

Finding Route Hotspots in Large Labeled Networks

Mingtao Lei, Xi Zhang, *Member, IEEE*, Lingyang Chu, Zhefeng Wang, Philip S. Yu, *Fellow, IEEE*, and Binxing Fang

Abstract—In many advanced network analysis applications, like social networks, e-commerce, and network security, hotspots are generally considered as a group of vertices that are tightly connected owing to the similar characteristics, such as common habits and location proximity. In this paper, we investigate the formation of hotspots from an alternative perspective that considers the routes along the network paths as the auxiliary information, and attempt to find the route hotspots in large labeled networks. A *route hotspot* is a cohesive subgraph that is covered by a set of routes, and these routes correspond to the same sequential pattern consisting of vertices' labels. To the best of our knowledge, the problem of *Finding Route Hotspots in Large Labeled Networks* has not been tackled in the literature. However, it is challenging as counting the number of hotspots in a network is #P-hard. Inspired by the observation that the sizes of hotspots decrease with the increasing lengths of patterns, we prove several anti-monotonicity properties of hotspots, and then develop a scalable algorithm called FastRH that can use these properties to effectively prune the patterns that cannot form any hotspots. In addition, to avoid the duplicate computation overhead, we judiciously design an effective index structure called *RH-Index* for storing the hotspot and pattern information collectively, which also enables incremental updating and efficient query processing. Our experimental results on real-world datasets clearly demonstrate the effectiveness and scalability of our proposed methods.

Index Terms—Graphs, sequential pattern, community detection, indexing

1 INTRODUCTION

FINDING communities from real-world networks has gained significant interests, which enjoys various applications involving social networks, e-commerce, and network security. Conventional community detection methods, such as dense subgraph detection [1], detect communities as groups of vertices that are densely interconnected, which focus on the graph structures of communities. More often than not, besides the topology information, the valuable route information [2] may also help to find meaningful communities.

The route on the graph, which is a sequence of consecutive edges, is a natural descriptive model for many real-world networks. The *route* defined here is similar to that defined in [3], which is also a sequence of edges and the consecutive edges must share a vertex. Examples in network security are shown in Figure 1, where each vertex corresponds to a tuple containing an IP address, a port and a specific vulnerability. And each edge indicates that the two connected systems are incident in the network. Hackers may find a sequence of exploits (i.e., a route) along the network path that can enable a successful attack. For example, the routes marked by the colored arrow lines indicate that some

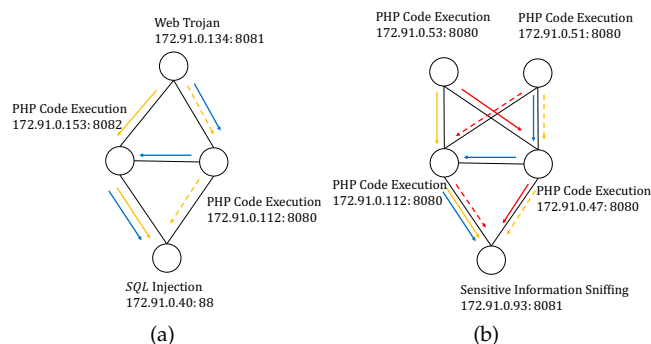


Fig. 1. Two examples of network security, where each vertex corresponds to a tuple containing an IP address, a port and a specific vulnerability. The colored arrow lines on edges represent routes, where a route is a sequence of consecutive edges, incidenting a chain of attacks. Fig. 1(a) shows that hackers frequently first exploit the vulnerabilities *Web Trojan* and *PHP Code Execution*, and then exploit the vulnerability *SQL Injection*. Fig. 1(b) indicates that hackers frequently first exploit the vulnerability *PHP Code Execution* on specific systems, then exploit the same kind of vulnerability (i.e., *PHP Code Execution*) on some other systems, and further exploit the vulnerability *Sensitive Information Sniffing*. These two examples show two types of successful attacking paths in network security.

- M. Lei, X. Zhang and B. Fang are with the Key Laboratory of Trustworthy Distributed Computing and Service (Beijing University of Posts and Telecommunications), Ministry of Education, Beijing 100876, China. E-mail: {leimingtao, zhangx, fangbx}@bupt.edu.cn
- L. Chu is with School of Computing Science, Simon Fraser University, Burnaby V5A 1S6, Canada. E-mail: lca117@sfu.ca
- Z. Wang was with School of Computer Science and Technology, University of Science and Technology of China, Hefei 230026, China. E-mail: zhefengwang@mail.ustc.edu.cn
- P.S. Yu is with the Department of Computer Science, University of Illinois at Chicago, Chicago, IL 60607, USA, and the Institute for Data Science, Tsinghua University, Beijing 100084, China. E-mail: psyu@uic.edu

Manuscript received April 19, 2005; revised August 26, 2015. (Corresponding author: Xi Zhang.)

hackers frequently first exploit the vulnerabilities *Web Trojan* and *PHP Code Execution*, and then exploit the vulnerability *SQL Injection* (Fig. 1(a)), while some other hackers frequently first exploit the vulnerability *PHP Code Execution* on some specific systems, then exploit the same kind of vulnerability (i.e., *PHP Code Execution*) on other systems, and further exploit the vulnerability *Sensitive Information Sniffing* (Fig. 1(b)). Another example lies in the anti-money laundering analysis [4], where each vertex is a user of a bank and an edge represents their money transfer relationship. Here a route can denote a sequence of transfer behaviours among a set of users. The routes may also exist in the collaboration




Fig. 2. An example network, where each vertex corresponds to a specific label, including DB, WS, MS and PS respectively. The colored arrow lines on edges represent routes, where a route \mathcal{T} is a sequence of consecutive edges. The dashed ellipse represents a route hotspot (i.e., the induced subgraph H of $\{v_4, v_6, v_7, v_9\}$) that can be marked by the sequential pattern $\langle PS, MS, DB \rangle$, which is covered by qualified routes that form this hotspot.

networks, such as DBLP [5], where each vertex is an author and an edge represents their collaboration relationship. In addition to collaboration relationships, we also consider the authors' citation relationships to build the citation sequence that is used as the route. Specifically, for every two consecutive collaborative authors A and B in the route, B should have cited a paper published by A. Such routes can indicate the trends of their research interests.

Informally, given a graph as well as routes on it, it is interesting to discover dense components which are covered by groups of routes that share the same patterns implying common behaviours. In the examples of network security that are shown in Figure 1, the detection of these components can facilitate in both identifying the security vulnerable areas in the networks and extracting the features of hacker behaviours, and these features can be used for hacker identification. In the example of the anti-money laundering analysis, the frequent money transfer behaviours among a specific group of users could be suspicious, and the detection of such users may indicate money laundering. In the third example given above, the detection can discover groups of closely cooperated scholars who hold the same continuous research interests in the collaboration networks. We call such dense components *route hotspots* and aim to discover the route hotspots in an efficient manner.

We can also understand the route hotspots detection problem from the perspective of the patterns. It would also be interesting to find the same patterns that are shared by a specific dense component. Although a few previous studies have attempted to mine frequent patterns from a large graph, a pattern that is frequent in the whole graph is not necessarily frequent in a dense component of the graph. Thus, the detection of route hotspots can also help to find meaningful patterns regarding to a specific component of the graph.

However, effectively *finding route hotspots in large labeled networks* is a non-trivial problem. It is challenging on how to define the route hotspot appropriately and how to develop an efficient solution for route hotspot finding.

1.1 Challenge 1: Suitable Model

We first review some related models defined in the previous work and then conclude the key attributes that are required to be possessed by our model.

On one hand, various community models have been proposed to describe dense subgraphs, such as k -core [1] and k -truss [6]. But these dense subgraphs cannot be adopted to model route hotspots as they mainly focus on the graph structures but ignore how routes perform. To better introduce the model, we give a more specific example of network security in Figure 2, where each vertex in the graph corresponds to a system vulnerability and is labeled as its type, such as DB (the database), WS (the website system), MS (the management system) and PS (the portal system). The induced subgraph of $\{v_4, v_7, v_8, v_9, v_{10}\}$ is a 4-truss. However, there are no routes traversing the edges (v_7, v_8) , (v_8, v_9) and (v_8, v_{10}) in this subgraph. A *coverage* relationship that requires the subgraph is covered by a group of routes may make more sense.

On the other hand, there also exist some concepts with the goal of discovering a group of routes that move together, such as *convoy* [7] and *trajectory gathering pattern* [8]. These concepts can be distinguished based on how the "group" is defined. However, using such group patterns to model route hotspots is limited by the following two issues. First, although these approaches are able to identify a group of objects moving together, they only consider spatial proximity for clustering a group of objects rather than more refined and meaningful interconnection relationships between pairwise objects. Moreover, the group of objects they find may not be densely interconnected and thus are not suitable for our problem settings. For example, as shown in Figure 2, the subgraph of $\{v_1, v_2, v_3, v_4, v_5\}$ is obtained by spatial clustering and is covered by $\{\mathcal{T}_1, \mathcal{T}_2, \mathcal{T}_4\}$, but it is not a dense structure. Second, the output of these approaches may be varied by using different distance measures to obtain clustering results. In contrast, our proposal working with k -truss to indicate dense groups is independent of distance measures.

With the insights given above, we regard a route hotspot as a cohesive subgraph covered by a set of routes containing the same sequential pattern. An appropriate model of route hotspot should possess the following key attributes.

- 1) (Cohesiveness) The value of the cohesiveness function that measures the structure of route hotspot is high.
- 2) (Participation) The route hotspot should contain at least min_sup subroutes (a subroute is one of the consecutive portions of the full route).
- 3) (Coverage) These subroutes contained by the route hotspot can cover the route hotspot.
- 4) (Correlation) Each route contained in the route hotspot should contain the same sequential pattern.

The cohesiveness condition is straightforward as hotspots are supposed to be dense subgraphs. Here we use k -truss as it has a stronger guarantee on cohesive structure than k -core. Furthermore, computing k -truss subgraphs takes low computational cost, that is, polynomial time [6]. The participation and coverage conditions are also straightforward as the hotspot is defined on the graph and route set. The correlation condition shows that, in the route hotspot, routes should contain the same sequential pattern, which is a sequence of consecutive labels, and thus we can mark the hotspot by this pattern. As shown in Figure 2, let $k = 3$ and $min_sup = 3$. The induced subgraph H of $\{v_4, v_6, v_7, v_9\}$ is a 4-truss, which contains $\mathcal{T}_5, \mathcal{T}_6, \mathcal{T}_7$ and one portion of

\mathcal{T}_8 , where each route contain the same sequential pattern $\langle PS, MS \rangle$ and these routes can cover H . Therefore, H is a route hotspot for $k = 3$ and pattern $\langle PS, MS \rangle$. We will give more formal definitions in Section 3.

1.2 Challenge 2: Efficient Detection Algorithm

It is natural to ask whether we can apply or extend the existing algorithms for sequential pattern mining, community detection, or group pattern discovery to discover route hotspots? Unfortunately, the answer is no due to the following challenges.

First, the conventional sequential pattern mining methods, like PrefixSpan [9], can only find the patterns that exist in the routes and describe how patterns are located on the graph. They don't consider dense graph structures.

Second, the routes may contain an exponential number of patterns. Since the conventional community detection methods, such as BULK [10], can only detect route hotspots of one pattern at a time, it is computationally intractable to use those methods for each of the exponential number of patterns.

