On-Line Discovery of Flock Patterns in Spatio-Temporal Data

Marcos R. Vieira
University of California
Riverside, CA 92507, USA
mvieira@cs.ucr.edu

Petko Bakalov
ESRI
Redlands, CA 92373, USA
pbakalov@esri.com

Vassilis J. Tsotras
University of California
Riverside, CA 92507, USA
tsotras@cs.ucr.edu

ABSTRACT

With the recent advancements and wide usage of location detection devices, large quantities of data are collected by GPS and cellular technologies in the form of trajectories. While most previous work on trajectory-based queries has concentrated on traditional range, nearest-neighbor and similarity queries, there is an increasing interest in queries that capture the "aggregate" behavior of trajectories as groups. Consider, for example, finding groups of moving objects that move "together", i.e. within a predefined distance to each other, for a certain continuous period of time. Such queries typically arise in surveillance applications, e.g. identify groups of suspicious people, convoys of vehicles, flocks of animals, etc. In this paper we first show that the on-line flock discovery problem is polynomial and then propose a framework and several strategies to discover such patterns in streaming spatio-temporal data. Experiments with real and synthetic trajectorial datasets show that the proposed algorithms are efficient and scalable.

Categories and Subject Descriptors

H.2 [**DATABASE MANAGEMENT**]: Database applications—*Data mining; Spatial databases and GIS*

Keywords

moving objects, spatio-temporal patterns

1. INTRODUCTION

Recent advances in the area of location-detection devices (RFID, GPS, etc.) and their widespread use have enabled the creation of complex tracking and situational awareness systems which continuously monitor the position of moving objects of interest. Examples include *AccuTracking* [1], *tracNET24* [12], Path Intelligence's *FootPath* [23], InSTEDD's *GeoChat* [9] and many others. This abundance of information, generated by those systems, motivates the need to develop efficient techniques for answering interesting queries about the past behavior of the moving objects like discovering similarity patterns among the object trajectories.

The existing methods for querying trajectories are mainly focused on answering simple single predicate range [25] or nearest neighbor queries [28]. Examples include queries like "find all moving objects that were in area A at 10 a.m. (in the past)" or "find the car which drove as close as possible to the location B during the time interval (10am:1pm)". Recently, a new group of similarity search querying methods have emerged [6][20][29]. The result of a similarity search query is a trajectory closest to the query trajectory according to some metric distance (e.g. Euclidean, Dynamic Time Warping, etc.). There are also work on spatio-temporal joins (e.g. [2][3]). Common to all the above methods is that the query answer is validated per trajectory. That is, a trajectory is reported to the user if its individual behavior satisfies the query predicate(s). In other words, all the above queries focus on the behavior of a trajectory as a single object and thus cannot be used to discover group patterns between the trajectories.

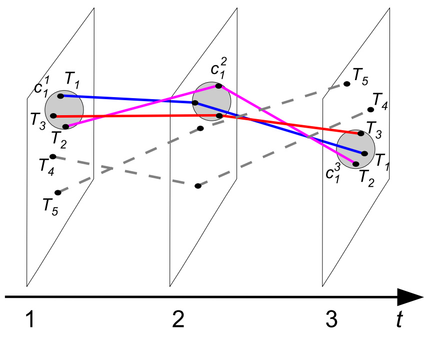

Figure 1: A flock pattern example: $\{T_1, T_2, T_3\}$

Recently there has been increased interest in querying patterns capturing "collaborative" or "group" behavior between moving objects. This includes queries like moving clusters [16, 14], convoy queries [15] and flocks patterns [5, 4, 10]. Such queries discover groups of moving objects that have a "strong" relationship in the space for a given time duration. The difference between all those patterns is the way they define the relationship between the moving objects and their duration in time. In this paper we consider the discovery of *flock* patterns among the moving objects, i.e., the problem of identifying all groups of trajectories that stay "together" for the duration of a given time interval. We consider moving objects to be "close" together if there exists a disk with given radius that covers all moving objects in the pattern (see Fig-

ACM GIS '09, November 4-6, 2009. Seattle, WA, USA (c) 2009 ACM ISBN 978-1-60558-649-6/09/11...$10.00.

ure 1). A trajectory satisfies the above pattern as long as "enough" other trajectories are contained inside the disk for the specified time interval; that is, the answer is based not only on a given trajectory's behavior but also on the ones near it. Such patterns are useful in security and monitoring applications, for example to potentially identify suspicious behavior within large number of people ("Identify all groups of five or more people that were always within a disk of 100 feet in the last 30 minutes") or to study patterns of animal behavior [7, 30, 24] (e.g. migration of sharks, whales, birds, etc.).

The example in Figure 1 shows a flock pattern containing 3 trajectories $\{T_1, T_2, T_3\}$ that are within a query defined disk for 3 consecutive time instances. Note that the location of the disk can freely "move" in the 2-dimensional space in order to accommodate all three moving objects and its center does not need to coincide with any moving object location for a given time instance. This makes the discovery of flock patterns difficult because there is an infinite number of possible placements of the disk at any time instance. It is that difficulty that makes the existing methods for flock pattern discovery [5, 4, 10] suffer from severe limitations. Such methods either find approximate solutions, or can be applied only for a single time instance of the problem (i.e. the solution does not support the minimum time duration in the query). To the best of our knowledge, our work is the first one to present exact solutions for reporting flock patterns in polynomial time. It is also the first one that does so for on-line environments. Our work is also different than clustering-based approaches (since clusters are not restricted to a specific shape); flocks are also different than convoy discovery [15]. More details of the previous methods are discussed in Section 2.

We start by providing a complexity analysis for the on-line flock problem. Our analysis reveals that polynomial time solution can be found through identifying a discrete number of locations to place the center of the flock disk inside the spatial universe. The number of such possible locations is polynomial in the total number of moving objects. Based on this analysis we propose several evaluation algorithms that can be used to find flock patterns in polynomial time. The first algorithm is based on time-joins, i.e., merging the results from one time instance to another. The other four algorithms use the *filter-and-refinement* paradigm with the purpose of reducing the total number of candidates and thus the overall computation cost of the algorithm. We evaluate our solutions using several real and synthetic moving object datasets.

The rest of the paper is organized as follows: Section 2 highlights related work while Section 3 formally defines the on-line flock pattern and provides a complexity analysis on the problem. Section 4 describes the proposed algorithms for flock pattern discovery. Section 5 presents the performance evaluation of our proposed algorithms and Section 6 concludes the paper.

2. RELATED WORK

Related work can be classified to (i) research on clustering moving objects, (ii) work on discovering convoys among trajectories and (iii) previous work on flock discovery. Various clustering algorithms have been proposed for static spatial datasets, with different strategies ranging from partitioning (e.g. k-medoids [22]), to hierarchical (e.g. *BIRCH* [31] and *CURE* [11]) and density-based (e.g. *DBSCAN* [8]). The DBSCAN algorithm works for arbitrary-shaped clusters based on the notion of density reachability. This method utilizes two parameters (maximum distance *eps* and minimum number of points *minPts*) to identify dense areas. It starts with an arbitrary starting point that has not been visited. Point that has more than *minPts* within *eps* distance is considered to be in a dense area and flagged as such. All points inside such dense area are processed recursively the same way. Otherwise, those points are considered not "reachable" from a dense area and are labeled as outliers.

