

LabelRankT: Incremental Community Detection in Dynamic Networks via Label Propagation

Jierui Xie
Rensselaer Polytechnic
Institute
Troy, New York 12180, USA
jierui.xie@gmail.com

Mingming Chen
Rensselaer Polytechnic
Institute
Troy, New York 12180, USA
chenm8@rpi.edu

Boleslaw K. Szymanski
Rensselaer Polytechnic
Institute
Troy, New York 12180, USA
szymab@rpi.edu

ABSTRACT

An increasingly important challenge in network analysis is efficient detection and tracking of communities in dynamic networks for which changes arrive as a stream. There is a need for algorithms that can *incrementally* update and monitor communities whose evolution generates huge real-time data streams, such as the Internet or on-line social networks. In this paper, we propose *LabelRankT*, an on-line distributed algorithm for detection of communities in large-scale dynamic networks through stabilized label propagation. Results of tests on real-world networks demonstrate that *LabelRankT* has much lower computational costs than other algorithms. It also improves the quality of the detected communities compared to dynamic detection methods and matches the quality achieved by static detection approaches. Unlike most of other algorithms which apply only to binary networks, *LabelRankT* works on weighted and directed networks, which provides a flexible and promising solution for real-world applications.

Keywords

social network, community detection, clustering, network evolution, dynamic network, temporal

1. INTRODUCTION

Communities are the basic structures in sociology in general and in social networks in particular. They have been intensively researched for more than a half of the century [17]. In sociology, community usually refers to a social unit whose members share common values and the identity of the members as well as their degree of cohesiveness depend on individuals' social and cognitive factors such as beliefs, preferences, or needs. The ubiquity of the Internet and social media eliminated spatial limitations on community range, enabling on-line communities to link people regardless of their physical location. The newly arising *computational sociology* relies on computationally intensive methods to analyze and model social phenomena [3], including communities and their detection.

Analysis of social networks became one of the basic tools of sociology [24] and has been used for linking micro and macro levels of sociological theory. The classical example of the approach is presented in [10] that elaborated the macro implications of one aspect of small-scale interaction, the strength of dyadic ties. Moreover, a lot of commercial applications, such as digital marketing, behavioral targeting and user preference mining, rely heavily on community analysis.

With the rapid growth of large-scale on-line social networks, e.g., Facebook connected a billion users in 2012, there is a high demand for efficient community detection algorithms that will be able to handle their evolution growth. Communities in on-line social networks are discovered by analyzing the observed and often recorded on-line interactions between people.

Numerous techniques have been developed for community detection. However, most of them require a *global* and often *static* view of the network and ignore temporal correlations between different snapshots over time. Such algorithms are not scalable enough to cope with dynamically evolving networks, especially when new data about them are generated continuously. Another limitation shared by most of the existing algorithms is that they are applicable only to networks with binary adjacency matrix, that is with undirected and unweighted edges.

Label propagation based community detection algorithms such as LPA [18, 26], COPRA [11] and SLPA¹ [27] have been shown to perform well in static networks. However, due to random tie breaking strategy, they produce different partitions in different runs. Such instability is highly undesirable when tracking the evolution of communities in a dynamic network.

The contributions of this paper are two-fold. First, we generalized the *LabelRank* algorithm introduced in [28] to incorporate important network features such as edge weights and directions. Second, built upon *LabelRank*, we introduce *LabelRankT* algorithm that incrementally detects evolving communities in dynamic networks. The new algorithm presented here delivers significant improvements over the existing solutions in both the quality of detected evolving communities and the speed of program execution.

¹Source codes: <https://sites.google.com/site/communitydetectionslpa/>

Algorithm 1 Generalized *LabelRank*

```
1: add self-loop to adjacency matrix  $W$ 
2: initialize label distribution  $P$  using Eq. 3
3: repeat
4:    $P' = W \times P$ 
5:    $P' = \Gamma_{in} P'$ 
6:    $P' = \Phi_r P'$ 
7:    $P = \Theta_q(P', P)$ 
8: until stop criterion satisfied
9: output communities
```

2. GENERALIZATION OF LABELRANK

To make the paper self-contained, this section first summarizes and generalizes the *LabelRank* algorithm introduced in [28] which is the basis for *LabelRankT*. Both algorithms are based on the idea of simulating the propagation of labels in the network.

