Balanced Loading

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

October 25, 1989; revised December 4, 1997

Abstract

We develop a heuristic for a problem motivated by the loading of aircraft or trucks: pack blocks into a bin so that their center-of-gravity is as close as possible to a target point. Our heuristic either produces good solutions or else signals that none is possible. It also works when loading non-homogeneous blocks into a bin of non-zero and possibly non-homogeneous mass.

How should cargo be stowed in an airplane? There are two conflicting objectives: to use limited space efficiently by filling the aircraft, and to balance the load so that the aircraft will fly more safely, fly faster, and use less fuel (Chapel, 1948; MACP 55-41, 1987). Although the first issue has been much-studied under the rubric of "bin-packing problems", there has been little work on the issue of balance. We model the problem of balancing the load as that of arranging n "blocks" (intervals) in a 1-dimensional bin so that the center-of-gravity of the packed blocks is as close as possible to a target point p. We assume that the bin is sufficiently small that we can treat acceleration due to the gravity of the earth as everywhere uniform and parallel. We ignore gravity other than that of the earth. In addition we assume that the bin is rigid. In this, the simplest formulation, we are interested only in the center-of-gravity of the packed blocks and so can ignore the mass distribution of the bin itself. Later we shall show how to account for a bin of possibly non-zero and non-homogeneous mass.

In our formalization all the blocks fit into the bin, so the only question is where to stow blocks, not which are to be stowed. By assuming that the blocks fit into the bin, we are free to concentrate on the issue of balance.

A 1-dimensional model is reasonable for the loading of freight aircraft because the challenge is to balance the load fore and aft. In general it is not a significant problem to balance the load from side to side because thwartship imbalances are usually small and ailerons give the aircraft sufficient stability to resist rolling. Chapel (1948, p. 32) says that "since there is little movement of the CG [center-of-gravity] along the vertical and lateral axes of an airplane, airplanes are almost always loaded with regard to the longitudinal (fore and aft) axis and little attention is paid to the possibility of a shift of CG along the other axes". Accordingly we consider balance in only one dimension, along the length of the airplane.

The load need not be perfectly balanced fore and aft; it is enough that the center-of-gravity of the load fall within a specified region. If the load is not perfectly balanced, the pilot can adjust the elevators or trim tabs of the airplane to correct the imbalance in flight. This, however, increases drag, which decreases speed and increases fuel consumption. In extreme cases, imbalance can render the airplane dynamically unstable (Chapel, 1948, p. 32 ff.).

The acceptable region for the center-of-gravity depends on the total weight of the load, and is generally available to the load planner as a chart called the "payload center-of-gravity flight window". For example, the acceptable region for a Boeing 747 is an interval whose diameter is the length of a standard pallet (MACP 55-41, 1987). For the Boeing 747 our heuristic will either produce a satisfactorily balanced load or else signal that it is impossible to achieve.

Currently, load planners rely on rules of thumb, not formalized procedures, for loading aircraft (see, for example, Chapel, 1948; MACP 55-41, 1987). Computerized systems such as those described by Cochard and Yost (1985) and Martin-Vega (1985) are used to help solve the higher level problem of assigning cargo to planes. Our work here could be incorporated into such planning systems.

Another application of our model is to the loading of trucks. To prevent damage to the roadway, trucks are forbidden to carry excessively heavy loads. This is formalized in federal and state laws that, among other restrictions, stipulate maximum weights allowed on any axle. For example, Title 23 United States Code Annotated § 127 (West 1990) sets a limit of 20,000 pounds per axle for any vehicle using the National System of Interstate and Defense Highways; and these restrictions are extended in the Official Code of Georgia Annotated § 32-6-26 (b) (Michie 1985 & Supplement 1990). Because the *distribution* of weight affects the weight measured at an axle, it can happen that axle weight, and not cargo capacity, limits the load carried on a truck. Stowing the cargo carefully can reduce the maximum axle weight and allow the truck to be more fully loaded.

We model a truck as a 1-dimensional beam with two simple supports (axles) arbitrarily placed. To remain within the weight limits, we want to load the truck so that the heaviest axle is as light as possible. A straightforward application of elementary mechanics shows that to minimize the maximum axle weight is equivalent to loading the truck so that the center-of-gravity is as close as possible to the point midway between the two axles. Thus our heuristic for loading airplanes applies equally well to the loading of trucks.

1 One-dimensional balanced loading

In the 1-dimensional version of the balanced loading problem, the bin and the blocks are intervals and p is collinear with the bin (but need not fall within the bin). The jth block is given by its length l_j and weight w_j, and the bin is of length $\sum_{j=1}^{n} l_j$ so that the blocks fill the bin exactly. The average density of the jth block is w_j/l_j. For convenience we assume that the blocks are labelled in non-decreasing order of average density, so that $w_1/l_1 \leq \ldots \leq w_n/l_n$. Initially we require that all the blocks be homogeneous, so that the center-of-gravity of each block is coincident with its geometric center. Later we will show how to relax this restriction.

