

Division of Research
School of Business Administration

May 14, 1988

LOAD PLANNING FOR SHIPMENTS
OF LOW DENSITY PRODUCTS

Working Paper #570

Robert W. Haessler
and
F. Brian Talbot
The University of Michigan

FOR DISCUSSION PURPOSES ONLY

None of this material is to be quoted or
reproduced without the expressed permission
of the Division of Research

Copyright 1988
University of Michigan
School of Business Administration
Ann Arbor Michigan 48109

Abstract

This paper presents a complex computer-based heuristic procedure for sizing customer orders and developing three dimensional load diagrams for rail and truck shipment of low density products. This heuristic procedure was developed for, and is in various phases of implementation at, a large multinational U.S.-based consumer products company. Products are shipped daily in high volume from inventory in corrugated containers of various sizes depending on the product package sizes and customer requirements. Vehicles used include railcars, truck trailers and tandem truck trailers, which also vary in size depending upon need and availability. In most cases, product volume or material handling considerations limit the amount of product loaded into vehicles before weight restrictions are met. Hence, the emphasis here is on low density products. The procedure developed has been demonstrated to significantly increase vehicle utilization, and improve customer service. It is fast and accurate enough to be used in real time during the order entry process. It has also been used successfully in a vehicle feasibility study of single versus tandem trailers.

Introduction

This paper describes a complex computer-based heuristic procedure for sizing customer orders and developing load patterns for truck and rail shipments of low density products. The density of the products is such that truck trailers as well as rail cars are loaded primarily based on volume rather than weight restrictions.

The procedure is designed to run on-line at order entry time. When an order is phoned in to the company, the product code and desired quantity of each item on the order are keyed into the computer. The order pricing system encourages the customer to order in full carload or truckload quantities. Because the products are all packed in corrugated boxes of varying size, it is difficult to know if any order will completely utilize the cubic volume of the shipment vehicle. Prior to the development of this system, general guidelines based on total order quantity were developed to help the customer determine how much should be ordered to fill the shipment vehicle for each vehicle size that might be used. These guidelines had to be conservative to insure that the order could be loaded regardless of the product mix. The clamp truck driver who loaded the order was responsible for deciding how to pack the order into the vehicle so that the order could be shipped complete without damage to the product. The order was not increased after it had been entered even if it became apparent at loading time that additional products would fit on the vehicle. This meant that in many situations cubic utilization was significantly less than it could have been.

With this new procedure, the customer order is submitted to the vehicle loading heuristic, along with some guidelines as to how the order can be increased or decreased and the dimensions of the shipment vehicle. The heuristic returns an adjusted order that "maximizes" the cube utilization of

the vehicle without undue risk of product damage, and a diagram that provides precise instructions as to how that order should be packed into the vehicle.

No explicit mathematical statement of the problem is developed and therefore no claim can be made for optimality. The development of an explicit model is made impossible because of the nature of the non-physical constraints. The desire is to maximize the size of the order in cubic feet, while at the same time

1. avoiding the risk of product damage.
2. limiting the time required to load the vehicle.
3. keeping each product in the order as close together as possible to minimize the effort in unloading inventory and storing the product at the receiving location.

It is clear that if individual corrugated boxes containing the product ordered were individually hand loaded into the vehicles, significant improvements in cube utilization would be obtained. It is also clear that the increased labor cost and loading time of doing this would far outweigh any benefits provided by a reduction in transportation cost.

Literature Review

There is a rich literature dealing with cutting stock and packing problems. The seminal work of Gilmore and Gomory [1961, 1963, 1965] demonstrated how linear programming could be used to maximize yield in one, two and three dimensional cutting stock or packing problems. More recently, Christofides and Whitlock [1977] and Beasley [1985] have developed sophisticated procedures for generating two dimensional patterns for rectangular shapes.

