

A Framework of Traveling Companion Discovery on Trajectory Data Streams

Lu-An Tang^{1,2}, Yu Zheng², Jing Yuan², Jiawei Han¹

Alice Leung³, Wen-Chih Peng⁴, Thomas La Porta⁵, Lance Kaplan⁶

¹Univ. of Illinois at Urbana-Champaign; ²Microsoft Research Asia; ³BBN Technologies

⁴National Chiao Tung University; ⁵Penn State Univ.; ⁶U.S. Army Research Laboratory

The advance of mobile technologies leads to huge volumes of spatio-temporal data collected in the form of trajectory data stream. In this study, we investigate the problem of discovering object groups that travel together (*i.e.*, *traveling companions*) from trajectory data streams. Such technique has broad applications in the areas of scientific study, transportation management and military surveillance. To discover traveling companions, the monitoring system should cluster the objects of each snapshot and intersect the clustering results to retrieve moving-together objects. Since both clustering and intersection steps involve high computational overhead, the key issue of companion discovery is to improve the efficiency of algorithms. We propose the models of closed companion candidates and smart intersection to accelerate data processing. A data structure termed *traveling buddy* is designed to facilitate scalable and flexible companion discovery from trajectory streams. The traveling buddies are micro-groups of objects that are tightly bound together. By only storing the object relationships rather than their spatial coordinates, the buddies can be dynamically maintained along trajectory stream with low cost. Based on traveling buddies, the system can discover companions without accessing the object details. In addition, we extend the proposed framework to discover companions on more complicated scenarios with spatial and temporal constraints, such as on the road network and battlefield. The proposed methods are evaluated with extensive experiments on both real and synthetic datasets. Experimental results show that our proposed buddy-based approach is an order of magnitude faster than the baselines and achieves higher accuracy in companion discovery.

1. INTRODUCTION

The technical advances in mobile devices and tracking technologies have generated huge amount of trajectory data recording the movement of people, vehicle, animal and natural phenomena in a variety of applications, such as social network, transportation management, scientific studies and military surveillance [Zheng and Zhou 2011]: (1) In Foursquare¹, the users check in the sequence of visited restaurants and shopping malls as trajectories. In many GPS-trajectory-sharing websites like Geolife [Zheng et al. 2010], people upload their travel or sports routes to share with friends. (2) Many taxis in major cities have been embedded with GPS sensors. Their locations are reported to the transportation system in the format of streaming trajectories [Yuan et al. 2010; Tang et al. 2011]. (3) Biologists solicit the moving trajectories of animals like migratory birds for their research². (4) The battlefield sensor network watches the designated area and collects the movement of possible intruders [Tang et al. 2010]. Their trajectories are watched by military satellites all the time.

In the above-mentioned applications, people usually expect to discover the object groups that move together, *i.e.*, *traveling companions*. For example, commuters want to discover people with the same route to share car pools. Scientists would like to study the pathways of species migration. Information about traveling companions can also be used for resource allocation, security management, infectious disease control and so on.

Despite of the wide applications, the discovery of traveling companion is not efficiently supported in existing systems, partly due to the following challenges:

¹<http://foursquare.com>

²<http://www.movebank.org>

- **Co-location:** Companions are objects that travel together. Here “travel together” means the objects are spatially close at the same time. Many state-of-the-art trajectory clustering methods, retrieving the object’s major moving direction from their trajectories, ignore the temporal information of objects [Lee et al. 2007; Har-Peled 2003; Li et al. 2004; Yang et al. 2009; Zhang and Lin 2004; Jensen et al. 2007]. Hence they cannot be directly used for companion discovery.
- **Incremental discovery:** In several applications like military surveillance, the system needs to monitor objects for a long time and discover companions as soon as possible. Hence the algorithm should report the companions in an incremental manner, *i.e.*, output the results simultaneously while receiving and processing the trajectory data stream.
- **Efficiency:** Most trajectories are generated in a format of data stream. Huge amounts of data arrive rapidly in a short period of time. The monitoring system has to cluster the data and intersect the clusters for companions. These steps involve high computational overhead. The algorithm should develop efficient data structures to process large scale data.
- **Effectiveness:** The number of companions is usually large. The system should report the large and long-lasting companions rather than small and short-time ones. The companion-discovery algorithm should be effective to select the most important results.
- **Spatio-temporal constraints:** In the real applications, the objects move with several spatial and temporal constraints, *e.g.*, the vehicles travel along the road network, the military objects need to follow certain orders to leave the team for short time. The algorithm should be adapted for such constraints to improve the system feasibility and applicability.

We are aware that several studies have retrieved object groups similar to the traveling companions, such as flock [Gudmundsson and Kreveld 2006], convoy [Jeung et al. 2008] and swarm [Li et al. 2010]. However, most of them are designed to work on static datasets on 2D Euclidean space, some methods need multiple scans of the data, or cannot output results in an incremental manner. Hence it is still desirable to provide high-quality but less costly techniques for companion discovery on trajectory stream with spatio-temporal constraints.

In this study, we investigate the models, principles and methodologies to discover traveling companions from trajectory streams. Since the objects keep on moving in the trajectory streams, it is hard to maintain an index for their spatial positions. However, the relationships among most objects are gradual evolutions rather than fierce mutations. The *traveling buddy* is proposed to store the relationship. Such model can be easily maintained along the data streams. Thus, in this paper, we explore the traveling-buddy-based companion discovery, which is able to discover companions without accessing the object details and significantly improve the system’s efficiency. The main contributions of this study include: (1) introducing the companion models to define the problem; (2) proposing the concepts of smart intersection and closed companions to accelerate data processing; (3) analyzing the bottleneck of the problem and proposing a traveling-buddy-based approach; (4) extending the proposed methods to complicated scenarios with spatio-temporal constraints, developing the methods to discover the road companions and loose companions; and (5) demonstrating the scalability and feasibility of the proposed methods by experiments on both real and synthetic datasets.

This paper substantially extends the version on ICDE 2012 conference [Tang et al. 2012], in the following ways: (1) introducing the concepts of *road companion* and *loose companion* to model the companion discovery problems on more compli-

cated scenarios; (2) analyzing the main bottleneck of road companion discovery and proposing a filtering-and-refinement-based framework; (3) designing new algorithms with the *road buddy* for efficient companion discovery on road network; (4) proposing the leaving time threshold and introducing the concept of *loose companion* to release the time constraints for more effective companion discovery; (5) carrying out the time complexity analysis for proposed algorithms; (6) providing complete formal proofs for lemmas and propositions; (7) covering the related work in more details and including recent ones; and (8) expanding our performance studies on datasets on road network and battlefield. The experimental results show that the new proposed methods are an order of magnitude faster than the old ones in [Tang et al. 2012].

The rest of the paper is organized as follows. Section 2 defines the problem; Section 3 introduces the general framework of companion discovery; Section 4 proposes the traveling-buddy-based method; Section 5 extends the proposed methods to discover companions on road networks; Section 6 discusses the techniques to discover companions with released temporal constraints; Section 7 evaluates the algorithms' performances; Section 8 gives a survey of the related work and finally in Section 9 we conclude the paper.

2. PROBLEM DEFINITION

In the various applications of traveling companion, there are some common principles shared in different scenarios. We illustrate the characteristics of companion discovery by the following example.

Example 1: Ten objects are tracked by a monitoring system. Fig.1 shows their positions in four snapshots. There are three key issues to discover the companions:

- **Cluster:** The companions are the objects that travel together, *i.e.*, in the same cluster. Since the people, vehicles and animals often move and organize in arbitrary ways, the companion shape is not fixed. In Fig.1, the objects are grouped in round shape in snapshots s_1 and s_2 , while in s_3 , they are moving in a queue and the companions are formed as thin and long ellipses.
- **Consistency:** The companions should be consistent enough to last for a few snapshots. This feature makes it possible to find the companions by intersecting the clusters of different snapshots.
- **Size:** Most users are only interested in the object groups that are big enough. They may have requirements on the companion's size. For example, if the user sets the size threshold as four and requires the companion to last for at least four snapshots, then $\{o_1, o_2, o_3, o_4\}$ is the result companion.

To discover the traveling companions with various shapes, we employ the concepts of density-based clustering [Ester et al. 1996] in this study.

Definition 1 (Direct Density Connection): Let O be the object set in a snapshot, ε be the distance threshold, μ be the density threshold and $N_\varepsilon(o_i) = \{o_j \in O \mid dist(o_i, o_j) \leq \varepsilon\}$. Object o_j is directly density connected from object o_i if $o_j \in N_\varepsilon(o_i)$ and $|N_\varepsilon(o_i)| \geq \mu$.

Definition 2 (Density Connection): Let O be the object set in a snapshot, object o_i is density connected to object o_j , if there is a chain of objects $\{o_1, \dots, o_n\} \in O$ where $o_1 = o_j$, $o_n = o_i$ such that o_{i+1} is directly density connected from o_i .

With the concepts of density connection, we formally define the traveling companion as follows.

Definition 3 (Traveling Companion): Let δ_s be the size threshold and δ_t be the duration threshold, a group of objects q is called *traveling companion* if:

- (1) The members of q are density connected by themselves for a period t where $t \geq \delta_t$;
- (2) q 's size $size(q) \geq \delta_s$.

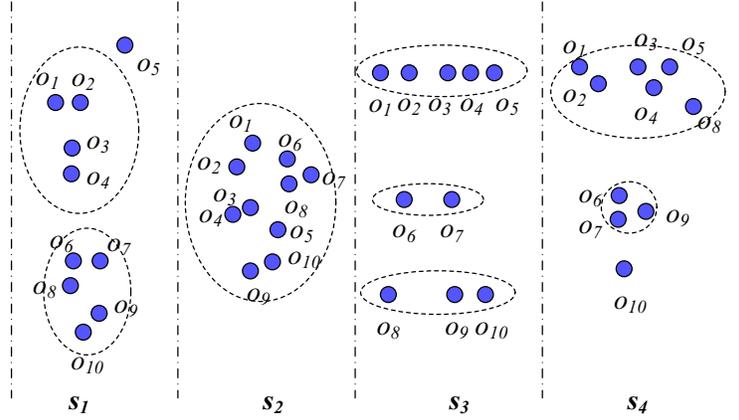


Fig. 1. Example: Discover Traveling Companions

Problem Definition: Let trajectory data stream S be denoted by a sequence of snapshots $\{s_1, s_2, \dots, s_i, \dots\}$. Each snapshot $s_i = \{(o_1, x_{1,i}, y_{1,i}), (o_2, x_{2,i}, y_{2,i}), \dots, (o_n, x_{n,i}, y_{n,i})\}$, where $x_{j,i}, y_{j,i}$ are the spatial coordinates of object o_j at snapshot s_i . When the data of snapshot s_i arrives, the task is to discover companion set Q that contains all the traveling companions so far.

We will introduce the framework and techniques for companion discovery in the next few sections. Fig. 2 lists the notations used throughout this paper.

Notation	Explanation	Notation	Explanation
S	the trajectory stream	s, s_b, s_j	the snapshots in stream
C	the cluster set	c_b, c_j	the clusters
Q	the companion set	q	the traveling companion
R	the candidate set	r_b, r_j	the companion candidates
B	the buddy set	b_b, b_j	the traveling buddies
O	the object set	o_1, o_2, o_i	the objects
ε	the distance threshold	μ	the density threshold
δ_s	the size threshold	δ_t	the duration threshold
δ_y	the buddy radius threshold	γ_b, γ_j	the buddy radius
M	the road network	δ_l	the leaving time threshold

Fig. 2. List of Notations

3. COMPANION DISCOVERY FRAMEWORK

3.1. The Clustering-and-Intersection Method

A general framework of *clustering-and-intersection* is proposed in [Gudmundsson and Kreveld 2006; Jeung et al. 2008] to retrieve the convoy patterns. This framework can also be adapted to discover companions on trajectory stream: The idea is to retrieve *companion candidates* by counting common objects in the clusters from different snapshots. The system keeps clustering the objects in coming snapshots and intersecting them with the stored candidates. In this way the candidates are gradually refined to become resulting companions.

Definition 4 (Companion Candidate): Let δ_s be the size threshold and δ_t be the duration threshold, a group of objects r is a companion candidate if:

- (1) The members of r are density connected by themselves for a period t where $t < \delta_t$;
- (2) $size(r) \geq \delta_s$.

