# Quasi-Polynomial Time Approximation Schemes for Target Tracking 

Matt Gibson Gaurav Kanade Erik Krohn Kasturi Varadarajan

July 7, 2009

## 1 Introduction

Target Tracking is the problem of keeping track of a set of specified targets by means of a given set of sensors. We study the target tracking problem in which the targets lie in the plane and the sensors are cameras also positioned in the plane. It requires two distinct cameras to estimate the position of a target. The quality of this estimation depends mainly on the relative position of the target with respect to that of the two cameras assigned to it [3, 2]. The field of view of a camera is a cone. A target can be tracked by a camera if it lies in this cone, and therefore a target tracked by a pair of cameras should lie in the intersection of their respective cones.

Tracking a target in this manner cannot in general provide accurate estimates of position. Hence it is important to carefully pick and assign pairs of cameras to different targets so as to minimize the error in estimation. With this setting our problem can be viewed as a resource allocation problem. For arbitrary error functions, this problem is NP-hard and hard to approximate - belonging to the class of Multi-Index Assignment Problems [6], but usually the error is some function of the geometry of the camera and target positions [4]. Other NP-hard versions of multi-index assignment problems also focus on geometry, such as those aiming to minimize the circumference or the area of a triangle formed by three assigned points in the plane [7].

The problem we thus consider is the Focus of Attention problem (FoA) which requires us to find a pairing of cameras and an assignment of camera pairs to a target in a manner that is optimum for some measure of tracking quality $[4,2]$. In our work we shall consider the constrained geometric setting in which the cameras are stationed on a line in the plane. Past work on this problem has also focused on this constrained setting $[4,2]$ and it is likely to model well the scenario in which targets are at relatively large distances from viable camera positions. Although the cameras cannot move they can rotate and freely choose their viewing direction. Both the cameras and the targets are represented by points.

We consider two geometric error metrics. The first is the "Aspect Ratio" which is the ratio of the vertical distance of the target from the camera line to the distance between the two cameras dedicated to it. (Refer Figure 1(a).) This metric can be used to gauge the error in stereo reconstruction and gives a good approximation if the cameras are not too close to the target. It was first considered by [4].

The second metric is the "Tracking Angle Deviation From Right Angles", studied by Gfeller et al. [2], who state that for a pair of cameras tracking a target, the tracking accuracy is best if the angle at the target - the tracking angle - is closest to $90^{\circ}$. Thus the considered metric is the


Figure 1: Here cameras $c_{i}$ and $c_{j}$ that are $b_{i, j}$ apart are assigned to target $t_{k}$ which is at a distance of $Z_{k}$ from the camera line


Figure 2: Tracking angles are all less than $90^{\circ}$ if targets are not too close
deviation of the angle from this desired value. (Refer Figure 1(b).) If the targets are not too close to the cameras, we can assume a scenario in which the tracking angles are all small i.e. less than $90^{\circ}$. In such a scenario minimizing the deviation from the optimum tracking angle is equivalent to maximizing the tracking angle. In this paper, we assume that all tracking angles are less than $90^{\circ}$ and hence our objective is to maximize the tracking angles.

We can assume without loss of generality that all the targets lie on one side of the line on which the cameras are placed. If they are not, we can consider their projections without affecting either of our optimization metrics.

Formally, we define the two versions of the FoA problem we consider as follows:
INPUT: A set $T$ of $n$ targets given as points in the plane and a set $C$ of $2 n$ cameras, given as collinear points on line $l$ in the same plane.
FEASIBLE SOLUTION: A camera assignment where each target is assigned to two cameras and each camera is assigned to exactly one target.
MEASURE: (1) An Aspect Ratio and (2) A Tracking Angle for every triple consisting of a target and two cameras.
GOAL: Find a feasible solution which is optimal for the sum of (1) aspect ratios and (2) tracking angles.

In this paper we consider the problem of (1) minimizing the sum of aspect ratios - the MIN-

SUMOFRATIOS problem and (2) maximizing the sum of tracking angles - the MAXSUMOFANGLES problem. For this second problem, we assume that the input has the property that the tracking angle of every triple is at most $90^{\circ}$, that is, every target lies outside the Thales' circle of any possible camera pair (Figure 2.)

Related Work. Target Tracking is an important research topic in the field of computer vision and image processing and has applications in environment surveillance and monitoring applications [8, 2].

Isler et al. [4] first studied target tracking by formulating the Focus of Attention (FoA) problem as a combinatorial optimization problem. The motivation behind their work was lowering the costs of optimum depth estimation. They showed that in a general setting (not in the plane) this comprises the classical NP-Hard 3-Dimensional Matching (3DM) problem as a special case. Therefore the focus in [4] is on the constrained geometric setting in which all cameras are restricted to lie on a single line $l$. The objective is the "aspect ratio" $\frac{Z_{k}}{b_{i, j}}$ where $Z_{k}$ is the distance of target $t_{k}$ from $l$ and $b_{i, j}$ is the distance between cameras $c_{i}$ and $c_{j}$ assigned to target $t_{k}$ (also called the baseline). They give a 2 -approximation for the problem of minimizing the sum of aspect ratios (that is, the MINSUMOFRATIOS problem) and for the problem of minimizing the maximum aspect ratio. Also, if the cameras are placed equidistantly on the line, they present a PTAS for the MINSUMOFRATIOS problem. They also consider cameras on a circle and targets inside the circle with tracking cost being $\frac{1}{\sin \theta}$, where $\theta$ is the tracking angle, and deliver a 1.42 -approximation for the problem of minimizing the the sum of tracking costs, and the maximum tracking cost.

Gfeller et al. [2] show that the problem of minimizing the sum of the deviations of tracking angles from $90^{\circ}$ (best tracking angle for accuracy) is NP-Hard, and that it admits no (multiplicative) approximation. For cameras on a line, they present a 2 -approximation algorithm for the problem of maximizing the sum of tracking angles (that is, MAXSUMOFANGLES) and maximizing the minimum tracking angle (or the bottleneck angle) under the assumption that all tracking angles are less than $90^{\circ}$. Also, if the cameras are placed equidistantly on the line, they present a PTAS for the MAXSUMOFANGLES problem. Arkin and Hassin [1] give a $2+\frac{1}{t}$ approximation for the problem of maximizing the sum of tracking angles with cameras lying on a line (Here, $t$ is the size of the local neighborhood in the local-search algorithm).

Our Contribution and Techniques. We consider the FoA problem that asks for camera assignment with minimum sum of Aspect Ratios (MINSUMOFRATIOS) and FoA that asks for camera assignment with maximum sum of tracking angles (MAXSUMOFANGLES). For cameras on a line, (and targets not lying in the Thales' circle of any camera pair) we present a QuasiPTAS for MAXSUMOFANGLES. For cameras on a line, we present a Quasi-PTAS also for the MINSUMOFRATIOS problem. A Quasi-PTAS is an algorithm that, for any constant $0<\epsilon<1$, returns a solution whose cost is within an additive $\epsilon$ factor of the optimal and has time complexity $n^{\text {polylog(n) }}$. Thus we improve upon the constant factor approximations known for these two problems, using quasi-polynomial rather than polynomial time.

