Heuristics for Balancing Turbine Fans

Samir V. Amiouny * John J. Bartholdi, III
John H. Vande Vate

April 20, 1997

Abstract

We develop heuristics for a problem that models the static balancing of turbine fans: Load point masses at regularly spaced positions on the periphery of a circle so that the residual unbalance about the center — which corresponds to the axis of rotation of the fan — is as small as possible. We prove that our heuristics provide the same worst-case guarantee in terms of residual unbalance as does total enumeration. Furthermore, computational tests show that our heuristics are orders of magnitude faster and not far from optimum on average.

The balancing of rotating elements in modern machinery is critical, and is done in some cases by sophisticated balancing machines. The presence of unbalance in a rotating machine results in vibrations and excess stresses on the bearings, both of which shorten its useful life. Static unbalance, the primary source of unbalance in narrow, disc-shaped rotors such as turbine fans, occurs when the center of gravity of the fan does not coincide with the axis of rotation. Static balancing is the process of reducing unbalance primarily by adding or removing

*The first author is at the Department of Mechanical Engineering of Concordia University, Montreal, CANADA. The others are at the School of Industrial and Systems Engineering of the Georgia Institute of Technology, Atlanta, GA USA.

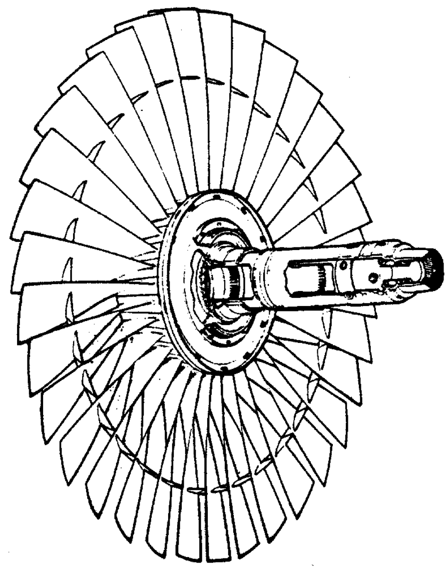

Figure 1: The blades of the low pressure compressor fan of the Rolls Royce RB211-535C jet engine are manufactured separately and then inserted into grooves around the periphery of a thin cylinder.

mass from the fan (Schenck Trebel, 1990). Spatial and structural constraints, however, limit the unbalance that can be corrected (Reiger, 1986).

In some cases, such as in the construction of hydraulic, steam or gas turbines, fan blades are manufactured separately and then welded or inserted into grooves at regularly spaced positions around the periphery of a cylinder as in Figure 1. Due to manufacturing imperfections, the blades are not identical: Their weights as well as the locations of their centers of gravity may vary. Static balancing in such cases proceeds in two stages: First find a "good" sequence of the blades around the cylinder, then attach counterweights to counteract the residual unbalance. For gas and steam turbines, this is necessary not only when the engine is first assembled, but also whenever it is overhauled, because the characteristics of the blades can change over its life.

We consider the combinatorial problem of sequencing n blades around a cylinder of radius r so that the resultant unbalance about the axis of rotation is as small as possible. In our opinion, exact optimization is not appropriate at this time, for the following reasons:

- It is too time-consuming for use in a mass production environment, such as the manufacture of jet engines, especially since jet engines have fans with 100–150 blades. Such problems apparently remain too large to solve quickly with current shop floor technology. (This cost may be inherent because, as we show in Section 6, the problem is NP-hard even in its simplest idealization.)

- The data is inaccurate: Human error, as well as vibrations from breezes and equipment on the shop floor, cause errors in the measurements of the blades. These errors are estimated to be about $\pm 0.2\%$ for gas turbine blades [16]; and our heuristics seem to get almost this close to optimum.

For these reasons we develop fast heuristics that perform well on average and in the worst case guarantee that the resulting unbalance is not "too large". This is consistent with the fact that any remaining unbalance can be corrected by adding counterweights as long as it is not too great.

In fact, we offer several heuristics, each requiring different levels of complexity to implement and computational resources to apply. This reflects the needs and practice of jet engine balancing: Some engine stages contain a large number of light blades for which simple, fast procedures are used. Other stages contain a smaller number of heavier blades for which more sophistication is needed. Where simple procedures have worked in practice, we offer better ones that are just as simple, and where more sophisticated procedures have been used or suggested, we offer faster ones that perform just as well. In addition, all of our heuristics come with worst case bounds providing guarantees of the quality of the solution in terms of the given weights.

We spend considerably more time and effort determining worst case bounds than we do evaluating average case performance. Although average performance more closely reflects experience with the heuristics, it is difficult to draw general conclusions about average performance without making restrictive assumptions about the distribution of the weights. Further, average performance describes what we should expect to see over the course of a number of turbines. It cannot, however, provide any assurances about the performance for a given turbine. A worst case bound, on the other hand, may be substantially worse than our experience with the heuristic, but it does provide a guarantee of how well the heuristic will perform on a given turbine.

Among the more attractive heuristics we suggest is ORDINAL PAIRING, which—when the number of blades is even but not a multiple of four—places the blades so that the unbalance of the entire assembly does not exceed $r\delta_{\max}$, where $\delta_{\max}$ is the maximum absolute difference in the weights of successive blades when the blades are sorted according to weight and r is the radial distance between the center of the disk and the center of gravity of a positioned blade.

In the worst case, any procedure – including complete enumeration – can lead to an unbalance of this magnitude (consider the instance in which one blade is heavy and all the rest are identically light). Thus ORDINAL PAIRING, which requires only slightly more effort than the simple procedures of current practice, offers the strongest possible worst case performance guarantee, which can be orders of magnitude better than that of current practice.

Table I summarizes the worst-case performance of current practice, our heuristics, and total enumeration in terms of r, $\delta_{\max}$, n, and m which is the number of weights in each group when the blades are placed in groups of equal size. Table I provides only lower bounds for the worst-case performance of current practice to illustrate that it degrades as the number of blades increases. In contrast, the worst-case performance of our heuristics improves or remains unchanged as the number of blades increases.

Even an optimum arrangement (described in Table I as ENUMERATION) can have some residual unbalance. We do not know how large this residual unbalance can be in general, but describe the worst case in terms of a function

f of the number of blades as $f(n)\, r\delta_{\max}$. The example in which only one blade has positive weight shows that $f(n) \geq 1$. Our heuristics provide rather tight upper bounds on $f(n)$. For example, when n is even, but not a multiple of 4, Ordinal Pairing shows that $f(n) = 1$ and when n is a multiple of 4, the same heuristic shows that $f(n) \leq \frac{1}{\cos(\pi/n)}$. The worst case performance of Greedy Grouping shows that when $m \geq 3$, $f(m) \geq f(km)$ for each positive integer k.

In most cases we are able to provide examples of weights for which the heuristics exhibit their worst case performance. This is not to suggest that the worst case bounds accurately reflect average performance. On the contrary, the pathological nature of these examples reinforces the conclusion that average performance is substantially better than the worst case bound. Our purpose is to show that no stronger bound is possible with the same parameters. We emphasize such fine distinctions in the worst case bounds because, under the tight tolerances and precise measurements of modern jet engine manufacturing, they can be significant.