LabelRank relies on four operators applied to the labels: (i) propagation, (ii) inflation, (iii) cutoff, and (iv) conditional update. The data structure that lies in the core of this algorithm is the sparse matrix of label distribution. Each node maintains a label distribution locally during the propagation. At the end of the algorithm, *LabelRank* ranks labels in each node. Nodes with the same highest probability label form a community.

LabelRank, a stabilized LPA, was initially introduced in [28] for binary networks. Here, we generalize the propagation operation over edges in order to take both edge direction and weight into account. Given a network $G = (E, V)$, where E is the set of edges and V is the set of nodes, this operator can be expressed in matrix form as:

$$W \times P, \quad (1)$$

where W is the $n \times n$ weight matrix and n is the number of nodes. For each w_{ij} , if there exist a directed edge $e_{ji} \in E$ from node j to i , w_{ij} takes on a positive value; otherwise it is 0. $w_{ij} > 0$ is the weight placed on the directed edge, which bears important and application specific information. P is the $n \times n$ label distribution matrix. Label is just a unique identifier. For simplicity, it usually takes on the same value as the node ID. P is composed of n $1 \times n$ row vectors P_i , one for each node i . Each element P_{ic} or $P_i(c)$ holds the current estimation of probability of node i observing label $c \in C$ taken from a finite set of alphabet C (here $|C| = n$).

Each node broadcasts the distribution to its neighbors at each iteration step and computes the new distribution P'_i simultaneously using the following equation:

$$P_{ic} = \frac{\sum_{j \in Nb(i)} w_{ij} P_{jc}}{\sum_{k \in Nb(i)} w_{ik}}, \forall c \in C, \quad (2)$$

where $Nb(i)$ is a set of neighbors of node i and the numerator sums up weights of all edges incoming node i . That is, a node sends out its information along outgoing edges to its neighbors and at the same time receives information along incoming edges from the neighborhood. Note that, P'_i is normalized to make proper probability distribution. One can show that the new distribution vector P'_i is the distribu-

tion that minimizes the KL divergence between any possible P'_i and P_i [22].

To initialize P , each node is assigned a distribution of probabilities of all incoming edges by assigning to each incoming edge the initial probability of seeing this neighbor's label proportional to the weight of this edge:

$$P_{ij} = \frac{w_{ij}}{\sum_{k \in Nb(i)} w_{ik}}, \forall j \in C \text{ s.t. } w_{ij} > 0. \quad (3)$$

We briefly characterize the three remaining operations of *LabelRank* as follows.

(1) The inflation operator Γ_{in} on P [23] is used to contract the propagation, where in is the parameter taking on real values. It operates on the label distribution matrix P (rather than to a stochastic matrix or adjacency matrix) to decouple it from the network structure. After applying $\Gamma_{in} P$, each P_{ic} is proportional to P_{ic}^{in} , i.e., P_{ic} rises to the in^{th} power.

(2) The full label propagation distribution induces a cost for memory. To alleviate this problem, the cutoff operator Φ_r on P is introduced to remove labels that are below threshold $r \in [0, 1]$. More importantly, Φ_r is shown empirically to reduce the space complexity efficiently, from quadratic to linear. The average number of labels in each node is typically less than 3.0 for $r = 0.1$.

(3) The conditional update Θ_q operator is used to trap the process in the quality space (e.g., modularity [14]) to avoid trivial network state where each node holds the same distribution. The algorithm updates a node only when it is significantly different from its neighbors in terms of labels. This allows us to preserve detected communities and detect termination based on scarcity of changes to the network. At each iteration, the change is accepted only by nodes that satisfy the following update condition:

$$\sum_{j \in Nb(i)} isSubset(C_i^*, C_j^*) \leq qk_i, \quad (4)$$

where C_i^* is the set of maximum labels which includes labels with the maximum probability at node i at the previous time step. Function $isSubset(s_1, s_2)$ returns 1 if $s_1 \subseteq s_2$, and 0 otherwise. k_i is the degree of node i , and q is a real number parameter chosen from the interval $[0, 1]$. Intuitively, $isSubset$ can be viewed as a measure of similarity between two nodes.

These four operators together with a post-processing that groups nodes whose highest probability labels are the same into a community form our algorithm (see Alg. 1).

The running time of the generalized *LabelRank* is the same as the original algorithm, as it is $O(m)$, linear with the number of edges m . The space complexity is $O(n)$ in practice because the number of labels in each node monotonically decreases and drops to a small constant in a few steps due to both cutoff and inflation operators. The P matrix is replaced by sparse matrix representation of n variable-length list of pairs (usually short) carried by each node; each pair contains label and its probability (with labels whose probability reduced to 0 not listed).