It is evident that the powerful geometric structure underlying the Focus of Attention Problem sets it apart from the more general assignment problems and makes it interesting. This geometry has been exploited in various ways by [4, 2] to obtain efficient approximation algorithms. In particular in both these efforts the cameras are divided into two sets of equal size and classified as "left" cameras and "right" cameras. Also in the special case of cameras placed equidistantly on the
line both $[4,2]$ use the technique of further partitioning both the left and right cameras and then guessing the number of camera pairs in the optimal solution for each pair of blocks of the partition.

We also make use of the concept of left and right cameras; we extend the partitioning technique to cameras spaced arbitrarily on the line by making use of an intelligent discretization process. The main idea is to partition the interval containing the left (and right) cameras into a small number of intervals called "buckets", and guess the number of camera pairs between every pair of buckets in the optimal solution.

In the MAXSUMOFANGLES problem, if we guess that there are $T$ camera pairs involving buckets $B$ and $B^{\prime}$, we cannot simply return $T$ arbitrary pairs involving cameras in $B$ and $B^{\prime}$. Some of these pairs can be too "sensitive" for this crude process. Luckily, we get around this difficulty via a geometric observation about angles which implies that a sensitive pair can be sensitive at its $B$-end or its $B^{\prime}$-end but not both.

For the MINSUMOFRATIOS, the issue of sensitive pairs does not arise, but the main difficulty is that there is no single "scale" of distances at which we can apply the discretization. (This difficulty would not arise if the ratio of the maximum inter-camera distance to the minimum intercamera distance is polynomially bounded.) So we first apply the discretization at the scale of the median inter-camera distance (in the optimal solution), and recursively obtain and solve two independent instances of the problem with size at most $n / 2$. A characterization due to Isler et al.[4] of the optimal solution for a fixed camera pairing with the targets turns out to be quite useful in making this approach work.

Organization of the paper. In Section 2, we introduce some notation, review some important observations made in previous work on these problems, and give a high level overview of our algorithms. We present our algorithms for MAXSUMOFANGLES and MINSUMOFRATIOS in Sections 3 and 4, respectively.

## 2 Preliminaries and Notation

Suppose we have some horizontal line $l$ and a set $C$ of $2 n$ cameras such that each camera lies on $l$. We call the $i$ th camera on the line $c_{i}$. For $c_{i}, c_{j} \in C$, we say $c_{i}<c_{j}$ if $c_{i}$ is to the left of $c_{j}$ on $l$. We assume that no two cameras have the same position, and thus we have $c_{1}<c_{2}<\ldots<c_{2 n}$. We are also given a set $T=\left\{t_{1}, t_{2}, \ldots, t_{n}\right\}$ of targets that lie in the plane. Without loss of generality, we can assume that all of the targets lie above $l$. The distance between cameras $c_{i}$ and $c_{j}$ is called the baseline of $c_{i}$ and $c_{j}$ and is denoted $b_{i, j}$. A camera pair is a set of two distinct cameras, and a camera pairing of a set $C^{\prime}$ of $2 m$ cameras is a set of $m$ camera pairs such that each camera in $C^{\prime}$ appears in exactly one of the camera pairs.

A pairing of cameras is all-overlapping if the baselines of any two camera pairs in the pairing intersect. Suppose there is a camera pairing that contains pairs $\left(c_{i}, c_{j}\right)$ and $\left(c_{i^{\prime}}, c_{j^{\prime}}\right)$ such that $c_{i}<c_{j}<c_{i^{\prime}}<c_{j^{\prime}}$ (i.e. the pairing is not all-overlapping). It is easily seen [2] that we can do better for both objective functions if we use the pairs $\left(c_{i}, c_{i^{\prime}}\right)$ and $\left(c_{j}, c_{j^{\prime}}\right)$. Refer Figure 3. This leads to the observation [2] that there is an optimal camera pairing where every camera pair $\left(c_{i}, c_{j}\right)$ has $i \in\{1,2, \ldots, n\}$ and $j \in\{n+1, n+2, \ldots, 2 n\}$. Let $M$ denote the midpoint of $b_{n, n+1}$. See Figure 4. Another way to view this observation is that there is an optimal solution with camera pairing $P$ such that for all pairs $\left(c_{i}, c_{j}\right) \in P, c_{i}$ is to the left of $M$ and $c_{j}$ is to the right of $M$. Our algorithms will only consider such camera pairings.


Figure 3: For both objective functions, we can do better if we overlap the camera pairs (see dotted lines).


Figure 4: In this example $n=3$, cameras $c_{1}, c_{2}, c_{3}$ are "left" cameras and the rest are "right" cameras. $M$ is the mid-point of baseline $b_{3,4}$ formed by cameras $c_{3}$ and $c_{4}$.

We shall now present an overview of our algorithms. Although historically the "Aspect Ratio" objective function was considered first [4], it was shown in [2] that the "Tracking Angle" is in general the most influential way of tracking quality directly. Hence we shall consider this metric first. Both algorithms use a discretization procedure to aid in computing our solution. However, a crude discretization procedure will not suffice for either problem, and we must use a more clever technique. In the case of MAXSUMOFANGLES, a geometric observation comes to our rescue. In the case of MINSUMOFRATIOS, a more sophisticated discretization process comes to our rescue.

First consider the MAXSUMOFANGLES problem. Our algorithm constructs a small number of camera "pairings" - (a pairing is simply the set formed when each camera is paired with exactly one other camera). Given this set of camera pairings, we then are able to determine which camera pairing is the best by using a polynomial time algorithm for minimum-weight perfect matching on a bipartite graph in which one set of vertices is the camera pairs [5] and the other is the targets , as observed by Gfeller et al. [2].

Suppose we fix some optimal solution that uses some all-overlapping camera pairing $P$. That is, every camera pair in $P$ involves one camera to the left of $M$ and one to the right of $M$. We discretize the line into a small number of buckets such that the length of each bucket is small compared to its distance from $M$. Consider a bucket on one side of $M$ and another bucket on the other side of $M$. Suppose we are able to correctly guess that the number of camera pairs in $P$ that have one camera in each of these buckets is $\mu$. We would like to be able to arbitrarily pick $\mu$ cameras from the first bucket and pair them with an arbitrary $\mu$ cameras in the second bucket and argue that the tracking angle formed by these camera pairings at the target will not be too small as compared to the one formed by the corresponding pair in $P$. Unfortunately, this approach will not work, as the arbitrary assignment might cause some tracking angle to be more than an $\epsilon$ factor smaller than the corresponding angle in $P$. To get around this obstacle, we state and prove a geometric lemma (Lemma 2) which allows us to handle these "sensitive" cases. The algorithm is then to guess the number of such "sensitive" cases that arise for each of our buckets. We handle those camera pairs (relying heavily on our lemma), and then we use the more crude technique on all of the remaining cameras falling in buckets. If any cameras still remain to be paired, we show that there is an easy way to pair them off (again making use of the same lemma).

