

“© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Efficient Radius-bounded Community Search in Geo-social Networks

Journal:	<i>Transactions on Knowledge and Data Engineering</i>
Manuscript ID	TKDE-2020-10-1186
Manuscript Type:	Regular
Keywords:	K-core, Geo-social Network, Community search, Diversification

SCHOLARONE™
Manuscripts

Cover Letter

Efficient Radius-bounded Community Search in Geo-social Networks

Kai Wang, Shuting Wang, Xin Cao, Wenjie Zhang and Lu Qin

Dear Editor and Fellow Reviewers,

We are writing to submit our paper “Efficient Radius-bounded Community Search in Geo-social Networks” to the prestigious journal, IEEE TKDE, for possible publication.

This submission is a substantial extension of the conference version “Efficient Computing of Radius-Bounded k -Cores” which is accepted by the International Conference on Data Engineering 2018 (IEEE ICDE 2018).

We summarize the major new materials in this submission as follows:

1. We have updated the abstract, introduction, related work, and conclusion for the new materials.
2. We have added Section 4 to discuss the *diversified radius-bounded k -core search* problem in geo-social networks which is commonly encountered in real applications. Compared with the existing problem -- *radius-bounded k -core search* studied in our conference paper, utilizing this new model can find distinctive radius-bounded k -cores with rich information. We have devised new algorithms to solve this problem which utilize the property of diversification. In addition, we not only have carefully analysed the theoretical guarantees of the new algorithms, but also have proposed several dedicated pruning techniques to further enhance the performance of them.
3. In Section 5, we have added 5 new experiments and the results are shown in Fig-6, Fig-13, Fig-14, Fig-15, and Fig-16. The effectiveness the diversified radius-bounded k -core model is evaluated in Fig-6. The efficiency of the new algorithms for solving the diversified radius-bounded k -core search problem is evaluated in Fig-13, Fig-14, Fig-15, and Fig-16 regarding different parameters.

Please kindly let us know if further information is required. Thank you for your valuable reviews and comments!

Best Regards,

Kai Wang, Shuting Wang, Xin Cao, Wenjie Zhang and Lu Qin

Efficient Radius-bounded Community Search in Geo-social Networks

Kai Wang, Shuting Wang, Xin Cao, Wenjie Zhang, and Lu Qin

Abstract—Driven by real-life applications in geo-social networks, we study the problem of computing radius-bounded k -cores (RB- k -cores) that aims to find communities satisfying both social and spatial constraints. In particular, the model k -core (i.e., the subgraph where each vertex has at least k neighbors) is used to ensure the social cohesiveness, and a radius-bounded circle is used to restrict the locations of users in an RB- k -core. We explore several algorithmic paradigms to compute RB- k -cores, including a triple-vertex-based paradigm, a binary-vertex-based paradigm, and a paradigm utilizing the concept of rotating circles. The rotating-circle-based paradigm is further enhanced by several pruning techniques to achieve better efficiency. In addition, to find representative RB- k -cores, we study the diversified radius-bounded k -core search problem, which finds t RB- k -cores to cover the most number of vertices. We first propose a baseline algorithm that identifies the distinctive RB- k -cores after finding all the RB- k -cores. Beyond this, we design algorithms that can efficiently maintain the top- t candidate RB- k -cores and also achieve a guaranteed approximation ratio. Experimental studies on both real and synthetic datasets demonstrate that our proposed techniques can efficiently compute (diversified) RB- k -cores. Moreover, our techniques can be used to compute the minimum-circle-bounded k -core and significantly outperform the existing techniques.

Index Terms—K-core, Geo-social network, Community search, Diversification

1 INTRODUCTION

With the wide availability of wireless communication techniques and GPS-equipped mobile devices (e.g., smartphones and tablets), people can now easily access the internet. This leads to the emergence of geo-social networks, such as Twitter and Foursquare, where social networks are combined with users' geo-spatial information. Consequently, retrieving subgraphs with high cohesiveness in geo-spatial social networks has become a popular research topic recently [14], [42], [44].

In this paper, we study the problem of efficiently computing *radius-bounded cohesive subgraphs* in a geo-spatial social network G (or abbreviated as a geo-social network). That is, given a vertex q in G and a radius r , find all cohesive subgraphs g of G such that g contains q and all vertices in g fall into a circle with the radius r . There are many types of cohesive subgraph models in the literature such as k -core [32], k -truss [8] and clique [25]. While our proposed framework generally works for various cohesive subgraphs, in this paper, we present our work restricted to a specific cohesive subgraph k -core [32] where each vertex has at least k neighbors.

Applications. The problem of computing *radius-bounded k -cores*, namely RB- k -cores, has many real real-life applications. On social platforms like Facebook and Twitter, personalized event recommendation is an essential part. For ex-

ample, “Events for you” is a valuable Facebook component that recommends events to users based on their locations and social connection. Nevertheless, the current technology cannot provide a service of Events-For-You based on users' arbitrary requests. For example, the Events-For-You based on Leo's following request cannot be accommodated by the current technology. In this example, Leo wants to hold a party (an Events-For-You activity) to play board games (e.g., Monopoly, Uno, and Risk) by gathering a group of people who are not living far away (say, bounded by a circle with a radius r) and each of whom has many friends in the group (say, at least k friends). Figure 1 shows a geo-social network where vertices represent users, edges represent friendships, and locations represent the home locations of users. If we set $r = 3$ and $k = 3$, there are two RB- k -cores recommended to Leo as illustrated by the shadow area, i.e., $\{\text{Leo, Ken, Jim, Adam}\}$ and $\{\text{Leo, Bill, Frank, Bob, Lee}\}$. This is a typical example of computing RB- k -cores. In addition, when users propose query requests, there may exist a large number of RB- k -cores that satisfy the query constraints. In these circumstances, the system only needs to recommend representative communities with rich information to users.

Moreover, as studied in [14], [27], people with close social relationships tend to purchase in places that are also close. For this application, the locations of users represent places in a geo-social network. To boost sales figures, advertisement messages can be sent to the RB- k -core of customers. For instance, if we want to promote an item A , the system can advertise A to the RB- k -core members (customers) using the query point (customers) who purchased A .

Existing Studies. Several studies over community retrieval exist in geo-social networks, but they are all different from our problem. In the literature, various models including k -core [32], k -truss [8], and clique [25] have been studied to

- Kai Wang, Xin Cao, and Wenjie Zhang are with the University of New South Wales, Australia. E-mail: kai.wang@unsw.edu.au, xin.cao@unsw.edu.au, zhangw@cse.unsw.edu.au.
- Shuting Wang is with Zhejiang Gongshang University, China. Email: wangshuting@zjgsu.edu.cn.
- Lu Qin is with the University of Technology Sydney, Australia. Email: lu.qin@uts.edu.au.




Figure 1: A geo-social network

retrieve cohesive subgraphs without considering the spatial information of users. Thus, these models are not applicable to compute RB- k -cores. For example, in Figure 1, Mark will be added into the community formed by {Leo, Ken, Adam, Jim} if we use the model k -core for $k = 3$ though Mark is far away from the other users. On the other hand, [16], [38] find a group of spatial objects without considering the network information, and thus they also cannot solve our problem RB- k -cores. In Figure 1, Roy will be added into the community formed by {Leo, Ken, Jim, Adam} using spatial information (bounded by a circle) only, but the network connections between Roy and the other people are fragile.

The most closely related works can be found in [14], [42], [44], which consider both social constraint (network structure) and spatial constraint in retrieving communities. In particular, they all use k -core to ensure the social (structure) cohesiveness of communities in a network. Zhang et al. [42] study the community detection problem, which uses pairwise similarity (distance) between each pair of vertices to ensure the spatial cohesiveness of communities while computing the maximum k -core or all maximal k -cores. As studied in [36], our problem is inherently different from the problem in [42]; that is, the generated results are very different. Moreover, our problem RB- k -cores is PTIME, while the problem in [42] is NP-hard. Zhu et al. in [44] study the problem of finding the maximum k -core in a given rectangle containing a query vertex. They also study the problem of finding the k -core with exact (or no less than) c vertices such that the longest distance from these vertices to q is minimized. These problems in [44] are also different from ours. Fang et al. in [14] study the problem of computing the k -core containing the query vertex covered by the smallest circle. While the problem in [14] is different from our problem RB- k -cores, as a byproduct, our techniques can be applied to the problem in [14] and can achieve a speed-up around twice.

Challenges. The main challenge of efficiently computing RB- k -cores is threefold.

- 1) The location of the radius-bounded circle of a RB- k -core is unknown. Therefore, it is a challenge to enumerate such circles efficiently.
- 2) In the process of finding all the RB- k -cores, it is cost-prohibitive to construct and verify the candidate subgraphs individually. Therefore, it is important to reuse intermediate computation results and explore possible cost-sharing, which is challenging.
- 3) Given a query request, there may exist many RB- k -cores

in the result set. Thus, when users only want to retrieve the representative RB- k -cores rather than all of them, it is also challenging to identify such distinctive RB- k -cores with rich information efficiently.

Contributions. We first explore three paradigms to retrieve all the RB- k -cores in this paper. The first paradigm is triple-vertex-based algorithm (TriV) inspired by [14]. It proposes to firstly generate all candidate circles containing q , secondly check the corresponding radius to verify the given radius bound, and then compute the maximum k -core for the vertices in each candidate circle. To avoid generating too many candidate circles or missing results, TriV generates all the candidate circles by enumerating all the triple-vertex- and binary-vertex-combinations.

To reduce the number of candidate circles, we further propose the binary-vertex-based algorithm (BinV). In BinV, we effectively use the given parameter r and only generate the circles with a radius r such that the circle arc passes a pair of vertices in G (for every pair). Generally, for each pair of vertices, we have at most two such circles. This guarantees to generate $O(n^2)$ candidate circles and reduces $O(n^3)$ candidate circles in TriV to $O(n^2)$.

We can observe that there are many reusable intermediate computation results in the process of finding RB- k -cores. The third paradigm is to share computation costs among the computation of RB- k -cores. To effectively share the computation, we design the rotating-circle-based algorithm (RotC) so that the computation in BinV can be shared among the “adjacent” circles. Specifically, we fix a vertex u for each vertex in G then check the remaining vertices v such that for each pair u and v , we use BinV to generate the two circles with radius r (maybe degenerate to one if r is half of the distance between u and v). Then, we will share the computation among the adjacent circles.

To find distinctive RB- k -cores, we study the diversified RB- k -core search problem to find t representative RB- k -cores that cover the most number of vertices. This problem can be recognized as the max k -cover problem [15], which is NP-hard. A simple greedy solution DivBS is that after obtaining all the RB- k -cores, we run t iterations and identify the RB- k -core, which covers the maximum number of uncovered vertices for each iteration. Although DivBS can have a good approximation ratio, the main drawback of DivBS is that it needs to compute all the RB- k -cores firstly, which isolate the computation of RB- k -cores and the finding of diversified top- t RB- k -cores. Towards this issue, we propose the DivRotC⁺ algorithm that maintains the top- t candidates in the RB- k -core computing process and achieves a guaranteed approximation ratio. Also, several useful pruning techniques are deployed into DivRotC⁺ to enhance the performance further.

Our principal contributions are summarized as follows.

- We propose the RB- k -core model and develop a novel paradigm to compute RB- k -cores to share computations among different RB- k -cores.
- We propose several new optimization techniques that speed up the computation of finding all the RB- k -cores.
- We propose efficient algorithms along with several dedicated pruning techniques to find distinctive RB- k -cores with rich information.

- Extending our algorithms to the problem in [14] can achieve a speed-up around twice.
- We conduct comprehensive experiments on real geo-social networks to evaluate our algorithms.

Organization. The rest of the paper is organized as follows. Section 2 presents the preliminaries. Section 3 introduces our techniques to solve the RB- k -core search problem. The study of the diversified RB- k -core search problem is presented in Section 4. Section 5 reports experimental results. Section 6 reviews the related work. Section 7 concludes the paper.

2 PROBLEM DEFINITION

In this section, we formally introduce fundamental concepts and definitions. Mathematical notations used throughout this paper are summarized in Table 1.

Table 1: The summary of notations

Notation	Definition
G	a geo-social graph
G_k^r, B	a RB- k -core
u, v	vertices in the geo-social graph
$deg_G(v)$	the degree of vertex v in G
$N_G(v)$	the set of neighbors of vertex v in G
$d(u, v)$	the Euclidean distance between u and v
$O(c, \gamma)$	a circle centered at c with radius γ
$g(c, \alpha)$	a square centered at c with side length α
X, S	a set of vertices
$G(S)$	an induced subgraph of S
\mathcal{R}	the result RB- k -core set
\mathcal{D}	a set of diversified RB- k -cores

Our problem is defined over a geo-social graph $G(V, E)$, where $V(G)$ denotes the vertex set, and $E(G)$ denotes the edge set. The vertices represent the social network users and the edges represent their relationships in geo-social networks. Each vertex $v \in V(G)$ has a location $(v.x, v.y)$ which denotes the position of v along x - and y -axis in a two-dimensional space and the vertices are static in our problem. The Euclidean distance between u and v is denoted as $d(u, v)$. We denote the set of neighbors of each vertex v in G by $N_G(v) = \{u \in V(G) \mid (v, u) \in E(G)\}$ and the degree of vertex v by $deg_G(v) = |N_G(v)|$. We denote a circle centered at c with radius γ as $O(c, \gamma)$. Given a set of vertices $S \subseteq V(G)$, we use $G(S)$ to denote an induced subgraph of G formed from S such that $G(S) = (S, \{(u, v) \in E(G) \mid u, v \in S\})$.

Before formally defining the problem, we first introduce the following critical concepts to describe the social constraint and the spatial constraint.

Definition 1 (k -Core). Given a graph G and a positive integer k , the k -core of G denoted as H_k is the maximal subgraph of G , where $deg_{H_k}(v) \geq k$, for each $v \in V(H_k)$.

Based on the k -core concept, we ensure the social constraint by restricting the minimal degree of vertices in a RB- k -core. Note that our proposed solutions can be easily adapted to other cohesive structure concepts (e.g., k -truss [8], clique [25]), which can be used to define the social constraint from different perspectives.

Definition 2 (Minimum Covering Circle (MCC)). Given a set of vertices S , the minimum covering circle of S is the circle, which encloses all the vertices $v \in S$ with

the smallest radius. We call the vertices which lie on the boundary of an MCC the boundary vertices.

After introducing k -core and MCC, we are ready to define the RB- k -core as follows.

Definition 3 (Radius-Bounded k -Core). Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k and a query radius r , a subgraph of G denoted by G_k^r is a Radius-Bounded k -Core, if it satisfies the following constraints:

- 1) **Connectivity constraint.** $G_k^r \subseteq G$ is connected and contains q ;
- 2) **Social constraint.** $\forall v \in V(G_k^r), deg_{G_k^r}(v) \geq k$;
- 3) **Spatial constraint.** The MCC of $V(G_k^r)$ has a radius $r' \leq r$;
- 4) **Maximality constraint.** There exists no other super-graph $G_k^{r'} \supset G_k^r$ satisfying (1), (2), and (3).

Problem Statement (RB- k -core Search) Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k , and a query radius r , our RB- k -core search problem aims to return all the RB- k -cores in G .

We then define coverage as follows.

Definition 4 (Coverage). Given a set of RB- k -cores $\mathcal{D} = \{B_1, B_2, \dots\}$ in G , the coverage of \mathcal{D} denoted by $cov(\mathcal{D})$ is the set of vertices covered by the RB- k -cores in \mathcal{D} , i.e., $cov(\mathcal{D}) = \bigcup_{B \in \mathcal{D}} V(B)$.

Problem Statement (Diversified top- t RB- k -core Search) Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, positive integers k and t , and a query radius r , our diversified top- t RB- k -core search problem (or short as diversified RB- k -core search problem) aims to return a set \mathcal{D} of RB- k -cores in G , such that (1) each $B \in \mathcal{D}$ is a RB- k -core, (2) $|\mathcal{D}| \leq t$, and (3) $|cov(\mathcal{D})|$ is maximized. The set \mathcal{D} contains the diversified top- t RB- k -cores.

Note that the diversified RB- k -core search problem is NP-hard since it can be recognized as the max k -cover problem [15] after obtaining all the RB- k -cores.

Remark. According to the spatial constraint in Definition 3, apparently, if the distance between a vertex v and the query vertex q is larger than $2r$, v cannot be included in any RB- k -cores. We call such vertices faraway vertices, and we can first remove all these vertices from G . We can also safely remove all the vertices that are not in the k -core of G containing q because of the social constraint in Definition 3. We use G_k to denote a connected subgraph of G which is a k -core containing q and for each vertex v in G_k , $d(q, v) \leq 2r$. We use $n = |V(G_k)|$ to represent the number of vertices and $m = |E(G_k)|$ to represent the number of edges of G_k in the following sections. Note that we use Euclidean distance to measure the proximity between two users in this paper, and it can be easily replaced by other measurements (e.g., the geographical distance).

Example 1. Consider a geo-social graph G in Figure 2(a). Suppose Q is the query vertex, given $k = 2$ and $r = 1$, to solve the RB- k -core search problem, we want to find all RB- k -cores from this geo-social graph. We can safely remove vertex A because $d(A, Q) > 2r = 2$ and vertex I because I is not in the 2-core of G . Then we can obtain the candidate geo-social subgraph G_k which is shown in Figure 2(b). As a result, we can find two




Figure 2: An example of the geo-social graph

RB- k -cores $G(S_1)$ and $G(S_2)$, where $S_1 = \{Q, C, J\}$ and $S_2 = \{Q, D, E, F\}$. Furthermore, given $t = 1$, the result of the diversified top- t RB- k -core search problem is S_2 .

3 RB- k -CORE SEARCH

In this section, we introduce our algorithms to solve the RB- k -core search problem.

3.1 The Triple-Vertex-Based Algorithm

Firstly we introduce the triple-vertex-based algorithm (TriV) which is designed based on the Exact algorithm in [14]. TriV lies on the following lemma.

Lemma 1. [12] Given a set S ($|S| \geq 2$) of vertices, the MCC of S can be determined by two or three vertices in S which lie on the boundary of the circle. If two vertices determine it, then the line segment connecting these two vertices must be the circle's diameter. If three vertices determine it, then the triangle consisting of those three vertices is not obtuse.

By Lemma 1, the MCC of a RB- k -core should have two or three vertices lying on its boundary, which are called boundary vertices. Thus, we can enumerate all candidate triple-vertex-combinations and binary-vertex-combinations, then check whether the subgraph enclosed by the circle fixed by the enumerated boundary vertices is a RB- k -core. The details of TriV can be found in [36].

Complexity Analysis. The time complexity of TriV is $O(n^3 \cdot m)$. This is because in TriV, we need to verify $O(n^3)$ candidate triple-vertex-combinations and $O(n^2)$ binary-vertex-combinations. For each combination, we need $O(m)$ time cost to verify the existence of the k -core. Thus, the total time cost of TriV is bounded by $O(n^3 \cdot m)$.

3.2 The Binary-Vertex-Based Algorithm

The major issue of TriV is that we need to verify $O(n^3 + n^2)$ candidate subgraphs based on all triple-vertex- and binary-vertex-combinations. Here we introduce a binary-vertex-based algorithm that only needs to verify $O(n^2)$ candidate subgraphs to solve the RB- k -core search problem.