Figure 1: The example network $G(0)$ with $n = 15$. Colors represent communities discovered by *LabelRankT*.




Figure 2: The example network $G(1)$ with $n = 15$ that change from $G(0)$ by splitting and merging. Nodes 10 and 14 moved from blue group into the red group. Three edges were deleted from and three were added. Colors represent communities discovered by *LabelRankT* after these events.




Figure 3: The example network $G(2)$ with $n = 15$ in which node 6 was removed and node 16 was born. The green group dissolved and its members merged into blue and red groups. Colors represent communities discovered by *LabelRankT*.

Algorithm 2 *LabelRankT*

- 1: input: snapshots $G([0, 1, \dots, T])$
 - 2: **for** $t=1:T$ **do**
 - 3: (a) Tracking the *changed* nodes in $G(t)$ due to the changes in edges they attach to since $t - 1$.
 - 4: (b) Initialize P^t . For node i that does not change since $t - 1$, we copy its label distributions, i.e., $P_i^t = P_i^{t-1}$. For *changed* nodes, we reinitialize their label distributions as in *LabelRank*.
 - 5: (c) Iteratively update only *changed* nodes' label distribution and assign them to the corresponding communities as in *LabelRank*.
 - 6: **end for**
-

Unlike in binary networks, there are various ways of adding a self-loop to each node to stabilize the results. It is interesting to see that the algorithm might perform slightly differently with different ways of defining self-loop. The most common ways include setting $w_{ii} = 1$, $w_{ii} = \max(w_{ik})$ or $w_{ii} = \sum_k w_{ik}$. In the experiments run for this paper, we use $w_{ii} = 1$. Yet, it is still an open question how to optimally select the selfloop weight for each node.

3. LABELRANKT: AN EXTENSION FOR DYNAMIC NETWORKS

The extended algorithm called *LabelRankT* is based on the generalized *LabelRank* introduced in Alg. 1. The description of *LabelRankT* is contained in Alg. 2. The main idea is to adjust our detection as the network structure changes. We take advantage of what we already obtain in previous snapshot for inferring the dynamics in the current time step. Since local structure information is encoded in the node label distributions, the evolving of communities is expected to be caught and reflected in these distributions.

LabelRankT can be viewed as a *LabelRank* with one extra conditional update rule by which only nodes involved any change accept the new distribution. Moreover, we only need to update nodes that are changed between two consecutive snapshots, including cases where an existing node adds or deletes links, or a node is removed from or newly joins the network. An example that shows different evolution events in three consecutive snapshots, $G(0)$, $G(1)$ and $G(2)$ is shown in Figs. 1, 2 and 3. During the evolution, nodes (edges) are added or removed, and communities split, merge and dissolve, all of which is captured by *LabelRankT*. To discover communities in these snapshots, we ran our algorithm with the same parameters for all three snapshots.

In our algorithm, all these cases are handled by simply comparing neighbors of a node i at two consecutive steps, $t - 1$ and t , i.e., $Nb^t(i)$ and $Nb^{t-1}(i)$. If $Nb^t(i)$ and $Nb^{t-1}(i)$ are not equal, then node i is called a *changed node*². For changed nodes, we reinitialize their label distributions (i.e., P^t) and update until the simulation stops as in *LabelRank*. However, since only changed nodes and their neighbors are involved (some neighbors only propagate labels but not update), *LabelRankT* is more efficient than *LabelRank*.

²Since it is often a case in practice, we assume here that all nodes in all steps are uniquely and consistently named.




Figure 4: AS-Internet Routers Graph. Degree distributions at the beginning, middle and the end of evolution (main plot). Average degree over time (inset).




Figure 5: The structure changes in AS-Internet Routers Graph, including the number of edges added (E_+) and deleted (E_-), as well as the number of nodes involved in changes (N_{+-}).

Such conditional update rule makes *LabelRankT* applicable to dynamic networks for which changes arrive as a stream. When a new edge arrives in the incoming stream, *LabelRankT* only updates the nodes that are attached to this edge. Thus, *LabelRankT* can efficiently update and monitor communities whose evolution generates huge real-time data streams, such as the Internet and on-line social networks.

The time complexity for *LabelRankT* can be derived as follows. It is easy to see for step (a) that we can track the number of changed nodes in $O(1)$. For steps (b) and (c), since at most n can be changed, we need to communicate across each edge at most twice in each iteration. Since the number of iterations required is a constant T (usually less than 50 iterations), the overall complexity for detecting evolving communities between two consecutive snapshots is $O(Tm)$, implying $O(m)$ in general.