Now consider the MINSUMOFRATIOS problem. Suppose we fix some optimal solution that uses some all-overlapping camera pairing $P$. Let $b^{i}$ denote the $i$ th baseline in $P$ when the baselines are indexed such that $b^{1} \leq b^{2} \leq \cdots \leq b^{n}$. The high level idea of our algorithm is that we will guess the length of the median baseline (i.e. $b^{n / 2}$ ) and then guess all camera pairs in $P$ whose baselines are within a polynomial factor of $b^{n / 2}$. To aid us in guessing, we use a discretization procedure. This procedure allows us to make all possible guesses in quasi-polynomial time while not doing too much worse than what the optimal solution would have done. Suppose that we make a correct guess. Then we will have correctly guessed all camera pairs in $P$ whose baselines fall within some interval $\left[b^{i}, b^{j}\right]$ for $b^{i} \leq b^{n / 2} \leq b^{j}$ up to an $\epsilon$ factor. We then recursively repeat this guessing procedure for baselines in the interval $\left[b^{1}, b^{i-1}\right]$ and the interval $\left[b^{j+1}, b^{n}\right]$. The correctness of the algorithm then follows by showing that the errors do not accumulate too much over the course of the recursion.

Fix a camera pairing $P$ and a set of targets $T$ such that $|P|=|T|$. The optimal association of camera pairs in $P$ with targets in $T$ is rather easy to compute, using a characterization given by Isler et al. [4]:

Lemma 1. Let $Z_{i}$ be the distances of targets in $T$ from line $l, Z_{1} \leq Z_{2} \leq \cdots \leq Z_{m}$ and $b^{i}$ be the baselines in $P$ sorted such that $b^{1} \leq b^{2} \leq \cdots \leq b^{m}$. There exists an optimal matching such that the
target at depth $Z_{i}$ is assigned to the pair with baseline $b^{i}$.
We denote the cost of such an assignment of a pairing $P$ to targets $T$ as $\operatorname{cost}(P, T)$ and we compute it according to the assignment in Lemma 1.

## 3 Maximizing the Sum of Angles

We now describe our approximation algorithm for the MAXSUMOFANGLES problem. We adopt the notation of Section $2-$ the camera locations are $c_{1}<c_{2}<\cdots<c_{2 n}$ on line $l$, with $c_{n}<$ $M<c_{n+1}$. The targets are points in the plane above the line, and denoted $t_{1}, \ldots, t_{n}$. Let $\epsilon>0$ be the given approximation parameter; we may assume that $\epsilon<1 / 2$. Our algorithm outputs a small number of camera pairings of $C$. One of the camera pairings that we output will have the property that it is possible to associate each of the camera pairs in it with targets in such a way that the sum of the angles at the targets is at least $(1-\epsilon)$ times that of OPT, where OPT is the value of the sum of tracking angles in an optimal solution. For a given camera pairing, we can compute the best association with the targets using a polynomial time algorithm for minimum-weight perfect matching [5], as observed by Gfeller et al. [2]. We can therefore evaluate all of the camera pairings that we output, and return the best one.

Our algorithm makes use of a subroutine for partitioning an interval to achieve a result of the following form: Given $0<\gamma_{1}<\gamma_{2}$, partition the interval [ $M+\gamma_{1}, M+\gamma_{2}$ ] (resp. $\left[M-\gamma_{2}, M-\gamma_{1}\right]$ ) into buckets (sub-intervals), so that the distance between the left (resp. right) endpoint of each bucket from $M$ is at least $\frac{1}{\epsilon^{2}}$ times the bucket length. We will call a partition of $\left[M+\gamma_{1}, M+\gamma_{2}\right]$ (resp. [M- $\left.\gamma_{2}, M-\gamma_{1}\right]$ ) with this property a conforming partition. Clearly, we can compute a conforming partition with $O\left(\frac{\log \left(\gamma_{2} / \gamma_{1}\right)}{\epsilon^{2}}\right)$ buckets: partition $\left[M+\gamma_{1}, M+2 \gamma_{1}\right]$ into $1 / \epsilon^{2}$ equal sized buckets, then partition $\left[M+2 \gamma_{1}, M+4 \gamma_{1}\right]$ into $1 / \epsilon^{2}$ equal sized buckets, and so on until we are past $M+\gamma_{2}$. We also adopt the following notation for the algorithm: $L$ will denote the cameras to the left of $M$, and $R$ the cameras to the right. During the course of the algorithm, some of the cameras get paired, and $L$ (resp. $R$ ) will always stand for the unpaired cameras. Given a bucket $B$ to the left (resp. right) of $M$, we will let $|B|$ denote the number of cameras in the current $L$ (resp. $R$ ) that fall in bucket $B$. Finally, let $a=M-c_{1}$, and $d=c_{2 n}-M$. In describing the algorithm and analysis, we will assume $a \leq d$; the other case is symmetric. See Figure 5 for an illustration.

The Algorithm. The following algorithm outputs one candidate pairing of the cameras $C$ for each combination of the functions $\sigma, \pi, \mu$, and $\lambda$ considered. We then evaluate each of these pairings based on the stated criterion (MAXSUMOFANGLES) and return the best one.

Running Time. We will now bound the running time of the algorithm by bounding the number of possibilities that the mappings $\sigma, \pi$, and $\mu$ consider. We can bound the number of cameras in any bucket by $n$. Thus we can bound the number of possibilities for $\sigma$ by $n^{O(k)}$, and we can bound the number of possibilities for $\pi$ by $n^{O(j)}$. We can bound the number of possibilities for $\mu$ by $n^{O(k j)}$.

Because the number of buckets $j$ and $k$ are bounded by $O\left(\frac{\log (n / \epsilon)}{\epsilon^{2}}\right)$, we have that the running time of the algorithm is bounded by

$$
n^{O\left(\frac{\log ^{2}(n / \epsilon)}{\epsilon^{4}}\right)} .
$$

This expression also bounds the number of candidate pairings of $C$ that are output


Figure 5: Illustration for a conforming partition (the intervals on $l$ are not drawn to scale). Each of the shown subintervals of $l$ (except for $\left[M-\gamma_{1}, M\right]$ and $\left[M, M+\gamma_{1}\right]$ ) are divided into $\frac{1}{\epsilon^{2}}$ buckets of equal length.


Figure 6: Illustration for Lemma 2. For $\alpha_{2}>\alpha_{1}, \frac{\left|x z_{2}^{\prime}\right|}{\left|x z_{2}\right|}>\frac{\left|x z_{1}^{\prime}\right|}{\left|x z_{1}\right|}$.

