paper.

## REFERENCES


- [1] K. S. Xu and A. O. Hero, III, "Dynamic stochastic blockmodels: Statistical models for time-evolving networks," in *Proc. 6th Int. Conf. Soc. Comput. Behav.-Cult. Model. Predict.*, 2013, pp. 201–210.
- [2] B. A. Miller, N. T. Bliss, and P. J. Wolfe, "Toward signal processing theory for graphs and non-Euclidean data," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2010, pp. 5414–5417.
- [3] A. Y. Mutlu and S. Aviyente, "Dynamic network summarization using convex optimization," in *Proc. IEEE Workshop Stats. Signal Process.*, 2012, pp. 117–120.
- [4] X. Zhu and M. Rabbat, "Approximating signals supported on graphs," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2012, pp. 3921–3924.
- [5] A. Sandryhaila and J. M. F. Moura, "Discrete signal processing on graphs," *IEEE Trans. Signal Process.*, vol. 61, no. 7, pp. 1644–1656, Jul. 2013.
- [6] Y. Park, C. E. Priebe, and A. Youssef, "Anomaly detection in time series of graphs using fusion of graph invariants," *IEEE J. Sel. Top. Signal Process.*, vol. 7, no. 1, pp. 67–75, Jan. 2013.
- [7] A. Goldenberg, A. X. Zheng, S. E. Fienberg, and E. M. Airoldi, "A survey of statistical network models," *Found. Trends. Mach. Learn.*, vol. 2, no. 2, pp. 129–233, 2009.