4. ALGORITHM PERFORMANCE EVALUATION

We tested the detection quality of *LabelRankT* in terms of modularity and the efficiency in two real-world datasets.

AS-Internet Routers Graph [12]. This is a communication network of who-talks-to-whom from the Border Gateway Protocol logs of routers in the Internet. The dataset contains 733 daily snapshots for 785 days from November 8 1997 to January 2 2000. The number of nodes in the largest snapshot is 6,477 (with 13,233 edges). The nodes and edges are added or removed over time. The structure at each snapshot could change dramatically as indicated by fluctuations in the average degree in Fig. 4 and structure change statistics in Fig. 5.

arXiv HEP-TH. High energy physics theory citation graph is from the e-print arXiv and covers all the citations within a dataset of 27,769 papers with 352,285 directed edges, each indicating that a paper at its tail cites the paper at its head. The data covers papers from January 1993 to April 2003 [12]. Unlike AS networks, it grows over time (i.e., edges and nodes are added but not removed).

The dataset is separated into snapshots by week (a total of 359 snapshots). The number of nodes in snapshots ranges from 12,917 to 27,769, while the number of edges ranges from 47,454 to 352,285. The structures over time are similar with a monotonic increase in average degree varying from 8 to 26 (see Fig. 6). The statistics of structure change over time are shown in Fig. 7.

We first compared the performance of *LabelRankT* with static algorithms MCL, using for all snapshots parameters that optimized performance for the first snapshot, and Infomap. Both run through each snapshot independently. Since a dynamic (especially incremental) algorithm like *LabelRankT* does not recompute the entire network, static algorithms might perform better. In fact, on AS Graph, see Fig. 8, three algorithms actually have close performance. Infomap slightly outperforms *LabelRankT* by about 5.03% in modularity on average, *LabelRankT* and MCL performance differs just by 0.43%. On arXiv HEP-TH (which is of much larger size than AS), as seen in Fig. 9, Infomap and *LabelRankT* perform within 0.88% of each other. However, *LabelRankT* outperforms MCL significantly by 15.37% (Note that the behavior of MCL is partially caused by its sensitivity to parameters). On the other hand, *LabelRankT* has benefit of efficiency. It runs 4 and 12 times faster than Infomap on AS Graph and arXiv HEP-TH respectively. And it is faster than MCL by a factor of 27 to 52 on AS Graph and arXiv HEP-TH, respectively.

We also compared *LabelRankT* with two publicly available dynamic algorithms that employ incremental detection methods : facetNet³ [13] and iLCD⁴ [7]. On AS Graph, see Fig. 10, facetNet and *LabelRankT* achieve performance within

³Since facetNet requires the number of communities as input, we assign it the value produced by *LabelRankT*.

⁴After detection, if a node belongs to more than one community, we assign it to the one with maximum size to be able to output only unique and disjoint partitions.

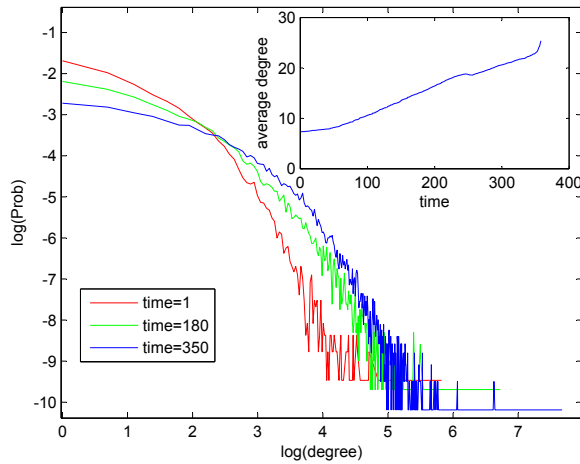


Figure 6: arXiv HEP-TH. Degree distributions at the beginning, middle and the end of evolution (main). Average degree over time (inset).




Figure 7: The structure changes in arXiv HEP-TH, including the number of edges added (E_+) and deleted (E_-), as well as the number of nodes involved in changes (N_{+-}).