Approximation Factor. Let $\alpha_{t}$ denote the target angle corresponding to target $t \in T$ in the optimal solution OPT. We now show that there is one choice of the functions $\sigma, \pi, \mu$, and $\lambda$ for which it is possible to associate the candidate pairing output by our algorithm to the targets in such a way that the target angle $\beta_{t}$ in this solution for any target $t$ satisfies

$$
\begin{equation*}
\beta_{t} \geq \alpha_{t}-4 \max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\} . \tag{1}
\end{equation*}
$$

Adding this inequality for each $t \in T$ gives $\sum_{t} \beta_{t} \geq(1-4 \epsilon) \sum_{t} \alpha_{t}$, establishing the algorithm's correctness. In order to establish this inequality, we will find it convenient to move the cameras as part of the analysis; this will be safe because we only move towards $M$. For each target $t$, let $l(t)$ and $r(t)$ initially denote the cameras associated with $t$ in OPT to the left and right of $M$, respectively. As we modify the association of targets with camera pairs in the analysis, $l(t)$ and $r(t)$ will change. At all times, $\beta_{t}$ will denote the angle $l(t) \operatorname{tr}(t)$. The following geometric lemma, whose proof we present in the appendix, plays a crucial role in the analysis. Refer Figure 6.

```
Algorithm 1
    Let \(B_{0}\) be the bucket \(\left[M-\frac{\epsilon a}{100 n^{2}}, M\right]\). Let \(B_{1}, \ldots, B_{k}\) be the buckets resulting from a conforming
    partition of \(\left[M-a, M-\frac{\epsilon a}{100 n^{2}}\right]\).
    for each map \(\sigma:\left\{B_{1}, \ldots, B_{k}\right\} \rightarrow \mathcal{Z}^{+}\)such that \(\sigma\left(B_{i}\right) \leq\left|B_{i}\right|\) do
        For each \(1 \leq i \leq k\), pair the \(\sigma\left(B_{i}\right)\) leftmost cameras in \(B_{i}\) with the \(\sigma\left(B_{i}\right)\) leftmost cameras
        in \(R\) arbitrarily.
        Let \(B_{0}^{\prime}\) be the bucket \(\left[M, M+\frac{\epsilon a}{100 n^{2}}\right]\). Let \(B_{1}^{\prime}, \ldots, B_{j}^{\prime}\) be the buckets resulting from a con-
        forming partition of \(\left[M+\frac{\epsilon a}{100 n^{2}}, M+\frac{a}{\epsilon^{2}}\right]\).
        for each map \(\pi:\left\{B_{1}^{\prime}, \ldots, B_{j}^{\prime}\right\} \rightarrow \mathcal{Z}^{+}\)such that \(\sigma\left(B_{i}^{\prime}\right) \leq\left|B_{i}^{\prime}\right|\) do
            For each \(1 \leq i \leq j\), pair the \(\pi\left(B_{i}^{\prime}\right)\) rightmost cameras in \(B_{i}^{\prime}\) with the \(\pi\left(B_{i}^{\prime}\right)\) rightmost
            cameras in \(L\) arbitrarily.
            for each map \(\mu:\left\{B_{0}, \ldots, B_{k}\right\} \times\left\{B_{0}^{\prime}, \ldots, B_{j}^{\prime}\right\} \rightarrow \mathcal{Z}^{+}\)such that \(\sum_{\ell} \mu\left(B_{i}, B_{\ell}^{\prime}\right) \leq\left|B_{i}\right|\) for
            every \(i \leq k\) and \(\sum_{\ell} \mu\left(B_{\ell}, B_{i}^{\prime}\right) \leq\left|B_{i}^{\prime}\right|\) for every \(i \leq j\) do
                    Go through the ( \(B_{i}, B_{\ell}^{\prime}\) ) pairs in any order and arbitrarily pair \(\mu\left(B_{i}, B_{\ell}^{\prime}\right)\) cameras from
                \(L \cap B_{i}\) with \(\mu\left(B_{i}, B_{\ell}^{\prime}\right)\) cameras from \(R \cap B_{\ell}^{\prime}\).
                    for each \(0 \leq \lambda \leq|B|\) where \(B\) is the bucket \([M-a, M]\) do
                    Pair the \(\lambda\) leftmost cameras in \(L \cap B=L\) with the \(\lambda\) leftmost cameras in \(R\) arbitrarily.
                    Pair the remaining cameras in \(L\) arbitrarily with cameras in \(R\).
```

Lemma 2. Let $x, y, t$ be three non-collinear points in the plane such that angle $(x t y)=\theta$. Let $z$ and $z^{\prime}$ be two points on the line $x y$ such that both $z$ and $z^{\prime}$ lie between $x$ and $y$ and angle $(x t z)=$ angle $\left(y t z^{\prime}\right)=\epsilon \cdot \theta$. Recall that $0<\epsilon<1 / 2$. Then the ratio $\frac{\left|x z^{\prime}\right|}{\mid x z} \leq \frac{1}{\epsilon^{2}}$ where $|x z|,\left|x z^{\prime}\right|$ are the lengths of segments $x z, x z^{\prime}$ respectively.

Fixing $\sigma$. In this step, let us call a target $t$ sensitive if $l(t)$ lies in bucket $B_{i}$ for $i \geq 1$ (so not in the first bucket $B_{0}$ ), and the angle $l(t) t p$ is at least $\epsilon$ times the angle $\beta_{t}=l(t) \operatorname{tr}(t)$, where $p$ is the right endpoint of bucket $B_{i}$. Since the buckets come from a conforming partition, Lemma 2 implies that for a sensitive $t$ the angle $l(t) t c$ for any camera $c$ to the right of $M$ is at least $(1-\epsilon) \beta_{t}$. (The upshot is that we have to be careful in changing $l(t)$, but we have considerable flexibility with $r(t)$.) Fix $\sigma$ so that $\sigma\left(B_{i}\right)$ is the number of sensitive targets $t$ with left endpoints within $B_{i}$. With this choice of $\sigma$, recall that our algorithm fixes $\sum_{i} \sigma\left(B_{i}\right)$ camera pairs at this stage. Now for each sensitive $t$, reset $l(t)$ so that it lies to the left of (or is the same as) the original $l(t)$ and is one of the $\sigma\left(B_{i}\right)$ leftmost points in its bucket $B_{i}$, and reset $r(t)$ to be the camera that is paired with the new $l(t)$ in the partial camera pairing that is fixed. It follows that now $\beta_{t} \geq(1-\epsilon) \alpha_{t} \geq \alpha_{t}-\max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\}$ for all sensitive $t$. The association of sensitive $t$ with the camera pairing we output is finalized now, and hence Inequality 1 holds for such $t$.

For the $t$ that are not sensitive, reset $l(t)$ to be some other camera in the same bucket as the original $l(t)$ and reset $r(t)$ to be a camera that is the same or to the right of the original $r(t)$. (During these resettings, we always ensure that each camera is $l(t)$ or $r(t)$ for exactly one $t$.) Now "move" the camera that is $l(t)$ for each such $t$ to the right endpoint of the bucket containing $l(t)$.

For a $t$ that is not sensitive and for which $l(t)$ does not lie in bucket $B_{0}$, it is clear that the new $\beta_{t}$ is at least $(1-\epsilon)$ times the original $\beta_{t}$. Now if $l(t)$ lies in $B_{0}$, letting $p_{1}$ and $p_{2}$ denote the left


Figure 7: A figure to illustrate that angle $\left(p_{1} t p_{2}\right)$ is very small compared to angle $\left(c_{1} t c_{2 n}\right)$.
and right endpoints of $B_{0}$, we argue that now

$$
\beta_{t} \geq \alpha_{t}-\operatorname{angle}\left(p_{1} t p_{2}\right) \geq \alpha_{t}-\frac{\epsilon}{n^{2}} \operatorname{angle}\left(c_{1} t c_{2 n}\right) \geq \alpha_{t}-\frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}
$$

For the second inequality, we show that angle $\left(p_{1} t p_{2}\right) \leq \frac{\epsilon}{n^{2}}$ angle $\left(c_{1} t c_{2 n}\right)$; we omit from this version the somewhat tedious argument for this claim, but note that it is here that we use the fact that $t$ is outside the circle with diameter $c_{1} c_{2 n}$. See Figure 7 for a figure to provide intuition for why the claim holds.

