A Recommender System Based on Belief Propagation over Pairwise Markov Random Fields

Erman Ayday[†], Jun Zou[††], Arash Einolghozati[††] and Faramarz Fekri[††]

[†]School of Comp. and Comm. Sciences, Ecole Polytechnique Federale de Lausanne, Lausanne, Switzerland
[††]School of Electrical and Comp. Eng., Georgia Institute of Technology, Atlanta, GA 30332, USA

***Abstract*—Recommender systems enable service providers to predict and address the individual needs of their customers so as to deliver personalized experiences. In this paper, we formulate the recommendation problem as an inference problem on a Pairwise Markov Random Field (PMRF), where nodes representing items are connected with each other to exploit item-based similarity. However, direct prediction of ratings has exponential time complexity, as it requires to compute marginal probabilities. Thus, we utilize the Belief Propagation (BP) algorithm to solve the problem with a complexity that grows linearly with the number of items in the system. The BP algorithm computes marginal probabilities by performing iterative probabilistic message-passing between item nodes on the PMRF. One desirable feature of the proposed scheme is that it does not require to solve the problem for all users if it wishes to update the recommendations for only a single active user. Moreover, its complexity remains linear per active user. Via computer simulations using 100K MovieLens dataset, we verify that the proposed algorithm outperforms the MovieAvg and Pearson Correlation Coefficient (PCC) algorithms in terms of both Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE).**

I. Introduction

Today, the quantity of available information grows rapidly, overwhelming consumers to discover useful information and filter out the irrelevant items. The explosive growth of the Internet has made this issue increasingly more serious. Thus, the user is confronted with a big challenge of finding the most relevant information or item in the short amount of time. Without some support, the process of filtering out irrelevant items and finally selecting the most appropriate one could be very difficult. This problem gives rise to the need for developing effective ways in order to deal with information overload and access to relevant information. Recommender systems are aimed at addressing this overload problem, suggesting to the users those items that meet their interests and preferences the best in a particular situation and context. These systems are used to predict a user's evaluation of a particular item. More generally, recommender systems can learn about user preferences and profile over time, based on data mining algorithms, and automatically suggest products (from a large space of possible options) that fit the user needs.

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1115199, and a gift from the Cisco University Research Program Fund, an advised fund of Silicon Valley Community Foundation.

E. Ayday was a graduate student at Georgia Institute of Technology at the time of this work.

Currently, recommender systems are used in a variety of application domains, e.g., books, movies, and music. Despite some advances in recommender systems, still scalable and robust recommender systems are very much in need. On one hand, unfortunately, recommender systems have to operate on incomplete profiles because users either do not like to disclose lots of personal information and preferences, and/or are not completely aware about their preferences. On the other hand, with the dynamic and rapid growth of information flow, an increasing number of applications require recommender systems to make predictions considering the temporal dynamics of users and concepts. Hence, new research needs to focus on algorithms which meet these challenges and yet maintain computational efficiency for large-scale deployments.

In this work, we solve the recommender system's problem on graphical models using efficient Belief Propagation (BP) algorithms. By modeling the recommender system on a Pairwise Markov Random Field (PMRF), we build proper similarity associations between various items, using which we compute the marginal probability distribution functions of the variables representing the recommendations to be predicted. However, we observe that computing the marginal probability functions is computationally prohibitive for large scale recommender systems. Therefore, we propose to utilize the BP algorithm to efficiently (in linear complexity) compute these marginal probability distributions. The BP-based recommendation algorithm offers a significant advantage on scalability while providing competitive accuracy for the recommender systems.

The rest of this paper is organized as follows. In Section II we summarize the related work. In Section III, we describe PMRF formulation of rating prediction in detail and discuss its computational complexity. Then, in Section IV, we evaluate the performance of the proposed algorithm. Finally, in Section V, we conclude our paper.