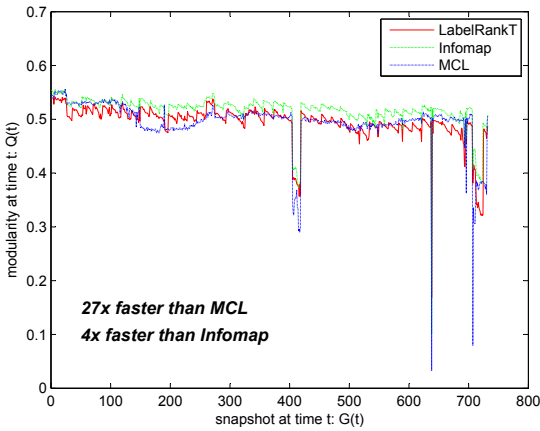


Figure 8: Comparison of modularity over time $Q(t)$ with static detection algorithms on AS-Internet Routers Graph.




Figure 9: Comparison of modularity over time $Q(t)$ with static detection algorithms on arXiv HEP-TH.

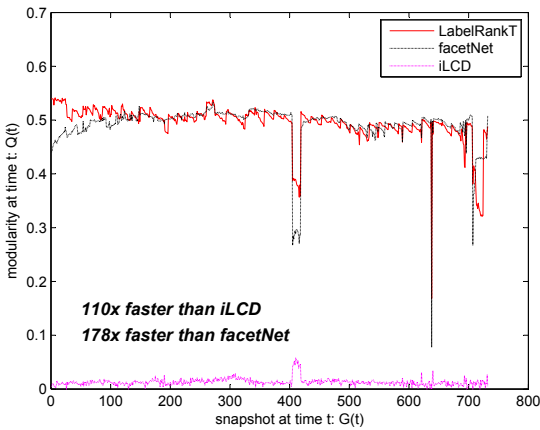


Figure 10: Comparison of modularity over time $Q(t)$ with dynamic detection algorithms on AS-Internet Routers Graph.

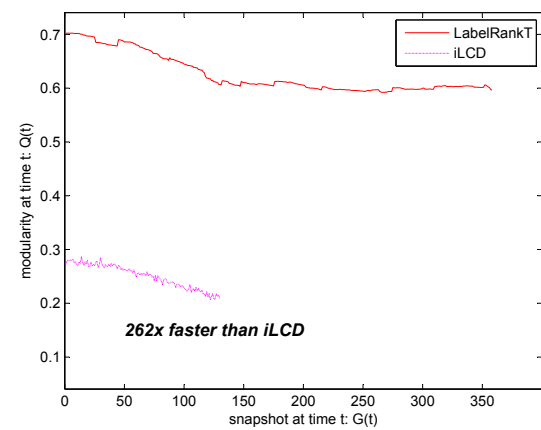


Figure 11: Comparison of modularity over time $Q(t)$ with dynamic detection algorithms on arXiv HEP-TH. We ran iLCD on only the first 130 snapshots due to the time complexity.

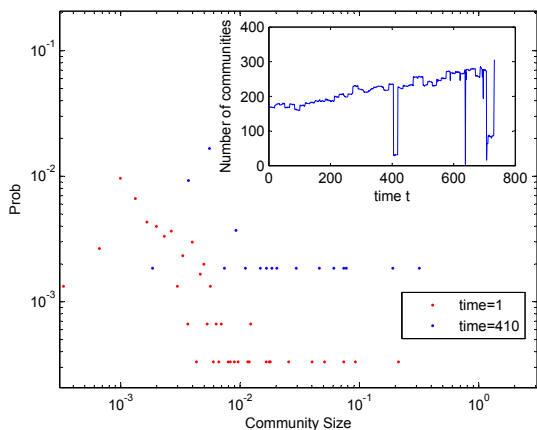


Figure 12: The community size distribution of AS-Internet Routers Graph tracked by *LabelRankT* (loglog plot). Results at time 1 and 410 (dramatic changes occur) are shown in the main plot. The inset shows the number of communities over time.

just 0.07% of each other, while iLCD fails to find strong community structure at all. As shown in Fig. 11, on arXiv HEP-TH, facetNet does not work due to the overflow in memory, while *LabelRankT* performs at least twice better than iLCD. Moreover, *LabelRankT* is more than 100 times faster than both facetNet and iLCD on the two datasets used here.