Intuitively, the companion candidates are the object groups with enough size but shorter duration. The candidate's size reduces when intersecting with the clusters from other snapshots, but its lasting time increases. Once a candidate's time grows longer than threshold, it will be reported as a traveling companion. Meanwhile, as soon as the candidate is not large enough, it is no longer qualified and should be removed from memory. Fig. 3 lists the steps of clustering-and-intersection algorithm.

Algorithm 1. Clustering-and-Intersection

Input: size threshold δ_s , duration threshold δ_t , distance threshold ϵ , density threshold μ , candidate set R and the trajectory data stream S

Output: every qualified companion q

1. **for** each coming snapshot s of S
 2. initialize new candidate set R' ;
 3. cluster the objects in s w.r.t to ϵ and μ ;
 4. **for** each candidate $r_i \in R$, **do**
 5. **for** each cluster $c_j \in s$, **do**
 6. new candidate $r_i' \leftarrow r_i \cap c_j$;
 7. $duration(r_i') = duration(r_i) + duration(s)$;
 8. **if** $size(r_i') \geq \delta_s$ **then**
 9. add r_i' to R' ;
 10. **if** $duration(r_i') \geq \delta_t$ **then**
 11. **output** r_i' as a qualified companion q ;
 12. add all the new clusters to R' ;
 13. $R \leftarrow R'$;
-

Fig. 3. Algorithm: The Clustering-and-intersection Method

Algorithm 1 first performs density-based clustering for all the objects in coming snapshot (Lines 1 – 3). Then the system refines companion candidates by intersecting them with new clusters (Lines 4 – 7). The intersection results with enough size are stored as new candidates (Lines 8 – 9). The ones with enough duration are reported as traveling companion (Lines 10 – 11). The new clusters are added to the candidate set (Line 12). At last the candidate set R is updated to process following snapshots (Line 13).

Proposition 1: Let n_1 be the size of objects and n_2 be the total size of candidate set R . The time complexity of Algorithm 1 is $O(n_1^2 + n_1 * n_2)$.

Proof: In the clustering step, the algorithm needs $O(n_1^2)$ time to generate density-based clusters³. In the intersection step, suppose there are average m_1 clusters and m_2 candidates, the system carries out $m_1 * m_2$ intersections, and the intersection takes $l_1 * l_2$ time, where l_1 is the average cluster size and l_2 is the average candidate size. Since $m_1 * l_1 = n_1$, $m_2 * l_2 = n_2$, thus the time complexity of intersection step is $O(m_1 * m_2 * l_1 * l_2) = O(n_1 * n_2)$ and the total time complexity is $O(n_1^2 + n_1 * n_2)$. ■

³The clustering process can be improved to $O(n_1 * \log n_1)$ with a spatial index, however it is costly to maintain such spatial index in each time snapshot [Lee et al. 2003].

Example 2: Fig. 4 shows the running process of clustering-and-intersection algorithm. Suppose each snapshot lasts for 10 minutes, the size threshold is 3 and the time threshold is 40 minutes. The objects are first clustered in each snapshot. Two clusters in s_1 are taken as the candidates, namely r_1 and r_2 . Then they are intersected with the clusters in s_2 , meanwhile, the cluster of s_2 is also added as a new candidate r_3 . The clustering and intersection steps are carried out in each snapshot. Finally, the algorithm reports $\{o_1, o_2, o_3, o_4\}$ as a traveling companion in s_4 . The total intersection times are 29, and the largest candidate set R appears in s_3 with 23 objects involved.

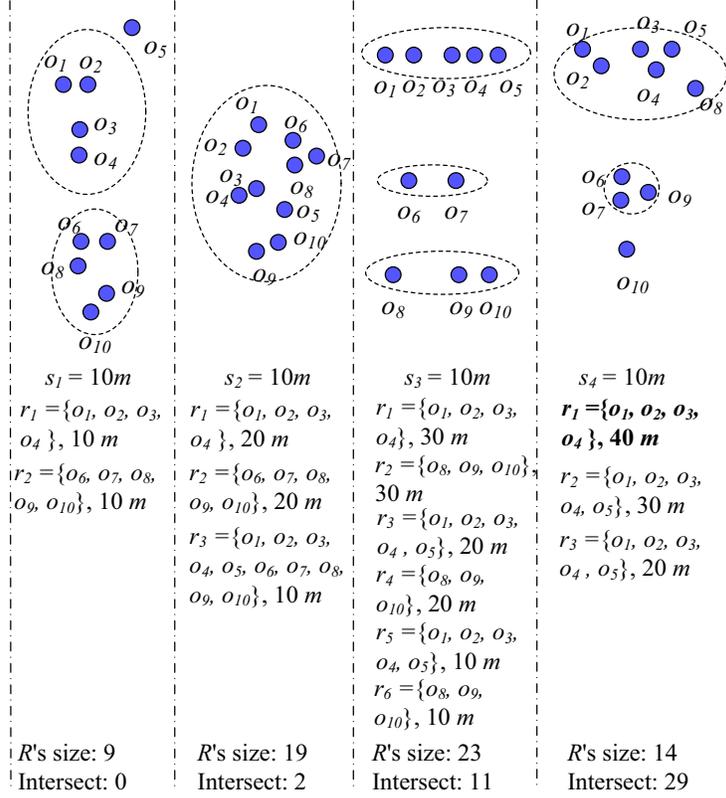


Fig. 4. Example: The Clustering-and-intersection Method

3.2. The Smart-and-Closed Algorithm

The computational overhead of clustering-and-intersection method is high in both time and space. In each snapshot, the intersection is carried out in every pair of candidate and cluster. However, most intersections cannot generate qualified results with enough size. In this subsection we introduce the methods to improve the efficiency: (1) the smart algorithm stops the intersection step early once it is impossible to generate qualified candidates, and (2) the closed candidate are used to help reduce the memory cost.

Lemma 1: Let r be a companion candidate and δ_s be the size threshold, if there are more than $size(r) - \delta_s$ objects of r already appearing in intersected clusters, continu-

ously intersecting r with remaining clusters will not generate any meaningful results with size larger than δ_s .

Proof: Since each object only appears once in a single snapshot and only belongs to one cluster⁴, if there are more than $size(r) - \delta_s$ objects appearing in already intersected clusters, even in the best case (all the remaining objects are in a single cluster), the intersection result will still be smaller than $size(r) - (size(r) - \delta_s) = \delta_s$. ■

Lemma 1 can be used to improve the candidate refining process with smart intersection. Once an object is found in the cluster, the algorithm removes it from the candidate. The intersection process will stop earlier if there are less than δ_s objects remaining in the candidate.

Another problem of clustering-and-intersection method is the space efficiency, if all new clusters are added as candidates, the size of the candidate set will increase rapidly as trajectory stream passes-by, such a huge candidate set is a burden for system memory. In the worst case, all the clusters stay constant in the series of snapshots, the intersection process cannot prune any existing candidates and all the new clusters are added to the candidate set. After m snapshots, the system needs to maintain a $m * n$ size candidate set, where n is the number of objects.

In Fig. 4, candidates r_3 and r_5 in s_3 contain the same objects with different lasting time. In such cases, the system only needs to store the one with longer time (e.g., r_3). Such candidates like r_3 are called *closed candidates*.

Definition 5 (Closed Candidate): For a companion candidate r_i , if there does not exist another candidate r_j such that $r_i \subseteq r_j$, and r_i 's duration is less than r_j 's duration, then r_i is a *closed candidate*.

Armed with Lemma 1 and Definition 5, we propose the smart-and-closed algorithm. The modifications are underlined in Fig. 5, the algorithm removes intersected objects from the candidate set and checks its remaining size before next intersection (Lines 5 and 9); when adding the new clusters to the candidate set, the algorithm always checks if there is already a candidate containing the same objects but with longer duration, only the ones passing the closeness check are added as new candidates (Lines 14 – 15).

In the worst case, Algorithm 2 cannot prune any candidates and the time complexity is the same as Algorithm 1. However, we find out that the smart-and-closed algorithm can save about 50% time and space in the experiments.

Example 3: Fig. 6 shows the running process of smart-and-closed algorithm. In snapshot s_3 , when making intersections for candidate r_1 with three clusters, the process ends early after the first round. Since the system only stores closed candidates, the largest candidate set size is only 19 in s_2 , and the total intersection time is 12, less than half of the cost in clustering-and-intersection.

4. TRAVELING BUDDY BASED DISCOVERY

Smart-and-closed algorithm improves the efficiency of intersection step to generate companions, but the system still has to cluster the objects in each snapshot. The density-based clustering costs $O(n^2)$ time without spatial index, where n is the number of the objects [Han and Kamber 2006]. Due to the dynamic nature of streaming trajectories (i.e., the objects' positions are always changing), maintaining traditional spatial indexes (such as R-tree or quad-tree) at each time snapshot incurs high cost [Lee et al. 2003]. In this section, we introduce a new structure, called *traveling buddy*, to maintain the relationship among objects and help discover companions.

⁴The clustering methods used in this study are all “hard-clustering”, i.e., an object can only belong to one cluster.

Algorithm 2. Smart-and-Closed Algorithm

Input: size threshold δ_s , duration threshold δ_t , distance threshold ε , density threshold μ , candidate set R and the trajectory data stream S

Output: every qualified companion q

1. **for** each coming snapshot s of S
2. initialize new candidate set R' ;
3. cluster the objects in s w.r.t to ε and μ ;
4. **for** each candidate $r_i \in R$, **do**
5. **for** each cluster $c_j \in s$, **do**
6. **if** r_i 's size is less than δ_s , **then break**;
7. new candidate $r_i' \leftarrow r_i \cap c_j$;
8. $duration(r_i') = duration(r_i) + duration(s)$;
9. **remove intersected objects from** r_i ;
10. **if** $size(r_i') \geq \delta_s$ **then**
11. add r_i' to R' ;
12. **if** $duration(r_i') \geq \delta_t$ **then**
13. output r_i' as a qualified companion q ;
14. **for** each cluster c_j **do**
15. **if** c_j is closed **then add to** R' ;
16. $R \leftarrow R'$;

Fig. 5. Algorithm: The Smart-and-closed Discovery

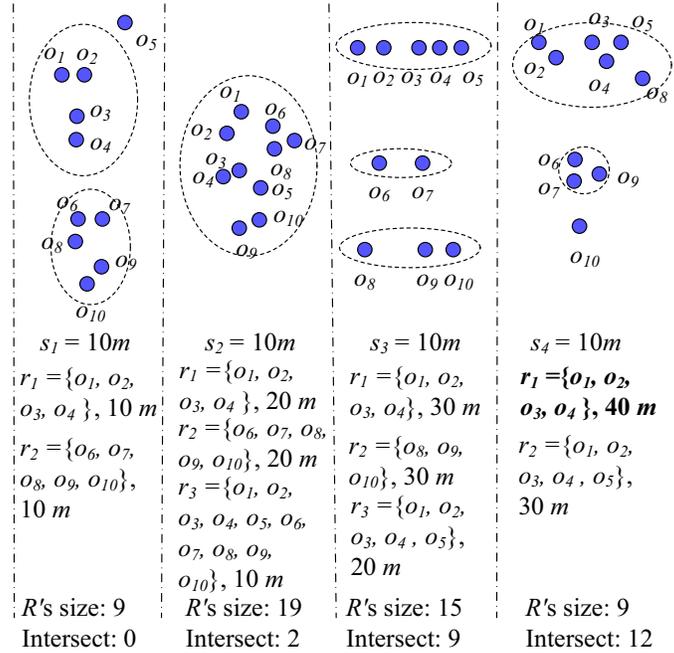


Fig. 6. Example: Smart-and-closed Algorithm

4.1. The Traveling Buddy

In streaming trajectories, the objects keep on moving and updating their positions, however, the changes of object relationships are gradual evolutions rather than fierce mutations. The object relationships are possible to be retained in a few snapshots, *i.e.*, the objects are likely to stay together with several members of the current cluster. It is attractive to reuse such information to speed up the clustering tasks. However, the system cannot reuse it directly. The major issue is about the intrinsic feature of density-based clustering. Unlike other types of clusters, the results of density-based clustering may be quite different due to minor position change of an individual object. This phenomenon is called *individual sensitivity* as illustrated in Example 4.