Thus, for the $t$ that are not sensitive, we have

$$
\beta_{t} \geq \alpha_{t}-\max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\} .
$$

We move on to the next stage with the targets $t$ whose association with camera pairings has not been finalized (the targets that were not sensitive in this stage), along with the cameras $l(t)$ and $r(t)$ for such targets (the cameras that have not been paired up by our algorithm).

Fixing $\pi$. This is quite symmetric to the previous step. Among the targets that have moved on to this step, let us call a target $t$ sensitive if $r(t)$ lies in bucket $B_{i}^{\prime}$ for $i \geq 1$ (so not in the first bucket $B_{0}^{\prime}$ ), and the angle $p \operatorname{tr}(t)$ is at least $\epsilon$ times the angle $\beta_{t}=l(t) \operatorname{tr}(t)$, where $p$ is the left endpoint of bucket $B_{i}^{\prime}$. Since the buckets come from a conforming partition, Lemma 2 implies that for a sensitive $t$ the angle $c \operatorname{tr}(t)$ for any camera $c$ to the left of $M$ is at least $(1-\epsilon) \beta_{t}$. Fix $\pi$ so that $\pi\left(B_{i}^{\prime}\right)$ is the number of sensitive targets $t$ with $r(t)$ within $B_{i}^{\prime}$. With this choice of $\pi$, recall that our algorithm now fixes $\sum_{i} \pi\left(B_{i}^{\prime}\right)$ camera pairs at this stage. Now for each sensitive $t$, reset $r(t)$ so that it lies to the right of (or is the same as) the original $r(t)$ and is one of the $\pi\left(B_{i}^{\prime}\right)$ rightmost points in its bucket $B_{i}^{\prime}$, and reset $l(t)$ to be the camera that is paired with the new $r(t)$ in the partial camera pairing that is now fixed. It follows that for all sensitive $t$, the current $\beta_{t}$ is at least $(1-\epsilon)$ times the $\beta_{t}$ at the end of the previous step, and thus $\beta_{t} \geq \alpha_{t}-2 \max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\}$. The association of sensitive $t$ with the camera pairing we output is finalized now, and hence Inequality 1 holds for such $t$.

For the $t$ that are not sensitive, reset $r(t)$ to be some other camera in the same bucket as the original $r(t)$ and reset $l(t)$ to be a camera that is the same or to the left of the original $l(t)$. Now "move" the camera that is $r(t)$ for each such $t$ to the left endpoint of the bucket containing $r(t)$.

For a $t$ that is not sensitive, we can argue along the same lines as in the previous step to conclude that $\beta_{t} \geq \alpha_{t}-2 \max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\}$.

We move on to the next stage with the targets $t$ whose association with camera pairings has not been finalized (the targets that were not sensitive in this stage), along with the cameras $l(t)$ and $r(t)$ for such targets (the cameras that have not been paired up by our algorithm).

Fixing $\mu$. For every pair $0 \leq i \leq k$ and $0 \leq i^{\prime} \leq j$, fix $\mu\left(i, i^{\prime}\right)$ to be the number of surviving targets $t$ with $l(t) \in B_{i}$ and $r(t) \in B_{i^{\prime}}^{\prime}$. With this choice of $\mu$, recall that our algorithm outputs $\mu\left(i, i^{\prime}\right)$ pairs of cameras at this stage with one endpoint in $B_{i}$ and the other in $B_{i^{\prime}}^{\prime}$. Reassign the $l(t)$ and $r(t)$ values for the surviving targets (this reassignment should not cause an $l(t)$ or $r(t)$ value to move to a different bucket), if necessary, so that these $\mu\left(i, i^{\prime}\right)$ pairs that are output by the algorithm are assigned to the targets that led to the definition of $\mu\left(i, i^{\prime}\right)$. The association of such targets with the camera pairing that we output is finalized at this stage. Notice that $\beta_{t}$ does not change for any target in this step, whether it is finalized or not. This is because the cameras have already been moved to endpoints of the bucket they belong to. In particular, this means that Inequality 1 holds for the targets that do not survive this stage.

We now move to the final stage with the targets that survive. Note that the $l(t)$ for such targets belong to the bucket $B=[M-a, M]$, and the $r(t) \geq M+\frac{a}{\epsilon^{2}}$.

Fixing $\lambda$. In this step, call a surviving target $t$ sensitive if angle $(l(t) t M) \geq \epsilon \cdot \operatorname{angle}(l(t) \operatorname{tr}(t))=\epsilon \beta_{t}$. Lemma 2 implies that for a sensitive $t$ the angle $l(t) t c$ for any camera $c$ to the right of $M+\frac{a}{\epsilon^{2}}$ is at least $(1-\epsilon) \beta_{t}$. Fix $\lambda$ to be the number of such sensitive targets $t$. With this choice of $\lambda$, recall that our algorithm fixes $\lambda$ camera pairs at this stage. Now for each sensitive $t$, reset $l(t)$ so that it lies to the left of (or is the same as) the original $l(t)$ and is one of the $\lambda$ leftmost points in its bucket $B$, and reset $r(t)$ to be the camera that is paired with the new $l(t)$ in the partial camera pairing that is fixed. It follows that for sensitive $t$, the new $\beta_{t}$ is at least $(1-\epsilon)$ times the $\beta_{t}$ before this step, and thus $\beta_{t} \geq \alpha_{t}-3 \max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\}$. The association of sensitive $t$ with the camera pairing we output is finalized now, and hence Inequality 1 holds for such $t$.

For the $t$ that are not sensitive, reset $l(t)$ to be some other camera in $B$, and reset $r(t)$ to be a camera that is the same or to the right of the original $r(t)$. Move the camera $l(t)$ to the point $M$. The new $\beta_{t}$ is at least $(1-\epsilon)$ times the $\beta_{t}$ before this step, and thus $\beta_{t} \geq \alpha_{t}-3 \max \left\{\epsilon \alpha_{t}, \frac{\epsilon}{n^{2}} \sum_{t} \alpha_{t}\right\}$.

By reassigning $l(t)$, if necessary, but not the $r(t)$, the camera pairs that the algorithm outputs now will have the form $(l(t), r(t))$ for each surviving $t$. This completes the association for all the targets. This reassignment does not change $\beta_{t}$ because each $l(t)$ is on point $M$, and Inequality 1 holds for every $t$.

This completes our argument about the approximation guarantee.
Theorem 3. There is an algorithm for the MAXSUMOFANGLES problem that, for any $0<\epsilon<1$, runs in quasi-polynomial time and returns a solution in which the sum of tracking angles is at least $(1-\epsilon)$ times that in the optimal solution.

