

Gateway finder in large graphs: problem definitions and fast solutions

Hanghang Tong · Spiros Papadimitriou · Christos Faloutsos · Philip S. Yu · Tina Eliassi-Rad

Received: 12 May 2011 / Accepted: 31 January 2012 / Published online: 16 February 2012
© Springer Science+Business Media, LLC 2012

Abstract Given a graph, how to find a small group of ‘gateways’, that is a small subset of nodes that are crucial in connecting the source to the target? For instance, given a social network, who is the best person to introduce you to, say, Chris Ferguson, the poker champion? Or, given a network of people and skills, who is the best person to help you learn about, say, wavelets? We formally formulate this problem in two scenarios: Pair-Gateway and Group-Gateway. For each scenario, we show that it is sub-modular and thus it can be solved near-optimally. We further give fast, scalable algorithms to find such gateways. Extensive experimental evaluations on real data sets demonstrate the effectiveness and efficiency of the proposed methods.

Keywords Social network · Gateway · Graph mining · Scalability · Sub-modularity

1 Introduction

What is the best gateway between a source node and a target node, in a network? This is a core problem that appears under several guises, with numerous generalizations. Motivating applications include the following:

H. Tong (✉)
IBM T.J. Watson Research Center, Yorktown, NY, USA
e-mail: htong@us.ibm.com

S. Papadimitriou
Google Inc, Mountain View, CA, USA

C. Faloutsos
Carnegie Mellon University, Pittsburgh, PA, USA

P. S. Yu
University of Illinois at Chicago, Chicago, IL, USA

T. Eliassi-Rad
Rutgers University, Newark, NJ, USA

1. In a corporate social network, which are the key people that bring or hold different groups together? Or, if seeking to establish a cross-division project, who are the best people to lead such an effort?
2. In an immunization setting, given a set of nodes that are infected, and a set of nodes we want to defend, which are the best few ‘gateways’ we should immunize?
3. Similarly, in a network setting, which are the gateway nodes we should best defend against an attack, to maximize connectivity from source to target.
4. Protein pathways: given a protein interaction network, we know that a certain protein group corresponds to an earlier type of flu (e.g., normal flu), and another group corresponds to a new type of flu (e.g., swine flu), and we want to know which other proteins play a critical role in developing normal flu towards swine flu.
5. Given a graph of co-workers and their skills (keywords), whom should you contact to learn more about, say, Linux? You want someone reasonably close to you and fairly well-versed in Linux, but not your secretary or Linus Torvalds himself.

The problem has several, natural generalizations: (a) we may be interested in the top k best gateways (in case our first few choices are unavailable); (b) we may have more than one source nodes, and more than one target nodes, as in the immunization setting above; (c) we may have a bi-partite graph with relationships (edges) between different node types, as in the last example above. Our main contributions in this paper are:

- A novel ‘gateway-ness’ score for a given source and target, that agrees with human intuition. Its generalization to the case where we have a group of nodes as the source and the target;
- Two algorithms to find a set of nodes with the highest ‘gateway-ness’ score, which (1) are fast and scalable; and (2) lead to *near-optimal* results;
- Extensive experimental results on real data sets, showing the effectiveness and efficiency of the proposed methods.

The rest of the paper is organized as follows: We give the problem definitions in Sect. 2; present ‘gateway-ness’ scores in Sect. 3; and deal with the computational issues in Sect. 4. We evaluate the proposed methods in Sect. 5. Finally, we review the related work in Sect. 6 and conclude in Sect. 7.

2 Problem definitions

Table 1 lists the main symbols we use throughout the paper. In this paper, we focus on directed weighted graphs. We represent the graph by its normalized adjacency matrix (\mathbf{A}). Following standard notation, we use capital bold letters for matrices (e.g., \mathbf{A}), lower-case bold letters for vectors (e.g., \mathbf{a}), and calligraphic fonts for sets (e.g., \mathcal{S}). We denote the transpose with a prime (i.e., \mathbf{A}' is the transpose of \mathbf{A}). We use arrowed lower-case letters for paths on the graph (e.g., \mathbf{p}), which are ordered sequences. We use parenthesized superscripts to represent source/target information for the corresponding variables. For example $\mathbf{p}^{(s,t)} = \{s = u_0, u_1, \dots, u_l = t\}$ is a path from the source node s to the target node t . If the source/target information is clear from the context, we omit the superscript for brevity. A sink node i on the graph is a node without out-links (i.e., $\mathbf{A}(:,i) = 0$). We use subscripts to denote the corresponding variable after setting the nodes indexed by the subscripts as sinks. For example, $\mathbf{p}_{\mathcal{I}}^{(s,t)}$ is the path from the source node s to the target node t , which does not go through any nodes indexed by the set \mathcal{I} (i.e., $u_i \notin \mathcal{I}, i = 0, \dots, l$). With the above notations, our problems can be formally defined as follows:

Table 1 Symbols

Symbol	Definition and description
$\mathbf{A}, \mathbf{B}, \dots$	Matrices (bold upper case)
$\mathbf{A}(i, j)$	The element at the i th row and j th column of matrix \mathbf{A}
$\mathbf{A}(i, :)$	The i th row of matrix \mathbf{A}
$\mathbf{A}(:, j)$	The j th column of matrix \mathbf{A}
\mathbf{A}'	Transpose of matrix \mathbf{A}
$\mathbf{a}, \mathbf{b}, \dots$	Column vectors
$\mathbf{p}, \mathbf{q}, \dots$	Ordered sequences
$\mathcal{S}, \mathcal{T}, \dots$	Sets (calligraphic)
n	Number of nodes in the graph
m	Number of edges in the graph
$g(s, t, \mathcal{I})$	The ‘Gateway-ness’ score for the subset of nodes \mathcal{I} wrt the source s and the target t
$g(\mathcal{S}, \mathcal{T}, \mathcal{I})$	The ‘Gateway-ness’ score for the subset of nodes \mathcal{I} wrt the source group \mathcal{S} and the target group \mathcal{T}
$\mathbf{r}(s, t)$	The proximity score from s to t
$\mathbf{r}_{\mathcal{I}}(s, t)$	The proximity score from s to t by setting the subset of nodes indexed by \mathcal{I} as sinks

Problem 1(Pair-Gateway)

Given: a weighted directed graph \mathbf{A} , a source node s , a target node t , and a budget (integer) k ;

Find: a set of at most k nodes which has the highest ‘gate-way-ness’ score wrt s and t .

Problem 2(Group-Gateway)

Given: a weighted directed graph \mathbf{A} , a group of source nodes \mathcal{S} , a group of target nodes \mathcal{T} , and a budget (integer) k ;

Find: a set of at most k nodes which has the highest ‘gate-way-ness’ score wrt \mathcal{S} and \mathcal{T} .

In both Problem 1 (Pair-Gateway) and Problem 2 (Group-Gateway), there are two sub-problems: (1) how to define the ‘gateway-ness’ score of a given subset of nodes \mathcal{I} ; (2) how to find the subset of nodes with the highest ‘gateway-ness’ score. In the next two sections, we present the solutions for each, respectively.

3 Proposed ‘Gateway-ness’ scores

In this section, we present our definitions for ‘Gateway-ness’. We first focus on the case of a single source s and a single target t (Pair-Gateway). We then generalize to the case where both the source and the target are a group of nodes (Group-Gateway)

3.1 Node ‘Gateway-ness’ score

Given a single source s and a single target t , we want to measure the ‘Gateway-ness’ score for a given set of nodes \mathcal{I} . We first give the formal definitions in such a setting and then provide some intuitions for our definitions.

Formal definitions For a graph \mathbf{A} , we can use random walk with restart to measure the proximity (i.e., relevance/closeness) from the source node s to the target node t , which is defined as follows: Consider a random particle that starts from node s . The particle iteratively transits to its neighbors with probability proportional to the corresponding edge weights. Also at each step, the particle returns to node s with some restart probability $(1 - c)$. The proximity score from node s to node t is defined as the steady-state probability $\mathbf{r}(s, t)$ that the particle will be on node t (Tong et al. 2008). Intuitively, $\mathbf{r}(s, t)$ is the fraction of time that the particle starting from node s will spend on node t of the graph, after an infinite number of steps.