Most of our analysis is based on the following ideas from elementary mechanics (see, for example, Yeh and Abrams, 1960). The moment exerted by a block about a point is the product of its weight and the distance between the point and the center-of-gravity of the block; similarly, the moment exerted by a set of blocks about a point is the sum of the moments of each block about that point. Then the resultant moment of a set of blocks about a point equals the product of the sum of the weights of the blocks and the distance between the point and the center-of-gravity C of the blocks. More formally, let the distance between points x and y be $d(x, y)$; then for any point x,

$$M(x) = d(x, C)(\sum_i w_i). \tag{1}$$

If the sense of a moment is to rotate the bin clockwise, then we consider its magnitude to be positive; and if counter-clockwise, then negative.

We assume a coordinate system superimposed on the bin in which the left-most end of the bin is the origin.

2 A heuristic of guaranteed accuracy

We develop a heuristic that guarantees nearly-balanced packings. Let $l_{\max}$ be the length of the longest block to be packed; our heuristic packs the bin so that the center-of-gravity of the packed blocks is either within $(1/2)l_{\max}$ of the target point p or else it is as close as possible. One might think of this as a "minimum regret" guarantee, since if there is a good packing, then our heuristic finds one; if there does not exist a good packing, then our heuristic produces the best possible. Thus our heuristic gives useful information even when it fails to produce a good solution: if it produces a packing for which the center-of-gravity is far from the target point, then we know that there is no better packing.

This guarantee holds regardless of the location of p and the weights of the blocks. In fact, the bound holds even when the blocks are not homogeneous.

The worst-case computational effort required by the heuristic is dominated by the $O(n \log n)$ work required to sort the blocks.

Our heuristic packs the least-dense blocks first, and fills the bin from the ends inward. It does this by the repeated alternation of two steps: 1. pack the next block; and 2. transform the problem to a similar one with a new target point and one fewer block to be packed. This strategy is based on the following informal reasoning. Some block must be packed at the far end of the bin and that block induces a moment M about the target point p. If the currently packed blocks are nearly balanced, then we want M small to minimize the disruption. Therefore pack the least-dense block at the far end of the bin. To offset the consequent disruption, compute a new target point for which a perfect balance of the remaining blocks will exactly countervail M.

We formalize the algorithm as follows.

Algorithm BALANCE**:** packs blocks so that their center-of-gravity lies close to target point p;

1 Sort the blocks from least dense to most dense;

2 Initialize variables: $p_0 \leftarrow p$, $M_0 \leftarrow 0$;

3 For $j = 1, 2, \ldots, n$ do:

Place the jth block so that its geometric center is as far as possible from

$$p_j \leftarrow p_{j-1} - M_{j-1} / (\sum_{i=j}^{n} w_i), \tag{2}$$

respecting the placement of previously-loaded blocks. Let M_j be the moment about p_j induced by the jth block;

Figure 1 illustrates the iterative step of Algorithm BALANCE. Figure 2 gives a time-expanded view of the progress of the heuristic for an example problem. Figure 3 shows that in the final packing the blocks are *not* sequenced according to weight. This is is worth remarking since most commercial load planners do pack by weight. They generally try, in an ad hoc manner, to place successively lighter items to each side of the heaviest. In other words, their packings tend to be unimodal with respect to weight. In contrast, our heuristic produces packings that are unimodal with respect to density. For the example in Figure 3, our

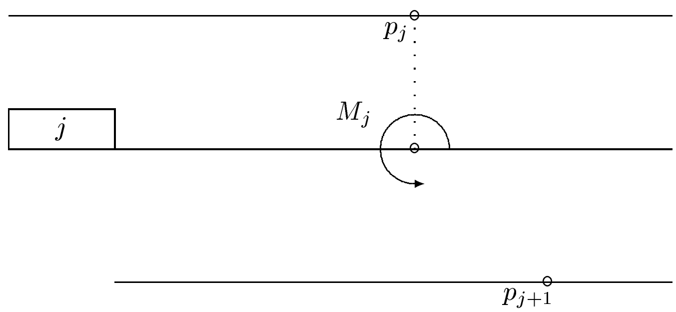

Figure 1: The jth iteration of the heuristic.

solution is very good, and, moreover, is unlikely to be generated by any technique that packs by weight.

Our heuristic repeatedly reduces the packing problem to a similar problem with one less block. Our first result shows that solving the reduced problem is equivalent to solving the problem from which it was derived.

Lemma 1 *In the completed packing the moment induced by blocks* $j, \ldots, n$ *about* p_j *equals the moment induced by blocks* $j-1, \ldots, n$ *about* p_{j-1}.