Although these developments provide useful insights and tools, they do not address directly the issues of loading industrial products into fixed three dimensional vehicles based on volume with restrictions to avoid damage

and localize products within the vehicle. There is relatively little literature available on loading low density products. Roberts and Taylor [1972] developed a procedure for loading full width racks of stamped metal parts into railcars. Because the racks had a limited number of base dimensions and had to be stacked securely, he was able to reduce the packing problem down to solving two one-dimensional bin packing problems. The first step involved building the minimum number of stacks up to the height of the boxcar, and the second step was to minimize the number of railcars required by packing the stacks into railcars along the length of the car.

Haessler [1979] reported a similar application involving loading large rolls of paper into railcars. This application was in one sense simpler than the one describe above because the rolls were loaded on end and the layout issue simply involved finding the number of circles of constant diameter that would fit into a rectangle. Individual stacks were constructed by solving a one dimensional trim problem. However, Haessler's procedure was more complicated in that it used freight rates by railcar size to determine the mix of cars to be used to minimize the total cost of shipping a multiple railcar order.

Literature searches carried out independently by us and the client organization did not uncover any research on the type of problem discussed here.

Background

The company has about one hundred different packaged products which it sells from inventory. The package size varies significantly with the largest size being about 10 times the volume of the smallest. The packaged product is shipped in cases (corrugated boxes) that also vary by a factor of 10 from largest to smallest. The cases are stored in the warehouse in what are called

unit loads. A unit load may contain as few as 10 to over 100 cases. The unit loads were designed so that each is as close as possible to being 40 inches wide, 48 inches long, and 60 inches high. One layer of a unit load is referred to as a tier. All the tiers in a unit load need not be the same heights or have the same number of cases. Because of variation in case sizes, the actual dimensions for a unit load may be off by anywhere from three to six inches from the target value. Figure 1 shows a sample unit load with the accompanying variation by dimension. The unit loads were originally designed to ship in railcars two high and two across the width of a car.

Insert Figure 1 About Here

Although customers are encouraged to order in full unit loads, they may order in both case and tier quantities. Prior to the development of this system, guidelines were developed as to how many units would fit in each size vehicle. These were disseminated to the customer through the sales force and full vehicle orders were entered according to these guidelines. Unless someone "knew" that a larger number of units would fit, perhaps because all the units being ordered were on the low side of the target unit size, the order was entered and shipped consistent with the established guidelines.

Because a customer could order any mix of products, the guidelines had to be established in a conservative manner to insure that the order could be loaded as entered. This resulted in the cube utilization of most vehicles being lower than necessary. It should be noted that adding to the order after it was loaded when it was apparent that more product would fit, would not necessarily give as good a result as first increasing the order and then loading that entire larger order into the vehicle. The best way to handle

the extra products may be to intermingle them with product on the original order, rather than load them by themselves in the gap that is left after the original order is loaded.

To summarize:

The objective was to develop a computer-based procedure to be run at order entry time, that will size the order and diagram the loading of that order so as to increase vehicle utilization and reduce freight cost.

The inputs to the computer system are

- initial customer order (including all product dimensions)
- guidelines for changing the orders (which products should be added or dropped first)
- dimensions of the shipment vehicle

The outputs from the system are

- suggested order changes
- load diagram

The guidelines for changing the order show the sequence in which product will be added to and dropped from the order. The original order can be either increased or decreased, not both.

Load Diagram

The easiest way to understand the loading problem is to view the top and side view of a load diagram for both railcars and trucks. A partial diagram for railcars is shown in Figure 2. In the top view, or footprint, there are two columns of stacks of product. Because no product has a base length greater than 52 inches, it is always possible to load two columns of stacks without concern for interference from side to side. Stacks 1, 4, and 6 represent one unit stacked on top of another unit. This is the most common situation in railcars. The units were designed originally to facilitate

loading units two high in railcars. Stack 2 is also 2 units, but with a tier of some product sandwiched between to make use of the available height in the railcar. Stack 3 is a number of tiers stacked on top of a unit load. Stack 5 is handloaded cases of a single product which is inserted to fill the gap in the doorway. This is done if there is a gap large enough to fill with loose cases.

