

The Focus of Attention Problem

Dries Goossens · Sergey Polyakovskiy ·
Frits C. R. Spieksma · Gerhard J. Woeginger

Received: 14 October 2013 / Accepted: 4 December 2014 / Published online: 11 December 2014
© Springer Science+Business Media New York 2014

Abstract We consider systems of small, cheap, simple sensors that are organized in a distributed network and used for estimating and tracking the locations of targets. The objective is to assign sensors to targets such that the overall expected error of the sensors' estimates of the target locations is minimized. The so-called *focus of attention* problem (FOA) deals with the special case where every target is tracked by one pair of range sensors. The resulting computational problem is a special case of the axial three-index assignment problem, a well-known fundamental problem in combinatorial optimization. We provide a complete complexity and approximability analysis of the FOA problem: we establish strong NP-hardness and the unlikeliness

A preliminary version of this work has appeared in the Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2010), 312–317. This research has been supported by the Netherlands Organisation for Scientific Research (NWO) under Grant 639.033.403, by BSIK Grant 03018 (BRICKS: Basic Research in Informatics for Creating the Knowledge Society), and by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office.

D. Goossens
Gent University, Ghent, Belgium
e-mail: dries.goossens@ugent.be

S. Polyakovskiy
School of Computer Science, The University of Adelaide, Adelaide, SA 5005, Australia
e-mail: sergey.polyakovskiy@adelaide.edu.au

F. C. R. Spieksma (✉)
Operations Research Group, KU Leuven, Naamsestraat 69, 3000 Leuven, Belgium
e-mail: frits.spieksma@kuleuven.be

G. J. Woeginger
Department of Mathematics, TU Eindhoven, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
e-mail: gwoegi@win.tue.nl

of an FPTAS, we identify polynomially solvable special cases, and we construct a sophisticated polynomial time approximation scheme for it.

Keywords Target tracking · Distributed sensors · Sensor assignment · Assignment problem · Complexity · Approximation · Intractable problem

1 Introduction

Sensors are everywhere. The presence and use of monitoring devices has become standard in Western societies. For example city centers, shopping malls, and other public places are continuously being monitored by cameras. Usually such sensors are organized in a network where they act jointly in order to perform a common task; such systems are referred to as *sensor networks*. Sensor networks typically consist of many small, inexpensive, low power, untethered devices that observe various environmental parameters. A sensor network is capable of a real-time, fine-grained monitoring of the surroundings. Such systems are relatively cheap, they are robust, and they are increasingly being deployed in practice. We refer the reader to Culler et al. [5] or Tubaishat and Madria [12] for an overview of the developments in sensor networks.


In order to realize the potential of sensor networks, there are at least two fundamental challenges that need to be addressed. The first challenge comes from the inherent limitations of individual sensors, as an individual sensor alone is incapable of estimating the state of a target. The second challenge arises as the measurements provided by the sensors are strongly corrupted by noise. To overcome these challenges, sensors must cooperate and groups of sensors are used to estimate the position of a single target. Central questions are which sensors should be assigned to which targets, and which measurements should be combined in order to get accurate estimates. These choices will directly determine the quality of the system.

Isler et al. [9] consider the following concrete (and of course highly simplified) scenario for target tracking in distributed systems. There are $2n$ sensors that are to be assigned in disjoint pairs to n targets. The sensors are located on a straight line, whereas the (static) targets are positioned somewhere in the plane. Without loss of generality the straight line is the x -axis, so that the $2n$ sensors are positioned in $2n$ points with coordinates $(x_i, 0)$ with $1 \leq i \leq 2n$; see Fig. 1 for an illustration. Isler et al. [9] discuss an error measure that is motivated by stereo reconstruction that mainly depends on the y -coordinates y_1, \dots, y_n of the n targets: if the i th and the j th sensor together are assigned to the k th target, then the corresponding incurred error cost amounts to

$$c_{ijk} = \frac{y_i}{|x_j - x_k|}. \quad (1)$$

Isler et al. [9] argue that the measure in (1) gives a good error approximation in case the targets are not too close to the sensors. For more information on this measure and in particular for its mathematical justification, we refer the reader to Appendix A of [9]; notice that they also consider other cost metrics as well. The objective in the *focus of attention* problem is to find an assignment of disjoint sensor pairs to targets such

Fig. 1 Six sensors/cameras
 C_1, \dots, C_6 track three targets
 t_1, t_2, t_3



that the sum of all error costs c_{ijk} is minimized. We denote this optimization problem as IKST-FOA, for short.

Isler et al. [9] derive a polynomial time 2-approximation algorithm for IKST-FOA. In the *equi-distant* special case of IKST-FOA, the sensors are at unit distances from each other in the $2n$ points $(i, 0)$ with $1 \leq i \leq 2n$. For this equi-distant special case of IKST-FOA, [9] design a non-trivial polynomial time approximation scheme (PTAS).

Formulation of problem p -FOA We will investigate a certain version of the three-dimensional assignment problem that contains problem IKST-FOA as a special case. This version is based on a real parameter p , and will throughout be denoted as p -FOA. An instance of p -FOA consists of $3n$ positive real numbers $a_1, \dots, a_n, b_1, \dots, b_n$, and c_1, \dots, c_n . The cost-coefficient corresponding to a triple (i, j, k) with $1 \leq i, j, k \leq n$ is defined as

$$c_{ijk} = \frac{a_i}{(b_j + c_k)^p}. \tag{2}$$

The goal in p -FOA is to group the $3n$ numbers into n triples (where each triple contains one a_i , one b_j and one c_k) such that the sum of the cost-coefficients corresponding to these triples becomes minimum. In Sect. 2 we show that for $p = 1$ this problem p -FOA coincides with the classic target tracking problem IKST-FOA as discussed above.