Proof Let blocks $j, \ldots, n$ induce moment M about p_j. Then blocks $j, \ldots, n$ must induce about p_{j-1} moment $M - (p_j - p_{j-1})(\sum_{i=j}^{n} w_i)$, which by expression 2 equals $M - M_{j-1}$. Now the moment induced about p_{j-1} by blocks $j-1, \ldots, n$ equals the sum of the moments induced by block $j-1$ and the moment induced by blocks $j, \ldots, n$ and so equals $M_{j-1} + M - M_{j-1} = M$. **QED**

Some instances of the balanced loading problem are hard to balance because the target point is too close to one end of the bin. Paradoxically, for such instances it is easy to do the best possible. For convenience we analyze this situation separately. Let C be the center-of-gravity of the blocks when they are packed in non-decreasing order of density. Then this packing is optimal for any target point greater than C. (That it is optimal in this special case to sequence the blocks by density is equivalent to the result of Smith (1956) that weighted mean flow time on a single machine is minimized by sequencing

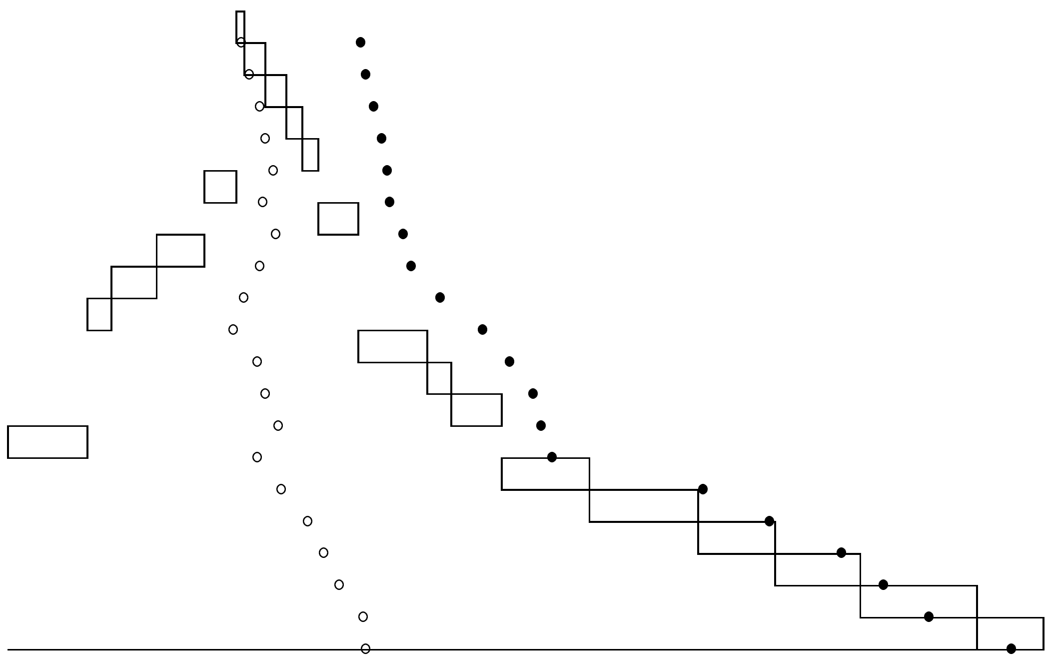

Figure 2: A time-expanded view of the heuristic in which successively-packed blocks are displaced upward so that the sequence of packing is apparent. (The higher blocks are denser and therefore were packed later.) The points labelled "$\circ$" are the successive target points p_j; the points labelled "$\bullet$" are the successive centers-of-gravity of the partial packings. The completed packing is nearly optimal since the last block is nearly centered on the final target point. The quality of the packing can also be seen by comparing the final center-of-gravity (the highest "$\bullet$") with the original target point (the lowest "$\circ$").

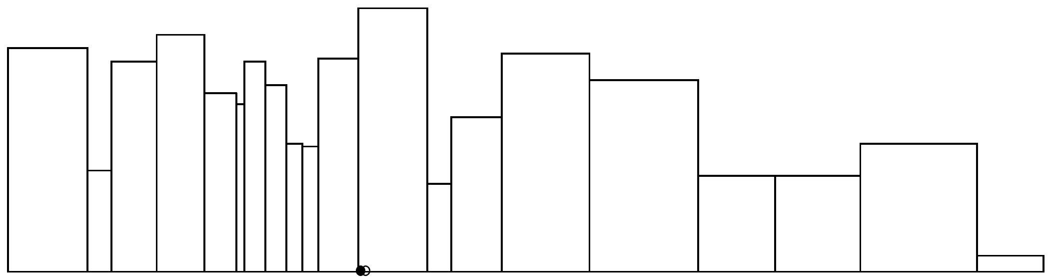

Figure 3: The final packing. The target point is indicated by "∘" and the resultant center-of-gravity by "•". The heights of the rectangles are proportional to their weights.

the jobs according to the "weighted shortest processing time" rule.) Similarly, for any target point less than $(\sum_{j=1}^{n} l_j) - C$ it is optimal to pack the blocks in non-increasing order of density. The following result explains how our algorithm recognizes this situation.

Lemma 2 *If all $M_j \geq 0$ or all $M_j \leq 0$ $(j = 1, \ldots, n)$, then the center-of-gravity of the packed blocks is as close as possible to the target point p.*

Proof If all the moments are of the same sense, then the blocks appear in non-decreasing order of density from left-to-right (or right-to-left). Within any sequence of blocks that fails to satisfy this condition there must exist some adjacent pair of blocks that are in strictly decreasing order of density. It is straightforward to verify that interchanging two such blocks in the packing will move the center-of-gravity of the packed blocks closer to p. **QED**

Our main result says that unless the heuristic packs the blocks from one side to the other as in Lemma 2, the resultant center-of-gravity is guaranteed to be close to the target point.