Third, the group pattern discovery methods, like CuTS [7], are not designed to process network topology information and thus the group of vertices they detect may not be densely connected from the perspective of the graph structure such as k -truss.

Last but not least, it is of practical importance to provide a fast query service to answer user queries for interesting route hotspots in real time. As the number of route hotspots is usually huge in a large labeled network, enumerating and indexing all of them efficiently is extremely challenging.

To tackle these challenges, we make the following contributions in this work.

First, to the best of our knowledge, we are the first to detect the route hotspots in the large labeled networks. We prove that counting the number of route hotspots is #P-hard and propose to solve the problem by exploring the graph structure and sequential patterns simultaneously.

Second, to reduce the computational cost of the straightforward greedy algorithms, we first prove the *pattern anti-monotonicity* and the *hotspot anti-monotonicity* properties on route hotspots, and then develop an efficient algorithm called *FastRH* via applying these properties. *FastRH* can also be parallelized since the detection procedure for different patterns is independent, and thus can be speeded up.

Third, to store the huge number of decomposed route hotspots, we design a tree index structure called *RH-Index*, which can efficiently index more than 8.2×10^5 hotspots for 6.7×10^5 patterns by using 2.3 hours and 39.9 GB main memory on a commodity PC with a large main memory as shown in Section 7 on CN dataset. We also develop an efficient querying method that takes only 0.12 seconds on average for each query pattern.

Fourth, we report an extensive experimental study on real-world datasets. The results clearly show that our method is accurate, efficient and scalable in finding and indexing route hotspots.

The rest of the paper is organized as follows. We review the related work in Section 2 and formulate the problem in Section 3. In Section 4 and 5, we present the greedy algorithm and the improved solution. The route hotspot indexing is introduced in Section 6. We report the evaluation results in Section 7, and conclude the paper in Section 8.

2 RELATED WORK

In this section, we review the related work, which can be categorized into sequential pattern mining, k -truss based community detection and group pattern discovery in spatial-temporal data.

2.1 Sequential pattern mining

Sequential pattern mining, which is first introduced by Agrawal and Srikant in [11], is a well-studied topic in data mining. Given a set of sequences, the aim of sequential pattern mining is to find all frequent subsequences. Priori-based algorithms, such as GSP [12] and SPADE [13], try to reduce the search space by using the anti-monotonicity of sequential patterns. To avoid generating candidate sequences, pattern growth-based algorithms, like FreeSpan [14] and PrefixSpan [9], are proposed to mine sequential patterns by recursively projecting the sequence database into a smaller database.

To deal with the large scale of data, approximated methods [15]–[19] and distributed methods [20], [21] are proposed. Riondato et al. [18] apply *Vapnik-Chervonenkis (VC) dimension* to provide tight bounds on the sample size, which is linearly dependent on the VC-dimension of a range space associated with the dataset to be mined. Progressive sampling methods [19] start from small samples and then increment the sampling sizes gradually until meeting the stopping conditions or hitting the upper bound. Li et al. [20] parallelize the FP-Growth on distributed machines in order to reduce memory use and computational cost. Ge et al. [21] develop a memory-efficient distributed dynamic programming approach to mine probabilistic frequent sequential patterns. A comprehensive review for this topic can be found in [22].

However, these approaches are not designed for graph data, and thus cannot be adopted for our problem.

2.2 k -truss based community detection

As there are various definitions on communities, here we mainly focus on the previous studies that are related to k -truss. Recently, Yang et al. [23] propose to find communities over attributed graphs based on k -cores. Zhang et al. [24] propose to find interesting (k, r) -cores, where each vertex in a (k, r) -core connects to at least k other vertices and the similarity between any two vertices of the (k, r) -core is bounded by a similarity threshold r . However, a connected k -core is not guaranteed to be 2-edge-connected, while k -truss is more rigorous than k -core owing to its triangle-based definition. Thus, k -truss based community models are proposed.

Huang et al. [25] propose to search k -truss communities in large-scale networks. To support the efficient search of k -truss communities, they design a compact index, which can be updated incrementally under a dynamic graph setting. They also extend the k -truss community detection models to probabilistic graphs [26] and attributed graphs [10]. In particular, for probabilistic graphs, to meet the requirements in real-world networks, they propose the notion of (k, γ) -truss, where the probability of each edge contained in at least $(k - 2)$ triangles is at least γ . For attributed graphs, motivated by balancing the attribute homogeneity and coverage, they propose the attributed truss community model and define the (k, d) -truss. The proposed (k, d) -truss is a

TABLE 1
Frequently used notations.

Notation	Description
$G(V, E, la)$	The labeled undirected graph.
Π	The union of labels in G .
\mathcal{D}	The route set, where each route contains p .
\mathcal{D}_p	A subset of \mathcal{D} .
min_sup	The user defined positive integer parameter.
$sup(e, G)$	The support of e in G .
\mathcal{T}, t	The route and the subroute of \mathcal{T} .
s	The sequence.
p	The sequential pattern.
$H_p(k)$	The route hotspot with fixed p and k .
n_p	The node of RH-Index for pattern p .
$\mathbb{L}_p, \mathbb{T}_p, \mathbb{E}_p$	The linked item, routes, and edges of n_p .

connected k -truss containing all query nodes and each node of the (k, d) -truss has a distance no larger than d from every query node. Such models are proven to be effective in retrieving meaningful communities in probabilistic graphs and attributed graphs.

In contrast to previous studies, our hotspot is uniquely identified by a pattern, which is not touched by the above methods. Thus, these methods cannot be applied straightforwardly to detect route hotspots.

2.3 Group pattern discovery in spatio-temporal data

Discovering groups of objects that travel together in spatio-temporal data is called *group pattern discovery* [8] and has enjoyed many applications, such as travel-time predication [27] and movement behaviour analysis [2]. Vieira et al. [28] propose to discover *flocks*, where a flock is a group of objects that travel together within a disc of some user-specified size for at least k consecutive timestamps. To deal with the so-called lossy-flock problem, the notion of *convoy* [7] is proposed, where the convoy is defined to have multiple objects that are densely connected with respect to a specific distance during k consecutive time points. Subsequently, to describe the common behaviours in dense groups, the trajectory gathering pattern [8] is proposed, where a gathering is regarded as a dense and continuing group of individuals. Zhang et al. [29] propose GMOVE, to model human mobility from massive geo-tagged social media, by grouping the users that share similar moving behaviors.

However, these methods only consider spatial proximity for clustering a group of objects rather than more refined and meaningful interconnection relationships between paired objects. Thus there is no straightforward way to apply these methods on route hotspots detection.

3 PROBLEM FORMULATION

In this section, we first define *route hotspot*, and then introduce the route hotspot finding problem.

3.1 Graph and k -truss

We consider an undirected graph $G = (V, E, la)$, where (1) V is a finite set of vertices, (2) $E = \{(u, v) \mid u, v \in V\} \subseteq V \times V$ is a set of edges; (3) la is a labeling function that

$$\begin{aligned} \mathcal{T}_1 &= (1, \langle (v_1, v_2), (v_2, v_5) \rangle) \\ \mathcal{T}_2 &= (2, \langle (v_1, v_3), (v_3, v_5) \rangle) \\ \mathcal{T}_3 &= (3, \langle (v_1, v_4), (v_4, v_7), (v_7, v_9) \rangle) \\ \mathcal{T}_4 &= (4, \langle (v_1, v_4), (v_4, v_5) \rangle) \\ \mathcal{T}_5 &= (5, \langle (v_6, v_4), (v_4, v_6) \rangle) \\ \mathcal{T}_6 &= (6, \langle (v_6, v_7), (v_7, v_9) \rangle) \\ \mathcal{T}_7 &= (7, \langle (v_6, v_7), (v_7, v_4), (v_4, v_9) \rangle) \\ \mathcal{T}_8 &= (8, \langle (v_6, v_7), (v_7, v_{10}) \rangle) \end{aligned}$$

Fig. 3. A route set \mathcal{D} of Figure 2, where \mathcal{T}_1 is a route with $tid = 1$, and (PS, WS, DB) is a sequence induced by \mathcal{T}_1 .

maps each vertex v to the unique label $L_v = la(v)$ in Π . Π denotes the set of labels. We denote the set of neighbors of v as $N(v)$, i.e., $N(v) = \{u \in V \mid (v, u) \in E\}$. The degree of v is denoted by $d(v) = |N(v)|$.

A *triangle* in G is a cycle of length 3. We denote a triangle as Δ_{uvw} , where $u, v, w \in V$ are the three vertices on the cycle. The *support of an edge* $e(u, v)$, denoted by $sup(e, G)$, is defined as $|\{\Delta_{uvw} : w \in V\}|$. We denote $e \in \Delta$, if e is an edge of Δ . A k -truss $H(V_H, E_H)$ is a subgraph of G , such that $\forall e \in E_H, sup(e, H) \geq (k - 2)$. H is said to be a *maximal k -truss*, if there does not exist another k -truss that can contain it.

3.2 Route on Graph

Let $es = \langle e_1, \dots, e_{h-1} \rangle$ be a sequence of edges, where $e_i = e(v_i, v_{i+1}) \in E$ ($1 \leq i < h$). A *route* is a tuple $\mathcal{T} = (tid, es)$, where tid is a unique route-id. We say that $e(v_i, v_{i+1})$ is contained by \mathcal{T} , denoted by $e(v_i, v_{i+1}) \in \mathcal{T}$. A route set \mathcal{D} is a set of routes, and $|\mathcal{D}|$ denotes the number of routes. The *length* of the route is $h - 1$, the number of edges.

A *subroute* t of \mathcal{T} , denoted by $t \in \mathcal{T}$, is one of the consecutive portions of \mathcal{T} . For example, $\langle (v_1, v_2), (v_2, v_3) \rangle$ and $\langle (v_2, v_3), (v_3, v_4) \rangle$ are subroutes of $\langle (v_1, v_2), (v_2, v_3), (v_3, v_4) \rangle$.

Given a subgraph $H = (V_H, E_H)$ of G , \mathcal{T} is said to be *route containment* with H , denoted by $\mathcal{T} \in H$, if $\forall e(u, v) \in \mathcal{T}$, such that $e \in E_H$. For example, given the subgraph H of G in Figure 2, where H is a vertex-induced subgraph induced by $\{v_4, v_6, v_7, v_9\}$, \mathcal{T}_7 in Figure 3 is route containment with H as each edge in \mathcal{T} is an edge of H . Similarly, the subroute $\langle (v_6, v_7), (v_7, v_4) \rangle$ of \mathcal{T}_7 is also route containment with H . \mathcal{D} is said to be *route set containment* with H , if $\forall \mathcal{T} \in \mathcal{D}$, such that $\mathcal{T} \in H$, denoted by $\mathcal{D} \in H$. For example, given $H_1 = (V_{H_1}, E_{H_1})$ induced from $\{v_1, v_2, v_3, v_5\}$ in Figure 2 and $\{\mathcal{T}_1, \mathcal{T}_2\}$ in Figure 3, H_1 contains not only the individual routes $\mathcal{T}_1, \mathcal{T}_2$, but also their route set $\{\mathcal{T}_1, \mathcal{T}_2\}$.