1 A 2-Dimensional Model

From elementary mechanics (for example, Yeh and Abrahams, 1960), the effect of each blade on the balance of the fan is identical to that of a point mass concentrated at the center of gravity of the blade. Thus, we model each blade as a point mass with known weight located at a distance r from the center of the fan.

One method for determining the weight of a blade is "moment weighing", which gives the magnitude of the moment that a blade creates about the center of the fan as the product of the weight of the blade and the distance from its center of gravity to the center of the fan. Because moment-weighing is time-consuming, blades are sometimes simply mass-weighed. Although there are no clear guidelines, it is generally true that blades whose radial dimension is large with respect to the radius of the cylinder, such as those found in the low pressure compressor of a gas turbine, are moment-weighed, while smaller blades, such as

	Worst Case Performance in units of $r\delta_{\max}$		
	n Even		n Odd
Heuristic	(multiple of 4)	(not multiple of 4)	
Current Practice			
Single Beam H/L decr.	$\geq 0.04n^2$	$\geq 0.04n^2$	$\geq 0.04n^2$
Double Beam H/L alt.	$\geq 0.15n$	$\geq 0.15n$	n/a
Double Beam H/L decr.	$\geq 0.15n$	$\geq 0.15n$	n/a
Double Beam decr.	$\geq 0.5n$	$\geq 0.5n$	n/a
Triple Beam H/L decr.	$\geq 0.25n$	$\geq 0.25n$	$\geq 0.25n$
Quad. Beam H/L decr.	$\geq 0.3n$	n/a	n/a
Quad. Beam F/R decr.	$\geq 0.3n$	n/a	n/a
Improved heuristics			
Greedy Pairing	$= \sqrt{2}$	$= \sqrt{2}$	n/a
Greedy m-Grouping	$= f(m)$	$= f(m)$	$= f(m)$
Ordinal Pairing	$= \frac{1}{\cos(\pi/n)}$	$= 1$	n/a
Enumeration	$= f(n)$	1	$= f(n)$

Table I: Comparison of the worst-case residual unbalance of heuristics for balancing n blades about the center of a circle. Values are given in units of $r\delta_{\max}$ where r is the radius of the circle and $\delta_{\max}$ is the largest absolute difference in weight between successive blades when the blades are sorted in order of weight. The function $f(\cdot)$ is defined so that $f(n)\, r\delta_{\max}$ is the worst-case unbalance of the best arrangement of n blades when the largest absolute difference in weight between successive blades is at most $\delta_{\max}$.

those found in the high pressure compressor of a gas turbine, are simply mass-weighed. This makes sense since smaller blades are usually more numerous and the consequences of error are smaller than with longer, heavier blades. Since mass-weighing gives no information about the location of the center of gravity, we assume in this case that the centers of gravity of all the blades are at the same distance r from the center of the fan.

In all cases, we assume that the centers of gravity of all the blades fall in a single plane orthogonal to the axis of rotation of the fan. This corresponds current practice in most cases and allows us to model the cylinder as a circle and the locations of the blades as equally spaced points on its periphery.

To achieve static balance, the center of gravity of the assembled fan should fall on its axis of rotation, which we assume coincides with the geometric axis of the cylinder. We formalize this as follows:

Balance around a circle: *Given n point masses with known weights w_i and a circle of radius r with n equally spaced locations on its periphery, find an assignment of the masses to the locations that minimizes residual unbalance about the center.*

The residual unbalance is the magnitude of the vector sum of the moments created by the individual blades about the center. Each moment can be broken down into two components along the axes of any given coordinate system. In our model, the magnitude of the moment created by each point mass about an axis in the plane of the circle is equal to the product of the magnitude of the weight and the orthogonal distance between the location of the weight and the axis.

For convenience, we assume a coordinate system in which the origin is at the center of the circle and the positive x axis goes through one of the n locations as in Figure 2. The locations are numbered in counterclockwise order starting with the one coincident with the positive x axis. So, the coordinates of location i are

$$\left(r\cos\left(\frac{2(i-1)\pi}{n}\right), r\sin\left(\frac{2(i-1)\pi}{n}\right)\right), i = 1, \ldots, n.$$

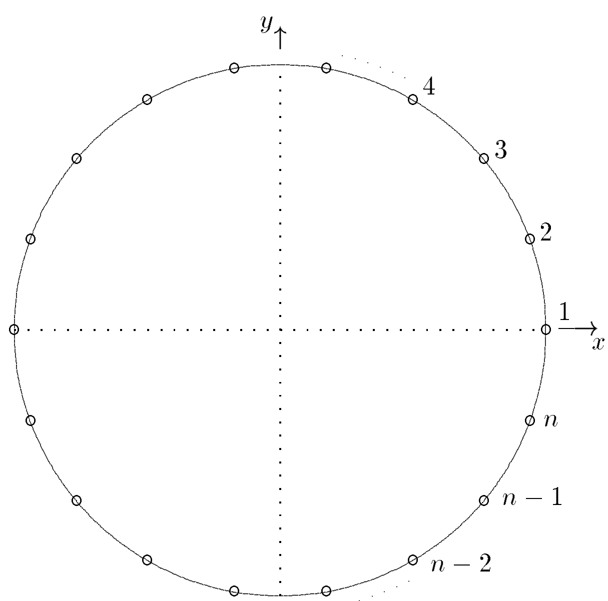

Figure 2: We model the problem of balancing blades on a fan as a problem of sequencing point masses at regular intervals around a circle. We adopt the convention of numbering the locations counter-clockwise.

2 Current Practice

Based on discussions with Pratt & Whitney and Delta Airlines, jet engine manufacturers and airline companies generally balance gas turbines using one of the following procedures (many of which are implemented in BLADIS, fan-balancing software distributed by Carl Schenck, Co.). All the methods we found in practice follow the same outline:

1. Weigh and sort the blades;
2. Group blades into "trips" (groups) of the same size (generally 1–4) so that all blades in the trip are of approximately equal weight.
3. Arrange the blades within each trip to minimize the resultant unbalance within the trip.
4. Arrange the blades of each trip so that:
   - They are equally spaced around the circle; and
   - The resultant unbalance within the trip is minimized.
5. Place the trips on the circle.

The heuristics differ in the size of the trips and in how the trips are placed on the circle. They can be used only when the total number of blades is a multiple of the size of the trip.

A variant of Step 3 is to find the best arrangement of each trip given the unbalance resulting from the placements of the previously placed trips (rather than fixing the arrangement of each trip separately).

In describing the algorithms we adopt the following shorthand notation. We say that a pair of blades is placed in location i to indicate that the heavier blade of that pair is placed in location i and its lighter blade is placed in the diametrically opposed location. Similarly, we say that a pair is placed in location $-i$ to indicate that its lighter blade is placed in location i and its heavier one in the diametrically opposed location. More generally, we say that a trip consisting of m blades is placed in arrangement $(i_1, i_2, \ldots, i_k)$ of equally spaced positions to indicate that the heaviest blade is in position i_1, the next heaviest is in position i_2, etc.