Example 4: Fig. 7 shows two consecutive snapshots of the trajectory stream. Suppose the density threshold μ is set to three. In snapshot s_1 , two clusters c_1 and c_2 are independent. However in s_2 , object o_1 moves a little to the south, and this movement makes the two clusters density connected and merged as one cluster c_3 . Such case may impose important meanings in real applications, for instance, in the scenario of infected disease monitoring, the people in the two clusters should then be watched together since the disease may spread among them.

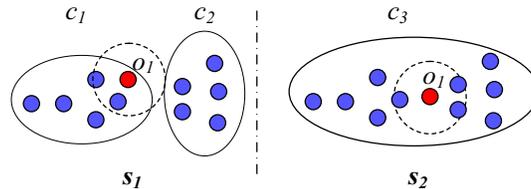


Fig. 7. Example: Individual Sensitivity Problem

The time cost of checking individual sensitivity is quadratic to the cluster size, and in many cases the system has to generate large clusters to produce meaningful companions. Hence high computational overhead is still involved in the clustering stage.

Then is it possible to explore a smaller and more flexible structure? In real world, there are some kinds of micro-groups in trajectory stream. For examples, couples would like to stay together on trips, military units operate in teams, families of birds, deer and other animals often move together in species migration. Such objects stay closer to each other than outside members. Even though they might not be as big as the companion, their information can be used to help clustering. Since they are way smaller than the cluster, their maintenance cost is much lower.

Definition 6 (Traveling Buddy): Let O be the object set and δ_γ be the buddy radius threshold, traveling buddy b is defined as a set of objects satisfying: (1) $b \subseteq O$; (2) for $\forall o_i \in b$, $dist(o_i, cen(b)) \leq \delta_\gamma$, where $cen(b)$ is the geometry center of b . The buddy's radius γ is defined as the distance from $cen(b)$ to b 's farthest member.

The traveling buddies can be initialized by incrementally merging the objects in two steps: (1) treating all objects as individual buddies; and (2) merging them with their nearest neighbors. This process stops if the buddy's radius is larger than γ . The initialization step costs $O(n^2)$ time for n objects. However, this step only needs to be carried out once and the traveling buddies are dynamically maintained along the stream.

There are two kinds of operations to maintain buddies on the data stream: namely *split* and *merge*, as shown in the following example.

Example 5: Fig. 8 shows the traveling buddies in two snapshots. Traveling buddy b_1 is split into three parts in snapshot s_2 . At the same time, b_2 , b_3 and a part of b_1 are merged as a new buddy in s_2 .

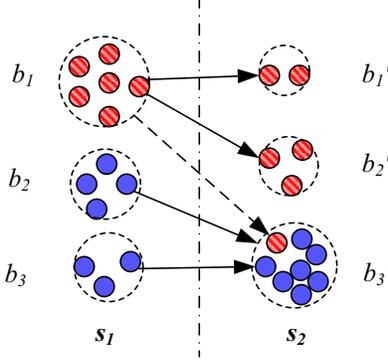


Fig. 8. Example: Merge and Split Buddies

When the data of a new snapshot s_{t+1} arrive, the maintenance algorithm first updates the center of each buddy b . For object $o_i \in b$, the system calculates the shift $(\Delta x_i, \Delta y_i)$ between s_{t+1} and s_t . And the new center is updated as:

$$cen_{t+1}(b) = cen_t(b) + \sum_{o_i \in b} (\Delta x_i, \Delta y_i)$$

Then every object $o_i \in b$ checks its distance to the buddy center; if the distance is larger than δ_γ , o_i will be split out as a new buddy. The $cen(b)$ is also updated by subtracting the shift of o_i .

The second operation is to merge the buddies that are close to each other. If two buddies b_i and b_j satisfy the following equation, they should be merged as a new buddy.

$$dist(cen(b_i), cen(b_j)) + \gamma_i + \gamma_j \leq 2\delta_\gamma$$

Suppose b_i has m_i objects and b_j has m_j objects, the new buddy b_k 's center is computed as $cen(b_k) = (m_i * cen(b_i) + m_j * cen(b_j)) / (m_i + m_j)$. Therefore, the system does not need to access the detailed coordinates of each object to merge buddies, the computation can be done with the information from the old buddy's center and size.

The detailed steps of buddy maintenance are shown in Fig. 9. When the data of a new snapshot arrives; the algorithm first updates the center of each buddy (Line 2). Then each buddy member is checked to see whether a split operation is needed (Lines 3 – 7). At last, the system scans the buddy set and merges the buddies that are close to each other (Lines 10 – 13).

Proposition 2: Let m be the average number of traveling buddies and n be the number of objects. The time cost of Algorithm 3 is $O(n + m^2)$.

Proof: The split operation needs to check each object and the time cost is $O(n)$. The merge operation has to check the buddies pairs with time complexity $O(m^2)$. Therefore the total maintenance cost is $O(n + m^2)$. ■

In the worst case, if the objects are sparse and each of them is an individual buddy, where $m = n$. The maintenance cost is still $O(n^2)$. However the number of m is usually much smaller than n and the algorithm is likely to strike a relatively high efficiency.

Algorithm 3. Traveling Buddy Maintenance

Input: the radius threshold δ_γ , the traveling buddy set B and the coming snapshot s

Output: updated buddy set B'

```

1.  for each  $b_i$  in  $B$  do
2.      update  $\text{cen}(b_i)$ ;
3.      for  $o_j$  in  $b_i$ , do
4.          if  $\text{dist}(o_j, \text{cen}(b_i)) > \delta_\gamma$  then // Split Operation
5.              split  $o_j$  out as a new buddy  $b_j$ ;
6.              add  $b_j$  to  $B'$ ;
7.              update  $\text{cen}(b_i)$ ;
8.          add  $b_i$  to  $B'$ ;
9.      //Merge Operation
10.     for each  $b_i, b_j$  in  $B'$ ,  $b_i \neq b_j$  do
11.         if  $\text{dist}(\text{cen}(b_i), \text{cen}(b_j)) + \gamma_i + \gamma_j \leq 2\delta_\gamma$  then
12.             merge  $b_i, b_j$  as  $b_k$ ;
13.             remove  $b_i, b_j$  and add  $b_k$  to  $B'$ ;
14.     return  $B'$ ;

```

Fig. 9. Algorithm: Buddy Maintenance

4.2. Buddy-based Clustering

In the clustering step, the system has to check the density connectivity for each object. The traveling buddies can help the clustering process avoid accessing those object details. To bring down computational overhead, we introduce following lemmas.

Lemma 2: Let b be a traveling buddy, ε be the distance threshold and μ be the density threshold. If b 's size is larger than $\mu + 1$ and the buddy radius $\gamma \leq \varepsilon/2$, then all the objects in b are directly density reachable to each other. Such a traveling buddy is called a *density-connected buddy*.

Proof: Note that $\gamma \leq \varepsilon/2$, thus for $\forall o_i, o_j \in b$, $\text{dist}(o_i, o_j) \leq 2\gamma \leq \varepsilon$. Then all the members of b are included in $N_\varepsilon(o_i)$. If b 's size is larger than $\mu + 1$, then $|N_\varepsilon(o_i)| \geq \mu$. By Definition 1, o_i and o_j are directly density reachable. ■

Lemma 2 shows that, if a traveling buddy is tight and large by itself, then all its members can be considered as density connected. Lemma 2 also gives the directions that the radius threshold δ_γ should not be set larger than $\varepsilon/2$.

Lemma 3: Let b_i and b_j be two traveling buddies with radius γ_i and γ_j , and ε be the distance threshold. If $\text{dist}(\text{cen}(b_i), \text{cen}(b_j)) - \gamma_i - \gamma_j > \varepsilon$, then the objects in b_i and b_j are not directly density reachable.

Proof: As shown in Fig. 10(a):

if $\text{dist}(\text{cen}(b_i), \text{cen}(b_j)) - \gamma_i - \gamma_j > \varepsilon$, then for $\forall o_i \in b_i, o_j \in b_j$, $\text{dist}(o_i, o_j) > \varepsilon$. Therefore, o_j does not belong to $N_\varepsilon(o_i)$ and they are not directly density reachable. ■

Lemma 3 tells us that, when searching for the directly density reachable objects for a traveling buddy, if another buddy is too far away, then the system can prune all its members without further computation. This lemma is very helpful. In the experiments it helps prune more than 80% of the objects.

For the traveling buddies that are close to each other, the detailed distance computation still needs to be carried out. But with the following lemmas, the system does not need to compute distances between all the pairs. Lemma 4 provides heuristics to speed up the computation.

Lemma 4: Let b_i, b_j be two density-connected buddies and ε be the distance threshold. If $\exists o_i \in b_i, o_j \in b_j$ such that $dist(o_i, o_j) \leq \varepsilon$, then all the objects of b_i and b_j are density connected.

Proof: As Fig. 10 (b) shows, since b_i is a density-connected traveling buddy and $|N_\varepsilon(o_i)| \geq \mu$, if $dist(o_i, o_j) \leq \varepsilon$, then o_i and o_j are directly density reachable. Since all the objects in b_i and b_j are directly density reachable from o_i and o_j , respectively. Therefore, all the objects in the two traveling buddies are density connected. ■

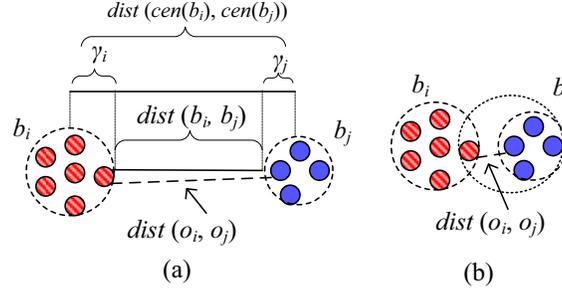


Fig. 10. Proof of Lemma 3 and 4

Based on Lemma 4, once the system finds a pair of objects close to each other, it ends the computation and considers the corresponding buddies density-connected. The detailed algorithm is listed in Fig. 11. The algorithm first updates the buddy set in a new snapshot (Line 1). Then it randomly picks a buddy and checks the density connectivity to others (Lines 2 – 4). The far-away buddies are filtered out (Lines 5 – 6). With the help of Lemma 4, the algorithm searches density reachable buddies and objects and adds them to the cluster (Lines 7 – 13). Finally, the algorithm outputs clustering results when all the buddies are processed (Line 15).

In the worst case, Algorithm 4 is still with $O(n^2)$ time complexity, where n is the number of objects. But in most cases, Lemmas 3 and 4 can prune majority buddies and save time for distance computation. The experiment results show that buddy-based clustering is an order of magnitude faster than the original clustering algorithm.

4.3. Companion Discovery with Buddies

The buddies are not only useful in clustering step, they are also helpful for the intersection process to generate companions. When intersecting a candidate with a cluster, the system needs to check whether each candidate's objects appear in the cluster or not. The information of traveling buddies can provide a shortcut to this process: If a buddy stays unchanged during the period, and it appears both in the candidate and the cluster, then the system can put all its members into the intersection result without accessing the detailed objects.

To efficiently utilize the buddy information, a buddy index is designed to keep the candidates dynamically updated with the buddies.

Definition 7 (Buddy Index): The buddy index is a triple $\{BID, ObjSet, CanIDs\}$, where BID is the buddy's ID, $ObjSet$ is the object members of the buddy, $CanIDs$ records the IDs of candidates containing the buddy.

As long as the buddy stays unchanged, the candidates only store the BID instead of detailed objects. While making intersections, the buddy is treated as a single object. When the buddy changes, the system updates all the candidates in $CanIDs$ and

Algorithm 4. Buddy-based Clustering

Input: the distance threshold ε , the density threshold μ , the coming snapshot s and the buddy set B .

Output: the cluster set C

```

1.  update buddy set  $B$ ; //Algorithm 3
2.  randomly pick a buddy  $b_i$  in  $B$ ;
3.  initialize cluster  $c_i \leftarrow b_i$ , add  $c_i$  to  $C$ ;
4.  for  $b_j$  in  $B$ ,  $b_j \neq b_i$ , do
5.      if  $\text{dist}(\text{cen}(b_i), \text{cen}(b_j)) - \gamma_i - \gamma_j > \varepsilon$ , then
6.          continue; // Lemma 3
7.      for each  $o_i$  in  $b_i$ ,  $o_j$  in  $b_j$ , do
8.          if  $\text{dist}(o_i, o_j) \leq \varepsilon$ , then
9.              if  $b_i, b_j$  are density connected then
10.                 add  $b_j$  to  $c_i$ ; //Lemma 4
11.                 break;
12.             else if  $o_j$  is density connected from  $o_i$  then
13.                 add  $o_j$  to  $c_i$ ;
14.  repeat steps 2 - 13 until all buddies are processed;
15. return the cluster set  $C$ ;

```

Fig. 11. Algorithm: Buddy-based Clustering

replaces BID with the corresponding objects in $ObjSet$. The buddy-based companion discovery algorithm is listed in Fig. 12.