\mathcal{D} is said to *cover* H , denoted by $H \sqsubseteq \mathcal{D}$, if each edge of H exists in at least one route of \mathcal{D} . Note that if \mathcal{D} covers H , H must contain the route set \mathcal{D} , but not vice versa. Continued with the above example, H is covered by $\{\mathcal{T}_1, \mathcal{T}_2\}$, but neither \mathcal{T}_1 nor \mathcal{T}_2 can cover H on its own.

3.3 Sequence and Pattern

We are interested in sequential patterns carried by routes in \mathcal{D} . Formally, a *sequence* $s = \langle L_1, \dots, L_h \rangle$ is a sequence of labels, which is *induced* by \mathcal{T} , where $L_i = L_{v_i}$ for $1 \leq i \leq h$. L_i is called the i -th *item* of s . The length of s is $len(s) = h$.

A sequence $p = \langle A_1, \dots, A_l \rangle$ is called a *subsequence* of s and s is called a *supersequence* of p , denoted as $p \sqsubseteq s$, if there

exists a sequence of consecutive integers $1 \leq B_1 < B_2 < \dots < B_l \leq h$ such that $A_1 = L_{B_1}, A_2 = L_{B_2}, \dots, A_l = L_{B_l}$. A route \mathcal{T} is said to *contain* p , if s is induced by \mathcal{T} and $p \sqsubseteq s$, which is denoted as $p \prec \mathcal{T}$ for simplicity. p is said to be a l -*pattern*, if $\text{len}(p) = l$.

The *support* of a sequence p in \mathcal{D} is the number of routes containing p , i.e., $f_{\mathcal{D}}(p) = |\{\mathcal{T} | p \prec \mathcal{T} \wedge \mathcal{T} \in \mathcal{D}\}|$. Given a positive integer min_sup as the *sequence support threshold*, a sequence p is called a *sequential pattern* in \mathcal{D} if $f_{\mathcal{D}}(p) \geq \text{min_sup}$.

Now we give an example to explain the concepts, including the graph, the route, the sequence and the pattern.

Example 1 (Concepts). Figure 2 and Figure 3 show a tiny graph and a set of routes. The graph G has 10 vertices $\{v_1, \dots, v_{10}\}$, 17 edges and $\Pi = \{PS, WS, MS, DB\}$. The route set $\mathcal{D} = \{\mathcal{T}_1, \dots, \mathcal{T}_8\}$. We can always find a sequence supporting a route. For example, $s_1 = \langle PS, WS, DB \rangle$ is induced by \mathcal{T}_1 . Given $\text{min_sup} = 3$, there are 4 frequent sequential patterns, whose lengths are greater than or equal to 2. The patterns are $\langle PS, MS \rangle$, $\langle PS, MS, DB \rangle$, $\langle PS, DB \rangle$ and $\langle MS, DB \rangle$. Suppose that H_2 is a subgraph induced by $\{v_4, v_6, v_7, v_8, v_9, v_{10}\}$, and it's obvious that H_2 is a 3-truss and is *maximal*.

3.4 Route Hotspot

With the commonly accepted desiderata of a good route hotspot described in Section 1, we now give a precise definition of *route hotspot*.

Definition 1 (Route Hotspot). Given a graph G , a set of routes \mathcal{D} , a pattern p , two parameters k and min_sup , $H_p(k)$ (denoted as H for short) is a route hotspot, if H satisfies the following conditions:

- 1) H is a k -truss and is connected.
- 2) H contains at least min_sup subroutes of routes in \mathcal{D} , where each subroute t contains the same pattern p . Formally, $\exists \mathcal{D}_p = \{t \in \mathcal{T} | p \prec t \wedge \mathcal{T} \in \mathcal{D}\}$ and $f_{\mathcal{D}_p}(p) \geq \text{min_sup}$, such that $\mathcal{D}_p \in H$.
- 3) \mathcal{D}_p can cover H and the length of each subroute t in \mathcal{D}_p should be as long as possible. Formally, $H \sqsubseteq \mathcal{D}_p$.
- 4) H is a maximal subgraph satisfying conditions (1), (2) and (3). That is, $\nexists H' \subseteq G$, such that $H \subset H'$, and H' satisfies conditions (1), (2) and (3).

Note that we don't consider the patterns whose lengths are equal to 1 since it's trivial and meaningless. An example of route hotspot H for $p = \langle PS, MS \rangle$ and $k = 3$ is given in Section 1.1, with the corresponding G in Figure 2, \mathcal{D} in Figure 3 and $\text{min_sup} = 3$ respectively.

Note that a vertex whose label doesn't belong to the pattern may also appear in a hotspot, in order to act as a "bridge" between other vertices. Specifically, suppose $p \prec \mathcal{T}$, the length of \mathcal{T} may be longer than the length of p and thus some vertices in \mathcal{T} may be associated with labels that don't exist in p . For example, in Figure 3, \mathcal{T}_6 contains the pattern $\langle PS, MS \rangle$, but $v_9 \in \mathcal{T}_6$ doesn't correspond to the label PS or MS . Nevertheless, v_9 is still involved in the hotspot as it may be indispensable in a practical network, e.g., it may act as a starting point or an endpoint of an attack sequence in the network security scenario. Moreover, longer routes are expected to cover route hotspots more easily than those with shorter lengths, and thus can help us find more hotspots.

3.5 The FRHLN problem

Now we are ready to define the problem of *Finding Route Hotspots in Large Labeled Networks (FRHLN)*.

Definition 2 (FRHLN). Given a graph G , a route set \mathcal{D} , a user-specific parameter min_sup , the problem of *finding route hotspots in large labeled networks* is to enumerate all route hotspots $H_p(k)$ for each p and k , where $k \geq 2$, $\text{len}(p) \geq 2$ and each $H_p(k)$ contains at least min_sup subroutes.

We prove in Theorem 1 that counting the number of the route hotspots is #P-hard.

Theorem 1 (Hardness). Given a graph G , a route set \mathcal{D} and a user input threshold min_sup , the problem of counting the number of route hotspots in G is #P-hard.

Sketch: We prove this theorem by a reduction from the *Sequential Pattern Mining (SPM)* problem, which is known to be #P-hard [30]. Given an instance of the SPM problem, where the input is a sequence database S and a minimal support min_sup , we are asked how many patterns p in S satisfy $f_S(p) \geq \text{min_sup}$. Here, $f_S(p)$ is the number of unique sequences in S that contain p . We denote the union of the items (i.e., the labels of the vertices) in S as \mathbb{I} , where the items follow an arbitrary fixed total order. Using S , we can construct a graph G and a route set \mathcal{D} as follows.

We assume that each item in \mathbb{I} represents a vertex (say v_i) in G and with a unique label (say \mathbb{I}_i). Therefore, the number of vertices and labels in G are both $|\mathbb{I}|$. To construct G , we add edges for every two vertices, and then every three vertices can form a triangle. For each sequence s in S , we build a route \mathcal{T} with the following two steps. Step 1: since s is an ordered list of items, let \mathcal{T} travel the edges with the order of items in s . For example, if $s = \langle \mathbb{I}_4, \mathbb{I}_2, \mathbb{I}_3 \rangle$, \mathcal{T} travels (v_4, v_2) and (v_2, v_3) sequentially. Step 2: let \mathcal{T} travel all the other edges which are not visited in Step 1. Apparently, G is ensured to be covered by \mathcal{T} . Then we obtain a route for each s , and put all these routes in \mathcal{D} . This reduction step takes polynomial time.

For any pattern $p \subset S$, since $f_S(p) \geq \text{min_sup}$, at least min_sup routes in \mathcal{D} contain p . As G is a complete graph, according to our route hotspot definition, there must exist a route hotspot $H_p(k)$ for a pattern p and a k , where k ranges from 2 to $|\mathbb{I}| - 1$. Therefore, the number of route hotspots in G is $|\mathbb{I}| - 2$ times the number of patterns in S , which is exactly the answer to the SPM problem. Consequentially, as SPM problem, the problem of counting the number of route hotspots in G is also #P-hard. \square

The above theorem shows that finding a polynomial-time exact algorithm for computing the number of route hotspots in G is hard. Thus we need to look for pruning strategies to tackle the efficiency issue.

Table 1 summarizes the frequently used notations for the rest of the paper.

4 THE GREEDY ROUTE HOTSPOT FINDING

In this section, we first define the trussness of the subgraph, edge, and route, and then introduce the greedy route hotspot detection method.

4.1 The Trussness of Subgraph, Edge and Route

The definitions of the trussness of subgraph, edge, and route are listed as follows.

Definition 3 (Subgraph Trussness [25]). The trussness of a subgraph $H \subseteq G$ represents the minimum support of an edge in H , denoted by $\tau(H)$. For the purpose of unifying the notion of k -truss, we make $\tau(H) = \min\{sup(e, H) | e \in E(H)\} + 2$. Thus, the trussness of a k -truss is equal to k .

Definition 4 (Edge Trussness [25]). The trussness of an edge $e \in E(H)$ is defined as $\tau_H(e) = \max_{H' \subseteq H} \{\tau(H') | e \in E(H')\}$.

Definition 5 (Route Trussness). The trussness of a route $\mathcal{T} \in H$ is the maximum trussness of the edge trussness, denoted by $\tau_H(\mathcal{T}) = \max\{\tau_H(e) | e \in \mathcal{T}\}$.

Consider the graph G in Figure 2 as an example. The trussness of G is $\tau(G) = 2$. According to Definition 4, the trussness of an edge e is equal to the maximum value of the trussness for different subgraphs that contain e . For example, the subgraph $H = (V_H, E_H)$ contains the edge $e(v_4, v_9)$, where $V_H = \{v_4, v_6, v_7, v_8, v_9, v_{10}\}$ and $E_H = \{e(v_4, v_6), e(v_4, v_7), e(v_4, v_9), e(v_6, v_7), e(v_6, v_8), e(v_7, v_8), e(v_7, v_9), e(v_7, v_{10}), e(v_8, v_{10})\}$. The trussness of H is 4, which is the largest value among that of subgraphs containing $e(v_4, v_9)$. Thus, the trussness of the edges $e(v_4, v_9)$ is 4. The trussness of the routes $\mathcal{T}_1, \mathcal{T}_5$ in Figure 3 are $\tau_G(\mathcal{T}_1) = 2, \tau_G(\mathcal{T}_5) = 4$.

The above definitions can be leveraged to prune unqualified edges and routes. We say that, an edge or a route is *unqualified* if it cannot be used to form a route hotspot. We also say that, a pair of pattern p and k is *unqualified* if $H_p(k) = \emptyset$. For instance, for a route hotspot $H_p(k)$, the support of each edge in $H_p(k)$ must be greater than or equal to k . Thus, if $\tau_G(\mathcal{T}) < k$, which indicates that the trussness of edges in \mathcal{T} is less than k , \mathcal{T} is thus unqualified for a route hotspot and can be pruned safely.