II. Related Work

The major approaches used in recommender systems can be classified into two categories: the content-based approach and the collaborative filtering (CF) approach [1]. The content-based approach builds profiles for users and items, and makes recommendations by matching user profiles with item profiles. The CF approach exploits the past transactions and rating records, which can be further divided into model-based methods and memory-based methods. The model-based method generates rating predictions from a model learned from the previous ratings [2], but the learning process is often time-consuming, which is not suitable for systems with frequent updates. The memory-based method, or the neighborhood method, aggregates ratings either from other users of the same item, known as user-based methods [3], or from other items rated by the active user, known as item-based methods [4]. In case of the user-based method, the new rating r_{ui} on item i for user u is predicted as

$$\hat{r}_{ui} = \bar{r}_u + \frac{\sum_{v \in U_i} s_{uv}(r_{vi} - \bar{r}_v)}{\sum_{v \in U_i} |s_{uv}|}, \tag{1}$$

where r_u is the average rating of user u in the past, s_{uv} is the similarity between user u and v, and U_i denotes the set of users who have rated item i. When using the item-based method, r_{ui} is predicted by

$$\hat{r}_{ui} = \bar{r}_i + \frac{\sum_{j \in I_u} s_{ij}(r_{uj} - \bar{r}_j)}{\sum_{j \in I_u} |s_{ij}|}, \tag{2}$$

where $\bar{r}_i$ is the average rating of item i from previous users, s_{ij} is the similarity between item i and j, and I_u denotes the items rated by user u.

The similarity computation of s_{uv} or s_{ij} is a key component for neighborhood methods. There are several well-known computation methods including Pearson Correlation Coefficient (PCC) [4], which considers the difference in average rating of users or items. In the case of user similarity, PCC computes s_{uv} as

$$s_{uv} = \frac{\sum_{i \in I_{uv}}(r_{ui} - \bar{r}_u)(r_{vi} - \bar{r}_v)}{\sqrt{\sum_{i \in I_{uv}}(r_{ui} - \bar{r}_u)^2}\sqrt{\sum_{i \in I_{uv}}(r_{vi} - \bar{r}_v)^2}}, \tag{3}$$

where $I_{uv} = I_u \cap I_v$. Although PCC is simple and easy to implement, its accuracy is not satisfactory for modern recommender systems. Recently, a more advanced algorithm based on solving least squares problems has been proposed in [5]. One of the contribution of our approach is that we compute user similarities s_{uv} on a factor graph [6], where we assume the user similarities are defined on a set of discrete values, and use BP to infer their similarity levels.

Since the user-based method ignores the information from item similarity, and vice versa, the hybrid method is suggested in [7], [8] to combine both user similarity and item similarity. As we will see later, our proposed approach via PMRF is a hybrid solution.

Graphical formulations of CF recommender systems using random fields are also studied in [9], [10]. In [9], the authors integrate content-based and CF methods on a conditional Markov random field, where each vertex represents a rating, each edge connects two ratings on the same item or two ratings by the same user, and a local evidence node at each vertex encodes user and item profiles. In this setting, item ratings for all users are predicted simultaneously on a single graph. In [10], the authors predict item ratings for one active user on a single graph in each run, but they do not make use of user relations. All of the above works require extensive training to learn the graphical models.

In [11], [12], we presented recommender systems on a factor graph and introduced a BP-based approach to solve for ratings. This approach was motivated by the successful application of BP for solving reputation systems on factor graphs [13]–[15]. The promise of our work relies on the fact that Belief Propagation (BP) [16] acts as a powerful engine to operate on statistical data encoded on some forms of large graphical models. BP algorithms are very powerful to solve inference problems, at least approximately. The key role of the BP algorithm is that we can use it to compute the marginal probability distributions of the unknown variables (e.g., variables representing the recommendations to be predicted) in the complexity that grows only linearly with the total number of variables. But the factor graph solution, despite its good performance, only captures the user similarities. In this article, we extended our previous work to incorporate the item similarities as well by using PMRF. Our solution departs from previous PMRF methods as it relies on the iterative BP algorithm for solving the rating problem.

III. Formulation of Recommender Systems on PMRF

In our previous work on recommender systems [12], we viewed the recommender system as an inference problem, solved on a factor graph, based on the similarities between the users only. In the following, we consider a different type of inference to solve for the recommendation problem, which leads to a better performance. As an alternative to our previous approach, we propose to formulate the recommender system by utilizing both user similarities and item similarities. Thus, we give a new formulation of the recommender system problem as finding the marginal probability distributions of the unknown variables on a PMRF which embeds the similarity relations between items as well. In the following, we describe how the proposed algorithm formulates the recommender system on a PMRF and computes the marginal probability distributions of the unknown variables using BP.