Algorithm 5. Buddy-based Companion Discovery

Input: Size threshold δ_s , duration threshold δ_t , candidate set R , buddy index BI and the trajectory data stream S

Output: every qualified companion q

```

1. for each coming snapshot  $s$  of  $S$ ;
2.  initialize new candidate set  $R'$ ;
3.  buddy based clustering; // Algorithm 4
4.  update  $BI$  and corresponding candidates;
5.  for each candidate  $r_i$  in  $R$ , do
6.      if  $\text{size}(r_i) < \delta_s$  then break;
7.      for each cluster  $c_j$  in  $s$ , do
8.           $r_i' \leftarrow \text{buddy-based-intersection}(r_i, c_j)$ ;
9.           $\text{duration}(r_i') = \text{duration}(r_i) + \text{duration}(s)$ ;
10.         remove intersected objects and buddies from  $r_i$ ;
11.         if  $\text{size}(r_i') \geq \delta_s$  then
12.             add  $r_i'$  to  $R'$ ;
13.             if  $\text{duration}(r_i') \geq \delta_t$  then
14.                 output  $r_i'$  as a qualified companion  $q$ ;
15.         for each cluster  $c_j$  do
16.             if  $c_j$  is closed then add to  $R'$ ;
17.   $R \leftarrow R'$ ;

```

Fig. 12. Algorithm: Buddy-based Companion Discovery

When a new snapshot arrives, the algorithm performs buddy-based clustering and updates the buddy index (Lines 2 – 4), then selects out the candidates with enough size (Lines 5 – 6). The candidates are interested with the generated clusters with the help of the buddy index (Lines 7 – 10). The candidate's duration and size are checked again

after the intersection, and the qualified ones are output as the companions (Lines 11 – 14). Finally, the closed candidates are added to the memory for further processing (Lines 15 – 17).

Example 6: Fig. 13 shows the running process for buddy-based companion discovery. There are four buddies initialized in snapshot s_1 . In the candidates, the buddy ID is stored instead of detailed objects. In snapshot s_2 , the four buddies stay the same and the algorithm makes intersections by only checking their *BIDs*. Although the total intersection time is not reduced, the time cost for each intersection operation has been brought down. It is common that different candidates contain the same objects, such as r_1 and r_3 in s_2 . The buddy index helps to keep only one copy of the objects and add only pointers (the *BIDs*) to candidates. Therefore, the space cost is further reduced. In s_3 , the buddy b_3 is no longer valid, then the system updates candidate r_2 , using the objects to replace the buddy's ID. In s_4 , traveling companion r_1 is discovered as $\{b_1, b_2\}$. With the help of buddy index, the system can easily look up detailed objects and output the companion as $\{o_1, o_2, o_3, o_4\}$.

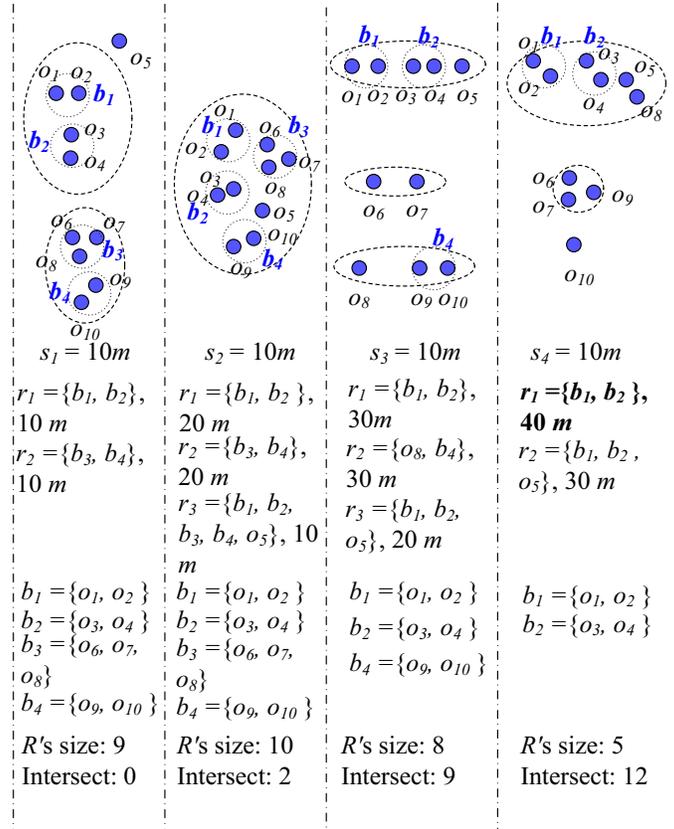


Fig. 13. Example: Buddy-based Discovery

5. ROAD COMPANION DISCOVERY

In the previous sections, we have investigated the problem of companion discovery on 2D Euclidean space. However, many objects move on the road networks in real

applications. There are several unique difficulties for companion discovery on the road network. In this section we explore the problem of discovering road companions.

5.1. Problem Formulation

Example 7: Fig. 14 shows the example of moving vehicles on the road network. There are several issues different from the companion discovery in 2D Euclidean space.

- **Distance computation:** In the road network, the distance between two objects should be the length of the shortest path connecting them, rather than a straight line between them. As shown in the figure, o_1 and o_2 are close to each other in the Euclidean space, but they are on different directions. The road network distance between them is actually much larger.
- **Density uselessness:** In most cases, the road companion move in the shape of a line. It is not necessary to require the vehicles be density-connected with others. In Fig.14, o_7, o_8 and o_9 have neighboring vehicles o_{10} and o_{11} , while o_4, o_5 and o_6 have no neighboring vehicle. Both of them should be considered as moving together. Therefore, the neighbors and density information should not influence the companion discovery.

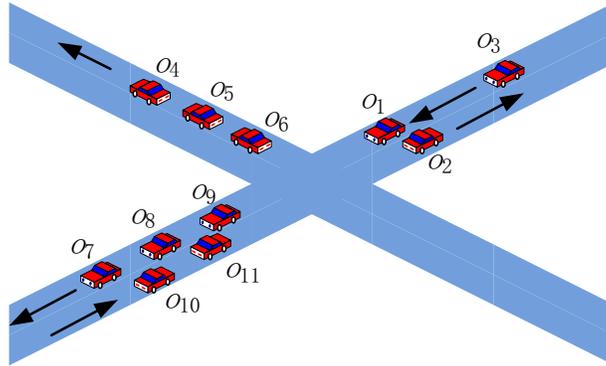


Fig. 14. Example: Traveling Companions on Road Network

Since the road companion discovery is a new type of problem, it is necessary to modify some basic concepts of the traveling companion and redefine the task with new constraints.

Definition 8 (Direct Road Connection): Let O be the object set in a snapshot, M be the road network, and ε be the distance threshold. Object o_j is *directly road connected* from object o_i on M if $netd(o_i, o_j) \leq \varepsilon$, where $netd(o_i, o_j)$ is the road network distance between o_i and o_j on M .

Note that, we remove the requirements about density and replace the Euclidean distance with the road network distance in Definition 8.

Definition 9 (Road Connection): Let O be the object set in a snapshot, M be the road network, object o_i is *road connected* to object o_j on M , if there is a chain of objects $\{o_1, \dots, o_n\} \in O$ where $o_1 = o_j$, $o_n = o_i$ such that o_{i+1} is directly road connected from o_i on M .

Based on the above definitions, we can formally define the task of road companion discovery as follows.

Definition 10 (Road Companion): Let M be the road network, δ_s be the size threshold and δ_t be the duration threshold, a group of objects q is called *road companion* if:

- (1) The members of q are road connected on M for a period t where $t \geq \delta_t$;
- (2) q 's size $size(q) \geq \delta_s$.

Problem Definition: Let trajectory data stream S be denoted by a sequence of snapshots $\{s_1, s_2, \dots, s_i, \dots\}$. Each snapshot $s_i = \{(o_1, x_{1,i}, y_{1,i}), (o_2, x_{2,i}, y_{2,i}), \dots, (o_n, x_{n,i}, y_{n,i})\}$, where $x_{j,i}, y_{j,i}$ are the spatial coordinates of object o_j at snapshot s_i , and all the objects move on a road network M . When the data of snapshot s_i arrives, the task is to discover the road companion set Q that contains all the road companions so far.

Note that, we assume the system can match the spatial coordinates of the moving objects to the road network efficiently. There are many state-of-the-art studies on this map-matching problem. In our previous studies, we have developed several methods for map-matching, the details can be found in [Yuan et al. 2010] and [Zheng et al. 2012].

5.2. The Discovery Framework

The general framework of clustering-and-intersection can be adapted to discover road companions. In each snapshot, the system first generates the road-connected clusters and intersects them with the *road companion candidates*. The candidates are gradually refined to be the road companions.

Definition 11 (Road Companion Candidate): Let M be the road network, δ_s be the size threshold and δ_t be the duration threshold, a group of objects q is called *road companion candidate* if:

- (1) The members of q are road connected on M for a period t where $t < \delta_t$;
- (2) $size(q) \geq \delta_s$.

Similarly, the ideas of smart-and-closed algorithm also works for this framework. To apply those algorithms on road network, the only difference is to replace the process of density-based clustering with the following algorithm of road-connection-based clustering.

Algorithm 6 first picks a random object as the seed to initialize an cluster (Lines 1 – 4), then expands the cluster (Lines 5 – 10). In the expansion process, the algorithm starts from the seed, adds in any objects that are directly road connected to the cluster member (Lines 7 – 10). Once a cluster is generated, the system compares its size with the threshold, only the ones with enough size are included in the final clustering results (Lines 11 – 13).

Proposition 3: Let n be the size of object set O and N be the total node number of road network M . The time complexity of Algorithm 6 is $O(n^2 * N)$.

Proof: There are three loops in Algorithm 6 (Lines 1, 5 and 7). In the worst case, no objects are road connected. Hence the algorithm has to run n times for the loops in Lines 1 and 7, and 1 time for the loop in Line 5 (each cluster only contains one object in such case). The total running number is $O(n^2)$. In each run, the system has to find the shortest path between objects o_i and o_j to compute their road network distance. The time cost of the shortest path searching step is determined to the detailed algorithm and heuristics [Pearl 1984]. In the worst case, the algorithm has to visit all the nodes of M to find out the shortest path, hence the time complexity of Algorithm 6 is $O(n^2 * N)$. ■

In many applications, the road network M contains millions of nodes, i.e., N is a quite large number. To make things worse, the system may not have enough memory to load in M in one time. Therefore the shortest path computation involves huge I/O

Algorithm 6. Road-connection-based Clustering

Input: the distance threshold ε , the size threshold δ_s , the object set O in a snapshot

Output: the road-connected cluster set C

```

1. for each unvisited object  $o$  of  $O$ , do
2.   mark  $o$  as visited;
3.   initialize a new cluster  $c$ ;
4.   add  $o$  to  $c$ ;
5.   for each unexpanded object  $o_i$  in  $c$ , do
6.     mark  $o_i$  as expanded;
7.     for each unvisited object  $o_j$  of  $O$ , do
8.       if  $\text{netd}(o_i, o_j) < \varepsilon$ , then
9.         mark  $o_j$  as visited;
10.        add  $o_j$  to  $c$ ;
11.   if  $\text{size}(c) \geq \delta_s$ , then
12.     add  $c$  to  $C$ ;
13. return  $C$ ;

```

Fig. 15. Algorithm: Road-connection-based Clustering

overhead. The time cost of Algorithm 6 is much larger than the density-based clustering, and it is not feasible for efficient road companion discovery on trajectory streams.

The bottleneck in Algorithm 6 is searching for the directly road-connected objects (Line 7 – 10). For a particular object o_i , the system has to find the shortest paths between o_i and all unvisited objects. This computation process is the most costly step of the algorithm. However, it is actually not necessary to compute all those shortest paths, the algorithm's time cost can be reduced significantly with the following lemma.