Lemma 3 *If the sense of the moments M_j changes, then the center-of-gravity of the packed blocks is within $(1/2)l_{\max}$ of p.*

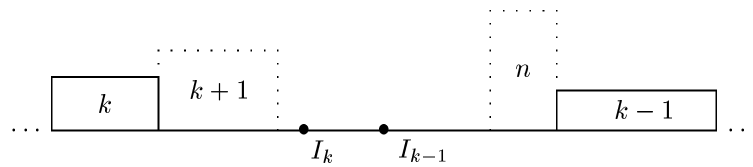

Figure 4: The height of each block is proportional to its density.

Proof Since the heuristic always packs as far as possible from p_j, until the last block has been packed there is exactly one unpacked interval in the bin. Let I_j be the center of the unpacked interval just before the jth block has been stowed. Let C_j be the center of gravity of the blocks $j, \ldots, n$.

Let the kth block be the last one whose placement changed the sense of the moment. Without loss of generality, assume that block $k-1$ was placed at the right end of the unpacked interval, as in Figure 4. This means that $M_{k-1} > 0$ and so $p_k = p_{k-1} - M_{k-1}/(\sum_{i=k}^{n} w_i) < p_{k-1}$. Furthermore, since the algorithm places each block j at the most distant available position from p_j, it must be that $p_{k-1} \leq I_{k-1}$.

Since block k was the last one which reversed the sense of the moment, it must be that blocks $k, \ldots, n$ have been packed together from left to right. Now since blocks $k, \ldots, n$ are in non-decreasing order of density, their center of gravity must lie on the right, so that $I_k \leq C_k$. Furthermore, since $M_k < 0$, and since block k was the last block for which the sense of the moment changed, it must be that $C_k < p_k$. Therefore we have established that

$$I_k \leq C_k < p_k < p_{k-1} \leq I_{k-1}. \tag{3}$$

Now let M be the moment created by blocks $k, \ldots, n$ about p_k. By Lemma 1 M is also the magnitude of the moment induced about p by the packed blocks.

By expression 3 we have that

$$\begin{aligned} |M| &= (\sum_{i=k}^{n} w_i)(|C_k - p_k|) \\ &< (\sum_{i=k}^{n} w_i)(|I_k - I_{k-1}|) \end{aligned}$$

$$= (\sum_{i=k}^{n} w_i)(l_{k-1}/2).$$

QED

Corollary 1 *If p is the midpoint of the bin, then algorithm* BALANCE *produces a packing for which the center-of-gravity of the packed blocks is within $(1/2)l_{\max}$ of p.*

Proof If p is the midpoint of the bin, then the second block will be placed on the side opposite the first-placed block and so M_1 and M_2 will be of different sign. **QED**

When p is the midpoint of the bin, not even an optimal packing can guarantee a better solution. To see this, consider packing two blocks, one of length 1 and great weight W and the other of great length L and weight 1. The only possible packing has its center-of-gravity at distance nearly $L/2$ from the midpoint of the bin.

Combining the preceding results give bounds on the performance of our heuristic.

Theorem 1 *Algorithm* BALANCE *produces a packing for which the center-of-gravity of the packed blocks is either within $(1/2)l_{\max}$ of p or else is as close as possible to p.*

Our algorithm can produce suboptimal packings. For example, consider three unit-length blocks of weights W, W, and 1, with $W > 1$. When our algorithm packs these blocks to place their joint center-of-gravity close to the midpoint of the bin, they appear in the sequence W - W - 1, and the center-of-gravity is distance almost 1/2 from the center of the bin for large values of W. In an optimal packing these blocks appear in the sequence W - 1 - W, and their joint center-of-gravity is exactly at the center of the bin. (Note, however, that any suboptimal packing produced by our heuristic must nevertheless be "good"; that is, the center-of-gravity must lie within $l_{\max}/2$ of the target point.)

2.1 Non-homogeneous blocks

With insignificantly more work, all of our bounds can be made to hold even when the blocks are non-homogeneous. In this case for each block we must determine not just a position in the bin, but also an orientation which tells whether the center-of-gravity of the block is to the left or right of its midpoint. One way to do this is to modify our heuristic slightly so that Step 3 becomes

> Place the jth block so that its geometric center O_j is as far as possible from p_j; then orient the block so that its center-of-gravity and p_j are on the same side of O_j.

It is straightforward to extend the arguments of Lemmas 2 and 3 to show that the bounds continue to hold for this more general version of the heuristic.

2.2 A more general problem

Our heuristic enables us to solve a more general problem of balance in which we are to pack the blocks to minimize the resultant moment about a given point p, about which there is an initial moment M. First transform the problem to one of minimizing the resultant moment about the point $p - M/(\sum_{i=1}^{n} w_i)$; then apply our heuristic. The resultant moment about M will then be either no greater than $(1/2)l_{\max}(\sum_{i=1}^{n} w_i)$ or else it will be as small as possible.

Note that the bound is independent of the structural details that induce the initial moment M. Such details include locations and magnitudes of any forces acting on the bin, how the bin is supported, and so on. For example, the initial moment might arise because the bin has non-zero and possibly non-homogeneous mass.