Based on the definition of RB- k -core, given a query radius r , an obvious observation is that for each RB- k -core in a geo-social graph G , it should be enclosed in at least one circle with radius r . A straightforward approach to finding all RB- k -cores verifies all the circles with radius r in the two-dimensional space. Obviously, there are too many circles with radius r sharing the same RB- k -core in this approach. In other words, for each RB- k -core, we need to ensure that

there is at least one circle with radius r enclosing it is checked. This can decrease the number of candidate circles significantly. Firstly, we define the binary-vertex-bounded circle as below.

Definition 5 (Binary-Vertex-Bounded Circle). Given two vertices u and v , we call all circles having u and v lying on the boundary the binary-vertex-bounded circles. A set of binary-vertex-bounded circles with radius r which takes u and v as bounded vertices is denoted as $W_r(u, v)$.

Lemma 2. [17] Given two vertices u and v and a radius r ($r \geq d(u, v)$), we have:

$$|W_r(u, v)| = \begin{cases} 1, & \text{iff. } d(u, v) = 2r, \\ 2, & \text{iff. } d(u, v) < 2r. \end{cases} \quad (1)$$

Lemma 3. [36] Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k , and a query radius r , for each RB- k -core G_k^r , all the vertices in $V(G_k^r)$ should be enclosed in at least one binary-vertex-bounded circle with radius r which takes u and v as the boundary vertices where $u, v \in V(G_k)$.

By Lemma 2, two vertices can bound one/two circles with a given radius r . Based on Lemma 3, we can get all the RB- k -cores in G by verifying all the binary-vertex-bounded circles bounded by vertices in $V(G)$ with radius r . Hence, a more efficient algorithm BinV can be designed by verifying $O(n^2)$ candidate subgraphs constructed from the corresponding binary-vertex-bounded circles, rather than $O(n^3 + n^2)$ candidate subgraphs as in TriV. The details of BinV can be found in [36].

Complexity Analysis. The time complexity of BinV is $O(n^2 \cdot m)$. This is because, in BinV, we need to verify all the binary-vertex-bounded circles generated from candidate binary-vertex-combinations, which needs $O(n^2 \cdot m)$ time in total.

3.3 The Rotating-Circle-Based Algorithms

Although the BinV algorithm improves TriV a lot by reducing the number of candidate subgraphs, it is still not efficient enough. Reviewing the process of BinV, we can observe that the corresponding candidate subgraph is constructed and verified individually for each binary-vertex-combination. There are $O(n^2)$ candidate graphs that need to be constructed, and each of the verification processes takes $O(m)$ time. This motivates us to develop a better algorithm to reduce these candidate subgraphs' construction and verification costs.

In this section, we first present the rotating-circle-based algorithm (RotC), which improves the BinV algorithm by exploring possible cost-sharing in the subgraph construction and verification process. Next, we employ non-trivial pruning techniques to improve the RotC algorithm and propose the optimized rotating-circle-based algorithm (RotC⁺).

3.3.1 The Algorithm RotC

Reviewing Lemma 3, we can find all the RB- k -cores in G by verifying all the candidate subgraphs constructed from corresponding binary-vertex-bounded circles. Considering Example 1, Figure 3(a) is a screenshot of all the binary-vertex-bounded circles which take F as one of the boundary

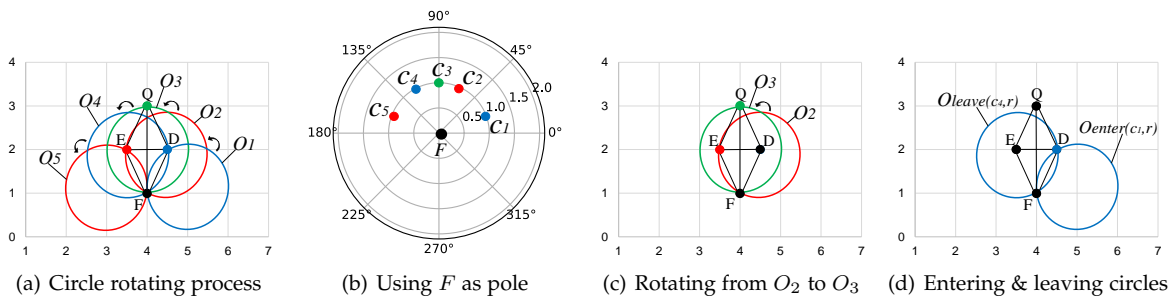


Figure 3: An example of Rotating-Circle-Based Algorithms (using vertex F as the pole)

vertices. In the BinV algorithm, we need to verify the candidate graphs enclosed by these binary-vertex-bounded circles one by one. Now we consider putting these binary-vertex-bounded circles into a polar coordinate system using F as the *pole*, and sorting these binary-vertex-bounded circles according to their centers' polar angles. Figure 3(b) shows the centers of binary-vertex-bounded circles in the polar coordinate system, and we can obtain a list of sorted circles $L = \{O_1, O_2, O_3, O_4, O_5\}$. Specifically, in Figure 3(c), O_2 and O_3 are two adjacent binary-vertex-bounded circles. We denote the vertex sets which O_2 and O_3 enclosed as X_2 and X_3 , respectively. We can observe that for these two induced subgraphs $G(X_2)$ and $G(X_3)$ where their binary-vertex-bounded circles are adjacent to each other, $V(G(X_2))$ is only one vertex (q) different from $V(G(X_3))$. Based on this observation, we devise a novel algorithm that shares the construction and verification cost for these candidate subgraphs.

In the construction step, we can construct the candidate graphs incrementally after sorting all the binary-vertex-bounded circles. In the verification process, the degree of vertices is easy to maintain dynamically because the difference of enclosed vertices between adjacent binary-vertex-bounded circles is only one vertex. We can divide the binary-vertex-bounded circles into two groups, entering circles and leaving circles. An entering circle denoted as $O_{enter}(c, r)$ is a circle which brings a new vertex in, and a leaving circle $O_{leave}(c, r)$ is a circle which takes an existing vertex out. For example in Figure 3(d), $O_{enter}(c_1, r)$ is an entering circles which brings vertex D in and $O_{leave}(c_4, r)$ is a leaving circle which takes D out of the candidate graph. So for an entering circle, we can avoid recomputing the degree of enclosed vertices when checking the k -core in a binary-vertex-bounded circle. For a leaving circle, we can maintain the degree of vertices and avoid the computation of checking the k -core because there cannot exist a new k -core while a vertex leaves. The detailed rotating-circle-based algorithm (RotC) is shown in Algorithm 1.

In RotC, we first run the core decomposition algorithm and obtain the k -core G_k of G containing q after removing all the faraway vertices in $V(G)$ (line 2). After that, for each vertex v in $V(G_k)$, we set it as the pole in a polar coordinate system P . For each pole v , we generate a candidate vertex set $Y = \{u \in V(G_k) \mid d(u, v) \leq 2r\}$. Then we combine v with other candidate vertices in Y and construct the corresponding binary-vertex-bounded circles based on Lemma 3. We also record whether it is an entering circle or a leaving

Algorithm 1: RotC

Input: $G(V, E)$: the input graph; q : the query vertex;
 k, r : constraint parameters
Output: \mathcal{R} : a set of RB- k -cores

```

1 initialize  $\mathcal{R} \leftarrow \emptyset$ 
2  $G_k \leftarrow$  the  $k$ -core of  $G$  containing  $q$  after removing
  faraway vertices
3 foreach node  $v \in V(G_k)$  do
4    $C \leftarrow \emptyset$ 
5   foreach node  $u \in V(G_k)$  do
6     if  $u \neq v \wedge d(u, v) \leq 2r$  then
7       compute  $W_r(u, v)$  using  $\{u, v\}$  and  $r$ 
8       put circles in  $W_r(u, v)$  into  $C$ 
9   sort  $C$  in ascending order of centers' polar angles
10  foreach  $O(c, r) \in C$  do
11     $X \leftarrow$  a set of vertices enclosed in  $O(c, r)$ 
12    maintain the degree of vertices in  $X$ 
13    if  $O(c, r)$  is an entering circle then
14      construct  $G(X)$  from  $X$ 
15      if exists a  $G_k^r$  in  $G(X)$  then
16         $\mathcal{R}.update(G_k^r)$ 
17 return  $\mathcal{R}$ 

```

circle for each binary-vertex-bounded circle. After that, we sort all the binary-vertex-bounded circles in ascending order of their centers' polar angles in P (line 8). Then, for each binary-vertex-bounded circle $O(c, r)$, we compute a set X that contains all the vertices enclosed in O and maintain the degrees of these vertices (lines 9-12). Note that we only need to insert/remove different vertices between O and its precedent binary-vertex-bounded circle, and the degrees of vertices in X can be updated correspondingly. If $O(c, r)$ is an entering circle, we construct a candidate graph $G(X)$, which is a subgraph of G_k induced by X (lines 13-14). After that, we verify whether there exists a k -core containing q in $G(X)$. Because the degrees of vertices in X is already maintained, in the cases such as $deg_{G(X)}(q) < k$, we can skip running a core decomposition to verify the existence of k -core in $G(X)$. Otherwise, if a k -core exists and it satisfies the maximality property, we put the k -core into the result set \mathcal{R} (lines 15-16). Finally, we get all the RB- k -cores in \mathcal{R} . As shown in [36], the time complexity of RotC is $O(n^2 \cdot (\log n + m'))$, where $m' \leq m$ and is much smaller than m in practice.

3.3.2 The Algorithm RotC⁺


We continue to introduce the optimized rotating-circle-based algorithm (RotC⁺), which improves RotC significantly by utilizing novel pruning techniques, including

grouping-based pre-process and the in-process pruning rules.

The Pre-Process Pruning. Firstly, we introduce the grouping-based pre-process pruning technique to partition the vertices into groups and filter out unpromising candidate vertices. The pruning technique is based on the following lemma.

Lemma 4. Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k and a query radius r , for each RB- k -core G_k^r in G , the center point c of the MCC $O(c, \gamma)$ of $V(G_k^r)$ should satisfy $d(c, q) \leq r$.

Proof. By Definition 3, for a RB- k -core G_k^r , the MCC $O(c, \gamma)$ of G_k^r should enclose q and satisfy $\gamma \leq r$. Hence we have $d(c, q) \leq r$ and complete the proof. ■



(a) Grouping $O(q, r)$, $\tau = 2$ (b) Verifying a group $g(c, \alpha)$

Figure 4: An example of grouping-based pre-process

Lemma 4 illustrates that all the centers of MCCs of RB- k -cores are in the circle $O(q, r)$. Apparently, the circle $O(q, r)$ can be partitioned into four groups which are squares with size $r \times r$. Similarly, the group with size $r \times r$ can also be partitioned into 4 smaller groups with size $\frac{r}{2} \times \frac{r}{2}$. Hence, given a grouping parameter τ , as shown in Figure 4(a), we can partition the circle from 4 groups with size $r \times r$ to $4 \lceil \frac{r}{\tau} \rceil^2$ groups with size $\tau \times \tau$ iteratively. In each iteration, we halve the group size and prune the groups which do not need further verification. For example, in Figure 4(a), if we set $\tau = \frac{r}{2}$, the pre-process will run 2 iterations in total and in each iteration, we need to verify at most 4 and 16 groups with size $= r$ and $\frac{r}{2}$, respectively.

We proceed to present the verification process of a given group of vertices denoted as $g(c, \alpha)$, where c is the center point and α is the side length. As shown in Figure 4(b), because the longest distance between c and the other points in $g(c, \alpha)$ is $\frac{\sqrt{2}}{2}\alpha$, we can use the circle $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ to enclose all the circles with radius r and centered at a point in $g(c, \alpha)$. In other words, for each circle $O(c', r)$ which centers at $g(c, \alpha)$ as shown in Figure 4(b), $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ can enclose it. Then, we can construct an induced subgraph $G(X)$ of G using X which contains all the vertices enclosed in the circle $O(c, r + \frac{\sqrt{2}}{2}\alpha)$. If there exists no k -core containing q in $G(X)$, we can prune the whole group $g(c, \alpha)$. Otherwise, if the MCC $O(c', \alpha')$ of the k -core $G(X)_k$ containing q has the radius $\alpha' \leq r$, we can mark $G(X)_k$ as a candidate result and prune the whole group $g(c, \alpha)$, because $G(X)_k$ is the only result that can be found using the vertices in $g(c, \alpha)$. Otherwise, if the MCC $O(c', \alpha')$ of $G(X)_k$ has the radius $\alpha' > r$, the RB- k -cores obtained from $g(c, \alpha)$ are subsets of $G(X)_k$. Thus, we add the vertices in $G(X)_k$ into a candidate vertex set and do further check for the group $g(c, \alpha)$. The details of the pre-processing is shown in Algorithm 2.

Algorithm 2: GROUPING-BASED PRE-PROCESS

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint parameters; τ : grouping parameter; \mathcal{R} : candidate result set

Output: G_k : a graph

```

1  $\alpha \leftarrow r; Y \leftarrow g(q, 2r)$ 
2  $G_k \leftarrow$  the  $k$ -core of  $G$  containing  $q$  after removing
   faraway vertices
3 while  $\alpha \geq \tau$  do
4   foreach group  $g(c, 2\alpha) \in Y$  do
5     partition  $g(c, 2\alpha)$  into four groups with size
        $\alpha \times \alpha$  and put them into  $Y_c$ 
6    $Y \leftarrow \emptyset; S \leftarrow \emptyset$ 
7   foreach group  $g(c, \alpha) \in Y_c$  do
8     construct graph  $G(X)$  using  $X$  which contains
       vertices enclosed in  $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ 
9     if exists a  $G_k^r$  in  $G(X)$  then
10       $\mathcal{R}.$ update( $G_k^r$ )
11     else if exists a  $k$ -core  $G(X)_k$  in  $G(X)$  then
12       $Y.$ insert( $g$ )
13      put vertices in  $V(G(X)_k)$  into  $S$ 
14     foreach node  $v \in V(G_k)$  do
15       if  $v \notin S$  then
16         remove vertex  $v$  from  $V(G_k)$ 
17      $\alpha = \alpha/2$ 
18 return  $G_k$ 

```

In-Process Pruning Techniques. During finding RB- k -cores from G_k , we introduce two in-process pruning rules to push the efficiency boundary further.

Pruning Rule 1: Overall Checking. Reviewing the process of the RotC algorithm, we choose a vertex v from $V(G_k)$ as the pole and generate a candidate vertex set $S = \{u \in V(G_k) \mid d(u, v) \leq 2r\}$. Then we construct an induced subgraph $G(S)$ using vertices in S and compute the k -core $G(S)_k$ of $G(S)$ containing q . If $G(S)_k$ does not exist or the vertices in $V(G(S)_k)$ are all enclosed in the MCC of a candidate RB- k -core in \mathcal{R} , we can prune the pole v .

Pruning Rule 2: Circle Filtering. In the RotC algorithm, after choosing the pole v and corresponding candidate vertices, we combine v with all candidate vertices and generate the binary-vertex-bounded circles.

Firstly, we can prune all the circles which exclude the query vertex q . After that, because there is only one vertex difference between two adjacent circles, we can compute the vertex difference between a circle and its precedent. We divide the circles into two groups, the entering circles and leaving circles, respectively. For each entering circle, we record the vertex it brings in, and for each leaving circle, we record the vertex it moves out. For the group of entering circles, we sort them in ascending order of their centers' polar angles and put them into a list L_{enter} . Then for each entering circle O_{enter} in L_{enter} , we compute a vertex set $\mathcal{V}(O_{enter})$ which contains all the vertices bringing from the entering circles that appear before O_{enter} in L_{enter} . This can be done by incrementally adding the vertices bringing from the first entering circle to the last entering circle, and the time complexity is $O(L_{enter})$. It is obvious that the number of vertices in $\mathcal{V}(O_{enter})$ monotonously increases with the index of O_{enter} in L_{enter} . Thus, we can use binary search to find the first entering circle O'_{enter} in L_{enter} such that we can construct a k -core from $\mathcal{V}(O'_{enter})$ containing q .

Algorithm 3: DivBS

Input: $G(V, E)$: the input graph; q : the query vertex;
 k, t, r : constraint parameters; τ : grouping
parameter

Output: \mathcal{D} : a set of diversified RB- k -cores

```

1  $\mathcal{R} \leftarrow \text{RotC}^+(G, q, k, r, \tau)$ 
2 initialize  $\mathcal{D} \leftarrow \emptyset$ 
3 initialize  $V^* \leftarrow \emptyset$ 
4 foreach  $B \in \mathcal{R}$  do
5    $V^* \leftarrow V^* \cup V(B)$ 
6 for  $i = 1 \dots t$  do
7    $cov_{max} \leftarrow 0$ 
8    $B^* \leftarrow \emptyset$ 
9   foreach  $B \in \mathcal{R}$  do
10    if  $|B \cap V^*| > cov_{max}$  then
11       $B^* \leftarrow B$ 
12       $cov_{max} \leftarrow |B \cap V^*|$ 
13     $V^* \leftarrow V^* \setminus V(B^*)$ 
14     $\mathcal{D} \leftarrow \mathcal{D} \cup B^*$ 
15 return  $\mathcal{D}$ 

```

The circles appearing before O'_{enter} can be safely discarded because they cannot contain a RB- k -core. Similarly, we sort in descending order of their centers' polar angles for all the leaving circles and put them into L_{leave} . In the same way, we can find the first leaving circle O'_{leave} such that we can construct a k -core from $\mathcal{V}(O'_{leave})$ containing q and discard all the circles before O'_{leave} in L_{leave} . In this way, we can reduce the number of binary-vertex-bounded circles, which need to be verified in the next stage.

The details of RotC⁺ with the above pruning techniques can be found in [36]. Note that the time complexity of RotC⁺ is $O(\lceil \frac{r}{\tau} \rceil^2 \cdot m \cdot (\log(\lceil \frac{r}{\tau} \rceil) + 1) + |F| \cdot m + |F_1| \cdot \log |F_1| \cdot (|F_1| + m) + |F_1| \cdot |F_2| \cdot m')$, where F denotes the candidate vertex set after pre-process pruning ($|F| \leq n$), F_1 is the vertex set obtained from F after the overall checking, F_2 is the set of circles that need to be verified ($|F_2| \leq n$), and m' is the average time cost of verifying the existence of a k -core ($m' \leq m$) [36].

4 DIVERSIFIED RB- k -CORE SEARCH

In the above section, we study how to find all the RB- k -cores efficiently. However, in many real-world applications, we only need to retrieve representative RB- k -cores for analysis. Motivated by this, we study how to solve the diversified (top- t) RB- k -core search problem in this section.

4.1 A Baseline Algorithm

As presented in Section 3.3, The algorithms RotC, and RotC⁺ can find all the RB- k -cores efficiently. Based on these algorithms, an algorithm DivBS to solve the diversified RB- k -core search problem can be naturally devised. Firstly, DivBS obtains all the RB- k -cores using the RotC⁺ algorithm, and then the diversified RB- k -core search problem becomes the max k -cover problem, which can be solved using a greedy approach with a guaranteed approximation ratio [15].

The details of the DivBS algorithm are shown in Algorithm 3. Firstly, the DivBS algorithm runs the algorithm

RotC⁺ to get the result set \mathcal{R} , which contains all the RB- k -cores. After that, it processes each RB- k -core in \mathcal{R} and puts the vertices in RB- k -core into V^* . Note that V^* contains all the different vertices covered by the RB- k -cores in \mathcal{R} (lines 4-5). Then, DivBS runs t iterations, and for each iteration, we get the RB- k -core B^* which covers the most number of uncovered vertices (lines 6-14). After running lines 6 - 14, \mathcal{D} is an approximate result of the set of diversified top- t RB- k -cores. Note that, the DivBS algorithm achieves the same approximation ratio (i.e., $(1 - 1/e) \approx 0.632$) as the best possible polynomial-time approximation algorithm for the max k -cover problem [15].