We assume two different sets in the system: i) the set of users $\mathbb{U}$ and ii) the set of items (products) $\mathbb{I}$. Users provide feedbacks, in the form of ratings, about the items for which they have an opinion. The main goal is to provide accurate recommendations for every user by predicting the ratings of the user for the items that he/she did not rate before (unseen item). Here, we consider an arbitrary user z (referred to as the active user) and compute the prediction of ratings for user z for unseen items to describe the algorithm. We assume M users and N items in the system (i.e., $|\mathbb{U}| = M$ and $|\mathbb{I}| = N$). Let $\mathbb{R}_z = \{r_{zi} : i \in \mathbb{I} \backslash I_z\}$ be the collection of variables representing the ratings of the items to be predicted[1] for the

[1] A subset of these variables are already known as the corresponding items were rated by user z. Hence, they do not require any prediction.

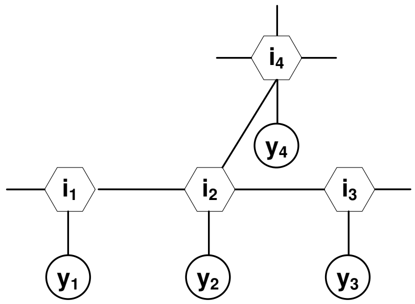

Fig. 1: Representation of item ratings on PMRF.

active user z, where I_z denotes the rated items by z. To be consistent with the most of existing recommender systems, we assume that the rating values are integers from the set $\Upsilon = \{1, 2, 3, 4, 5\}$.

Our objective is to formulate the recommendation problem as making statistical inference about the ratings of users for unseen items based on observations. That is, given the past data evidence, what would be the likelihood (probability) that the rating takes a particular value? Here, the probability is the degree of belief to which the prediction of the rating is supported by the available evidence. This requires finding the marginal probability distributions of the variables in $\mathbb{R}_z$ conditioned on some observed preferences.

However, the complexity of direct computing these marginal probability distributions grows exponentially with the number of variables, which is computationally prohibitive for large-scale recommender systems. Thus, we utilize the BP algorithm to estimate these marginal functions via message passing between the items on PMRF with complexity linear in the number of items.

A. PMRF Representation of Recommender Systems

A PMRF consists of hidden and observed nodes [17], in which the statistical dependencies between two connected hidden nodes are represented by compatibility functions, and the dependency relations between the observed nodes and hidden nodes are represented by local evidence functions. We formulate the recommender system on a PMRF as in Fig. 1. For each item, whose rating is to be predicted, we assign a hidden node, or item node, shown as a hexagon, which takes values from a set of discrete ratings in Υ. Further, to incorporate local evidence, i.e., ratings from other users, we connect each item node i to a unique local evidence node y_i (observed node) shown as a circle, which represents the average item rating from other users weighted by the user similarities between other users and the active user as

$$y_i = \bar{r}_z + \frac{\sum_{v \in \hat{U}_i} s_{zv}(r_{vi} - \bar{r}_v)}{\sum_{v \in \hat{U}_i} |s_{zv}|}, \tag{4}$$

where $\hat{U}_i$ denotes a set of K_z users that are most similar to the active user z among users who have rated item i. Further, the user similarity s_{zv} is computed on a factor graph, the details of which is deferred to [18] due to page limit.

We use an edge to connect two items i and j if and only if

- The similarity s_{ij} between them is above a predefined threshold σ_{th}, and
- At least one of them is among the K most similar items of the other, if not both.

Here, we compute item similarities using the adjusted cosine similarity measure [4] as below:

$$s_{ij} = \frac{\sum_{v \in U_{ij}} (r_{vi} - \bar{r}_v)(r_{vj} - \bar{r}_v)}{\sqrt{\sum_{v \in U_{ij}} (r_{vi} - \bar{r}_v)^2} \sqrt{\sum_{v \in U_{ij}} (r_{vj} - \bar{r}_v)^2}}, \tag{5}$$

where $U_{ij} = U_i \cap U_j$.

We define a local evidence function $\phi_i(r_{zi})$ to capture the relations between item i and its local evidence node y_i. $\phi_i(r_{zi})$ is a probability distribution over possible ratings in Υ, and it satisfies that

- $\sum_{r \in \Upsilon} \phi_i(r_{zi} = r) = 1$, and
- $\sum_{r \in \Upsilon} r\phi_i(r_{zi} = r) = y_i$.