3 Alternative heuristics

To help better understand the performance of our algorithm we compared it with two alternative heuristics. The first, which was suggested by an anonymous referee, is disarmingly simple. Furthermore it can be shown that this algorithm has the same worst-case performance as our heuristic.

Algorithm PERMUTE**:** packs blocks so that their center-of-gravity lies close to target point p;

1 Sort the blocks from least dense to most dense; assume that the blocks are labelled so that this sequence is $(1, 2, \ldots, n)$;

2 Choose the best of the following n sequences:

$$(1, \ldots, n), (2, \ldots, n, 1), (3, \ldots, n, 2, 1), \ldots, (n, n-1, \ldots, 1)$$

Of course this algorithm can be implemented to construct the candidate sequences only as long as each successive one has center-of-gravity closer to the target point than the previous one. In any case, the worst-case effort is dominated by the $O(n \log n)$ effort to sort the blocks by density.

We tested algorithms BALANCE and PERMUTE on randomly generated problems and found, as expected, that they required about the same computational effort. However, even though they both provide the same guarantee (a packing for which the center-of-gravity is either within $(1/2)l_{\max}$ of the target point or else is as close as possible), in practice algorithm BALANCE was much more effective than PERMUTE. We discovered this in repeated trials of the following experiment: Generate an instance of the 1-dimensional balanced loading problem; solve it for each integral target point between 0 and the bin length and measure the directed distance from the target point to the resultant center-of-gravity. Figure 5 shows a typical outcome for algorithm PERMUTE. The structure is apparent: the graph is piecewise linear, with each piece corresponding to one of the n solutions produced by the heuristic. For comparison, Figure 6 shows the outcome when the same set of blocks are repeatedly packed by algorithm BALANCE. Its performance is clearly superior. We attribute this to the fact that algorithm BALANCE is capable of generating many more sequences than algorithm PERMUTE and so can be much more responsive to changes in the target point.

Another heuristic that is simple but capable of generating many sequences is 2-INTERCHANGE. This is a standard procedure for performing local search

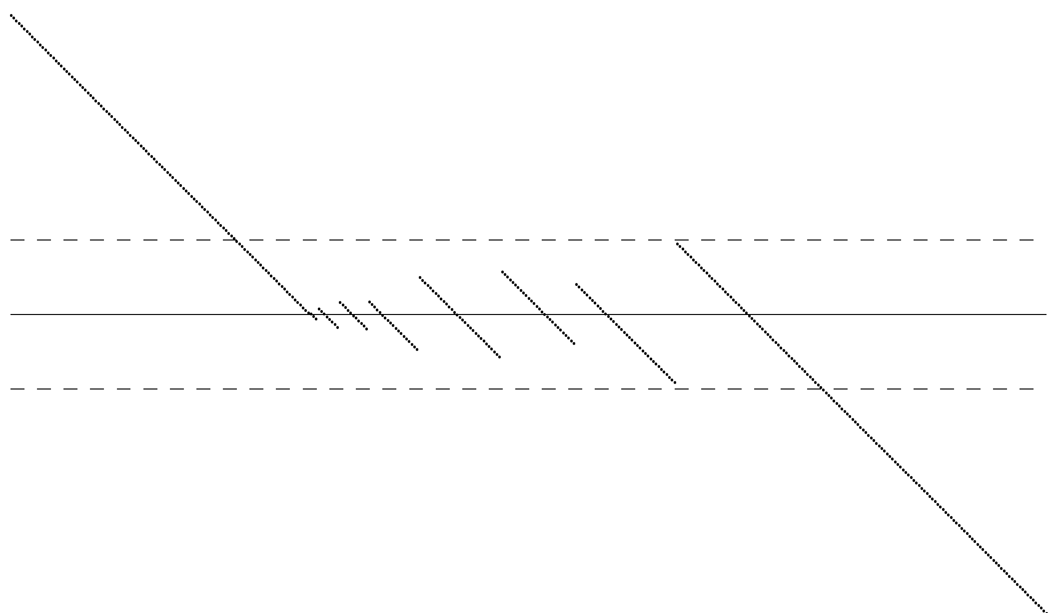

Figure 5: Performance of algorithm Permute on a fixed set of blocks for many different target points. The horizontal axis gives the location of the target point within the bin and the vertical axis gives the directed distance from the target point to the resultant center-of-gravity when the bin is packed by Permute. The dashed horizontal lines are at heights $\pm(1/2)l_{\max}$.