Time Complexity of DivBS. Firstly, DivBS needs $O(T_{\text{RotC}^+})$ time to get all the RB- k -cores where T_{RotC^+} is the time complexity of RotC⁺. After that, DivBS needs $O(t \cdot \sum_{B \in \mathcal{R}} |V(B)| \cdot cov(\mathcal{R}))$ time to get the diversified top- t RB- k -cores. This is because DivBS need to run t iterations and in each iteration, we need to compute cov_{max} using $O(\sum_{B \in \mathcal{R}} |V(B)| \cdot cov(\mathcal{R}))$ time. Totally, the time complexity of DivBS is $O(T_{\text{RotC}^+} + t \cdot \sum_{B \in \mathcal{R}} |V(B)| \cdot cov(\mathcal{R}))$.

4.2 Maintenance-Based Solutions

Motivation. Although DivBS achieves a good approximation ratio, it isolates the RB- k -core finding process from the diversified top- t RB- k -core search process. Thus, DivBS cannot have much pruning ability when finding the RB- k -core, and it keeps all the RB- k -cores in memory. Motivated by the above observations, we devise new algorithms that can maintain the diversified top- t candidates with the following advantages. Firstly, the new algorithms only need to maintain t candidate RB- k -cores in \mathcal{D} rather than keeping all the RB- k -cores in the memory and checking for the diversification. Secondly, they have a higher pruning ability than DivBS. As the diversification checking process is integrated into finding RB- k -cores, we can develop pruning techniques to early terminate unpromising processing and reduce the searching space. Thirdly, the new algorithms can achieve a bounded approximation ratio of 0.25, as shown in the algorithm analysis.

4.2.1 The Algorithm DivRotC⁺

Before introducing the new algorithm, we first define the private-vertex-coverage of a RB- k -core.

Definition 6 (Private-Vertex-Coverage). Given a set of RB- k -cores $\mathcal{D} = \{B_1, B_2, \dots\}$ in G , the private-vertex-coverage of each $B \in \mathcal{D}$ denoted by $pvcov(B)$ is the set of vertices in B that are not covered by the other RB- k -cores in \mathcal{D} , i.e., $pvcov(B, \mathcal{D}) = V(B) \setminus cov(\mathcal{D} \setminus B)$.

Based on the definition of private-vertex-coverage, we can have the following definition of Min-cover RB- k -core.

Definition 7 (Min-Cover RB- k -core). Given a set of RB- k -cores $\mathcal{D} = \{B_1, B_2, \dots\}$ in G , the min-cover RB- k -core of \mathcal{D} , denoted by $B_{min}(\mathcal{D})$, is the RB- k -core $B \in \mathcal{D}$ with the smallest $pvcov(B, \mathcal{D})$, i.e., $B_{min}(\mathcal{D}) = \arg \min_{B \in \mathcal{D}} |pvcov(B, \mathcal{D})|$.

With the definition of Min-cover RB- k -core, we propose the DivRotC⁺ algorithm. DivRotC⁺ is based on RotC⁺, and a new result updating function is used on RotC⁺ to ensure only t RB- k -cores are maintained in \mathcal{D} . We show the details

of DivRotC⁺ in Algorithm 4. We first initialize the result set \mathcal{D} and run the grouping-based pre-process (lines 1-2). In the grouping-based pre-processing, the updating of \mathcal{R} is replaced with the updating of \mathcal{D} . In addition, since candidate RB- k -cores can be retrieved in this process, we replace the update process in line 10 with a new updating function **DivUpdate**. After that, for each vertex v in $V(G_k)$, we set it as the pole in a polar coordinate system P (line 3). For each pole v , we generate candidate binary-vertex-bounded circles and corresponding candidate graphs following the same framework as RotC⁺ (lines 4-21). For each candidate subgraph, if there exists a RB- k -core G_k^r in it, we run the function **DivUpdate**(G_k^r, \mathcal{D}) to check whether \mathcal{D} can be updated (lines 22-24). In the function **DivUpdate**, we directly insert G_k^r into \mathcal{D} if $|\mathcal{D}| < t$. Otherwise, we generate \mathcal{D}' by replacing $B_{min}(\mathcal{D})$ with G_k^r . If G_k^r meets the condition that $|pvcov(G_k^r, \mathcal{D}')| > \delta$, we replace \mathcal{D} with \mathcal{D}' . Here δ is the updating threshold which is equal to $|pvcov(B_{min}(\mathcal{D}), \mathcal{D})| + \frac{cov(\mathcal{D})}{t}$. Finally, we get a set of RB- k -cores in \mathcal{D} .

Threshold maintenance for early termination. Deriving from the candidate updating condition, we can have the following lemma.

Lemma 5. Given a geo-social graph $G(V, E)$, a subgraph G' of G , and a set of t candidate RB- k -cores in \mathcal{D} , there is no RB- k -core in G' can be included in \mathcal{D} by Algorithm DivRotC⁺ if $|V(G')| \leq \delta$.

Proof. According to the candidate update process in Algorithm 4 lines 27-33, a RB- k -core G_k^r can be included in \mathcal{D} if the private-vertex-coverage after replacing $B_{min}(\mathcal{D})$ with G_k^r is increased by more than $\frac{cov(\mathcal{D})}{t}$. Since adding G_k^r increases the private-vertex-coverage by at most $|G_k^r|$, we can get that if $G_k^r \in V(G')$, the private-vertex-coverage cannot be increased by more than $|V(G')|$. Thus, this lemma holds. ■

According to lemma 5, we can maintain the threshold δ each time we update the candidate RB- k -cores in \mathcal{D} . After that, a candidate subgraph in the finding process can be pruned directly if the size of it is less than δ . In Algorithm 4, the early termination checking according to the threshold can be applied in the following stages: 1) checking the size of $G(X)$ in line 8 in the grouping-based pre-process (i.e., Algorithm 2); 2) checking the size of $G(X)$ in line 12; 3) checking the number of vertices in X in line 19. Note that the number of vertices in X can be dynamically maintained since there is at most one vertex is included/excluded between adjacent circles.

Analysis of DivRotC⁺. We first show the theoretical guarantee of the result of Algorithm 4, and then the time complexity of Algorithm 4.

Theoretical guarantee of the result. Suppose \mathcal{D} is the result provided by Algorithm 4 and \mathcal{D}^* is the set of optimal diversified top- t RB- k -cores. We have $|cov(\mathcal{D})| \geq 0.25 \times |cov(\mathcal{D}^*)|$. This can be easily extended by the theoretical result shown in [4] which analyses an online approximate algorithm for maximum k -coverage problem.

Time complexity. The time complexity of DivRotC⁺ is $O(T_{RotC^+} + \sum_{B \in \mathcal{R}} (|V(B)| + |V(B_{max})|))$. Here, T_{RotC^+} is the time complexity of RotC⁺, \mathcal{R} is the set of all the RB- k -cores, and B_{max} denotes the RB- k -core with the largest

Algorithm 4: DivRotC⁺

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint parameters; τ : grouping parameter

Output: \mathcal{D} : a set of RB- k -cores

```

1 initialize  $\mathcal{D} \leftarrow \emptyset$ 
2  $G_k \leftarrow PreProcess(G, q, k, r, \tau, \mathcal{D})$ , replace  $\mathcal{R}$  with  $\mathcal{D}$ ,
  update  $\mathcal{D}$  using function DivUpdate in line 10, early
  terminate the processing of  $X$  if  $|X| \leq \delta$ 
3 foreach node  $v \in V(G_k)$  do
4    $C \leftarrow \emptyset$ ;  $X \leftarrow \emptyset$ 
5   foreach node  $u \in V(G_k)$  do
6     if  $u \neq v \wedge d(u, v) \leq 2r$  then
7       put  $u$  into  $X$ 
8       compute  $W_r(u, v)$  using  $\{u, v\}$  and  $r$ 
9       put circles in  $W_r(u, v)$  into  $C$ 
10  if OverallChecking( $X$ ) = false then
11    continue
12  if  $|\mathcal{D}| > t \wedge |X| \leq \delta$  then
13    continue
14  sort  $C$  in ascending order of centers' polar angles
15  employ circle filtering to  $C$ 
16  foreach  $O(c, r) \in C$  do
17     $X \leftarrow$  a set of vertices enclosed in  $O(c, r)$ 
18    maintain the degree of vertices in  $X$ 
19    if  $|\mathcal{D}| > t \wedge |X| \leq \delta$  then
20      continue
21    if  $O(c, r)$  is an entering circle then
22      construct  $G(X)$  from  $X$ 
23      if exists a  $G_k^r$  in  $G(X)$  then
24        | DivUpdate( $G_k^r, \mathcal{D}$ )
25 return  $\mathcal{D}$ 
26
27 DivUpdate( $G_k^r, \mathcal{D}$ )
28 if  $|\mathcal{D}| < t$  then
29   |  $\mathcal{D} \leftarrow \mathcal{D} \cup G_k^r$ 
30 else
31   |  $\mathcal{D}' \leftarrow (\mathcal{D} \setminus B_{min}(\mathcal{D})) \cup G_k^r$ 
32   | if  $|pvcov(G_k^r, \mathcal{D}')| > |pvcov(B_{min}(\mathcal{D}), \mathcal{D})| + \frac{cov(\mathcal{D})}{t}$ 
33   | then
34   | |  $\mathcal{D} \leftarrow \mathcal{D}'$ 

```

number of vertices in \mathcal{R} . Note that, the second term is the time complexity of candidate updating. For each $B \in \mathcal{R}$, we need to compute the private-vertex-coverage after replacing $B_{min}(\mathcal{D})$ with B which needs $O(|V(B)| + |V(B_{max})|)$ time.

4.2.2 The Algorithm AdvDivRotC⁺

Following the framework of DivRotC⁺, here we propose optimization techniques to improve the efficiency of DivRotC⁺. Specifically, we first propose a tighter upper bound for the private-vertex-coverage of candidate subgraphs. Secondly, we explore the effect of the processing order of poles.

A tighter upper bound. In Lemma 5, we use $|V(G')|$ as the upper bound of the private-vertex-coverage increasing. Although $|V(G')|$ is a correct upper bound, it cannot reflect the number of private vertices in the possible RB- k -cores in the candidate subgraph G' and may not be very tight. In order to obtain a tighter upper bound, we maintain a $O(n)$ hash table J , which contains all the vertices in $\mathcal{D} \setminus B_{min}(\mathcal{D})$. When we need to check a candidate subgraph G' , we obtain the upper bound of the number of uncovered vertices provided by G' by computing $|V(G') \setminus V(J)|$. Then,

we can early terminate the checking if $|V(G') \setminus V(J)| \leq \delta$. Note that the value $|V(G') \setminus V(J)|$ can be easily computed in $O(V(G'))$ time, and it can also be dynamically maintained when processing the circles in line 15 in Algorithm 4.

The processing order of poles. Apart from the above bound, we can also consider the processing order of poles. This is because a result set with large private-vertex-coverage can increase the chance to prune the unpromising candidate subgraphs. We rank the candidate poles according to the following equation and choose the pole with the highest rank in each iteration.

$$\text{rank}(v) = \alpha(v) + \frac{|V(G^*) \setminus V(J)|}{|V(G_k)|} \quad (2)$$

Here, we use $\alpha(v)$ to ensure that we first process the pole which is not in the candidate result set. We set $\alpha(v) = 0$ if the vertex v is contained in a RB- k -core in the current result set \mathcal{D} . Otherwise, we set $\alpha(v) = 1$. Secondly, we prefer the pole which may expand more vertices which are not covered by the RB- k -cores in the current result set. We use G^* to estimate this term. For each $u \in V(G^*)$, it should satisfy the following constraints: 1) u is a neighbor of v ; 2) $d(u, v) \leq r$; Thus, we compute the second term of $\text{rank}(v)$ by $\frac{|V(G^*) \setminus V(J)|}{|V(G_k)|}$. Here, G_k is the candidate subgraph after running the pre-processing and J is the set of vertices in $\mathcal{D} \setminus B_{\min}(\mathcal{D})$. Note that, for each vertex v , $\alpha(v)$ can be obtained in $O(\text{cov}(\mathcal{D}))$ time. G^* can be obtained in $O(\text{deg}_G(v))$ time and we can compute $\frac{|V(G^*) \setminus V(J)|}{|V(G_k)|}$ in $O(G^*)$ time.

The AdvDivRotC⁺ algorithm. Utilizing the above strategies, we propose the AdvDivRotC⁺ algorithm as shown in Algorithm 5. We first initialize the result set \mathcal{D} and run the grouping-based pre-processing (lines 1-2). In the grouping-based pre-processing, comparing with DivRotC⁺, we use advanced early termination condition after line 8 in Algorithm 2. After that, for each vertex v in $V(G_k)$ with the highest $\text{rank}(v)$, we set it as the pole in a polar coordinate system P (line 3). The rank value is computed according to Equation 2. For each pole v , we first get all the candidate vertices and check whether we can skip the process of v . Then, we generate candidate binary-vertex-bounded circles and corresponding candidate subgraphs following the same framework as DivRotC⁺. For each candidate subgraph, if it satisfies the upper bound checking and there exists a RB- k -core G_k^r in it, we run the function **AdvDivUpdate**(G_k^r, \mathcal{D}) to check whether \mathcal{D} can be updated. In the function **AdvDivUpdate**, we directly insert G_k^r into \mathcal{D} if $|\mathcal{D}| < t$. Otherwise, we generate \mathcal{D}' by replacing $B_{\min}(\mathcal{D})$ with G_k^r . If G_k^r meets the condition that $|pvcov(G_k^r, \mathcal{D}')| > |pvcov(B_{\min}(\mathcal{D}), \mathcal{D})| + \frac{\text{cov}(\mathcal{D})}{t}$, we replace \mathcal{D} with \mathcal{D}' . Finally, we get a set of RB- k -cores in \mathcal{D} .

5 EXPERIMENTS

In this section, we report the evaluation of the effectiveness of our model and the efficiency of our algorithms.

5.1 Experimental Settings

Algorithms. In the experimental study, we implement and evaluate four algorithms to solve the RB- k -core search problem: the triple-vertex-based algorithm TriV in Section 3.1,

Algorithm 5: AdvDivRotC⁺

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint parameters; τ : grouping parameter

Output: \mathcal{D} : a set of RB- k -cores

```

1 initialize  $\mathcal{D} \leftarrow \emptyset$ 
2  $G_k \leftarrow \text{PreProcess}(G, q, k, r, \tau, \mathcal{D})$ , replace  $\mathcal{R}$  with  $\mathcal{D}$ ,
  early terminate the processing of  $X$  if
   $|\mathcal{D}| > t \wedge X \setminus V(J) \leq \delta$ , update  $\mathcal{D}$  using function
  AdvDivUpdate in line 10
3 foreach node  $v \in V(G_k)$  with the highest rank do
4    $C \leftarrow \emptyset$ ;  $S \leftarrow \emptyset$ ;  $J \leftarrow \emptyset$ ;  $\delta = 0$ 
5   run Algorithm 5 lines 5-24, replace the condition
     ( $|\mathcal{D}| > t \wedge X \leq \delta$ ) with ( $|\mathcal{D}| > t \wedge X \setminus V(J) \leq \delta$ ) in
     lines 12 and 19;
6 return  $\mathcal{D}$ 
7
8 AdvDivUpdate( $G_k^r, \mathcal{D}$ )
9 if  $|\mathcal{D}| < t$  then
10    $\mathcal{D} \leftarrow \mathcal{D} \cup G_k^r$ 
11 else
12    $\mathcal{D}' \leftarrow (\mathcal{D} \setminus B_{\min}(\mathcal{D})) \cup G_k^r$ 
13   if  $|pvcov(G_k^r, \mathcal{D}')| > |pvcov(B_{\min}(\mathcal{D}), \mathcal{D})| + \frac{\text{cov}(\mathcal{D})}{t}$ 
     then
14      $\mathcal{D} \leftarrow \mathcal{D}'$ 
15      $\delta \leftarrow |pvcov(B_{\min}(\mathcal{D}), \mathcal{D})| + \frac{\text{cov}(\mathcal{D})}{t}$ 
16      $J \leftarrow V(\mathcal{D} \setminus B_{\min}(\mathcal{D}))$ 

```

the binary-vertex-based algorithm BinV in Section 3.2, the rotating-circle-based algorithm RotC in Section 3.3, and the optimized rotating-circle-based algorithm RotC⁺ in Section 3.3. We also extend our RotC⁺ algorithm to solve the SAC (spatial-aware community) search problem proposed in [14]. We also evaluate the algorithms DivBS, DivRotC⁺ and AdvDivRotC⁺ to solve the diversified top- t RB- k -core search problem in Section 4.

The algorithms are implemented in C++, and the experiments are run on a Linux server with Intel Xeon E5-2687W (3.4GHz, 8 Cores) processor and 64GB main memory. We randomly select 200 query vertices and report the average result for these queries. We terminate an algorithm if the running time is more than three hours.

Table 2: Summary of Datasets

Dataset	$ V $	$ E $	d_{avg}
Brightkite	51,406	197,167	7.67
Gowalla	107,092	456,830	8.53
Flickr	214,698	2,096,306	19.5
Foursquare	2,127,093	8,640,352	8.12
Synthetic	4,000,000	40,000,000	20

Datasets. We use four real datasets in our experiments including Brightkite, Gowalla, Flickr, and Foursquare. In the four datasets, we consider each user associated with a geo-location coordinate (latitude and longitude) as a vertex and the friendship between two users as an edge. Helmert transformation [37] is adopted to transform geo-location coordinates of vertices to Cartesian coordinates.


We also conduct experiments on a synthetic dataset Synthetic. We first generate a non-spatial graph using a

well-known graph generator GTGraph¹. The degrees of the vertices in the graph follow a power-law distribution, as often used in the study of social networks. After generating the graph, we randomly generate the locations of the vertices in a square of side 300km.

Parameters. The experiments are conducted using different settings on 4 parameters: k (the minimum degree), r (the maximal radius), τ (the parameter used in the pre-process of RotC⁺), and n (the percentage of vertices). We vary k from 4 to 16 and set 4 as the default value. We vary r from 1km to 40km and set r to 5km by default. When varying the graph size, we randomly sample 20% to 100% vertices of the original graphs and construct the induced subgraphs using these vertices. The parameter n is varied from 20% to 100%, representing the percentage of the vertices we use in each dataset. The parameter τ is varied from r to $\frac{r}{16}$ which controls the number of iterations of the pre-processing in RotC⁺. We vary t , which is a parameter used to control the number of diversified top- t RB- k -cores from 1 to 9 and set 5 as the default value.

5.2 Effectiveness Evaluation

In this section, we show the effectiveness of our RB- k -core model and the diversified RB- k -core model.