Intuitively, a set of nodes \mathcal{I} are good gateways wrt s and t if they play an important role in the proximity measure from the source to the target. Therefore, our ‘Gateway-ness’ score can be defined as follows:

$$g(s, t, \mathcal{I}) \triangleq \Delta \mathbf{r}(s, t) \triangleq \mathbf{r}(s, t) - \mathbf{r}_{\mathcal{I}}(s, t) \tag{1}$$

where $\mathbf{r}_{\mathcal{I}}(s, t)$ is the proximity score from source s to t after setting the subset of nodes indexed by \mathcal{I} as sinks.

Intuitions Here, we provide some intuition of the ‘Gateway-ness’ score defined by (1), using the running example in Fig. 1.

In Fig. 1, each solid arrowed line is a path from node 1 to node 20, which can be denoted by an ordered sequence. For example, the path marked by the red line can be denoted by $\mathbf{p}^{(1,20)} = \{1, 3, 4, 5, 12, 14, 20\}$. For each path $\mathbf{p}^{(s,t)} = \{s = u_0, u_1, \dots, u_l = t\}$, we can define its score by (2), where $\prod_{i=0}^l \mathbf{A}(u_{i-1}, u_i)$ is the probability that the random particle will traverse this path, and $(1 - c)^l$ penalizes the length of the path. For example, the red path ($\mathbf{p}^{(1,20)} = \{1, 3, 4, 5, 12, 14, 20\}$), has score $(1 - c)^6 \mathbf{A}(3, 1) \mathbf{A}(4, 3) \mathbf{A}(5, 4) \mathbf{A}(12, 5) \mathbf{A}(14, 12) \mathbf{A}(20, 14)$.

$$\text{score}(\mathbf{p}^{(s,t)}) \triangleq (1 - c)^l \prod_{i=0}^l \mathbf{A}(u_{i-1}, u_i) \tag{2}$$

where \mathbf{A} is the normalized adjacency matrix of the graph.

With the above definitions for the path score, we have the following lemma:

Lemma 1 Sum of Weighted Path Scores *Let \mathbf{P} be the set of all the paths from the source node s to the target node t , and \mathbf{Q} be the set of all the paths from the source node s*




Fig. 1 Running example (best viewed in color). (Color figure online)

to the target node t which go through at least one node indexed by the subset \mathcal{I} . Let $\mathbf{r}(s, t)$ be the proximity score defined by random walk with restart and $g(s, t, \mathcal{I})$ be the ‘Gateway-ness’ score defined by eq. (1). Then we have

$$\mathbf{r}(s, t) = \sum_{\mathbf{p}^{(s,t)} \in \mathbf{P}} \text{score}(\mathbf{p}^{(s,t)}); \quad g(s, t, \mathcal{I}) = \sum_{\mathbf{p}^{(s,t)} \in \mathbf{Q}} \text{score}(\mathbf{p}^{(s,t)}) \tag{3}$$

Proof Omitted for brevity. □

By (3), the ‘Gateway-ness’ score for a given set of nodes \mathcal{I} accounts for all the paths from the source node s to the target node t which pass through one or more nodes in \mathcal{I} . For example, given the source node 1 and the target node 20 in Fig. 3, the ‘Gateway-ness’ score for $\mathcal{I} = \{2\}$ is the sum of the scores of all the paths from node 1 to node 20 that go through node 2 (e.g., the green path, the yellow path, and so on).

3.2 Group ‘Gateway-ness’ score

Here we consider the case where the source and/or target consist of more than one nodes. Suppose we have a group of source nodes \mathcal{S} and a group of target nodes \mathcal{T} . Then, the ‘Gateway-ness’ score for a given set of nodes \mathcal{I} can be defined in a similar way:

$$g(\mathcal{S}, \mathcal{T}, \mathcal{I}) \triangleq \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \Delta \mathbf{r}(s, t) \triangleq \sum_{s \in \mathcal{S}, t \in \mathcal{T}} (\mathbf{r}(s, t) - \mathbf{r}_{\mathcal{I}}(s, t)) \tag{4}$$

where $r_{\mathcal{I}}(s, t)$ is the proximity score from s to t by setting the subset of nodes indexed by \mathcal{I} as sinks (i.e., delete all out-edges, by setting $\mathbf{A}(:, i) = 0$ for all $i \in \mathcal{I}$).

Intuitively, the score defined by (4) accounts for all the paths from the source group to the target group¹ which go through at least one node in \mathcal{I} . For example, given $\mathcal{S} = \{1\}$ and $\mathcal{T} = \{19, 20\}$ in Fig. 1, the group ‘Gateway-ness’ score for $\mathcal{I} = \{5, 8\}$ corresponds to all the paths from node 1 to 19 or 20 (e.g., red, yellow and green solid lines, purple and blue dashed lines and so on).

4 BASSET: proposed fast solutions

In this section, we address how to quickly find a subset of nodes of the highest ‘Gateway-ness’ score. We start by showing that the straight-forward methods (referred to as ‘Com-RWR’) are computationally intractable. Then, we present the proposed BASSET (BASSET-N for Pair-Gateway and BASSET-G for Group-Gateway). For each case, we first present the algorithm and then analyze its effectiveness as well as its computational complexity.

4.1 Computational challenges

Here, we present the computational challenges and the way we tackle them. For the sake of succinctness, we mainly focus on BASSET-N.

There are two main computational challenges in order to find a subset of nodes with the highest ‘Gateway-ness’ score. First of all, we need to compute the proximity from the

¹ A path from the source group to the target group is a path which starts from a node of the source group and ends at a node of the target group.

source to the target on different graphs, each of which is a perturbed version of the original graph. This essentially means that we cannot directly apply some powerful pre-computational method to evaluate the proximity from the source to the target (after setting the subset of nodes indexed by \mathcal{I} as sinks). Instead, we have to rely on on-line iterative methods, whose computational complexity is $O(m)$. The challenges are compounded by the need to evaluate $g(s, t, \mathcal{I})$ (1) or $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ (4) an exponential number of times $\binom{n}{k}$. Putting these together, the straightforward way to find k nodes with the highest ‘Gateway-ness’ score is $O\left(\binom{n}{k}m\right)$. This is computationally intractable. Suppose on a graph with 1,000,000 nodes, we want to find the best $k = 5$ gateway nodes. If computing each proximity score takes 0.001 s, then 2.64×10^{17} years are needed to find the gateways. This is much longer than the age of the universe.²

To tackle such challenges, we resort to two main ideas, which are summarized in Theorem 1. According to Theorem 1, in order to evaluate the ‘Gateway-ness’ score of a given set of nodes, we do not need to actually set these nodes as sinks and compute the proximity score on the new graph. Instead, we can compute it from the original graph. In this way, we can utilize methods based on pre-computation to accelerate the process. Furthermore, since $g(s, t, \mathcal{I})$ and $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ are sub-modular wrt \mathcal{I} , we can develop some greedy algorithm to avoid exponential enumeration, and still get some *near-optimal* solution. In Theorem 1, \mathbf{A} is the normalized adjacency matrix of the graph. It is worth pointing out that the proposed methods (BASSET-N and BASSET-G) we will introduce are orthogonal to the specific way of normalization. For simplicity, we use column-normalization throughout this paper. Also, $\mathbf{Q}(\mathcal{I}, \mathcal{I})$ is a $|\mathcal{I}| \times |\mathcal{I}|$ matrix, containing the elements in the matrix \mathbf{Q} which are at the rows/columns indexed by \mathcal{I} . Similarly, $\mathbf{Q}(t, \mathcal{I})$ is a row vector with length $|\mathcal{I}|$, containing the elements in the matrix \mathbf{Q} which are at the t^{th} row and the columns indexed by \mathcal{I} . $\mathbf{Q}(\mathcal{I}, s)$ is a column vector with length $|\mathcal{I}|$, containing the elements in the matrix \mathbf{Q} which are at the s^{th} column and the rows indexed by \mathcal{I} .

Theorem 1 Core Theorem *Let \mathbf{A} be the normalized adjacency matrix of the graph, and $\mathbf{Q} = (\mathbf{I} - c)(\mathbf{I} - c\mathbf{A})^{-1}$. For a given source s and target t , the ‘Gateway-ness’ score of a subset of nodes \mathcal{I} defined in (1) satisfies the properties P1 and P2. For a given source group \mathcal{S} and target group \mathcal{T} , the ‘Gateway-ness’ score of a subset of nodes \mathcal{I} defined in (4) satisfies the properties P3 and P4, where $s \neq t, s, t \notin \mathcal{I}, \mathcal{S} \cap \mathcal{T} = \emptyset, \mathcal{S} \cap \mathcal{I} = \emptyset$, and $\mathcal{T} \cap \mathcal{I} = \emptyset$.*

- P1. $g(s, t, \mathcal{I}) = \mathbf{Q}(t, \mathcal{I})\mathbf{Q}(\mathcal{I}, \mathcal{I})^{-1}\mathbf{Q}(\mathcal{I}, s)$;
- P2. $g(s, t, \mathcal{I})$ is sub-modular wrt the set \mathcal{I} .
- P3. $g(\mathcal{S}, \mathcal{T}, \mathcal{I}) = \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{Q}(t, \mathcal{I})\mathbf{Q}(\mathcal{I}, \mathcal{I})^{-1}\mathbf{Q}(\mathcal{I}, s)$;
- P4. $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ is sub-modular wrt the set \mathcal{I} .