Our results We completely analyze the complexity and approximability behavior of problem p -FOA for every value of the parameter p . Sections 3 and 4 provide the following complexity classification of problem p -FOA:

- For every real p with $-1 \leq p \leq 0$, problem p -FOA is polynomially solvable.
- For every real p with $p < -1$ or $p > 0$, problem p -FOA is strongly NP-hard.
- Even the equi-distance special case of IKST-FOA is strongly NP-hard. This settles a question left open in [9].

On the approximation side, Sect. 5 presents a fast and simple approximation algorithm for p -FOA with worst case performance guarantee of 2^p for the cases where $p \geq 0$; this result is based on the methodology developed in [9]. As our main contribution, we derive an approximability result in Sect. 6 and a complementary inapproximability result in Sect. 3 that together fully resolve the approximability status of p -FOA:

- For every real p with $p < -1$ or $p > 0$, problem p -FOA possesses a PTAS.

- For every real p with $p < -1$ or $p > 0$, problem p -FOA does not possess an FPTAS unless $P=NP$.

In many cases, the development of a (fully) polynomial time approximation scheme has nowadays become a straightforward exercise. Indeed, when an optimization problem satisfies certain conditions (see Woeginger [13]), the existence of such a scheme follows automatically. We stress however that our problems here do not fall into that category; in fact, the design of our PTAS is quite intricate, and introduces a number of new ideas to the area.

We also stress that the proof of our inapproximability result is not done by routine methods. The literature contains a number of negative results (see for instance Theorem 6.8 in Garey and Johnson [7]) showing that a strongly NP-hard and sufficiently well-behaved optimization problem cannot have an FPTAS unless $P=NP$. Here well-behaved means that (i) all solution values are positive integers and that (ii) the value of an optimal solution is polynomially bounded in the size of a unary encoding of the instance. Note that the theorem cannot be applied directly to problem p -FOA, as our objective values in general are not integral and hence violate condition (i). Our inapproximability proof is based on the standard approach from the literature, but on top of that introduces an additional trick for working around integrality.

Further links to the literature The literature contains a number of results on target tracking where sensors are to be assigned in pairs to targets. There are various ways of modeling the measurement errors and the resulting error costs, and we only mention two results that have a strong algorithmic component. Gfeller et al. [6] discuss scenarios where the error mainly depends on the intersection angle of the two viewing cones subtended by a pair of sensors. Al-Hasan et al. [1] consider a related scenario with moving sensors; they introduce an intricate cost model for the movements of the sensors, and they develop a GRASP routine for cost minimization in this model.

Assignment problems have received much attention in the literature; see for instance the book by Burkard et al. [2]. IKST-FOA is related to the (axial) three-index assignment problem (3AP), see Chapter 10 of [2]. Our results here fall into the research branch that concentrates on the algorithmic behavior of strongly structured special cases. Let us mention some results that discuss complexity and approximability of 3AP. Burkard et al. [3] discuss an NP-hard special case where the cost-coefficients are given by $c_{ijk} = a_i b_j c_k$, for positive real numbers $a_1, \dots, a_n, b_1, \dots, b_n$, and c_1, \dots, c_n . Crama and Spieksma [4] design polynomial time $\frac{4}{3}$ -approximation algorithms for special cases where the cost coefficients are decomposable and satisfy a certain type of triangle inequality, and Spieksma and Woeginger [11] establish NP-hardness of the corresponding Euclidean special case. Queyranne and Spieksma [10] exhibit a 2-approximation algorithm for a related problem, again assuming a certain type of triangle inequality.

2 The Connection Between IKST-FOA and 1-FOA

Consider an instance of IKST-FOA that is specified by $2n$ real numbers x_1, \dots, x_{2n} and by n real numbers y_1, \dots, y_n , with costs defined as in (1). Assume that the sensors

on the x -axis are ordered as

$$x_1 \leq x_2 \leq \dots \leq x_{2n}.$$

A feasible solution is called *left-right separating*, if it matches every sensor from the left half $1, \dots, n$ with one sensor from the right half $n + 1, \dots, 2n$ (and with some target). We stress that the essence of the following Lemma 2.1 is due to Isler et al. [9].

Lemma 2.1 *There exists an optimal solution for IKST-FOA that is left-right separating.*

Proof A feasible solution is specified by a permutation π of $1, \dots, 2n$, such that sensors $\pi(2k - 1)$ and $\pi(2k)$ are assigned to target k for $1 \leq k \leq n$. Among all optimal solutions, consider one solution π that maximizes the auxiliary function $\sum_{k=1}^n |\pi(2k - 1) - \pi(2k)|$. If π is not left-right separating, it must—at least once—match two sensors from the left half (say $\pi(1)$ and $\pi(2)$), and it must at least once match two sensors from the right half (say $\pi(3)$ and $\pi(4)$). Assume without loss of generality that

$$x_{\pi(1)} \leq x_{\pi(2)} \leq x_{\pi(3)} \leq x_{\pi(4)}.$$

Then $|x_{\pi(2)} - x_{\pi(1)}| \leq |x_{\pi(3)} - x_{\pi(1)}|$ and $|x_{\pi(4)} - x_{\pi(3)}| \leq |x_{\pi(4)} - x_{\pi(2)}|$. Therefore switching the values $\pi(2)$ and $\pi(3)$ in π will not worsen the objective value, whereas it does increase the auxiliary function. This contradiction completes the argument. \square

Next, let x^* with $x_n \leq x^* \leq x_{n+1}$ be a real number that separates the sensors in the left half from the sensors in the right half. Then the IKST-FOA instance can be rewritten as an instance of 1-FOA in the following way: Let a_1, \dots, a_n denote the positive real numbers y_1, \dots, y_n ; let b_1, \dots, b_n denote the positive real numbers $x^* - x_1, \dots, x^* - x_n$; let c_1, \dots, c_n denote the positive real numbers $x_{n+1} - x^*, \dots, x_{2n} - x^*$. Define the cost-coefficient corresponding to a triple (i, j, k) as in (2).