(a) Case study on Gowalla ($q=1396, k=3, r=0.76\text{km}$) (b) Case study on Flickr ($q=111419, k=3, r=1.67\text{km}$)

Figure 5: Case studies

Case study. We present two case studies to show the result of RB- k -core search on Gowalla and Flickr in Figure 5(a) and Figure 5(b), respectively. The query vertices are marked by question mark symbols. Under setting $q=1396, k=4$ and $r=0.76\text{km}$ on Gowalla, we can get two RB- k -cores containing q as shown in Figure 5(a). We mark the vertices and the MCC of these two RB- k -cores in black color and grey color, respectively. We can see that, the social constraint and the spatial constraint both contribute to the construction of these two RB- k -cores. For example, if the social constraint is ignored, the black vertices enclosed by the grey circle will be included in the grey RB- k -core. On the other hand, all the vertices in Figure 5(a) will be united into one community if the radius constraint is not being considered. Figure 5(b) shows the result of RB- k -core search on Flickr using $q=111419, k=3$ and $r=1.67\text{km}$ contains two retrieved communities. Using the same q and k , the SAC search (i.e., a similar model in [14]) will provide the communities with black color as shown in Figure 5(a) and Figure 5(b). The radiuses of the black circles are 0.74km and 1.67km

1. <http://www.cse.psu.edu/~kxm85/software/GTgraph/>

in Figure 5(a) and Figure 5(b), respectively. Compared to the SAC search, in Figure 5(a), our RB- k -core search can give users more options by slightly increasing the minimum radius (i.e., from 0.74km to 0.76km) for the same q and k . In Figure 5(b), the RB- k -core search is able to provide users more than one selection for the same q, k , and minimum r . **Effectiveness of the diversified RB- k -core model.** Here, we show the effectiveness of the diversified RB- k -core search by comparing the total coverage of the result sets provided by RotC⁺, DivBS, DivRotC⁺, and AdvDivRotC⁺. We vary the parameter t , which controls the number of different RB- k -cores returned by the algorithms.




Figure 6: Evaluate the total coverage

We can see that in Figure 6, the total coverage increases when t increases for all the algorithms DivBS, DivRotC⁺, and AdvDivRotC⁺. This is obvious since the number of RB- k -cores in the result sets increases when t increases. We can also observe that the coverages of the algorithms DivBS, DivRotC⁺, and AdvDivRotC⁺ are only slightly smaller than the RotC⁺ algorithm, especially when $t > 1$. This validates the effectiveness of our diversified RB- k -core model. In addition, the total coverage of AdvDivRotC⁺ is smaller than DivBS. This is because AdvDivRotC⁺ maintains only top- t RB- k -cores in the computation and is more efficient.

5.3 Efficiency Evaluation of the algorithms to solve the RB- k -core search problem

In this section, we first evaluate the efficiency of the proposed four algorithms to solve the RB- k -core search problem on all the datasets. Then we evaluate the effect of k and the scalability of the proposed algorithm. After that, the parameter τ used in the RotC⁺ algorithm is evaluated. Finally, we extend our RotC⁺ algorithm to solve the SAC search problem [14] and compare the performance.





Figure 7: Performance on different datasets

Evaluating the performance of all algorithms on different datasets. In Figure 7, we show the performance of our RB- k -core search algorithms on five datasets. We set k as default and r to 1km, 5km, 10km, 20km, 40km on Brightkite,


Gowalla, Flickr, Foursquare and Synthetic, respectively. In the meantime, we keep the other parameters fixed to default values. We can observe that BinV is more efficient than TriV on Brightkite, Gowalla, and Flickr. The algorithms RotC and RotC⁺ using the rotating circle strategy are more efficient than TriV and BinV on the three datasets because RotC and RotC⁺ can compute the RB- k -cores in an incremental manner, which significantly reduces the computation cost. On Foursquare and Synthetic, we can see that only RotC⁺ can return the results within the time limit. Foursquare is much larger than the first three datasets, which means many more candidate vertices to be processed. On Synthetic, the vertices are more densely distributed over the space than the other datasets, and thus the candidate circles contain more vertices. In summary, as shown in Figure 7, our RotC⁺ algorithm significantly outperforms the other three algorithms on all datasets.



(a) Gowalla, varying k (b) Foursquare, varying k

Figure 8: Effect of k

Evaluating the effect of k . Figure 8 evaluates the effect of k for four algorithms on Gowalla and Foursquare. We vary k from 4 to 16 and fix the other parameters as default values. In Figure 8(a), we can observe that the time cost of all four algorithms drops when k increases because of the number of vertices in the k -core of the original graph (selected as candidate vertices) decreases. Similar trends can be observed in Figure 8(b). As expected, RotC and RotC⁺ significantly outperform TriV and BinV on both datasets due to the usage of the rotating circle technique. For example, on both datasets, RotC is about one order of magnitude faster than TriV and BinV, and RotC⁺ is at least two orders of magnitude faster than TriV and BinV.




(a) Gowalla, varying n (b) Foursquare, varying n

Figure 9: Effect of graph size

Scalability. (1) *Evaluating the effect of graph size.* Figure 9 shows the scalability of four algorithms by varying the graph size from 20% to 100% in all datasets. We can observe that, on Gowalla, all these four algorithms are scalable and their running time increases as the percentage of vertices increases. As shown in Figure 9(b), on Foursquare, TriV


and BinV can only finish within the time limit when $n=20\%$, while RotC and RotC⁺ have similar trends as in Figure 9(a). As discussed before, RotC⁺ is more efficient than the other three algorithms.



(a) Gowalla, varying r (b) Foursquare, varying r

Figure 10: Effect of r

(2) *Evaluating the effect of r .* Figure 10 illustrates the effect of r on Gowalla and Foursquare. We vary r from 1km to 40km and fix the other parameters as default values. In Figures 10(a) and 10(b), the time cost increases as r becomes larger because the number of vertices in circle $O(q, 2r)$ grows when r increases. We can also see that, on Gowalla, both RotC and RotC⁺ are several orders of magnitude faster than TriV and BinV. On Foursquare, TriV and BinV can only compute the result when $r = 1$ km and RotC can get the results when r is no more than 10km within reasonable time. As expected, the RotC⁺ algorithm significantly outperforms the other three algorithms on Foursquare, and the time cost is stable when r is large on both datasets.



(a) Gowalla, varying τ (b) Foursquare, varying τ

Figure 11: Effect of τ

Evaluating the effect of τ . Figure 11 illustrates the effect of τ , which is a parameter used in the grouping-based pre-processing in RotC⁺. Because the value of τ is related to r , we set r to 1km, 5km, 10km, 20km, and 40km on both Gowalla and Foursquare. As discussed before, as τ increases, the time cost of pre-processing increases, and the number of candidate vertices decreases. We can observe that the running time is not very sensitive to τ when τ is relatively large on the two datasets. The time cost starts to increase from $\tau = \frac{r}{4}$ in most cases because the number of vertices that can be pruned increases slowly, and the time cost of pre-processing begins to dominate the cost of RotC⁺. Hence we set $\tau = \frac{r}{4}$ in our experiments on all datasets.

Extend to solve the SAC search problem [14]. As discussed before, the SAC search problem can be solved by slightly modifying our RotC⁺ algorithm using the binary search. In Figure 12, we study the performance of the SAC-RotC⁺ algorithm, which is extended from RotC⁺ to solve the SAC search problem, and we compare its performance with the state-of-the-art exact algorithm SAC-Exact⁺ proposed in




Figure 12: Extend to solve SAC search problem

[14]. Fang et al. [14] implemented the SAC-Exact⁺ algorithm in JAVA, while we implement the SAC-Exact⁺ algorithm in C++ for the fairness of comparison.

The SAC-Exact⁺ algorithm includes two phases. Firstly, it conducts the quad-tree-based vertex pruning phase, which can reduce the number of potential vertices. Next, it conducts a triple-vertex-based algorithm, which is similar to the TriV algorithm in this paper. In the RB- k -core search problem, we have analyzed that the triple-vertex-based algorithm is time-consuming, and it can be improved by the rotating circle strategy to compute the result incrementally. We can do the same thing in the SAC search problem. In our SAC-RotC⁺ algorithm, we also conduct the vertex pruning phase, but we adopt the rotating-circle-based algorithm in the second phase. Note that the in-process pruning technique in RotC⁺ can also be applied in SAC-RotC⁺, but the pre-process pruning technique cannot be used because of the model difference.


We vary the parameter ϵ , which controls the number of iterations in the vertex pruning phase, and the number of iterations decreases with an increase of ϵ . From Figure 12(a) and Figure 12(b), we can observe that the time cost of SAC-RotC⁺ and SAC-Exact⁺ is almost the same when ϵ is very small because the cost of processing the vertex pruning phase dominates the cost in the second phase. On Foursquare, SAC-RotC⁺ outperforms SAC-Exact⁺ when ϵ is larger than 10^{-3} . Also, on Gowalla, SAC-RotC⁺ is about one order of magnitude faster than SAC-Exact⁺ when ϵ is larger than 10^{-4} . This is because our SAC-RotC⁺ algorithm obtains the result incrementally and significantly outperforms the triple-vertex-based algorithm in the second phase, which incurs the dominating time cost as ϵ gets large. Comparing the two algorithms' minimal time cost on both datasets, we can conclude that SAC-RotC⁺ can achieve a speed-up around twice.

5.4 Efficiency Evaluation of the algorithms to solve the diversified RB- k -core search problem


In this section, we evaluate the efficiency of the proposed algorithms (i.e., DivBS, DivRotC⁺, and AdvDivRotC⁺) to solve the diversified top- t RB- k -core search problem under different settings of parameters. Firstly, we evaluate the effect of t . Then, we evaluate the effect of k . We also evaluate the scalability (i.e., the effect of graph size and the effect of r) of these algorithms.

Evaluating the effect of t in diversified RB- k -core search.

In Figure 13, we evaluate the effect of t for the three diversified algorithms on Gowalla and Foursquare. We vary t from 1 to 9 and fix the other parameters as default


Figure 13: Diversified RB- k -core search - effect of t

values. We can observe that when $t = 1$, the algorithms DivRotC⁺ and AdvDivRotC⁺ are much faster than DivBS on both datasets. This is because these two algorithms only need to maintain one RB- k -core in the result set, while DivBS needs to maintain all the RB- k -cores. When $t \geq 3$, the time cost of all the algorithms slightly increases when t increases. As expected, AdvDivRotC⁺ outperforms DivBS and DivRotC⁺ on both datasets because it only maintains top- t RB- k -cores with advanced pruning strategies.


Figure 14: Diversified RB- k -core search - effect of k

Evaluating the effect of k in diversified RB- k -core search.

In Figure 14, we evaluate the effect of k for the algorithms DivBS, DivRotC⁺ and AdvDivRotC⁺ on Gowalla and Foursquare. We vary k from 4 to 16 and fix the other parameters as default values. We can observe that the time cost of all these algorithms drops when k increases. This is because the number of candidate vertices is reduced with an increase of k . Also, as an advanced solution, AdvDivRotC⁺ outperforms DivRotC⁺ and DivBS on both datasets.

Figure 15: Diversified RB- k -core search - effect of graph size

Scalability. (1) *Evaluating the effect of graph size in diversified RB- k -core search.* Figure 15 shows the scalability of four algorithms by varying n from 20% to 100% in all datasets. We can observe that, on both datasets, our advanced algorithm AdvDivRotC⁺ is scalable and its computation cost increases when the percentage of vertices increases. As discussed before, AdvDivRotC⁺ is more efficient than the other two algorithms.

Figure 16: Diversified RB- k -core search - effect of r

(2) *Evaluating the effect of r in diversified RB- k -core search.* Figure 16 shows the effect of r on Gowalla and Foursquare. We vary r from 1km to 40km and fix the other parameters as default values. In Figures 16(a) and 16(b), the time cost of these three algorithm increase as r increases and AdvDivRotC⁺ is much faster than DivBS and DivRotC⁺ when r is large. For instance, in Figure 16(b), AdvDivRotC⁺ is more than three times faster than DivBS and DivRotC⁺ when $r = 40$ km.

6 RELATED WORK

Community retrieval has been widely studied and used in many applications such as location-aware marketing [27], influence analysis [22], and event recommendation [24].

Community retrieval considering social connections. Prior works study various models such as k -core [21], [26], [32], k -truss [8], [19], [31], [35], [43], and clique [25] to retrieve communities based on users' social connections. Based on k -core, [5], [9], [33] study algorithms for k -core community search. Based on k -truss, Huang et al. [20] study the closest model and the triangle-connected model for community search are studied in [2], [18]. In [40], Yuan et al. propose algorithms to solve the densest clique percolation community search problem. However, the geo-locations of users are not considered in the above works.

Community retrieval considering spatial locations. In spatial databases, several works study the group objects retrieval problem based on users' spatial locations such as [16], [30], [38] and [11]. Guo et al. [16] study the spatial keyword query which retrieves a group of objects close to each other and cover a set of keywords together. Wu et al. [38] adapt the densest subgraph model to the spatial community search problem on dual networks. The work [30] proposes localitySearch which retrieves top- k sets of spatial web objects by integrating spatial distance, textual relevance, and a "co-locality" measure into one ranking function. The work [11] focuses on context-aware search over social media data. It analyses the data-centric challenges in temporal, spatial, and spatio-temporal contexts. These proposals do not consider the social connections of users, and thus they are different from our problem.

Community retrieval considering both social connections and spatial locations. On geo-social networks, recently, some works study the community retrieval problem [7], [14], [38], [42], [44] considering both the spatial and social features. The works [7], [42] mainly focus on analyzing and understanding the complexity networks rather than online community search. The most closely related work of radius-bounded k -core computation is that Zhu et al. study finding a community within a given rectangle in [44]. Their study

is different from our work because what we consider is restricting the size of community spatially instead of within a given rectangle. Fang et al. [14] propose both exact and approximate algorithms to find a community covered by the smallest circle for a given query vertex. In their work, the radius of a circle is not given by users and only one community covered by the smallest circle is returned to users, and thus it cannot provide more options for users as done by our work.

Diversified top- t search. In the literature, many works [1], [3], [6], [10], [13], [23], [28], [29], [34], [39], [41] study to find diversified top- t answers according to a specific problem. In these works, Yuan et al. [39] aim to find diversified top- k cliques. [1], [3] studied diversified top- t document retrieval. Lin et al. [23] focus on the t most representative skyline problem. The diversified top- t graph pattern matching problem is studied in [13]. The diversified (k, r) -core search problem is studied in [41]. However, since the problems and models are different, none of them can be directly used to solve the diversified top- t RB- k -core search problem.

7 CONCLUSION

In this paper, we study the RB- k -core search problem. We propose a triple-vertex-based algorithm and a binary-vertex-based algorithm as benchmark algorithms to find all the RB- k -cores. We propose a rotating-circle-based algorithm which can find possible cost sharing opportunities. The rotating-circle-based algorithm is further enhanced by critical pruning techniques. In addition, we study the diversified RB- k -core search problem which aims to find representative RB- k -cores with rich information. We conduct extensive experiments on both real and synthetic datasets and the experimental result shows that our rotating-circle-based algorithm significantly outperforms the benchmark algorithms.

REFERENCES

- [1] R. Agrawal, S. Gollapudi, A. Halverson, and S. Ieong. Diversifying search results. In *Proceedings of the second ACM international conference on web search and data mining*, pages 5–14. ACM, 2009.
- [2] E. Akbas and P. Zhao. Truss-based community search: a truss-equivalence based indexing approach. *PVLDB*, 10(11):1298–1309, 2017.
- [3] A. Angel and N. Koudas. Efficient diversity-aware search. In *SIGMOD*, pages 781–792. ACM, 2011.
- [4] G. Ausiello, N. Boria, A. Giannakos, G. Lucarelli, and V. T. Paschos. Online maximum k -coverage. *Discrete Applied Mathematics*, 160(13–14):1901–1913, 2012.
- [5] N. Barbieri, F. Bonchi, E. Galimberti, and F. Gullo. Efficient and effective community search. *Data mining and knowledge discovery*, 29(5):1406–1433, 2015.
- [6] A. Borodin, H. C. Lee, and Y. Ye. Max-sum diversification, monotone submodular functions and dynamic updates. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 155–166. ACM, 2012.
- [7] Y. Chen, J. Xu, and M. Xu. Finding community structure in spatially constrained complex networks. *IJGIS*, 29(6):889–911, 2015.
- [8] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. *NSATR*, 16, 2008.
- [9] W. Cui, Y. Xiao, H. Wang, and W. Wang. Local search of communities in large graphs. In *SIGMOD*, pages 991–1002, 2014.
- [10] T. Deng and W. Fan. On the complexity of query result diversification. *PVLDB*, 6(8):577–588, 2013.


- [11] L. R. Derczynski, B. Yang, and C. S. Jensen. Towards context-aware search and analysis on social media data. In *Proceedings of the 16th international conference on extending database technology*, pages 137–142. ACM, 2013.
- [12] J. Elzinga and D. W. Hearn. Geometrical solutions for some minimax location problems. *Transportation Science*, 6(4):379–394, 1972.
- [13] W. Fan, X. Wang, and Y. Wu. Diversified top-k graph pattern matching. *PVLDB*, 6(13):1510–1521, 2013.
- [14] Y. Fang, R. Cheng, X. Li, S. Luo, and J. Hu. Effective community search over large spatial graphs. *PVLDB*, 10(6):709–720, 2017.
- [15] U. Feige. A threshold of $\ln n$ for approximating set cover. *JACM*, 45(4):634–652, 1998.
- [16] T. Guo, X. Cao, and G. Cong. Efficient algorithms for answering the m-closest keywords query. In *SIGMOD*, pages 405–418. ACM, 2015.
- [17] R. Hartshorne. *Geometry: Euclid and beyond*. Springer Science & Business Media, 2013.
- [18] X. Huang, H. Cheng, L. Qin, W. Tian, and J. X. Yu. Querying k-truss community in large and dynamic graphs. In *SIGMOD*, pages 1311–1322, 2014.
- [19] X. Huang and L. V. Lakshmanan. Attribute-driven community search. *PVLDB*, 10(9):949–960, 2017.
- [20] X. Huang, L. V. Lakshmanan, J. X. Yu, and H. Cheng. Approximate closest community search in networks. *PVLDB*, 9(4):276–287, 2015.
- [21] W. Khaouid, M. Barsky, V. Srinivasan, and A. Thomo. K-core decomposition of large networks on a single pc. *PVLDB*, 9(1):13–23, 2015.
- [22] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse. Identification of influential spreaders in complex networks. *arXiv*, 2010.
- [23] X. Lin, Y. Yuan, Q. Zhang, and Y. Zhang. Selecting stars: The k most representative skyline operator. In *ICDE*, pages 86–95. IEEE, 2007.
- [24] X. Liu, Q. He, Y. Tian, W.-C. Lee, J. McPherson, and J. Han. Event-based social networks: linking the online and offline social worlds. In *SIGKDD*, pages 1032–1040. ACM, 2012.
- [25] R. D. Luce and A. D. Perry. A method of matrix analysis of group structure. *Psychometrika*, 14(2):95–116, 1949.
- [26] D. W. Matula and L. L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *JACM*, 30(3):417–427, 1983.
- [27] D. McKenzie-Mohr. *Fostering sustainable behavior: An introduction to community-based social marketing*. New society publishers, 2011.
- [28] E. Minack, W. Siberski, and W. Nejdl. Incremental diversification for very large sets: a streaming-based approach. In *SIGIR*, pages 585–594. ACM, 2011.
- [29] L. Qin, J. X. Yu, and L. Chang. Diversifying top-k results. *PVLDB*, 5(11):1124–1135, 2012.
- [30] Q. Qu, S. Liu, B. Yang, and C. S. Jensen. Efficient top-k spatial locality search for co-located spatial web objects. In *MDM*, volume 1, pages 269–278. IEEE, 2014.
- [31] K. Saito, T. Yamada, and K. Kazama. Extracting communities from complex networks by the k-dense method. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, 91(11):3304–3311, 2008.
- [32] S. B. Seidman. Network structure and minimum degree. *Social networks*, 5(3):269–287, 1983.
- [33] M. Sozio and A. Gionis. The community-search problem and how to plan a successful cocktail party. In *SIGKDD*, pages 939–948, 2010.
- [34] M. R. Vieira, H. L. Razente, M. C. Barioni, M. Hadjieleftheriou, D. Srivastava, C. Traina, and V. J. Tsotras. On query result diversification. In *ICDE*, pages 1163–1174. IEEE, 2011.
- [35] J. Wang and J. Cheng. Truss decomposition in massive networks. *PVLDB*, 5(9):812–823, 2012.
- [36] K. Wang, X. Cao, X. Lin, W. Zhang, and L. Qin. Efficient computing of radius-bounded k-cores. In *2018 IEEE 34th International Conference on Data Engineering (ICDE)*, pages 233–244. IEEE, 2018.
- [37] G. Watson. Computing helmert transformations. *JCAM*, 197(2):387–394, 2006.
- [38] Y. Wu, R. Jin, X. Zhu, and X. Zhang. Finding dense and connected subgraphs in dual networks. In *ICDE*, pages 915–926. IEEE, 2015.
- [39] L. Yuan, L. Qin, X. Lin, L. Chang, and W. Zhang. Diversified top-k clique search. *VLDBJ*, 25(2):171–196, 2016.
- [40] L. Yuan, L. Qin, W. Zhang, L. Chang, and J. Yang. Index-based densest clique percolation community search in networks. *TKDE*, 30(5):922–935, 2017.
- [41] F. Zhang, X. Lin, Y. Zhang, L. Qin, and W. Zhang. Efficient community discovery with user engagement and similarity. *VLDBJ*, pages 1–26, 2019.
- [42] F. Zhang, Y. Zhang, L. Qin, W. Zhang, and X. Lin. When engagement meets similarity: efficient (k, r)-core computation on social networks. *PVLDB*, 10(10):998–1009, 2017.
- [43] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *ICDE*, pages 1049–1060. IEEE, 2012.
- [44] Q. Zhu, H. Hu, C. Xu, J. Xu, and W.-C. Lee. Geo-social group queries with minimum acquaintance constraints. *VLDBJ*, pages 1–19, 2014.