Clustering for moving objects was examined in [16], where the *DBSCAN* algorithm is performed for every time instance of the dataset. Then clusters that have been found for two consecutive time instances $t-1$ and t are joined. The clusters can be joined only if the number of common objects among them are above the predefined parameter θ. A cluster is reported if no other new cluster can be joined to it. This process is applied each time for all time instances in the dataset. Other works on clustering moving objects also include [14, 27, 21, 19, 18]. In [14] techniques were proposed to incrementally update clusters of moving objects based on the cluster centers. The object movements were used to predict the cluster evolution over time. The *MONIC* framework [27] deals with transitions in moving clusters, e.g. disappearance and splitting. [21] presented the *microclustering* technique that groups moving objects that are not only close to each other at a specific time instance, but are also expected to move together in the "near" future. Recently, [18, 19] proposed to segment trajectories into line segments. Then line segments are grouped together to build the clusters. However, time is not consider in [18, 19], which makes some line segments to be clustered together even though they are not "close" when time is considered. Nevertheless, such approaches for clustering moving objects cannot solve the flock pattern query since: **(1)** they use different criteria when joining the moving object clusters for two consecutive time instances; **(2)** they employ clustering algorithms, and therefore no strong relationship among *all* elements are enforced; **(3)** moving clustering does not require the *same* set of moving objects to stay in a cluster *all the time* for the specified minimum duration.

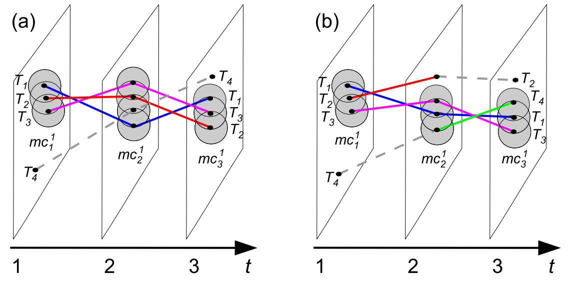

Figure 2: Clustering vs. flock patterns

Related to discovering collaborative behavior between trajectories is the work of finding *convoy patterns* in trajectory archives [15]. A convoy query is defined as a "dense" cluster of trajectories that stay "together" at least for a predefined continuous time. This type of query has four parameters: *eps* and *minPts* (the same used by the *DBSCAN* algorithm), θ' (threshold used to join clusters), and δ' (minimum duration time). *Convoy patterns* are closely related to moving clustering since both use clustering algorithms as the base of their algorithms. The main difference between these two methods is on the criteria for how clusters are joined between two consecutive timestamps. However, neither of them can solve the flock pattern query since clusters do not assume any shape restriction. For example, in Figure 2(a) convoy query returns trajectories $\{T_1, T_2, T_3\}$ for $\theta = 3$ and for 3 time instances, while in Figure 2(b) it does not return nothing. For the moving cluster, if $\theta = 1$ then moving clusters return nothing in both Figure 2(a) and (b). On the other hand, if $\theta = 1/2$ then it returns $\{T_1, T_2, T_3\}$ in Figure 2(a) and $\{T_1, T_3, T_4\}$ in Figure 2(b), but the last one is not a convoy query. Both examples return results based on the density of the objects, but for the flock pattern it would return nothing in either examples. The reason is that in both examples the objects belong to dense areas but they do not have "strong" interaction among them.

Flock pattern query was first introduced in [5, 4, 17], without the notion of minimum lasting time. Later [10] introduced the minimum duration as a parameter of the pattern. Unlike the convoy patterns in a flock the cluster has a predefined shape – a disk with radius r. A set of moving objects is considered a flock if there is a disk with radius r which covers all of them and there are at least some predefined number of objects in the disk. It is shown in [10] that the discovery of the "longest" duration flock pattern is an *NP-hard* problem. As a result, [10] presents only approximation algorithms. To the best of our knowledge our paper is the first which proposes a polynomial time solution to the flock problem with a "predefined" time duration. Moreover our algorithms can be applied in a streaming environment for on-line discovery of the flock patterns.

3. PRELIMINARIES

We assume that object O_{id} is uniquely identified by identifier id. Its movement is represented by a trajectory T_{id} which is defined as an ordered sequence of n multidimensional points $T_{id} = \{p(t_1), p(t_2), \ldots, p(t_n)\}$. Here t_i is a timestamp and $p(t_i)$ is the location of object O_{id} in the two dimensional space $\mathbb{R}^2$ as recorded at timestamp t_i ($t_i \in \mathbb{N}$, $t_{i-1} < t_i$, and $0 < i \leq n$). For simplicity when we discuss the current time instance, t_i is omitted, and we just use p_{id} to denote the object location.

Given two object locations $p_a^{t_i}$ and $p_b^{t_i}$ in a specific time instance t_i from trajectories T_a and T_b respectively, $d(p_a^{t_i}, p_b^{t_i})$ denotes the L_2 Euclidean distance between p_a and p_b. Even though here in this paper we only use the L_2 distance, our methods can be generalized to other metrics as well. A flock pattern query $Flock(\mu, \epsilon, \delta)$ is defined as follows:

Definition 1. *Given are a set of trajectories $\mathcal{T}$, a minimum number of trajectories $\mu > 1$ ($\mu \in \mathbb{N}$), a maximum distance $\epsilon > 0$ defined over the distance function d, and a minimum time duration $\delta > 1$ ($\delta \in \mathbb{N}$). A flock pattern $Flock(\mu, \epsilon, \delta)$ reports all maximal size collections $\mathcal{F}$ of trajectories where: for each f_k in $\mathcal{F}$, the number of trajectories in f_k is greater or equal than μ ($|f_k| \geq \mu$) and there exist δ consecutive time instances such that for every $t_i \in [f_k^{t_1} .. f_k^{t_1+\delta}]$, there is a disk with center $c_k^{t_i}$ and radius $\epsilon/2$ covering all points in $f_k^{t_i}$. That is: $\forall f_k \in \mathcal{F}, \forall t_i \in [f_k^{t_1} .. f_k^{t_1+\delta}], \forall T_j \in f_k : |f_k^{t_i}| \geq \mu, d(p_j^{t_i}, c_k^{t_i}) \leq \epsilon/2$*

The $c_k^{t_i}$ is called the center of the flock f_k at time t_i.

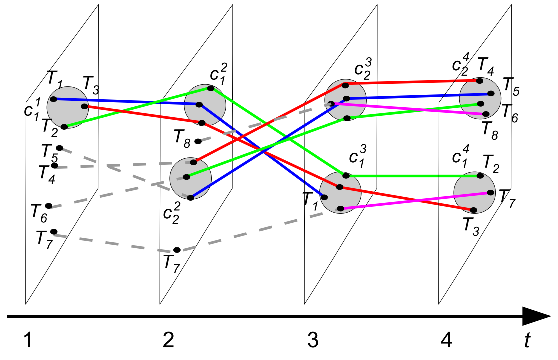

Figure 3: Flock pattern example

In the above definition, a flock pattern can be viewed as a "tube" shape formed by the centers c and expanded with diameter ϵ, and having length δ (consecutive time instants) such that there are at least μ trajectories which stay inside the tube all the time, as shown in Figure 3. For $Flock(\mu = 3, \epsilon, \delta = 3)$, the flocks $\mathcal{F}$ reported are $f_1 = \{T_1, T_2, T_3\}$ (from time instance t_1 to t_3 and disks c_1^1, c_1^2, and c_1^3) and $f_2 = \{T_4, T_5, T_6\}$ from (from time instance t_2 to t_4 and disks c_2^2, c_2^3, and c_2^4).

Having this formal definition we proceed with the complexity analysis of the flock pattern. The major challenge in this type of queries is the fact that the center of the flock pattern $c_k^{t_i}$ may not belong to any of the trajectories. Hence we cannot iterate over the discrete number of trajectory locations stored in the database and check if each one of them is a center of a flock or not. Since any point in the spatial domain can be a center of a flock there is an infinite number of possible locations to test.

Nevertheless, we show using the following theorem 1 that there is a limited and discrete number of locations where we can look for flocks among the infinite number of options.