Single Beam, heavy/light adjacent, decreasing Place blades counterclockwise around the disc with the heaviest one first, the lightest one second, the next heaviest third, the next lightest fourth, and so on [14]. (This may be considered to operate with trips of size 1.)

Double Beam, heavy/light adjacent, decreasing Group blades of nearly equal weight into pairs and place the heaviest pair in location 1, the next heaviest pair in location -2, the third heaviest pair in location 3, and so on.

Double Beam, heavy/light adjacent, alternating Pair blades of nearly equal weight; then place the heaviest pair in location 1 and the lightest pair in location -2; the next heaviest pair in location 3 and the next lightest pair in location -4; and so on.

Double Beam Decreasing Group blades of nearly equal weight into pairs and place the heaviest pair is placed in location 1, the next heaviest pair in location 2, the third heaviest pair in location 3, and so on.

Triple Beam, heavy/light adjacent, decreasing Group blades of nearly equal weight into triples and place the heaviest trip in arrangement $(1,\ k+1,\ 2k+1)$, the second heaviest trip in arrangement $(k+2,\ 2k+2,\ 2)$, the third heaviest trip in arrangement $(2k+3,\ 3,\ k+3)$, and so on.

Quadruple Beam, heavy/light adjacent,decreasing Group blades of nearly equal weight into pairs and place the heaviest pair in position 1, the next heaviest pair in position $k+1$, the next pair in position $2k+2$, the next pair in position $3k+2$, the next pair in position 3, the next pair in position $k+3$, and so on.

Quadruple Beam, forward/reverse, decreasing Group blades of nearly equal weight into pairs and place the heaviest pair in position 1, the next heaviest pair in position $k+1$, the next pair in position $3k+2$, the next pair in position $2k+2$, the next pair in position 3, the next pair in position $k+3$, and so on.

These procedures, that we refer to collectively as "current practice", are usually implemented as described for the stages of gas turbines with a relatively large number of light blades: They do not require computerization and are therefore easily implemented on the shop floor. For the stages with a smaller number of heavier blades, they are used to obtain an initial solution that is improved upon by a computerized improvement procedure, usually pairwise interchange.

3 Previous Research

Mosevich (1986) used the same model ("Balance around a circle") for the static balancing problem in hydraulic turbine runners. The terminology is different for hydraulic turbines: "runner" is used instead of "fan", and "bucket" is sometimes used instead of "blade". Mosevich used a Monte Carlo approach to select the best of a large number of randomly generated sequences and reported an eighty percent reduction in the average weight of the necessary correction masses compared to balances generated manually.

Krozenjak and Batagelj (1987) proposed a pairwise interchange heuristic that iteratively improves the sequence by interchanging the positions of two blades until no further improvement is possible.

Laporte and Mercure (1988) modeled the same problem as a quadratic assignment problem. They also used a heuristic based on an interchange algorithm and observed that it performs better on average than the random search method of Mosevich.

Fathi and Ginjupalli (1993) also modeled the problem as a quadratic assignment problem and proposed two families of heuristics for it. The first family of heuristics performs better with a small number of blades, e.g. fewer than 15 blades. It is based on the Placement Heuristic, which places the blades in order of weight – the heaviest blade first – choosing for each blade the available position that brings the resulting center of gravity as close as possible to the center of the disc. They also suggested generalizations of this procedure that required significantly more computational effort.

The second family of heuristics, based on a divide-and-conquer approach called the Rotational Heuristic, is designed for problems with larger numbers of blades. The Rotational Heuristic divides the blades into equal-sized subsets, finds good sequences for the smaller problems of balancing with only the blades in each subset and then interleaves the sequences. This heuristic and its generalizations are competitive with Korenjak and Batagelj's pairwise interchange procedure for problems with up to 24 blades. It also employs the same general approach as our heuristics of finding good arrangements for groups of blades and then placing the arrangements.

Mason and Rönnqvist (1997) tested several implementations of pairwise and three-way interchange algorithms. In each implementation, they used different randomly generated arrangements as starting points for the interchanges and selected the best among the resulting arrangements. They found that for a given total computational effort average performance improves when a relatively small effort is spent on improving a large number of initial solutions as opposed to when more effort is spent on improving a smaller number of initial solutions. In particular, they recommend using a "next descent" pairwise interchange procedure.

Although the interchange procedures suggested by the different authors far outperform current practice, they require computational effort that increases quickly with the number of blades. Even on a modern workstation, the procedure of Laporte and Mercure runs overnight without completion when used for the 90 blades of the sixth stage turbine disc of the Pratt & Whitney PW 2000 jet engine. In contrast, our heuristics require essentially the same computational effort as current practice (a fraction of a second of CPU time for $n = 90$), provide worst case performance guarantees, and produce balances that are nearly as good as those produced by Laporte and Mercure's procedure. Furthermore, the difference in performance is within the accuracy of current industrial data. So, although the procedures of Mosevich and of Laporte and Mercure may be suited for hydraulic turbines, they seem impractical for balancing jet engines: a jet engine manufacturer typically has scores of engines in production at once, each involving 6–10 rotors with 30–130 blades apiece. The size and number of the balance problems together with the need for real-time solutions requires fast, effective heuristics such as ours.

4 Improved heuristics

4.1 Greedy Grouping

We begin by considering the problem in which n, the number of blades, is a multiple of m. In this case, we divide the blades into k groups of m blades, find the best arrangement for each group and orient the groups so that their unbalances counteract. Since the residual unbalance of a group of blades generally increases with $\delta_{\max}$, the maximum absolute difference between the weight of successive blades when they are sorted in order of weight, it makes sense to form the groups on the basis of weight: forming the first group from the m lightest blades, and so on.

Algorithm Greedy Grouping: loads $n = mk$ point masses around the periphery of a circle so that the residual unbalance is small.

1. Sort the blades in order of weight and form k groups of m successive blades.

2. Find an arrangement of the m blades in each group at equally spaced positions around the periphery of the circle that minimizes the residual unbalance of the group.

3. Choose an available set of equally spaced positions around the cylinder and a group to assign to them. Orient the arrangement of the group in the positions so that its residual unbalance is most nearly diametrically opposed to the residual unbalance of the partial assembly consisting of the previous groups.

Experience indicates that choosing the groups in decreasing order of their unbalances leads to better average performance, but does not affect worst case performance. The heuristic can, however, easily accommodate other concerns that arise in assembling turbines. For example, the *clapper tip angle* measures the twist of a blade with respect to its base; and the blades in the low pressure compressor fan of the Rolls Royce RB211-535C jet engine (shown in Figure 1) must be placed so that the angles of adjacent blades differ by less than one half of a degree.

To observe this restriction while balancing the fan, we can both form the groups and choose the next group to place according to the clapper tip angles. The ability to accommodate other concerns like this is an advantage of greedy heuristics such as ours; but it is generally purchased at some cost in balance. For example, forming the groups without regard for the weights can increase the residual unbalance of the worst group.