Insert Figure 2 About Here

A typical trailer load diagram is shown in Figure 3. In the top view or footprint, there are several common configurations across the width of the truck. Section A shows 2 units with their long side along the width of the truck. This is possible in a 95" wide trailer if the unit base length is less than the target value of 48" and is the generally the best utilization of space. Section B shows 4 units in a box-like arrangement that is also quite efficient in terms of space utilization. Section C is sometimes the only way to layout the units with larger base dimensions and is generally the least efficient utilization of space. In the side view, Stack 1 shows one unit on top of the other. This can occur when the unit height of some product is at the very low end of the height range. Stack 2 shows a common arrangement in which 1/2 of a unit is placed on top of a full unit. This generally does not completely utilize the height of the vehicle, but is very efficient from a material handling standpoint. Stack 3 is the same as Stack 2 except a tier of some product has been sandwiched between a unit and 1/2 unit to better utilize the available trailer height. Stacks 4 and 5 show tiers of different products on top of a full unit. This also may lead to better vehicle height utilization. Finally, stack 6 is made up of tiers of a variety of products.

This also tends to yield good vehicle height utilization. Loose cases may also be hand loaded to fill any large gaps at the rear of the truck.

Insert Figure 3 About Here

Solution Procedure

The solution procedure begins by estimating the number of stacks that will fit into the vehicle. The order is then adjusted by adding or deleting units so that the adjusted order requirements will generate the estimated number of stacks. The next step is to determine if a footprint can be found that locates the base-dimensions of the stacks in the rectangular area of the vehicle floor. If a footprint can be found, then the number of stacks is increased by one and the process is repeated. If no footprint is found, the number of footprints is reduced by 1 and the process repeated. The procedure stops when the footprint containing the largest number of stacks is found. A detailed description of this procedure follows:

Step 1. Estimate the number of stacks, NS.

For railcars,

$$NS = \left[\frac{2*VL}{AUW} \right]$$

where VL - vehicle length,

AUW - average width of all units on order,

[a] - largest integer less than or equal to a.

For trailers,

$$NS = \left[\frac{.95 * TFA}{AUBA} \right]$$

where TFA - trailer floor area,

AUBA - average unit base area.

Step 2. Calculate the number of equivalent units, NEU, that can be included in NS stocks.

For railcars,

$$NEU = 2 * NS$$

Tier quantity orders are grouped into unit size blocks for comparing units on order, UOO, to NEU. Maximum tier block height is defined to be one half of the vehicle height.

For trailers,

$$NEU = 1.5 * NS + .5 * [NSU/2]$$

where NSU is the number of short units where short is defined to be less than one half of trailer height.

Tier quantity orders are conceptually grouped into either full unit, half unit, or full vehicle height stacks for comparing UOO, to NEU.

Step 3. Adjust the order up or down based on the difference between UOO and NEU. If NEU is greater than UOO, units or half units are added to the order. If UOO is greater than NEU, units are deleted from the order. Adding or deleting is done based on a predetermined list that controls the sequence in which changes are made. This list is a function of the particular customer's usage and ordering patterns.

Step 4. Build NS stocks for the ordered items as adjusted in Step 3. There are two primary considerations in building stacks in addition to the constraint that the stack height not exceed the vehicle height.

- very heavy products must be on the bottom of the stack.
- all the products in the stack should have identical or at least similar base dimensions. The ideal situation is to have only one

product in each stack. If this is not possible, then the difference should be small and the larger dimension along the length of the vehicle should be on top. This will minimize damage due to shifting of product along the length of the vehicle as the vehicle accelerates and stops.

Step 5. Determine if a footprint can be found for the NS stacks generated in Step 4.

For railcars, finding a footprint reduces to solving a knapsack problem. Let SW_j be the width of stack j . Clearly to be feasible,

$$T = \sum_j SW_j \leq 2 * VL$$

where VL is the usable length of the vehicle. NS stacks will fit in the railcar if values of $X_j = 0,1$ can be found such that

$$T - VL \leq \sum_j SW_j X_j \leq VL$$

Once again there are soft constraints that should be satisfied if possible. These are,

- Localize all of each product in some section of the vehicle to facilitate checking the order as it is loaded and unloaded.
- Minimize the height difference between adjacent stacks so that the top tiers are effectively blocked in and cannot move very much if they become loose.
- Save short, lightweight stacks for the doorway. This is the most difficult area to load because of limited maneuverability of the clamp trucks.