Theorem 1. *If for a given time instance t_i there exists a point in the space $c_k^{t_i}$ such that:*

$$\forall T_j \in f, d(p_j^{t_i}, c_k^{t_i}) \leq \epsilon/2$$

then there exists another point in the space $c'^{t_i}_k$ such that

$$\forall T_j \in f, d(p_j^{t_i}, c'^{t_i}_k) \leq \epsilon/2$$

and there are at least trajectories $T_a \in f$ and $T_b \in f$ such that

$$\forall T_j \in \{T_a, T_b\}, d(p_j^{t_i}, c'^{t_i}_k) = \epsilon/2$$

Theorem 1 states that if there is a disk $c_k^{t_i}$ with diameter ϵ that covers all trajectories in the flock f at time instance t_i then there exists another disk with the same diameter but with different center $c'^{t_i}_k$ that also covers all trajectories covered by the first one and has at least two common points on its circumference. Theorem 1 can be easily proved by construction.

Proof Sketch. Assume that we have a disk with diameter ϵ and center c_k that covers all trajectories in the flock at given time instance t_i as shown in Figure 4(a). Assume for simplicity that there is no trajectory point on the circumference of the disk defined by c_k and ϵ, i.e. $\forall T_j \in f, d(T_j, c_k) < \epsilon/2$. We can find another disk with the same properties but with different center by using a combination of translation

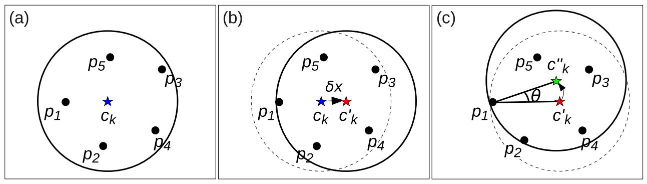

Figure 4: Finding disks to cover set of points

and rotation of the disk with center c_k. As a first step of the construction the center of the disk c_k is moved along the x axis until the first of the trajectory points inside lies on the circumference of the disk. For example in Figure 4(b) the first point which falls on the circumference after the horizontal move of the disk center is p_1. The new center of the disk is point c'_k. All points in the flock are covered by the new disk with center c'_k and diameter ϵ. Otherwise, there would be a contradiction to the assumption that p_1 is the first point on the circumference. The next step of the construction rotates the new disk using as pivot the first point on the circumference (p_1). The disk is rotated until another point falls on its circumference. In the example of Figure 4(c) the disk is rotated until point p_2 is on the circumference of disk c''_k. All points in the flock are still covered by the new disk with center c''_k and diameter ϵ (otherwise there will be a contradiction to the assumption that p_2 is the first one to be on the circumference of the disk during the rotation process). The new disk c''_k has at least two points on its circumference (points p_1 and p_2) $\square$

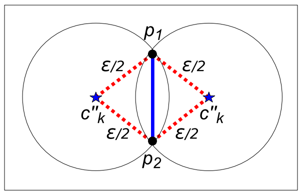

Figure 5: Disks for $\{p_1, p_2\}$, $d(p_1, p_2) \leq \epsilon$

Theorem 1 has great impact on the search for flock patterns because it limits the number of locations inside the spatial domain where to look for flocks. For a database of $|\mathcal{T}|$ trajectories there are $|\mathcal{T}|^2$ possible pairs of point combinations at any given time instance. For each such pair there are exactly two disks with radius $\epsilon/2$ that have those points on their circumference (Figure 5). We test those disks to find if they have the required minimum number of μ trajectories inside. For each time instance of the time-interval δ we have to perform $2|\mathcal{T}|^2$ tests for flock pattern. The total number of possible flock patterns that need to be tested is $2|\mathcal{T}|^{2\delta}$. In order to solve the problem, the algorithm has to not only consider each such sequence of disks (a possible flock pattern), but also to identify the trajectories that match it. The check if the trajectory stays within the sequence of disks can be done in $O(\delta)$ time. For the whole database it takes $O(|\mathcal{T}|\delta)$ time, and the total running time of the algorithm will be $O(|\mathcal{T}|^{(2\delta)}|\mathcal{T}|\delta) = O(\delta|\mathcal{T}|^{(2\delta)+1})$. As a result, the flock problem with fixed time duration has polynomial time complexity $O(\delta|\mathcal{T}|^{(2\delta)+1})$.

4. REPORTING FLOCK PATTERNS

In this section we describe a grid-based structure and some optimizations in order to efficiently compute flock disks and report flocks. We also describe five *on-line* algorithms to process spatio-temporal data in an incremental fashion.

The grid-based structure employed for all proposed algorithms is based on grid cells with edges of ϵ distance. Each trajectory location $p_{id}^{t_i}$ reported for a specific time instance t_i is inserted in a specific grid cell. The cell is determined by its components' location latitude and longitude. Thus, each location is inserted in only one cell. The total number of cells in the index is thus affected by the trajectory distribution in the each specific time instance t_i and the ϵ. The smaller the value of ϵ, the larger number of grid cells are needed. In our implementation, grid cells that are empty, i.e. there is no trajectory location in them, are not allocated. Other structures, e.g. *k-d-trees*, could be employed for organizing all trajectory locations in each cell grid. However, since for small ϵ the number of locations within each cell is relatively small, and given its access simplicity we used a list for each cell. The organization of this index is shown in Figure 6.

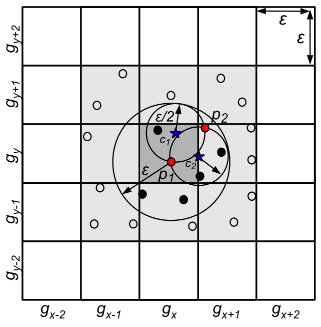

Figure 6: A grid-based index example.

Once the grid structure is built for t_i, disks can be processed using the Algorithm 1. For each grid cell $g_{x,y}$, only the 9 adjacent grid cells, including itself, are analyzed. Algorithm 1 first process every point in $g_{x,y}$ and every point in $[g_{x-1,y-1}...g_{x+1,y+1}]$ in order to find pair of points p_r, p_s whose distances satisfy: $d(p_r, p_s) \leq \epsilon$. Because all cells in the grid index have ϵ distance, there is no need to analyze points further away of the range $[g_{x-1,y-1}...g_{x+1,y+1}]$ cells for points in a particular cell $g_{x,y}$. Pairs that have not been processed yet and are within ϵ to each other are further used to compute the two disks c_1 and c_2. In case that the pairs are exactly at distance $d(p_r, p_s) = \epsilon$, c_1 and c_2 have the same center and only one has to be further processed.

It should be noted that not all points in $[g_{x-1,y-1}...g_{x+1,y+1}]$ have to be "paired" with each point in $g_{x,y}$: only those that have distance $d(p_r, p_s) \leq \epsilon$ (as illustrated in Figure 6). Another optimization involves the points that have to be checked whether they are inside each disk computed in the previous step. Figure 7 illustrates these situations. For each point $p_r \in g_{x,y}$ (point p_1 in Figure 7(a)), a range query with radius ϵ is performed over all 9 grids $[g_{x-1,y-1}...g_{x+1,y+1}]$ to find points that can be "paired" with p_r, that is $d(p_r, p_s) \leq \epsilon$ holds. The result of such range search with more or equal points than μ ($|\mathcal{H}| \geq \mu$) is stored in the list $\mathcal{H}$ that is used to check for each disk computed. For those valid pairs, at most 2 disks are generated. For each of them, points in the list $\mathcal{H}$ are checked if they are inside the disk (Figure 7(b)). Disks that have less than μ points are further discarded and only the ones that $|c_k| \geq \mu$ holds are kept. In Figure 7(c) disk c_1 is discarded and c_2 is considered a valid disk. Because we are interested only in maximal instances of flock patterns, a

Algorithm 1: Computing disks in grid-based index

```
input  : set of points $\mathcal{T}[t_i]$ for timestamp $t_i$
output: sets of maximal disks $\mathcal{C}$
$\mathcal{C} \leftarrow \emptyset$
$Index.Build(\mathcal{T}[t_i], \epsilon)$
for each non-empty cell $g_{x,y} \in Index$ do
    $P_r \leftarrow g_{x,y}$
    $P_s \leftarrow [g_{x-1,y-1} \cdots g_{x+1,y+1}]$
    if $|P_s| \geq \mu$ then
        for each $p_r \in P_r$ do
            $\mathcal{H} \leftarrow Range(p_r, \epsilon)$, $|\mathcal{H}| \geq \mu$, $d(p_r, p_s) \leq \epsilon, p_s \in P_s$
            for each $p_j \in \mathcal{H}$ do
                if $\{p_r, p_j\}$ not yet computed then
                    compute disks $\{c_1, c_2\}$ defined by $\{p_r, p_j\}$
                    and diameter $\epsilon$
                    for each disk $c_k \in \{c_1, c_2\}$ do
                        $c \leftarrow c_k \cap \mathcal{H}$
                        if $|c| \geq \mu$ then $\mathcal{C}.Add(c)$
                    end
                end
            end
        end
    end
end
```

valid disk is further checked whether another disk has a superset of instances that the current disk has just computed. In this particular case, disks that have subset of instances are discarded and only those ones stored in $\mathcal{C}$ that have the maximal instances are returned by Algorithm 1.