Theorem 1 shows that when there are three or more blades in a group, Greedy Grouping produces a final assembly with unbalance no greater than that of the worst group.

Theorem 1 *When* $m \geq 3$ Greedy Grouping *produces an assembly with residual unbalance no greater than the largest residual unbalance of a group.*

Proof: Since the m blades in each group are equally spaced around the periphery of the circle, we can orient the next group so that its moment is at an angle $0 \leq \theta \leq \pi/m$ from being diametrically opposed to the moment created by the partial assembly.

Let M_p be the residual unbalance from the partial assembly and let M_g be the residual unbalance for the new group. Then the residual unbalance after placing the new group is:

$$M = \sqrt{(M_p - M_g \cos\theta)^2 + (M_g \sin\theta)^2} = \sqrt{M_p^2 + M_g^2 - 2M_pM_g\cos\theta}.$$

Clearly, M is largest when $\theta = \pi/m$. But, since $m \geq 3$, $\cos\theta \geq 1/2$ and so $M \leq \sqrt{M_p^2 + M_g^2 - M_pM_g} \leq \max\{M_p, M_g\}$. It follows by induction that the residual unbalance for the final assembly is at most that of the worst group. □

We refer to Greedy Grouping with $m = 2$ as Greedy Pairing. The unbalance from any pair of successive weights is at most $r\delta_{\max}$. If we could extend Theorem 1 to groups consisting of only two blades, we would have a heuristic with the strongest possible performance guarantee for even numbers of blades. Unfortunately, the simple example with the four weights $(0, w, w, 2w)$ shows that this is not possible: Even an optimum arrangement of these blades has residual unbalance of $\sqrt{2}rw$ although $\delta_{\max}$, the maximum difference in successive weights, is only w.

Greedy Pairing can produce a sequence with residual unbalance as large as $\sqrt{2}r\delta_{max}$. But, it can do no worse than this: By virtue of the way we orient each pair, the magnitude of the moment about each axis cannot exceed the largest possible moment about the center resulting from the placement of a single pair, which is $r\delta_{\max}$.

When Greedy Grouping is used with $m = 1$, it is identical to the Placement Heuristic of Fathi and Ginjupalli (1993). This may be the only option when the smallest divisor of n is too large. This however seems to be an unlikely situation: Among the multitudes of turbine engines we have seen, none has such a stage; and almost none of the heuristics used in practice can accommodate it. Turbine designers tend to avoid stages with an odd number of blades because they have less symmetry [9].

4.2 Ordinal Pairing

We introduce a heuristic, called Ordinal Pairing, for the case in which there is an even number of blades. This heuristic is similar to current practice in that we form pairs of similar weight blades and then place the blades of each pair in diametrically opposed positions. The main difference between current practice and ORDINAL PAIRING is that we place the pairs in order of the *difference* rather than the *sum* of the weights of the two blades.

Current practice procedures are informationally parsimonious in that the sequence in which the blades are placed depends only on the order of the weights, not on their magnitudes. As a consequence there is no need to evaluate any characteristic of the partial assemblies. This offers a practical advantage over procedures like GREEDY GROUPING: they do not require a measuring or computing device to determine the unbalance of the partial assemblies.

ORDINAL PAIRING, illustrated in Figure 3 uses only ordinal information as well, but requires a little extra computation to sort the pairs of blades in order of the *difference* in the weights of the two blades that form a pair. Then we systematically alternate the placement of the pairs with respect to the center of the circle. This procedure has the same practical advantage as current practice in that it does not require evaluation of the partial assemblies. The small additional effort to sort the pairs, however, leads not only to a better worst case performance guarantee, but also to better average performance.

Algorithm Ordinal Pairing: loads $n = 2k$ point masses around the periphery of a circle so that the residual unbalance about the center of the circle is small.

1. Sort the weights in order of magnitude and form the pairs $p_i, i = 1, 2, \ldots, k$ of consecutive weights; Let δ_i be the difference in weights in pair p_i;
2. Renumber the pairs so that $\delta_{\max} = \delta_1 \geq \delta_2 \geq \cdots \geq \delta_k$;
3. Place the pairs in order in locations: 1, $-n$, -2, $(n-1)$, 3, $-(n-2)$, -4, $(n-3)$, 5, $\ldots$.

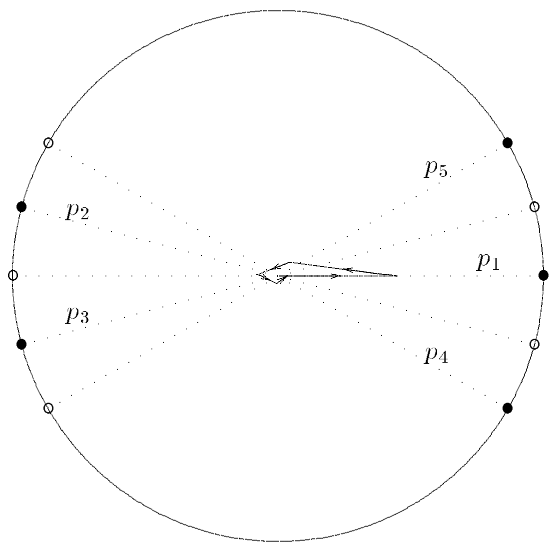

Figure 3: An illustration of Algorithm ORDINAL PAIRING showing the placement of the first five pairs. In each pair, $\circ$ represents the location of the lighter mass and $\bullet$ represents the location of the heavier mass. The solid polyline tracks the location of the center of gravity (scaled up by a factor of 40 to be visible) after each pair is placed. Note that subsequent pairs generally counteract the remaining unbalance, which tends to diminish.

Theorem 2 *When there is an even number of blades,* Ordinal Pairing *places them so that the residual unbalance about the center is no greater than*

$$\begin{cases} r\delta_{\max} & \textit{if } n \textit{ is not a multiple of four} \\ \frac{1}{\cos(\pi/n)} r\delta_{\max} & \textit{if } n \textit{ is a multiple of four.} \end{cases}$$

Proof: We outline the proof in the Appendix. □

Compare this with the two Double Beam, heavy/light heuristics. When the pairs are sorted in order of combined weight, it is possible for each pair with heavier mass on the right side of the y axis to have $\delta_i = \delta_{\max}$ and each pair with heavier mass on the left side of the y axis to have $\delta_i = 0$. In this case, the unbalance can be as large as

$$\left(1 + 2\cos\left(\frac{4\pi}{n}\right) + 2\cos\left(\frac{8\pi}{n}\right) + \dots\right) r\delta_{\max} \quad (1)$$

Plotting (1) shows that as a function of n it grows approximately as $0.15nr\delta_{\max}$. As an illustration, in our previous example with $n = 90$, Ordinal Pairing guarantees an unbalance of at most $r\delta_{\max}$. Under the circumstances described above, the residual unbalance from Double Beam, heavy/light adjacent can be more than 14 times as large.