- [8] J. S. Katz, "Scale-independent bibliometric indicators," *Meas.: Interdiscip. Res. Perspect.*, vol. 3, no. 1, pp. 24–28, 2005.
- [9] J. Leskovec, J. Kleinberg, and C. Faloutsos, "Graph evolution: Densification and shrinking diameters," *ACM Trans. Knowl. Discov. Data*, vol. 1, no. 1, p. 2, 2007.
- [10] P. J. Mucha, T. Richardson, K. Macon, M. A. Porter, and J.-P. Onnela, "Community structure in time-dependent, multiscale, and multiplex networks," *Science*, vol. 328, no. 5980, pp. 876–878, 2010.
- [11] D. Greene, D. Doyle, and P. Cunningham, "Tracking the evolution of communities in dynamic social networks," in *Proc. Int. Conf. Adv. Soc. Netw. Anal. Min.*, 2010, pp. 176–183.
- [12] K. S. Xu, M. Klinger, and A. O. Hero, III, "Adaptive evolutionary clustering," *Data Min. Knowl. Discov.*, vol. 28, no. 2, pp. 304–336, 2014.
- [13] P. W. Holland, K. B. Laskey, and S. Leinhardt, "Stochastic blockmodels: First steps," *Soc. Netw.*, vol. 5, no. 2, pp. 109–137, 1983.
- [14] S. Haykin, *Kalman Filtering and Neural Networks*. New York, NY, USA: Wiley-Interscience, 2001.
- [15] C. E. Priebe, J. M. Conroy, D. J. Marchette, and Y. Park, "Scan statistics on Enron graphs," *Comput. Math. Organ. Theory*, vol. 11, no. 3, pp. 229–247, 2005.
- [16] C. E. Priebe, J. M. Conroy, D. J. Marchette, and Y. Park, "Scan statistics on Enron graphs," 2009 [Online]. Available: <http://cis.jhu.edu/parky/Enron/enron.html>
- [17] B. Karrer and M. E. J. Newman, "Stochastic blockmodels and community structure in networks," *Phys. Rev. E*, vol. 83, no. 1, p. 016107, 2011.
- [18] K. Nowicki and T. A. B. Snijders, "Estimation and prediction for stochastic blockstructures," *J. Amer. Statist. Assoc.*, vol. 96, no. 455, pp. 1077–1087, 2001.
- [19] Y. Zhao, E. Levina, and J. Zhu, "Consistency of community detection in networks under degree-corrected stochastic block models," *Ann. Statist.*, vol. 40, no. 4, pp. 2266–2292, 2012.
- [20] K. Rohe, S. Chatterjee, and B. Yu, "Spectral clustering and the high-dimensional stochastic blockmodel," *Ann. Statist.*, vol. 39, no. 4, pp. 1878–1915, 2011.
- [21] D. L. Sussman, M. Tang, D. E. Fishkind, and C. E. Priebe, "A consistent adjacency spectral embedding for stochastic blockmodel graphs," *J. Amer. Statist. Assoc.*, vol. 107, no. 499, pp. 1119–1128, 2012.
- [22] F. Guo, S. Hanneke, W. Fu, and E. P. Xing, "Recovering temporally rewiring networks: A model-based approach," in *Proc. 24th Int. Conf. Mach. Learn.*, 2007, pp. 321–328.
- [23] P. Sarkar and A. W. Moore, "Dynamic social network analysis using latent space models," *ACM SIGKDD Explor. Newsl.*, vol. 7, no. 2, pp. 31–40, 2005.
- [24] P. Sarkar, S. M. Siddiq, and G. J. Gordon, "A latent space approach to dynamic embedding of co-occurrence data," in *Proc. 11th Int. Conf. Artif. Intell. Statist.*, 2007.
- [25] P. D. Hoff, "Hierarchical multilinear models for multiway data," *Comput. Statist. Data Anal.*, vol. 55, no. 1, pp. 530–543, 2011.
- [26] N. H. Lee and C. E. Priebe, "A latent process model for time series of attributed random graphs," *Statist. Inference Stoch. Process.*, vol. 14, no. 3, pp. 231–253, 2011.
- [27] P. O. Perry and P. J. Wolfe, "Point process modelling for directed interaction networks," *J. R. Statist. Soc.: Ser. B (Statist. Methodol.)*, vol. 75, no. 5, pp. 821–849, 2013.
- [28] E. P. Xing, W. Fu, and L. Song, "A state-space mixed membership blockmodel for dynamic network tomography," *Ann. Appl. Statist.*, vol. 4, no. 2, pp. 535–566, 2010.
- [29] Q. Ho, L. Song, and E. P. Xing, "Evolving cluster mixed-membership blockmodel for time-varying networks," in *Proc. 14th Int. Conf. Artif. Intell. Statist.*, 2011, pp. 342–350.
- [30] T. Yang, Y. Chi, S. Zhu, Y. Gong, and R. Jin, "Detecting communities and their evolutions in dynamic social networks—A Bayesian approach," *Mach. Learn.*, vol. 82, no. 2, pp. 157–189, 2011.
- [31] J. Durbin and S. J. Koopman, *Time Series Analysis by State Space Methods*, 2nd ed. Oxford, U.K.: Oxford Univ. Press, 2012.
- [32] L. Ljung, *System Identification: Theory for the User*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 1999.
- [33] S. J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd ed. Upper Saddle River, NJ, USA: Prentice-Hall, 2003.
- [34] A. T. Nelson, "Nonlinear estimation and modeling of noisy time-series by dual Kalman filtering methods," Ph.D. dissertation, Oregon Graduate Inst. of Sci. and Technol., Portland, OR, USA, 2000.
- [35] E. A. Wan and A. T. Nelson, "Dual Kalman filtering methods for nonlinear prediction, smoothing, and estimation," in *Adv. Neural Inf. Process. Syst.* 9, 1996, pp. 793–799.
- [36] Z. Ghahramani and S. T. Roweis, "Learning nonlinear dynamical systems using an EM algorithm," in *Adv. Neural Inf. Process. Syst.*, 1998, vol. 11.
- [37] F. Gustafsson and G. Hendeby, "Some relations between extended and unscented Kalman filters," *IEEE Trans. Signal Process.*, vol. 60, no. 2, pp. 545–555, Feb. 2012.
- [38] M. E. J. Newman and M. Girvan, "Finding and evaluating community structure in networks," *Phys. Rev. E*, vol. 69, no. 2, p. 026113, 2004.
- [39] L. Hubert and P. Arabie, "Comparing partitions," *J. Classif.*, vol. 2, no. 1, pp. 193–218, 1985.
- [40] N. Eagle, A. S. Pentland, and D. Lazer, "Inferring friendship network structure by using mobile phone data," in *Proc. Natl. Acad. Sci.*, 2009, vol. 106, no. 36, pp. 15 274–15 278.
- [41] T. Tylenda, R. Angelova, and S. Bedathur, "Towards time-aware link prediction in evolving social networks," in *Proc. 3rd Workshop Soc. Netw. Min. Anal.*, 2009, p. 9.
- [42] D. Liben-Nowell and J. Kleinberg, "The link-prediction problem for social networks," *J. Amer. Soc. Inf. Sci.*, vol. 58, no. 7, pp. 1019–1031, 2007.
- [43] C. Cortes, D. Pregibon, and C. Volinsky, "Computational methods for dynamic graphs," *J. Comput. Graph. Statist.*, vol. 12, no. 4, pp. 950–970, 2003.
- [44] K. S. Xu, M. Klinger, and A. O. Hero, III, "A shrinkage approach to tracking dynamic networks," in *Proc. IEEE Statist. Signal Process. Workshop*, 2011, pp. 517–520.
- [45] "ENRON: The Smartest Guys in the Room—Enron Timeline," 2007 [Online]. Available: <http://www.pbs.org/independentlens/enron/timeline.html>