Vice versa, every instance of 1-FOA can be rewritten as an instance of IKST-FOA, if a_1, \dots, a_n play the role of y_1, \dots, y_n , and if b_1, \dots, b_n together with $-c_1, \dots, -c_n$ play the role of x_1, \dots, x_{2n} . Notice that these constructions map a solution to an instance of IKST-FOA to a solution of an instance of 1-FOA with the same value, and vice versa; we call such a pair of instances equivalent. We summarize our findings in the following theorem.

Theorem 2.2 *The problems IKST-FOA and 1-FOA are equivalent in the following sense: there exist simple linear time reductions that translate an instance of one problem into an instance of the other problem with the same optimal objective value. Furthermore, any polynomial time approximation algorithm for one problem yields a polynomial time approximation algorithm for the other problem with the same worst case guarantee.*

3 Hardness and Inapproximability

In this section we establish strong NP-hardness and inapproximability (with respect to fully polynomial time approximation schemes) of p -FOA. We first recall the following

formulation (3) of the Hölder inequality; see for instance Theorem 13 in the book [8] of Hardy, Littlewood & Pólya. For a non-zero real number q with $q < 1$, and for $2n$ positive real numbers $\alpha_1, \dots, \alpha_n$ and β_1, \dots, β_n we have

$$\sum_{\ell=1}^n \alpha_\ell^{1/q} \beta_\ell^{(q-1)/q} \geq \left(\sum_{\ell=1}^n \alpha_\ell \right)^{1/q} \left(\sum_{\ell=1}^n \beta_\ell \right)^{(q-1)/q} \tag{3}$$

Most importantly, equality holds in (3) if and only if the sequences α_i and β_i are proportional, that is, if and only if there exists a real number λ such that $\alpha_i = \lambda\beta_i$ for $1 \leq i \leq n$. We will use these facts in the following hardness proofs.

3.1 Hardness and Inapproximability of p-FOA

Throughout this section, let p be some fixed real number with $p < -1$ or $p > 0$. Our reduction is from the strongly NP-hard problem Numerical Matching with Target Sums (NMTS); see Garey and Johnson [7]. Given target sums A_k ($1 \leq k \leq n$), and given positive integers B_i ($1 \leq i \leq n$) and C_j ($1 \leq j \leq n$), can we find a collection of n triples (i, j, k) such that $A_k = B_i + C_j$ holds for each triple and such that each element is used exactly once? Without loss of generality, we assume that the sum $S := \sum A_k$ equals $\sum B_i + \sum C_j$.

We consider an instance of NMTS, and transform it into an instance I of p -FOA by setting $a_k := A_k^{p+1}$, and $b_i := B_i$, and $c_j := C_j$ for all $1 \leq i, j, k \leq n$. The optimal objective value of instance I is denoted $\text{OPT}(I)$. We claim that $\text{OPT}(I) \leq S$ if and only if the NMTS instance has answer YES.

Lemma 3.1 *If the NMTS instance has answer YES, then $\text{OPT}(I) \leq S$.*

Proof We interpret the triples in the solution for NMTS as a feasible solution for p -FOA. Then any triple (i, j, k) with $A_k = B_i + C_j$ in this feasible solution contributes $A_k^{p+1} / (B_i + C_j)^p = A_k$ to the objective value. Hence, the corresponding objective value for p -FOA equals S . □

Lemma 3.2 *If $\text{OPT}(I) \leq S$, then the NMTS instance has answer YES.*

Proof We interpret the triples in the feasible solution for p -FOA with cost at most S as a feasible solution for NMTS. We use inequality (3) with $q = 1/(p + 1)$; note that for $p < -1$ and for $p > 0$, the corresponding q indeed satisfies $q < 1$. Furthermore, we set $\alpha_\ell = A_k$ and $\beta_\ell = B_i + C_j$ in (3), where j and k are the indices that occur together with index i in the ℓ th triple (i, j, k) in the feasible solution. For the objective value this then yields

$$\begin{aligned} S &\geq \sum A_k^{p+1} (B_i + C_j)^{-p} \\ &\geq \left(\sum A_k \right)^{p+1} \left(\sum B_i + C_j \right)^{-p} = S^{p+1} \cdot S^{-p} = S. \end{aligned} \tag{4}$$

Hence all inequalities in this chain are actually equalities. As we are dealing with the case of equality in (3), the values $\alpha_\ell = A_k$ and $\beta_\ell = B_i + C_j$ must be proportional to each other. Since $\sum \alpha_\ell = \sum \beta_\ell$, the factor λ of proportionality is $\lambda = 1$. This yields $A_k = B_i + C_j$ for all triples (i, j, k) in the feasible solution, so that the NMTS instance has answer YES. \square

Lemmas 3.1 and 3.2 establish the correctness of our reduction from NMTS, and hence yield the following theorem.