We then propose Algorithm 1 to detect hotspots for a pair of inputs p and k . The general idea is that, we iteratively prune unqualified edges and routes based on the trussness information, and use the remaining edges to form the route hotspots. Specifically, let $H_p(V_p, E_p)$ denote a subgraph induced by \mathcal{D}_p . We first set up an empty queue Q to store unqualified edges (Line 1). We compute the support for each edge and put unqualified edges, whose supports are less than $k - 2$, into Q (Lines 2-3). For each e in Q , we decrement the supports of those related edges who can form triangles with e in H_p (Lines 6-9). Then e will be removed from the routes that contain e (Lines 10-12). Upon the removal of e , some routes may become unqualified and thus will be removed from \mathcal{D}_p (Lines 13-15). After that, e is removed from E_p (Line 16). For each $e' \in E_p$, if it is not contained in any \mathcal{T} , the support of e' will be set as zero (Lines 17-19). We then find all the edges that are unqualified again and put them into Q (Line 20). This process iterates until Q becomes empty. The remaining edges in E_p are used to form connected subgraphs. If the subgraph contains at least min_sup routes that contain p , it would be a route hotspot and put into $\Phi_p(k)$ (Lines 22-26). Finally, $\Phi_p(k)$ is returned (Line 27).

The correctness of Algorithm 1 is apparent since the algorithm essentially detects hotspots by definition. The time

Algorithm 1: Hotspot Detection For p And k

Input: $H_p(V_p, E_p), \mathcal{D}_p, min_sup, p, k$
Output: A set of hotspots $\Phi_p(k)$

- 1: $Q \leftarrow \emptyset$;
- 2: Compute $sup(e, H_p)$ for each edge $e \in E_p$;
- 3: Push each $e(u, v)$ into Q if $sup(e, H_p) < (k - 2)$;
- 4: **while** $Q \neq \emptyset$ **do**
- 5: $e(u, v) = Q.dequeue()$;
- 6: **for** Each $w \in N(u)$ and $(v, w) \in E_p$ **do**
- 7: $sup((u, w), H_p) \leftarrow sup((u, w), H_p) - 1$;
- 8: $sup((v, w), H_p) \leftarrow sup((v, w), H_p) - 1$;
- 9: **end for**
- 10: **while** $\exists \mathcal{T} \in \mathcal{D}_p$, such that $e \in \mathcal{T}$ **do**
- 11: Remove e from \mathcal{T} ;
- 12: **end while**
- 13: **while** $\exists \mathcal{T} \in \mathcal{D}_p$, such that $\tau_{H_p}(\mathcal{T}) < k \vee p \notin \mathcal{T}$ **do**
- 14: Remove \mathcal{T} from \mathcal{D}_p ;
- 15: **end while**
- 16: Remove e from E_p ;
- 17: **while** $\exists e' \in E_p$, such that $e' \notin$ any \mathcal{T} in \mathcal{D}_p **do**
- 18: $sup(e', H_p) = 0$;
- 19: **end while**
- 20: Push each e' into Q if $sup(e', H_p) < (k - 2)$;
- 21: **end while**
- 22: **if** $E_p \neq \emptyset$ **then**
- 23: **while** \exists a connected subgraph $H_p(k)$ of H_p , such that $\mathcal{D}_p(k) \subseteq H_p(k)$ and $|\mathcal{D}_p(k)| \geq min_sup$ **do**
- 24: $\Phi_p(k) \leftarrow \Phi_p(k) \cup H_p(k)$;
- 25: **end while**
- 26: **end if**
- 27: **return** $\Phi_p(k)$

complexity of Algorithm 1 is determined by the enumeration of all triangles for each edge in $H_p(V_p, E_p)$. The time complexity of this enumeration is $\mathcal{O}(\sum_{(v_i, v_j) \in E_p} (d(v_i) + d(v_j))) = \mathcal{O}(|V_p| d_{pmax}^2)$, where $|V_p|$ is the volume of V_p and d_{pmax} denotes the maximum degree of the vertices in V_p .

4.2 The Greedy Route Hotspot Detection (GreedyRH)

Algorithm 1 can be used iteratively to find all route hotspots, but it requires all feasible patterns and feasible values of k as prior inputs. To obtain such inputs, we denote $S = \{s \prec \mathcal{T} | \mathcal{T} \in \mathcal{D}\}$ as a set of sequences, and denote P as a set of patterns mined from S with the support threshold min_sup . We also denote $k_{max} - truss$ as the *maximum truss* in G , indicating that $\exists H \in G$, such that $\forall e \in E_H, sup(e, H) \geq (k_{max} - 1)$.

Observation 1. If there exists a route hotspot $H_p(k)$, p must belong to P .

Observation 2. For each route hotspot $H_p(k)$, k must be less than or equal to k_{max} .

Observation 1 implies that P is a superset of qualified patterns. Observation 2 implies that k_{max} is the upper bound of the maximum value of k . Based on these observations, we propose Algorithm 2, called *GreedyRH*, that runs iteratively for each pattern p and each k ($2 \leq k \leq k_{max}$) and returns all route hotspots. P can be obtained by applying a sequential pattern mining method, like PrefixSpan [9]. And k_{max} can be obtained by applying a truss decomposition method, like [6].

The time complexity of GreedyRH is related to the number of patterns and k_{max} . We denote the maximal length of routes as h_{max} . Thus, the number of patterns is at most $|\Pi|^{h_{max}}$. For each pattern, the Algorithm 1 will be performed for at most $(k_{max} - 1)$ times. Therefore, the time complexity of GreedyRH is $\mathcal{O}(|\Pi|^{h_{max}} \times (k_{max} - 1) \times |V| \times d_{max}^2)$, where $|V|$ is the volume of V and d_{max} denotes the maximum degree of the vertices in V .

Algorithm 2: Greedy Hotspot Detection (GreedyRH)

Input: Graph $G, \mathcal{D}, min_sup, P, k_{max}$
Output: A set of hotspots Φ

```

1: for Each  $p \in P$  do
2:   for  $k = 2 \rightarrow k_{max}$  do
3:     Initialize  $\mathcal{D}_p$  based on  $\mathcal{D}$  and induce  $H_p(V_p, E_p)$  based on  $\mathcal{D}_p$ ;
4:      $\Phi_p(k) :=$  Call Algorithm 1 by using  $H_p, \mathcal{D}_p, min\_sup, p, k$  as input;
5:      $\Phi \leftarrow \Phi \cup \Phi_p(k)$ ;
6:   end for
7: end for
8: return  $\Phi$ 

```

The greedy algorithm, i.e., GreedyRH, is correct due to the following reasons. First, it calls Algorithm 1 iteratively to find all route hotspots, where Algorithm 1 detects a specific route hotspot for a pair of inputs p and k by definition. Second, according to Observations 1 and 2, we can obtain all the patterns and the maximum value of k . Thus, GreedyRH can find all qualified route hotspots for all the patterns and all possible values of k , ensuring the correctness of this method. However, running GreedyRH is impractical. We then analyze the unnecessary computational overhead.

- 1) **Not every k -truss can form a hotspot:** The upper bound k_{max} for k may be too loose. When increasing k , some edges may be removed from G and the remaining part may not be covered by at least min_sup routes. Thus, checking every k becomes unnecessary.
- 2) **Not every pattern is frequent in a k -truss:** In a k -truss of G , the number of routes containing the same p may be less than min_sup . Thus, checking unfrequent patterns in k -trusses also becomes unnecessary.

To overcome these drawbacks, we develop an improved algorithm and describe it precisely in the next section.

5 THE NOVEL ROUTE HOTSPOT FINDING

In this section, we first prove the anti-monotonicity and independence properties that can be used to avoid unnecessary checking. With these properties, we then propose a novel efficient finding method.

5.1 Anti-Monotonicity And Independence Properties

Theorem 2 (Pattern Anti-monotonicity on Hotspots). For patterns p_1, p_2 , where $|p_1| \geq 2$ and $p_1 \sqsubseteq p_2$, and a fixed k , where $k \geq 2$, the following two properties hold.

- 1) If there exists a hotspot $H_{p_2}(k) \neq \emptyset$, there must exist a hotspot $H_{p_1}(k) \neq \emptyset$.
- 2) If $H_{p_1}(k) = \emptyset, H_{p_2}(k) = \emptyset$.

Proof: Suppose that $H_{p_2}(k) = (V_{p_2}(k), E_{p_2}(k))$ is a hotspot for p_2 . We build a subgraph $H_{p_1} = (V_{p_1}, E_{p_1})$, where $V_{p_1} = V_{p_2}(k)$ and $E_{p_1} = E_{p_2}(k)$. There must exist min_sup routes, denoted by \mathcal{D}' , containing p_2 . Since $p_1 \sqsubseteq p_2$, for $\forall T \in \mathcal{D}', p_1 \prec T$. Thus, \mathcal{D}' can also cover H_{p_1} . For p_1 and k , H_{p_1} and \mathcal{D}' satisfy the conditions (1), (2) and (3) of Definition 1, respectively. We just need to determine whether there exists a subgraph $H_{p_1}(k)$ ($H_{p_1} \subset H_{p_1}(k)$) satisfying condition (4) of Definition 1. As $p_1 \sqsubseteq p_2$, the size of H_{p_1} must be larger than or equal to the size of $H_{p_2}(k)$, i.e., $H_{p_2}(k) \subseteq H_{p_1}$. Meanwhile, it is obvious that $H_{p_1} \subseteq G$. Thus, we can always find a maximal H_{p_1} (say $H_{p_1}(k)$) which

is exactly a route hotspot for p_1 and k , where $H_{p_1} \subseteq H_{p_1}(k)$ and $H_{p_1}(k) \subseteq G$. The property 1 holds.

The second property is the contraposition of the first one, and thus it holds as well. \square

Theorem 3 (Hotspot Anti-monotonicity with Patterns). For a pattern p and a fixed k ($k > 2$), if there exists a hotspot $H_p(k)$, there must exist a hotspot $H_p(k - 1)$.

Proof: Suppose that $H_p(k) = (V_p(k), E_p(k))$ is a hotspot for p and k . We build a subgraph $H_p = (V_p, E_p)$, where $V_p = V_p(k)$ and $E_p = E_p(k)$. H_p is a k -truss, where $sup(e, H) \geq (k - 1) \geq (k - 2)$. Thus, H_p is also a $(k - 1)$ -truss. Since $H_p(k)$ is a route hotspot, there must exist min_sup routes, denoted by \mathcal{D}' , containing p . \mathcal{D}' can also cover H_p . For p and $k - 1$, H_p and \mathcal{D}' satisfy the conditions (1), (2) and (3) of Definition 1, respectively. We just need to determine whether H_p satisfies condition (4) of Definition 1, i.e., whether H_p is a maximal subgraph for p and $(k - 1)$. However, the size of H_p is at least equal to the size of $H_p(k)$. Similar to the proof for Theorem 2, there must exist a H_p satisfying condition (4) of Definition 1, which is exactly a hotspot $H_p(k - 1)$ for p and $(k - 1)$. \square

The above two properties ensure that we can prune unqualified patterns and values of k efficiently and safely. Theorem 2 indicates how to prune unqualified edges for a new pattern based on existing patterns. For any pattern p_1 and p_2 , if $p_1 \sqsubseteq p_2$ and p_1 is unqualified, then p_2 is also unqualified. Theorem 3 indicates how to prune unqualified edges for the same pattern with different values of k . For any p , if k is unqualified, then $(k + 1)$ is also unqualified.

We then give an example to show how the above two theorems work. Continued with Example 1, we consider $min_sup = 3$. According to Theorem 2, $\langle WS \rangle$ cannot be grown to $\langle WS, DB \rangle$, since there are only 2 routes containing $\langle WS, DB \rangle$, which doesn't meet the condition of min_sup . Then, there are no route hotspots for $\langle WS, DB \rangle$ or other patterns grown from $\langle WS, DB \rangle$. According to Theorem 3, given the pattern $\langle PS, MS \rangle$, as there are no route hotspots for $k = 4$, there will be no route hotspots for $k = 5$ or even larger values of k .