There are many potential choices of designing $\phi_i(r_{zi})$. In this work, we distribute the probability on the integer ratings in Υ close to y_i, i.e.,

- If $1 \le y_i \le 5$
$$\lfloor y_i \rfloor \phi_i(r_{zi} = \lfloor y_i \rfloor) + \lceil y_i \rceil \phi_i(r_{zi} = \lceil y_i \rceil) = y_i,$$
- If $y_i < 1$, $\phi_i(r_{zi} = 1) = 1$,
- If $y_i > 5$, $\phi_i(r_{zi} = 5) = 1$.

The local evidence function represents how one item should be rated according to the local evidence in (4). We also define a compatibility function between two connected item nodes i_a and i_b as follows

$$\psi_{ab}(r_{zi_a}, r_{zi_b}) = \frac{1}{Z_{ab}} \exp\left\{ -\frac{(r_{zi_a} - r_{zi_b})^2}{\sigma_{ab}^2} \right\} \tag{6}$$

where σ_{ab} is adjusted according to item similarity computed using (5). The compatibility function penalizes the incompatible ratings from the connected items if they deviate from each other. Moreover, the compatibility between the ratings of connected items can be adjusted according to item similarity. The overall joint distribution of all item ratings $r_{zi}, i \in \mathbb{I} \backslash I_z$, on the PMRF is defined by

$$P(\{r_{zi}, i \in \mathbb{I} \backslash I_z\}) = \frac{1}{Z} \prod_{i_a, i_b} \psi_{ab}(r_{zi_a}, r_{zi_b}) \prod_i \phi_i(r_{zi}), \tag{7}$$

where Z is a normalization constant. To compute item ratings, we first compute the marginal probabilities. However, direct computation incurs an exponential complexity. We therefore apply the BP algorithm to infer marginal probabilities.

B. Probabilistic Message-Passing

The message-passing on PMRF is illustrated in Fig. 2. The message $m_{a,b}(r_{zb})$ from item i_a to its connected item i_b is

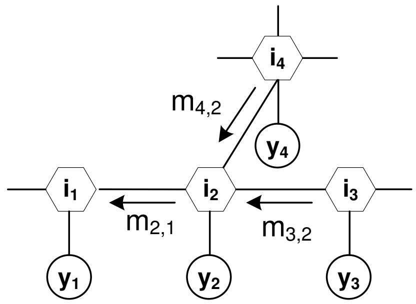

Fig. 2: Message exchange between item nodes on PMRF.

given by

$$m_{a,b}(r_{zi_b}) = \frac{1}{Z_{a,b}} \sum_{r_{zi_a}} \psi_{ab}(r_{zi_a}, r_{zi_b}) \phi_{i_a}(r_{zi_a}) \prod_{c \in N_a \setminus b} m_{c,a}(r_{zi_a}) \quad (8)$$

where $Z_{a,b}$ is some normalization constant, and N_a denotes the set of neighbor item nodes connected to node i_a. The message in (8) represents how one item should be rated according to the items it is connected to in PMRF. The m-messages are iteratively exchanged between item nodes. During each iteration, the outgoing messages at each item node are updated using incoming messages from the last iteration. Finally, every item i calculates its predicted rating value as the expectation of the marginal probability distribution. We summarize the BP algorithm for rating predictions in Algorithm 1.

Algorithm 1 Belief Propagation Algorithm over PMRF

- Initialization. Initialize $m_{a,b}(r_{zi_b})$ as $m_{a,b}^{(0)}(r_{zi_b}) = \frac{1}{|\Upsilon|}$, and set iteration counter $n = 1$.
- Iterative message-passing until convergence.
  - (1) Update $m_{a,b}(r_{zi_b})$ at all item nodes:

    $$m_{a,b}^{(n)}(r_{zi_b}) = \frac{1}{Z_{a,b}} \sum_{r_{zi_a}} \psi_{ab}(r_{zi_a}, r_{zi_b}) \phi_{i_a}(r_{zi_a}) \prod_{c \in N_a \setminus b} m_{c,a}^{(n-1)}(r_{zi_a})$$

  - (2) $n = n + 1$. Repeat (1) until convergence.
- Compute marginal item rating probability:

  $$P(r_{zi}) = \frac{1}{Z_i} \phi_i(r_{zi}) \prod_{j \in N_i} m_{j,i}(r_{zi}),$$

  where Z_i is some normalization constant.
- Predict item rating as the expectation of the marginal probabilities:

  $$\bar{r}_{zi} = \sum_{r \in \Upsilon} r P(r_{zi} = r).$$

We further note that the beliefs are exact if the PMRF has no loops. However, the graphical model for recommender systems has cycles. Due to the statistical dependencies between the messages in this loopy case, the marginal probabilities of the variables in $\mathbb{R}_z$ are approximations rather than exact values. However, it is shown that BP often works well even in the loopy graphs [17].