Theorem 3.3 *For all real $p < -1$ and for all real $p > 0$, problem p -FOA is strongly NP-hard.*

Now let us turn to the inapproximability result. Lemma 3.2 essentially states that if the NMTS instance has answer NO, then $\text{OPT}(I) > S$. By looking a little bit deeper into the proofs of the above lemmas, we will establish the following strengthening.

Lemma 3.4 *If the NMTS instance has answer NO, then $\text{OPT}(I) > S + S^{-2p}$.*

Proof Let π and σ denote two permutations of $1, \dots, n$ that yield the optimal objective value for the p -FOA instance I .

$$\text{OPT}(I) = \sum_{k=1}^n \frac{a_k}{b_{\pi(k)} + c_{\sigma(k)}} = \sum_{k=1}^n \frac{A_k^{p+1}}{(B_{\pi(k)} + C_{\sigma(k)})^p} > S. \tag{5}$$

If $A_k = B_{\pi(k)} + C_{\sigma(k)}$ holds for all k , then we get the contradiction $\text{OPT}(I) = S$ from the proof of Lemma 3.1. Hence, we will assume without loss of generality that $A_1 < B_{\pi(1)} + C_{\sigma(1)}$. To simplify notation, we set $x := A_1$ and $y := B_{\pi(1)} + C_{\sigma(1)}$ and we note that $1 \leq x < y < S$. Then $\sum_{k=2}^n A_k = S - x$ and $\sum_{k=2}^n (B_{\pi(k)} + C_{\sigma(k)}) = S - y$, and in an analogous fashion as in (4) the Hölder inequality (3) with $q = 1/(p + 1)$ yields

$$\sum_{k=2}^n \frac{A_k^{p+1}}{(B_{\pi(k)} + C_{\sigma(k)})^p} \geq \frac{(S - x)^{p+1}}{(S - y)^p}. \tag{6}$$

We conclude from (5) and (6) that

$$\text{OPT}(I) \geq \frac{x^{p+1}}{y^p} + \frac{(S - x)^{p+1}}{(S - y)^p}. \tag{7}$$

Let us consider the right hand side of (7) as a function $f(x)$ of variable x . It is easily seen that this function is strictly convex, and that it attains its unique minimum in the point $x = y$ with the value $f(y) = S$. Now in our situation x is an integer from the range $1 \leq x < y$, which implies $f(x) \geq f(y - 1) > S$. Finally, we note that the value $f(y - 1)$ is a rational number whose denominator is at most $y^p(S - y)^p < S^{2p}$. Since the smallest rational number above S with such a denominator is greater than $S + S^{-2p}$, we conclude $f(y - 1) > S + S^{-2p}$. Then (7) yields the desired lower bound $\text{OPT}(I) > S + S^{-2p}$. \square

Theorem 3.5 *For all real $p < -1$ and for all real $p > 0$, problem p -FOA does not possess an FPTAS (unless $P=NP$).*

Proof Suppose that p -FOA does possess an FPTAS. We take a p -FOA instance I as in the above NP-hardness proof, and we execute the FPTAS with an approximation guarantee of $\varepsilon = S^{-2p-1}$. The time complexity of the resulting algorithm is polynomially bounded in the instance size and in $1/\varepsilon = S^{2p+1}$. As p is a fixed real number, this time complexity is pseudo-polynomially bounded in the instance size.

We let V denote the objective value found by the FPTAS. We claim that $V \leq S + S^{-2p}$ if and only if the underlying instance of NMTS has answer YES. Indeed, if the NMTS instance has answer YES, then $\text{OPT}(I) \leq S$ by Lemma 3.1. Since the objective value V is at most a factor $1 + \varepsilon$ above the optimal objective value, we conclude that

$$V \leq (1 + \varepsilon) \text{OPT}(I) \leq (1 + S^{-2p-1}) S = S + S^{-2p}.$$

On the other hand if the NMTS instance has answer NO, then Lemma 3.4 implies

$$V \geq \text{OPT}(I) > S + S^{-2p}.$$

Hence, by looking at the output of the FPTAS we could decide in pseudo-polynomial time whether the NMTS instance has answer YES (in case $V \leq S + S^{-2p}$) or answer NO (in case $V > S + S^{-2p}$). This yields a pseudo-polynomial time decision algorithm for a strongly NP-hard problem, and consequently implies $P=NP$. \square

3.2 Hardness and Inapproximability of IKST-FOA

Next, let us discuss the equi-distant special case of IKST-FOA where the sensors are at unit distances from each other in the $2n$ points $(i, 0)$ with $1 \leq i \leq 2n$. The equivalent instance of 1-FOA has $b_i = c_i = i - \frac{1}{2}$ for $i = 1, \dots, n$.

We use a similar reduction and the same notation as in the preceding section. Yu et al. [14] have shown that Numerical Matching with Target Sums is NP-hard even if $B_i = C_i = i$ holds for $i = 1, \dots, n$. We start with an NMTS instance A_k, B_i, C_j ($1 \leq i, j, k \leq n$) of this particular form, and define a new (equivalent) NMTS instance with $A'_k = A_k - 1$, $B'_i = B_i - \frac{1}{2}$ and $C'_j = C_j - \frac{1}{2}$. Then $B'_i = C'_i = i - \frac{1}{2}$ holds for all i , and the reduction in Theorem 3.3 for $p = 1$ yields the desired NP-hardness argument. Also the inapproximability argument goes through in the same way as before in Lemma 3.4 and Theorem 3.5.