We also propose the independence property for route hotspot detection.

Proposition 1 (Computation Independence). The detection process for different patterns is independent.

Proposition 1 allows us to speed up the route hotspot detection in a parallel manner. Next, based on Theorem 2 and 3 as well as Proposition 1, we will introduce the improved hotspot finding algorithm compared with GreedyRH.

5.2 The Fast Route Hotspot Scanner (FastRH)

Algorithm 3 shows the procedure of the improved route hotspot detection, i.e., *FastRH*, that can prune unnecessary checking for unqualified patterns and values of k . Denote p^i as the length- i pattern. For initialization, we first calculate the set of length-2 patterns P^2 (Line 1) and put them into a pattern queue Q_p (Line 2). For each pattern p^i extracted from Q_p (Line 4), we initialize the route set \mathcal{D}_{p^i} for p^i and build a graph H_{p^i} based on \mathcal{D}_{p^i} (Line 5). We iteratively call Algorithm 1 to obtain the route set $\Phi_{p^i}(k)$. According to Theorem 3, instead of enumerating k from 2 to k_{max} in Algorithm 2, here we increment k until all the edges have been removed (Lines 6-10), which can reduce the number of

Algorithm 3: Fast Hotspot Scanner (FastRH)

Input: G, \mathcal{D}, \min_sup
Output: The set of hotspots Φ
 1: Initialize: $P^2, Q_p \leftarrow \emptyset$;
 2: Put all $p^2 \in P^2$ into Q_p ;
 3: **while** $Q_p \neq \emptyset$ **do**
 4: $p^i = Q_p.dequeue()$, $k \leftarrow 2$;
 5: Initialize \mathcal{D}_{p^i} based on \mathcal{D} and induce $H_{p^i}(V_{p^i}, E_{p^i})$ based on \mathcal{D}_{p^i} ;
 6: **while** $E_{p^i} \neq \emptyset \wedge |\mathcal{D}_{p^i}| \geq \min_sup$ **do**
 7: $\Phi_{p^i}(k) :=$ Call Algorithm 1 by using $H_{p^i}, \mathcal{D}_{p^i}, \min_sup, p^i, k$ as input;
 8: $\Phi \leftarrow \Phi \cup \Phi_{p^i}(k)$;
 9: $k \leftarrow k + 1$;
 10: **end while**
 11: **if** $\Phi_{p^i}(2) \neq \emptyset$ **then**
 12: Generate P^{i+1} based on p^i ;
 13: Put all $p^{i+1} \in P^{i+1}$ into Q_p ;
 14: **end if**
 15: **end while**
 16: **return** Φ

iterations. According to Theorem 2, instead of enumerating patterns from P in Algorithm 2, here we generate P^{i+1} only when the hotspots for p^i exist, and put all candidate length- $(i + 1)$ patterns into Q_p ; otherwise, we stop the pattern growth from p^i immediately (Lines 11-14). This process iterates until Q_p becomes empty. Finally, the exact set of route hotspots Φ is returned. Please note that lines 3-15 can be parallelized according to Proposition 1. That is, we use the multithreading technique and process multiple patterns at a time in each iteration.

Compared with GreedyRH, the scanner is improved in efficiency due to two reasons. First, it can reduce unnecessary checking for unqualified patterns since a large number of patterns that cannot form route hotspots have been pruned (Line 11-14). Theoretically, we can have up to $|\Pi|^{h_{max}}$ patterns, where h_{max} is the maximum length of patterns. E.g., the datasets GW in Section 7 has 10^7 patterns, but only 8.6×10^5 patterns (in Table 5) can form route hotspots. That means FastRH prunes 91.4% patterns safely. Second, FastRH saves unnecessary checking for unqualified values of k . For instance, Table 5 demonstrates that the maximum k for GW is 6, less than $k_{max} = 16$ in Table 2, indicating the unnecessary computation with k ranging from 7 to 16. Please note that FastRH and its parallelized version PFastRH can produce the correct results of route hotspot finding, as they both adopt Theorems 2 and 3 that can prune unqualified patterns and values of k safely.

Although the computation cost can be saved significantly with FastRH, when a user inputs a new \min_sup , the route hotspot detection procedures have to be performed from scratch, even if the same detection has been performed before. To avoid the duplicate detection, it is essential to store the detected hotspots to enable the quick answering for new user queries. In the next section, we develop a novel framework to index all the detected route hotspots.

6 ROUTE HOTSPOT INDEXING

In this section, we propose an efficient indexing structure for the route hotspots, which is named as the *Route Hotspot Index*, or RH-Index for short.

The *RH-Index* follows the pattern growth property, which is first introduced in FreeSpan [14]. However, the existing pattern growth-based tree structures like FreeSpan can only mine and index the sequential patterns whose

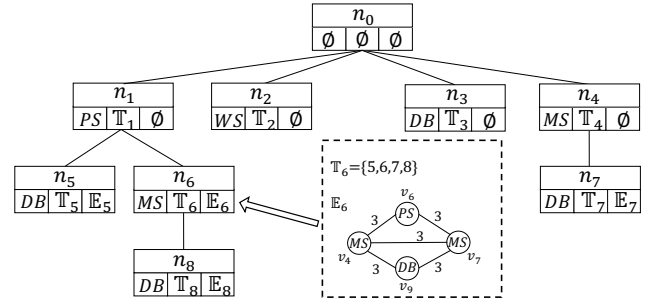


Fig. 4. An example of RH-Index. Note that, although the first layer nodes, such as n_2 and n_4 , cannot form any route hotspots, we still retain these nodes to support index updating.

information is quite simple, but fall short in storing the rich information associated with route hotspots. In particular, they are not able to store the hotspot information and thus cannot be applied to recover the route hotspots from the index for answering the user queries. Moreover, their index construction processes only consider the pattern anti-monotonicity, but ignore the hotspot anti-monotonicity, which may result in low efficiency due to the redundant computing for unqualified values of k .

To overcome these challenges, RH-Index is proposed by introducing the hotspot information and qualified values of k in the index structure. Compared to previous pattern growth-based tree structures such as FreeSpan, the benefits are two-fold. First, it can recover the indexed route hotspots with the linked edges stored in the index. Moreover, the recovery procedure can be quite efficient by using the stored trussness information of edges. Specifically, the RH-Index can directly use the edges, whose trussness is greater than or equal to k , to recover the indexed route hotspots, avoiding extracting unqualified edges whose trussness is smaller than k . Second, following the pattern and hotspot anti-monotonicity properties, the RH-Index can update incrementally when new routes are inserted or the existing routes are deleted, minimizing the overhead needed for index reconstruction.

6.1 Route Hotspot Index Construction

We first define RH-Index. Denote \mathbb{L} as the set of all items in the patterns, and we map each item in \mathbb{L} to Π by using a function $\chi : \mathbb{L} \rightarrow \Pi$. The i -th node of the RH-Index is denoted by n_i , consisting of the item \mathbb{L}_i , the list of routes \mathbb{T}_i and the list of edges \mathbb{E}_i . The i -th pattern p_i is the inverted sequence traveling from n_i to n_0 along the tree structure. \mathbb{T}_i stores the routes which can cover all the hotspots for p_i . The edges of the found hotspots are stored in \mathbb{E}_i . The edges in \mathbb{E}_i are the tuples, each of which consists of an edge e and the edge weight $w(e)$, where $w(e)$ represents the maximal trussness of the hotspot containing e . Given G, \mathcal{D} in Example 1 and $\min_sup = 3$, we can build the corresponding RH-Index consisting of 8 nodes, as shown in Figure 4. For node n_0 , $\mathbb{L}_0 = \emptyset, \mathbb{T}_0 = \emptyset, \mathbb{E}_0 = \emptyset$. For the 1-pattern p , $\mathbb{T}_p = \emptyset, \mathbb{E}_p = \emptyset$, since we only consider the patterns whose lengths are greater than 1. n_6 corresponds to indexed hotspots for $p_6 = \langle PS, MS \rangle$. $\mathbb{L}_6 = MS, \mathbb{T}_6 = \{5, 6, 7, 8\}, \mathbb{E}_6 = \{((v_4, v_6), 3), ((v_4, v_7), 3), ((v_4, v_9), 3), ((v_6, v_7), 3), ((v_7, v_9), 3)\}$.

Algorithm 4: RH-Index Construction

Input: G, \mathcal{D} , The set of route hotspots Φ
Output: A RH-Index \mathbb{T}

- 1: $Q \leftarrow \emptyset, n_0, \mathbb{L}_0 \leftarrow \emptyset, n_0, \mathbb{E}_0 \leftarrow \emptyset$ and $n_0, \mathbb{T}_0 \leftarrow \emptyset$;
- 2: **for** Each item $L_i \in L$ **do**
- 3: Assign $n_i, \mathbb{L}_i \leftarrow L_i, n_i, \mathbb{E}_i \leftarrow \emptyset$ and calculate n_i, \mathbb{T}_i ;
- 4: $n_i, p_i \leftarrow L_i$;
- 5: $n_0.addChild(n_i); Q.inqueue(n_i)$;
- 6: **end for**
- 7: **while** $Q \neq \emptyset$ **do**
- 8: $n_i = Q.dequeue()$;
- 9: **if** $n_i, \mathbb{E}_i \neq \emptyset$ **then**
- 10: **for** Each item $L_j \in L$ **do**
- 11: $p \leftarrow n_i, p_i \cup L_j$;
- 12: $\mathbb{L}_j = L_j$;
- 13: $\mathbb{T}_j \leftarrow \mathcal{D}_{H_p(2)}$, where $H_p(2) \in \Phi_p(2)$;
- 14: $\mathbb{E}_j \leftarrow \{(e, k) | e \in E_{H_p(k)} \wedge H_p(k) \in \Phi_p(k)\}$;
- 15: $n_i.addChild(n_j); Q.inqueue(n_j)$;
- 16: **end for**
- 17: **end if**
- 18: **end while**
- 19: **return** \mathbb{T}

We then introduce how to construct the RH-Index with the given G, \mathcal{D} and the returned route hotspots Φ . The index is built in a top-down manner by using Algorithm 4. We first initialize a queue Q and construct the root node of \mathbb{T} , where $\mathbb{L}_0 = \emptyset, \mathbb{T}_0 = \emptyset, \mathbb{E}_0 = \emptyset$ (Line 1). Then, we generate the first layer nodes of \mathbb{T} , which represent length-1 patterns, and put all nodes into Q (Lines 2-6). Here n_i, \mathbb{L}_i and n_i, \mathbb{T}_i refer to \mathbb{L}_i and \mathbb{T}_i of node n_i , respectively. The rest of RH-Index nodes are built iteratively as the lengths of patterns increase. For each node n_i dequeued from Q , we first obtain the pattern p represented by n_i and then grow p by appending L_j (Line 11). Then $\mathbb{L}_j, \mathbb{T}_j$ and \mathbb{E}_j of n_j will be built correspondingly (Lines 12-15). Finally, the RH-Index \mathbb{T} is successfully built and returned. Please note that the construction of RH-Index don't need to wait until FastRH has detected all the route hotspots. Instead, they may interleave to improve the efficiency. In particular, each time when the set of route hotspots for a specific pattern p is found, a new tree node n_p can thus be built on the tree.