Kai Wang received the BEng degree in Computer Science from Zhejiang University in 2016, and the PhD degree in Computer Science from the University of New South Wales in 2020. He is currently a research associate in the School of Computer Science and Engineering, University of New South Wales. His research interests lie in big data analytics, especially for the graph/network and spatial data.



Shuting Wang is currently working at Zhejiang Gongshang University. She is also conducting research as a volunteer researcher at the East China Normal University in data science Lab. She received Bachelor degree in 2014 from Jiangxi Normal university and Research-Master degree in 2017 from Zhejiang Gongshang University.



Xin Cao is currently a senior lecturer in the School of Computer Science and Engineering (CSE) at The University of New South Wales (UNSW). His research interests include Databases, Data Mining, and Artificial Intelligence. He received the Ph.D. degree from Nanyang Technological University (NTU), Singapore, in 2014.



Wenjie Zhang received the PhD degree in computer science and engineering from the University of New South Wales, in 2010. She is currently an associate professor and ARC DECRA (Australian Research Council Discovery Early Career Researcher Award) fellow in the School of Computer Science and Engineering, the University of New South Wales, Australia. Her research interests lie in big data management and processing.



Lu Qin received the BEng degree from the Department of Computer Science and Technology, Renmin University of China, in 2006, and the PhD degree from the Department of Systems Engineering and Engineering Management, Chinese University of Hong Kong, in 2010. He is currently an associate professor with the Centre for Artificial Intelligence, University of Technology, Sydney. His research interests include big graph analytics and graph query processing.

Efficient Computing of Radius-Bounded k -Cores

Kai Wang[†], Xin Cao[†], Xuemin Lin[†], Wenjie Zhang[†], Lu Qin[‡]

[†]University of New South Wales, [‡]University of Technology Sydney

kai.wang@unsw.edu.au, xin.cao@unsw.edu.au, lxue@cse.unsw.edu.au, wenjie.zhang@unsw.edu.au, lu.qin@uts.edu.au

Abstract—Driven by real-life applications in geo-social networks, in this paper, we investigate the problem of computing the radius-bounded k -cores (RB- k -cores) that aims to find cohesive subgraphs satisfying both social and spatial constraints on large geo-social networks. In particular, we use k -core to ensure the social cohesiveness and we use a radius-bounded circle to restrict the locations of users in a RB- k -core. We explore several algorithmic paradigms to compute RB- k -cores, including a triple-vertex-based paradigm, a binary-vertex-based paradigm, and a paradigm utilizing the concept of rotating circles. The rotating-circle-based paradigm is further enhanced with several pruning techniques to achieve better efficiency. The experimental studies conducted on both real and synthetic datasets demonstrate that our proposed rotating-circle-based algorithms can compute all RB- k -cores very efficiently. Moreover, it can also be used to compute the minimum-circle-bounded k -core and significantly outperforms the existing techniques for computing the minimum-circle-bounded k -core.

I. INTRODUCTION

With the wide availability of wireless communication techniques and GPS-equipped mobile devices (e.g., smart phones and tablets), people now can easily access the internet. This leads to the emergence of the geo-social networks, such as Twitter and Foursquare, where the social networks are combined with users' geo-spatial information. Consequently, retrieving subgraphs with high cohesiveness in geo-spatial social networks has become a popular research topic in recent years [1, 2, 3].

In this paper, we study the problem of efficiently computing *radius-bounded cohesive subgraphs* in a geo-spatial social network G (or geo-social network in short). That is, given a vertex q in a geo-social network G and a radius r , find all cohesive subgraphs g of G such that g contains q and all vertices in g fall into a circle with the radius r . There are many types of cohesive graphs in the literature [4, 5, 6]. While our proposed framework works generally for various cohesive subgraphs such as k -truss [5] and clique [6], in this paper we present our work restricted to a specific cohesive subgraph k -core [4] where each vertex has at least k neighbours.

Applications. The problem of computing *radius-bounded k -cores*, namely RB- k -cores, can have many real real-life applications. Indeed, in applications such as Facebook, Twitter, and Google+, personalized event recommendation is a very important part. Specifically, Events-For-You is a valuable part of Facebook, which can recommend events to users based on their personal locations and social relationships. Nevertheless, the current technology cannot provide a service of Events-For-You based on users' arbitrary requests. For example, the Events-For-You based on the following request by Leo cannot be accommodated by the current technology. In this example, Leo wants to hold a party (an Events-For-You activity) to play

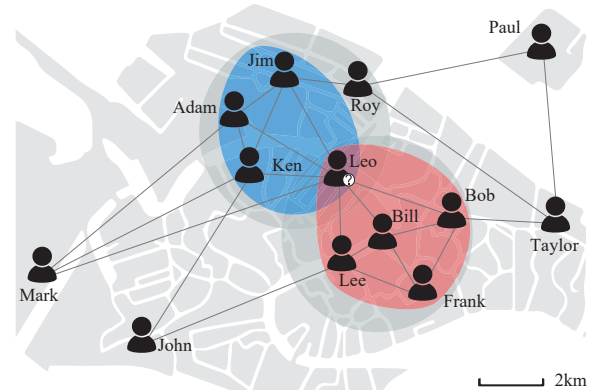


Figure 1: A geo-social network

board games (e.g., Monopoly, Uno, Risk, etc.) by gathering a group of people who are not living far away (say, bounded by a circle with radius r) and each of whom has many friends in the group (say, at least k friends). Figure 1 shows such a geo-social network where vertices represent users, edges represent friendships, and locations represent the home locations of users. If we set $r = 3$ and $k = 3$, there are two RB- k -cores recommended to Leo as illustrated by the shadow area, i.e., $\{\text{Leo, Ken, Jim, Adam}\}$ and $\{\text{Leo, Bill, Frank, Bob, Lee}\}$. This is a typical example of computing RB- k -cores.

Moreover, as studied in [3, 7], people with close social relationships tend to purchase in places that are also physically close. For this application, in a geo-social network, the locations of users represent places. To boost sales figures, advertisement messages can be sent to the RB- k -core of customers. For instance, if we want to promote an item A, the system can advertise A to the RB- k -core members (customers) using the query point (customers) who purchased A.

Existing Studies. There exist several studies over community retrieval in geo-social networks, but they are all different from our RB- k -cores problem. In the literature, various models including k -core [4], k -truss [5], and clique [6] have been studied to retrieve cohesive subgraphs without considering the spatial information of users. Thus, these models are not applicable to RB- k -cores. For example, in Figure 1, Mark will be added into the community formed by $\{\text{Leo, Ken, Adam, Jim}\}$ if we use the model k -core for $k = 3$ though Mark is far away from the other users. On the other hand, [8, 9] found a group of spatial objects without considering the network information, and thus they also cannot solve our problem RB- k -cores. In Figure 1, Roy will be added into the community formed by $\{\text{Leo, Ken, Jim, Adam}\}$ using spatial information (bounded by a circle) only but the network connections between Roy and the other people are very weak.

The most closely related works can be found in [1, 2, 3]

that consider both the social constraint (network structure) and the spatial constraint in retrieving communities. In particular, they all use k -core to ensure the social (graph) cohesiveness of communities in a network. [1] studied the community detection problem which uses pairwise similarity (distance) between each pair of vertices to ensure the spatial cohesiveness of communities while computing the maximum k -core or all maximal k -cores. Our experimental results in Section VI demonstrated that our problem is inherently different from the problem in [1]; that is, the generated results are very different. Moreover, our problem RB- k -cores is PTIME, while the problem in [1] is NP-hard. Zhu et al. in [2] studied the problem of finding the maximum k -core in a given rectangle containing a query vertex. [2] also studied the problem of finding the k -core with exact (or no less than) c vertices such that the longest distance from these vertices to q is minimized. These problems in [2] are also different from ours. Fang et al. in [3] studied the problem of computing the k -core containing the query vertex covered by the smallest circle. While the problem in [3] is different from our problem RB- k -cores, as a byproduct our techniques can be applied to the problem in [3] and can achieve a speed-up around twice.

Challenges. The main challenge of efficiently computing RB- k -cores is twofold.

- 1) The location of the radius-bounded circle of a RB- k -core is unknown. Therefore, it is a challenge to efficiently enumerate such circles.
- 2) During the process of finding RB- k -cores, it is time-consuming to construct and verify the candidate subgraphs individually. Therefore, it is important to reuse intermediate computation results and explore possible cost sharing which is challenging.

Contributions. We explore three paradigms to compute the RB- k -cores in this paper. The first paradigm is triple-vertex-based algorithm (TriV) inspired by [3]. It proposes to firstly generate all candidate circles containing q , secondly check the corresponding radius to verify the given radius bound, and then compute the maximum k -core for the vertices in each candidate circle. To avoid generating too many candidate circles or missing results, TriV generates two types of candidate circles: 1) determined by any three vertices in G (see O_2 in Figure 2(a)), and 2) determined by the two vertices u_1 and u_2 (see O_1 in Figure 2(a)); that is, using the distance of u_1 and u_2 as the diameter of O_1 and the centroid of the line u_1 and u_2 as the center of O_1 . Note that, it is necessary to consider the circles determined by two vertices, e.g., in Figure 2(a), if we set $q = u_1$, $r = r_1$, and $k = 2$, the RB- k -core contains $\{u_1, u_2, u_3\}$ will be missed if only considering the circle (O_2) determined by three vertices because of $r_2 > r_1$.

In TriV, we need to generate $O(n^3)$ candidate circles. To reduce the number of candidate circles, we propose the binary-vertex-based algorithm (BinV). In BinV, we effectively use the given parameter r and only generate the circles with r as the radius such that the circle arc passes a pair of vertices in G (for every pair). Generally, for each pair of vertices, we have at most two such circles (see O_1 and O_2 in Figure 2(b)). This guarantees to generate $O(n^2)$ candidate circles and reduces $O(n^3)$ candidate circles in TriV to $O(n^2)$.

We can observe that there are many reusable intermediate computation results in the process of finding RB- k -cores. The

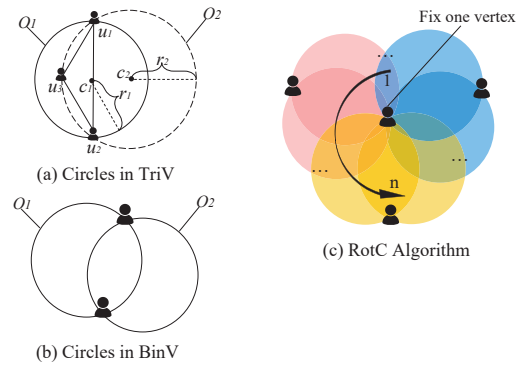


Figure 2: Illustrating proposed algorithms

third paradigm is to share computation costs among the computation of RB- k -cores. To effectively share the computation, we design the rotating-circle-based algorithm (RotC) so that the computation in BinV can be shared among the “adjacent” circles. Specifically, as shown in Figure 2(c), we fix a vertex u for each vertex in G then check the remaining vertices v such that for each pair u and v , we use BinV to generate the two circles with radius r (maybe degenerate to one if r is half of the distance between u and v). Then, we will share the computation among the adjacent circles.

Our principal contributions are summarized as follows.

- We proposed the model of RB- k -core and developed a novel paradigm to compute RB- k -cores with the aim to share computation.
- We proposed several new optimization techniques to speed up the computation.
- Extending our algorithms to the problem in [3] can achieve a speed-up around twice.
- We conducted comprehensive experiments on real geo-social networks to evaluate our algorithms.

Organization. The rest of the paper is organized as follows. Section II presents the preliminaries. Section III and section IV introduce the solutions based on TriV and BinV, respectively. Techniques based on rotating circles are presented in Section V. Section VI reports extensive experiments. Section VII reviews the related work. Section VIII concludes the paper.

II. PROBLEM DEFINITION

In this section, we formally introduce the fundamental concepts and definitions. Mathematical notations used throughout this paper are summarized in Table I.

Notation	Definition
G	a geo-social graph
$deg_G(v)$	the degree of vertex v in G
u, v	vertices in the geo-social graph
$d(u, v)$	the Euclidean distance between u and v
$O(c, \gamma)$	a circle centered at c with radius γ
$g(c, \alpha)$	a square centered at c with side length α
X, S	a set of vertices
$G(S)$	an induced subgraph of G formed from a set of vertices S
$W_\gamma(u, v)$	a set of binary-vertex-bounded circles with radius γ
\mathcal{R}	the result RB- k -core set

Table I: The summary of notations

Our problem is defined over a geo-social graph $G(V, E)$, where $V(G)$ denotes the vertex set, and $E(G)$ denotes the

edge set. The vertices represent the social network users and the edges represent their relationships in geo-social networks. Each vertex $v \in V(G)$ has a location $(v.x, v.y)$ which denotes the position of v along x - and y -axis in a two-dimensional space and the vertices are static in our problem. The Euclidean distance between u and v is denoted as $d(u, v)$. We denote the set of neighbors of each vertex v in G by $N_G(v) = \{u \in V(G) \mid (v, u) \in E(G)\}$ and the degree of vertex v by $deg_G(v) = |N_G(v)|$. We denote a circle centered at c with radius γ as $O(c, \gamma)$. Given a set of vertices $S \subseteq V(G)$, we use $G(S)$ to denote an induced subgraph of G formed from S such that $G(S) = (S, \{(u, v) \in E(G) \mid u, v \in S\})$.

Before formally defining the problem, we first introduce the following critical concepts to describe the social constraint and the spatial constraint.

Definition 1 (k -Core [4]). Given a graph G and a positive integer k , the k -core of graph G denoted as H_k is the maximal subgraph of G , where $deg_{H_k}(v) \geq k$, for each $v \in V(H_k)$.

Based on the k -core concept, we ensure the social constraint by restricting the minimal degree of vertices in a RB- k -core. Note that our proposed solutions can be easily adapted to other cohesive structure concepts (e.g., k -truss [5], clique [6]) which can be used to define the social constraint from different perspectives.

Definition 2 (Minimum Covering Circle (MCC)). Given a set of vertices S , the minimum covering circle of S is the circle which encloses all the vertices $v \in S$ with the smallest radius. We call the vertices which lie on the boundary of a MCC the boundary vertices.

After introducing k -core and MCC, we are ready to define the RB- k -core as below.

Definition 3 (Radius-Bounded k -Core). Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k and a query radius r , a subgraph of G denoted by G_k^r is a Radius-Bounded k -Core, if it satisfies the following constraints:

- 1) **Connectivity constraint.** $G_k^r \subseteq G$ is connected and contains q ;
- 2) **Social constraint.** $\forall v \in V(G_k^r), deg_{G_k^r}(v) \geq k$;
- 3) **Spatial constraint.** The MCC of $V(G_k^r)$ has a radius $r' \leq r$;
- 4) **Maximality constraint.** There exists no another RB- k -core $G_k^{r'} \supseteq G_k^r$ satisfying (1), (2), and (3).

Problem Statement: Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k , and a query radius r , our RB- k -core search problem aims to return all RB- k -cores in G .

Given a geo-social graph G , a query vertex q , and a query radius r , according to the spatial constraint in Definition 3, apparently, if the distance between a vertex v and the query vertex q is larger than $2r$, v cannot be included in any RB- k -cores. We call such vertices faraway vertices and we can first remove all these vertices from the given graph G . Given a positive integer k , we can also safely remove all the vertices which are not in the k -core of G containing q because of the social constraint in Definition 3. We use G_k to denote a connected subgraph of G which is a k -core containing q and for each vertex v in G_k , $d(q, v) \leq 2r$. We use $n = |V(G_k)|$ to

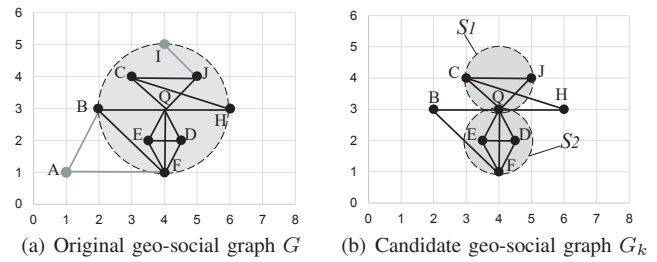


Figure 3: An example of geo-social network

represent the number of vertices and $m = |E(G_k)|$ to represent the number of edges of G_k in the following sections.

Example 1. Consider a geo-social graph G in Figure 3(a). Suppose Q is the query vertex, given $k = 2$ and $r = 1$, we want to find all RB- k -cores from this geo-social graph. We can safely remove vertex A because $d(A, Q) > 2r = 2$ and vertex I because I is not in the 2-core of G . Then we can obtain the candidate geo-social subgraph G_k which is shown in Figure 3(b). As a result, we can find two RB- k -cores $G(S_1)$ and $G(S_2)$, where $S_1 = \{Q, C, J\}$ and $S_2 = \{Q, D, E, F\}$.

III. TRIPLE-VERTEX-BASED ALGORITHM

In this section, we introduce the triple-vertex-based algorithm (TriV) which is designed based on the Exact algorithm in [3]. This algorithm lies on the following lemma.

Lemma 1. [10] Given a set S ($|S| \geq 2$) of vertices, the MCC of S can be determined by at most three vertices in S which lie on the boundary of the circle. If it is determined by only two vertices, then the line segment connecting those two vertices must be a diameter of the circle. If it is determined by three vertices, then the triangle consisting of those three vertices is not obtuse.