The performance bounds given by Theorem 2 are tight as we can see from the following examples. For n not a multiple of four, consider an instance of the problem in which $w_1 = 1$ and $w_i = 1 - \delta$ for all other blades. Then our pairing gives $\delta_1 = \delta$ and $\delta_i = 0$ for all other pairs. Using Ordinal Pairing, we get a residual unbalance of δr. In this case however, the sequence is optimal because all except one of the weights are equal and the unbalance is the same for all sequences.

For n a multiple of four, consider an instance of the problem in which $w_1 = 1$ and $w_i = w_{i-1} + \delta, i = 2, \dots, n$. In this case, our pairing gives $\delta_i = \delta$ for all pairs and Ordinal Pairing produces an arrangement with residual unbalance $\frac{1}{\cos(\pi/n)} r\delta$, but this sequence need not be optimal. This may have implications for engine design, as we discuss in the Conclusions.

When n is odd, consider the n pairs formed by pairing each blade with a fictional one of weight equal to that of the lightest blade. Now assume that we sort these pairs and place them according to the sequence of ORDINAL PAIRING: If w_1 and w_n respectively denote the weights of the heaviest and lightest blades, we are guaranteed that the unbalance of the $2n$ blades is no more than $r(w_1 - w_n)$ since $(w_1 - w_n)$ is the maximum difference in weights of any pair. Noting that the original blades fall in n equally spaced positions around the circle and that the fictitious blades produce no unbalance of their own since they all have equal weights, we conclude that we have obtained a valid placement of the original blades that has a residual unbalance of $r(w_1 - w_n)$. Although this unbalance can be as large as $rn\delta_{\max}$, it is an order of magnitude better than the only heuristic of current practice that is valid when the smallest divisor of n is too large.

5 Computational results

Mosevich (1986) reported that the weights of the blades of hydraulic turbines vary by as much as 5% around the average weight. The weights (or moment-weights) of blades from jet engines of our experience closely fit a normal distribution and vary by about 3% around the average value. To evaluate the average performances of our heuristics, we randomly generated sets of weights by sampling from a normal distribution with a mean of 100 and we varied the standard deviation between 1 and 5. The standard deviation had little effect on the performances. Consequently, we report the results for a standard deviation of 5/3 in order to be consistent with the tests of Laporte and Mercure (1988).

Figure 4 summarizes the results. The vertical axis in Figure 4 represents the average distance between the resulting center of gravity and the center of a circle of radius $r = 100$. We chose to present the results in terms of this distance rather than in terms of residual unbalance in order to avoid having to make assumptions about the weights of the blades. Furthermore, the two measures are equivalent: The residual unbalance is equal to the total weight of the blades multiplied by the distance between their center of gravity and the center of the circle. For each value of n, we tested the heuristics for 1000 randomly generated sets of weights. The heuristics from current practice that are omitted from the figure all had practically indistinguishable performances slightly worse than that of Double Beam, heavy/light adjacent, alternating.

Although the more elaborate procedure of Laporte and Mercure performs better on average, it seems unsuited for the problem of balancing jet engines. First, it requires considerably more computational effort than other heuristics. On a Sun workstation, it required a few minutes of CPU time for one run with $n = 20$ and over an hour for $n = 50$. All the other heuristics required less than half a second of CPU time, even for n as large as 100. Second, the accuracy of the blade weights and the approximations inherent to the model cannot justify the additional effort required by the Laporte and Mercure procedure in this context.

In order to gauge the effect of the errors in the measurements of the blade weights (estimated to be about $\pm 0.2\%$ for gas turbine blades [16]), we ran a Monte-Carlo simulation as follows: we assumed that the blades were perfectly balanced around the center of the circle and we introduced errors in the weights that follow a normal distribution with a mean of 0 and a standard deviation that corresponds to having 99.86% of the blades within $\pm 0.2\%$ of the average blade weight. The resulting deviation from the center of the circle is shown as a dashed line on Figure 4. Any performance below this line is likely to be overshadowed by the errors in blade weight measurements.

Figure 5 displays the effects of group size and pre-sorting of groups on the performance of Greedy Grouping. Clearly, placing groups in non-increasing order of their unbalances results in significant improvement in performance. As we noted earlier however, this may not be possible when there are other concerns to accommodate. The effect of group size depends on whether or not the groups are pre-sorted: If they are, using with the smallest manageable group size works best; Otherwise, a larger group size is generally better.

There is obviously a trade-off in average performance between the number of blades in each group and the number of groups. Larger groups generally

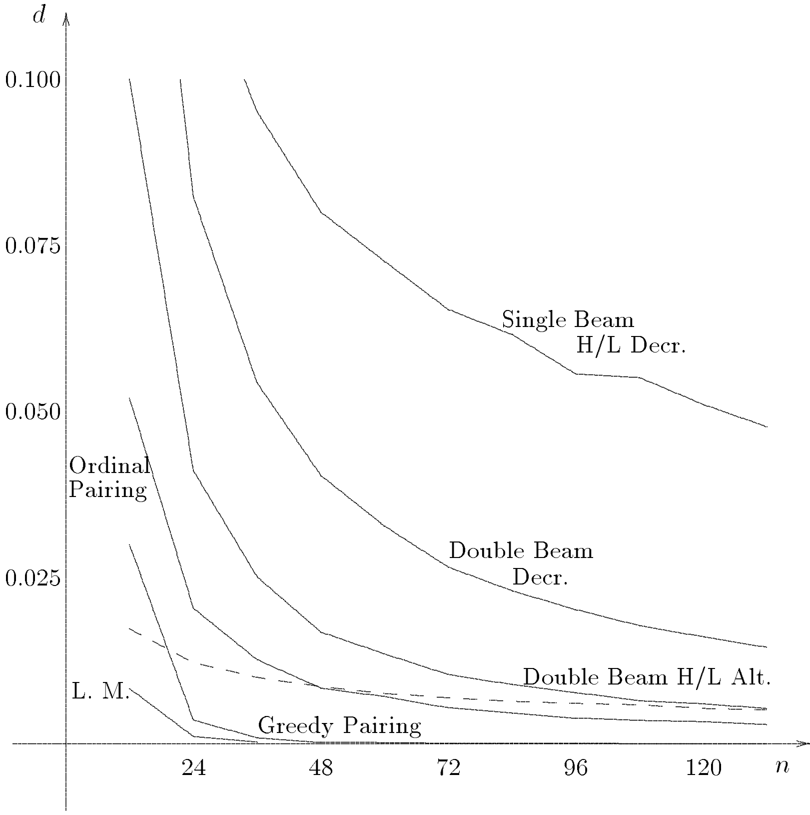

Figure 4: Average performance of different heuristics for balancing n blades about the center of a disc of normalized radius $r = 100$. On the vertical axis, d is the resulting distance between the center of mass of the blades and the center of the disc. The dashed line represents the average deviation due to errors in weight measurements for a seemingly perfect balance. L. M. stands for the procedure of Laporte and Mercure (1988).