When new routes are inserted or existing routes are deleted, the RH-Index is required to be updated. Inserting new routes may lead to (1) updating some existing nodes with the revised trussness information and edges; or (2) inserting some new nodes as some new patterns become qualified. Let \mathcal{D}_{new} denote the set of inserted routes, and Algorithm 1 is called to determine the nodes that represent the patterns contained by \mathcal{D}_{new} and make the corresponding updating operations. Note that the pattern and hotspot anti-monotonicity properties introduced in Theorem 2 and Theorem 3 can be used to prune unqualified patterns and values of k , avoiding unnecessary updates. This updating process continues until pattern growth ends. We omit the route deletion process as it is similar to the insertion case.

6.2 Query Processing Using RH-Index

We can easily retrieve the hotspots with the given p and k , since the nodes of RH-Index store the trussness information as the weights of edges, which indicate the largest values of k for each hotspot marked by p . Here we propose a query processing method for RH-Index. The answer to query (p, k) is the set corresponding to hotspots $\Phi_p(k)$.

The query processing algorithm is outlined in Algorithm 5. The general idea is that it traverses the RH-Index in

Algorithm 5: Query Processing with RH-Index

Input: \mathbb{T}, p, k
Output: The set of hotspots $\Phi_p(k)$

- 1: $Q \leftarrow n_0, q \leftarrow 0$;
- 2: **while** $Q \neq \emptyset \wedge q < \text{len}(p)$ **do**
- 3: $n_i = Q.dequeue()$; $q \leftarrow q + 1$;
- 4: **if** $\exists n_j \in n_i.children$, such that n_j, \mathbb{L}_j is equal to q -th item of p **then**
- 5: $Q.inqueue(n_j)$;
- 6: **end if**
- 7: **end while**
- 8: $n_p = Q.dequeue()$;
- 9: Collect edges $E_p(k)$ from n_p, \mathbb{E}_p where $\forall e, w(e) \geq k$;
- 10: Induce $\Phi_p(k)$ from $E_p(k)$;
- 11: **return** $\Phi_p(k)$

a breadth-first order to locate the specific node representing the input pattern, and collects the appropriate edges from this node to form the subgraphs, i.e., the route hotspots for answering the queries. Specifically, we first initialize a queue Q with the root node n_0 and a parameter q with 0 (Line 1), where q counts the number of layers that we have traversed along the RH-Index tree. The algorithm first visits a node n_i of RH-Index (n_i starts from n_0). In each step, it increments q and visits the children of n_i in the next layer to select the appropriate node (say n_j) that follows the label sequence of p (Lines 3-4), and put n_j into Q (Line 5). Then it continues to find the appropriate children of n_j . This process iterates until Q is empty or q reaches the length of p (Line 2). If Q is not empty, n_p is dequeued from Q , which represents the pattern p . It then selects the edges from \mathbb{E}_p whose edge weights are greater than or equal to k (Line 9). Finally, the output is the connected subgraphs that are built with the selected edges (Lines 10-11).

The following example illustrates the quick query answering. Given the RH-Index in Figure 4, suppose there is a query $(\langle PS, MS \rangle, 3)$. It first visits the first layer nodes and obtain the node n_1 , where $n_1, \mathbb{L}_1 = PS$. Then it visits the children of n_1 , and obtains the node n_6 , where $n_6, \mathbb{L}_6 = MS$. From n_6, \mathbb{E}_6 , it selects the edges whose weights are greater than or equal to 3, and finally it recovers the route hotspots for $\langle PS, MS \rangle$ and $k = 3$ based on the selected edges.

In summary, RH-Index is simple to construct and efficient to query. We will evaluate its effectiveness and scalability in the experimental section.

7 EXPERIMENTS

In this section, we evaluate the effectiveness and efficiency of our proposed algorithms on real-world networks.

7.1 Experimental Setup

7.1.1 Evaluation Datasets

The following 4 real-world data sets are used for evaluation.

- CN [31] is a collaboration network extracted from DBLP [5]. We regard each author as a vertex and each collaboration paper between two authors as an undirected edge. The label of a vertex represents the author's main research topic. In addition to collaborative relationships that are used to construct the network, we also consider the authors' citation relationships to build the citation sequence that is used as the route. Specifically, for every two consecutive authors A and B in the sequence, B should have

TABLE 2

Graph statistics. #Vertices, #Edges and #Routes indicate the number of vertices, the number of edges and the number of routes respectively. k_{max} is the maximum value of k for the graph. #AT is the average length of the routes.

	#Vertices	#Edges	#Routes	#AT	k_{max}
GW	5.0×10^4	1.3×10^6	3.5×10^5	6.31	16
WB	2.0×10^5	4.9×10^6	1.0×10^6	7.90	31
CN	1.0×10^6	3.9×10^6	5.0×10^6	5.03	58
YELP	1.4×10^5	5.6×10^6	4.9×10^5	4.00	61

cited at least one paper published by A. In CN, the detection of the route hotspot is to identify a group of closely collaborated authors that are covered by a set of routes (i.e., citing sequences), indicating they have common evolving research interests.

- **GW** [32] is a location-based dataset collected from Gowalla where users share their locations by checking-in. To construct GW, we first run k-means [33] to cluster proximate locations as a common vertex, and then connect two vertices if there exists a user moving between them. The labels of vertices are generated randomly as positive integers ranging from 1 to 100. A route represents the path of a user in chronological order. In GW, the detection can discover groups of friends sharing the same traveling routes, indicating they have the common movement and mobility pattern.
- **WB** [34] is a microblogging network crawled from *Sina Weibo*. We regard each user as a vertex and each reciprocal following relationship between two users as an undirected edge. The label of a vertex represents the location of the user. A route represents a sequence of users along the network path according to their following time. For example, for three consecutive users A, B and C in the sequence, B follows A first and then C follows B subsequently. In WB, the detection can discover groups of *Sina Weibo* users who are closely connected and meanwhile have the same following behavior.
- **YELP** [35] is built by treating each business location as a vertex associated with a unique label. If two businesses are reviewed by at least one common user, there would be an edge between the two vertices. The routes follow the order of reviews according to their published time. In YELP, the detection can identify clusters of businesses as well as the common reviewing sequences carried by these businesses.

The network statistics are shown in Table 2.

7.1.2 Comparison Methods

- **CuTS** [7] aims to find convoys which are groups of objects traveling together for a certain period of time. As it is not designed to process graph data, to compare with our proposal, we make a minor modification, that is, we use NG-DBSCAN [36] instead of DBSCAN [37] as the clustering algorithm. Note that we consider the convoys as the route hotspots when using CuTS.
- **CuTSG** extends CuTS in that it first finds convoys by using CuTS, and then determines whether a convoy

can also cover a connected maximal k -truss with as many vertices, edges and routes as possible.

- **GreedyRH** is the greedy algorithm described in Algorithm 2.
- **FastRH** is our core algorithm described in Algorithm 3. We apply both the properties of pattern anti-monotonicity and hotspot anti-monotonicity on GreedyRH.
- **PFastRH** is the parallel version of FastRH, which is implemented by using multithreading.
- **GreedyRHP** is an improved algorithm based on GreedyRH, where we apply the property of pattern anti-monotonicity on GreedyRH.
- **GreedyRHH** is the other improved algorithm based on GreedyRH, where we apply the property of hotspot anti-monotonicity on GreedyRH.

7.1.3 Evaluation Metrics

- **Precision Score (PS)**. \mathcal{C} denotes the complete set of hotspots obtained by GreedyRH and $\hat{\mathcal{C}}$ denotes the set of hotspots obtained by other methods, such as FastRH, GreedyRHP, GreedyRHH and PFastRH. PS is computed as $PS = \frac{|\mathcal{C} \cap \hat{\mathcal{C}}|}{|\hat{\mathcal{C}}|}$. This metric is used to verify our proposal works correctly with the anti-monotonicity and independence properties (i.e., Theorem 2, Theorem 3 and Proposition 1). Please note that, as CuTS and CuTSG don't use these properties to prune unqualified route hotspots or achieve parallel speedup, they are not evaluated in terms of PS.
- **Time Cost and Space Cost**. "Time Cost" is the running time for the detection or indexing process. "Space Cost" is the main memory used for indexing.
- **Number of Patterns (#NP) and Number of Hotspots (#NH)**. "Number of Patterns" and "Number of Hotspots" refer to the total number of patterns that are used to support hotspots and the total number of hotspots that are obtained with our methods respectively. We use #CuNP, #CuGNP and #RHNP (resp. #CuNH, #CuGNH and #RHNH) to represent the total number of patterns (resp. number of hotspots) obtained with CuTS, CuTSG and PFastRH.

7.1.4 Implementation Details

All algorithms are implemented in Java and all of the experiments are conducted on a commodity PC with Linux 16.04, Core-i7 6700K CPU (4.00GHz) and 64 GB main memory. We use 8 threads for running PFastRH. To evaluate the impacts of min_sup , we set $min_sup = 100, 200, 300, 400, 500$ for CN, and $min_sup = 10, 20, 30, 40, 50$ for GW, WB and YELP respectively.

7.2 Experimental results

7.2.1 Effectiveness Analysis

We evaluate the effectiveness of these methods on WB, GW, CN and YELP, but omit the results of GreedyRH as it is too slow on large datasets.

Table 3 shows the precision scores with varying min_sup for all compared methods. We can see that the PS values of FastRH, GreedyRHP, GreedyRHH, and PFastRH on GW, WB, CN, and YELP are always equal to 100%. It indicates that, FastRH, GreedyRHP, GreedyRHH, and PFastRH won't

TABLE 3
The PS on the GW, WB, YELP and CN.

GW/WB/YELP					
min_sup	10	20	30	40	50
PS (FastRH)	100%	100%	100%	100%	100%
PS (GreedyRHP)	100%	100%	100%	100%	100%
PS (GreedyRHH)	100%	100%	100%	100%	100%
PS (PFastRH)	100%	100%	100%	100%	100%
CN					
min_sup	100	200	300	400	500
PS (FastRH)	100%	100%	100%	100%	100%
PS (GreedyRHP)	100%	100%	100%	100%	100%
PS (GreedyRHH)	100%	100%	100%	100%	100%
PS (PFastRH)	100%	100%	100%	100%	100%

change the number of hotspots, which validates the correctness of our pattern and hotspot anti-monotonicity properties as well as the independent properties.

7.2.2 Analysis on Patterns and Route Hotspots

We analyze the total number of patterns and route hotspots obtained with PFastRH, CuTS and CuTSG by varying min_sup , and the results are shown in Figure 5.

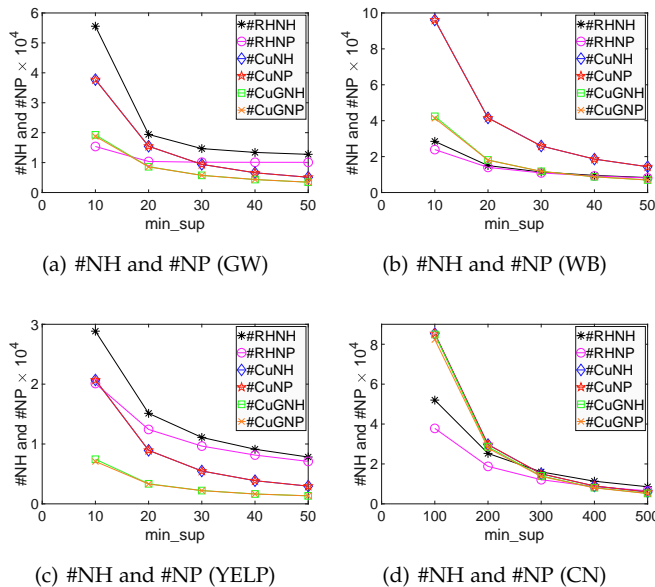


Fig. 5. #NH and #NP on the GW, WB, YELP, and CN.