Theorem 3.6 *The equi-distant special case of IKST-FOA is strongly NP-hard, and it does not possess an FPTAS (unless $P=NP$).*

4 A Polynomial Time Result for p -FOA

In this section we discuss the parameter range $-1 \leq p \leq 0$ for p -FOA with cost coefficients of the form (2). These problems are almost trivial and can essentially be

solved by sorting. The following lemma settles the case with input sequences of length $n = 2$.

Lemma 4.1 *Let $a_1 \leq a_2$ and $b_1 \geq b_2$, $c_1 \geq c_2$ be six positive real numbers that form an instance of p -FOA with $-1 \leq p \leq 0$. Then the matching a_1, b_1, c_1 and a_2, b_2, c_2 forms an optimal solution.*

Proof For any real $s \geq 0$, the function $f(x) = a_1x^{-p} + a_2(s - x)^{-p}$ is concave on the range $0 < x \leq s$. This implies that in the p -FOA instance the minimum cost is attained on the boundary of the domain, and that the optimal solution either matches a_1, b_1, c_1 and a_2, b_2, c_2 , or a_1, b_2, c_2 and a_2, b_1, c_1 . Since $-1 \leq p \leq 0$ implies

$$(a_2 - a_1) \left(\frac{1}{(b_1 + c_1)^p} - \frac{1}{(b_2 + c_2)^p} \right) \geq 0,$$

the first one of these two candidate solutions gives the minimum cost. □

Repeated application of Lemma 4.1 now yields the following theorem.

Theorem 4.2 *Let $a_1 \leq \dots \leq a_n$ together with $b_1 \geq \dots \geq b_n$ and $c_1 \geq \dots \geq c_n$ form an instance of p -FOA with $-1 \leq p \leq 0$. Then an optimal solution is given by the triples (i, i, i) with $1 \leq i \leq n$.*

5 A Simple Approximation Algorithm for p -FOA

In this section we discuss the approximability of p -FOA with cost coefficients of the form (2). Without loss of generality we assume that the numbers a_1, \dots, a_n are in non-decreasing order

$$a_1 \leq a_2 \leq \dots \leq a_n. \tag{8}$$

The well-known *rearrangement inequality* (see for instance Theorem 368 in the book [8] of Hardy, Littlewood & Pólya) states the following: If two finite sequences $\langle \alpha_i \rangle$ and $\langle \beta_i \rangle$ of equal length, are given except in arrangement, then the value of the sum $\sum_i \alpha_i \beta_i$ is minimum if the two sequences are monotonic in opposite order. An immediate consequence of the rearrangement inequality and of (8) is that any reasonable feasible solution of p -FOA with triples $(k, \pi(k), \sigma(k))$ for $1 \leq k \leq n$ satisfies

$$b_{\pi(1)} + c_{\sigma(1)} \leq b_{\pi(2)} + c_{\sigma(2)} \leq \dots \leq b_{\pi(n)} + c_{\sigma(n)}. \tag{9}$$

Indeed, if one of the inequalities in (9) would be violated, then rearranging the sums $b_{\pi(k)} + c_{\sigma(k)}$ into non-decreasing order would improve the objective value. In particular, any optimal solution will satisfy (9).

Isler et al. [9] consider the following simple polynomial time approximation algorithm for 1-FOA, and they establish that it has a worst case performance guarantee of 2:

1. For $k = 1, \dots, n$ match the k th largest number among b_1, \dots, b_n with the k th smallest number among c_1, \dots, c_n . Denote the n resulting sums $b_i + c_j$ by $s_1 \geq s_2 \geq \dots \geq s_n$.

2. Match the sums s_1, \dots, s_n according to the rearrangement inequality with the numbers a_1, \dots, a_n .

This algorithm can also be applied to instances of the general p -FOA problem with $p \geq 0$. The only minor modification is that in Step 2 we now assign the sums $s_1^p \geq s_2^p \geq \dots \geq s_n^p$ according to the rearrangement inequality to the numbers a_1, \dots, a_n . Note that, compared to the case $p = 1$, only the objective value changes, whereas the set of triples $(k, \pi(k), \sigma(k))$ with $1 \leq k \leq n$ (and hence the feasible solution) remains the same.

Now consider an optimal solution for some p -FOA instance, and denote by $s_1^* \geq s_2^* \geq \dots \geq s_n^*$ the corresponding sums $b_i + c_j$ that the optimal solution matches with the numbers a_1, \dots, a_n . The following result is essentially due to [9].

Lemma 5.1 *For $1 \leq k \leq n$, these sums satisfy the inequality $s_k^* \leq 2s_k$.*

Proof Let the multi-set T_k contain the k largest values among b_1, \dots, b_n and c_1, \dots, c_n , and let t_k denote the smallest value in T_k . Then the first step of the algorithm ensures $s_k \geq t_k$. If one of the k sums s_1^*, \dots, s_k^* matches two values in T_k with each other, then another such sum must match two values outside T_k with each other; this sum is at most $2t_k$. If each of the k sums s_1^*, \dots, s_k^* matches some value in T_k with some value outside T_k , then the smallest such sum is at most $2t_k$. In either case we have the desired inequality $s_k^* \leq 2t_k \leq 2s_k$. \square

Theorem 5.2 *For every $p \geq 0$, the above polynomial time approximation algorithm for p -FOA has a worst case performance guarantee of 2^p . This bound is tight.*

Proof Whenever the optimal solution uses a cost coefficient $a_i/(s_k^*)^p$, the approximate solution uses a cost coefficient $a_i/(s_k)^p \leq a_i/(s_k^*/2)^p$. By summing these inequalities, we get that the approximate objective value is at most a factor of 2^p above the optimal objective value.