In order to meet these constraints, an attempt is made to solve the knapsack problem heuristically. If this attempt is unsuccessful, then a mathematical solution to the problem is found, if one exists, using a lexicographic search algorithm.

For trailers finding a footprint is more complex. This is discussed in detail in the next two sections of the paper.

Step 6. If a footprint is found for NS stacks, increase NS by 1 and go to Step 2.

If no footprint is found, reduce NS-1 and go to Step 2 if NS-1 has not been previously checked. If a footprint has been found for NS-1, go to Step 7.

Step 7. Fill in major gaps in the vehicle with hand loaded cases. Such gaps smaller than the base dimensions of a unit can occur in any load configuration.

In railcars there may be a gap on either or both sides that is significant but less than the width of a stack. This space is intentionally positioned in the doorway in the loading diagram and filled in by hand loading cases. Alternatively, air bags may be used to fill this space.

In trucks the gap is forced to occur at the back of the trailer and may be on one or both sides. Again, cases of a single product are hand loaded to utilize as much of that gap as possible.

Trailer Footprint Procedure

The trailer footprint procedure is more complex. This is due to a combination of factors including greater variety of possible loading patterns, inventory handling and staging concerns, product damage concerns, axle weight limitations, multiple customers and tandem trailers. First, as with rail, at most two stacks can fit side by side across the width of a trailer. However, unlike rail, each stack can be oriented with either its long side or short side parallel to the long side of the trailer. The former is referred to as a rotated orientation. As indicated in Figure 3, this leads to a number of

possible loading patterns. Inventory and staging considerations are generally more complicated with trucks than with rail, due to the fact that the load contains a greater variety of product types, which translates into more complex stack building and material handling. In order to reduce staging complexity, the footprint procedures give high priority to positioning stacks containing identical products close together. A secondary but important benefit of this is a reduction in clerical counting errors at both shipment and receiving.

Concerns about minimizing product damage are operationalized in several ways. One already mentioned is keeping identical products close together, which has the effect of reducing loading dock congestion and consequential handling damage. Other, often conflicting, mechanisms are: (1) to place the highest stacks in the front of the vehicle to avoid having tall stacks tip over shorter ones during sudden stops; (2) to follow a prespecified stack orientation preference (rotated or non-rotated) which is defined prior to loading as being the only acceptable way to avoid in-transit damage to a particular stack. (This could be due to the unevenness of a vertical face of a stack, or due to the stack load bearing geometry resulting from how cases are built into tiers, and thence, into unit loads); (3) to place stacks such that the front/back vertical faces remain parallel during sudden stops (There is a need to avoid having a following stack rub against just the corner of the predecessor stack on the opposite side of the trailer.); (4) to tightly load stacks to minimize side to side movement. This is most important at the back of a trailer load, where loose products could potentially move sideways or backwards into the empty gap that may exist.

Unlike rail shipments where volume constraints are always met before weight constraints, truck shipments can exceed axle weight limits. Hence, an

attempt is made to develop load patterns within vehicle axle weight limits. Unlike rail shipments, it is possible with trucks to specify the order in which product groupings for a given customer are to be dropped off, and it is possible to ship several customer orders in a given trailer. In these situations, products within a group or for a given customer must be loaded contiguously to avoid clerical errors and extra material handling at the drop-off points. Finally, truck tractors can pull either a single trailer or two trailers in tandem. The latter is more complicated than simply taking a loading pattern for a larger trailer and splitting it in two. All of the above concerns, flush faces, height, orientation, axle weight, etc., come into play for each trailer. In addition, if like product (or product group, or customer group) cannot totally fit in the back of the leading trailer, the carryover should go in the back of the tandem so as to facilitate material handling at the drop off point. (The two trailers are loaded and unloaded simultaneously.)

A very simplified 0-1 integer formulation of the single trailer footprint problem follows.