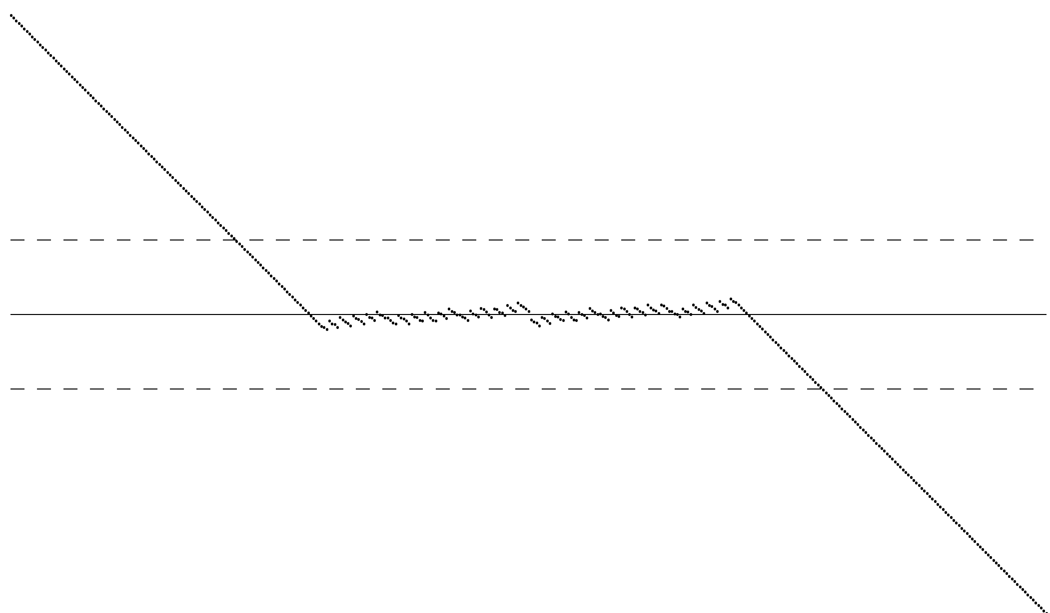

Figure 6: Performance of algorithm BALANCE on a fixed set of blocks for many different target points. The horizontal axis gives the location of the target point within the bin and the vertical axis gives the directed distance from the target point to the resultant center-of-gravity when the bin is packed by BALANCE. The dashed horizontal lines are at heights $\pm(1/2)l_{\max}$.

in combinatorial problems. Its natural implementation here is to begin with an arbitrary sequence and repeatedly interchange any adjacent pair of blocks for which the balance is improved. As pointed out by an anonymous referee, 2-INTERCHANGE has the same worst-case performance as our heuristic. This can be established by elaborating the following interchange argument: Consider a sequence at which 2-INTERCHANGE terminates and assume without loss of generality that the resultant center-of-gravity is to the left of the target point. Searching from the left, find an adjacent pair of blocks for which the denser block is to the left. If there is no such pair then the sequence is optimal. Otherwise interchanging those blocks would move the center-of-gravity no closer to the target point. From the algebraic formalization of this statement it follows that the loading produced by 2-INTERCHANGE must have the center-of-gravity no farther from the target point than one-half the length of the longest block.

Algorithm 2-INTERCHANGE can fail to load the blocks optimally, as shown in the following example. Consider four blocks, of lengths 1, 1, 2, 3 and of weights M, M, 2, 3 (where $M \gg 1$). Then for a target point of 3.5 (the midpoint of the bin), 2-INTERCHANGE cannot improve the sequence of weights 2 - M - M - 3, but M - 2 - 3 - M is optimal.

In testing we found the solutions produced by 2-INTERCHANGE remarkably good, as illustrated in Figure 7. However, despite its unreasonably good performance, 2-INTERCHANGE has some disadvantages. First its running time exceeded that of the other heuristics by an order of magnitude or more. In fact we doubt whether 2-INTERCHANGE can be guaranteed to terminate in polynomial time. In addition there is the aesthetic objection that 2-INTERCHANGE does not give any insight into the problem of balanced loading. Nevertheless, our tests suggest that 2-INTERCHANGE ought to be among those algorithms considered for practical use.