Lemma 5: In the road network M , if the Euclidean distance between two objects o_i and o_j is larger than the distance threshold ε , o_i and o_j are not directly road connected.

Proof: In the Euclidean space, the shortest path between o_i and o_j is a straight line connection them. Since the road network M is also in the same Euclidean space, the Euclidean distance must be less than or equal to the road network distance: $\text{dist}(o_i, o_j) \leq \text{netd}(o_i, o_j)$. If $\text{dist}(o_i, o_j) > \varepsilon$, then $\text{netd}(o_i, o_j) > \varepsilon$. According to Definition 8, o_i and o_j are not directly road connected. ■

Lemma 5 can help accelerate the road-connection clustering process. We develop a new clustering algorithm with the *filtering-and-refinement* strategy, as listed in Fig.16. The main difference is at Line 8. When searching for the directly road-connected objects for object o_i , the system first computes the Euclidean distance $\text{dist}(o_i, o_j)$, if $\text{dist}(o_i, o_j)$ is already larger than the threshold ε , o_j will be pruned directly.

Proposition 4: Let n be the size of object set O , N be the total node number of road network M and m be the number of objects that pass the filtering process. The time complexity of Algorithm 7 is $O(n^2 + mN)$.

Proof: With the filtering-and-refinement strategy, the algorithm only needs to compute road network distances for the m objects which pass the filtering process. Therefore the total time complexity is $O(n^2 + mN)$. ■

Note that, m is much smaller than n with a reasonable distance threshold ε . And the Euclidean distance computation does not need to access the road network M . The computation time and I/O overhead are reduced dramatically.

Algorithm 7. Clustering with Filtering-and-refinement

Input: the distance threshold ε , the size threshold δ_s , the object set O in a snapshot

Output: the road-connected cluster set C

-
1. **for** each unvisited object o of O , **do**
 2. mark o as visited;
 3. initialize a new cluster c ;
 4. add o to c ;
 5. **for** each unexpanded object o_i in c , **do**
 6. mark o_i as expanded;
 7. **for** each unvisited object o_j of O , **do**
 8. **if** $\text{dist}(o_i, o_j) > \varepsilon$, **then continue**;
 9. **if** $\text{netd}(o_i, o_j) < \varepsilon$, **then**
 10. mark o_j as visited;
 11. add o_j to c ;
 12. **if** $\text{size}(c) \geq \delta_s$, **then**
 13. add c to C ;
 14. **return** C ;
-

Fig. 16. Algorithm: Clustering with Filtering-and-Refinement

5.3. The Road-buddy-based Approach

The road-connection-based clustering algorithm also has the problem of individual sensitivity. The similar idea of traveling buddy can be applied to improve the algorithm's efficiency. The *road buddy* is thus proposed to maintain the small groups of objects moving together along the roads.

Definition 12 (Road Buddy): Let M be the road network, O be the object set and δ_γ be the buddy radius threshold, the road buddy b is defined as a set of objects satisfying: (1) $b \subseteq O$; (2) $\forall o_i \in b$, $\text{netd}(o_i, \text{netcen}(b)) \leq \delta_\gamma$, where $\text{netcen}(b)$ is the projection of the geometry center of b on the road network M . The buddy's radius γ is defined as the road network distance from $\text{netcen}(b)$ to b 's farthest member.

To obtain $\text{netcen}(b)$, the system needs to first compute the geometry center of b , then employ a map matching algorithm to project the geometry center to the nearest road. In this study, we use the map-matching algorithm developed in our previous works [Yuan et al. 2010].

The road buddy has the same operations of *split* and *merge* as the traveling buddy. The initialization of them are also similar. Their major difference is at the maintenance process. Because it is costly to compute the road network distance from $\text{netcen}(b)$ to each member, the maintenance algorithm employs the filtering-and-refinement strategy to reduce time cost, as listed in Fig. 17.

When the data of a new snapshot arrives; Algorithm 8 first computes the network center of each buddy (Lines 2–3), then checks each road buddy to see whether a split operation is needed (Lines 4 – 11), finally scans the buddy set and merges the ones that are close to each other (Lines 12 – 17). The key steps of filtering-and-refinement are at Lines 5, 6, 14 and 15. Before computing the road network distance between two points, the algorithm checks whether their Euclidean distance passing the threshold and only carries out further computation on the qualified pairs.

The road buddy can be used to improve the efficiency of road-connection-based clustering and companion generation by avoiding accessing the object details. Similar to

Algorithm 8. Road Buddy Maintenance

Input: the road network M , the radius threshold δ_γ , the road buddy set B and the coming snapshot s

Output: updated buddy set B'

```

1.  for each  $b_i$  in  $B$  do
2.    update  $cen(b_i)$ ;
3.    match  $cen(b_i)$  to  $M$  and compute  $netcen(b_i)$ ;
4.    for  $o_j$  in  $b_i$ , do
5.      if  $dist(o_j, cen(b_i)) > \delta_\gamma$  then  $isSplit \leftarrow true$ ;
6.      else if  $netd(o_j, cen(b_i)) > \delta_\gamma$  then  $isSplit \leftarrow true$ ;
7.      if  $isSplit = true$ , then //Split Operation
8.        split  $o_j$  out as a new buddy  $b_j$ ;
9.        add  $b_j$  to  $B'$ ;
10.       update  $netcen(b_i)$ ;
11.     add  $b_i$  to  $B'$ ;
12.   //Merge Operation
13.   for each  $b_i, b_j$  in  $B'$ ,  $b_i \neq b_j$  do
14.     if  $dist(netcen(b_i), netcen(b_j)) + \gamma_i + \gamma_j \leq 2\delta_\gamma$  then
15.       if  $netd(netcen(b_i), netcen(b_j)) + \gamma_i + \gamma_j \leq 2\delta_\gamma$  then
16.         merge  $b_i, b_j$  as  $b_k$ ;
17.         remove  $b_i, b_j$  and add  $b_k$  to  $B'$ ;
18.   return  $B'$ ;

```

Fig. 17. Algorithm: Road Buddy Maintenance

the traveling buddy, we propose several lemmas that are helpful for road companion discovery.

Lemma 6: Let b be a road buddy, ε be the distance threshold. If the buddy radius $\gamma \leq \varepsilon/2$, then all the objects in b are directly road connected to each other. Such a road buddy is called a *road-connected buddy*.

Proof: Note that $\gamma \leq \varepsilon/2$, thus for $\forall o_i, o_j \in b$, $netd(o_i, netcen(b)) \leq \gamma$ and $netd(o_j, netcen(b)) \leq \gamma$. Hence there exists a path ζ by-passing $netcen(b)$ that connects o_i and o_j , and $length(\zeta) \leq 2\gamma \leq \varepsilon$. Therefore $netd(o_i, o_j) \leq length(\zeta) \leq \varepsilon$. According to Definition 8, o_i and o_j are directly road connected. ■

Lemma 7: Let b_i, b_j be two road-connected buddies and ε be the distance threshold. If $\exists o_i \in b_i, o_j \in b_j$ such that $netd(o_i, o_j) \leq \varepsilon$, then all the objects of b_i and b_j are network connected.

Proof: If $netd(o_i, o_j) \leq \varepsilon$, then o_i and o_j are directly road connected. Since all the objects in b_i and b_j are directly road connected from o_i and o_j , respectively. Therefore, all the objects in the two traveling buddies are road connected. ■

Lemma 6 and 7 can be used to speed up the road-connection-based clustering. The lemmas show that if two buddies are tight by themselves and close to each other, the system can consider all their members as road connected without further computation.

Lemma 8: Let b_i and b_j be two road buddies with radius γ_i and γ_j , and ε be the distance threshold. If $dist(netcen(b_i), netcen(b_j)) \geq \gamma_i + \gamma_j + \varepsilon$, then the objects in b_i and b_j are not directly road connected.

Proof: As Lemma 5 shows, the Euclidean distance is the lower-bound of road network distance, $netd(netcen(b_i), netcen(b_j)) \geq dist(netcen(b_i), netcen(b_j)) \geq \gamma_i + \gamma_j + \varepsilon$, then

for $\forall o_i \in b_i, o_j \in b_j, netdist(o_i, o_j) \geq \varepsilon$. Therefore, o_i and o_j are not directly network connected. ■

Lemma 8 is helpful to prune most of the unconnected buddies in road-connection-based clustering. Especially the lemma does not require the system to compute any road network distance on M . The system only needs the network center of buddies and their radius as input (which are already computed), the huge I/O cost could be saved.

The detailed algorithm is listed in Fig.18. Algorithm 9 first calls Algorithm 8 to update the road buddies with new data (Line 1), then randomly picks a road buddy as the seed and searches for the ones that are directly road connected (Lines 2 – 13). The buddies that are distant from the seed are filtered out directly without detailed distance computation (Lines 5 – 6). The algorithm searches road-connected buddies with Lemmas 6 and 7 (Lines 7 – 13). Finally, the algorithm outputs the clustering results when all road buddies are processed (Line 15).

Algorithm 9. Road-Buddy-based Clustering

Input: the distance threshold ε , the coming snapshot s and the buddy set B .

Output: the cluster set C

1. update the buddy set B ; //Algorithm 8
 2. randomly pick a buddy b_i in B ;
 3. initialize cluster $c_i \leftarrow b_i$, add c_i to C ;
 4. **for** b_j in $B, b_j \neq b_i$, **do**
 5. **if** $dist(netcen(b_i), netcen(b_j)) - \gamma_i - \gamma_j > \varepsilon$, **then**
 6. **continue**; // Lemma 8
 7. **for** each o_i in b_i, o_j in b_j , **do**
 8. **if** $netd(o_i, o_j) \leq \varepsilon$, **then**
 9. **if** b_i, b_j are road connected **then**
 10. add b_j to c_i ; //Lemmas 6 and 7
 11. **break**;
 12. **else if** o_j is road connected from o_i **then**
 13. add o_j to c_i ;
 14. repeat steps 2 - 13 until all buddies are processed;
 15. **return** the cluster set C ;
-

Fig. 18. Algorithm: Road-buddy-based Clustering

The buddy index can be retrieved from road buddies and help companion generation. Because this technique is actually independent from the metrics and distance computation, Algorithm 5 can be applied directly on road buddies.

6. LOOSE COMPANION DISCOVERY

In many applications such as military object monitoring, several members may temporarily leave the group and go back in short time. The companion discovery algorithm will miss such companions if strictly following the time constraints.

Example 8: Fig.19 shows the trajectory streams of a small team of military troops. At snapshot s_1 , the team members move together. They send out a member o_1 to scout around at s_2 and o_1 returns to the team at s_3 . The team then splits to two parts at s_4 to conduct a “pincer attack” against enemies. Finally they reunion at s_5 . Suppose the size threshold is 6 and the duration threshold δ_t is set as 30 minutes. The system cannot discover any companion from the data if strictly following the constraints.

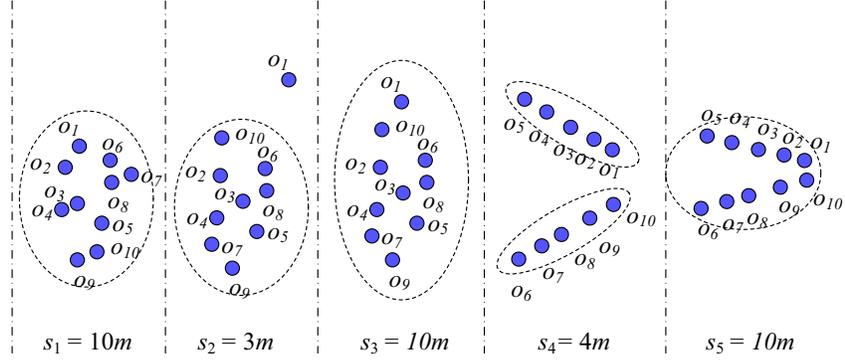


Fig. 19. Example: Movement of Military Troops

In most cases, the rigid time constraints may lead to no result or not the best results of discovered traveling companion. It is necessary to release the constraints for more effective discovery. To this end, we introduce the concept of *loose companion* as follows.