Proof See the “Appendix”. □

Intuition Here, we provide some intuition why $g(s, t, \mathcal{I})$ and $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ are sub-modular. According to Lemma 1, for a given source s and a given target $t, g(s, t, \mathcal{I} \cup \mathcal{K}) - g(s, t, \mathcal{I})$ accounts for the scores of all the paths from s to t , which go through some nodes in \mathcal{K} but none of the nodes in \mathcal{I} . Therefore, for a given set \mathcal{K} , if we already have a bigger subset \mathcal{J} , the additional benefit $(g(s, t, \mathcal{J} \cup \mathcal{K}) - g(s, t, \mathcal{J}))$ will be relatively small, compared to the case where we have a smaller subset $\mathcal{I} (g(s, t, \mathcal{I} \cup \mathcal{K}) - g(s, t, \mathcal{I}))$. For example, in Fig. 1, let $s = 1, t = 20$, and $\mathcal{I} = \{5\}$,

² According to Wikipedia, (http://en.wikipedia.org/wiki/Age_of_the_universe), the age of the universe is about 1.4×10^{10} years.

Algorithm 1 BASSET-N

Input: the normalized adjacency matrix \mathbf{A} , the source node s , the target node t , the budget k and the parameter c

Output: a set of nodes \mathcal{I} , where $|\mathcal{I}| \leq k$.

```

1: initialize  $\mathcal{I}$  to be empty.
2: compute the proximity score  $\mathbf{r}(s, t)$  from the source node  $s$  to the target node  $t$ .
3: find  $i_0 = \operatorname{argmax}_i \frac{\mathbf{r}(s,i)\mathbf{r}(i,t)}{\mathbf{r}(i,i)}$ , where  $i = 1, \dots, n$  and  $i \neq s, i \neq t$ . add  $i_0$  to  $\mathcal{I}$ .
4: for  $j = 2$  to  $k$  do
5:   for  $i = 1$  to  $n$ , and  $i \neq s, i \neq t$  and  $i \notin \mathcal{I}$  do
6:     let  $\mathcal{J} = \mathcal{I} \cup i$ .
7:     compute  $\mathbf{v}(i) = \mathbf{r}(\mathcal{J}, t)' \mathbf{r}(\mathcal{J}, \mathcal{J})^{-1} \mathbf{r}(s, \mathcal{J})'$ 
8:   end for
9:   if  $\max_i \mathbf{v}(i) \leq \mathbf{r}(s, t)$  then
10:    find  $i_0 = \operatorname{argmax}_i \mathbf{v}(i)$ ; add  $i_0$  to  $\mathcal{I}$ .
11:   else
12:    break;
13:   end if
14: end for
15: return  $\mathcal{I}$ 

```

$\mathcal{J} = \{2, 5\}$. Then, if we have a new subset $\mathcal{K} = \{8\}$, the additional benefit for subset \mathcal{I} accounts for all the paths from $s = 1$ to $s = 20$ which go through node 8, but not node 5 (e.g., the green path, etc). While the additional benefit for subset \mathcal{J} is 0, since all the paths from $s = 1$ to $t = 20$ which go through node 8 must also go through some node in \mathcal{J} (node 2).

4.2 BASSET-N for problem 1

4.2.1 BASSET-N algorithm

Our fast solution for Problem 1 is summarized in Algorithm 1. In Algorithm 1, after initialization (step 1), we first pick a node i_0 with the highest $\frac{\mathbf{r}(s,i)\mathbf{r}(i,t)}{\mathbf{r}(i,i)}$ (step 3). Then, in steps 4–14, we find the rest of the nodes in a greedy way. That is, in each outer loop, we try to find one more node while keeping the current \mathcal{I} unchanged. According to P1 of theorem 1, $\mathbf{v}(i)$ computed in step 7 is the gateway score for the subset \mathcal{J} .³ If the current subset of nodes \mathcal{I} can completely disconnect the source and the target (by setting them as sinks), we will stop the algorithm (step 12). Therefore, Algorithm 1 always returns no more than k nodes. It is worth pointing out that in Algorithm 1, all the proximity scores are computed from the original graph \mathbf{A} . Therefore, we can utilize some powerful methods based on pre-computation to accelerate the whole process. To name a few, for a medium size graph \mathbf{A} (e.g., a few thousands of nodes), we can pre-compute and store the matrix $\mathbf{Q} = (\mathbf{I} - c)(\mathbf{I} - c\mathbf{A})^{-1}$; for large unipartite graphs and bipartite graphs, we can use the NB_LIN and BB_LIN algorithms, respectively (Tong et al. 2008).

³ This is because in random walk with restart, we have $\mathbf{r}(i, j) = \mathbf{Q}(j, i)$ for any i, j (Tong et al. 2008).

4.2.2 Analysis of BASSET-N

In this subsection, we analyze the effectiveness and the efficiency of Algorithm 1. First, the effectiveness of the proposed BASSET-N is guaranteed by the following lemma. According to Lemma 2, although BASSET-N is a greedy algorithm, the results it outputs are *near-optimal*.

Lemma 2 Effectiveness of BASSET-N *Let \mathcal{I} be the subset of nodes selected by Algorithm 1 and $|\mathcal{I}| = k_0$. Then, $g(s, t, \mathcal{I}) \geq (1 - 1/e)\max_{|\mathcal{J}|=k_0} g(s, t, \mathcal{J})$, where $g(s, t, \mathcal{I})$, e is the base of the natural logarithm, and $g(s, t, \mathcal{J})$ are defined by (1).*

Proof It is easy to verify that the node i_0 selected in step 10 of Algorithm 1 satisfies $i_0 = \operatorname{argmax}_{j \notin \mathcal{I}, j \neq s, j \neq t} g(s, t, \mathcal{I} \cup j)$. Also, we have $g(s, t, \phi) = 0$, where ϕ is an empty set. On the other hand, according to Theorem 1, $g(s, t, \mathcal{I})$ is sub-modular wrt the subset \mathcal{I} . Therefore (Nemhauser et al. 1978), we have $g(s, t, \mathcal{I}) \geq (1 - 1/e)\max_{|\mathcal{J}|=k_0} g(s, t, \mathcal{J})$, which completes the proof. \square

Next, we analyze the efficiency of BASSET-N, which is given in Lemma 3⁴. We can draw the following two conclusions, according to Lemma 3: (1) the proposed BASSET-N achieves a significant speedup over the straight-forward method ($O(n \cdot k^4)$ vs. $O(\binom{n}{k}m)$). For example, in the graph with 100 nodes and 1,000 edges, in order to find the gateway with $k = 5$ nodes, BASSET-N is more than 6 orders of magnitude faster, and the speedup quickly increases wrt the size of the graph; (2) the proposed BASSET-N is applicable to large graphs since it is linear wrt the number of the nodes.

Lemma 3 Efficiency of BASSET-N *The computational complexity of Algorithm 1 is upper bounded by $O(n \cdot k^4)$.*

Proof The cost for steps 1–2 is constant. The cost for step 3 is $O(n)$. At each inner loop (steps 6–7), the cost is $O(nj^3 + nj^2)$. The cost for steps 9–13 is $O(n)$. The outer loop has no more than $k - 1$ iterations. Putting these together, the computational cost for BASSET-N is:

$$\begin{aligned} \text{Cost}(\text{BASSET-N}) &\leq n + \sum_{j=1}^k (nj^3 + nj^2 + n) \\ &= n + nk + n \frac{k(k+1)(2k+1)}{6} + n \frac{k^2(k+1)^2}{4} = O(nk^4) \end{aligned} \tag{5}$$

which completes the proof. \square

4.3 BASSET-G for problem 2

4.3.1 BASSET-G algorithm

Our fast solution for Problem 2 is summarized in Algorithm 2. It works in a similar way as Algorithm 1: after initialization (step 1), we first pick a node i_0 with the highest $\sum_{s \in \mathcal{S}, t \in \mathcal{T}} \frac{r(s,i)r(i,t)}{r(i,i)}$ (step 3). Then, in steps 4–14, we find the rest of the nodes in a greedy way. That is, in each outer-loop, we try to find one more node while keeping the current \mathcal{I} unchanged. If the current subset of the nodes \mathcal{I} can completely disconnect the source group

⁴ Here, we assume that the cost to get one proximity score is constant, which can be achieved with pre-computation methods (Tong et al. 2008).