Tightness of the bound 2^p can be seen from the instance $a_1 = b_1 = c_1 = 1$, $a_2 = t^2$, and $b_2 = c_2 = t$ for some huge number t . Then the approximation algorithm matches the numbers as $(1, 1, t)$ and $(t^2, t, 1)$, whereas the optimal solution matches the numbers as $(1, 1, 1)$ and (t^2, t, t) . As t tends to infinity, the ratio between the two objective values tends to 2^p . \square

For parameter values $p < -1$, the above approximation algorithm does not have a finite performance guarantee. For the instance $a_1 = b_1 = c_1 = 1$ and $a_2 = b_2 = c_2 = t$ with huge t , the optimal solution would match the numbers as $(t, 1, 1)$ and $(1, t, t)$ with objective value $t2^{-p} + (2t)^{-p}$. The approximation algorithm matches the numbers as $(1, 1, t)$ and $(t, 1, t)$ with objective value $(t + 1)^{1-p}$. A similar example shows that even for the polynomially solvable cases of p -FOA with $-1 \leq p < 0$, the approximation algorithm does not have a finite performance guarantee.

6 An Approximation Scheme for p -FOA

In this section we derive a polynomial time approximation scheme for p -FOA. To keep the analysis simple, we will throughout concentrate on the cases with $p > 0$.

In Sect. 6.3 we briefly sketch how to settle the remaining cases with negative p by a similar approach.

Hence consider an arbitrary instance I of p -FOA (where cost coefficients are of the form (2)). Without loss of generality we assume that the numbers a_1, \dots, a_n are in non-decreasing order (8). Recall from the preceding section that any reasonable solution of p -FOA will consist of triples $(k, \pi(k), \sigma(k))$ with properties as described in (9).

The worst case guarantee in our PTAS will be of the form $(1 + \varepsilon)^{2p}$, where ε with $0 < \varepsilon < 1/2$ is a fixed real number that can be chosen arbitrarily close to zero. We introduce L as the smallest integer satisfying

$$\varepsilon (1 + \varepsilon)^{L-1} \geq 1. \tag{10}$$

Some straightforward calculations show that L is of order $O((1/\varepsilon) \ln 1/\varepsilon)$. Since ε is a constant whose value does not depend on the input, all expressions that only depend on ε and L will also be fixed constants that are independent of the size of the input.

We start with a rounding phase, in which we round down all the numbers b_1, \dots, b_n and c_1, \dots, c_n in instance I to the next integer power of $1 + \varepsilon$. This rounding is harmless, since it changes the objective value by at most a factor of $(1 + \varepsilon)^p$. Define K as the largest integer for which $(1 + \varepsilon)^K$ occurs among these rounded values b_1, \dots, b_n and c_1, \dots, c_n . We stress that the value of K is polynomially bounded in the input size and in the reciprocal value of ε : If z is the maximum value among the b_1, \dots, b_n and c_1, \dots, c_n , then K is $O(\ln(z)/\varepsilon)$.

6.1 Definition of the Auxiliary Instances

We introduce a family of auxiliary instances I' that encode certain useful sub-instances of the original instance I . This family has two crucial properties. First, the family is small: It contains only a polynomial number of auxiliary instances. Secondly, the auxiliary instances in this family are easy to approximate: Every instance in the family can be approximated by reducing it to several smaller instances in the family. The appropriate choice of these auxiliary instances is rather delicate, and constitutes the main step in deriving the PTAS.

Part of the structure of an auxiliary instance I' is determined by a quadruple (m, k, β, γ) which is called the *type* of instance I' . The quadruple consists of:

- An integer m with $1 \leq m \leq n$.
- An integer k with $0 \leq k \leq K$.
- Two non-negative integers β and γ with $0 \leq \beta, \gamma \leq m$.

In the following, a real number x will be called *k-small* if $x < \varepsilon(1 + \varepsilon)^{k-1}$, and it will be called *k-medium* if $\varepsilon(1 + \varepsilon)^{k-1} \leq x \leq (1 + \varepsilon)^k$. Every auxiliary instance I' of type (m, k, β, γ) consists of $3m$ real numbers $a'_1, \dots, a'_m, b'_1, \dots, b'_m$, and c'_1, \dots, c'_m that satisfy the following:

- The numbers $a'_1 \leq \dots \leq a'_m$ coincide with a_1, \dots, a_m , that is, they form the m smallest elements in the enumeration (8).

- The list $b'_1 \leq \dots \leq b'_m$ consists of the β largest k -small elements among b_1, \dots, b_n , together with $m - \beta$ arbitrarily chosen k -medium elements from b_1, \dots, b_n .
- The list $c'_1 \leq \dots \leq c'_m$ consists of the γ largest k -small elements among c_1, \dots, c_n , together with $m - \gamma$ arbitrarily chosen k -medium elements from c_1, \dots, c_n .
- At least one of b'_m and c'_m equals $(1 + \varepsilon)^k$.

We note that for some of the types there is no corresponding auxiliary instance, as sequence b_1, \dots, b_n or sequence c_1, \dots, c_n do not contain sufficiently many k -small and k -medium elements. We also stress that the original instance I occurs among the auxiliary instances.