Definition 13 (Loose Companion): Let δ_s be the size threshold, δ_t be the duration threshold and δ_l be the leaving time threshold, a group of objects q is called *loose companion* if:

- (1) Let T be the total time that the members of q are density connected, $T \geq \delta_t$;
- (2) q 's size $size(q) \geq \delta_s$;
- (3) For each member o of q , let t be the maximum period that o is not density-connected with other members of q , $t \leq \delta_l$.

The loose companion allows the member objects temporarily leaving the companion, as long as the leaving time is less than the threshold δ_l . In Fig.19, if we set δ_l as 5 minutes, the military team could be discovered as a companion.

Similarly we propose the definition of *loose buddy*.

Definition 14 (Loose Buddy): Let s be a snapshot of the trajectory stream, δ_γ be the buddy radius threshold and δ_l be the leaving time threshold, *loose buddy* b is defined as a set of objects, for $\forall o_i \in b$,

- (1) $dist(o_i, cen(b)) \leq \delta_\gamma$, where $cen(b)$ is the geometry center of b ;
- or (2) $dist(o_i, cen(b)) > \delta_\gamma$, but the total time of $dist(o_i, cen(b)) > \delta_\gamma$ is less than δ_l .

To discover the loose companions and maintain the loose buddies, the system can follow the same frameworks proposed in previous sections. Only minor modifications need to be carried out in the intersection and split operations. When an object leaves the companion candidate or buddy, the system does not remove that object or split the buddy immediately, instead puts the object/buddy in a buffer to be removed/split after a time period of δ_l . If the object returns in δ_l , the remove/split command will be canceled. Such modification does not influence the general frameworks of companion discovery. The other steps of the clustering-and-intersection algorithm, smart-and-closed method and the buddy-based approach remains the same for loose companion discovery, hence we omit the details here due to space limitation.

7. PERFORMANCE EVALUATION

7.1. Experiment Setup

Datasets: We evaluate the proposed methods on both real and synthetic trajectory datasets. The taxi dataset (D_1) is retrieved from the Microsoft GeoLife and T-Drive projects [Yuan et al. 2010; Zheng et al. 2010] with the road network of Beijing. The trajectories are generated from GPS devices installed on 500 taxis in the city of Beijing.

The dataset is available to public⁵. The military trajectory dataset (D_2) is retrieved from the CBMANET project [Krout 2007], in which an infantry battalion of 780 units, divided as 30 teams, moves from Fort Dix to Lakehurst for a mission on two routes in 3 hours. Meanwhile, to test the algorithm’s performance in large datasets, we also generate two synthetic datasets (D_3 and D_4), being comprised of 1,000 to 10,000 objects, with more than 10 million data records.

Baselines: The proposed Smart-and-Closed algorithm (SC) and Buddy-based discovery algorithm (BU) are compared with Clustering-and-Intersection method (CI), which is used as the framework to find convoy patterns [Jeung et al. 2008]; and two state-of-the-art algorithms: (1) The Swarm pattern (SW) [Li et al. 2010] that captures the objects moving within arbitrary shape of clusters for certain snapshots that are possibly non-consecutive; (2) The TraClu algorithm (TC) [Lee et al. 2007] that discovers the common sub-trajectories with a density-based line-segment clustering algorithm.

Environments: The experiments are conducted on a PC with Intel 6400 Dual CPU 2.13G Hz and 2.00 GB RAM. The operating system is Windows 7 Enterprise. All the algorithms are implemented in Java on Eclipse 3.3.1 platform with JDK 1.6.0. The parameter settings are listed in Fig. 20.

Dataset	Obj. #	Duration	Sample Freq.	Snapshot#	Record#
Taxi (D_1)	500	4.2 hours	5 minutes	50	25,000
Military (D_2)	780	3 hours	1 minute	180	140,400
Syn 1 (D_3)	1,000	24 hours	1 minute	1,440	1.44 M
Syn 2 (D_4)	10,000	24 hours	1 minute	1,440	14.4 M
The companion size threshold δ_s : 5 – 40, default 10					
The companion duration threshold δ_t : 3 – 15, default 10					
The clustering parameter ε and μ are set according to different datasets.					
The buddy radius threshold δ_r : $\varepsilon/2 - \varepsilon/10$, default $\varepsilon/2$.					
The leaving time threshold δ_l : 0 – 6, default 0					

Fig. 20. Experiment Settings

7.2. Comparisons in Discovery Efficiency

In this subsection we conduct experiments to evaluate the efficiency of companion discovery algorithms in Euclidean space. Since both SW and TC cannot output the results incrementally, we take the running time of entire dataset as the measure for time cost. The size of candidate set (# of objects) is used to measure the space cost of companion computation. The only exception is TC, since the algorithm only carries out the sub-trajectory clustering task and does not store any companion candidates, TC’s space cost is not included in the experiment.

We first evaluate the algorithm’s time and space costs on different datasets with default settings. Fig. 21 shows the experiment results. Note that the y-axes are in logarithmic scale. BU achieves the best performances on all the datasets. In the largest dataset D_4 , BU is an order of magnitude faster than CI and SW. BU’s space cost is only 20% of SW and less than 5% of CI.

Figure 22 illustrates the influences of companion size threshold δ_s in the experiments. The experiment is carried on dataset D_3 . Based on default settings, we evaluate the algorithms with different values of δ_s . Generally speaking, when the size threshold

⁵GeoLife GPS Trajectories Datasets. Released at: <http://research.microsoft.com/en-us/downloads/b16d359d-164-469e-9fd4-daa38f2b2e13/default.aspx>

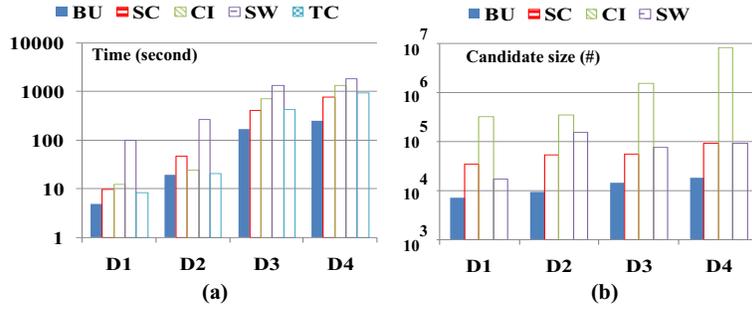
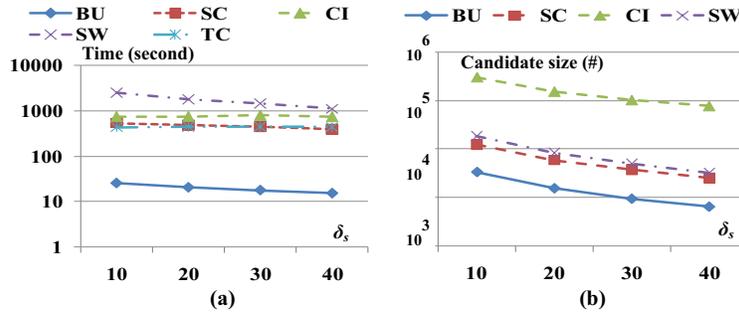


Fig. 21. Efficiency: (a) time, (b) space on diff. datasets

grows larger, the filtering mechanism is more effective to prune more companion candidates in each snapshot. The space costs reduce significantly, and the running times also decrease for fewer intersections.

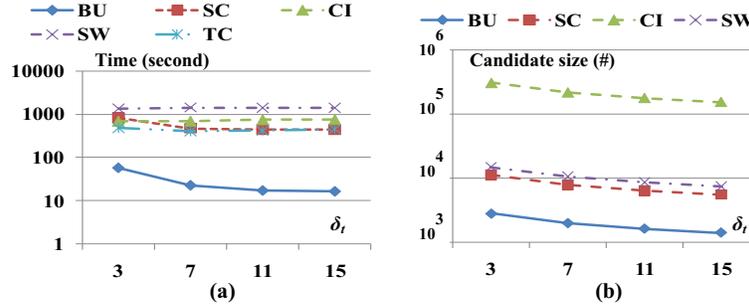
Fig. 22. Efficiency: (a) time, (b) space vs. δ_s

We also study the influence of duration threshold δ_t . Based on default settings, the experiments are conducted on dataset D_3 . The value of δ_t is changed from 3 to 15, the algorithm's performances are shown in Fig. 23. BU, SC and CI are all faster when δ_t grows larger, because many companion candidates are not consistent enough to last for a long time. When setting δ_t as 15 snapshots, BU can process the dataset in less than 20 seconds (Fig. 23 (a)). It is almost an order of magnitude faster than SC and CI. TC is not influenced by δ_s and δ_t , since it is only a clustering algorithm and does not generate any companion candidates. Beside TC, SW also could not improve the performance when δ_t increases, the reason is SW utilizes the *object-growth* strategy to prune candidates. Such heuristics could only work with the size threshold δ_s , but cannot benefit from larger δ_t .

7.3. Efficiency Analysis for Buddy-based Discovery

Why is the buddy-based discovery algorithm more efficient? BU outperforms other methods in the efficiency evaluations, especially in the scenarios of long lasting stream with large number of objects. In this subsection we carry out the experimental analysis to reveal the advantages of buddy-based discovery method.

In the beginning, we tune the parameters of BU to study the factors that influence its efficiency. With δ_s and δ_t set as default values, we test BU with different buddy radius threshold δ_γ from $\varepsilon/10$ to $\varepsilon/2$, and record the average buddy size $|b|$, buddy number and algorithm's running time. Their relationships are demonstrated in Fig. 24. One

Fig. 23. Efficiency: (a) time, (b) space vs. δ_t

can clearly learn from Fig. 24 (a) that the total buddy number is inversely proportional to the average buddy size $|b|$. In addition, the number of unchanged buddies decreases rapidly as $|b|$ grows larger. However, as shown in Fig. 24 (b), the running time of both buddy-based clustering (B-Cluster) and BU decreases with larger $|b|$. This phenomenon can be explained by Proposition 2, the cost of buddy's maintenance algorithm is $O(n + m^2)$, where n is the number of objects and m is the number of buddies. If n is fixed, then m is inversely proportional to $|b|$. Hence BU costs less time if $|b|$ is larger. Based on the efficiency analysis, we recommend setting the buddy radius as a relatively large value (such as $\varepsilon/2$). Fig. 24 (b) also records the time cost of DBSCAN clustering algorithm as a reference. Even if less than 20% buddies stay unchanged (which is rare for real-world objects), as long as the average size of the buddies is larger than 3, the buddy-based clustering algorithm can still outperform DBSCAN. The experiment results show that BU is especially feasible for processing a trajectory stream with dense object clusters.

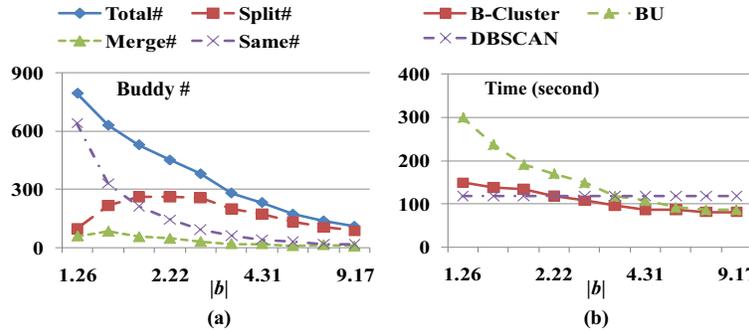


Fig. 24. Efficiency Analysis: (a) buddy number, (b) time vs. buddy size

BU has three steps, namely the maintenance step (M-step, Algorithm 3), clustering step (C-step, Algorithm 4) and intersection step (I-step, Algorithm 5). To study the time cost of each step, the system carries out BU on the four datasets and record the time costs of each step, as well as their proportions in the total running time, as shown in Fig. 25. The results denote that the clustering step is actually the most efficient in the three, it costs less than 5% of the total running time, compared to the DBSCAN clustering which usually takes 40-50% of the total running time of SC. BU spends an extra 10% to 15% time in maintaining the buddies to save more time from the clustering task. The two lemmas, especially Lemma 3, utilize the buddy information to filter out many objects without accessing their details. In addition, the buddy index

helps to reduce the size of the candidate set, and so decreases the intersection times of companion discovery.