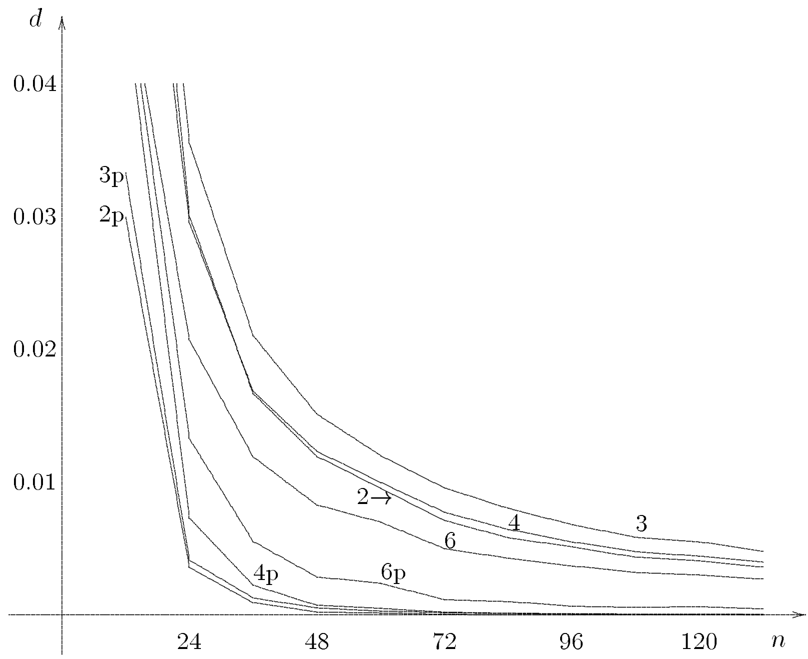

Figure 5: Comparison of the average performance of different implementations of Greedy Grouping. The number on each line indicates the group size used, and a letter p next to the number indicates that the groups were pre-sorted in non-increasing order of their unbalances.

have somewhat better balance within each group (at the cost of dramatically greater computational effort), but the more groups there are, the better they can counteract each others' unbalances, and more so if the groups are pre-sorted in non-increasing order of their unbalances.

A surprising property of all of these heuristics is the very large variance in the residual unbalance about the circle, even when the variance in the distribution of the weights of the blades is small. This implies that their worst case performance is a better testimony to the usefulness of the heuristics than their average performance: Although one heuristic may have significantly better average performance, because of the large variance it may also produce assemblies with unacceptable unbalance more frequently. We believe that the large variance is also a property of the optimal solution.

6 Complexity of balancing around a circle

The problem BALANCE AROUND A CIRCLE is NP-hard by reduction from the partition problem, which is known to be NP-complete (Garey and Johnson, 1979). An instance of the partition problem consists of a set of indices $J = \{1, 2, \ldots, k\}$ and a set of positive integers $\{l_j\}_{j \in J}$; the question is whether there exists a partition J_1, J_2 such that $\sum_{j \in J_1} l_j = \sum_{j \in J_2} l_j$. For convenience, we assume that the integers have been sorted in non-decreasing order of magnitude, so that $l_1 \leq l_2 \leq \cdots \leq l_k$. Given such an instance, create an instance of BALANCE AROUND A CIRCLE as follows. There are $n = 2k + 1$ point masses with weights

$$\begin{aligned} w_{2i} &= iw, \\ w_{2i+1} &= iw - \frac{l_i}{\sin(2i\pi/n)}, \text{ for } i = 1, \ldots, k \\ w_1 &= -\sum_{i=1}^{k} (w_{2i} + w_{2i+1}) \cos\left(\frac{2i\pi}{n}\right) \end{aligned}$$

The radius of the circle is $r = 1$, and the target point is the center of the circle. We will show that for $w > l_k / sin(\pi/n)$, the weights can be sequenced to balance about the center if and only if there is a partition J_1, J_2 such that $\sum_{j \in J_1} l_j = \sum_{j \in J_2} l_j$.

Consider a coordinate system with the x axis going through the location of w_1. Any sequence of the point masses that is balanced about the center should be such that the resulting moments about the x and y axes are both zero. If we project the locations onto the x axis, we see that the weight w_1 is so large that the only way to counteract the moment it creates about the y axis is to have the remaining masses appear from left to right in non-increasing order of weight. By our choice of weights, this configuration will exactly balance about the y axis. Now, to balance about the x axis, we have to orient k pairs of masses since the locations of weights w_{2i} and $w_{2i+1}, i = 1, \cdots, k$ have the same projection onto the x axis. The moment created by each pair about the x axis can be either positive or negative, depending on whether the heavier or the lighter mass goes on the positive side of the axis. In either case, however, the magnitude of the moment is $l_i = |w_{2i} - w_{2i-1}| \sin(2i\pi/n)$. Therefore, the weights can be sequenced to balance about the x axis if and only if there is a partition J_1, J_2 such that $\sum_{j \in J_1} l_j = \sum_{j \in J_2} l_j$.

The above argument establishes weak NP-completeness for the case of an odd number of blades. We can easily extend the argument for n even by adding a mass of zero weight.

The proof above appeared in Amiouny (1993). A slightly different proof, based on a reduction from the even-odd partition problem, appeared in Burkard, Çela, Rote, and Woeginger (1995).

7 Conclusions

Our heuristics are similar to those used in practice. The modifications we recommend, however, lead to significantly better worst case performance, as proven above, as well as significantly better average performance, as shown in computational tests.

Our worst case bounds are all given in terms of the magnitude of the difference between successive weights when the weights are sorted in order of magnitude. This kind of bound allows us to set the manufacturing tolerances for the blades at exactly the level required to guarantee a desired quality of balance in the final assembly. For example, the sixth stage turbine disc of the Pratt & Whitney PW 2000 jet engine must be statically balanced at a minimum of 900 rpm to within 1.0 ounce-inch without adding counterweights [5]. Our heuristic, Ordinal Pairing, guarantees a residual unbalance of no more than 1.0 ounce-inch if the difference between successive weights is no more than 0.08 ounces. To provide the same guarantee when the blades are sequenced according to current practice, the difference between successive weights cannot exceed *0.008 ounces* — a requirement that is an order of magnitude more stringent.

Engine designers determine the number of blades at each stage (fan) of a jet engine based on considerations of pressure changes and other issues of physics and mechanical engineering. In particular, they do not explicitly consider ease of balance; yet balancing the fans is the bottleneck to production and to overhaul of jet engines.

Our analysis suggests how engine designers might accommodate issues of balance. For example, it is generally easier to balance a fan with many blades (and our worst-case bounds reflect this). Also, there is a compelling reason to have a number of blades that is divisible by 2, but not 4: Our performance guarantees for procedures that require only ordinal information about the weights of the blades are unavoidably and strictly weaker when the number of blades is a multiple of four. On reflection this makes sense: Each pair of blades on the fan has another pair that is orthogonal to it and so we cannot use the unbalance of the second pair to correct for the unbalance of the first. Therefore it is as if we must balance two independent fans, each with half the number of blades; and fewer blades means a weaker performance guarantee.