## 4 Minimizing the Sum of Ratios

We now present our algorithm for the MINSUMOFRATIOS problem. Recall that in this problem the cost of assigning a camera pair $\left(c_{i}, c_{j}\right)$ to target $t_{k}$ is $\frac{Z_{k}}{b_{i, j}}$ where $Z_{k}$ is the vertical distance
between $t_{k}$ and the line $l$ containing the cameras, and $b_{i, j}$ is the baseline corresponding to (distance between) $c_{i}$ and $c_{j}$. We would like to minimize the sum of these costs. In this section, we present an algorithm that returns a solution whose cost (sum of aspect ratios) is no worse than ( $1+\epsilon$ ) times the same in the optimal solution, in quasi-polynomial time for any $\epsilon>0$.

The Algorithm. Our algorithm is recursive and is formally defined as Algorithm 2. The algorithm takes as input a set of $2 m$ cameras $X$, a set of $m$ targets $Y$, and two positive real numbers $\ell$ and $u$. The cameras are indexed $x_{1}, x_{2}, \ldots x_{2 m}$ such that $x_{1}<x_{2}<\cdots<x_{2 m}$ and $x_{m}<M<x_{m+1}$ and the targets are indexed $y_{1}, y_{2}, \ldots y_{m}$ such that $Z_{1} \leq Z_{2} \leq \cdots \leq Z_{m}$. The algorithm will return a camera pairing $P$ of the cameras in $X$ such that the baselines of all of the pairs in $P$ are at least $\ell$ and are at most $u$. If no such pairing is possible, our algorithm returns a "dummy" camera pairing $I$; for ease of description, we define $\operatorname{cost}(P, T)=\infty$ for any pairing $P$ that contains $I$. Our algorithm breaks up the line $l$ into intervals called buckets. We will use $|d|$ to denote the number of cameras in $X$ that lie in bucket $d$. We assume $0<\epsilon<1 / 2$.

```
Algorithm 2 minRatioPair \((X, Y, \ell, u)\)
    If \(|X|=0\), return the empty set.
    Let \(\mathcal{L}\) be the set of all baselines \(b_{i, j}\) such that \(1 \leq i \leq m, m+1 \leq j \leq 2 m\), and \(\ell \leq b_{i, j} \leq u\).
    Set local variable best \(\leftarrow \infty\), and \(P_{\text {best }} \leftarrow I\).
    for each \(\beta \in \mathcal{L}\) do
        Partition the interval \([M-2 n \beta, M]\) into \(O\left(\frac{\log n}{\epsilon}\right)\) buckets in the following way. Initialize
        \(r=M\).
        for \(i=-1,0,1, \ldots, 2 \log n\) do
            Discretize the interval \(\left[r-2^{i} \frac{\beta}{n}, r\right]\) into \(\frac{2}{\epsilon}\) buckets of equal length.
            \(r \leftarrow 2^{i} \frac{\beta}{n}\)
        Symmetrically partition the interval \([M, M+2 n \beta]\) into \(O\left(\frac{\log n}{\epsilon}\right)\) buckets.
        Let \(\mathcal{B}_{1}\) denote the set of all buckets to the left of \(M\), and let \(\mathcal{B}_{2}\) denote the set of buckets to
        the right of \(M\).
        for each map \(\mu: \mathcal{B}_{1} \times \mathcal{B}_{2} \rightarrow \mathcal{Z}^{+}\)and each map \(\sigma: \mathcal{B}_{1} \cup \mathcal{B}_{2} \rightarrow \mathcal{Z}^{+}\)such that \(\sigma(B)+\)
        \(\sum_{B^{\prime} \in \mathcal{B}_{2}} \mu\left(B, B^{\prime}\right) \leq|B|\) for each \(B \in \mathcal{B}_{1}, \sigma\left(B^{\prime}\right)+\sum_{B \in \mathcal{B}_{1}} \mu\left(B, B^{\prime}\right) \leq\left|B^{\prime}\right|\) for each \(B^{\prime} \in \mathcal{B}_{2}\),
        and \(\sum_{B \in \mathcal{B}_{1}} \sigma(B)=\sum_{B^{\prime} \in \mathcal{B}_{2}} \sigma\left(B^{\prime}\right)\) do
            \(X_{\text {short }} \leftarrow \emptyset, X_{\text {mid }} \leftarrow \emptyset, X_{\text {long }} \leftarrow \emptyset, P_{\text {mid }} \leftarrow \emptyset\)
            Go through each \(\left(B, B^{\prime}\right) \in \mathcal{B}_{1} \times \mathcal{B}_{2}\) in any order and pair the \(\mu\left(B, B^{\prime}\right)\) cameras in \(B\) that
            are closest to \(M\) with the \(\mu\left(B, B^{\prime}\right)\) cameras in \(B^{\prime}\) that are closest to \(M\) and place these
            pairs in \(P_{\text {mid }}\). Add these cameras to \(X_{\text {mid }}\) and remove them from their respective buckets.
            For each bucket \(B \in \mathcal{B}_{1} \cup \mathcal{B}_{2}\) : Of the cameras that remain in \(B\), place the \(\sigma(B)\) cameras
            that are farthest from \(M\) into \(X_{\text {short }}\). Place all other remaining cameras in \(B\) into \(X_{\text {long }}\).
            Denoting \(m_{s} \equiv \frac{\left|X_{\text {short }}\right|}{2}\) and \(m_{\text {mid }} \equiv \frac{\left|X_{\text {mid }}\right|}{2}\), let \(Y_{\text {short }}=\left\{y_{1}, y_{2}, \ldots, y_{m_{s}}\right\}, Y_{\text {mid }}=\)
        \(\left\{y_{m_{s}+1}, y_{m_{s}+2}, \ldots, y_{m_{s}+m_{\text {mid }}}\right\}\), and \(Y_{\text {long }}=\left\{y_{m_{s}+m_{m i d}+1}, \ldots, y_{m}\right\}\).
        if \(\left|Y_{\text {short }}\right| \leq m / 2\) and \(\left|Y_{\text {long }}\right| \leq m / 2\) then
            \(P \leftarrow P_{\text {mid }} \cup\) minRatioPair \(\left(X_{\text {short }}, Y_{\text {short }}, \ell,(1+\epsilon) \frac{\beta}{2 n}\right) \cup \operatorname{minRatioPair}\left(X_{\text {long }}, Y_{\text {long }},(1-\right.\)
            є) \(2 n \beta, u(1+\epsilon))\)
            If \(\operatorname{cost}(P, Y)<\) best, then best \(\leftarrow \operatorname{cost}(P, Y)\) and \(P_{\text {best }} \leftarrow P\).
    Return \(P_{\text {best }}\).
```



Figure 8: Illustration for the discretization process. The value in each interval denotes the length of the interval. Each interval is divided into $\frac{2}{\epsilon}$ buckets of equal length.