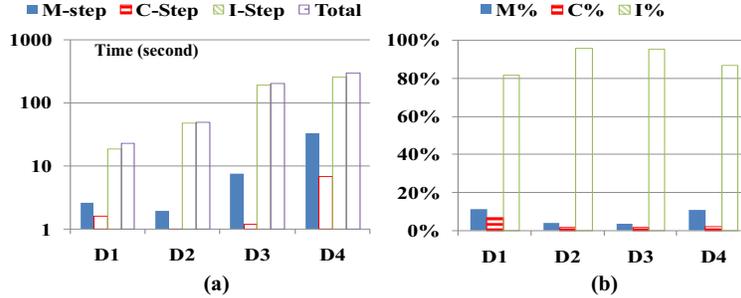


Fig. 25. Efficiency Analysis: (a) running time, (b) percentage of BU steps on diff. datasets

7.4. Evaluations on Algorithm's Effectiveness

The third part of the experiment is to evaluate the quality of the retrieved companions. In dataset D_2 , an infantry battalion of 780 units moves from Fort Dix to Lakehurst for a mission on two routes in 3 hours. The objects are organized in 30 teams, each team has 25 to 30 units. The information of team partitioning is retrieved as the ground truth. The algorithm's outputs are matched to the ground truth and the measures of precision and recall are calculated as follows.

Precision: The proportion of true companions over all the retrieved results of the algorithm. It represents the algorithm's selectivity in finding out meaningful companions.

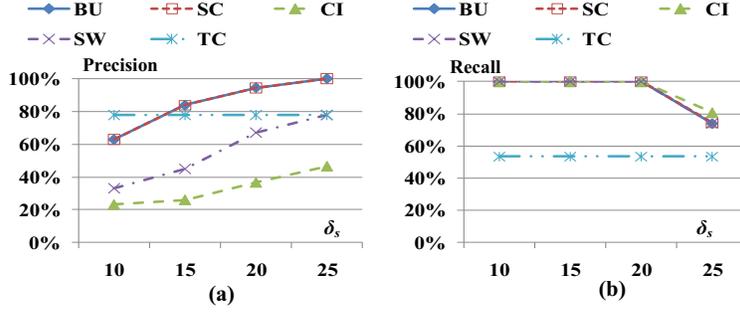
Recall: The proportion of detected true companions over the ground truth. This criteria shows the algorithm's sensitivity for detecting traveling companions.

We conduct experiments with different values of the size threshold δ_s . The results of effectiveness evaluation are shown in Fig. 26. BU and SC have same precision and recall since they output identical companions. They have about 20% precision improvement over SW, and near 40% precision improvement over CI. SW generates the swarm patterns of frequently meeting objects, which is actually a super set of the companions. The swarm pattern is highly sensitive to help find out all the companions (*i.e.*, 100% recall), but SW also generates more false positives that bring down the algorithm's selectivity. CI has the same problem with even lower precision. Since there are many redundant and non-closed companions in the results, more than half of CI's results are not useful.

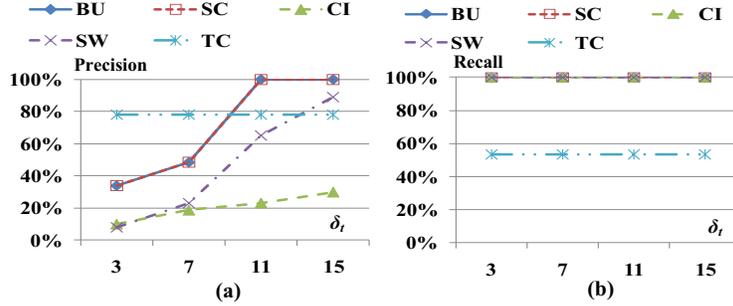
Again, TC is not affected by the parameters of δ_s and δ_t . TC takes the movement direction as an important measure to compute sub-trajectory clusters; its results reflect the major directions of the object movements. However, such clusters may not capture the information of companions, because the companion member's moving direction might be different. As an illustration, please go back to Fig. 1. From snapshot s_2 to s_3 , the moving directions of o_8 and o_9 are different, hence they may be put in different sub-trajectory clusters.

Another interesting observation is that, in Fig. 26, BU, SC, CI and SW's precisions all increase when δ_s becomes larger, since fewer companions can pass a higher size threshold. However, if δ_s is set too high (more than 25), several true companions will also be filtered out and the algorithm cannot achieve 100% recall.

In the next experiment, we study the influence of time threshold δ_t . Fig. 27 shows the precision and recall of the five algorithms with different δ_t on D_2 . BU and SC achieve better performance than SW and CI. When increasing δ_t , the algorithm's precision

Fig. 26. Effectiveness: (a) precision, (b) recall vs. δ_s

increases, but they can still keep a high recall. Since all the true companions last for a long period in D_2 . If we set δ_t greater than 11, both BU and SC can achieve 100% precision and recall. Hence we suggest that in real applications, the user should set a relatively high time threshold to filter out false positives, but a moderate size threshold to guarantee the algorithm's sensitivity.

Fig. 27. Effectiveness: (a) precision, (b) recall vs. δ_t

7.5. Experiments on Road Companion Discovery

To test the efficiency of road companion discovery, we perform the evaluation on dataset D_1 with the road network of Beijing, which has 106,579 road nodes and 141,380 road segments. The default size threshold δ_s is set as 8 and the time threshold δ_t is set as 11. In this experiment, we compare the performance of four methods: (1) The Clustering-and-Intersection framework with road network distance computation (CI); (2) The Smart-and-Closed algorithm with road network distance computation (SC); (3) The smart-and-closed algorithm with Filtering-and-Refinement strategy (FR); and (4) The Road-Buddy based method (RB).

We first evaluate the time and space costs of road companion discovery. The number of accessed road nodes is used as the measure for I/O cost. Based on default settings, we evaluate the algorithms with different values of δ_s . Figure 28 shows the running time and accessed node number. Generally speaking, when the size threshold grows larger, both running time and I/O costs decrease. The computation cost of road companion discovery is much larger than the traveling companion discovery on Euclidean space. This is mainly caused by the high I/O overhead in road network distance computation. Since the road network distance computation becomes the major cost, SC cannot save

much time comparing to CI. However, FR and RB are an order of magnitude faster than SC and CI, because they utilize the filtering-and-refinement strategy to avoid most unnecessary road network distance computations. The effects of RB is better, since RB groups the objects in small buddies and limits the distance computation in a small region with lower I/O overhead.

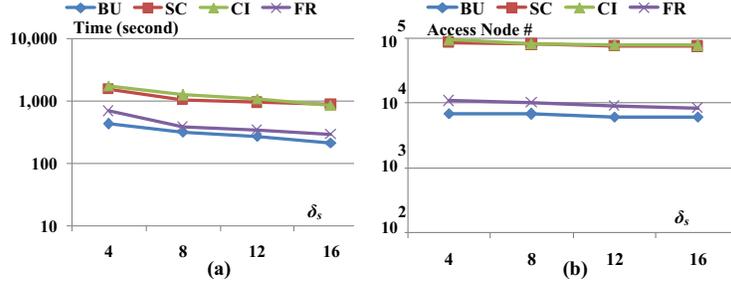


Fig. 28. Efficiency: (a) time, (b) I/O of road companion discovery vs. δ_s

The influence of duration threshold δ_t is also studied in our experiment. Based on default settings, the value of δ_t is changed from 3 to 15, the algorithm's performances are shown in Fig. 29. All the algorithms run faster when δ_t grows larger, because fewer road companion candidates can last for a long time. Again, RB and FR only cost 20% to 50% time as CI and SC. From the experiments we can see clearly that, the main bottleneck of road companion discovery is at the distance computation stage, RB and FR reduce the times of unnecessary distance computation, therefore they can outperform other algorithms.

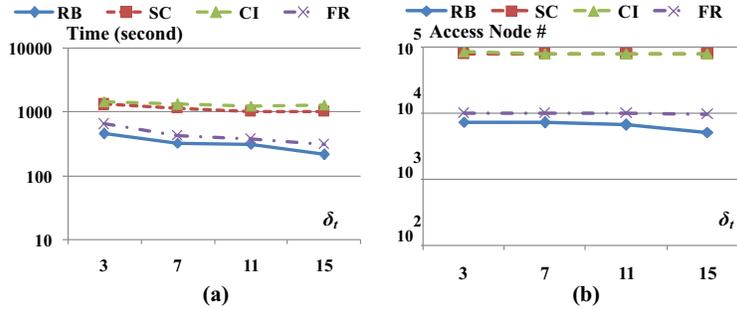
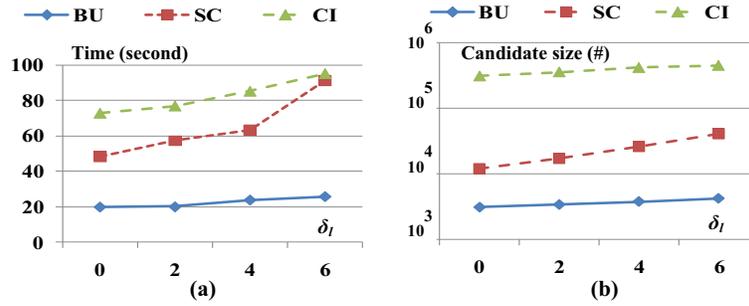


Fig. 29. Efficiency: (a) time, (b) I/O of road companion discovery vs. δ_t

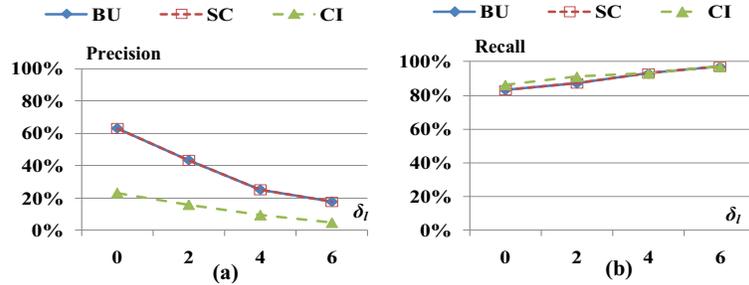
7.6. Evaluations on Loose Companion Discovery

In the previous experiments, we set the leaving threshold δ_l as 0. In this subsection, we conduct experiments on loose companion discovery. We run the algorithms of BU, SC and CI on dataset D_3 by tuning δ_l from 0 to 6 snapshots.

Fig. 30 shows the algorithms' time and space costs. With larger δ_l , all the algorithms' space costs increase rapidly since they cannot prune the candidates if several objects temporarily leave the companion, hence the system has to spend more time in making intersections with a larger candidate set. However, even with large δ_l , BU still can discover the loose companions in about 20 seconds.

Fig. 30. Efficiency: (a) time, (b) space vs. δ_l

Finally we carry out the effectiveness experiment on the military dataset D_2 . δ_l is changed from 0 to 6 snapshots, and other parameters are set as the default values. As shown in Fig. 31 (a), the precision of companion discovery decreases with larger δ_l , since more companions are generated and inevitably the number of false positive increases. However, the good news is that the recall increases as δ_l grows (Fig. 31 (b)). With a released time constraint, BU and SC can discover more meaningful companions and achieve a high recall more than 90%, which increases the system's feasibility in real applications.

Fig. 31. Effectiveness: (a) precision, (b) recall vs. δ_l

8. RELATED WORK

According to the methodologies, the related works of traveling companion discovery can be loosely classified into two categories: trajectory clustering and movement pattern discovery.

8.1. Trajectory Clustering

The works in this category focus on developing efficient algorithms to cluster moving objects. Gaffney *et al.* first proposed the fundamental principles of clustering moving objects based on the theories of probabilistic modeling [Gaffney and Smyth 1999; Cadez *et al.* 2000]. Many distance functions, such as DTW [Yi *et al.* 1998] and LCSS [Gunopoulos 2002] are proposed. Lee *et al.* proposed a novel partition-and-group framework to find the clusters based on sub-trajectories [Lee *et al.* 2007].

In [Har-Peled 2003], Har-Peled shows that the moving objects can be clustered when the resulting clusters are competitive at any time during the motion. Yang *et al.* proposed the idea of neighbor-based pattern detection method for windows [Yang *et al.* 2009]. Ester *et al.* made the progress to generate incremental clusters

[Ester et al. 1998]. Li *et al.* propose a micro-cluster [Li et al. 2004] based schema to cluster moving objects. Zhang and Lin use the k-centre clustering algorithm [Gonzalez 1985] for histogram construction. A distance function combining velocity and position differences is proposed in their work [Zhang and Lin 2004]. More recently, Jensen *et al.* utilize the velocity features to cluster objects for the current and near future positions [Jensen et al. 2007].