Algorithm 2 BASSET-G

Input: the normalized adjacency matrix \mathbf{A} , the source group \mathcal{S} , the target group \mathcal{T} , the budget k and the parameter c

Output: a set of nodes \mathcal{I} , where $|\mathcal{I}| \leq k$.

- 1: initialize \mathcal{I} to be empty.
- 2: compute the proximity score $\sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{r}(s, t)$ from the source group \mathcal{S} to the target group \mathcal{T} .
- 3: find $i_0 = \operatorname{argmax}_i \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \frac{\mathbf{r}(s, i)\mathbf{r}(i, t)}{\mathbf{r}(i, i)}$, where $i = 1, \dots, n$ and $i \neq s, i \neq t$; add i_0 to \mathcal{I} .
- 4: **for** $j = 2$ to k **do**
- 5: **for** $i = 1$ to n , and $i \neq s, i \neq t$ and $i \notin \mathcal{I}$ **do**
- 6: let $\mathcal{J} = \mathcal{I} \cup i$.
- 7: compute $\mathbf{v}(i)$ as $\mathbf{v}(i) = \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{r}(\mathcal{J}, t)'\mathbf{r}(\mathcal{J}, \mathcal{J})^{-1}\mathbf{r}(s, \mathcal{J})'$
- 8: **end for**
- 9: **if** $\max_i \mathbf{v}(i) \leq \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{r}(s, t)$ **then**
- 10: find $i_0 = \operatorname{argmax}_i \mathbf{v}(i)$; add i_0 to \mathcal{I} .
- 11: **else**
- 12: **break**;
- 13: **end if**
- 14: **end for**
- 15: return \mathcal{I}

and the target group (by setting them as sinks), we will stop the algorithm (step 10). As in Algorithm 1, all the proximity scores are computed from the original graph \mathbf{A} . Therefore, we can again utilize those powerful pre-computation based methods to accelerate the whole process.

4.3.2 Analysis of BASSET-G

The effectiveness and efficiency of the proposed BASSET-G are given in Lemma 4 and Lemma 5, respectively. Similar as BASSET-N, the proposed BASSET-G is (1) *near-optimal*; and (2) fast and scalable for large graphs.

Lemma 4 Effectiveness of BASSET-G *Let \mathcal{I} be the subset of nodes selected by Algorithm 2 and $|\mathcal{I}| = k_0$. Then, $g(\mathcal{S}, \mathcal{T}, \mathcal{I}) \geq (1 - 1/e)\max_{|\mathcal{J}|=k_0} g(\mathcal{S}, \mathcal{T}, \mathcal{J})$, where $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$, and $g(\mathcal{S}, \mathcal{T}, \mathcal{J})$ are defined by (4).*

Proof Similar as for Lemma 2. Omitted for brevity. □

Lemma 5 Efficiency of BASSET-G *The computational complexity of Algorithm 2 is upper bounded by $O(n \cdot (\max(k, |\mathcal{S}|, |\mathcal{T}|))^4)$.*

Proof Similar as for Lemma 3. Omitted for brevity. □

5 Experimental evaluations

In this section we present experimental results. All the experiments are designed to answer the following questions:

- 1 *Effectiveness*: how effective are the proposed ‘Gateway-ness’ scores in real graphs?
- 2 *Efficiency*: how fast and scalable are the proposed BASSET-N and BASSET-G?

Table 2 Summary of the data sets

Name	n	m	Weight
<i>Karate</i>	34	152	No
<i>PolBooks</i>	105	882	No
<i>MovieLens</i>	2,399	55,375	No
<i>AC</i>	421,807	2,133,632	No
<i>AA</i>	418,236	2,753,798	Yes
<i>NetFlix</i>	2,667,199	56,919,190	No

5.1 Experimental setup

Data sets We used six real data sets, which are summarized in Table 2.

The first data set (*Karate*) is an un-weighted unipartite graph, which describes friendship among the 34 members of a karate club at a US university (Zachary 1977). Each node is a member in the karate club and the existence of the edge indicates that the two corresponding members are friends. Overall, we have $n = 34$ nodes and $m = 156$ edges.

The second data set (*PolBooks*) is a co-purchasing book network.⁵ Each node is a political book and there is an edge between two books if purchased by the same person. Overall, we have $n = 105$ nodes and $m = 882$ edges.

The third data set (*MovieLens*) is from GroupLens project⁶. It contains 100,000 rating information from 943 users on 1,682 movies. Each user has rated at least 20 movies from 1 (strongly unsatisfactory) to 5 (strongly satisfactory). We use this data set to construct a user-movie bipartite graph. We connect a user with a particular movie if s/he has given some positive ratings (4 or 5) to this movie. As the result, users who give only negative ratings and movies which receive only negative ratings are neglected. On the whole, there are 2,399 nodes (942 users, 1,447 movies) and 55,375 edges. The fourth data set (*AC*) and the fourth data set (*AA*) are both from DBLP.⁷ The third data set (*AC*) is an un-weighted bipartite graph. We have two types of nodes: author and conference. The existence of the edge indicates that the corresponding author has published in the corresponding conference. Overall, we have 421,807 nodes and $m = 2,667,199$ edges.

The fifth data set (*AA*) is a co-authorship network, where each node is an author and the edge weight is the number of the co-authored papers between the two corresponding persons. Overall, we have $n = 418,236$ nodes and $m = 2,753,798$ edges.

The last data set (*NetFlix*) is from the Netflix prize⁸. Rows represent users and columns represent movies. If a user has given a particular movie positive ratings (4 or 5), we connect them with an edge. In total, we have 2,667,199 nodes (2,649,429 users and 17,770 movies), and 56,919,190 edges.

Parameter settings and machine configurations There is one parameter in BASSET-N and BASSET-G, the probability c for random walk with restart. We set $c = 0.95$, as suggested in (Tong et al. 2008). For the computational cost, we report the wall-clock time. All the experiments ran on the same machine with four 2.4GHz AMD CPUs and 48GB memory, running Linux (2.6 kernel). For each experiment, we run it 10 times and report the average.

⁵ <http://www.orgnet.com/>

⁶ <http://www.movielen.org/>

⁷ <http://www.informatik.uni-trier.de/~ley/db/>

⁸ <http://www.netflixprize.com/>

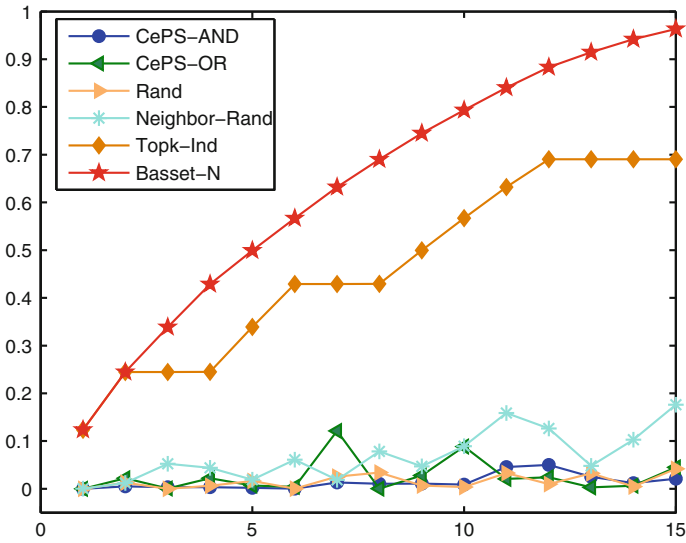


Fig. 2 Effectiveness comparison between BASSET-N and alternatives. Normalized decay of proximity versus k . Higher is better. The proposed BASSET-N (red star) is the best. (Color figure online)

5.2 Effectiveness

Here, we evaluate the effectiveness of the proposed ‘Gateway-ness’ scores. We first compare with several candidate methods in terms of separating the source from the target. And then, we present various case studies.