C. Complexity Analysis

For each active user, the complexity for computing user similarities between other users and the active user on a factor graph is linear in the number of total users [18]. Then to predict all item ratings on a PMRF using BP, the complexity in terms of multiplications is $\mathcal{O}(NK^2)$, where K is the item neighborhood size on PMRF, which is a fixed parameter independent of M and N. Thus, the complexity is linear in the total number of items, N. Further, the proposed algorithm converges quickly, on the average in 10 iterations. Hence, we did not include the number of iterations in the complexity measure as it only introduces a small constant factor in the total complexity. We therefore conclude that the proposed algorithm is scalable, which is critical for its successful application to large-scale systems.

IV. Evaluation

We evaluated the performance of the proposed iterative PMRF-based algorithm using the 100K MovieLens dataset[2]. The dataset contains $100,000$ ratings from 943 users on 1682 items (movies) in which each user has rated at least 20 items. Further, the rating values are integers from 1 to 5. We randomly divided the users into two disjoint set, 80% for training and 20% for testing. Specifically, for each user in the test set, we keep a subset of its ratings as known (15 ratings in our setup), and remove the rest of its ratings, which are used later for evaluation purpose. We implemented all the experiments via MATLAB on a 2 GHz PC with 4 GB RAM.

We evaluated the rating prediction accuracy of the proposed algorithm in terms of both Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) metrics over the predicted ratings. The RMSE metric is more sensitive to large errors than MAE. We compute the MAE and RMSE as below:

$$\text{MAE} = \frac{1}{N_t} \sum_{r_{ui} \in \mathbb{R}_t} |r_{ui} - \bar{r}_{ui}|,$$

$$\text{RMSE} = \sqrt{\frac{1}{N_t} \sum_{r_{ui} \in \mathbb{R}_t} (r_{ui} - \bar{r}_{ui})^2},$$

where $\mathbb{R}_t$ is the set of all test ratings for users in the test dataset, N_t is the total number of test ratings, r_{ui} is the actual value of the rating provided by user u on the item i in the test dataset, and $\bar{r}_{ui}$ is the predicted rating value by the algorithm.

Here, we summarize the parameters of the proposed algorithm used in our evaluation. To compute user similarities on

[2]Available at: http://www.grouplens.org/node/73.

TABLE I: Prediction Accuracy Performance

Algorithm	MAE		RMSE	
	$K_z = 10$	$K_z = 30$	$K_z = 10$	$K_z = 30$
MovieAvg	0.7920		0.9913	
PCC	0.7622	0.7484	0.9688	0.9550
Proposed	0.7261	0.7224	0.9451	0.9380

a factor graph, we assume the user similarities are defined on $\mathcal{S} = \{1, 2, 3, 4, 5\}$. Further, we set $D = 4$ and $\sigma = 0.5$, where D is the degree at each item node in the complexity-reduction factor graph, and σ is used to control the sharpness of the check constraint function at factor nodes as discussed in [18]. To predict item ratings using BP over the PMRF, we set the parameters as $K = 5$ and $s_{th} = 0.35$, and adjust the compatibility between two connected items as follows

$$\sigma_{ab} = \frac{1}{\sqrt{2}} \left(\frac{1 - s_{i_a i_b}}{1 - s_{th}} + 1 \right).$$

We compare our algorithm with the PCC method [4] and the MovieAvg algorithm (which computes the average rating of each item). The accuracy results are presented in Tab. I, where K_z is the number of similar users used to compute local evidence in (4). The achieved accuracy improvement of our algorithm over PCC is approximately 4% in MAE and 2% in RMSE. The state-of-the-art neighborhood method in [5] achieves $2 \sim 3\%$ improvement in RMSE over PCC on Netflix data. So roughly speaking, our algorithm has competitive performance compared to the existing work. Also note that, the work in [5] needs to solve a $K \times K$ system of equations for each rating prediction, where K is the neighborhood size.