This paper is the third in a series on combinatorial mechanics in which we consider physical issues in loading problems. We have already studied balance (Amiouny, Bartholdi, Vande Vate, and Zhang 1992) as well as deflection and bending in one-dimensional structures (Amiouny, Bartholdi and Vande Vate, 1993).

Appendix

Theorem 3 *When there is an even number of blades,* ORDINAL PAIRING *places them so that the residual unbalance about the center is*

$$\leq \begin{cases} r\delta_{\max} & \text{if } n \text{ is not a multiple of four} \\ \frac{1}{\cos(\pi/n)} r\delta_{\max} & \text{if } n \text{ is a multiple of four.} \end{cases}$$

Proof: The final moments about the coordinate axes are explicit functions of the differences in the weights in each pair. Since the angle between adjacent locations is π/k, the final moment about the x axis is

$$M_x \quad = \quad r \sum_{i=1}^{k} (-1)^{\lfloor (i-1)/2 \rfloor} \, \delta_i \sin\left(\left\lfloor \frac{i}{2} \right\rfloor \frac{\pi}{k} \right); \tag{2}$$

the final moment about the y axis is

$$M_y \quad = \quad r \sum_{i=1}^{k} (-1)^{\lfloor i/2 \rfloor} \, \delta_i \cos\left(\left\lfloor \frac{i}{2} \right\rfloor \frac{\pi}{k} \right); \tag{3}$$

and the final moment M about the center satisfies

$$M^2 = M_x^2 + M_y^2.$$

We outline the proof for the case when $M_x \geq 0$ and $M_y \leq 0$. The other cases can be treated in a similar fashion and yield the same bound (the worst cases occur when M_x and M_y are of opposite sign).

The moment M_x is a linear function of the differences δ_i. Since the pairs are numbered so that $\delta_{\max} = \delta_1 \geq \delta_2 \geq \cdots \geq \delta_k$, the maximum value of M_x satisfies

$$\begin{aligned} M_x \quad &\leq \quad r \left[(\delta_1 - \delta_6) \sin(\pi/k) + (\delta_6 - \delta_{10}) \sin\left(\frac{3\pi}{k} \right) + \cdots \right] \\ &= \quad r \left[\delta_1 \sin(\pi/k) + \delta_6 \left(\sin\left(\frac{3\pi}{k} \right) - \sin(\pi/k) \right) + \cdots \right]. \end{aligned}$$

In terms of the same δ_i variables, the most negative value of M_y in (3) is bounded by

$$M_y \geq r \left[\delta_1 \left(1 - 2\cos(\pi/k)\right) + 2\delta_6 \left(\cos\left(\frac{2\pi}{k} \right) - \cos\left(\frac{3\pi}{k} \right) \right) + \cdots \right].$$

In the bounds for M_x and M_y, the coefficients of $\delta_6, \delta_{10}, \ldots$ are all positive. So, for a given value of δ_1, M_x^2 and M_y^2 are both bounded by convex functions of $\delta_6, \delta_{10}, \ldots$ so that $M^2 = M_x^2 + M_y^2$ is bounded by a function that achieves its maximum when $\delta_6, \delta_{10}, \ldots$ are either equal to δ_1 or to 0. Since $\delta_1 \geq \delta_2 \geq \cdots \geq \delta_k$, we let δ_{2j} be one of $\delta_6, \delta_{10}, \ldots$ such that $\delta_1 = \delta_2 = \cdots = \delta_{2j}$ and $\delta_{2j+1} = \delta_{2j+2} = \cdots = \delta_k = 0$. Then we get

$$M^2 \leq r^2 \delta_1^2 \left[\left(\sin\left(\frac{j\pi}{k} \right) \right)^2 + \left(1 - 2 \sum_{i=1}^{j} (-1)^{i+1} \cos\left(\frac{i\pi}{k} \right) \right)^2 \right]. \tag{4}$$

We prove in Lemma 1 that unless $j = k/2$,

$$-\cos\left(\frac{j\pi}{k}\right) \leq 1 - 2\sum_{i=1}^{j}(-1)^{i+1}\cos\left(\frac{i\pi}{k}\right) \leq \cos\left(\frac{j\pi}{k}\right). \tag{5}$$

When k is odd, or equivalently, when n is not a multiple of four, j is at most $(k-1)/2$ and (5) holds for all j. In this case (4) becomes $M^2 \leq r^2\delta_1^2$.

When k is even, however, or equivalently when n is a multiple of four, j can be as large as $k/2$ at which point (5) no longer holds. From standard mathematics textbooks, such as Vygodski (1973), we have:

$$\sum_{i=1}^{j}\cos i\alpha = \frac{\sin\left(\frac{(2j+1)\alpha}{2}\right) - \sin\left(\frac{\alpha}{2}\right)}{\sin(\alpha)}. \tag{6}$$

Using (6), we obtain

$$\begin{aligned}
1 - 2\sum_{i=1}^{k/2}(-1)^{i+1}\cos\left(\frac{i\pi}{k}\right) &= 1 - 2\sum_{i=1}^{k/2}\cos i\pi/k + 4\sum_{i=1}^{\lfloor k/4\rfloor}\cos\left(\frac{2i\pi}{k}\right) \\
&= \frac{2\sin\left(\frac{(k+2)\pi}{n}\right)}{\sin(\pi/k)} - \frac{\sin\left(\frac{(k+1)\pi}{n}\right)}{\sin(\pi/n)} \\
&= \frac{\cos(\pi/k) - 1}{\sin(\pi/k)} \\
&= -\tan(\pi/n)
\end{aligned}$$

so that (4) becomes

$$\begin{aligned}
M^2 &\leq r^2\delta_1^2\left[(\tan(\pi/n))^2 + 1\right] \\
&= \left(\frac{r\delta_1}{\cos(\pi/n)}\right)^2
\end{aligned}$$

and so

$$M \leq \frac{1}{\cos(\pi/n)}r\delta_{\max}.$$

□

Lemma 1 *For $j < n/4$,*

$$-\cos\left(\frac{j\pi}{k}\right) \leq 1 - 2\sum_{i=1}^{j}(-1)^{i+1}\cos\left(\frac{i\pi}{k}\right) \leq \cos\left(\frac{j\pi}{k}\right), \tag{7}$$

Proof: Our proof is by induction on j. Clearly (7) holds for $j = 0$ and $j = 1$. Now we assume that it holds for $j - 1$ and j (j odd) and we show that it holds for $j + 1$ and $j + 2$. First observe that

$$\begin{aligned} & \cos\left(\frac{j\pi}{k}\right) - 2\cos\left(\frac{(j+1)\pi}{k}\right) + \cos\left(\frac{(j+2)\pi}{k}\right) \\ = \quad & \cos[\frac{(j+1)\pi}{k} - \frac{\pi}{k}] + \cos[\frac{(j+1)\pi}{k} + \frac{\pi}{k}] - 2\cos\left(\frac{(j+1)\pi}{k}\right) \\ = \quad & 2(\cos(\pi/k) - 1)\cos\left(\frac{(j+1)\pi}{k}\right). \end{aligned}$$