5.2.1 Quantitative comparisons

The basic idea of the proposed ‘Gateway-ness’ scores is to find a subset of nodes which collectively play an important role in measuring the proximity from the source node (or source group) to the target node (or target group). Here, we want to validate this basic assumption. We compare it with the following alternative choices: (a) selecting k nodes with the highest center-piece AND score (CePS-AND) (Tong and Faloutsos 2006); (b) selecting k nodes with the highest center-piece OR score (CePS-OR) (Tong and Faloutsos 2006); (c) randomly selecting k nodes (Rand); (d) randomly selecting k nodes from the neighboring nodes of the source node and the target node (Neighbor-Rand); (e) selecting k nodes with the highest $\frac{r(s,i)r(i,t)}{r(i,i)}$ (Topk-Ind). We randomly select a source node s and a target node t ,⁹ and then use the different methods to select a subset \mathcal{I} with k nodes. Figure 2 presents the comparison results, where the x-axis is the number of nodes selected (k), and the y-axis is the normalized decay in terms of the proximity score from the source node s to the target node t ($\frac{r(s,t)-r_T(s,t)}{r(s,t)}$). The resulting curves are averaged over 1,000 randomly chosen source-target pairs. From Fig. 2, we can see that (1) the proposed BASSET-N performs best in terms of separating the source from the target; (2) Topk-Ind, where we simply select k nodes with highest $\frac{r(s,i)r(i,t)}{r(i,i)}$, does not perform as well as BASSET-N, where

⁹ The result when source and target are a group of nodes is similar, and omitted for brevity.

we want to find a subset of k nodes which *collectively* has the highest score $\mathbf{r}(\mathcal{I}, t)' \mathbf{r}(\mathcal{I}, \mathcal{I})^{-1} \mathbf{r}(s, \mathcal{I})'$.

5.2.2 Case studies

Next, we will show some case studies, to demonstrate the effectiveness of BASSET-N and BASSET-G.

Karate We start with *Karate* graph, which is widely used in social network analysis. In Fig. 3, there are two different communities in the graph (shaded). In each community, there are some ‘hub’ nodes (e.g., nodes 33 and 34 in the left community; and nodes 1 and 4 in the right community). The two communities are connected by some ‘bridging nodes’ (e.g., nodes 3, 10, 19, 20). Table 3 presents the resulting gateways of BASSET-N with the budget $k = 5$ for a few source-target pairs. The results are consistent with human intuition. The gateways either are the local center of the community that the source/target node belongs to, or are bridging nodes that connect the two communities when the source node and the target node belong to different communities. For example, if $s = 1$ and $t = 33$, the resulting nodes 3, 10, 11 are bridging nodes, while node 34 is the local center for the left community. Note that, we always return less than or equal to $k = 5$ nodes. For example, if $s = 15$ and $t = 34$, we only output one node (node 1) as the gateway. This is because all the paths from node 15 to node 34 must go through node 1.

PolBooks For this data set, the nodes are political books and the existence of the edge indicates the co-purchasing (by the same person) of the two books. Each book is annotated by one of the following three labels: ‘liberal’, ‘conservative’ and ‘neutral’. We pick a ‘liberal’ book (‘The Price of Loyalty’) as the source node, and a ‘conservative’ book (‘Losing Bin Laden’) as the target node. Then, we ran the proposed BASSET-N to find the gateway with 10 nodes. The result is presented in Table 4. The result is again consistent with human intuition, - the resulting gateway books are either popular books in one of the two communities (‘conservative’ vs. ‘liberal’) such as, ‘Bush country’ from ‘conservative’, ‘Back up suck up’ from ‘liberal’, etc; or those ‘neutral’ books which are likely to be purchased by readers from both communities (e.g., ‘Sleeping with the devil’, etc).




Fig. 3 Karate graph

Table 3 BASSET-N on *Karate* graph

Source (<i>s</i>)	Target (<i>t</i>)	Gateways (\mathcal{I})
24	31	{33,34}
15	34	{1}
1	33	{3, 10, 11, 21, 34}

Table 4 BASSET-N on *Pol-Books* graph

Node Index	Book title	Label
10	Bush country	c
13	Off with their heads	c
103	Back up such up	l
5	Sleeping with the devil	n
8	Ghost wars	n
77	Plan of attack	n
78	Bush at war	c
59	Rise of the vulcanes	c
52	Allies	c
42	The bushes	c

c conservative, *l* liberal, *n* neutral

MovieLens For this data set, we pick up a user who likes comedy but not horror movies (i.e., she has given positive ratings to some comedy movies, but has given negative ratings to some horror movies, or has never rated any horror movies) as the source node. We also pick a user who likes horror but not comedy movies as the target node. Then, we ran BASSET-N to find the gateway with $k = 10$ nodes. Table 5 lists the resulting gateways for two such pairs of source-target. We can see that the resulting movies are either popular comedies (e.g., “Shakespeare in Love”) or horror/thrill movies (e.g., “Silence of the Lambs”). Interestingly, “Ghostbusters”, whose genre is both comedy and horror (therefore might be liked by both the source user and the target user), shows up in Table 5a.

AC This is a bipartite graph. Given a source conference/author and a target conference/author, we can run BASSET-N to find either the gateway conferences or the gateway authors (Table 6). Table 7 gives one such example when the source is ‘VLDB’ and the target is ‘NIPS’. Conceptually, we treat an $n_1 \times n_2$ bipartite graph as a $(n_1 + n_2) \times (n_1 + n_2)$ unipartite graph, and we further restrict the search to the desired node type. Again, we can see that the results make sense. The resulting gateway authors are either productive in one of the two fields: databases vs. statistics, (e.g., Prof. Michael I. Jordan in statistics, Prof. Hector Garcia-Molina in databases, etc); or productive in data mining (e.g., Dr. Rakesh Agrawal, Prof. Jiawei-Han), which is an intersection field between statistics and databases. We have similar observations for the resulting gateway conferences. For example, ‘SIGMOD’ and ‘UAI’ are isomorphic (i.e., have very similar neighbor sets) to ‘VLDB’ and ‘NIPS’, respectively; and ‘KDD’ is one major conference in data mining, which is a highly plausible major connection from ‘VLDB’ (databases) to ‘NIPS’ (statistics / machine learning).

AA We use this data set to perform case studies for the proposed BASSET-G. We choose (1) a group of people from a certain field (e.g., ‘text’, ‘theory’, etc) as the source group \mathcal{S} ; and (2) another group of people in some other field (e.g., ‘databases’, ‘bioinformatics’, etc) as the target group \mathcal{T} . Then, we ran the proposed BASSET-N to find the

Table 5 BASSET-N on *Movie-Lens* network

Movie title	Genre
(a) Source user id: 499, target user id: 350	
American beauty	c,d
Silence of the lambs	d,t
Shakespeare in love	c,r
Speed	a,r,t
Four weddings and a funeral	c,r
Dances with wolves	ac,d,w
As good as it gets	c,d
There's something about mary	c
Matrix	ac,s,t
Sixth sense	t
(b) Source user id: 725, target user id: 425	
American beauty	c,d
Annie hall	c,r
There's something about mary	c
Shakespeare in love	c,r
Raiders of the lost ark	ad,ac
Sixth sense	t
Back to the future	c,s
Groundhog day	c,r
Ghostbusters	c,h
Airplane!	c

c comedy, *d* drama, *h* horror, *t* thrill, *r* romance, *ac* action, *ad* adventure, *w* western, *s* sci-fi

gateway with $k = 10$ nodes. Table 8 lists some results. They are all consistent with human intuition, - the resulting authors are either productive authors in one of the two fields, or multi-disciplinary, who have close collaborations to both the source and the target groups of authors.

5.3 Efficiency

We will study the wall-clock running time of the proposed BASSET-N and BASSET-G here. Basically, we want to answer the following two questions:

1. (*Speed*) What is the speedup of the proposed BASSET-N and BASSET-G over the straightforward methods?
2. (*Scalability*) How do BASSET-N and BASSET-G scale with the size of the graph (n and m)?