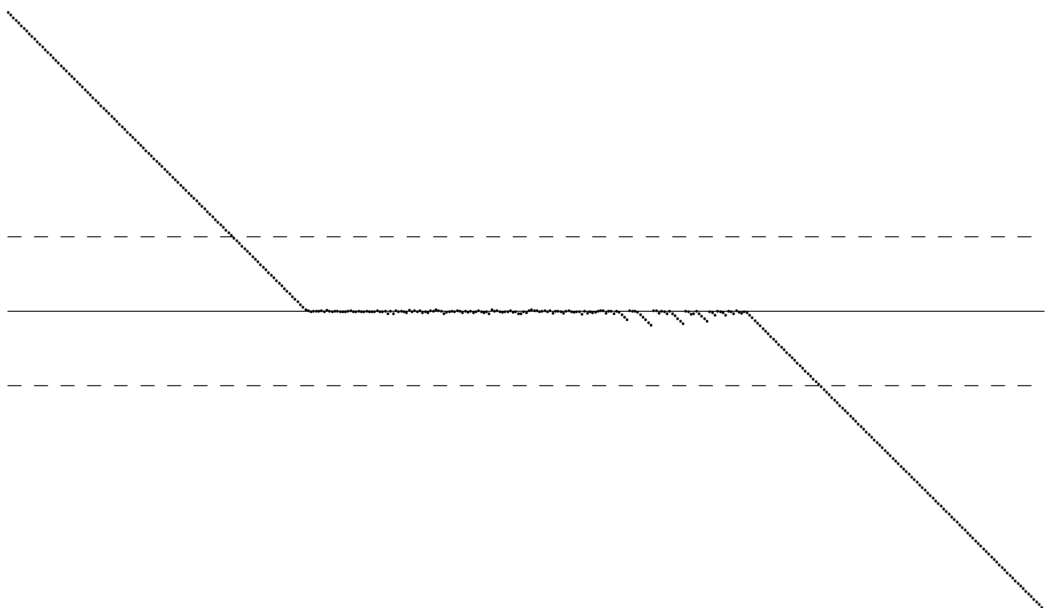

Figure 7: Performance of algorithm 2-INTERCHANGE on a fixed set of blocks for many different target points. The horizontal axis gives the location of the target point within the bin and the vertical axis gives the directed distance from the target point to the resultant center-of-gravity when the bin is packed by 2-INTERCHANGE. The dashed horizontal lines are at heights $\pm(1/2)l_{\max}$.

4 Related work

Mosevich (1986) and Laporte and Mercure (1988) have studied the issue of balance in building hydraulic turbines. Their problem was to assign turbine blades to positions around the perimeter of a wheel so that the resultant center-of-gravity of the wheel is as close as possible to its geometric center. This problem is 2-dimensional and apparently more difficult than ours. Also this application required extremely accurate solutions and it made sense to invest considerable effort toward this end.

The problem of 1-dimensional balanced loading is essentially a question of sequencing and so is similar to 1-machine scheduling. The problem of balanced loading is superficially similar to the problem of scheduling n jobs on a single machine to minimize weighted deviation from a common due date, which has been studied by Hall, Kubiak, and Sethi (1989), Hall and Posner (1989), and Hoogenveen and van de Velde (1989). In this problem, earliness costs are assessed against all jobs completed before the common due date and tardiness costs are assessed against all those completed after. The problem is to minimize the sum of costs. If we interpret the balanced loading problem as a scheduling problem, the target point is the common due date, moments about the target point are costs, and we want to minimize the *difference* between earliness and tardiness costs.

Despite such tantalizing similarities, "earliness-tardiness" scheduling and balanced loading seem to be fundamentally different. In a sense, earliness-tardiness scheduling is easier, since optimal schedules are highly structured: Let w_j be the (metaphorical) "weight" associated with the jth job and let l_j be its processing time; then in any optimal schedule all jobs completed at or before the common due date are scheduled in order of nondecreasing w_j/l_j and all jobs started at or after the due date are scheduled in order of non-increasing w_j/l_j. By exploiting this structure Hoogenveen and van de Velde (1989) have constructed a pseudo-polynomial time algorithm for earliness-tardiness scheduling with a common due date.

In contrast, no similar structure seems to hold for optimally-balanced loadings; and in fact, as we show in the next section, apparently there can be no pseudo-polynomial time algorithm for 1-dimensional balanced packing.

5 Complexity of balanced packing

Marc Posner of Ohio State University has observed in private communication that the problem of 1-dimensional balanced packing is strongly NP-complete. His reduction is from the 3-partition problem, which is known to be strongly NP-complete (Garey and Johnson, 1979).

3-partition

GIVEN: $3n$ items, where item i has "size" a_i and $b/4 < a_i < b/2$ for $i = 1, \ldots, 3n$, and $b = (\sum_{i=1}^{3n} a_i)/n$.

QUESTION: Does there exist a partition of the items into sets $A_1, \ldots, A_n$ where $\sum_{i \in A_j} a_i = b$?

Now given any instance of 3-partition, one can in polynomial time construct an instance of 1-dimensional balanced packing for which there is an exact balance if and only if there exists a 3-partition. The construction is as follows. There will be a total of $4n$ blocks. The first n blocks are all of unit length, but are especially heavy. Letting $K = bn^2$, block 1 is of weight $w_1 = ((b+1)n^2 - n) \sum_{i=2}^{n} K^i(2b(i-1) + 2i - 3)$; and the weight of block i is $w_i = K^i$, for $i = 2, \ldots, n$. Each of the remaining $3n$ blocks has the same density: block $n+i$ has $l_{n+i} = w_{n+i} = a_i$, $i = 1, \ldots, 3n$. The length of the bin is $(n+1)b$ and the target point is 1.

Now we determine necessary and sufficient conditions for there to a sequence of blocks with center-of-gravity at the target point 1. First observe that because block 1 is so heavy, its left corner must be at position 0 in the bin, where it induces moment $w_1/2$ about the target point 1. The only block that can offset the K^n term of this moment is block n, which must be placed with its left corner at position $(n-1)(b+1)$. If block n is placed at any other

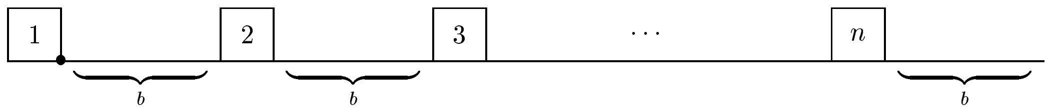

Figure 8: To balance about the target point •, the heaviest blocks must be placed as shown.