Let,

W = floor width of trailer in inches.

T = floor length of trailer in inches ($t=1, \dots, T$
where $t=1$ indicates front of trailer).

w_{jm} = footprint width in inches of stack j in orientation m .
($m=1, 2$ indicating rotated or not.
 $j=1, \dots, N$, where N is a dummy stack with zero footprint.)

$$x_{jmt} = \begin{cases} 1 & \text{if the following edge of stack } j \text{ in orientation} \\ & m \text{ is positioned at location } t \\ 0 & \text{otherwise} \end{cases}$$

Then the problem of minimizing the usable floor space is:

$$\text{Min.} \quad \sum_{t=1}^T tx_{N1t} \quad (1)$$

$$\text{S.T.} \quad \sum_{m=1}^2 \sum_{t=1}^T x_{jmt} = 1 \text{ for } j=1, \dots, N \quad (2)$$

$$\sum_{j=1}^N \sum_{m=1}^2 \sum_{q=t}^{t+w_{jm}-1} w_{jm} x_{jmq} \leq W \text{ for } t=1, \dots, T \quad (3)$$

$$\sum_{t=1}^T tx_{N1t} - \sum_{t=1}^T tx_{jmt} \geq 0 \text{ for } j=1, \dots, N-1, m=1, 2 \quad (4)$$

Constraints (2) insure that each stack is loaded only once. Constraints (3) insure that the vehicle width is not violated, and (4) permits the objective function to force loading towards the front of the trailer. It should be pointed out that (3) is valid in this loading environment because no more than two stacks can fit side by side across the width. Constraints (3) would have to be modified if more than two stacks could fit side by side across the trailer width in order to maintain contiguity of the space available for loading a stack.

Early in this research process two 0-1 integer implicit enumeration algorithms were developed and tested for solving this problem. For the following reasons, however, these algorithms were dropped in favor of the heuristic procedures to be described. First, it was quickly discovered that optimality was impossible to verify within an acceptable period of time, and the quality of nonoptimal solutions found during partial enumeration was difficult to control. As mentioned earlier, the client organization required that the loading diagram we prepared while their customer was still on the phone with their order entry person. This was translated into a contractual specification of generating the diagram within two seconds of CPU time on their order entry computer. Testing indicated that two seconds was not

sufficient time for our optimization codes to be of much value, especially given all the data manipulation and stack building that had to also occur within those two seconds. We developed the second optimization code, because, as the formulation indicates, this 0-1 approach is sensitive to the metric used. When the client decided that solution accuracy had to be increased from inches to tenths of an inch, the size of the formulation grew significantly. Hence, in the second code, data structures and logic were developed to make the solution procedure computationally metric free. This significantly reduced solution times, but not enough to make the optimal seeking approach practical.

Other major reasons for selecting a heuristic are: (1) The difficulty of explicitly modeling the client's notion of "closeness" of like product (recall that each stack can contain multiple product types), and (2) the objective function and constraints are not firm. In reality, the problem is multicriteria. For example, some of the prespecified stack orientation preferences in a load might be relaxed in order to achieve a tight back end on the trailer. Another example of this is that when customers do not want a full load, then it becomes important to spread the product across the trailer floor to minimize product damage, rather than minimize floor space as is implied by (1)-(4).

Trailer Footprint Heuristic

The trailer footprint heuristic has two stages, and can involve more than one pass of Stage Two. The first stage develops a list of all stacks that specifies the order in which stacks should be considered for loading. No effort is made to actually load in Stage One. Stacks are ordered on this list into contiguous groups. The hierarchy of grouping from most to least inclusive is: (1) all stacks; (2) by customer; (3) by drop off sequence; (4)

by likeness of product (two levels used). Within likeness groups, stacks are ordered from tallest to shortest.

If unusual height variation occurs within a common product group, then the unusually shorter stacks are split off and treated as separate groups. A similar sequencing takes place across drop off groups and customer groups, using median measures of height, without violating any firm ordering constraints. Analogously, to deal with axle weight constraints, stacks are sequenced by alternating the heaviest and lightest stacks within the group without violating firm drop off ordering or height restrictions. The effect of this is to approximate an even weight distribution. The ordered sequence of stacks obtained from the above procedures is then passed to Stage Two.