By Lemma 1, the MCC of a RB- k -core should have two or three vertices lying on its boundary, which are called boundary vertices. Thus, we can enumerate all candidate triple-vertex-combinations and binary-vertex-combinations, then check the subgraph enclosed by the circle fixed by the enumerated boundary vertices to see whether it is a RB- k -core. The details of TriV are shown in Algorithm 1.

Given a geo-social graph G , a query vertex q and parameters k and r , we first remove all the faraway vertices in $V(G)$. Then, we do a core decomposition in G using existing algorithms (e.g., [11]), and find the connected k -core G_k of G containing q . After obtaining G_k , we enumerate all candidate triple-vertex-combinations in $V(G_k)$ and compute the corresponding MCC $O(c, \gamma)$ (lines 2-8). If γ is smaller than r , we obtain a set of vertices $X = \{v \in G_k \mid d(v, c) \leq r\}$ and construct an induced connected subgraph $G(X) \subseteq G_k$ formed from X . If there exists a k -core $G(X)_k$ containing q in $G(X)$ and it satisfies the maximality constraint (for each $G' \in \mathcal{R}, G' \not\supseteq G(X)_k$), we put $G(X)_k$ into the result set \mathcal{R} (lines 9-12). We next enumerate all candidate binary-vertex-combinations in $V(G_k)$ to verify the circles which use the distance between the two vertices as the diameter. The verification process is similar with the triple-vertex-combination cases (lines 13-19). Finally, we obtain all RB- k -cores in \mathcal{R} .

Remark. Utilizing the spatial constraint and the maximality constraint, we skip the verification of a triple/binary-vertex-combination if it satisfies one of the following conditions: (1)

Algorithm 1: ALGORITHM TriV

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint perimeters

Output: \mathcal{R} : a set of RB- k -cores

```

1 initialize  $\mathcal{R} \leftarrow \emptyset$ 
2  $G_k \leftarrow$  the  $k$ -core of  $G$  containing  $q$  after removing faraway
  vertices
3 foreach node  $v \in V(G_k)$  do
4   foreach node  $u \in V(G_k)$  do
5     if  $u \neq v \wedge d(u, v) \leq 2r$  then
6       foreach node  $w \in V(G_k)$  do
7         if  $w \neq u \wedge d(w, u) \leq 2r \wedge w \neq$ 
8            $v \wedge d(w, v) \leq 2r$  then
9           compute MCC  $O(c, \gamma)$  of  $\{u, v, w\}$ 
10          if  $\gamma \leq r$  then
11             $X \leftarrow$  vertices enclosed in  $O(c, \gamma)$ 
12            construct  $G(X)$  from  $X$ 
13            if exists a  $G_k^r$  in  $G(X)$  then
14               $\mathcal{R}.$ update( $G_k^r$ )
15           $\gamma = \frac{d(u, v)}{2}$ ;
16          compute MCC  $O(c, \gamma)$  of  $\{u, v\}$ 
17           $X \leftarrow$  vertices enclosed in  $O(c, \gamma)$ 
18          construct  $G(X)$  from  $X$ 
19          if exists a  $G_k^r$  in  $G(X)$  then
20             $\mathcal{R}.$ update( $G_k^r$ )
21 return  $\mathcal{R}$ 

```

The distance between any pair of vertices in it is larger than $2r$. (2) The radius of its MCC is larger than r . (3) The candidate vertices enclosed by its MCC are all enclosed by the MCC of a candidate RB- k -core in the result set \mathcal{R} .

Theorem 1. *The time complexity of TriV is $O(n^3 \cdot (n + m))$.*

Proof. In Algorithm 1, we need to verify all candidate triple-vertex- and binary-vertex-combinations. For the triple-vertex-combinations, there are three for-loops which cost $O(n^3)$ to enumerate all the candidate circles. And we need $O(n)$ time cost to construct the induced subgraph $G(X)$ and $O(m)$ time cost to verify the existence of the k -core in $G(X)$. Hence, verifying the triple-vertex-combinations takes $O(n^3 \cdot (n + m))$. Similarly, verifying the binary-vertex-combinations takes $O(n^2 \cdot (n + m))$. In total, the time cost of TriV is $O(n^3 \cdot (n + m))$. ■

IV. BINARY-VERTEX-BASED ALGORITHM

The major issue of TriV is that we need to verify $O(n^3 + n^2)$ candidate subgraphs based on all triple-vertex- and binary-vertex-combinations. In this section, we introduce a binary-vertex-based algorithm which only needs to verify $O(n^2)$ candidate subgraphs to solve the RB- k -core search problem.

Based on the definition of RB- k -core, given a query radius r , an obvious observation is that for each RB- k -core in a geo-social graph G , it should be enclosed in at least one circle with radius r . A straightforward approach to find all RB- k -cores is verifying all the circles with radius r in the two-dimensional space. Obviously, there are too many circles with radius r sharing the same RB- k -core in this approach. In other words, for each RB- k -core, we just need to ensure that there is at least one circle with radius r enclosing it is checked. This can decrease the number of candidate circles significantly. Before introducing the algorithm, we give the definition of the binary-vertex-bounded circle as below.

Definition 4 (Binary-Vertex-Bounded Circle). *Given two vertices u and v , we call all circles having u and v lying on the boundary the binary-vertex-bounded circles. A set of binary-vertex-bounded circles with radius γ which takes u and v as bounded vertices is denoted as $W_\gamma(u, v)$.*

Lemma 2. [12] *Given two vertices u and v and a radius r ($r \geq d(u, v)$), we have:*

$$|W_r(u, v)| = \begin{cases} 1, & \text{iff. } d(u, v) = 2r, \\ 2, & \text{iff. } d(u, v) < 2r. \end{cases} \quad (1)$$

Lemma 3. *Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k , and a query radius r , for each RB- k -core G_k^r , all the vertices in $V(G_k^r)$ should be enclosed in at least one binary-vertex-bounded circle with radius r which takes u and v as the boundary vertices where $u, v \in V(G_k)$.*

Proof. By Definition 3, for a RB- k -core G_k^r in G_k , the MCC $O(c, \gamma)$ of G_k^r should have $\gamma \leq r$. (1) If $\gamma = r$, by Lemma 1, there should exist at least two vertices on the boundary of $O(c, \gamma)$, and thus Lemma 3 holds. (2) If $\gamma < r$, we take Figure 4 as an example to explain the proof process. We can choose any vertex v which lies on the boundary of $O(c, \gamma)$, and draw a circle $O_1(c_1, r)$ which internally tangents with $O(c, \gamma)$ at vertex v . Apparently, $O_1(c_1, r)$ can enclose all the vertices of $V(G_k^r)$ (enclosed in $O(c, \gamma)$). Then we can always rotate $O_1(c_1, r)$ anticlockwise and stop at the position (i.e., $O_2(c_2, r)$) when $O_1(c_1, r)$ first meets a vertex u in $O(c, \gamma)$. Because during this process all vertices of $V(G_k^r)$ are always in the intersection of two circles, $O_2(c_2, r)$, a binary-vertex-bounded circle which takes u and v as the boundary vertices, encloses all vertices of $V(G_k^r)$, and thus Lemma 3 holds. ■




Figure 4: Illustrating Lemma 3

By Lemma 2, two vertices can bound one/two circles with a given radius r . Based on Lemma 3, we can get all the RB- k -cores in G by verifying all the binary-vertex-bounded circles bounded by vertices in $V(G)$ with radius r . Hence, a more efficient algorithm BinV can be designed by verifying $O(n^2)$ candidate subgraphs constructed from the corresponding binary-vertex-bounded circles, rather than $O(n^3)$ candidate subgraphs as in TriV. We introduce the algorithm BinV which is shown in Algorithm 2.

We first obtain the k -core G_k of G containing q after removing all the faraway vertices in $V(G)$ (line 2). Then we enumerate all candidate binary-vertex-combinations with their distance $d \leq 2r$ in $V(G_k)$ as boundary vertices. For each binary-vertex-combination $\{u, v\}$, we compute $W_r(u, v)$ which contains one or two binary-vertex-bounded circles bounded by u and v (lines 3-6). For each binary-vertex-bounded circle $O(c, r) \in W_r(u, v)$, we obtain a set of vertices $X = \{x \in G_k \mid d(x, c) \leq r\}$ and construct a candidate subgraph $G(X)$ which is an induced subgraph of G_k formed from X . We then do a core decomposition to $G(X)$ and verify whether there exists a RB- k -core. If there exists, we put the

Algorithm 2: ALGORITHM BinV

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint perimeters

Output: \mathcal{R} : a set of RB- k -cores

```

1 initialize  $\mathcal{R} \leftarrow \emptyset$ 
2  $G_k \leftarrow$  the  $k$ -core of  $G$  containing  $q$  after removing faraway
  vertices
3 foreach node  $v \in V(G_k)$  do
4   foreach node  $u \in V(G_k)$  do
5     if  $u \neq v \wedge d(u, v) \leq 2r$  then
6       compute  $W_r(u, v)$  using  $\{u, v\}$  and  $r$ 
7       foreach  $O(c, r) \in W_r(u, v)$  do
8          $X \leftarrow$  vertices enclosed in  $O(c, r)$ 
9         construct  $G(X)$  from  $X$ 
10        if exists a  $G_k^r$  in  $G(X)$  then
11           $\mathcal{R}.update(G_k^r)$ 
12 return  $\mathcal{R}$ 

```

RB- k -core into the result set \mathcal{R} (lines 7-11). Finally, we can get all the RB- k -cores in \mathcal{R} .

Remark. Utilizing the spatial constraint and the maximality constraint, we skip the verification of a binary-vertex-combination if it satisfies one of the following conditions. (1) The distance between the two vertices is larger than $2r$. (2) The candidate vertices enclosed by its derived binary-vertex-bounded circle are all enclosed by the MCC of a candidate RB- k -core in the result set \mathcal{R} .

Theorem 2. *The time complexity of BinV is $O(n^2 \cdot (n + m))$.*

Proof. In Algorithm 2, we need to verify all the binary-vertex-bounded circles generated from candidate binary-vertex-combinations. There are two for-loops which cost $O(n^2)$ to enumerate all candidate binary-vertex-bounded circles in total. In addition, we need $O(n)$ time cost to construct the induced subgraph $G(X)$ and $O(m)$ time cost to find the k -core of $G(X)$. Therefore, the total time cost of BinV is $O(n^2 \cdot (n + m))$. ■

V. ROTATING-CIRCLE-BASED ALGORITHMS

Although the BinV algorithm improves TriV a lot by reducing the number of candidate subgraphs, it is still not efficient enough. Reviewing the process of BinV, we can observe that for each binary-vertex-combination, the corresponding candidate subgraph $G(X)$ is constructed and verified individually. There are $O(n^2)$ candidate graphs need to be constructed and each of the verification process takes $O(m)$ time. This motivates us to develop a better algorithm which can reduce the construction and verification cost of these candidate subgraphs.

In this section, we first present the rotating-circle-based algorithm (RotC) which improves the BinV algorithm by exploring possible cost sharing in the subgraph construction and verification process. Next, we employ some non-trivial pruning techniques to improve the RotC algorithm and propose the optimized rotating-circle-based algorithm (RotC⁺).

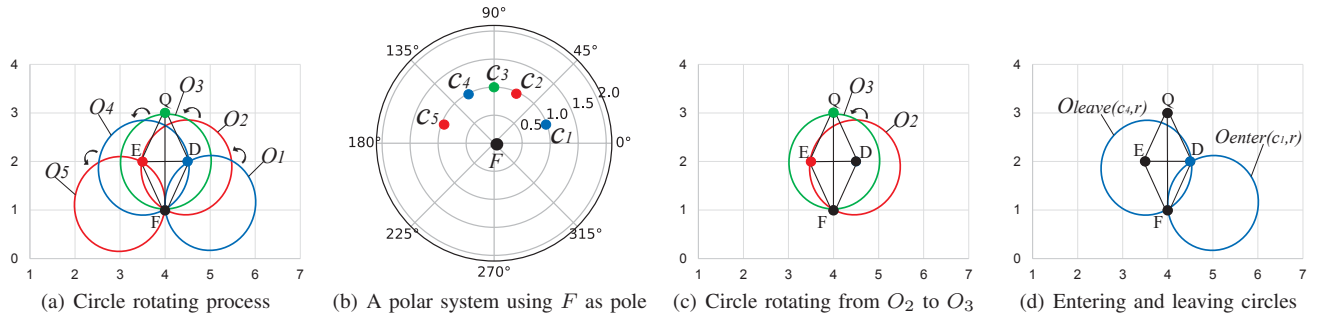
A. Algorithm RotC

Reviewing Lemma 3, we can find all the RB- k -cores in a geo-social graph G by verifying all the candidate subgraphs constructed from corresponding binary-vertex-bounded circles. Considering Example 1, Figure 5(a) is a screenshot of all

the binary-vertex-bounded circles which take F as one of the boundary vertices. In the BinV algorithm, we need to verify the candidate graphs enclosed by these binary-vertex-bounded circles one by one. Now we consider putting these binary-vertex-bounded circles into a polar coordinate system using F as the *pole*, and sorting these binary-vertex-bounded circles according to their centers' polar angles. Figure 5(b) shows the centers of binary-vertex-bounded circles in the polar coordinate system, and we can obtain a list of sorted circles $L = \{O_1, O_2, O_3, O_4, O_5\}$. Specifically, in Figure 5(c), O_2 and O_3 are two adjacent binary-vertex-bounded circles. We denote the vertex sets which O_2 and O_3 enclosed as X_2 and X_3 , respectively. We can observe that for these two induced subgraphs $G(X_2)$ and $G(X_3)$ where their binary-vertex-bounded circles adjacent to each other, $V(G(X_2))$ is only one vertex (vertex q) difference from $V(G(X_3))$. Based on this observation, we devise a novel algorithm which shares the construction and verification cost for these candidate subgraphs.

In the construction step, we can construct the candidate graphs incrementally after sorting all the binary-vertex-bounded circles. In the verification process, the degree of vertices are easy to maintain dynamically because the difference of enclosed vertices between adjacent binary-vertex-bounded circles is only one vertex. We can divide the binary-vertex-bounded circles into two groups, entering circles and leaving circles. An entering circle denoted as $O_{enter}(c, \gamma)$ is a circle which brings a new vertex in and a leaving circle $O_{leave}(c, \gamma)$ is a circle which takes an existing vertex out. For example in Figure 5(d), $O_{enter}(c_1, r)$ is an entering circles which brings vertex D in and $O_{leave}(c_4, r)$ is a leaving circle which takes D out of the candidate graph. So for an entering circle, we can avoid recomputing the degree of enclosed vertices when checking the k -core in a binary-vertex-bounded circle. For a leaving circle, we can just maintain the degree of vertices and avoid the computation of checking the k -core because there cannot exist a new k -core while a vertex leaves. We present the detailed rotating-circle-based algorithm (RotC) in Algorithm 3.

In RotC, we do a core decomposition and obtain the k -core G_k of G containing q after removing all the faraway vertices in $V(G)$ (line 2). After that, for each vertex v in $V(G_k)$, we set it as the pole in a polar coordinate system P . For each pole v , we generate a candidate vertex set $Y = \{u \in V(G_k) \mid d(u, v) \leq 2r\}$. Then we combine v with other candidate vertices in Y and construct the corresponding binary-vertex-bounded circles based on Lemma 3 and record whether it is an entering circle or a leaving circle for each binary-vertex-bounded circle (lines 3-7). Then, we sort all the binary-vertex-bounded circles in ascending order of their centers' polar angles in P (line 8). After sorting, for each binary-vertex-bounded circle $O(c, r)$, we compute a set X which contains all the vertices enclosed in O and maintain the degrees of these vertices (lines 9-12). Note that, we just need to insert/remove different vertices between O and its precedent binary-vertex-bounded circle, and the degrees of vertices in X can be updated correspondingly. If $O(c, r)$ is an entering circle, we need to construct a candidate graph $G(X)$ which is an induced subgraph of G_k formed from X (lines 13-14). After that, we verify whether there exists a k -core containing q in $G(X)$. Because the degrees of vertices in X is already maintained, in some cases such as $deg_{G(X)}(q) < k$, we can just skip doing a core decomposition in $G(X)$. Otherwise, if

Figure 5: An example of Rotating-Circle-Based Algorithm (using vertex F as the pivot)**Algorithm 3: ALGORITHM RotC**

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint perimeters

Output: \mathcal{R} : a set of RB- k -cores

```

1 initialize  $\mathcal{R} \leftarrow \emptyset$ 
2  $G_k \leftarrow$  the  $k$ -core of  $G$  containing  $q$  after removing faraway vertices
3 foreach node  $v \in V(G_k)$  do
4    $C \leftarrow \emptyset$ 
5   foreach node  $u \in V(G_k)$  do
6     if  $u \neq v \wedge d(u, v) \leq 2r$  then
7       compute  $W_r(u, v)$  using  $\{u, v\}$  and  $r$ 
8       put circles in  $W_r(u, v)$  into  $C$ 
9   sort  $C$  in ascending order of centers' polar angles
10  foreach  $O(c, r) \in C$  do
11     $X \leftarrow$  a set of vertices enclosed in  $O(c, r)$ 
12    maintain the degree of vertices in  $X$ 
13    if  $O(c, r)$  is an entering circle then
14      construct  $G(X)$  from  $X$ 
15      if exists a  $G_k^r$  in  $G(X)$  then
16         $\mathcal{R}.update(G_k^r)$ 
17 return  $\mathcal{R}$ 

```

a k -core exists and it satisfies the maximality property, we put the k -core into the result set \mathcal{R} (lines 15-16). Finally, we will get all the RB- k -cores in \mathcal{R} .

Example 2. Considering the geo-social graph in Example 1, suppose Q is the query vertex, $r = 1$ and $k = 3$, Figure 5 is an example of the RotC algorithm when it takes vertex F as the pole. As shown in Figure 5(a), we first get a set of five binary-vertex-bounded circles which take F as boundary vertex. Then we sort these circles in ascending order of their centers' polar angles as shown in Figure 5(b) and get a list of sorted circles $L = \{O_1, O_2, O_3, O_4, O_5\}$. After that, we start the rotating process from the first circle O_1 and get an induced subgraph $G(X_1)$ from $X_1 = \{D, F\}$. Because there is no k -core containing q in $G(X_1)$, we keep the rotating process and find that there is also no k -core containing q in $G(X_2)$ where $X_2 = \{D, E, F\}$. After rotating to O_3 , we can obtain a k -core in $G(X_3)$ where $X_3 = \{D, F, E, Q\}$. And there is no other result in O_3 and O_4 . So after verifying all the binary-vertex-bounded circles, we can get a candidate result set $R = \{G(X)\}$ where $X = \{D, F, E, Q\}$.

Theorem 3. The time complexity of RotC is $O(n^2 \cdot (\log n + m'))$, where $m' \ll m$.

Proof. In Algorithm 3, We first need to enumerate all the candidate vertices as poles, so there is one for-loop which costs $O(n)$. For each pole v , we combine it with the other vertices and generate all the binary-vertex-bounded circles which takes v as one of the boundary vertices. This process also costs

$O(n)$. Then we need to take $O(n \log n)$ to sort these circles in ascending order of their centers' polar angles. After that, the construction of the induced subgraph $G(X)$ will only cost $O(1)$ because the incremental maintenance of X . The degree of vertices also calculated incrementally and in some cases, the verification process just costs $O(1)$, so in average, finding the k -core of $G(X)$ costs $O(m')$ where $m' \ll m$. The total time cost of RotC is $O(n^2 \cdot (\log n + m'))$. ■

B. Algorithm RotC⁺

We continue to introduce the optimized rotating-circle-based algorithm (RotC⁺) which improves RotC significantly by utilizing some non-trivial pruning techniques. Specifically, we develop two types of pruning techniques, the grouping-based pre-process and the in-process pruning rules. In the pre-process part, we develop a grouping-based algorithm to reduce the number of candidate vertices. In the in-process part, we employ some critical pruning rules to optimize the rotating-circle-based algorithm.