V. Conclusion

In this paper, we solve the recommender system problem using Belief Propagation (BP) algorithm on a Pairwise Markov Random Field (PMRF). The proposed approach formulates the recommendation problem as making statistical inference about the ratings of users for unseen items based on observations. Then all item ratings are represented on a PMRF, where similar items are connected with each other to exploit item similarity to further improve accuracy. The BP algorithm is utilized to efficiently estimate the marginal probability distributions of the item ratings via iterative message-passing between item nodes. The complexity of the proposed algorithm remains linear per active user, making it very attractive for large-scale systems. Further, it can update the recommendations for each active user instantaneously using the most recent data (ratings) and without solving the recommendation problem for all users. Via computer simulations using the 100K MovieLens dataset, we showed that the proposed algorithm achieves improved accuracy over MovieAvg and PCC algorithms. Therefore, we claim that the BP-based approach on PMRF for recommendation problem is very promising and can result in a new class of accurate and scalable recommender systems.

References

[1] G. Adomavicius and A. Tuzhilin, "Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions," *IEEE Transactions on Knowledge and Data Engineering*, vol. 17, pp. 734–749, Jun. 2005.

[2] D. Billsus and M. Pazzani, "Learning collaborative information filters," *Proc. International Conf. Machine Learning (ICML)*, 1998.

[3] P. Resnick, N. Iakovou, M. Sushak, P. Bergstrom, and J. Riedl, "Grouplens: An open architecture for collaborative filtering of netnews," *Proc. 1994 Computer Supported Cooperative Work Conf.*, 1994.

[4] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl, "Item-based collaborative filtering recommendation algorithms," *Proc. 10th International WWW Conf.*, 2001.

[5] R. Bell and Y. Koren, "Scalable collaborative filtering with jointly derived neighborhood interpolation weights," *Proc. of Seventh IEEE International Conference on Data Mining*, pp. 43–52, 2007.

[6] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.

[7] J. Wang, A. D. Vries, and M. Reinders, "Unifying user-based and item-based collaborative filtering approaches by similarity fusion," *Proc. of ACM SIGIR*, pp. 501–508, 2006.

[8] H. Ma, I. King, and M. Lyu, "Effective missing data prediction for collaborative filtering," *Proc. of ACM SIGIR*, pp. 39–46, 2007.

[9] T. T. Truyen, D. Q. Phung, and S. Venkatesh, "Preference networks: Probabilistic models for recommendation systems," *Proc. of Sixth Australasian Data Mining Conference (AusDM 2007)*, pp. 191–198, 2007.

[10] A. J. Defazio and T. S. Caetano, "A graphical model formulation of collaborative filtering neighbourhood methods with fast maximum entropy training," *Proc. of the 29th International Conference on Machine Learning (ICML)*, 2012.

[11] E. Ayday and F. Fekri, "A belief propagation based recommender system for online services," *Proceedings of the fourth ACM conference on Recommender systems*, pp. 217–220, 2010.

[12] E. Ayday, A. Einolghozati, and F. Fekri, "BPRS: belief propagation based iterative recommender system," *Proceedings of ISIT '12: IEEE International Symposium on Information Theory*, 2012.

[13] E. Ayday and F. Fekri, "BP-ITRM: belief propagation for iterative trust and reputation management," *Proceedings of ISIT '11: IEEE International Symposium on Information Theory*, 2011.

[14] ——, "BP-P2P: a belief propagation-based trust and reputation management for p2p networks," *Proceedings of IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks - SECON*, 2012.

[15] ——, "Belief propagation based iterative trust and reputation management," *IEEE Transactions on Dependable and Secure Computing*, vol. 9, no. 3, pp. 375–386, May/June 2012.

[16] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers, Inc., 1988.

[17] J. S. Yedidia, W. T. Freeman, and Y. Weiss, "Understanding belief propagation and its generalizations," *In Exploring artificial intelligence in the new millennium*, pp. 239–269, 2003.

[18] J. Zou, A. Einolghozati, E. Ayday, and F. Fekri, "Iterative similarity inference via factor graphs for collaborative filtering," [on-line] available at http://users.ece.gatech.edu/jzou8/publications/Zou_CF_2012.pdf.