First, we compare BASSET-N and BASSET-G with two straightforward methods: (1) 'Com-RWR', where we use combinatorial enumeration to find the gateway and, for each enumeration, we compute the proximity from the *new* graph; and (2) 'Com-Eval', where we use combinatorial enumeration to find the gateway, and for each enumeration, we compute the proximity from the *original* graph. Figure 4 shows the comparison on two real data sets. We can draw the following conclusions. (1) Straightforward methods ('Com-RWR' and 'Com-Eval') are computationally intractable even for a small graph. For example, on the *Karate* data set with only 34 nodes, it takes more than 20,560 seconds and

Table 6 BASSET-N on AC graph

Gateway authors	Michael I. Jordan, Philip S. Yu, Jiawei Han, Geoffrey E. Hinton, H. V. Jagadish, Christos Faloutsos, Sebastian Thrun, Rakesh Agrawal, Hector Garcia-Molina, Raghu Ramakrishnan
Gateway conferences	SIGMODE, ICDE, ICML, IJCAI, KDD, AAAI, CIKM, ICANN, SAC, UAI
From the source ‘VLDB’ to the target ‘NIPS’	

Table 7 BASSET-N on AC network

Gateway authors	Gateway conferences
Michael I. Jordan	Sigmod
Philip S. Yu	ICDE
Jiawei Han	ICML
Geoffrey E. Hinton	IJCAI
H. V. Jagadish	KDD
Christos Faloutsos	AAAI
Sebastian Thrun	CIKM
Rakesh Agrawal	ICANN
Hector Garcia-Molina	SAC
Raghu Ramakrishnan	UAI

Table 8 BASSET-G on AA network

Source group	Target group	Gateway [$k = 10$]
(a) A group of people in ‘text’ to a group of people in ‘databases’		
Tom M. Mitchell, William W. Cohen, Jaime G. Carbonell	David J. Dewitt, Hector Garcia-Molina, H. V. Jagadish	Sunita Sarawagi, Christos Faloutsos, James P. Callan, Rakesh Agrawal, Andrew Y. Ng, Yiming Yang, Rich Caruana, Andrew Mccallum, Chengxiang Zhai, Sebastian Thrun,
(b) A group of people in ‘theory’ to a group of people in ‘bioinformatics’		
Manuel Blum, Christos H. Papadimitriou	Vipin Kumar, Wei Wang, George Karypis, Mohammed J. Zaki	Philip S. Yu, Mihalis Yannakakis, Hui Xiong, Hongjun Lu, Sally A. Goldman, Jiawei Han, Moni Naor, Mitsunori Ogihara, Richard M. Karp, Prabhakar Raghavan

100,000 seconds to find the $k = 10$ gateway by ‘Com-Eval’ and by ‘Com-RWR’, respectively. (2) The speedup of the proposed BASSET-N and BASSET-G over both ‘Com-Eval’ and ‘Com-RWR’ is significant - in most cases, we achieve *several (up to 6) orders of magnitude* speedups. (3) The speedup of the proposed BASSET-N and BASSET-G over both ‘Com-RWR’ and ‘Com-Eval’ quickly increases wrt the size of the gateway k . Note that we stop running the program if it takes more than 100,000 seconds (i.e., longer than a day).

Next, we evaluate the scalability of the proposed BASSET-N and BASSET-G wrt the size of the graph, using the largest data set (*NetFlix*). From Fig. 5, we can make the following conclusions: (1) if we fix the number of nodes (n) in the graph, the wall-clock




Fig. 4 Comparison of speed. Wall-clock time versus k . Lower is better. Time is in logarithm scale. The proposed BASSET-N and BASSET-G (red star) are significantly faster. (Color figure online)

time of both BASSET-N and BASSET-G is almost *constant* wrt the number of edges (m); and (2) if we fix the number of edges (m) in the graph, the wall-clock time of both BASSET-N and BASSET-G is *linear* wrt the number of nodes (n). Therefore, they are suitable for large graphs.

6 Related work

In this section, we review the related work, which can be categorized into four parts:

Betweenness centrality The proposed ‘Gateway-ness’ scores relate to measures of betweenness centrality, both those based on the shortest path (Freeman 1977), as well as those based on random walk (Newman 2005). When the gateway set size is $k = 1$, the proposed ‘Gateway-ness’ scores can be viewed as *query-specific* betweenness centrality measures. Moreover, in the proposed BASSET-N and BASSET-G, we aim to find a subset of nodes *collectively*, wherein traditional betweenness centrality, we usually calculate the score for each node *independently* (and then might pick k nodes with the highest individual scores).

Connection subgraphs In the proposed BASSET-N, the idea of finding a subset of nodes wrt the source/target is also related to the concept of connection subgraphs, such as (Faloutsos et al. 2004; Koren et al. 2006; Tong and Faloutsos 2006). However, in connection subgraphs, we aim to find a subset of nodes which have *strong* connections among

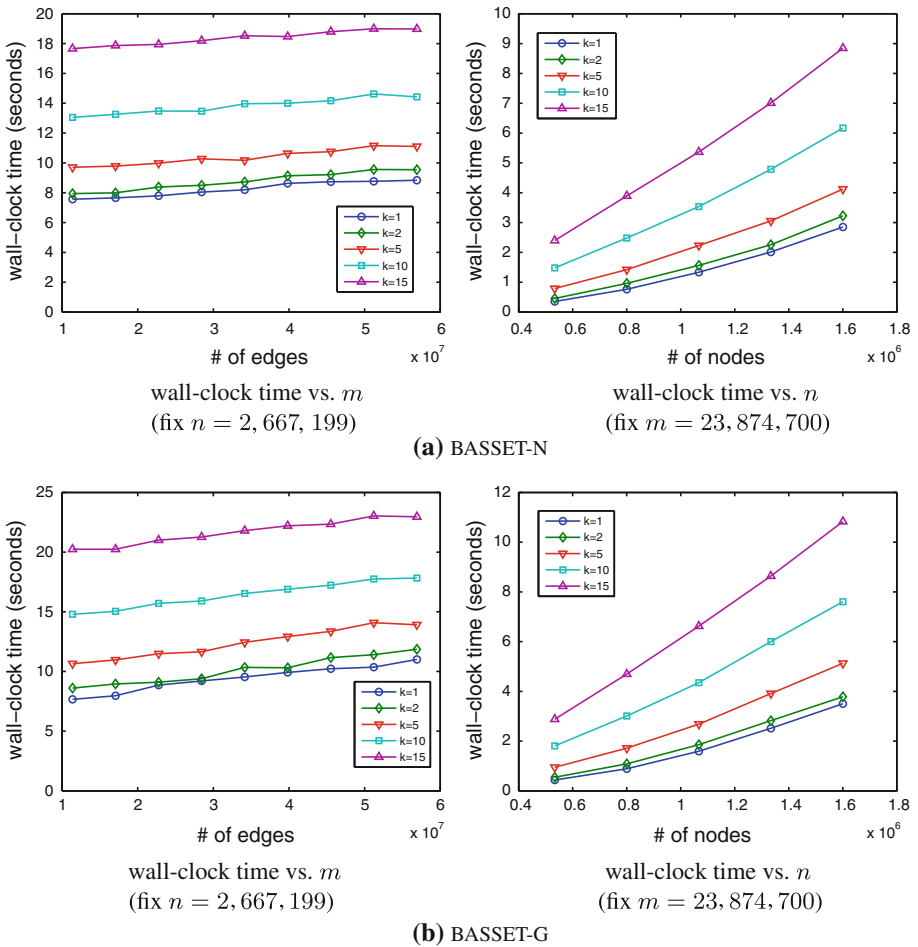


Fig. 5 Scalability of BASSET. Wall-clock time versus the size of the graph. Lower is better. $|S| = |T| = 5$

themselves for the purpose of visualization. While in the proposed BASSET-N, we implicitly encourage the resulting subset of nodes to be disconnected with each other so that they are able to *collectively disconnect* the target node from the source node to the largest extent (if we set them as sinks). It is interesting to notice that, if we want to find the gateway with $k = 1$ for BASSET-N, it can be viewed as a *normalized directed* version of CePS-AND score (Tong and Faloutsos 2006).¹⁰ Moreover, We allow the more general case where the source/target is a group of nodes in the proposed BASSET-G; however in connection subgraphs, the source/target is always a single node.

Graph proximity The basic idea of the proposed BASSET-N and BASSET-G is to find a subset of nodes which will bring the largest decrease of the proximity score from the source node (or the source group) to the target node (or the target group). Graph proximity

¹⁰ To see this, notice that in the case $k = 1$, in BASSET-N, we want to find the node with the highest $\frac{r(s,i)r(i,t)}{r(i,i)}$; while in CePS-AND (Tong and Faloutsos 2006), it picks the nodes with the highest $\frac{r(s,i)}{r(i,i)}$, where $i = 1, \dots, n$ and $i \neq s, i \neq t$.

itself is an important building block in many graph mining settings. Representative work includes the BANKS system (Aditya et al. 2002), link prediction (Liben-Nowell and Kleinberg 2003), content-based image retrieval (He et al. 2004), cross-modal correlation discovery (Pan et al. 2004), pattern matching (Tong et al. 2007), ObjectRank (Balmin et al. 2004), RelationalRank (Geerts et al. 2004), etc.