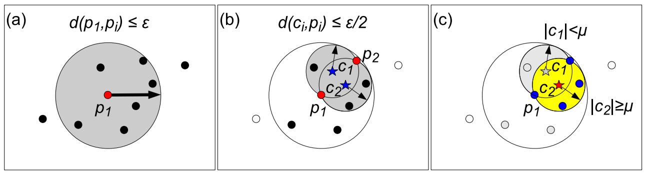

Figure 7: Steps on finding flocks for time t

The process that Algorithm 1 employs to keep only the maximal disks is based on the center of the disk and the total number of common elements that each disk has. Disks are checked only with the ones that are "close" to each other, that is, disk c_1 is checked with c_2 only if $d(c_1, c_2) \leq \epsilon$. If $d(c_1, c_2) > \epsilon$, we can safely state that they do not have any elements in common. To efficient process the operations described above, we store disks in $\mathcal{C}$ using a *k-d-tree* where the center of each disk along with its radius $\epsilon/2$ are stored. When checking for a particular entry c_1, we only need to check entries in the *k-d-tree* that "intersect" with the new one. Only those disks that cannot be pruned are further verified to check their contents. Because we store entries that belong to each disk in a binary tree, we can efficiently check if one disk has supersets/subsets elements than the other disk. Therefore, we only need to count common elements in both disks by scanning each entry in each disk once. If the cardinality of common elements are $|c_1 \cap c_2| = |c_1|$ then c_1 is subset of c_2 disk, or they have all common elements when $|c_1| = |c_2|$. Therefore, c_1 can be discarded and only c_2 is kept in $\mathcal{C}$. When $|c_1 \cap c_2| = |c_2|$, c_2 can be discarded. Otherwise we can safely say that one is not maximal than the other disk and we have to keep both c_1 and c_2 in $\mathcal{C}$.

In the following subsection we describe the basic flock pattern evaluation algorithm which combines the candidate disks generated for each time instance into flock patterns. Later in this section we describe four variations of the basic algorithm which use different filtering heuristics in order to reduce the number of candidate disks which have to be analyzed.

4.1 The Basic Flock Evaluation Algorithm

In the basic flock pattern evaluation algorithm *BFE*, we generate the candidate disks for every time instance t_i, starting with the first one t_1 and moving one time instance at a time. Every candidate disk generated in given time instance t_i is analyzed and joined with potential flocks generated in the previous time instance t_{i-1}. Only those potential flocks that are successfully augmented with disk in the current time instance are kept for further processing in the next time instance. This method reports flock patterns as soon as they satisfy the temporal constraint δ (e.g. we have at least δ candidate disks successfully joined in a flock).

As it was mentioned in previous section, we use a grid-based index to find disks for the current time instance t_i. For the first time instance t_1, all disks returned by the grid-based index are stored as potential flocks (we can view a candidate disk as a partial flock with length 1) in the list of candidate flocks for this time instance $\mathcal{F}^{t_i}$. In the following time instances all disks returned by the grid-based index are stored in their candidate flock lists $\mathcal{F}^{t_i}$ and then "joined" with the candidate flocks from the previous time instance $\mathcal{F}^{t_{i-1}}$. The "join" condition used for this operation is $|c \cap f| \geq \mu$, i.e. the total number of common elements between the candidate flock and the disk has to be greater or equal to μ in order to be joined. If the condition is satisfied then we move the join result into the list of candidate flocks for the current time instance t_i. A flock is found if there are at least δ join operations applied over the candidate flock, i.e. $u.t_{end} - u.t_{start} = \delta$. In this case, the flock pattern is immediately reported to the user and its $u.t_{start}$ attribute is updated and reinserted in $\mathcal{F}^{t_i}$ to be further joined with other disks in the following time instance.

It should be noted that $\mathcal{F}^{t_i}$ only maintains potential flocks starting at some previous time instance $t_{start} > t_i - \delta$ and ending in the current time instance $t_{end} = t_i$. Entries that cannot be joined in the next time instance are discarded and not transferred into the list of candidate flocks for the next time instance.

Algorithm 2: *BFE*: Basic Flock Evaluation

```
input  : parameters $\mu$, $\epsilon$ and $\delta$
output: flocks patterns
$\mathcal{F}^{t_0} \leftarrow \emptyset$
for each new time instance $t_i$ do
    $\mathcal{F}^{t_i} \leftarrow \emptyset$,   $\mathcal{C} \leftarrow Index.Disks(\mathcal{T}[t_i])$
    for each $c \in \mathcal{C}$ do
        for each $f \in \mathcal{F}^{t_{i-1}}$ do
            if $|c \cap f| \geq \mu$ then
                $u \leftarrow c \cap f$
                $u.t_{start} \leftarrow f.t_{start}$
                $u.t_{end} \leftarrow t_i$
                if $(u.t_{end} - u.t_{start}) = \delta$ then
                    report flock pattern $u$ from $u.t_{start}$ to
                    $u.t_{end}$
                    update $u.t_{start}$
                end
                $\mathcal{F}^{t_i} \leftarrow \mathcal{F}^{t_i} \cup u$
            end
        end
        $\mathcal{F}^{t_i} \leftarrow \mathcal{F}^{t_i} \cup c$
    end
end
```

One advantage of the *BFE* Algorithm is that for each time instance being processed, the algorithm stores only the trajectory *ids* in $\mathcal{F}^{t_i}$. There is no need to keep the actual locations of moving objects in $\mathcal{F}^{t_i}$ since they do not participate in the join condition. Another advantage is that trajectory locations for each time instance are processed only once, that is, there is no need to buffer trajectory data for a time window with length δ like our other algorithms explained later in this section.

4.2 Filtering Heuristics

The number of candidate disks in a given time instance can be quite large and the cost to join those candidate disks in a flock pattern can be quite expensive. In order to improve the performance of the *BFE* algorithm we propose a set of four different heuristics used to limit the number of generated candidate disks. These heuristics are described next.

4.2.1 Top Down Evaluation

The first heuristic is a "Top Down Evaluation" (*TDE*). It differs from the basic algorithm in the fact that the construction of the flocks is not done in a *bottom-up* approach (by extending flock patterns one candidate disk at a time until they become at least δ time instances long) but in a *top-down* fashion. Here we compare the candidate disks for time instances which are δ time instances apart. This is based on the assumption that the difference between the candidate disks in two consecutive time instances will be small (thus resulting in a large number of short flocks which still have to be kept as candidates until it becomes clear that they do not have the required length), while the differences between candidate disks from time instances which are δ time instances apart will be significant (and will result in smaller set of candidate flocks).

This heuristic buffers trajectory locations for time window w which has length δ time instances. It also performs a different strategy on joining the candidate disks in this time window w. First the algorithm calculates the candidate disks $\mathcal{C}^1$ for the first time instance $t_{i-\delta+1}$ in the window w. Then, disks for the last time instance t_i in w are calculated and joined with the ones in $\mathcal{C}^1$. The candidate flocks for time window w generated as a result of this step are then verified using the basic flock pattern evaluation algorithm.