**Kevin S. Xu** received the B.A.Sc. degree in electrical engineering from the University of Waterloo (2007) and the M.S.E. (2009) and Ph.D. (2012) degrees in electrical engineering: Systems from the University of Michigan. He is currently a researcher at the Technicolor Palo Alto Research Center. From 2012–2013, he was a Senior Research Scientist at the 3M Corporate Research Laboratory, St. Paul, MN. Kevin Xu was a recipient of the Natural Sciences and Engineering Research Council of Canada (NSERC) Postgraduate Master's and Doctorate Scholarships.

His main research interests are in statistical signal processing and machine learning with applications to network science and human dynamics. He is particularly interested in the development of statistical models and computational methods for analyzing time-dependent data in various forms including dynamic social networks and data collected from wearable sensors.



**Alfred O. Hero III** received the B.S. (summa cum laude) from Boston University (1980) and the Ph.D. from Princeton University (1984), both in electrical engineering. Since 1984 he has been with the University of Michigan, Ann Arbor, where he is the R. Jamison and Betty Williams Professor of Engineering. His primary appointment is in the Department of Electrical Engineering and Computer Science and he also has appointments, by courtesy, in the Department of Biomedical Engineering and the Department of Statistics. From 2008–2013 he was held the Digiteo Chaire d'Excellence at the Ecole Supérieure d'Electricité, Gif-sur-Yvette, France. He is a Fellow of the IEEE, and several of his research articles have received best paper awards. Alfred Hero was awarded the University of Michigan Distinguished Faculty Achievement Award (2011). He received the IEEE Signal Processing Society Meritorious Service Award (1998), the IEEE Third Millennium Medal (2000), and the IEEE Signal Processing Society Technical Achievement Award (2014). Alfred Hero was President of the IEEE Signal Processing Society (2006–2008) and was on the Board of Directors of the IEEE (2009–2011) where he served as Director of Division IX (Signals and Applications).

Alfred Hero's recent research interests are in statistical signal processing, machine learning and the analysis of high dimensional spatio-temporal data. Of particular interest are applications to networks, including social networks, multi-modal sensing and tracking, database indexing and retrieval, imaging, and genomic signal processing.