And $2(\cos(\pi/k) - 1)\cos\left(\frac{(j+1)\pi}{k}\right) \leq 0$ as long as $\cos\left(\frac{(j+1)\pi}{k}\right) \geq 0$, that is, as long as $\frac{(j+1)\pi}{k} \leq \pi/2$ or, equivalently, $j + 1 \leq k/2 = n/4$. So,

$$\cos\left(\frac{j\pi}{k}\right) - 2\cos\left(\frac{(j+1)\pi}{k}\right) + \cos\left(\frac{(j+2)\pi}{k}\right) \leq 0 \tag{8}$$

holds for all $j < n/4$. In particular, adding (7) and (8) evaluated at $j - 1$, we obtain

$$1 - 2\sum_{i=1}^{j+1} (-1)^{i+1} \cos\left(\frac{i\pi}{k}\right) \leq \cos\left(\frac{(j+1)\pi}{k}\right),$$

and adding (7) and (8) evaluated at j, we obtain

$$-\cos\left(\frac{(j+2)\pi}{k}\right) \leq 1 - 2\sum_{i=1}^{j+2} (-1)^{i+1} \cos\left(\frac{i\pi}{k}\right).$$

To establish that (7) holds for $j + 1$, we still need to show that

$$-\cos\left(\frac{(j+1)\pi}{k}\right) \leq 1 - 2\sum_{i=1}^{j+1} (-1)^{i+1} \cos\left(\frac{i\pi}{k}\right). \tag{9}$$

Since j is odd,

$$1 - 2\sum_{i=1}^{j+1} (-1)^{i+1} \cos\left(\frac{i\pi}{k}\right) = 1 - 2\sum_{i=1}^{j} (-1)^{i+1} \cos\left(\frac{i\pi}{k}\right) + 2\cos\left(\frac{(j+1)\pi}{k}\right).$$

So (9) holds as long as

$$\cos\left(\frac{j\pi}{k}\right) - \cos\left(\frac{(j+1)\pi}{k}\right) \leq 2\cos\left(\frac{(j+1)\pi}{k}\right),$$

that is, as long as

$$\cos\left(\frac{j\pi}{k}\right) - 3\cos\left(\frac{(j+1)\pi}{k}\right) \leq 0.$$

The function $\cos\left(\frac{j\pi}{k}\right) - 3\cos\left(\frac{(j+1)\pi}{k}\right)$ increases with $\frac{j\pi}{k}$. When n is not a multiple of 4, $\frac{(j+1)\pi}{k}$ is at most $\frac{\pi}{2} - \frac{\pi}{2k}$ so that

$$\cos\left(\frac{j\pi}{k}\right) - 3\cos\left(\frac{(j+1)\pi}{k}\right) \leq \sin\left(\frac{3\pi}{2k}\right) - 3\sin\left(\frac{\pi}{2k}\right) \leq 0.$$

When n is a multiple of 4, $\frac{(j+1)\pi}{k}$ can be as large as $\pi/2$ for which

$$\cos\left(\frac{j\pi}{k}\right) - 3\cos\left(\frac{(j+1)\pi}{k}\right) = \sin(\pi/k) > 0,$$

so that (7) does not hold for $j = k/2 = n/4$. However, we still have $\cos\left(\frac{j\pi}{k}\right) - 3\cos\left(\frac{(j+1)\pi}{k}\right) \leq 0$ for all $j < n/4$, and consequently, (7) holds for $j < n/4$.

Finally, to establish that (7) holds for $j+2$, we still need to show that

$$1 - 2\sum_{i=1}^{j+2}(-1)^{i+1}\cos\left(\frac{i\pi}{k}\right) \leq \cos\left(\frac{(j+2)\pi}{k}\right).$$

The argument is similar to the one we used to show (9) and yields the same result. □

Acknowlegments

S. Amiouny was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada (OGP0155986). J. Bartholdi was supported in part by a grant from the Air Force Office of Scientific Research (F49620-94-1-0232) and by the office of Naval Research (N00014-89-J-1571).

References

[1] Amiouny, S.V. 1993. *Combinatorial mechanics*, doctoral thesis, Georgia Institute of Technology, Atlanta, GA 30332.

[2] Amiouny, S.V., Bartholdi, J.J. III, Vande Vate, J.H. and J. Zhang. 1992. "Balanced loading", *Operations Research* **40**(2):238–246.

[3] Amiouny, S.V., Bartholdi, J.J. III and J.H. Vande Vate. 1992. "Minimizing deflection and bending moment in a beam with end supports", *Mechanics of Structures and Machines* **21**(2):167–184.

[4] Anonymous. 1990. *Fundamentals of Balancing*, Third Edition. Published by Schenck Trebel, A Unit of Schenck Corporation, 535 Acorn Street, Deer Park, NY 11729.

[5] Blue Prints for the Pratt & Whitney 2000 Jet Engine. Drawing No. 8A1026 page 4. Drawn December 1, 1983.

[6] Burkard, R.E., Çela, E., Rote, G. and J.G. Woeginger. 1995. "The quadratic assignment problem with an Anti-Monge matrix and a Toeplitz matrix: easy and hard cases", *Mathematical Programming.* To appear.

[7] Fathi, Y. and K.K. Ginjupalli. 1993. "A mathematical model and a heuristic for the turbine balancing problem", *European Journal of Operations Research* **63**:336–342.

[8] Garey, M. R. and D. S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*, W. H. Freeman and Co., San Fransisco.

[9] Kauth, K., Advanced Design, Inc. Powder Springs, GA (USA). Personal communication.

[10] Korenjak, S. and V. Batagelj. 1987. "Turbine balancing problem", technical report, Department of Mathematics, University of Edvard Kardelj of Eljubljana, Yugoslavia.

[11] Laporte, G. and H. Mercure. 1988. "Balancing hydraulic turbine runners: A quadratic assignment problem", *European Journal of Operations Research* **35**:378–381.

[12] Mason, A. and M. Rönnqvist. 1997. "Solution methods for the balancing of jet turbines", *Computers & Operations Research* **24**(2):153–167.

[13] Mosevich, J. 1986. "Balancing hydraulic turbine runners: a discrete combinatorial optimization problem", *European Journal of Operations Research* **26**:202–204.

[14] Reader, H., Pratt & Whitney Canada, Longueuil, QC (Canada). Personal Communication.

[15] Rieger, N.F. 1986. *Balancing of Rigid and Flexible Rotors*, The Shock and Vibration Monograph Series **SVM-12**, The Shock and Vibration Information Center, United States Department of Defense.

[16] Root, B., Assembly Engineer at Pratt & Whitney, East Hartford, CT (USA). Personal communication.

[17] Vygodski, M. 1973. *Aide mémoire de mathématiques supérieures.* Published by Editions Mir, 2 Pervi Rijski pereoulok, Moscow, I-110, GSP, Russia.