(integral) position, then a balanced solution is not possible. If it is placed to the left, then the total moment due to all the remaining blocks must be less than K^n and so it is not possible to balance about the target point. Similarly, if it is placed to the right, then it induces a moment greater than that due to block 1 and again the desired balance is impossible. By similar analysis, to countervail the K^i term of the moment induced by block 1, it is necessary that block i be placed with its left corner at position $(i-1)(b+1)$, as shown in Figure 5. The remaining blocks must then be placed in the n "gaps" of length b, where they would contribute a moment about the target point of $b/2+(b+1+b/2)+2(b+1)+b/2+\cdots+(n-1)(b+1)+b/2 = nb/2+(b+1)(n-1)n/2$. Therefore, to balance the $4n$ blocks about target point 1, it is necessary for there to exist a 3-partition of the $\{a_i\}$. This condition is obviously sufficient as well.

Posner has also observed that 1-dimensional balanced loading is NP-complete (in the weaker sense) even for the special case in which the target point p is the midpoint of the bin. This reduction is 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 $J = 1, 2, \ldots, n$ 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 $\sum_{j \in J_1} l_j = \sum_{j \in J_2} l_j$. Given such an instance, create an instance of balanced packing as follows. There are n blocks, with block j of length l_j and weight l_j, and there is an additional block of length $l_{n+1} = 1$ and weight 0. The target point is $(1/2)\sum_{j=1}^{n+1} l_j$, the midpoint of the bin. Now the center of gravity of a packing will fall at the target point if and only if block $n+1$ is centered over the target point; but this means the indices of the blocks to the left and right of block $n+1$ must determine a partition of the $\{l_j\}_{j \in J}$.

6 Conclusions

The 1-dimensional balanced loading problem is an example of a more general class that asks how a load should be distributed on a given structure. This is complementary to the traditional question of mechanical design, which asks for the structure to bear a given load. This class of problems seems rich and challenging and sufficiently new that we have taken the liberty of naming it "combinatorial mechanics". This is the first of a planned sequence of papers exploring these issues. Common themes, which may be observed in this paper, seem to be the NP-hardness of the problems and the design of heuristics to exploit a central fact about Newtonian mechanics: that force systems can be separated into their component parts and considered independently.

Acknowledgements

We thank Capt. Mike Ackley, USAF for helpful discussions about loading aircraft, and Marc Posner and Nick Hall for discussions about earliness-tardiness scheduling. We especially appreciate Marc Posner's sharing his result on the NP-completeness of 1-dimensional balanced packing. We also thank two anonymous referees for helpful suggestions, including the observation that 2-INTERCHANGE has the same worst-case bound as our heuristic. We also thank Marian Burge for help in legal citations.

S. Amiouny and J. Bartholdi were supported in part by the Office of Naval Research (N00014-85-K-0147). J. Bartholdi was also supported in part by a Presidential Young Investigator Award from the National Science Foundation (ECS-8351313). J. Vande Vate was supported in part by the National Science Foundation (DMC-8504466). J. Zhang was supported in part by the National Science Foundation (ECS-8658157) while at the Division of Applied Sciences,

Harvard University.

J. Zhang is at the Nanyang Technological Institute, Singapore 2263. The other authors are at the School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332, USA.

References

[1] Chapel, C. E. 1948. *Aircraft Weight, Balance & Loading*, Aero Publishers, Inc.

[2] *Civil Reserve Air Fleet Load Planning Guide*, MACP 55-41, 30 September 1987; Department of the Air Force, Headquarters Military Airlift Command, Scott Air Force Base, Illinois 62225-5001.

[3] Cochard, D. D. and K. A. Yost. 1985. "Improving utilization of Air Force cargo loading", *Interfaces*, **15**, 53–68.

[4] 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 Francisco.

[5] Hall, N. G., W. Kubiak, and S. P. Sethi. 1989. "Earliness-tardiness scheduling problems II: Deviation of completion times about a restrictive common due date", manuscript, Faculty of Management Sciences, Ohio State University, Columbus, Ohio 43210-1399, USA. To appear in *Operations Research.*

[6] Hall, N. G. and M. E. Posner. 1989. "Earliness-tardiness scheduling problems I: Weighted deviation of completion times about a common due date", manuscript, Faculty of Management Sciences, Ohio State University, Columbus, Ohio 43210-1399, USA. To appear in *Operations Research.*

[7] Hoogeveen, J. A. and S. L. van de Velde. 1989. "Scheduling around a small common due date", Report BS-89xx, Centre for Mathematics and Computer Science, PO Box 4079, 1009 AB Amsterdam, The Netherlands.

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

[9] Martin-Vega, L. A. 1985. "Aircraft load planning and the computer: description and review", *Computers and Industrial Engineering*, **4**, 357–369.

[10] Mosevich, J. 1986. "Balancing hydraulic turbine runners—A discrete combinatorial optimization problem", *European Journal of Operational Research* **26**, 202–204.

[11] Smith, W. E. 1956. "Various optimizers for single-stage production" *Naval Research Logistics Quarterly* **28** 643–651.

[12] Yeh, H. and J. I. Abrams. 1960. *Principles of Mechanics of Solids and Fluids*, Volume I: *Particle and Rigid Body Dynamics*, McGraw-Hill Series in Engineering Sciences, McGraw-Hill Book Company, Inc.