Stage Two sequentially "loads" the stacks from the ordered list in much the same way as the stacks are physically placed in the trailer, starting at the front and filling in towards the rear. In Stage Two actual dimensions, orientation preferences, heights and other product damage concerns are taken into account. Each stack is taken from the list and is positioned either on the right or left side of the trailer immediately behind the existing loaded left or right stack. Hence, as each stack is considered for loading, two decisions have to be made: stack orientation (rotate or not) and whether it should be placed on the left or right side of the trailer. From the most desirable to least, one would generally prefer the following results: (a) two non rotated stacks side by side; (b) boxing of four stacks; (c) one non rotated and one rotated side by side; and finally, (d) two non rotated side by side. (See illustrations in Figure 3.) The major factors that might mitigate against reaching the ideal utilization pattern in Stage Two include concerns of prespecified stack orientation preference, product closeness, insufficient vertical faces between adjacent stacks, a loose back end, and, of course,

stack footprint dimensions. (Preprocessing in Stage One largely takes care of height and axle weight concerns, and simplifies the process of keeping like product, or customer groups, together in Stage Two.)

As each stack is loaded, an analysis is made of the current partial solution, and characteristics of the stack to be loaded are noted along with unusual characteristics of unloaded stacks. (An example of the latter might be that all remaining unloaded stacks must be rotated due to their sizes.) An effort is made to orient and position the stack such that the most likely outcome is (a) then (b), (c) or (d), respectively. As the procedure begins to fill the vehicle, increasingly more attention is paid to the tightness of the back end, and outcomes (a) - (d) are viewed as being equivalently desirable.

If the above process is successful in creating a feasible load, and that is all that is required, then the procedure goes to Step 6. If the quality of the solution is unsatisfactory, or if the objective is to maximize the number of stacks that will fit in a vehicle, then there could be two more passes through Stage Two. During the second pass, orientation preferences are relaxed for those stacks in the last half of the load that could likely be oriented in either direction without incurring in-transit damage. In pass three, orientation preferences for all such stacks are relaxed. The procedure selects the best of these three solutions, and returns it to the stack building part of the procedure (Step 6).

Conclusion

Industrial packing problems are complex, and have numerous soft constraints. As such, they do not have concise mathematical formulations and corresponding optimization solution procedures. Sequential heuristic procedures, on the other hand, can incorporate preferences and soft constraints as the search is carried out and selections made. Sequential

heuristics are also, generally much faster than optimization algorithms. This is extremely important in this application because of the large number of orders processed each day and the requirements for immediate feedback during the order entry process.

The methods presented in this paper have been used by the company sponsoring this research in loading railcars since late 1987, and parallel testing of the truck portion has been ongoing since early 1988, both with very favorable results. The major benefits recognized are significantly higher vehicle utilizations, and improved customer service. Included in the latter are all the benefits associated with increasing the certainty that orders will arrive as originally specified at order entry time. In addition, the methods presented are able to be used at loading time to quickly and accurately adjust to last minute changes, such as having to replace a vehicle with a different size trailer, or to change load patterns because of customer dictated needs. Finally, these procedures have been used outside their intended operational setting to perform a detailed feasibility study of single versus tandem trailers for another manufacturing/distribution region for the company sponsoring this research. Without these procedures, it would have been virtually impossible for this organization to accurately carry out such a study in a timely and economical manner.

Computerizing the process of preparing loading diagrams is a major hurdle that has been overcome as this company continues its drive towards automating all material handling.