Pre-Process Pruning. Firstly, we introduce the grouping-based pre-process pruning technique and it is based on the following lemma 4.

Lemma 4. Given a geo-social graph $G(V, E)$, a vertex $q \in V(G)$, a positive integer k and a query radius r , for each RB- k -core G_k^r in G , the center point c of the MCC $O(c, \gamma)$ of $V(G_k^r)$ should satisfy $d(c, q) \leq r$.

Proof. By Definition 3, for a RB- k -core G_k^r , the MCC $O(c, \gamma)$ of G_k^r should enclose q and satisfy $\gamma \leq r$. Hence we have $d(c, q) \leq r$ and we complete the proof. ■




Figure 6: An example of grouping-based pre-process

Lemma 4 illustrates that all the centers of MCCs of RB- k -cores are in the circle $O(q, r)$. Apparently, the circle $O(q, r)$ can be partitioned into four groups which are squares with size $r \times r$. Similarly, the group with size $r \times r$ can also be partitioned into 4 smaller groups with size $\frac{r}{2} \times \frac{r}{2}$. Hence, given a grouping parameter τ , as shown in Figure 6(a), we can partition the circle from 4 groups with size $r \times r$ to $4 \lceil \frac{r}{\tau} \rceil^2$ groups with size $\tau \times \tau$ iteratively. In each iteration, we halve the group size and prune the groups which do not need further verification. For

Procedure 1: GROUPING-BASED PRE-PROCESS

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint parameters; τ : grouping perimeter; \mathcal{R} : candidate result set

Output: G_k : a graph

- 1 $\alpha \leftarrow r$; $Y \leftarrow g(q, 2r)$
- 2 $G_k \leftarrow$ the k -core of G containing q after removing faraway vertices
- 3 **while** $\alpha \geq \tau$ **do**
- 4 **foreach** group $g(c, 2\alpha) \in Y$ **do**
- 5 partition $g(c, 2\alpha)$ into four groups with size $\alpha \times \alpha$ and put them into Y_c
- 6 $Y \leftarrow \emptyset$; $S \leftarrow \emptyset$
- 7 **foreach** group $g(c, \alpha) \in Y_c$ **do**
- 8 construct graph $G(X)$ using X which contains vertices enclosed in $O(c, r + \frac{\sqrt{2}}{2}\alpha)$
- 9 **if** exists a G_k^r in $G(X)$ **then**
- 10 $\mathcal{R}.\text{update}(G_k^r)$
- 11 **else if** exists a k -core $G(X)_k$ in $G(X)$ **then**
- 12 $Y.\text{insert}(g)$
- 13 put vertices in $V(G(X)_k)$ into S
- 14 **foreach** node $v \in V(G_k)$ **do**
- 15 **if** $v \notin S$ **then**
- 16 remove vertex v from $V(G_k)$
- 17 $\alpha = \alpha/2$
- 18 **return** G_k

example, in Figure 6(a), if we set $\tau = \frac{r}{2}$, the pre-process will run 2 iterations in total and in each iteration, we need to verify at most 4 and 16 groups with size = r and $\frac{r}{2}$, respectively.

We proceed to present the verification process of a given group of vertices denoted as $g(c, \alpha)$, where c is the center point and α is the side length. As shown in Figure 6(b), because the longest distance between c and the other points in $g(c, \alpha)$ is $\frac{\sqrt{2}}{2}\alpha$, we can use the circle $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ to enclose all the circles with radius r and centered at a point in $g(c, \alpha)$. In other words, for each circle $O(c', r)$ which centers at $g(c, \alpha)$ as shown in Figure 6(b), $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ can enclose it. Then we can construct a induced subgraph $G(X)$ of G using X which contains all the vertices enclosed in the circle $O(c, r + \frac{\sqrt{2}}{2}\alpha)$. If there exists no k -core containing q in $G(X)$, we can prune the whole group $g(c, \alpha)$. Otherwise, if the MCC $O(c', \alpha')$ of the k -core $G(X)_k$ containing q has the radius $\alpha' \leq r$, we can mark $G(X)_k$ as a candidate result and prune the whole group $g(c, \alpha)$, because $G(X)_k$ is the only result that can be found using the vertices in $g(c, \alpha)$. Otherwise, if the MCC $O(c', \alpha')$ of $G(X)_k$ has the radius $\alpha' > r$, the RB- k -cores obtained from $g(c, \alpha)$ are subsets of $G(X)_k$, and thus we add the vertices in $G(X)_k$ into a candidate vertex set and do further check for the group $g(c, \alpha)$.

We show the pre-processing in Procedure 1. We first put the largest group $g(q, 2r)$ into Y and compute the k -core G_k of G containing q after removing all the faraway vertices (lines 1-2). For each group in Y , we partition it into four smaller groups and put them into Y_c (lines 3-5). After that, we empty Y and S . For each group $g(c, \alpha)$ in Y_c , we check the induced subgraph formed from the vertices in $O(c, r + \frac{\sqrt{2}}{2}\alpha)$ and put $g(c, \alpha)$ into Y if it needs further checking. Furthermore, we put the vertices which need further checking into S (lines 7-13). Next, we prune the vertices not in S from $V(G_k)$ (lines 14-16). Then we halve α and check whether $\alpha \geq \tau$ (line 17). If $\alpha < \tau$, we return the graph G_k after pruning.

In-Process Pruning. We next review the rotating-circle-based algorithm and introduce two in-process pruning rules.

Pruning Rule 1: Overall Checking. Reviewing the process of the RotC algorithm, we choose a vertex v from $V(G_k)$ as the pole and generate a candidate vertex set $S = \{u \in V(G_k) \mid d(u, v) \leq 2r\}$. Then we construct an induced subgraph $G(S)$ using vertices in S and compute the k -core $G(S)_k$ of $G(S)$ containing q . If $G(S)_k$ doesn't exist or the vertices in $V(G(S)_k)$ are all enclosed in the MCC of a candidate RB- k -core in \mathcal{R} , we can prune the pole v .

Pruning Rule 2: Circle Filtering. In the RotC algorithm, after choosing the pole v and corresponding candidate vertices, we combine v with all candidate vertices and generate the binary-vertex-bounded circles. Firstly, we can prune all the circles which exclude the query vertex q . After that, because there is only one vertex difference between two adjacent circles, we can compute the vertex difference between a circle and its precedent. We divide the circles into two groups, the entering circles and leaving circles, respectively. For each entering circle, we record the vertex it brings in and for each leaving circle, we record the vertex it moves out.

For the group of entering circles, we sort them in ascending order of their centers' polar angles and put them into a list L_{enter} . Then for each entering circle O_{enter} in L_{enter} , we compute a vertex set $\mathcal{V}(O_{enter})$ which contains all the vertices bringing from the entering circles appear before O_{enter} in L_{enter} . This can be done by incrementally adding the vertices bringing from the first entering circle to the last entering circle and the time complexity is $O(L_{enter})$. It is obvious that the number of vertices in $\mathcal{V}(O_{enter})$ monotonously increase with the index of O_{enter} in L_{enter} . Thus, we can use binary search to find the first entering circle O'_{enter} in L_{enter} such that we can construct a k -core from $\mathcal{V}(O'_{enter})$ containing q . The circles appear before O'_{enter} can be safely discarded because they cannot contain a RB- k -core. Similarly, for all the leaving circles, we sort in descending order of their centers' polar angles and put them into L_{leave} . In the same way, we can find the first leaving circle O'_{leave} such that we can construct a k -core from $\mathcal{V}(O'_{leave})$ containing q and discard all the circles before O'_{leave} in L_{leave} . In this way, we can reduce the number of binary-vertex-bounded circles which need to be verified in the next stage.

Algorithm 4 shows the optimized rotating-circle-based algorithm (RotC⁺). We first initialize the result set \mathcal{R} and do the grouping-based pre-process (lines 1-2). After that, for each vertex v in $V(G_k)$, we set it as the pole in a polar coordinate system P . For each pole v , we generate a candidate vertices set $S = \{u \in V(G_k) \mid d(u, v) \leq 2r\}$. We generate binary-vertex-bounded circles using the combinations between v and all candidate vertices in S (lines 3-9) and record whether it is an entering circle or a leaving circle. Then we apply pruning rule 1 to do a overall checking for the graph constructed from S (lines 10-11). Next, we sort all the binary-vertex-bounded circles in ascending order of their centers' polar angles in P . After sorting, we use pruning rule 2 to reduce the number of circles in C (lines 12-13). Then for each binary-vertex-bounded circle $O(c, r) \in C$, we compute a set X which contains all the vertices enclosed in $O(c, r)$ and maintain the degrees of these vertices. If $O(c, r)$ is an entering circle, we construct a candidate graph $G(X)$ which is an induced subgraph of G_k formed from X (lines 14-18). After that, if

Algorithm 4: ALGORITHM RotC⁺

Input: $G(V, E)$: the input graph; q : the query vertex; k, r : constraint perimeters; τ : grouping perimeter

Output: \mathcal{R} : a set of RB- k -cores

```

1 initialize  $\mathcal{R} \leftarrow \emptyset$ 
2  $G_k \leftarrow PreProcess(G, q, k, r, \tau, \mathcal{R})$ 
3 foreach node  $v \in V(G_k)$  do
4    $C \leftarrow \emptyset$ ;  $S \leftarrow \emptyset$ 
5   foreach node  $u \in V(G_k)$  do
6     if  $u \neq v \wedge d(u, v) \leq 2r$  then
7       put  $u$  into  $S$ 
8       compute  $W_r(u, v)$  using  $\{u, v\}$  and  $r$ 
9       put circles in  $W_r(u, v)$  into  $C$ 
10  if  $OverallChecking(S) = false$  then
11    continue ▷ Pruning Rule 1
12  sort  $C$  in ascending order of centers' polar angles
13  employ circle filtering to  $C$  ▷ Pruning Rule 2
14  foreach  $O(c, r) \in C$  do
15     $X \leftarrow$  a set of vertices enclosed in  $O(c, r)$ 
16    maintain the degree of vertices in  $X$ 
17    if  $O(c, r)$  is an entering circle then
18      construct  $G(X)$  from  $X$ 
19      if exists a  $G_k^r$  in  $G(X)$  then
20         $\mathcal{R}.update(G_k^r)$ 
21 return  $\mathcal{R}$ 

```

there exists a RB- k -core in $G(X)$, we put it into the result set \mathcal{R} (lines 19-20). Finally, we can get all the RB- k -cores in \mathcal{R} .

Theorem 4. *The time complexity of RotC⁺ is $O(\lceil \frac{r}{\tau} \rceil^2 \cdot m \cdot (\log(\lceil \frac{r}{\tau} \rceil) + 1) + |F| \cdot m + |F_1| \cdot \log |F_1| \cdot (|F_1| + m) + |F_1| \cdot |F_2| \cdot m')$, where F denote the candidate vertex set after pre-process pruning ($|F| < n$), F_1 is the vertex set obtained from F after the overall checking, F_2 is the set of circles need to be verified ($|F_2| < n$), and m' is the average time cost of verifying the existence of a k -core ($m' \ll m$).*

Proof. The RotC⁺ algorithm consists of four phases: (1) In the grouping-based pre-process, it needs to run $(\log(\lceil \frac{r}{\tau} \rceil) + 1)$ iterations and for each iteration, at most $O(\lceil \frac{r}{\tau} \rceil^2)$ groups are verified and the verification takes $O(m)$ for each group. In total, the pre-process costs $O(\lceil \frac{r}{\tau} \rceil^2 \cdot m \cdot (\log(\lceil \frac{r}{\tau} \rceil) + 1))$. (2) We denote the candidate vertex set after pre-process pruning as F where $|F| < n$. The overall checking in pruning rule 1 will cost $O(m)$ for each vertex in F . (3) After the overall checking, some vertices will be pruned, and we put the remaining vertices into F_1 . for each vertex in F_1 , we need to use $O(|F_1| \cdot \log |F_1|)$ time cost to sort the binary-vertex-bounded circles and $O(\log(|F_1| \cdot m))$ time cost to do the circle filtering pruning. (4) Finally, for each vertex in F_1 , we put the binary-vertex-bounded circles still need to be verified into F_2 ($|F_2| < n$) and the average time of verifying the existence of k -core costs $O(m')$ where $m' \ll m$. In total, The time complexity of the RotC⁺ algorithm is $O(\lceil \frac{r}{\tau} \rceil^2 \cdot m \cdot (\log(\lceil \frac{r}{\tau} \rceil) + 1) + |F| \cdot m + |F_1| \cdot \log |F_1| \cdot (|F_1| + m) + |F_1| \cdot |F_2| \cdot m')$. ■

VI. EXPERIMENTS

In this section, we report the evaluation of the effectiveness of our model and the efficiency of our algorithms.

A. Experiments Setting

Algorithms. To our best knowledge, there is no existing works that can solve the problem of finding RB- k -cores in geo-social graphs. In the experimental study, we implement and

evaluate four algorithms: The triple-vertex-based algorithm TriV in Section III, the binary-vertex-based algorithm BinV in Section IV, the rotating-circle-based algorithm RotC in Section V, and the optimized rotating-circle-based algorithm (RotC⁺) in Section V.

In addition, we extend our RotC⁺ algorithm to solve the SAC (spatial-aware community) search problem proposed in [3]. Given a query vertex q and a parameter k , a SAC is a connected k -core containing q enclosed in a circle with the minimum radius. If there exists a connected k -core containing q enclosed in a circle with a given radius r , our algorithms can always find it out. Thus, the minimum radius can be found by employing the binary search strategy using our algorithms. Given q , r , and k , if there exists no RB- k -core in a geo-social network, the radius of the SAC's MCC must be larger than r , and we need to increase r in the subsequent search. Otherwise, if we find a RB- k -core G_k^r , by definition we know that the MCC of $V(G_k^r)$ has a radius no larger than r , and thus r is an upper bound of the radius of the SAC's MCC, and we can decrease r in the subsequent search. We use our best algorithm RotC⁺ in the binary search process to find SACs.

The algorithms are implemented in C++ and the experiments are run on a Linux server with Intel Xeon E5-2687W(3.4GHz, 8 Cores) processor and 64GB main memory. We randomly select 200 query vertices and report the average result for these queries. We terminate an algorithm if the running time is more than three hours.

Dataset	$ V $	$ E $	d_{avg}
Brightkite	51,406	197,167	7.67
Gowalla	107,092	456,830	8.53
Flickr	214,698	2,096,306	19.5
Foursquare	2,127,093	8,640,352	8.12
Synthetic	4,000,000	40,000,000	20

Table II: Summary of Datasets

Datasets. We use four real datasets in our experiments including Brightkite, Gowalla, Flickr, and Foursquare. In the four datasets, we consider each user associated with a geo-location coordinate (latitude and longitude) as a vertex and the friendship between two users is represented by an edge. Helmert transformation [13] is adopted to transform geo-location coordinates of vertices to Cartesian coordinates. The original data of Brightkite and Gowalla was downloaded from <http://snap.stanford.edu/>. The geo-locations are extracted from the check-in data recorded by the Brightkite and Gowalla services, and the friendship network is obtained using the public API of the two services. The Flickr dataset is downloaded from <https://www.flickr.com/>. In Flickr, the photos taken by users are associated with geo-locations. We use the location where a user takes the most number of photos as his geo-coordinate. The social graph is constructed using the friendship information in Flickr. The Foursquare dataset [14] is downloaded from <https://archive.org/> which contains the data extracted from the Foursquare application. Each user in Foursquare has a unique id and a geo-location. The social graph represents the user relationships of the Foursquare users. Because of the limitation of space, we only show the comparison of the performance on Gowalla (the most commonly used dataset among the four datasets) and Foursquare (the largest dataset among the four datasets) in some of our experiments. Table II shows the number of vertices ($|V|$), number of edges ($|E|$) and average degree (d_{avg}) of vertices in these four datasets.

Parameter	Range	Default
k	4,7,10,13,16	4
r	1,5,10,20,40	5
n	20%,40%,60%,80%,100%	100%
τ	$r, \frac{r}{2}, \frac{r}{4}, \frac{r}{8}, \frac{r}{16}$	$\frac{r}{4}$

Table III: Summary of Parameters

We also conduct experiments on a synthetic dataset *Synthetic*. We first generate a non-spatial graph using a well-known graph generator GTGraph downloaded from <http://www.cse.psu.edu/~kxm85/software/GTgraph/>. The distribution of the degrees in the graph follows a power-law distribution which is often used in the study of social networks. After generating the graph, we generate the locations of the vertices randomly in a square with size $[0, 300]\text{km} \times [0, 300]\text{km}$.

Parameters. The experiments are conducted using different settings on 4 parameters: k (the minimum degree), r (the maximal radius), τ (the parameter used in the pre-process of RotC⁺), and n (the percentage of vertices). Table III shows the ranges and default values of these parameters. We vary k from 4 to 16 and set 4 as the default value. We vary r from 1km to 40km and set r to 5km by default. When varying the graph size, we randomly sample 20% to 100% vertices of the original graphs, and construct the induced subgraphs using these vertices. The parameter n is varied from 20% to 100% which represents the percentage of the vertices we use in each dataset. The parameter τ is varied from r to $\frac{r}{16}$ which controls the number of iterations of the pre-processing in RotC⁺.

B. Effectiveness

In this section, we first conduct two case studies on Gowalla and Flickr to show the effectiveness of our RB- k -core model. Then we do a comparison with the (k,r) -core model to show the difference between our RB- k -core model and the (k,r) -core model.

Figure 7: Case study on Gowalla ($q=1396$, $k=3$, $r=0.76\text{km}$)

Case study. We present two case studies to show the result of RB- k -core search on Gowalla and Flickr in Figure 7 and Figure 8, respectively. The query vertices are marked by question mark symbols. Under setting $q=1396$, $k=4$ and $r=0.76\text{km}$ on Gowalla, we can get two RB- k -cores containing q as shown in Figure 7. We mark the vertices and the MCC of these two RB- k -cores in black color and grey color, respectively. We can see that, the social constraint and the spatial constraint both contribute to the construction of these two RB- k -cores. For example, if the social constraint is ignored, the black vertices enclosed by the grey circle will be included in the grey RB- k -core. On the other hand, all the vertices in Figure 7 will be united into one community if the



Figure 8: Case study on Flickr ($q=111419$, $k=3$, $r=1.67\text{km}$) radius constraint is not being considered. Figure 8 shows the result of RB- k -core search on Flickr using $q=111419$, $k=3$ and $r=1.67\text{km}$ there are also two communities can be retrieved. Using the same q and k , the SAC search (i.e., a similar model in [3]) will provide the communities with black color as shown in Figure 7 and Figure 8. The radius of the black circles are 0.74km and 1.67km in Figure 7 and Figure 8, respectively. Comparing with the SAC search, in Figure 7, our RB- k -core search can give users more options by slightly increasing the minimum radius (i.e., from 0.74km to 0.76km) for the same q and k . In Figure 8, the RB- k -core search is able to provide users more than one selection for the same q , k , and minimum r .