Other related work in graph mining In recent years, graph mining is a very hot research topic. Representative work includes pattern and law mining (Albert et al. 1999; Broder et al. 2000), frequent substructure discovery (Jin et al. 2005; Xin et al. 2005), influence propagation (Kempe et al. 2003), fraud and anomaly detection (Neville et al. 2005; Noble and Cook 2003), recommendation (Agarwal and Merugu 2007; Cheng et al. 2007), community mining and graph partition (Backstrom et al. 2006; Chen et al. 2009, Gibson et al. 1998; Girvan and Newman ; Karypis and Kumar 1999; Qian et al. 2009), near-clique detection (Pei et al. 2005), etc.

7 Conclusion

In this paper, we study how to find good ‘gateway’ nodes in a graph, given one or more source and target nodes. Our main contributions are: (a) we formulate the problem precisely; (b) we develop BASSET-N and BASSET-G, two fast (up to 6 orders of magnitude of speedup) and scalable (*linear* wrt the number of the nodes in the graph) algorithms to solve it in a provably near-optimal fashion, using sub-modularity. We applied the proposed BASSET-N and BASSET-G on real data sets to validate the effectiveness and efficiency.

Acknowledgments This material is based upon work supported by the National Science Foundation under Grants No. IIS1017415, IIS 0905215, and DBI-0960443. Research was sponsored by the Defense Threat Reduction Agency under contract No. HDTRA1-10-1-0120 and by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053. It is continuing through participation in the Anomaly Detection at Multiple Scales (ADAMS) program sponsored by the U.S. Defense Advanced Research Projects Agency (DARPA) under Agreements No. W911NF-11-C-0200 and W911NF-11-C-0088. This work is also partially supported by an IBM Faculty Award and Google Mobile 2014 Program. The content of the information in this document does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, or other funding parties.

Appendix

A proofs of the core theorem

Here, we provide the detailed proof of Theorem 1 for completeness.

Proof of P1 WLOG, we assume that $\mathcal{I} = \{n - k + 1, \dots, n\}$. Let \mathbf{A} and $\tilde{\mathbf{A}}$ be the normalized adjacency matrices of the graph before/after we set the subset of nodes in \mathcal{I} as sinks. Write \mathbf{A} and $\tilde{\mathbf{A}}$ in block form:

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_{1,1} & \mathbf{A}_{1,2} \\ \mathbf{A}_{2,1} & \mathbf{A}_{2,2} \end{pmatrix}, \tilde{\mathbf{A}} = \begin{pmatrix} \tilde{\mathbf{A}}_{1,1} & \tilde{\mathbf{A}}_{1,2} \\ \tilde{\mathbf{A}}_{2,1} & \tilde{\mathbf{A}}_{2,2} \end{pmatrix} = \begin{pmatrix} \mathbf{A}_{1,1} & \mathbf{0} \\ \mathbf{A}_{2,1} & \mathbf{0} \end{pmatrix} \quad (6)$$

where $\mathbf{0}$ is a matrix with all zero elements.

Let $\tilde{\mathbf{Q}} = (1 - c)(\mathbf{I} - c\tilde{\mathbf{A}})^{-1}$. We can also write $\tilde{\mathbf{Q}}$ and \mathbf{Q} in block form:

$$\mathbf{Q} = (1 - c)(\mathbf{I} - c\mathbf{A})^{-1} = \begin{pmatrix} \mathbf{Q}_{1,1} & \mathbf{Q}_{1,2} \\ \mathbf{Q}_{2,1} & \mathbf{Q}_{2,2} \end{pmatrix} = (1 - c) \begin{pmatrix} \mathbf{I} - c\mathbf{A}_{1,1} & -c\mathbf{A}_{1,2} \\ -c\mathbf{A}_{2,1} & \mathbf{I} - c\mathbf{A}_{2,2} \end{pmatrix}^{-1}$$

$$\tilde{\mathbf{Q}} = \begin{pmatrix} \tilde{\mathbf{Q}}_{1,1} & \tilde{\mathbf{Q}}_{1,2} \\ \tilde{\mathbf{Q}}_{2,1} & \tilde{\mathbf{Q}}_{2,2} \end{pmatrix} = (1 - c) \begin{pmatrix} \mathbf{I} - c\mathbf{A}_{1,1} & \mathbf{0} \\ -c\mathbf{A}_{2,1} & \mathbf{I} \end{pmatrix}^{-1}$$

Applying the block matrix inverse lemma (Piegorsch 1990) to $\tilde{\mathbf{Q}}$ and \mathbf{Q} , we get the following equations:

$$\begin{aligned} \tilde{\mathbf{Q}}_{1,1} &= (1 - c)(\mathbf{I} - c\mathbf{A}_{1,1})^{-1}, \quad \tilde{\mathbf{Q}}_{1,2} = \mathbf{0} \\ \tilde{\mathbf{Q}}_{2,1} &= c(1 - c)\mathbf{A}_{2,1}(\mathbf{I} - c\mathbf{A}_{1,1})^{-1}, \quad \tilde{\mathbf{Q}}_{2,2} = (1 - c)\mathbf{I} \\ \mathbf{Q}_{1,1} &= (1 - c)(\mathbf{I} - c\mathbf{A}_{1,1})^{-1} + c^2(\mathbf{I} - c\mathbf{A}_{1,1})^{-1}\mathbf{A}_{1,2}\mathbf{Q}_{2,2}\mathbf{A}_{2,1}(\mathbf{I} - c\mathbf{A}_{1,1})^{-1} \quad (7) \\ \mathbf{Q}_{1,2} &= c(\mathbf{I} - c\mathbf{A}_{1,1})^{-1}\mathbf{A}_{1,2}\mathbf{Q}_{2,2} \\ \mathbf{Q}_{2,1} &= c\mathbf{Q}_{2,2}\mathbf{A}_{2,1}(\mathbf{I} - c\mathbf{A}_{1,1})^{-1} \end{aligned}$$

Therefore, we have

$$\tilde{\mathbf{Q}}_{1,1} = \mathbf{Q}_{1,1} - \mathbf{Q}_{1,2}\mathbf{Q}_{2,2}^{-1}\mathbf{Q}_{2,1} \quad (8)$$

On the other hand, based on the properties of random walk with restart (Tong et al. 2008), we have $\mathbf{r}(i, j) = \mathbf{Q}(j, i)$, and $\mathbf{r}_{\mathcal{I}}(i, j) = \tilde{\mathbf{Q}}(j, i)$, $(i, j = 1, \dots, n)$. Together with (8), we have

$$\begin{aligned} \mathbf{g}(s, t, \mathcal{I}) &= \mathbf{r}(s, t) - \mathbf{r}_{\mathcal{I}}(s, t) = \mathbf{Q}_{1,1}(t, s) - \tilde{\mathbf{Q}}_{1,1}(t, s) \\ &= \mathbf{Q}_{1,2}(t, :) \mathbf{Q}_{2,2}^{-1} \mathbf{Q}_{1,2}(:, s) \end{aligned} \quad (9)$$

which completes the proofs of P1. □

Proof of P3 Since P1 holds, we have

$$\begin{aligned} \mathbf{g}(\mathcal{S}, \mathcal{T}, \mathcal{I}) &= \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \Delta r(s, t) = \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{g}(s, t, \mathcal{I}) \\ &= \sum_{s \in \mathcal{S}, t \in \mathcal{T}} \mathbf{Q}(t, \mathcal{I}) \mathbf{Q}(\mathcal{I}, \mathcal{I})^{-1} \mathbf{Q}(\mathcal{I}, s) \end{aligned} \quad (10)$$

which completes the proofs of P3. □

Proof of P2 Let $\mathcal{I}, \mathcal{J}, \mathcal{K}$ be three subsets and $\mathcal{I} \subseteq \mathcal{J}$. We will first prove by induction that, for any integer power j , the following inequality holds element-wise.

$$\mathbf{A}_{\mathcal{I}}^j - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^j \geq \mathbf{A}_{\mathcal{J}}^j - \mathbf{A}_{\mathcal{J} \cup \mathcal{K}}^j \quad (11)$$