It can be observed that #NP and #NH decrease when min_sup increases. Generally, #NH would be greater than or equal to #NP as one pattern may correspond to multiple hotspots, but we can observe in Figure 5 that #NH is almost equal to #NP when min_sup is large. The reason is that larger min_sup may make some route hotspots become unqualified as their numbers of routes are now less than min_sup . Note that, #CuNH is equal to #CuNP, as CuTS considers a pattern (i.e., a convoy) as a hotspot.

Also, as shown in Figure 5, #CuNP and #CuNH are larger than #CuGNP and #CuGNH, because the definition of k -truss-based convoys defined in CuTSG is more rigorous than that defined in CuTS. In the datasets WB and CN, #CuNP and #CuGNP (resp. #CuNH and #CuGNH)

are greater than #RHNP (resp. #RHNH) when min_sup is small. This is because the definition of route hotspot defined in Definition 1 uses both pattern and graph information, which is able to avoid a large amount of unqualified patterns. However, when min_sup increases, #CuNP and #CuGNP become smaller than #RHNP, and #CuNH and #CuGNH are also smaller than #RHNH. This is because CuTS and CuTSG only use route information but ignore the sequential patterns in the routes. In the datasets GW and YELP, #RHNH is greater than #CuNH and #CuGNH with different min_sup . We can also observe that, for each dataset, the gap between #RHNH and #RHHP is larger than that between #CuGNH and #CuGNP, as well as that between #CuNH and #CuNP. The reason is that, our method can find more densely interconnected subgraphs with larger values of k .

7.2.3 Efficiency of Hotspot Finding

We evaluate the efficiency of FastRH, GreedyRHP, GreedyRHH, PFastRH, CuTS and CuTSG by varying min_sup . The performances in terms of time cost are shown in Figure 6, where the time cost is the running time for a specific algorithm. Please note that as GreedyRH is too slow, we omit its results.

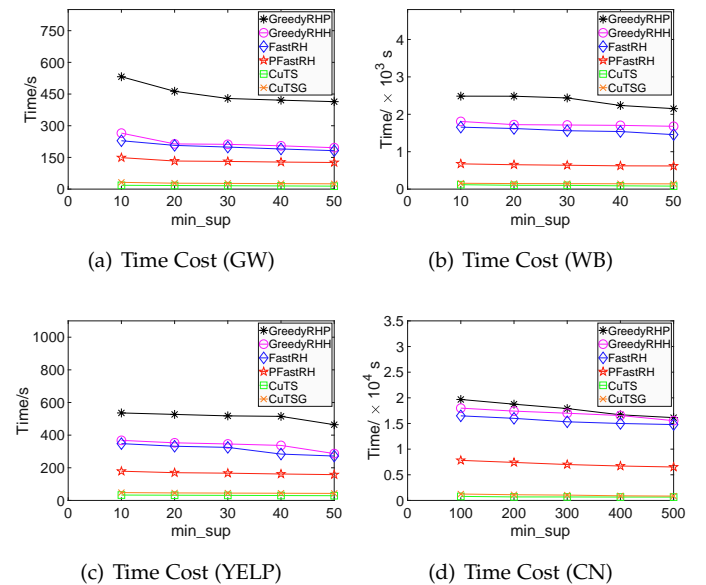


Fig. 6. The Time Cost on the GW, WB, YELP, and CN.

As shown in Figure 6, when min_sup increases, the time cost decreases due to the decrease of #NP and #NH shown in Figure 5. We can observe that GreedyRHH is more efficient than GreedyRHP. The reason is that GreedyRHH uses the hotspot anti-monotonicity to prune unnecessary candidate values of k , and this pruning policy can prevent a significant number of checking operations. This is validated in Table 6 that the number of route hotspots drops sharply with the increase of k . FastRH is faster than GreedyRHH and GreedyRHP, since FastRH uses two pruning rules to speed up the hotspot finding process. PFastRH is even faster than FastRH because PFastRH is a parallel method. CuTS is the fastest one since it only uses the route information. CuTSG is faster than GreedyRH and its variations, as it doesn't consider the patterns in routes and thus avoids the

TABLE 4
An Example of NRH-Index.

ID	Pattern	k	Edges
1	$\langle a, b \rangle$	2	$(v_1, v_2), (v_1, v_3), (v_2, v_3), (v_2, v_4), (v_2, v_5), (v_3, v_4), (v_3, v_5)$
2	$\langle a, b \rangle$	3	$(v_1, v_2), (v_1, v_3), (v_2, v_3), (v_2, v_4), (v_2, v_5), (v_3, v_4), (v_3, v_5)$
...

sequential pattern mining process. Please note that another reason why CuTS and CuTSG are faster than PFastRH is that we only use 8 threads in our experiments (according to our machine), which limits the speed-up ratio of the parallelized algorithm.

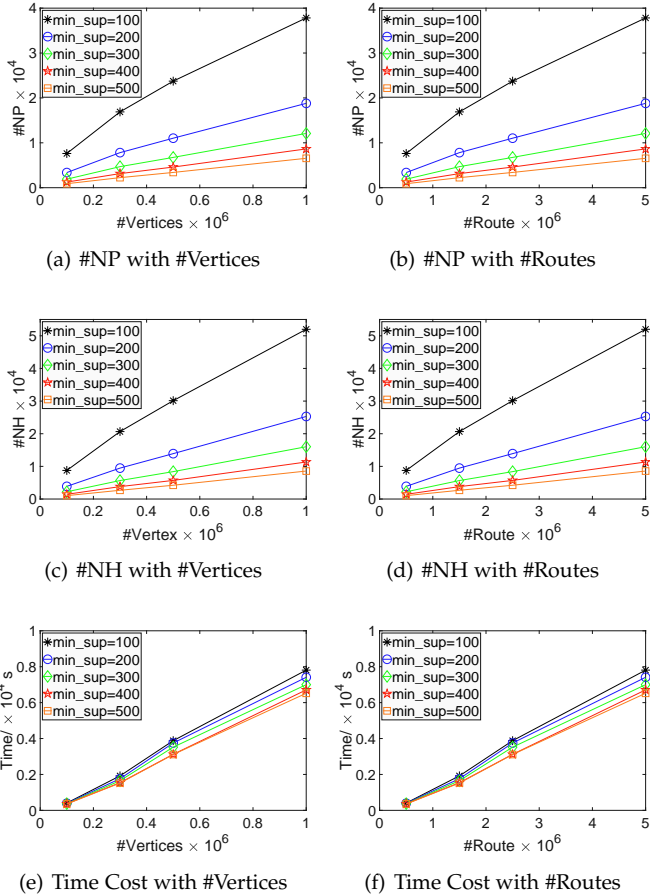


Fig. 7. The #NP, #NH and Time Cost on sampled subgraphs of CN with different values of min_sup , where $\#Vertices=1.0 \times 10^5, 3.0 \times 10^5, 5.0 \times 10^5, 1.0 \times 10^6$ and $\#Routes=5.0 \times 10^5, 1.5 \times 10^6, 2.5 \times 10^6, 5.0 \times 10^6$ respectively. The Time Cost is obtained by PFastRH.

As shown in Figure 7, with the increasing number of vertices and routes for CN, the time cost increases almost linearly. This is because a greater number of vertices and routes bring in a greater number of patterns and route hotspots, i.e., #NP and #NH, and thus it would take longer time. The near-linear increasing demonstrates that the distribution of route hotspots in the network is near-uniform.

7.2.4 Indexing and Query Processing

We evaluate the scalability and the efficiency of RH-Index. Since it's trivial for the situation when a hotspot owns just one route, we fix $min_sup = 2$. Please note that a larger min_sup can result in smaller RH-Index. For comparing

TABLE 5
Index statistics (Size in Megabytes and Time in seconds)

	GW	WB	YELP	CN
#MNRH	47,924	44,035	27,102	39,918
#MRH	13,336	28,181	19,776	15,403
#TNRH	382	865	945	8,392
#TRH	294	691	934	8,166
#NP	8.6×10^5	1.6×10^5	1.8×10^5	6.7×10^5
#NH	2.1×10^6	3.8×10^6	5.4×10^5	8.2×10^5
#AVH	20.6	95.5	42.6	15.5
#AEH	25.0	131.2	51.3	17.6
#ARH	4.0	19.7	9.5	32.2
#ALH	6.31	2.82	4.00	3.39
$max\{k\}$	6	7	20	12

TABLE 6
#NH for different values of k on all datasets

k	GW	WB	YELP	CN
$k = 2$	1.7×10^6	3.6×10^5	5.2×10^5	7.8×10^5
$k = 3$	4.6×10^5	1.6×10^5	1.6×10^4	3.1×10^4
$k = 4$	1.6×10^4	76	118	6,163
$k = 5$	370	11	29	1,707
$k = 6$	9	4	18	602

with RH-Index, we propose a naive method called NRH-Index, where each hotspot is stored as a set of edges in the row of a big table. NRH-Index can also support querying by patterns via traversing rows in that big table. Table 4 shows an example of NRH-Index with pattern $\langle a, b \rangle$ and $k = 2, 3$. Please note that the IDs in this example are assigned randomly.

The indexing performance of RH-Index is shown in Table 5, where "#MNRH" and "#MRH" are the cost of the main memory for NRH-Index and RH-Index respectively; "#TNRH" and "#TRH" denote the time cost to build NRH-Index and RH-Index respectively, including both detection time and indexing time; "#AVH", "#AEH", "#ARH", "#ALH" are the average number of vertices, edges, routes and the length of patterns for each hotspot. As shown in Table 5, over all the datasets, the memory cost of NRH-Index (i.e., #MNRH) is much larger than that of RH-Index (i.e., #MRH), and similarly, the time cost of NRH-Index (i.e., #TNRH) is also greater than that of RH-Index (i.e., #TRH). The main reason is that RH-Index uses a tree-based structure to index the route hotspots, which can save both space and time cost compared with the table-based structure.

Table 6 shows how k affects the performance on different datasets. Since the evaluated RH-Index is obtained by fixing $min_sup = 2$, one can obtain the maximum number of hotspots when $k = 2$, and this number decreases when k increases. The maximum value of k is 12 for CN. It can also be observed that the number of hotspots drops quickly when k increases, which is beneficial for GreedyRHH.

We compare the average query time of RH-Index and NRH-Index in Table 7. The query time includes the time to find patterns and the time to recover hotspots. It is shown that over all datasets, the time cost for querying RH-Index is much lower than that for NRH-Index, indicating that RH-Index is much more efficient even with far less memory

TABLE 7
Query time(In milliseconds)

DataSet	GW	WB	YELP	CN1	CN2	CN
With NRH-Index	0.098	9.468	0.424	1.789	3.786	4.715
With RH-Index	0.043	0.864	0.224	0.067	0.090	0.122

requirement. In addition, the query time generally increases with the size of the networks.