We also analyzed the number of communities and the distribution of community sizes relative to n (i.e., the probability of seeing a community with certain size), produced by *LabelRankT*. As shown in Fig. 12, AS Graph does not evolve smoothly all the time. The abrupt drop in the number of communities at time 410 signals a dramatic change in structure, which is verified by a completely different distribution of community sizes in comparison with the beginning one. Although this violates our assumption, *LabelRankT* still worked well as evidenced by consistency of its results with the results of static algorithm Infomap. In contrast, arXiv HEP-TH exhibits a fairly smooth pattern shown in Fig. 13. The distributions of community sizes at time 1 and 350 (near the end of evolution) obey power laws with essentially identical exponents. Small size communities grow faster as more and more papers are published as indicated by the downward shift in these distributions. Some communities grow relatively faster than the others and the largest communities expand as indicated by the shift to the right (see the inset).

5. EFFECT OF EDGE WEIGHT AND DIRECTION

In this section, we demonstrate that incorporating the edge weight and direction into the network description allows *LabelRankT* to identify communities better. The experiments were conducted on a dataset including weighted and directed networks.

Reality Mining Bluetooth Scan Data [8]. This dataset was created from the records of Bluetooth Scans generated among the 94 subjects in Reality Mining study conducted

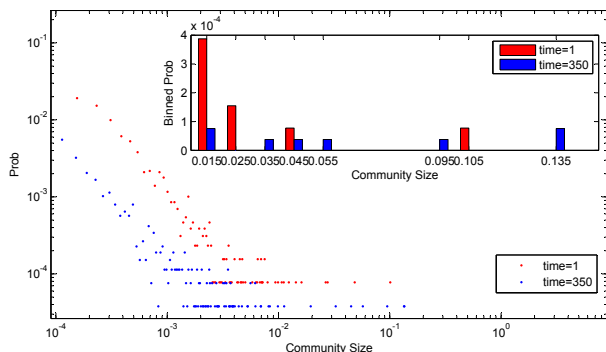


Figure 13: The community size distribution of arXiv HEP-TH tracked by *LabelRankT* (loglog plot). Results at time 1 and 350 (near the end of evolution) are shown. In the inset, the tails of the distributions in main plot are binned with a width 0.01 to show the shift in large size communities.

from 2004-2005 at the MIT Media Laboratory. In the network, nodes represent the subjects and the directed edges correspond to the Bluetooth Scan records and the weight of each edge represent the number of directed Bluetooth scans between the two subjects. In the comparison described below, we only adopted the records from August 02, 2004 (Monday) to May 29, 2005 (Sunday) and we divided them weekly snapshots, so each snapshot represents scans collected during the corresponding week. There are total of 43 snapshots.

We compared the community detection results produced by *LabelRankT* on the Reality Mining Bluetooth Scan network with and without edge weight and direction. By varying the parameter q (from 0.05 to 0.95) of the conditional update Θ_q operator, we calculated the *average modularity differences*, shown in Table 1, between the weighted and directed version and the unweighted and undirected version of *LabelRankT*. All the *average modularity differences* in Table 1 are positive. This demonstrates that including the edge weight and direction improves the performance of our algorithm. Further, Fig. 14 presents the modularity of *LabelRankT* with and without edge weight and direction on all the 43 snapshots with the conditional update parameter $q = 0.6$ when the *average modularity difference* is the largest. In conclusion, these results demonstrate that inclusion of edge weight and direction of the network improves the quality of communities detected by *LabelRankT*.

6. RELATED WORK

Label propagation and random walk based algorithms are most relevant to our work. LPA [18, 26] identifies disjoint groups as nodes with the same label. COPRA [11] and SLPA [27] extend LPA to detection of overlapping communities by allowing multiple labels. However, none of these algorithm resolves the LPA randomness issue, where different communities may be detected in different runs over the same network. Markov Cluster Algorithm (MCL) proposed in [23] is based on simulations of flow (random walk). MCL exe-

Table 1: The average modularity differences between the weighted and directed version and the unweighted and undirected version of *LabelRankT* with different values of the conditional update parameter q on all the 43 snapshots of the Reality Mining Bluetooth Scan data.