Algorithm 3: *TDE*: Top Down Evaluation

```
for each new time instance $t_i$ do
    let $\mathcal{L}$ be trajectories in windows size $|w| = \delta$ $(t_{i-\delta}...t_i)$
    $\mathcal{F} \leftarrow \emptyset, \mathcal{U} \leftarrow \emptyset$
    $\mathcal{C}^1 \leftarrow Index.Disks(\mathcal{L}[1]),\quad \mathcal{C}^w \leftarrow Index.Disks(\mathcal{L}[w])$
    for each $c^1 \in \mathcal{C}^1$ do
        for each $c^w \in \mathcal{C}^w$ do
            if $|c^1 \cap c^w| \geq \mu$ then $\mathcal{U} \leftarrow \mathcal{U} \cup \{c^1 \cap c^w\}$
        end
    end
    for each $u \in \mathcal{U}$ do
        $\mathcal{L}' \leftarrow u,\quad \mathcal{F}^1 \leftarrow u^1$
        for $t \leftarrow 2$ to $|w| - 1$ do
            $\mathcal{F}^t \leftarrow \emptyset,\quad \mathcal{C}^t \leftarrow Index.Disks(\mathcal{L}'[t])$
            for each $c \in \mathcal{C}^t$ do
                for each $f \in \mathcal{F}^{t-1}$ do
                    if $|c \cap f| \geq \mu$ then $\mathcal{F}^t \leftarrow \mathcal{F}^t \cup \{c \cap f\}$
                end
            end
            if $|\mathcal{F}^t| = 0$ then break
        end
        for each $f \in \mathcal{F}^{w-1}$ do
            for each $c^w \in \mathcal{C}^w$ do
                if $|f \cap c^w| \geq \mu$ then $\mathcal{F} \leftarrow \mathcal{F} \cup \{f \cap c^w\}$
            end
        end
    end
    report flocks $\mathcal{F}$
end
```

4.2.2 The Pipe Filter Evaluation

Our second heuristic, the *Pipe Filtering Evaluation* (*PFE*), also employs the *filter-and-refine* paradigm. It first filters all trajectories that have at least μ objects within distance ϵ of them for a duration of at least δ time instances. Then in a refinement step performed over the trajectories returned by the filtering step we search for flock patterns using the basic flock pattern evaluation algorithm. Figure 8 illustrates a pipe for trajectory T_2 with radius ϵ. Trajectories $\{T_1, T_2, T_3, T_4\}$ are in the pipe for all δ times stamps and are further processed in the refinement step.

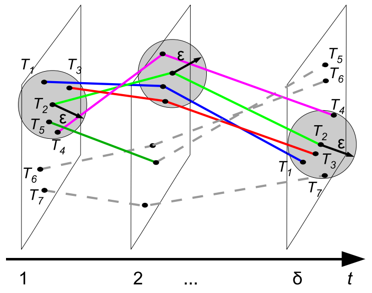

Figure 8: Pipe filtering δ for T_2 and radius ϵ.

The Pipe Filtering algorithm, first builds a grid-based index for the first time instance $t_{i-\delta}$ in the w window. Then, for each trajectory location T_j in $t_{i-\delta}$ a range search is issued. The purpose of this range query is to examine how many other object locations are within distance ϵ from the trajectory being processed. If the cardinality of the result set is greater or equal than the threshold μ, then we continue with the same check for time instances $t_{i-\delta+1}$ to t_i. If the total number of trajectories inside the "pipe" for given trajectory T_j is $|\mathcal{U}| \geq \mu$, then the set of trajectories qualifies and it is stored in the list of candidates $\mathcal{M}$, to be further processed in the refinement step of the algorithm.

The refinement step employs the basic flock pattern evaluation algorithm. The difference however is that now it process only the trajectory locations returned as a result of the filtering step $\mathcal{M}$ instead of using the whole trajectory database. This approach is beneficial in cases where a large number of trajectories will be pruned by the pipe filtering step and the computationally expensive candidate disk generation and flock construction will be performed over a limited subset of trajectories $m \in \mathcal{M}$.

4.2.3 The Continuous Refinement Evaluation

As the name implies, the *Continuous Refinement Evaluation* (*CRE*) heuristic continuously refines the set of trajectories which can participate in a flock pattern. This approach uses the candidate disk generation step for time instance t_i as a filtering step for time instance t_{i+1}. Only trajectories that are associated with the candidate disk in time t_i are analyzed in t_{i+1}. This approach can be used in cases where the selectivity of the candidate disks is high, e.g. there exists

Algorithm 4: *PFE*: Pipe Filter Evaluation

```
for each new time instance $t_i$ do
    let $\mathcal{L}$ be trajectories in windows size $|w| = \delta$, $(t_{i-\delta} \dots t_i)$
    $\mathcal{F} \leftarrow \emptyset$
    for each $T_j \in \mathcal{L}$ do
        $\mathcal{L}' \leftarrow Index.Range(T_j, \epsilon)$
        if $|\mathcal{L}'| \geq \mu$ then
            $\mathcal{U} \leftarrow \emptyset$
            for each $T_k \in \mathcal{L}'$ do
                if $\forall t_i \in w, p_k^{t_i} \in T_k, p_j^{t_i} \in T_j, d(p_k^{t_i}, p_j^{t_i}) \leq \epsilon$
                then $\mathcal{U} \leftarrow \mathcal{U} \cup T_k$
            end
            if $|\mathcal{U}| \geq \mu$ then $\mathcal{M} \leftarrow \mathcal{M} \cup \mathcal{U}$
        end
    end
    for each $m \in \mathcal{M}$ do
        $\mathcal{F}^1 \leftarrow Index.Disks(m^1)$
        for $t \leftarrow 2$ to $|w|$ do
            $\mathcal{F}^t \leftarrow \emptyset$,   $\mathcal{C} \leftarrow Index.Disks(m^t)$
            for each $c \in \mathcal{C}$ do
                for each $f \in \mathcal{F}^{t-1}$ do
                    if $|c \cap f| \geq \mu$ then $\mathcal{F}^t \leftarrow \mathcal{F}^t \cup \{c \cap f\}$
                end
            end
            if $|\mathcal{F}^t| = 0$ then break
        end
        $\mathcal{F} \leftarrow \mathcal{F} \cup \mathcal{F}^t$
    end
    report flocks $\mathcal{F}$
end
```

a relatively small number of candidate disks and the number of trajectories in them is low.

In its first step, the *CRE* algorithm finds disks $\mathcal{C}^1$ using locations $\mathcal{L}[1]$ for time instance $t_{i-\delta}$. Then, for each disk $c^1 \in \mathcal{C}^1$, all trajectories associated with it are further processed from time instance $t_{i-\delta+1}$ to t_i.

At the first time instance, disks $\mathcal{C}^1$ for time instance $t_{i-\delta}$ are stored in $\mathcal{F}^1$ (potential flocks of length 1). Then, each instance of c^1 is further processed to compute disks and is "merge-joined" with the previous ones stored in $\mathcal{F}^t$. If $\mathcal{F}^t$ has no potential flock at time t, then the processing of c^1 can be discarded. After this second step, flock patterns are reported from time $t_{i-\delta}$ to t_i.