Figure 9: Illustration for the step in algorithm where the bucket pair ( $B, B^{\prime}$ ) contributes to $X_{\text {mid }}$. The filled in cameras will be paired up and placed into $X_{\text {mid }}$ for $\mu\left(B, B^{\prime}\right)=2$

For solving the input instance of MINSUMOFRATIOS, we invoke minRatioPair $\left(C, T, b_{n, n+1}, b_{1,2 n}\right)$. For this invocation, note that the depth of the recursion is at most $\log n$, because the size of $Y$ falls by a factor of 2 with each recursive call. It will also be useful to note that when we make a recursive call, the lower bound $\ell$ for the recursive call is not smaller than the original $\ell$; the upper bound $u$ for the recursive call can be larger than the original, but only by a factor of $(1+\epsilon)$.

Running Time. We will now show that Algorithm 2 runs in quasi-polynomial time when the instance minRatioPair $\left(C, T, b_{n, n+1}, b_{1,2 n}\right)$ is invoked. The set $\mathcal{L}$ contains $O\left(m^{2}\right)$ elements. The number of buckets in $\mathcal{B}_{1} \cup \mathcal{B}_{2}$ is $O\left(\frac{\log n}{\epsilon}\right)$. This means that the number of choices of $\mu$ and $\sigma$ is at most $(m+1)^{O\left(\frac{\log ^{2} n}{\epsilon^{2}}\right)}$. Thus we make $(m+1)^{O\left(\frac{\log ^{2} n}{\epsilon^{2}}\right)}=n^{O\left(\frac{\log ^{2} n}{\epsilon^{2}}\right)}$ direct recursive calls. Since the depth of the recursions is at most $\log n$, this gives us our quasi-polynomial running time.

Approximation Ratio. We now give an argument for the approximation factor guaranteed which, though informal, highlights the main issues. Consider the optimal pairing $P_{O P T}$ for the input instance with cameras $C$ and targets $T$, and the optimal association of targets in $T$ with cameras in $P_{O P T}$. We show we can associate each target $t$ with a camera pair in such a way that (a) these camera pairs form a camera pairing of $C$, (b) the baseline of the camera pair associated with each $t$ is at least $(1-5 \epsilon)$ times the corresponding baseline in $P_{O P T}$, and (c) our algorithm outputs a solution that is at least as good as this association. Thus, we obtain our $(1+O(\epsilon))$ approximation factor.

Let us describe how this special association is constructed. Let us start with the optimal
association of each $t \in T$ with the corresponding camera pair in $P_{O P T}$, and consider the invocation minRatioPair $\left(C, T, b_{n, n+1}, b_{1,2 n}\right)$. Consider the situation where the algorithm chooses $\beta$ to be the median baseline in $P_{O P T}$. With this $\beta$ it computes bucket sets $\mathcal{B}_{1}$ and $\mathcal{B}_{2}$. The pairs in $P_{O P T}$ can be split into three sets: those whose baselines are strictly smaller than $\beta / 2 n$ (the short baselines), those whose baselines are between $\beta / 2 n$ and $2 n \beta$ (the medium baselines), and those whose baselines are strictly larger than $2 n \beta$ (the long baselines). Consider the algorithm's choice of $\mu$ so that $\mu\left(B, B^{\prime}\right)$ equals the number of medium baselines with endpoints in $B$ and $B^{\prime}$, and the choice of $\sigma$ so that $\sigma(B)$ equals the number of cameras in $B$ that are endpoints of short baselines. Notice that with this choice of $\mu$ the algorithm constructs sets $Y_{\text {short }}, Y_{\text {mid }}$, and $Y_{\text {long }}$ of sizes $m_{s}=\sum_{B} \sigma(B)$, $m_{\text {mid }}=\sum_{\left(B, B^{\prime}\right)} \mu\left(B, B^{\prime}\right)$, and $m-m_{s}-m_{\text {mid }}$, respectively, and sets $X_{\text {short }}, X_{\text {mid }}$, and $X_{\text {long }}$ of sizes $2 m_{s}, 2 m_{\text {mid }}$, and $2\left(m-m_{s}-m_{\text {mid }}\right)$, respectively. It also constructs a pairing $P_{\text {mid }}$ of the points in $X_{\text {mid }}$. We show that from $P_{O P T}$ we can obtain also a pairing $P_{\text {short }}$ of $X_{\text {short }}$ and a pairing $P_{\text {long }}$ of $X_{\text {long }}$, and modify our initial association of targets so that:

1. Each target in $Y_{\text {mid }}$ is associated with a pair in $P_{\text {mid }}$ which is at most $(1-\epsilon)$ times shorter than the pair it was initially associated with.
2. Each target in $Y_{\text {short }}$ is associated with a pair in $P_{\text {short }}$ that is at least as long as the pair it was associated with.
3. Each target in $Y_{\text {long }}$ is associated with a pair in $P_{\text {long }}$ that is shorter than the original pair it was associated with by at most an additive $\epsilon 2 n \beta$.
(In addition, baselines in $P_{\text {short }}$ have length at most $(1+\epsilon) \frac{\beta}{2 n}$, and this explains the upper bound in the recursive call minRatioPair $\left(X_{\text {short }}, Y_{\text {short }}, \ell,(1+\epsilon) \frac{\beta}{2 n}\right)$. Similarly for the other recursive call.)

At this point, we have finalized our special association for targets in $Y_{\text {mid }}$. This is the association in (1) above; notice that we lose only a $(1-\epsilon)$ factor. For the targets in $Y_{\text {short }}$, we "recursively" construct the special association in this way starting with the new association with $P_{\text {short }}$ and following the recursive call minRatioPair $\left(X_{\text {short }}, Y_{\text {short }}, \ell,(1+\epsilon) \frac{\beta}{2 n}\right)$. Similarly for the targets in $Y_{\text {long }}$.

Notice that the baselines associated with targets in $Y_{\text {short }}$ have not shrunk; the baselines associated with targets in $Y_{\text {long }}$ may have shrunk, but by at most an additive $\epsilon 2 n \beta$. Now for a given target $t$, how much shrinkage can it experience as we recursively construct our special association? This is no more than the sum of the $\epsilon 2 n \beta$ terms over all the $\beta$ 's that it sees and for which it is in $Y_{\text {long }}$. Now, Lemma 4 implies that this can be bounded by a geometric series that is at most $4 \epsilon n \beta$ for the largest $\beta$ it is in $Y_{\text {long }}$ for. And this in turn means that in our special association, $t$ is associated with a baseline that is at least $(1-5 \epsilon)$ times the corresponding baseline in $P_{O P T}$.

Lemma 4. Consider two recursive calls made by the algorithm where the second recursive call is nested (possibly by several levels) within the first. Let $\beta_{1}$ denote the choice of $\beta$ in the first call within which the second call is contained, and let $\beta_{2}$ denote any choice of $\beta$ within the second call. Then either $\beta_{2} \geq 2 \beta_{1}$ or $\beta_{2} \leq \beta_{1} / 2$.

Proof. The immediate recursive calls that we make from the first call with $\beta_{1}$ either have an upper bound of $(1+\epsilon) \frac{\beta_{1}}{2 n}$ or a lower bound of $(1-\epsilon) 2 n \beta_{1}$. Now lower bounds do not decrease with recursion in our algorithm, so $\beta_{2}>(1-\epsilon) 2 n \beta_{1} \geq 2 \beta_{1}(\epsilon<1 / 2)$ if the second call is nested within an immediate call of the latter type. Upper bounds may increase with recursion but only by a factor
of $(1+\epsilon)$, and since the depth of recursion is less than $\log n$, we have $\beta_{2} \leq(1+\epsilon)^{\log n} \frac{\beta_{1}}{2 n} \leq \beta_{1} / 2$ if the second call is nested within an immediate call of the former type.