Conditional update parameter q	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Average modularity difference	0.101	0.13	0.169	0.174	0.181	0.197	0.206	0.201	0.199	0.193	0.19

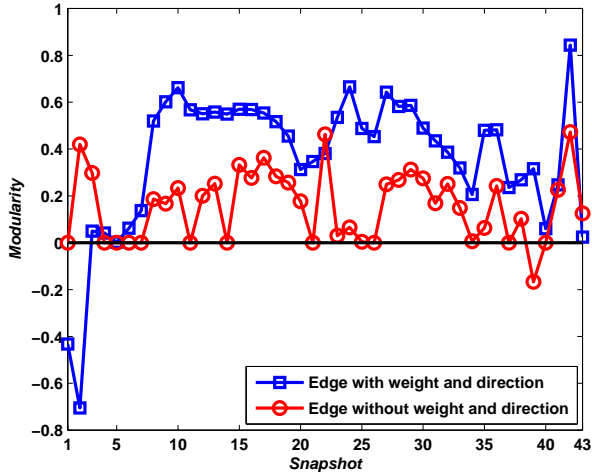


Figure 14: The modularity of *LabelRankT* with and without edge weight and direction on Reality Mining Bluetooth Scan data with the conditional update parameter $q = 0.6$.

cutes repeatedly matrix multiplication followed by inflation operator.

LabelRankT, like its predecessor *LabelRank* (see [28]) differs from MCL in at least two aspects. First, *LabelRankT* applies the inflation to the label distributions and not to the matrix M . Second, the update of label distributions on each node in *LabelRankT* requires only local information. Thus it can be computed in a decentralized way. Regularized-MCL [19] also employs a local update rule of label propagation operator. Despite that, the authors observed that it still suffers from the scalability issue of the original MCL. To remedy, they introduced Multi-level Regularized MCL, making it complex. In contrast, we address the scalability by introducing new operator, conditional update, and the novel stopping criterion, preserving the speed and simplicity of the LPA based algorithms. Moreover, neither MCL nor Regularized-MCL is suitable for dynamic networks.

For dynamic networks, there has been work that focus on exploring the properties of evolving communities that could be used to guide the detection algorithms. Palla et al. [16] developed an algorithm based on the clique percolation method and investigated the time dependence of overlapping communities to uncover basic relationships characterizing community evolution. Tantipathananandh and Berger-Wolf [21] extended their previous social cost model to arbitrary dynamic networks and approximately solved the optimization problem using semidefinite programming relaxation and a rounding heuristic. Bassett et al. [5] proposed an approach

to construct representative partitions. This approach adopts a null model to correct for statistical noise in sets of partitions to improve robustness of community detection result in time-dependent networks.

Typically, an incremental detection [2] considers the stream of changes between snapshots explicitly as opposed to applying static algorithms to each static snapshot [6, 1]. Ning et al. [15] proposed an incremental spectral clustering that continuously updates the eigenvalues and eigenvectors by computing approximations of the generalized eigenvalue system of the normalized cut. facetNet [13] extends the non-negative matrix factorization algorithm. Its drawback is that it requires the number of communities to be provided as an input. GraphScope [20] is a parameter-free algorithm where the minimum description length principle is used to extract communities and to detect the changes. It does not consider the deletion of nodes. Bansal et al. [4] extended CNM algorithm. However, their approach is limited to networks that change very little from snapshot to snapshot. Similarly, Gorke et al. [9] modified both global and local modularity optimization based CNM algorithm and Louvain algorithm, but this modification can handle only small changes between snapshots. iLCD [7] updates the existing community by adding a new node to it if the node's number of second neighbors and number of robust second neighbors are greater than expected values. A limitation is that it can not add two new nodes and a link in-between them at the same time.

7. CONCLUSIONS

Aiming at a highly efficient and general on-line detection algorithm, we introduce *LabelRankT* for incremental detection of evolving communities in large-scale dynamic networks through label propagation. *LabelRankT* is based on the introduced generalized *LabelRank*, in which each node requires only local information during label propagation processing. *LabelRankT* is also able to detect communities in various network types including networks with directed/undirected and weighted/unweighted edges in linear time.

In future work, we plan to apply *LabelRankT* to self-organizing applications such as ad-hoc mobile networks and P2P networks, where each node corresponds to a physical platform. By taking into account temporal and spatial (communities) correlations, we will attempt to construct efficient distributed social-based message routing algorithm on top of *LabelRankT*. We also plan to extend *LabelRankT* to overlapping community detection [25].

Acknowledgment

This work was supported in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053 and by the Office of Naval Research Grant No. N00014-09-1-0607. The views and conclusions contained in

this document are those of the authors and should not be interpreted as representing the official policies either expressed or implied of the Army Research Laboratory, the Office of Naval Research, or the U.S. Government.