Algorithm 5: *CRE*: Continuous Refinement Evaluation

```
for each new time instance $t_i$ do
    let $\mathcal{L}$ be trajectories in windows size $|w| = \delta$, $(t_{i-\delta} \dots t_i)$
    $\mathcal{F} \leftarrow \emptyset$,   $\mathcal{C}^1 \leftarrow Index.Disks(\mathcal{L}[1])$
    for each $c^1 \in \mathcal{C}^1$ do
        let $\mathcal{L}'$ be the trajectories in $c^1$ with length $w$
        $\mathcal{F}^1 \leftarrow c^1$
        for $t \leftarrow 2$ to $|w|$ do
            $\mathcal{F}^t \leftarrow \emptyset$,   $\mathcal{C}^t \leftarrow Index.Disks(\mathcal{L}'[t])$
            for each $c \in \mathcal{C}^t$ do
                for each $f \in \mathcal{F}^{t-1}$ do
                    if $|c \cap f| \geq \mu$ then $\mathcal{F}^t \leftarrow \mathcal{F}^t \cup \{c \cap f\}$
                end
            end
            if $|\mathcal{F}^t| = 0$ then break
        end
        $\mathcal{F} \leftarrow \mathcal{F} \cup \mathcal{F}^t$
    end
    report flocks $\mathcal{F}$
end
```

4.2.4 The Cluster Filtering Evaluation

The last proposed heuristic, *Cluster Filtering Evaluation* (*CFE*), has two phases: **(1)** the first phase applies the

Algorithm 6: *CFE*: Clustering Filtering Evaluation

```
$\mathcal{I}^{t_i} \leftarrow \emptyset$
for each new time instance $t_i$ do
    $\mathcal{U} \leftarrow \emptyset$,   $\mathcal{L} \leftarrow \mathcal{T}[t_i]$
    $\mathcal{Q} \leftarrow DBSCAN(\mathcal{L}, \mu, \epsilon)$
    for each $q \in \mathcal{Q}$ do
        for each $f \in \mathcal{I}^{t_{i-1}}$ do
            if $|q \cap f| \geq \mu$ then
                $u \leftarrow \{q \cap f\}$
                $u.t_{start} \leftarrow f.t_{start}$
                $u.t_{end} \leftarrow t$
                if $(u.t_{end} - u.t_{start}) = \delta$ then
                    $\mathcal{F}^1 \leftarrow Index.Disks(u^1)$
                    for $t \leftarrow 1$ to $|w|$ do
                        $\mathcal{F}^t \leftarrow \emptyset$,   $\mathcal{C} \leftarrow Index.Disks(m^t)$
                        for each $c \in \mathcal{C}$ do
                            for each $f \in \mathcal{F}^{t-1}$ do
                                if $|c \cap f| \geq \mu$ then
                                $\mathcal{F}^t \leftarrow \mathcal{F}^t \cup \{c \cap f\}$
                            end
                        end
                        if $|\mathcal{F}^t| = 0$ then break
                    end
                    $\mathcal{F} \leftarrow \mathcal{F} \cup \mathcal{F}^t$
                    update $u.t_{start}$
                end
                $\mathcal{U} \leftarrow \mathcal{U} \cup u$
            end
        end
        $\mathcal{U} \leftarrow \mathcal{U} \cup q$
    end
    $\mathcal{I}^{t_i} \leftarrow U$
end
```

DBSCAN clustering algorithm with parameters *eps*=ϵ and *minPts*=μ for each time instance t_i; **(2)** clusters reported for a given time instance t_i by the *DBSCAN* algorithm are further joined with clusters found in the previous time instance t_{i-1}. The join criteria is that the clusters should have at least μ trajectories in common, and then only the resulted intersection of trajectories are maintained. If a cluster u can be augmented in this way for δ consecutive time instances ($u.t_{end} - u.t_{start} = \delta$), then it is saved as a candidate which has to be verified in the second phase, using the basic flock pattern evaluation algorithm.

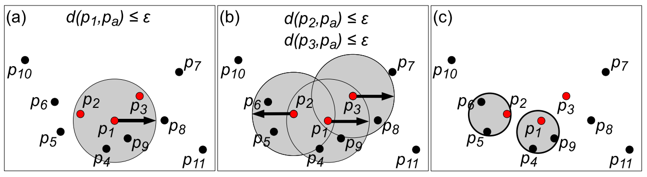
(a)
$d(p_1,p_a) \leq \varepsilon$
(b)
$d(p_2,p_a) \leq \varepsilon$
$d(p_3,p_a) \leq \varepsilon$
(c)
p_1
p_2
p_3
p_4
p_5
p_6
p_7
p_8
p_9
p_{10}
p_{11}

Figure 9: *CFE* steps to find flock patterns

Figure 9 illustrates the steps performed by the *CFE* algorithm. In Figure 9(a), the *DBSCAN* is applied to a specific object location p_1 with parameters *eps*=ϵ and *minPts*=μ. Then, in Figure 9(b), shows the propagation of the *DBSCAN* algorithm over p_1's neighbors. Object locations that do not belong to any cluster are discarded. The final two clusters reported by the *DBSCAN* algorithm in Figure 9(c) ($\{p_2, p_5, p_6\}$ and $\{p_1, p_4, p_9\}$) are then further processed in the refinement step of the *CFE* Algorithm.

5. EXPERIMENTAL RESULTS

In order to evaluate the performance of the proposed methods, we run several sets of experiments with various trajectorial datasets and using different flock pattern parameters. In

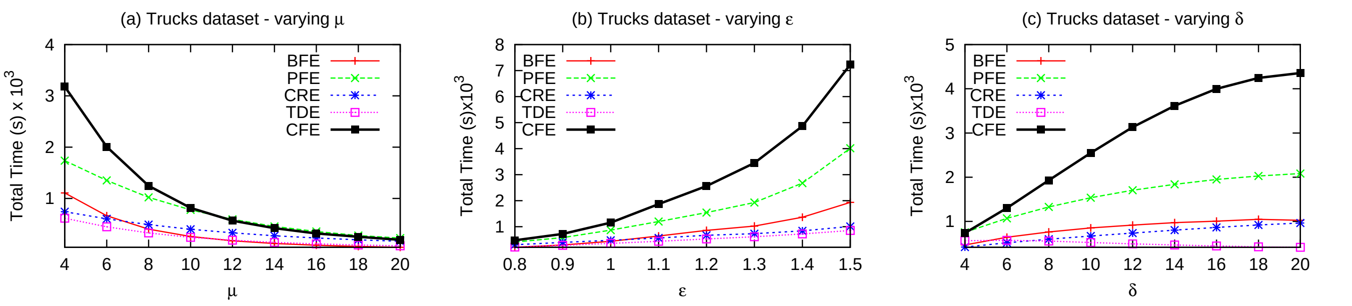

Figure 10: Total time (s) when varying (a) μ, (b) δ and (c) ϵ for the *Trucks* dataset

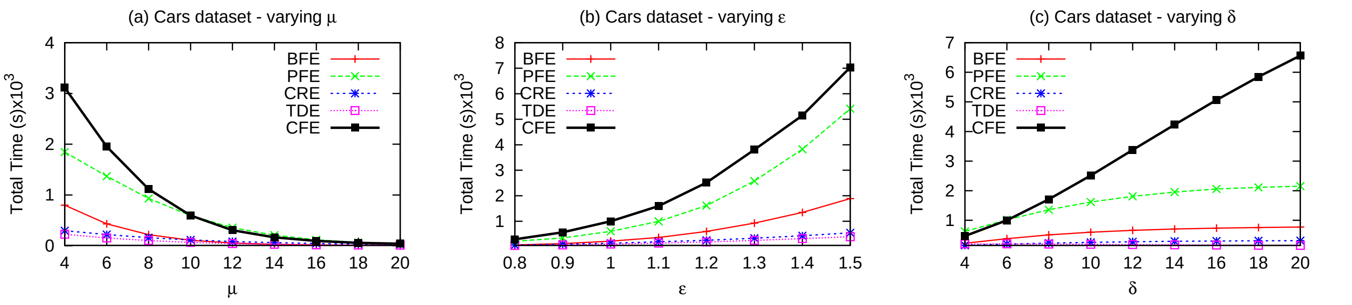

Figure 11: Total time (s) when varying (a) μ, (b) δ and (c) ϵ for the *Cars* dataset

particular we show the results for four real datasets – *Trucks*, *Buses*, *Cars* and *Caribous* – and one synthetic dataset – *SG*. The first two real datasets *Trucks* and *Buses* [26] contain 112,203 and 66,096 moving object locations generated from 276 and 145 moving trucks and buses, respectively, collected in the great metropolitan area of Athens, Greece. The third dataset *Cars* [13] contains 134,263 object locations collected from 183 private cars moving in Copenhagen, Denmark. The last real trajectorial dataset *Caribous* [24] is generated from the analysis of the migration movements of 43 caribous in northwestern states of Canada. The size of the dataset is 15,796 object locations.