However, as pointed out in [Jeung et al. 2008], most of the above methods can not be used directly for traveling companion discovery. The major problem is that those algorithms tend to generate clusters for the entire trajectory dataset, instead of each snapshot. Hence the detailed object relationships and evolving companion patterns are all lost. In addition, some algorithms require the object's velocity in advance and need to scan the data for multiple times. Such requirements are not fit for the trajectory streams.

8.2. Movement Pattern Discovery

Movement pattern discovery is a hot topic in recent years. The problem has been variously referred to as the search for *flocks* [Gudmundsson and Kreveld 2006], *moving clusters* [Kalnis et al. 2005], *spatial-tempo joins* [Bakalov et al. 2005], *spatial co-locations* [Yoo and Shekhar 2004], *meetings* [Gudmundsson et al. 2004], *convoy* [Jeung et al. 2008], *moving groups* [Aung 2008], *swarms* [Li et al. 2010] and so on.

One of the earliest works is *flock* discovery [Gudmundsson et al. 2004]. A flock is defined as a group of objects moving together within a circular region [Gudmundsson and Kreveld 2006]. There are several variations of this model: *Variable flock* permits the members to change during the time span [Benkert et al. 2008], *meeting* is a circle similar to flock but fixed in a single location all the time [Gudmundsson and Kreveld 2006]. However, such shapes are restricted to circles and the results are also sensitive to the parameter of radius.

Li *et al.* designed a flow scan algorithm for hot route mining [Li et al. 2007]. Liu *et al.* mined frequent trajectory patterns by using RF tag arrays. Their work successfully demonstrated the feasibility and the effectiveness of movement patterns in real life [Liu et al. 2007]. Tao *et al.* proposed the technique of spatio-temporal aggregation using sketch index. This method can process the queries an order of magnitude faster than the previous works [Tao et al. 2004]. Giannotti *et al.* proposed the *interest region* based mining algorithm [Giannotti et al. 2007]. Horvitz *et al.* propose the models of using groups of mobile users to discover congestions in urban areas [Horvitz et al. 2005].

Zhang *et al.* propose the techniques to produce intersections of streaming moving objects [Zhang et al. 2008]. This method is a big improvement from existing algorithms by the speed-up of several orders of magnitude. Nutanong *et al.* use a safe region to report objects that do not change over time [Nutanong et al. 2008]. The proposed V*-Diagram has much smaller I/O and computation costs than previous methods. It outperforms the best existing technique by two orders of magnitude.

However, since the above methods focus more on discovering hot spots, regions or routes rather than object groups, they cannot be used directly for companion discovery.

Kalnis *et al.* proposed the first study to automatic extract *moving clusters* from large spatial datasets [Kalnis et al. 2005]. In a recent work, Jeung *et al.* proposed the framework of *convoy* query [Jeung et al. 2008]. It is a significant step forward in the works of movement pattern mining, since it allows the objects to organize in arbitrary shapes. Li *et al.* further released the constraints of convoy and proposed the *swarm pattern* to discover object groups in a sporadic way [Li et al. 2010].

The concepts of convoy and swarm patterns are similar to traveling companion. The major differences are about the discovery algorithms. The convoy algorithm needs to scan the entire trajectory into memory to make trajectory simplification, and the sys-

tem also needs to load the whole dataset into memory to search for swarms. Hence it is impractical to use them in a data stream environment. In addition, the swarm pattern is a frequent itemset-based concept. Since it is difficult to detect large size frequent itemsets [Zhu et al. 2007], the swarm pattern has limited applicability for datasets with large scale objects.

Fig. 32 compares the features of some related methods with the proposed algorithms to discovery traveling companions, road companions and loose companions.

Name	Pattern Shape	Object Number	Partnership Discover	Incremental Output	Released Time Constraints
<i>TraCluster</i>	arbitrary	multiple	No	No	No
<i>Flock</i>	circle	multiple	Yes	No	No
<i>Meeting</i>	circle	multiple	Yes	No	No
<i>Hot Route</i>	road segment	multiple	No	No	No
<i>Swarm</i>	arbitrary	multiple	Yes	No	Yes
<i>Convoy</i>	arbitrary	multiple	Yes	No	No
<i>Traveling companion</i>	arbitrary	multiple	Yes	Yes	No
<i>Road companion</i>	along the roads	multiple	Yes	Yes	No
<i>Loose companion</i>	arbitrary	multiple	Yes	Yes	Yes

Fig. 32. The Comparison with Related Works

9. CONCLUSION AND FUTURE WORK

In this study we investigate the problem of traveling companion discovery on trajectory data streams. We propose the algorithms of smart-and-closed discovery to efficiently generate companions from trajectory data. The model of traveling buddy is proposed to help improve both the clustering and intersection processes for companion discovery. The proposed methods are extended to more complex scenarios for road companion and loose companion discovery. We evaluate the proposed algorithms in extensive experiments on both real and synthetic datasets. The buddy-based method is shown to be an order of magnitude faster than existing approaches on both Euclidean space and road networks. The effectiveness of buddy-based algorithm also outperforms other competitors in terms of precision and recall.

In the future, we are going to integrate the companion discovery methods to real application services such as battlefield monitoring systems and traffic analysis services.

10. ACKNOWLEDGEMENTS

The work was supported in part by U.S. NSF grants IIS-0905215, CNS-0931975, CCF-0905014, IIS-1017362, the U.S. Army Research Laboratory under Cooperative Agreement No. W911NF-09-2-0053 (NS-CTA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

REFERENCES

AUNG, H.-H. 2008. Discovering moving groups of tagged objects. In *Technique Report, National University of Singapore*.

- BAKALOV, P., HADJIELEFTHERIOU, M., AND TSOTRAS, V. J. 2005. Time relaxed spatiotemporal trajectory joins. In *ACM GIS*.
- BENKERT, M., GUDDMUNDSSON, J., HUBNER, F., AND WOLLE, T. 2008. Reporting flock patterns. *Comput. Geom. Theory Appl.* 41, 3, 111–125.
- CADEZ, I. V., GAFFNEY, S., AND SMYTH, P. 2000. A general probabilistic framework for clustering individuals and objects. In *SIGKDD*.
- ESTER, M., KRIEGEL, H.-P., SANDER, J., WIMMER, M., AND XU, X. 1998. Incremental clustering for mining in a data warehousing environment. In *VLDB*.
- ESTER, M., KRIEGEL, H.-P., SANDER, J., AND XU, X. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *SIGKDD*.
- GAFFNEY, S. AND SMYTH, P. 1999. Trajectory clustering with mixtures of regression models. In *SIGKDD*.
- GIANNOTTI, F., NANNI, M., PEDRESCHI, D., AND PINELLI, F. 2007. Trajectory pattern mining. In *SIGKDD*.
- GONZALEZ, T. 1985. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, 293–306.
- GUDDMUNDSSON, J. AND KREVELD, M. V. 2006. Computing longest duration flocks in trajectory data. In *ACM GIS*.
- GUDDMUNDSSON, J., KREVELD, M. V., AND SPECKMANN, B. 2004. Efficient detection of motion patterns in spatio-temporal data sets. In *ACM GIS*.
- GUNOPOULOS, D. 2002. Discovering similar multidimensional trajectories. In *ICDE*.
- HAN, J. AND KAMBER, M. 2006. *Data Mining: Concepts and Techniques Second Edition*. Morgan Kaufmann.
- HAR-PELED, S. 2003. Clustering motion. *Discrete and Computational Geometry* 31, 4, 545–565.
- HORVITZ, E., APACIBLE, J., SARIN, R., AND LIAO, L. 2005. Prediction, expectation, and surprise: Methods, designs, and study of a deployed traffic forecasting service. In *Twenty-First Conference on Uncertainty in Artificial Intelligence, UAI*.
- JENSEN, C. S., LIN, D., AND OOI, B. C. 2007. Continuous clustering of moving objects. *IEEE TKDE* 19, 9, 1161–1174.
- JEUNG, H., YIU, M. L., ZHOU, X., JENSEN, C. S., AND SHEN, H. T. 2008. Discovery of convoys in trajectory databases. In *VLDB*.
- KALNIS, P., MAMOULIS, N., AND BAKIRAS, S. 2005. On discovering moving clusters in spatial-temporal data. In *SSTD*.
- KROUT, T. 2007. Cb manet scenario data distribution. In *Technique Report of BBN*.
- LEE, J.-G., HAN, J., AND WHANG, K.-Y. 2007. Trajectory clustering: a partition-and-group framework. In *SIGMOD*.
- LEE, M., HSU, W., JENSEN, C. S., CUI, B., AND TEO, K. 2003. Supporting frequent updates in r-trees: A bottom-up approach. In *VLDB*.
- LI, X., HAN, J., LEE, J.-G., AND GONZALEZ, H. 2007. Traffic density based discovery of hot routes in road networks. In *SSTD*.
- LI, Y., HAN, J., AND YANG, J. 2004. Clustering moving objects. In *SIGKDD*.
- LI, Z., DING, B., HAN, J., AND KAYS, R. 2010. Swarm: Mining relaxed temporal moving object clusters accurate discovery of valid convoys from moving object trajectories. In *VLDB*.
- LIU, Y., CHEN, L., PEI, J., CHEN, Q., AND ZHAO, Y. 2007. Mining frequent trajectory patterns for activity monitoring using radio frequency tag arrays. In *IEEE PerCom*.
- NUTANONG, S., ZHANG, R., TANIN, E., AND KULIK, L. 2008. The v*-diagram: A query dependent approach to moving knn queries. In *VLDB*.
- PEARL, J. 1984. *Heuristics: intelligent search strategies for computer problem solving*. Addison-Wesley Longman Publishing Co., Inc.
- TANG, L.-A., YU, X., KIM, S., HAN, J., HUNG, C.-C., AND PENG, W.-C. 2010. Tru-alarm: Trustworthiness analysis of sensor networks in cyber-physical systems. In *ICDM*.
- TANG, L.-A., ZHENG, Y., XIE, X., YUAN, J., YU, X., AND HAN, J. 2011. Retrieving k-nearest neighboring trajectories by a set of point locations. In *SSTD*.
- TANG, L.-A., ZHENG, Y., YUAN, J., HAN, J., LEUNG, A., HUNG, C.-C., AND PENG, W.-C. 2012. On discovery of traveling companions from streaming trajectories. In *ICDE*.
- TAO, Y., KOLLIOS, G., CONSIDINE, J., LI, F., AND PAPADIAS, D. 2004. Spatio-temporal aggregation using sketches. In *ICDE*.
- YANG, D., RUNDENSTEINER, E. A., AND WARD, M. O. 2009. Neighbor-based pattern detection for windows over streaming data. In *EDBT*.

- YI, B., JAGADISH, H. V., AND FALOUTSOS, C. 1998. Efficient retrieval of similar time sequences under time warping. In *ICDE*.
- YOO, J. S. AND SHEKHAR, S. 2004. A partial join approach for mining co-location patterns. In *ACM GIS*.
- YUAN, J., ZHENG, Y., ZHANG, C., XIE, W., XIE, X., SUN, G., AND HUANG, Y. 2010. T-drive: driving directions based on taxi trajectories. In *GIS*.
- ZHANG, Q. AND LIN, X. 2004. Clustering moving objects for spatial-temporal selectivity estimation. In *ADC*.
- ZHANG, R., LIN, D., RAMAMOHANARAO, K., AND BERTINO, E. 2008. Continuous intersection joins over moving objects. In *ICDE*.
- ZHENG, K., ZHENG, Y., XIE, X., AND ZHOU, X. 2012. Reducing uncertainty of low-sampling-rate trajectories. In *ICDE*.
- ZHENG, Y., XIE, X., AND MA, W. 2010. *GeoLife: A Collaborative Social Networking Service among User, location and trajectory*. IEEE Data Engineering Bulletin.
- ZHENG, Y. AND ZHOU, X. 2011. *Computing with Spatial Trajectories*. Springer.
- ZHU, F., YAN, X., HAN, J., YU, P. S., AND CHENG, H. 2007. Mining colossal frequent patterns by core pattern fusion. In *ICDE*.