Theorem 5. There is an algorithm for MINSUMOFRATIOS that, for any parameter $0<\epsilon<1$, runs in quasi-polynomial time and returns a solution whose cost is at most $(1+\epsilon)$ times that of the optimal solution.

## References

[1] Esther M. Arkin and Refael Hassin. On local search for weighted k-set packing. In Rainer E. Burkard and Gerhard J. Woeginger, editors, ESA, volume 1284 of Lecture Notes in Computer Science, pages 13-22. Springer, 1997.
[2] Beat Gfeller, Matús Mihalák, Subhash Suri, Elias Vicari, and Peter Widmayer. Angle optimization in target tracking. In Joachim Gudmundsson, editor, SWAT, volume 5124 of Lecture Notes in Computer Science, pages 65-76. Springer, 2008.
[3] R. I. Hartley and A. Zisserman. Multiple View Geometry in Computer Vision. Cambridge University Press, ISBN: 0521623049, 2000.
[4] Volkan Isler, Sanjeev Khanna, John R. Spletzer, and Camillo J. Taylor. Target tracking with distributed sensors: The focus of attention problem. Computer Vision and Image Understanding, 100(1-2):225-247, 2005.
[5] Eugene Lawler. Combinatorial Optimization: Networks and Matroids. Dover Publications (Originally published in 1976), 2001.
[6] Frits C. R. Spieksma. Chapter 1: Multi index assignment problems: Complexity, approximation, applications.
[7] Frits C. R. Spieksma and Gerhard J. Woeginger. Geometric three-dimensional assignment problems. European Journal of Operational Research, 91(3):611-618, June 1996.
[8] Alper Yilmaz, Omar Javed, and Mubarak Shah. Object tracking: A survey. ACM Comput. Surv., 38(4):13, 2006.

## A Proof of Lemma 2

Denote angle $(t x y)$ by $\alpha$. For greater values of $\alpha$, the rays $\overrightarrow{t y}$ and $\overrightarrow{x y}$ meet further apart and in the limiting case i.e. for $\alpha=\pi-\theta$ these become parallel. Thus, we are interested in the range $0<\alpha<\pi-\theta$. First we shall show that within this range as $\alpha$ increases, the ratio $\frac{\left|x z^{\prime}\right|}{|x z|}$ strictly increases. Then we shall show that in the limiting case - i.e. when $\alpha \rightarrow \pi-\theta$, the ratio $\frac{\left|x z^{\prime}\right|}{|x z|} \rightarrow \frac{1}{\epsilon^{2}}$. This will complete the proof.

For the first part, refer back to Figure 6. Here we consider two different values of $\alpha$ such that $0<\alpha_{1}<\alpha_{2}<\pi-\theta$. Let $y_{1}, z_{1}, z_{1}^{\prime}$ and $y_{2}, z_{2}, z_{2}^{\prime}$ denote the corresponding points as defined before for these two cases.


Figure 10: Illustration for Lemma 2. In the limiting case, the ratio $\frac{\left|x z^{\prime}\right|}{|x z|}$ is at most $\frac{1}{\epsilon^{2}}$.

Now consider the line through $z_{1}^{\prime}$ that is parallel to $x y_{2}$. Let this line intersect the line $t z_{2}$ at point $w$. Now clearly $\triangle z_{1} x z_{2}$ and $\triangle z_{1} z_{1}^{\prime} w$ are similar since two of their angles are equal. Hence, $\frac{\left|z_{1}^{\prime} w\right|}{\left|x z_{2}\right|}=\frac{\left|z_{1} z_{2}^{\prime}\right|}{\left|z_{1} x\right|}$. But clearly $\left|z_{2} z_{2}^{\prime}\right|>\left|z_{1}^{\prime} w\right|$. Hence, we have $\frac{\left|z_{2} z_{2}^{\prime}\right|}{\left|x z_{2}\right|}>\frac{\left|z_{1} z_{1}^{\prime}\right|}{\left|x z_{1}\right|}$. This clearly implies $\frac{\left|x z_{2}^{\prime}\right|}{\left|x z_{2}\right|}>\frac{\left|x z_{1}^{\prime}\right|}{\left|x z_{1}\right|}$. Thus as $\alpha$ increases our ratio increases.

For the second part, refer to Figure 10. We consider the limiting case in which $\alpha=\pi-\theta$. As we observed before in this case rays $\overrightarrow{t y}$ and $\overrightarrow{x y}$ become parallel. Let $y^{\prime}$ be the point on $t y$ at a distance $\left|x z^{\prime}\right|$ from $t$, so that we have the parallelogram $t x z^{\prime} y^{\prime}$. In this scenario, $\Delta t x z$ is similar to $\Delta t y^{\prime} z^{\prime}$ (The Figure is not to scale). This is because angle $(x t z)=\operatorname{angle}\left(y^{\prime} t z^{\prime}\right)=\epsilon \cdot \theta$ and angle $(t x z)=$ angle $\left(t y^{\prime} z^{\prime}\right)$ since these are opposite angles of parallelogram $t x z^{\prime} y^{\prime}$. Now, due to similarity of triangles $\triangle t x z$ and $\Delta t y^{\prime} z^{\prime}, \frac{|x z|}{\left|y^{\prime} z^{\prime}\right|}=\frac{|t x|}{\left|t y^{\prime}\right|}$. Hence, $\left|t y^{\prime}\right|=\frac{|t x|\left|y^{\prime} z^{\prime}\right|}{|x z|}$. But note that $\left|t y^{\prime}\right|=\left|x z^{\prime}\right|$ and $\left|y^{\prime} z^{\prime}\right|=|t x|$ since these are opposite sides of the parallelogram.

After plugging in these values we have $\frac{\left|x z^{\prime}\right|}{|x z|}=\frac{|t x|^{2}}{|x z||x z|}=\frac{|t x|^{2}}{|x z|^{2}}$. By applying the sine rule in $\triangle t x z$, we have $\frac{|t x|}{\sin \text { angle }(t z x)}=\frac{|x z|}{\sin \text { angle }(x t z)}$. But note that angle $(x t z)=\epsilon \cdot \theta$ and angle $(t z x)=$ angle $(z t y)=(1-\epsilon) \cdot \theta$ since these form a pair of alternate angles between parallel lines $t y^{\prime}$ and $x z^{\prime}$ cut by transversal $t z$.

Hence we have $\left(\frac{|t x|}{|x z|}\right)^{2}=\left(\frac{\sin (1-\epsilon) \theta}{\sin \epsilon \cdot \theta}\right)^{2} \leq\left(\frac{1-\epsilon}{\epsilon}\right)^{2} \leq \frac{1}{\epsilon^{2}}$.
Thus $\frac{\left|x z^{\prime}\right|}{|x z|} \leq \frac{1}{\epsilon^{2}}$