Because the real datasets that we could find in public domain are relatively small, in order to test the scalability of our methods we use synthetic dataset *SG* as well. This dataset is generated by simulating the movement of 50,000 vehicles on road network of Singapore. Those moving objects have different velocities and their starting locations were randomly placed in the road network. The size of the synthetic dataset is 2,548,084 moving object locations.

In our experiments we use several values for the flock parameters μ, ϵ and δ. The ranges of values for each dataset are shown in Table 1, where in bold we show the default values for each parameter. Those default values are used when the value of the parameter is not explicitly specified for a given experimental set. The total number of patterns discovered for the minimum and maximum value of each parameter, taken from Table 1, are shown in Table 2.

Figures 10-14 show the results when varying parameters μ (first column), ϵ (second column) and δ (third column) respectively, for the different datasets. All plots show the total time in seconds, needed to process the whole dataset, including building the grid index. As it can be seen, when increasing μ, decreasing ϵ, or decreasing δ, the total time needed to discover the flock patterns for all methods decreases. This is expected since the flock queries becomes

Table 1: Parameters values for each dataset

Dataset	μ **[default]**	ϵ **[default]**	δ **[default]**
	min #traj.	*max* dist.	*min* time
Trucks	4, 6,...,20 [**5**]	0.8, 0.9,...,1.5 [**1.2**]	4, 6,...,20 [**10**]
Cars	4, 6,...,20 [**5**]	0.8, 0.9,...,1.5 [**1.2**]	4, 6,...,20 [**10**]
Caribous	2, 3,...,10 [**5**]	0.1, 0.2,...,0.8 [**1.6**]	4, 6,...,20 [**10**]
Buses	4, 6,...,20 [**5**]	0.4, 0.5,...,1.1 [**1.2**]	4, 6,...,20 [**10**]
SG	4, 6,...,20 [**5**]	2.2, 2.6,...,5.0 [**3.4**]	4, 6,...,20 [**10**]

more selective and we have to maintain fewer candidate trajectories during the query evaluation.

For the *Trucks* and *Cars* datasets, the *TDE* and *CRE* methods have significantly better performance compared to the other methods. The gap between those methods and the rest increases when the selectivity of the queries becomes low (for small μ and big ϵ). This is due to the large number of partial intermediate results which have to be maintained by the other three methods (*BFE*, *PFE* and *CFE*) and the increase of the total time needed to process those partial results. Similar behavior can be observed for large values of the flock duration δ, but only for the *PFE* and *CFE* methods. This is due to the fact that these two methods keep the trajectory history in a time window w before computing the disks for each timestamp. Similar behaviors are found for the *Buses* dataset.

For the *Caribous* dataset the *BFE* algorithm has the best performance, closely followed by the *TDE* and *CRE*. Our analysis shows that the *BFE* algorithm performed well in this dataset because the data characteristics. The data in this dataset seems to be very well correlated, e.g. all 43 caribous have similar migration patterns and stay close together, well grouped in herds during their movement. Because of this other methods are not able to prune a lot of trajectories in their filtering phases. The fact that the data in the *Caribous* dataset is well clustered in terms of flocks

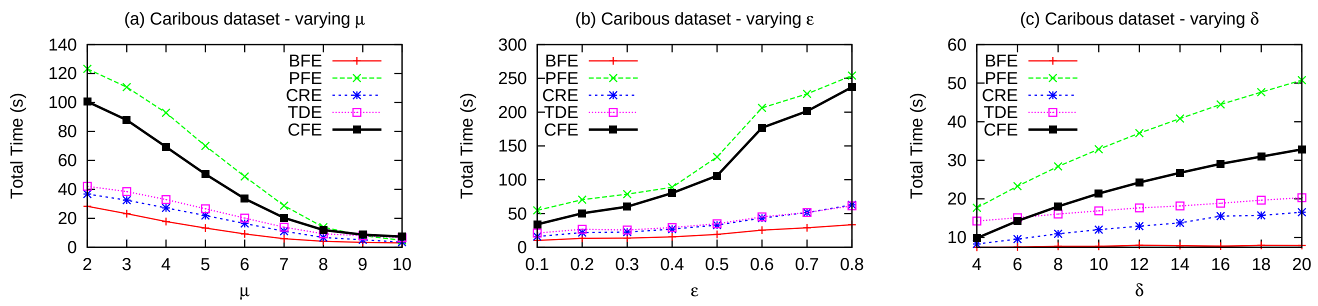

Figure 12: Total time (s) when varying (a) μ, (b) δ and (c) ϵ for the *Caribous* dataset

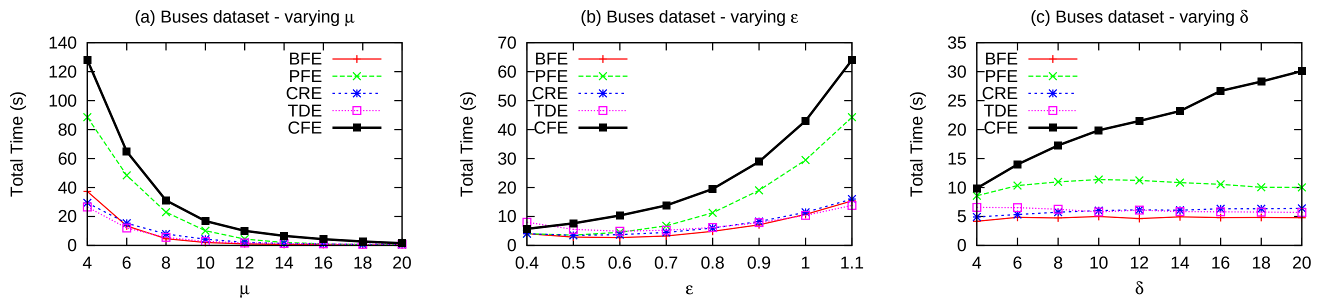

Figure 13: Total time (s) when varying (a) μ, (b) δ and (c) ϵ for the *Buses* dataset

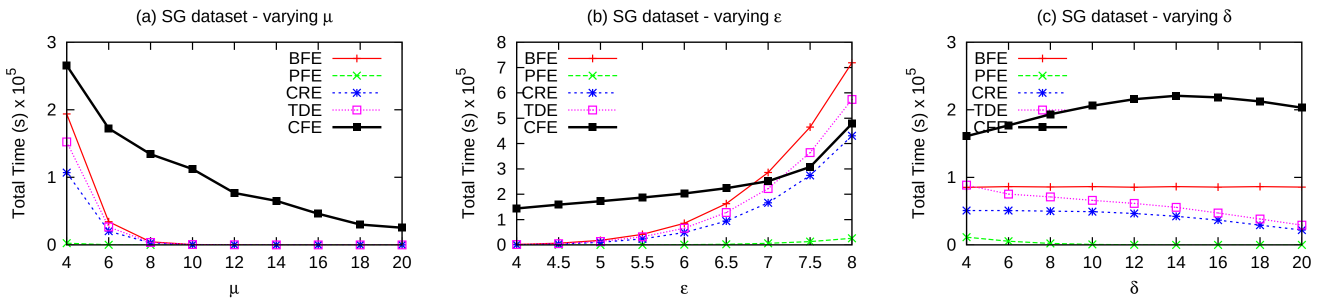

Figure 14: Total time (s) when varying (a) μ, (b) δ and (c) ϵ for the *SG* dataset

can also be observed by the total number of flocks discovered for this dataset (see Table 2). The number of discovered flocks is quite large for a dataset with only 15,796 moving object locations.