It is easy to verify the base case (i.e. $j = 1$) for (11) holds. Next, assume that (11) holds for $j = 1, \dots, j_0$, and we want to prove that it also holds for the case $j = j_0 + 1$:

$$\begin{aligned}
 \mathbf{A}_{\mathcal{I}}^{j_0+1} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0+1} &= \mathbf{A}_{\mathcal{I}}^{j_0+1} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0} \mathbf{A}_{\mathcal{I}} + \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0} \mathbf{A}_{\mathcal{I}} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0+1} \\
 &= (\mathbf{A}_{\mathcal{I}}^{j_0} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0}) \mathbf{A}_{\mathcal{I}} + \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0} (\mathbf{A}_{\mathcal{I}} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}) \\
 &\geq (\mathbf{A}_{\mathcal{I}}^{j_0} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0}) \mathbf{A}_{\mathcal{I}} + \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0} (\mathbf{A}_{\mathcal{I}} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}) \\
 &\geq (\mathbf{A}_{\mathcal{I}}^{j_0} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0}) \mathbf{A}_{\mathcal{I}} + \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0} (\mathbf{A}_{\mathcal{I}} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}) \\
 &= \mathbf{A}_{\mathcal{I}}^{j_0+1} - \mathbf{A}_{\mathcal{I} \cup \mathcal{K}}^{j_0+1}
 \end{aligned} \tag{12}$$

In (12), the first inequality holds because of the induction assumption. The second inequality holds because $\mathbf{A}_{\mathcal{I}} \geq \mathbf{A}_{\mathcal{I} \cup \mathcal{K}} \geq 0$ holds element-wise, and $\mathbf{A}_{\mathcal{I} \cup \mathcal{K}} \geq \mathbf{A}_{\mathcal{I}} \geq 0$ holds element-wise.

Since $\tilde{\mathbf{Q}} = (1 - c)(\mathbf{I} - c\tilde{\mathbf{A}})^{-1} = (1 - c) \sum_{j=0}^{\infty} (c\tilde{\mathbf{A}})^j$, we have

$$\begin{aligned}
 g(s, t, \mathcal{I} \cup \mathcal{K}) - g(s, t, \mathcal{I}) &= (1 - c) \sum_{j=0}^{\infty} ((c\mathbf{A}_{\mathcal{I}})^j(t, s) - (c\mathbf{A}_{\mathcal{I} \cup \mathcal{K}})^j(t, s)) \\
 &\geq (1 - c) \sum_{j=0}^{\infty} ((c\mathbf{A}_{\mathcal{I}})^j(t, s) - (c\mathbf{A}_{\mathcal{I} \cup \mathcal{K}})^j(t, s)) \\
 &= g(s, t, \mathcal{I} \cup \mathcal{K}) - g(s, t, \mathcal{I})
 \end{aligned} \tag{13}$$

Therefore, $g(s, t, \mathcal{I})$ is sub-modular, which completes the proof of P2. □

Proof of P4 Since $g(\mathcal{S}, \mathcal{T}, \mathcal{I}) = \sum_{s \in \mathcal{S}, t \in \mathcal{T}} g(s, t, \mathcal{I})$ (In other words, $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ is a non-negative linear combination of sub-modular functions), according to the linearity of sub-modular functions (Krause and Guestrin 2005), we have that $g(\mathcal{S}, \mathcal{T}, \mathcal{I})$ is also sub-modular, which completes the proof of P4. □

References

Aditya, B., Bhalotia, G., Chakrabarti, S., Hulgeri, A., Nakhe, C., & Parag, S. S. (2002). Banks: Browsing and keyword searching in relational databases. In *VLDB* (pp. 1083–1086).

Agarwal, D., & Merugu, S. (2007). Predictive discrete latent factor models for large scale dyadic data. In *KDD* (pp. 26–35).

Albert, R., Jeong, H., & Barabasi, A. -L. (1999). Diameter of the world wide web. *Nature*, *401*, 130–131.

Backstrom, L., Huttenlocher, D. P., Kleinberg, J. M., & Lan, X. (2006). Group formation in large social networks: membership, growth, and evolution. In *KDD* (pp. 44–54).

Balmin, A., Hristidis, V., & Papakonstantinou, Y. (2004). Objectrank: Authority-based keyword search in databases. In *VLDB* (pp. 564–575).

Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, A., & Wiener, J. (2000). Graph structure in the web: experiments and models. In *WWW conference*.

Chen, J., Zaïane, O. R., & Goebel R. (2009). Detecting communities in social networks using max-min modularity. In *SDM* (pp. 978–989).

Cheng, H., Tan, P.-N., Sticklen, J., & Punch, W. F. (2007). Recommendation via query centered random walk on k-partite graph. In *ICDM* (pp. 457–462).

Faloutsos, C., McCurley, K. S., & Tomkins, A. (2004). Fast discovery of connection subgraphs. In *KDD* (pp. 118–127).

Freeman, L. C. (1977). A set of measures of centrality based on betweenness. *Sociometry*, *40*, 35–41.

Geerts, F., Mannila, H., & Terzi, E. (2004). Relational link-based ranking. In *VLDB* (pp. 552–563).

- Gibson, D., Kleinberg, J., & Raghavan, P. (1998). Inferring web communities from link topology. In *Ninth ACM conference on hypertext and hypermedia* (pp. 225–234), New York.
- Girvan, M., & Newman, M. E. J. (2002). Community structure in social and biological networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(12), 7821–7826.
- He, J., Li, M., Zhang, H.-J., Tong, H., & Zhang, C. (2004). Manifold-ranking based image retrieval. In *ACM Multimedia* (pp. 9–16).
- Jin, R., Wang, C., Polshakov, D., Parthasarathy, S., & Agrawal, G. (2005). Discovering frequent topological structures from graph datasets. In *KDD* (pp. 606–611).
- Karypis, G., & Kumar, V. (1999). Multilevel k -way hypergraph partitioning. In *DAC* (pp. 343–348).
- Kempe, D., Kleinberg, J., & Tardos, E. (2003). Maximizing the spread of influence through a social network. *KDD*.
- Koren, Y., North, S. C., & Volinsky, C. (2006). Measuring and extracting proximity in networks. In *KDD* (pp. 245–255).
- Krause, A., & Guestrin, C. (2005). Near-optimal nonmyopic value of information in graphical models. In *UAI* (pp. 324–331).
- Liben-Nowell, D., & Kleinberg, J. (2003). The link prediction problem for social networks. In *Proceedings of CIKM*.
- Nemhauser, G. L., Wolsey, L. A., & Fisher, M. L. (1978). An analysis of approximations for maximizing submodular set functions-i. *Mathematical Programming*, 14, 265–294.
- Neville, J., Simsek, Ö., Jensen, D., Komoroske, J., Palmer, K., & Goldberg, H. G. (2005). Using relational knowledge discovery to prevent securities fraud. In *KDD* (pp. 449–458).
- Newman, M. (2005). A measure of betweenness centrality based on random walks. *Social Networks*, 27, 39–54.
- Noble, C. C., & Cook, D. J. (2003). Graph-based anomaly detection. In *KDD* (pp. 631–636).
- Pan, J. -Y., Yang, H.-J., Faloutsos, C., & Duygulu, P. (2004). Automatic multimedia cross-modal correlation discovery. In *KDD* (pp. 653–658).
- Pei, J., Jiang, D., & Zhang, A. (2005). On mining cross-graph quasi-cliques. In *KDD* (pp. 228–238).
- Piegorsch W., & Casella, G. E. (1990). Inverting a sum of matrices. In *SIAM Review*, vol. 32 (pp. 470–470).
- Qian, T., Srivastava, J., Peng, Z., & Sheu, P. C.-Y. (2009). Simultaneously finding fundamental articles and new topics using a community tracking method. In *PAKDD* (pp. 796–803).
- Tong, H., & Faloutsos, C. (2006). Center-piece subgraphs: problem definition and fast solutions. In *KDD* (pp. 404–413).
- Tong, H., Faloutsos, C., Gallagher, B., & Eliassi-Rad, T. (2007). Fast best-effort pattern matching in large attributed graphs. In *KDD* (pp. 737–746).
- Tong, H., Faloutsos, C., & Pan, J.-Y. (2008). Random walk with restart: Fast solutions and applications. *Knowledge and Information Systems: An International Journal (KAIS)*.
- Xin, D., Han, J., Yan, X., & Cheng, H. (2005). Mining compressed frequent-pattern sets. In *VLDB* (pp. 709–720).
- Zachary, W. W. (1977). An information flow model for conflict and fission in small groups. pp. 452–473.