REFERENCES

1. J. E. Beasley, "An Exact Two-Dimensional Non-Guillotine Cutting Tree Search Procedure," Operations Research, 33, 49-64, (1985).
2. N. Christofides and C. Whitlock, "An Algorithm for Two-Dimensional Cutting Problems," Operations Research, 25, 30-35, (1977).
3. P. C. Gilmore and R. E. Gomory, "A Linear Programming Approach to the Cutting-Stock Problem," Operations Research, 9, 849-859, (1961).
4. P. C. Gilmore and R. E. Gomory, "Linear Programming Approach in the Cutting-Stock Problem - Part II," Operations Research, 11, 863-88 (1963).
5. P. C. Gilmore and R. E. Gomory, "Multistage Cutting Stock Problems of Two or More Dimensions," Operations Research, 13, 94-120, (1965).
6. R. W. Haessler, "Solving a Boxcar Loading Problem," Presented at the joint ORSA/TIMS Conference, May, 1979.
7. D. Roberts and D. Taylor, "A Distribution and Railcar Loading System," paper presented at the Joint ORSA/TIMS Conference, November, 1972.

Figure 1

Illustration of a Typical Unit Load Configuration




Figure 2
Section of Typical
Railcar Load Diagram

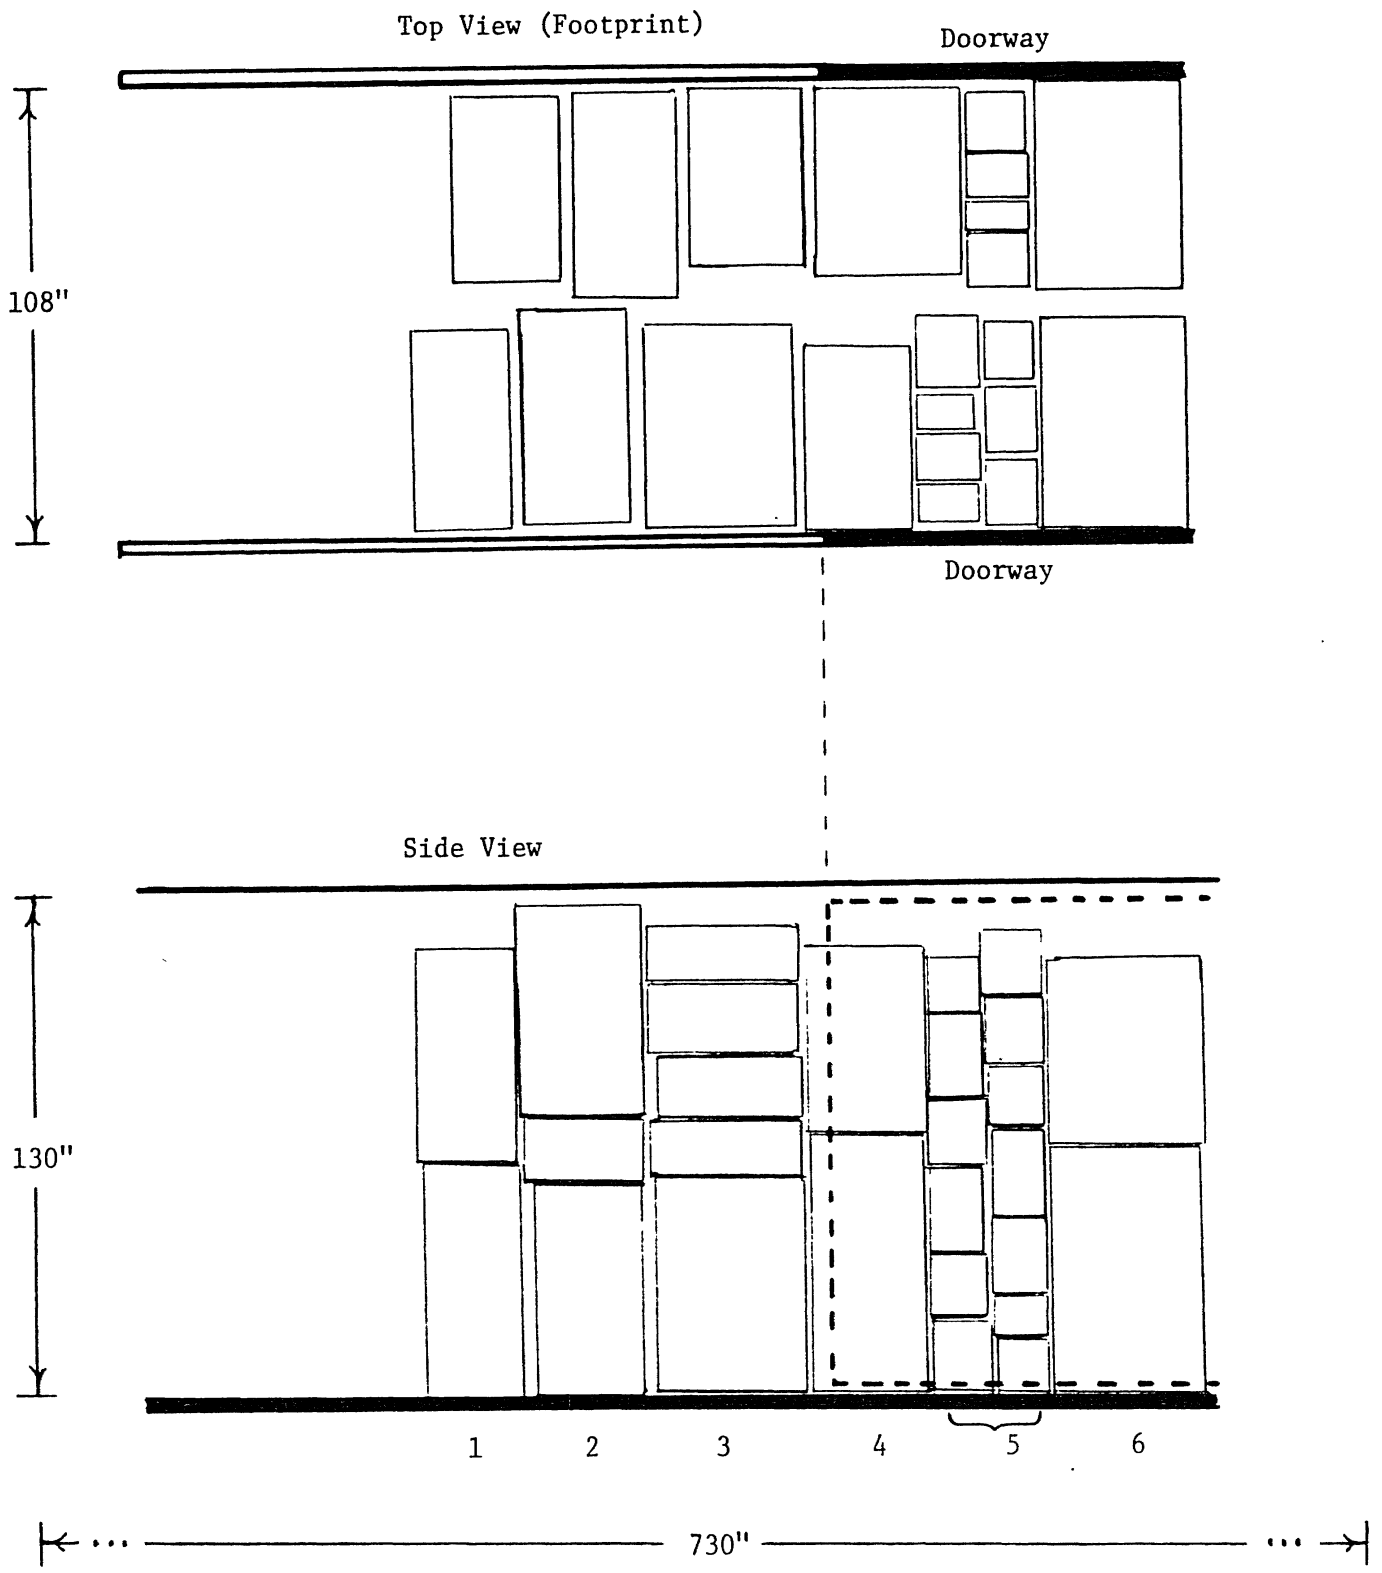



Figure 3
Section of Typical
Trailer Load Diagram

Top View (Footprint)



Front
of
Trailer

Side View