In the next group of experiments we use the synthetic dataset *SG*. As it can be seen from the plots the *PFE* algorithm is by far the best algorithm for this particular dataset. The main reason for this behavior is that even though the total number of potential flocks for each timestamp can be quite big (see Table 3 for details), this approach performs more holistic filtering compared to the other solutions. This heuristic is checking the minimum occupancy criteria (the number of trajectories closer than the threshold ϵ to a given

Table 2: Number of flock patterns discovered

Dataset	μ - *min* #traj.		ϵ - *max* dist.		δ - *min* time	
	min	*max*	*min*	*max*	*min*	*max*
Trucks	309	14,935	3,741	15,608	2,045	23,222
Cars	62	18,451	3,218	23,440	3,149	24,211
Caribous	124	9,480	5,292	6,915	3,364	4,598
Buses	0	2,988	16	1,021	55	1,730
SG	0	1,304	53	741	112	385

trajectory should be more than μ) for the whole duration of the flock pattern δ. The other four methods try to join candidate disks for two consecutive timestamps, without having holistic view of the trajectories. Therefore, the first filtering phase of the *PFE* has a higher pruning capability compared to the other methods for the *SG* dataset. We should note that in the real datasets, trajectories follow similar patterns, while in the *SG* dataset objects follow random patterns and though they might be close together in one time instance, they tend not to follow similar patterns for several consecutive timestamps.

As it can be seen from the plots, for most of the datasets the *CFE* algorithm has the worst performance among all methods. This is due to the fact that the filtering step in this approach employs clustering which can be very expensive for large datasets. This approach however works significantly better when the datasets are relatively small and the moving objects in those datasets have similar moving patterns (see the results for the *Caribous* dataset). In scenarios like those, the cost for clustering is not that high which explains the improved performance.

In our next set of experiments, we measure the minimum and the maximum number of disks computed for each time instance using our grid-based index. The results can be depicted in Table 3. As it can be seen even for big values of the parameters μ, ϵ and δ, the maximum number of disks computed per timestamp is relatively small compared with the number of trajectories. This shows the efficiency of our grid-based index structure.

Table 3: Min/Max Number of disks per time

Dataset	μ - *min* #traj.		ϵ - *max* dist.		δ - *min* time	
	min	*max*	*min*	*max*	*min*	*max*
Trucks	505	1,257	812	1,547	1,237	1,237
Cars	72	294	142	387	279	279
Caribous	393	235	587	342	309	309
Buses	7	236	27	183	105	105
SG	1,343	12,894	1,232	2,916	10,934	10,934

6. CONCLUSIONS

Recently there has been increased interest in queries that capture the collaborative behavior of spatio-temporal data (e.g. convoys, flocks, etc.). In particular, a flock contains a group of at least μ moving objects all of them "enclosed" by a disk of diameter ϵ for at least δ consecutive time periods. Discovering flock patterns on line is useful for several applications ranging from tracking suspicious activities to migrations of animals. Previous related works either cannot apply on finding flock patterns, work only for archived datasets and/or find approximate results. We first show that flock discovery under a fixed time duration can be solved in polynomial time. We then present a framework that uses a lightweight grid-based structure in order to efficiently and incrementally process the trajectory locations. Using this framework we provide various on-line flock discovery algorithms. Experiments on real and synthetic trajectorial datasets show that our methods can efficiently report flock patterns even for large datasets and for different variations of the flock parameters (μ, ϵ and δ). As future work we will examine cost models to enable the user pick the most efficient algorithm based on the data distribution.

Acknowledgements: This research was partially supported by NSF grants IIS-0534781 and IIS-0803410. M. Vieira's work has been funded by a CAPES (Brazilian Federal Agency for Post-Graduate Education)/Fulbright Ph.D. fellowship. We would like to thank N. Young for the many discussions on this work and C. S. Jensen and H. Jeung for providing us with their datasets.

7. REFERENCES

[1] AccuTracking Inc. www.accutracking.com.

[2] S. Arumugam and C. Jermaine. Closest-point-of-approach join for moving object histories. In *IEEE ICDE*, pages 86–86, 2006.

[3] P. Bakalov, M. Hadjieleftheriou, and V. Tsotras. Time relaxed spatiotemporal trajectory joins. In *ACM GIS*, pages 182–191, 2005.

[4] M. Benkert, J. Gudmundsson, F. Hübner, and T. Wolle. Reporting flock patterns. In *Annual European Symposium on Algorithms (ESA)*, 2006.

[5] M. Benkert, J. Gudmundsson, F. Hübner, and T. Wolle. Reporting flock patterns. *Comput. Geom. Theory Appl.*, 41(3):111–125, 2008.

[6] Y. Cai and R. Ng. Indexing spatio-temporal trajectories with Chebyshev polynomials. In *SIGMOD Conference*, pages 599–610, 2004.

[7] ES. www.environmental-studies.de.

[8] M. Ester, H.-P. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, 1996.

[9] GeoChat. www.instedd.org/geochat.

[10] J. Gudmundsson and M. van Kreveld. Computing longest duration flocks in trajectory data. In *ACM GIS*, pages 35–42, 2006.

[11] S. Guha, R. Rastogi, and K. Shim. Cure: An efficient clustering algorithm for large databases. In *SIGMOD Conference*, pages 73–84, 1998.

[12] iSECUREtrac. tracNET24. www.isecuretrac.com.

[13] C. Jensen. Daisy. www.daisy.aau.dk.

[14] C. Jensen, D. Lin, and B. Ooi. Continuous clustering of moving objects. *IEEE Trans. Knowl. Data Eng*, 19(9):1161–1174, 2007.

[15] H. Jeung, M. L. Yiu, X. Zhou, C. S. Jensen, and H. T. Shen. Discovery of convoys in trajectory databases. *PVLDB*, 1(1):1068–1080, 2008.

[16] P. Kalnis, N. Mamoulis, and S. Bakiras. On discovering moving clusters in spatio-temporal data. In *SSTD*, pages 364–381, 2005.

[17] P. Laube, M. Kreveld, and S. Imfeld. Finding REMO: Detecting relative motion patterns in geospatial lifelines. *Dev. in Spatial Data Handling*, 2006.

[18] J.-G. Lee, J. Han, X. Li, and H. Gonzalez. *TraClass*: trajectory classification using hierarchical region-based and trajectory-based clustering. *PVLDB*, 2008.

[19] J.-G. Lee, J. Han, and K.-Y. Whang. Trajectory clustering: a partition-and-group framework. In *SIGMOD Conference*, pages 593–604, 2007.

[20] S.-L. Lee, S.-J. Chun, D.-H. Kim, J.-H. Lee, and C.-W. Chung. Similarity search for multidimensional data sequences. In *ICDE*, pages 599–608, 2000.

[21] Y. Li, J. Han, and J. Yang. Clustering moving objects. In *ACM SIGKDD*, pages 617–622, 2004.

[22] R. T. Ng and J. Han. Efficient and effective clustering methods for spatial data mining. In *VLDB*, 1994.

[23] PathIntelligence. www.pathintelligence.com.

[24] PCMB. www.taiga.net/satellite.

[25] D. Pfoser, C. Jensen, and Y. Theodoridis. Novel approaches in query processing for moving object trajectories. In *VLDB*, pages 395–406, 2000.

[26] RTreePortal. www.rtreeportal.org.

[27] M. Spiliopoulou, I. Ntoutsi, Y. Theodoridis, and R. Schult. Monic: modeling and monitoring cluster transitions. In *KDD*, pages 706–711, 2006.

[28] Y. Tao, D. Papadias, and Q. Shen. Continuous nearest neighbor search. In *VLDB*, pages 287–298, 2002.

[29] M. Vlachos, G. Kollios, and D. Gunopulos. Discovering similar multidimensional trajectories. In *IEEE ICDE*, pages 673–684, 2002.

[30] WhaleNet. whale.wheelock.edu.

[31] T. Zhang, R. Ramakrishnan, and M. Livny. Birch: An efficient data clustering method for very large databases. In *SIGMOD Conference*, 1996.