7.3 Case Study

Two case studies from CN and YELP are given to explain the effectiveness of our method.




Fig. 8. Case Study for (DB, DM, SP) and $k = 3$ in CN, where the dashed rectangles represent the research areas and the colored and dashed lines represent routes in the hotspot. We omit the edges.

Figure 8 shows a route hotspot for pattern $\langle DB, DM, SP \rangle$ and $k = 3$, where DB, DM and SP stand for Database, Data Mining and Sequential Pattern respectively. This hotspot shows how the research highlights evolved in the 2000s. During the time, the techniques for database were applied to data mining, and then applied to sequential pattern mining, and the detected scholars in these domains had close cooperations.

Figure 9 shows interesting relationships among dance, event planning, and buffets. Dance1 and Dance2 denote two different vertices but have the same attribute dance, and so is for Hotel1 and Hotel2 affiliated with the same attribute Hotel. From the observed sequential pattern $\langle Dance, Event Planning, Buffets \rangle$, we can know that dance, event planning and buffets are tightly coupled activities, and people who participate these activities usually first dance, then followed by the planning on the subsequent activities, and finally go on for dinner. Note that, vertices whose attributes don't exist in this pattern may also appear in the hotspot (e.g., attributes Hotel and Night Life don't appear in the pattern but appear in the hotspot). The reason is that our definition is relaxed in the attributes so that triangles can be formed easily.




Fig. 9. Case Study for (Dance, Event Planning, Buffets) and $k = 3$ in YELP, where the colored lines represent routes in the hotspot. We omit the edges.

In contrast to the route hotspot for $\langle Dance, Event Planning, Buffets \rangle$ and $k = 3$, as shown in Figure 9, we cannot

find the route hotspot for $\langle Dance, Buffets, Event Planning \rangle$ and $k = 3$. The reason is that, people don't exhibit such kind of behaviors on YELP.

8 CONCLUSION

In this paper, we study the problem of *finding route hotspots in large labeled networks*. We introduce the definition of the route hotspot and propose a novel approach *FastRH* to find them, based on pattern anti-monotonicity and hotspot anti-monotonicity properties. To preserve the detected hotspots and support fast user querying by patterns, we design a novel index, called *RH-Index*. Extensive experiments performed on large real-world networks demonstrate the effectiveness and efficiency of the proposed methods.

ACKNOWLEDGMENTS

This work was supported by the National Key Research and Development Program of China (No. 2017YFB0803301) and Natural Science Foundation of China (No. 61976026, No. U1836215) and 111 Project (B18008).

REFERENCES

- [1] Y. Peng, Y. Zhang, W. Zhang, X. Lin, and L. Qin, "Efficient probabilistic k-core computation on uncertain graphs." in *ICDE*, 2018, pp. 1192-1203.
- [2] Z. Shang, G. Li, and Z. Bao, "Dita: A distributed in-memory trajectory analytics system." in *SIGMOD*, 2018, pp. 1681-1684.
- [3] J. Dai, B. Yang, C. Guo, and Z. Ding, "Personalized route recommendation using big trajectory data," in *ICDE*, 2015, pp. 543-554.
- [4] E. Lopez-Rojas, A. Elmir, and S. Axelsson, "Paysim: A financial mobile money simulator for fraud detection," in *EMSS*, 2016, pp. 249-255.
- [5] DBLP, "Dblp computer science bibliography," <http://dblp.uni-trier.de>.
- [6] J. Wang and J. Cheng, "Truss decomposition in massive networks." *PVLDB*, pp. 812-823, 2012.
- [7] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen, "Discovery of convoys in trajectory databases," *PVLDB*, pp. 1068-1080, 2008.
- [8] K. Zheng, Y. Zheng, N. J. Yuan, and S. Shang, "On discovery of gathering patterns from trajectories," in *ICDE*, 2013, pp. 242-253.
- [9] J. Pei, J. Han, B. Mortazavi-Asl, H. Pinto, Q. Chen, U. Dayal, and M.-C. Hsu, "Pre-fixspan: Mining sequential patterns efficiently by prefix-projected pattern growth." in *ICDE*, 2001, pp. 215-224.
- [10] X. Huang and L. V. Lakshmanan, "Attribute-driven community search," in *PVLDB*, 2017, pp. 949-960.
- [11] R. Agrawal and R. Srikant, "Mining sequential patterns." in *ICDE*, 2005, pp. 3-14.
- [12] R. Srikant and R. Agrawal, "Mining sequential patterns: Generalizations and performance improvements." in *EDBT*, 1996, pp. 3-17.
- [13] M. J. Zaki, "Spade: An efficient algorithm for mining frequent sequences," *Machine Learning*, vol. 42, no. 1/2, pp. 31-60, 2001.
- [14] J. Han, J. Pei, B. Mortazavi-Asl, Q. Chen, U. Dayal, and M.-C. Hsu, "Freespan: Frequent pattern-projected sequential pattern mining." in *SIGKDD*, 2000, pp. 355-359.
- [15] M. Riondato and E. M. Kornaropoulos, "Fast approximation of betweenness centrality through sampling." in *WSDM*, 2014, pp. 413-422.
- [16] T. Calders, C. Garboni, and B. Goethals, "Efficient pattern mining of uncertain data with sampling." in *PAKDD*, 2010, pp. 480-487.
- [17] A. Pietracaprina, M. Riondato, E. Upfal, and F. Vandin, "Mining top-k frequent itemsets through progressive sampling." *DMKD*, pp. 310-326, 2010.
- [18] M. Riondato and E. Upfal, "Efficient discovery of association rules and frequent itemsets through sampling with tight performance guarantees." in *ECML PKDD*, 2012, pp. 25-41.
- [19] —, "Mining frequent itemsets through progressive sampling with rademacher averages," in *KDD*, 2015, pp. 1005-1014.
- [20] H. Li, W. Yi, Z. Dong, Z. Ming, and Y. C. Edward, "Pfp: parallel fp-growth for query recommendation," in *RecSys*, 2008, pp. 107-114.

[21] J. Ge and Y. Xia, "Distributed sequential pattern mining in large scale uncertain databases," in *PAKDD*, 2016, pp. 17–29.

[22] G. Dong and J. Pei, *Sequence data mining*. Springer, 2007, vol. 33.

[23] Y. Fang, R. Cheng, S. Luo, and J. Hu, "Effective community search for large attributed graphs," in *PVLDB*, 2016, pp. 1233–1244.

[24] F. Zhang, Y. Zhang, L. Qin, W. Zhang, and X. Lin, "When engagement meets similarity: efficient (k, r) -core computation on social networks," in *PVLDB*, 2017, pp. 998–1009.

[25] X. Huang, H. Cheng, L. Qin, W. Tian, and J. Yu, "Querying k-truss community in large and dynamic graphs," in *SIGMOD*, 2014, pp. 1311–1322.

[26] X. Huang, W. Lu, and L. Lakshmanan, "Truss decomposition of probabilistic graphs: Semantics and algorithms," in *SIGMOD*, 2016, pp. 77–90.

[27] Z. Wang, M. Lu, X. Yuan, J. Zhang, and H. Van De Wetering, "Visual traffic jam analysis based on trajectory data," *TVCG*, vol. 19, no. 12, pp. 2159–2168, 2013.

[28] M. R. Vieira, P. Bakalov, and V. J. Tsotras, "On-line discovery of flock patterns in spatio-temporal data," in *SIGSPATIAL*, 2009, pp. 286–295.

[29] C. Zhang, K. Zhang, Q. Yuan, L. Zhang, T. Hanratty, and J. Han, "Gmove: Group-level mobility modeling using geo-tagged social media," in *SIGKDD*, 2016, pp. 1305–1314.

[30] G. Dimitrios, K. Roni, M. Heikki, S. Sanjeev, T. Hannu, and S. R. Sewak, "Discovering all most specific sentences," *TODS*, pp. 140–174, 2003.

[31] J. Tang, J. Zhang, L. Yao, L. Zhang, and Z. Su, "Arnetminer: Extraction and mining of academic social networks," in *SIGKDD*, 2008, pp. 990–998.

[32] E. Cho, S. Myers, and J. Leskovec, "Friendship and mobility: User movement in location-based social networks," in *SIGKDD*, 2011, pp. 1082–1090.


[33] S. Lloyd, "Least squares quantization in pcm," *TIT*, pp. 129–137, 1982.

[34] J. Zhang, B. Liu, J. Tang, T. Chen, and J. Li, "Social influence locality for modeling retweeting behaviors," in *IJCAI*, 2013, pp. 2761–2767.

[35] YELP, "Yelp dataset challenge," https://www.yelp.com/dataset_challenge.

[36] A. Lulli, M. Dell'Amico, P. Michiardi, and L. Ricci, "Ng-dbscan: scalable density-based clustering for arbitrary data," *PVLDB*, pp. 157–168, 2016.

[37] M. Ester, H.-P. Kriegel, J. Sander, X. Xu *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *KDD*, 1996, pp. 226–231.



on Circuits and Systems for Video Technology.

Lingyang Chu received the doctoral degree from the University of Chinese Academy of Sciences, Beijing, China, in 2015. He is currently a postdoctoral fellow in the School of Computing Science, Simon Fraser University, Canada. His research interests include large scale network analysis, information retrieval, and cross modality information processing. He has published several papers in refereed conference proceedings and journals, such as KDD, VLDB, IEEE Trans. on MultiMedia, and the IEEE Transactions



gineering.

Zhefeng Wang received the BE degree from the University of Science and Technology of China, China, in 2012. He received the PhD degree in the School of Computer Science and Technology, University of Science and Technology of China. His research interests include social network, social media analysis, machine learning, and text mining. He has published several papers in refereed conference proceedings and journals such as SIGIR, KDD, IJCAI, and the IEEE Transactions on Knowledge and Data En-



Innovation Award, a Research Contributions Award from the IEEE International Conference on Data Mining (2003), and a Technical Achievement Award from the IEEE Computer Society (2013). He is a fellow of the IEEE and the ACM.

Philip S. Yu received the PhD degree in electrical engineering from Stanford University. He is a distinguished professor in computer science at the University of Illinois, Chicago, and is the Wexler chair in Information Technology. His research interests include big data, data mining, data stream, database, and privacy. He was the editor-in-chief of the IEEE Transactions on Knowledge and Data Engineering and the ACM Transactions on Knowledge Discovery from Data. He received the ACM SIGKDD 2016

Mingtao Lei received the bachelor degree in Information Security from Beijing University of Posts and Telecommunications in 2012. He is currently a PHD student in Key Laboratory of Trustworthy Distributed Computing and Service (BUPT). His research interests include data mining, graph mining and pattern mining.



Xi Zhang received the PhD degree in Computer Science from Tsinghua University. He is an associate professor in Beijing University of Posts and Telecommunications, and is also the vice director of Key Laboratory of Trustworthy Distributed Computing and Service, Ministry of Education, China. He was a visiting scholar at the University of Illinois at Chicago. His research interests include data mining and computer architecture. He is a member of IEEE.



Binxing Fang received his PhD from Harbin Institute of Technology, China in 1989. He is a member of the Chinese Academy of Engineering and a professor in School of Cyberspace Security at Beijing University of Posts and Telecommunications. He was the chief scientist of State Key Development Program of Basic Research of China. His current interests include big data and Cybersecurity.