Comparison with (k,r) -core [1]. We conduct experiments on Gowalla to show the difference between our RB- k -core model and the (k,r) -core model. A (k,r) -core is defined as a k -core where the distance between each pair of vertices is no more than r . Zhang et al. [1] only solve a community detection problem which enumerates all the (k,r) -cores without considering the query vertices. We consider the (k,r) -core search problem based on the (k,r) -core enumeration problem by adding a query vertex q . We compare the result of our RB- k -core search problem and the (k,r) -core search problem to show the model difference.


The query results of the two problems are two sets of k -cores. We use the set-similarity proposed in [15] to measure the similarity of the results, described as below. This measure first defines a similarity function to compute the similarity between two elements x and y (i.e., vertex sets of k -cores in our problem). Given a similarity threshold β , the similarity function ϕ_β is defined as:

$$\phi_\beta(x, y) = \begin{cases} \frac{|x \cap y|}{|x \cup y|}, & \text{if } \frac{|x \cap y|}{|x \cup y|} \geq \beta, \\ 0, & \text{if } \frac{|x \cap y|}{|x \cup y|} < \beta. \end{cases} \quad (2)$$

Given two sets R and S and a similarity function ϕ_β , the set-similarity of this measure is defined as:

$$\text{similar}_{\phi_\beta}(R, S) = \frac{|R \tilde{\cap}_{\phi_\beta} S|}{|R| + |S| - |R \tilde{\cap}_{\phi_\beta} S|} \quad (3)$$

Using ϕ_β , a bipartite graph \mathcal{H} can be constructed between R and S . In \mathcal{H} , vertices represent the elements in R or S and edges are weighted using the similarity function ϕ_β . Then we employ the maximum weighted bipartite matching [16] algorithm to get the maximum matching of \mathcal{H} and sum up the weights of the edges in the maximum matching as the value $|R \tilde{\cap}_{\phi_\beta} S|$. Then, we can get the similarity between R and S by equation 3.

Figure 9: Comparison with (k,r) -core

The experimental results are as shown in Figure 9. We vary r and k on Gowalla in Figure 9(a) and Figure 9(b), respectively, and we fix the other parameters as the default value. We randomly select 200 query vertices from the result of the (k,r) -core enumeration problem and report the average result for these vertices. Note that, given a radius r , the largest distance between two vertices in a RB- k -core is $2r$, so we set the similarity threshold to $2r$ in the (k,r) -core search problem because it considers the pair-wise similarity. In Figure 9, we study the impact of β and show the corresponding similarity between the results of RB- k -core search and (k,r) -core search. We can observe that the similarities are smaller with a larger β in both Figures 9(a) and 9(b), since more elements are considered as similar. The most important observation is that the similarities are all less than 0.65 even we set the β to 0.8 which is a very small threshold to measure the similarity of two vertex sets. This demonstrates that our RB- k -core model differs from the (k,r) -core model, because (k,r) -core uses the pairwise similarity. Such a strong spatial constraint makes the problem NP-hard, but it is not necessary for many applications, especially in community search problems.

C. Efficiency

In this section, we first evaluate the efficiency of the proposed four algorithms. Then we evaluate the effect of the pruning techniques and the parameter τ used in the RotC⁺ algorithm. Finally, we extend our RotC⁺ algorithm to solve the SAC search problem [3] and compare the performance.





Figure 10: Effect of different datasets

Evaluating the performance of all algorithms on different datasets. In Figure 10, we show the performance of our RB- k -core search algorithms on five datasets. We set k as default and r to 1km, 5km, 10km, 20km, 40km on Brightkite, Gowalla, Flickr, Foursquare and Synthetic, respectively, and we fix the other parameters as the default value. We can observe that BinV is efficient than TriV on Brightkite, Gowalla, and Flickr. The algorithms RotC and RotC⁺ using the rotating circle strategy are more efficient than TriV and BinV on the three datasets because RotC and RotC⁺ can compute the RB- k -cores in an incremental manner, which significantly reduces the computation cost. On Foursquare and Synthetic, we can see that only RotC⁺ is able to return the results within the timeout threshold. Foursquare is much larger than the first three

datasets, and there exist many candidate vertices that need to be processed. On Synthetic, both the size and the density of vertices is much larger than the other datasets, and thus the candidate circles contain many vertices. In summary, Figure 10 demonstrates the efficiency of our RotC⁺ algorithm because it significantly outperforms the other three algorithms on all datasets.

(a) Gowalla, varying k (b) Foursquare, varying k Figure 11: Effect of k

Evaluating the effect of k . Figure 11 evaluates the effect of k for four algorithms on Gowalla and Foursquare. We vary k from 4 to 16 and fix the other parameters as the default value. In Figure 11(a), we can observe that the time cost of all four algorithms drops when k increases because the number of vertices in the k -core of the original graph (selected as candidate vertices) decreases. Similar trends can be observed in Figure 11(b). As expected, RotC and RotC⁺ significantly outperform TriV and BinV on both datasets because of using the rotating circle technique. For example, on both datasets, RotC is about one order of magnitude faster than TriV and BinV and RotC⁺ is at least two orders of magnitude faster than TriV and BinV.



(a) Gowalla, varying n (b) Foursquare, varying n

Figure 12: Effect of graph size

Scalability. (1) Evaluating the effect of graph size. Figure 12 studies the scalability of four algorithms by varying the graph size from 20% to 100% in all datasets. We can observe that, on Gowalla, all these four algorithms scale almost linearly and the computation cost of them all increases when the percentage of vertices increases. In Figure 12(a), we can see that TriV and BinV can only get the result when $n = 20\%$ in the given timeout threshold while RotC and RotC⁺ have similar trends with that in Figure 12(a) on Foursquare. As discussed before, RotC⁺ is more efficient than the other three algorithms.

(a) Gowalla, varying r (b) Foursquare, varying r Figure 13: Effect of r

(2) *Evaluating the effect of r .* Figure 13 studies the effect of

r on Gowalla and Foursquare. We vary r from 1km to 40km and fix the other parameters as the default value. In Figures 11(a) and 11(b), the time cost increases as r becomes larger because the number of vertices in circle $O(q, 2r)$ grows when r increases. We can also see that, on Gowalla, both RotC and RotC⁺ are several orders of magnitude faster than TriV and BinV. On Foursquare, TriV and BinV can only compute the result when $r = 1$ km and RotC can get the results when r is no more than 10km in a reasonable time period. As expected, the RotC⁺ algorithm significantly outperforms the other three algorithms on Foursquare and the time cost is stable when r is large on both datasets.




Figure 14: Effect of Pruning Rules

Evaluating the pruning techniques. We evaluate the efficiency of our pre-process and in-process pruning techniques on Gowalla and Foursquare in Figure 14. On Gowalla, we vary r from 1km to 40km and fix the other parameters. On Foursquare, we vary k from 4 to 16 and fix the other parameters. The RotC-IP represents the RotC algorithm with in-process pruning techniques. Then by employing the grouping-based pre-process pruning techniques, we get the RotC⁺ algorithm. In Figure 14(a) and Figure 14(b), we can see that, the in-process pruning technique significantly reduces the computation cost while the pre-process pruning technique also enhances the performance of our RotC⁺ algorithm.




Figure 15: Effect of τ

Evaluating the effect of τ . Figure 15 studies the effect of τ which is a parameter used in the grouping-based pre-processing in RotC⁺. Because the value of τ is related to r , we set r to 1km, 5km, 10km, 20km, and 40km on both Gowalla and Foursquare. As discussed before, as τ increases, the time cost of pre-processing increases and the number of candidate vertices decreases. We can observe that, the running time is not very sensitive with τ when τ is relatively large on the two datasets. The time cost starts to increase from $\tau = \frac{r}{4}$ in most cases, because the number of vertices that can be pruned increase slowly and the time cost of pre-processing begins to dominate the cost of RotC⁺. Hence we set $\tau = \frac{r}{4}$ in our experiments on all datasets.

Extend to solve the SAC search problem [3]. As discussed before, the SAC search problem can be solved by slightly modifying our RotC⁺ algorithm using binary search. In Figure 16, we study the performance of the SAC-RotC⁺ algorithm which is extended from RotC⁺ to solve the SAC search




Figure 16: Extend to solve SAC search problem

problem, and we compare its performance with the state-of-the-art exact algorithm SAC-Exact⁺ proposed in [3]. Fang et al. [3] implemented the SAC-Exact⁺ algorithm in JAVA, while we implement the SAC-Exact⁺ algorithm in C++ for the fairness of comparison.

The SAC-Exact⁺ algorithm includes two phases. Firstly, it conducts the quad-tree-based vertex pruning phase which can reduce the number of potential vertices. Next, in the second phase, it conducts a triple-vertex-based algorithm which is similar with the TriV algorithm in this paper. In the RB- k -core search problem, we have analyzed that the triple-vertex-based algorithm is time-consuming and it can be improved by the rotating circle strategy to incrementally compute the result. We can do the same thing in the SAC search problem. In our SAC-RotC⁺ algorithm, we also conduct the vertex pruning phase, but we adopt the rotating-circle-based algorithm in the second phase. Note that, the in-process pruning technique in RotC⁺ can also be applied in SAC-RotC⁺, but the pre-process pruning technique cannot be used because of the model difference.

We vary the parameter ϵ which controls the number of iterations in the vertex pruning phase, and the number of iterations decreases with the increase of ϵ . From Figure 16(a) and Figure 16(b), we can observe that the time cost of SAC-RotC⁺ and SAC-Exact⁺ is almost the same when ϵ is very small because the cost of processing the vertex pruning phase dominates the cost in the second phase. On Foursquare, SAC-RotC⁺ outperforms SAC-Exact⁺ when ϵ is larger than 10^{-3} . Also, on Gowalla, SAC-RotC⁺ is about one order of magnitude faster than SAC-Exact⁺ when ϵ is larger than 10^{-4} . This is because the time cost on the second phase begins to dominate the first phase when ϵ is large, and our SAC-RotC⁺ algorithm computes the result in an incremental manner which significantly outperforms the triple-vertex-based algorithm in the second phase. Comparing the minimal time cost of the two algorithms on both datasets in our experiments, we can conclude that SAC-RotC⁺ can achieve a speed-up around twice.

VII. RELATED WORK

Community retrieval has been widely studied and used in many applications such as location-aware marketing [7], influence analysis [17], and event recommendation [18].

Prior works studied various models such as k -core [4], k -truss [5], and clique [6] to retrieve communities based on users' social connections. In the studies of k -core, efficient algorithms have been proposed in [11, 19] for core decomposition. Huang et al. [20] proposed algorithms to compute the community based on k -truss and various algorithms for clique computation have been studied such as in [21, 22].

Based on these models, existing works considering social cohesiveness of users can be categorized into two types,

i.e., community detection [23, 24] and community search [25, 26, 27, 28, 29]. In the studies of the community detection problem, Lee et al. [23] proposed a model (k, d)-core which uses k to control the edge density and d to control the number of common neighbours of two vertices of an edge. Cai et al. [24] studied the concept of community profiling based on community detection to describe the profile of a community using published contents and diffusion links of users. Different from community detection, a query user is given in the community search problem. Several studies have been proposed to retrieve communities containing the query vertex in an online manner. For example, Sozio et al. [25] proposed algorithms to find a community for a set of given query vertices. The local search strategy was proposed by Cui et al. [26] to evaluate the core number of vertices in communities. Li et al. [27] studied the influential community search problem based on the weighted graphs to capture the influence of communities. In addition, some works [28, 29] studied the community search problem on attributed graphs which use a set of keywords as the attributes. However, the geo-locations of users were not considered in these works.

In spatial databases, several works studied the group objects retrieval problem based on users' spatial locations such as [8, 9, 30] and [31]. Guo et al. [8] studied the spatial keyword query which retrieves a group of objects close to each other and cover a set of keywords together. Wu et al. [9] adapted the densest subgraph model to the spatial community search problem on dual networks. The work [30] proposed localitySearch which retrieves top- k sets of spatial web objects by integrating spatial distance, textual relevance, and a "co-locality" measure into one ranking function. The work [31] focused on context-aware search over social media data. It analysed the data-centric challenges in temporal, spatial, and spatio-temporal contexts. These proposals did not consider the social connections of users, and thus they are different from our problem.

Recently, some works studied the community retrieval problem [1, 2, 3, 9, 32] considering both the spatial and social features. The works [1, 32] solving the community detection problem mainly focused on analyzing and understanding the complexity networks rather than online community search. The most closely related work in [2] is that Zhu et al. studied finding a community within a given rectangle. Their study is different from our work because what we consider is restricting the size of community spatially instead of within a given rectangle. Fang et al. [3] proposed both exact and approximate algorithms to find a community covered by the smallest circle for a given query vertex. In their work, the radius of circle is not given by users and only one community covered by the smallest circle is returned to users, and thus it cannot provide more options for users as done by our work.

VIII. CONCLUSION

In this paper, we study the RB- k -core search problem. We propose a triple-vertex-based algorithm and a binary-vertex-based algorithm as benchmark algorithms. We propose a rotating-circle-based algorithm which can find possible cost sharing when solving the RB- k -core search problem. The rotating-circle-based algorithm is further enhanced with critical pruning techniques. We conduct extensive experiments on both real and synthetic datasets and the experimental result shows that our rotating-circle-based algorithm significantly

outperforms the benchmark algorithms. In the future, we plan to study the RB- k -core search problem when users are moving.

IX. ACKNOWLEDGMENT

Xuemin Lin is supported by NSFC 61672235, DP170101628 and DP180103096, Wenjie Zhang is supported by ARC DP180103096, Lu Qin is supported by ARC DP160101513.

REFERENCES

- [1] F. Zhang, Y. Zhang, L. Qin, W. Zhang, and X. Lin, "When engagement meets similarity: efficient (k, r)-core computation on social networks," *PVLDB*, vol. 10, no. 10, pp. 998–1009, 2017.
- [2] Q. Zhu, H. Hu, C. Xu, J. Xu, and W.-C. Lee, "Geo-social group queries with minimum acquaintance constraints," *VLDBJ*, pp. 1–19, 2014.
- [3] Y. Fang, R. Cheng, X. Li, S. Luo, and J. Hu, "Effective community search over large spatial graphs," *PVLDB*, vol. 10, no. 6, pp. 709–720, 2017.
- [4] S. B. Seidman, "Network structure and minimum degree," *Social networks*, vol. 5, no. 3, pp. 269–287, 1983.
- [5] J. Cohen, "Trusses: Cohesive subgraphs for social network analysis," *NSATR*, vol. 16, 2008.
- [6] R. D. Luce and A. D. Perry, "A method of matrix analysis of group structure," *Psychometrika*, vol. 14, no. 2, pp. 95–116, 1949.
- [7] D. McKenzie-Mohr, *Fostering sustainable behavior: An introduction to community-based social marketing*. New society publishers, 2011.
- [8] T. Guo, X. Cao, and G. Cong, "Efficient algorithms for answering the m -closest keywords query," in *SIGMOD*. ACM, 2015, pp. 405–418.
- [9] Y. Wu, R. Jin, X. Zhu, and X. Zhang, "Finding dense and connected subgraphs in dual networks," in *ICDE*. IEEE, 2015, pp. 915–926.
- [10] J. Elzinga and D. W. Hearn, "Geometrical solutions for some minimax location problems," *Transportation Science*, vol. 6, no. 4, pp. 379–394, 1972.
- [11] W. Khaouid, M. Barsky, V. Srinivasan, and A. Thomo, "K-core decomposition of large networks on a single pc," *PVLDB*, vol. 9, no. 1, pp. 13–23, 2015.
- [12] R. Hartshorne, *Geometry: Euclid and beyond*. Springer Science & Business Media, 2013.
- [13] G. Watson, "Computing helmert transformations," *JCAM*, vol. 197, no. 2, pp. 387–394, 2006.
- [14] M. Sarwat, J. J. Levandoski, A. Eldawy, and M. F. Mokbel, "Lars*: An efficient and scalable location-aware recommender system," *TKDE*, vol. 26, no. 6, pp. 1384–1399, 2014.
- [15] D. Deng, A. Kim, S. Madden, and M. Stonebraker, "Silkmoth: An efficient method for finding related sets with maximum matching constraints," *arXiv*, 2017.
- [16] Z. Galil, "Efficient algorithms for finding maximum matching in graphs," *CSUR*, vol. 18, no. 1, pp. 23–38, 1986.
- [17] M. Kitsak, L. K. Gallos, S. Havlin, F. Liljeros, L. Muchnik, H. E. Stanley, and H. A. Makse, "Identification of influential spreaders in complex networks," *arXiv*, 2010.
- [18] X. Liu, Q. He, Y. Tian, W.-C. Lee, J. McPherson, and J. Han, "Event-based social networks: linking the online and offline social worlds," in *SIGKDD*. ACM, 2012, pp. 1032–1040.
- [19] V. Batagelj and M. Zaversnik, "An $o(m)$ algorithm for cores decomposition of networks," *arXiv*, 2003.
- [20] X. Huang, L. V. Lakshmanan, J. X. Yu, and H. Cheng, "Approximate closest community search in networks," *PVLDB*, vol. 9, no. 4, pp. 276–287, 2015.
- [21] J. Cheng, L. Zhu, Y. Ke, and S. Chu, "Fast algorithms for maximal clique enumeration with limited memory," in *SIGKDD*. ACM, 2012, pp. 1240–1248.
- [22] J. Wang, J. Cheng, and A. W.-C. Fu, "Redundancy-aware maximal cliques," in *SIGKDD*. ACM, 2013, pp. 122–130.
- [23] P. Lee, L. V. Lakshmanan, and E. Milios, "Cast: A context-aware story-teller for streaming social content," in *CIKM*. ACM, 2014, pp. 789–798.
- [24] H. Cai, V. W. Zheng, F. Zhu, K. C.-C. Chang, and Z. Huang, "From community detection to community profiling," *PVLDB*, vol. 10, no. 7, pp. 817–828, 2017.
- [25] M. Sozio and A. Gionis, "The community-search problem and how to plan a successful cocktail party," in *SIGKDD*. ACM, 2010, pp. 939–948.
- [26] W. Cui, Y. Xiao, H. Wang, and W. Wang, "Local search of communities in large graphs," in *SIGMOD*. ACM, 2014, pp. 991–1002.
- [27] R.-H. Li, L. Qin, J. X. Yu, and R. Mao, "Influential community search in large networks," *PVLDB*, vol. 8, no. 5, pp. 509–520, 2015.
- [28] X. Huang and L. V. Lakshmanan, "Attribute-driven community search," *PVLDB*, vol. 10, no. 9, pp. 949–960, 2017.
- [29] Y. Fang, R. Cheng, S. Luo, and J. Hu, "Effective community search for large attributed graphs," *PVLDB*, vol. 9, no. 12, pp. 1233–1244, 2016.
- [30] Q. Qu, S. Liu, B. Yang, and C. S. Jensen, "Efficient top- k spatial locality search for co-located spatial web objects," in *Mobile Data Management (MDM), 2014 IEEE 15th International Conference on*, vol. 1. IEEE, 2014, pp. 269–278.
- [31] L. R. Derczynski, B. Yang, and C. S. Jensen, "Towards context-aware search and analysis on social media data," in *Proceedings of the 16th international conference on extending database technology*. ACM, 2013, pp. 137–142.
- [32] Y. Chen, J. Xu, and M. Xu, "Finding community structure in spatially constrained complex networks," *IJGIS*, vol. 29, no. 6, pp. 889–911, 2015.