Let us estimate the overall number of auxiliary instances: there are $O(n^3 K)$ quadruples that describe a type. For every type (m, k, β, γ) all values a'_i , all the k -small values b'_i , and all the k -small values c'_i in any instance of that type are fixed. The k -medium values b'_i are integer powers of $1 + \varepsilon$ that lie between the bounds $\varepsilon(1 + \varepsilon)^{k-1}$ and $(1 + \varepsilon)^k$. Inequality (10) yields that they must occur among the $L + 1$ numbers

$$(1 + \varepsilon)^{k-L}, (1 + \varepsilon)^{k-L+1}, \dots, (1 + \varepsilon)^k.$$

Hence there are only $O(n^L)$ possible choices for the k -medium values b'_i . An analogous argument shows that there are only $O(n^L)$ possible choices for the k -medium values c'_i . Altogether this yields a polynomial upper bound of $O(K \cdot n^{2L+3})$ on the number of auxiliary instances.

6.2 Approximation of the Auxiliary Instances

Throughout we denote by $\text{OPT}(I)$ the optimal objective value of instance I . For every auxiliary instance I' , we will compute in polynomial time an approximate objective value $f(I')$ that satisfies

$$\text{OPT}(I') \leq f(I') \leq (1 + \varepsilon)^p \cdot \text{OPT}(I').$$

The computation is done in the order of increasing values of m : Whenever we are handling an auxiliary instance with $3m$ numbers, all auxiliary instances with $3(m - 1)$ numbers have already been settled. The computation of $f(I')$ in the cases with $m = 1$ is trivial.

Now consider an auxiliary instance I' of type (m, k, β, γ) with $m \geq 2$. An optimal solution matches element a'_m with two partners b^* and c^* , and the rearrangement inequality and (9) tell us that the sum $b^* + c^*$ of these two partners must be relatively large. Since (by the definition of an auxiliary instance) at least one of b'_m and c'_m takes the value $(1 + \varepsilon)^k$, we certainly have $b^* + c^* \geq (1 + \varepsilon)^k$, and this means that at least one of b^* and c^* is a k -medium element. Our strategy is to enumerate many cases, and to try out all possibilities for such a k -medium partner b^* or c^* . The case checking covers two possible scenarios.

In the first scenario both partners b^* and c^* are k -medium. Hence we check all $O(L^2)$ possibilities for b^* and c^* . In every check, we remove the corresponding three numbers a'_m, b^*, c^* from the instance I' and thus create a residual instance I'' of type

$(m - 1, k', \beta', \gamma')$ for appropriate integers k', β', γ' . Then $f(I'') + a'_m/(b^* + c^*)^p$ yields a $(1 + \varepsilon)^p$ -approximation for the objective value of the best solution that matches a'_m with b^* and c^* .

In the second scenario one partner, say the partner b^* , is k -small (while c^* is k -medium). Then $b^* + c^* \geq (1 + \varepsilon)^k$ and $b^* < \varepsilon(1 + \varepsilon)^{k-1}$ together imply $c^* > (1 + \varepsilon)^{k-1}$. We conclude $b^* < \varepsilon c^*$, and hence

$$b^* + c^* \leq (1 + \varepsilon) c^* \leq (1 + \varepsilon) (b'_1 + c^*),$$

where b'_1 is the minimum of b'_1, \dots, b'_m . Rewriting this last inequality yields

$$\frac{a'_m}{(b'_1 + c^*)^p} \leq (1 + \varepsilon)^p \frac{a'_m}{(b^* + c^*)^p}. \tag{11}$$

We construct instance I'' by removing a'_m, b'_1, c^* from instance I' , and we construct instance I''' by removing a'_m, b^*, c^* from instance I' . From $b'_1 \leq b^*$ we derive $\text{OPT}(I'') \leq \text{OPT}(I''')$. This yields

$$f(I'') \leq (1 + \varepsilon)^p \cdot \text{OPT}(I'') \leq (1 + \varepsilon)^p \cdot \text{OPT}(I'''). \tag{12}$$

Now how do we proceed in this second scenario? We check all $O(L)$ possibilities for a k -medium partner $c^* > (1 + \varepsilon)^{k-1}$. In every single check, we match element a'_m with the elements c^* and with b'_1 . The residual instance I'' then is of type $(m - 1, k', \beta', \gamma')$ for appropriate integers k', β', γ' . The inequalities (11) and (12) show that

$$f(I'') + \frac{a'_m}{(b'_1 + c^*)^p} \leq (1 + \varepsilon)^p \cdot \left(\text{OPT}(I''') + \frac{a'_m}{(b^* + c^*)^p} \right) = (1 + \varepsilon)^p \cdot \text{OPT}(I').$$

Therefore the value $f(I'') + a'_m/(b'_1 + c^*)^p$ yields a $(1 + \varepsilon)^p$ -approximation for the objective value of the best solution that matches a'_m with b^* and c^* .

In the end, the value $f(I')$ is defined as the best approximation detected in all the explored cases under both scenarios.

6.3 The Approximation Scheme

Let us now summarize the main steps of the approach outlined in the above paragraphs. Consider an arbitrary instance I of p -FOA with $p > 0$.

1. Round down all b_1, \dots, b_n and c_1, \dots, c_n to the next integer power of $1 + \varepsilon$.
2. Enumerate all possible auxiliary instances I' of all possible types (m, k, β, γ) .
3. Determine the value $f(I')$ for every auxiliary instance I' .
4. Output the approximate objective value $f(I')$ for the auxiliary instance I' that coincides with the original instance I .