8. REFERENCES

- [1] S. Asur, S. Parthasarathy, and D. Ucar. An event-based framework for characterizing the evolutionary behavior of interaction graphs. *TKDD*, 3(4):16:1–16:36, 2009.
- [2] T. Aynaud, J.-L. Guillaume, Q. Wang, and E. Fleury. Communities in evolving networks: definitions, detections and analysis techniques. 2011.
- [3] W. S. Bainbridge. Computational sociology. In *Blackwell Encyclopedia of Sociology*, 2007.
- [4] S. Bansal, S. Bhowmick, and P. Paymal. Fast community detection for dynamic complex networks. In *Proc. of the Second Workshop on Complex Networks*, 2010.
- [5] D. S. Bassett, M. A. Porter, N. F. Wymbs, S. T. Grafton, J. M. Carlson, and P. J. Mucha. Robust detection of dynamic community structure in networks. *Chaos*, 23, 2013.
- [6] T. Y. Berger-Wolf and J. Saia. A framework for analysis of dynamic social networks. In *SIGKDD*, pages 523–528, 2006.
- [7] R. Cazabet, F. Amblard, and C. Hanachi. Detection of overlapping communities in dynamical social networks. In *SOCIALCOM*, pages 309–314, 2010.
- [8] N. Eagle, A. Pentland, and D. Lazer. Inferring social network structure using mobile phone data. *Proceedings of the National Academy of Sciences (PNAS)*, 106(36):15274–15278, 2009.
- [9] R. Görke, P. Maillard, C. Staudt, and D. Wagner. Modularity-Driven Clustering of Dynamic Graphs. In *SEA*, volume 6049, pages 436–448, 2010.
- [10] M. S. Granovetter. The strength of weak ties. *American journal of sociology*, pages 1360–1380, 1973.
- [11] S. Gregory. Finding overlapping communities in networks by label propagation. *New J. Phys.*, 12:103018, 2010.
- [12] J. Leskovec, J. Kleinberg, and C. Faloutsos. Graphs over time: densification laws, shrinking diameters and possible explanations. In *SIGKDD*, pages 177–187, 2005.
- [13] Y.-R. Lin, Y. Chi, S. Zhu, H. Sundaram, and B. L. Tseng. Analyzing communities and their evolutions in dynamic social networks. *TKDD*, 3(2):8:1–8:31, 2009.
- [14] M. E. J. Newman. Fast algorithm for detecting community structure in networks. *Phys. Rev. E*, 69:066133, 2004.
- [15] H. Ning, W. Xu, Y. Chi, Y. Gong, and T. Huang. Incremental spectral clustering with application to monitoring of evolving blog communities. In *SIAM*, 2007.
- [16] G. Palla, A.-L. Barabasi, and T. Vicsek. Quantifying social group evolution. *Nature*, 446(7136):664–667, 2007.
- [17] R. E. Park. *Human communities: The city and human ecology*. Free Press, New York, NY, 1952.
- [18] U. N. Raghavan, R. Albert, and S. Kumara. Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E*, 76:036106, 2007.
- [19] V. Satuluri and S. Parthasarathy. Scalable graph clustering using stochastic flows: applications to community discovery. In *ACM SIGKDD*, pages 737–746, 2009.
- [20] J. Sun, C. Faloutsos, S. Papadimitriou, and P. S. Yu. Graphscope: parameter-free mining of large time-evolving graphs. In *SIGKDD*, pages 687–696, 2007.
- [21] C. Tantipathananandh and T. Y. Berger-Wolf. Finding communities in dynamic social networks. In *Proceedings of IEEE International Conference on Data Mining (ICDM)*, 2011.
- [22] K. Tsuda. Propagating distributions on a hypergraph by dual information regularization. *ICML '05*, pages 920–927, 2005.
- [23] S. van Dongen. A cluster algorithm for graphs. Technical Report INS-R0010, National Research Institute for Mathematics and Computer Science, 2000.
- [24] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*. Cambridge University Press, Cambridge, U.K., 1994.
- [25] J. Xie, S. Kelley, and B. K. Szymanski. Overlapping community detection in networks: the state of the art and comparative study. *ACM Computing Surveys*, 45(4), 2013.
- [26] J. Xie and B. K. Szymanski. Community detection using a neighborhood strength driven label propagation algorithm. In *IEEE NSW 2011*, pages 188–195, 2011.
- [27] J. Xie and B. K. Szymanski. Towards linear time overlapping community detection in social networks. In *PAKDD*, pages 25–36, 2012.
- [28] J. Xie and B. K. Szymanski. Labelrank: A stabilized label propagation algorithm for community detection in networks. In *Proc. IEEE Network Science Workshop*, West Point, NY, pages 138–143, April 2013.