The running time of this approach is polynomial: The overall number of auxiliary instances is polynomially bounded by $O(K \cdot n^{2L+3})$, and every single value $f(I')$ can

be computed in polynomial time. Also the approximation guarantee $(1 + \varepsilon)^{2p}$ is easy to see: The rounding in Step #1 introduces a multiplicative error of at most $(1 + \varepsilon)^p$, and the computation of the function values $f(I')$ introduces another factor of at most $(1 + \varepsilon)^p$. All in all, this yields the following theorem.

Theorem 6.1 *For all real p , problem p -FOA possesses a PTAS.* □

Finally let us briefly discuss the cases of p -FOA with $p < 0$, for which a PTAS can be constructed in a very similar fashion. We modify the above PTAS in the following way: First, we reverse all inequality-signs in (8). Secondly, in the rounding phase instead of rounding *down* we round all the numbers *up* to the next integer power of $1 + \varepsilon$. Thirdly, in the definition of the auxiliary instances we perform two changes: For the list $b'_1 \leq \dots \leq b'_m$ we now choose the β smallest (and not the β largest) k -small elements among b_1, \dots, b_n , and for the list $c'_1 \leq \dots \leq c'_m$ we choose the γ smallest k -small elements among c_1, \dots, c_n . Finally, in the second scenario in Sect. 6.2 we do not match element a'_m with the elements c^* and the smallest k -small element b'_1 , but we match a'_m with c^* and with the largest k -small element. The rest of the analysis goes through just as before, and all the (straightforward) details are left to the reader. We conclude that for all real p (positive or negative) problem p -FOA allows a PTAS.

7 Conclusion

We have provided a complete complexity and approximability analysis of the focus of attention problem p -FOA. In a nutshell, the problem is strongly NP-hard and possesses a polynomial time approximation scheme, but does not allow a *fully* polynomial time approximation scheme (unless P=NP).

One can consider different variants of p -FOA. Our problem p -FOA is a special case of the axial three-dimensional assignment problem. Next, a natural generalization to a four-dimensional assignment problem takes four positive real sequences a_1, \dots, a_n , b_1, \dots, b_n , c_1, \dots, c_n , and d_1, \dots, d_n . The cost-coefficient for a quadruple (i, j, k, ℓ) with $1 \leq i, j, k, \ell \leq n$ is defined as

$$c_{ijkl} = \frac{a_i}{(b_j + c_k + d_\ell)^p}.$$

The goal in the generalization is to group the $4n$ numbers into n quadruples (each containing one a_i , one b_j , one c_k , and one d_ℓ) such that the sum of the cost-coefficients of these quadruples is minimized. We note without proof that all our results for p -FOA can be carried over to this four-dimensional generalization (and also to appropriately defined higher-dimensional generalizations as long as the number of dimensions is a fixed constant): the generalized problem is strongly NP-hard, possesses a PTAS, but does not allow an FPTAS.

Another variant is the bottleneck version of p -FOA where the goal is to minimize the maximum cost over all triples in a feasible solution. We restrict ourselves to noting here that the resulting problem is (strongly) NP-hard (by a similar reduction), and that the PTAS described here can be adapted to deal with this variant.

References

1. Al-Hasan, M., Ramachandran, K.K., Mitchell, J.E.: Optimal placement of stereo sensors. *Optim. Lett.* **2**, 99–111 (2008)
2. Burkard, R.E., Dell’Amico, M., Martello, S.: *Assignment Problems*. SIAM, Philadelphia (2009)
3. Burkard, R.E., Rudolf, R., Woeginger, G.J.: Three-dimensional axial assignment problems with decomposable cost coefficients. *Discret. Appl. Math.* **65**, 123–139 (1996)
4. Crama, Y., Spieksma, F.C.R.: Approximation algorithms for three-dimensional assignment problems with triangle inequalities. *Eur. J. Oper. Res.* **60**, 273–279 (1992)
5. Culler, D., Estrin, D., Srivastava, M.: Guest editors’ introduction: overview of sensor networks. *IEEE Comput.* **37**, 41–49 (2004)
6. Gfeller, B., Mihalak, M., Suri, S., Vicari, E., Widmayer, P.: Angle optimization in target tracking. In: *Proceedings of the 11th Scandinavian Workshop on Algorithm Theory (SWAT’2008)*, Springer LNCS, vol. 5124, pp. 65–76 (2008)
7. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, San Francisco (1979)
8. Hardy, G.H., Littlewood, J.E., Pólya, G.: *Inequalities*. Cambridge University Press, Cambridge (1934)
9. Isler, V., Khanna, S., Spletzer, J.R., Taylor, C.J.: Target tracking with distributed sensors: the focus of attention problem. *Comput. Vis. Image Underst.* **100**, 225–247 (2005)
10. Queyranne, M., Spieksma, F.C.R.: Approximation algorithms for multi-index transportation problems with decomposable costs. *Discret. Appl. Math.* **76**, 239–253 (1997)
11. Spieksma, F.C.R., Woeginger, G.J.: Geometric three-dimensional assignment problems. *Eur. J. Oper. Res.* **91**, 611–618 (1996)
12. Tubaishat, M., Madria, S.: Sensor networks: an overview. *IEEE Potentials* **22**, 20–23 (2003)
13. Woeginger, G.J.: When does a dynamic programming formulation guarantee the existence of a fully polynomial time approximation scheme (FPTAS)? *INFORMS J. Comput.* **12**, 57–74 (2000)
14. Yu, W., Hoogeveen, H., Lenstra, J.K.: Minimizing makespan in a two-machine flow shop with delays and unit-time operations is NP-hard. *J. Sched.* **7**, 333–348 (2004)