THE ALLOCATION OF SWITCHING WORK IN A SYSTEM OF CLASSIFICATION YARDS*

M. Beckmann   C. B. McGuire
T. C. Koopmans   C. B. Winsten

P-448

3 December 1953

The RAND *Corporation*
1700 MAIN ST. · SANTA MONICA · CALIFORNIA

The Allocation of Switching Work in a System
of Classification Yards*⌀

Martin Beckmann, Tjalling C. Koopmans, C. B. McGuire, and C. B. Winsten

A week or so ago the newspapers carried reports of the General Motors "Rail-Truck Coordination Plan" which is mainly concerned with the operation of piggy-back cars. The reports would have us believe that the operation of these new railroad cars, which carry loaded truck trailers, will add the short-haul economies of trucking to the long-haul economies of rail transport. Without entering into the merits of the plan itself, we would like to look a little closer at one of the proposals involved.

Rail transport is slow, according to the report, partly because loaded cars spend so much "dead" time in yards. Big investments go into providing fancy equipment which speeds the actual breaking-up of incoming trains, but once into the yard the cars just sit there. The way to avoid this wasted time, says the report, is to keep traffic out of yards altogether. Run only "solid" trains, that is, trains containing cars for only one destination. Trains can then haul their load from where it originates to its destination without passing throught the very time-consuming switching operations at intermediate yards. Well, how about it?

Put baldly, the plan perhaps attributes more naivete to its proposers than they deserve. At any rate, it is an interesting and profitable question to examine. If it is silly, why is it silly? And is there a good suggestion in it someplace?

First it must be recognized that loaded cars do in fact spend a large part of their time standing in yards. Though nobody seems to have a firm estimate of what this proportion is, one can find reports for particular roads of one-

third dead time. At least we know that it is big enough to warrant serious study.

Then of course we are interested in transit time not only because the shippers value time, but also because the same job can be done with less rolling stock if overall speed is increased. The importance of these yard delays is underlined by the fact that the subject is attacked anew each year by such organizations as the American Association of Railroad Superintendents which has gone on record as disapproving of this "unproductive" yard time -- unproductive because it produces no ton-miles. Now in outlining a very preliminary analysis of this problem we hope to make clear that this yard time is by no means completely unproductive. We want to reduce it as far as possible, of course, subject to other costs, but after a certain point it costs us more to reduce this dead time than it's worth. It's even conceivable that this point has already been reached, if the railroads have been as efficient as possible in their operations. This is doubtful though, as almost anyone will admit, however friendly he is to the railroads. The problem is simply too difficult for us to believe that the solution has been reached.

Let us ask ourselves first why the problem arises at all. Why are classification yards necessary? No one has complained that loaded highway transports spend too much time standing idle. And, loading docks aside, there are no big truck classification yards for people to look at and wonder about all the time being wasted.

The first reason for yards and the basic cause of these classification problems, is that the size of the ordinary shipment is smaller by far than the size of a train. The second reason is that setting out or picking up a car or two is a costly procedure which consumes time, fuel, wages, etc. If the size of a train were exactly equal to the size of a shipment, there would

be no reason for switching, and yards would have no function whatsoever. Locomotives would simply pick up their load from a shipper and take it directly to the consignee; yards would not facilitate its journey. This in fact is just about the situation we have in truck transport and it explains the absence of a classification problem in that business.

Suppose however, that on the average it takes several shipments to fill the ordinary sized train. Then if the higher costs of short trains are to be avoided, two alternatives are open to us:

(1) Single trains may carry shipments bound for different destinations, and cars may be shifted from train to train at points called classification yards, or

(2) Shipments may be allowed to accumulate at their point of origin until whole trainloads are ready.

Notice that the second alternative is precisely the "solid train" proposal of the General Motors plan. The only trouble with it is that it brings about quite a bit of dead time, if not in classification yards proper, then on the shippers' tracks while the necessary accumulation of cars takes place. So we are left with the choice of either making use of classification yards or suffering these very large delays of accumulation. It is to this problem of the use of classification yards that we now turn our attention.

Probably no one in the audience, and particularly no operating man, will deny that the problem is one of a high order of complexity. However, we hope to show that the possibilities of analytic methods in this area are by no means small. We should also like to emphasize that the approach throughout is an economic one; we are asking how the most product can be produced with given resources, or, turning the statement around, how a given product can be produced with the least amount of resources. We want to distribute work over a system of yards so that another dollar spent will result in an equal value

of transportation-product no matter where in the system this dollar is applied, so long as it is properly applied.

To simplify our terminology, let us think of the tracks from which cars are originally picked up and on which they are finally set off, as small flat yards. In this way, we can think of all traffic as being yard-to-yard. Now as we have seen there are two extreme ways of running a system of yards. One is to let traffic bound for exactly the same destination to accumulate at the originating yard until it fills a train. This is obviously wasteful due to the large accumulation delays which would build up with the usual sized shipments. The other extreme is to cut down on accumulation delays by sending off a train the minute enough traffic going in that general direction is available. This is also wasteful for it means that little or no work has been done on the train to facilitate set-offs. Mr. Lake mentioned some of these troubles a few minutes ago.

It is clear that some happy medium between these extremes must be found. The means by which this middle ground is reached is prior classification, sometimes called grouping or blocking. It means that some yards of the system do more work, in the way of grouping and lining up the train, than is absolutely necessary at that yard. Nearly all yards try to group their local trains and put the groups in station order. The way we have defined yards, to include destination tracks, means that we have to regard this station ordering of locals a case of prior classification, so in this sense nearly all yards of any size do some prior classification.

What we are really interested in however, is the prior classification done by the larger yards. We would like to know the costs incurred by the yard that does it, and how much work it saves other yards farther down the line. In a sense, prior classification is a device used to shift work from

yard to yard. If the costs, in terms of both money and time, are known at each of the yards, and if the way these costs change with changes in the degree of prior classification, then it should be possible to distribute the work over the whole system of yards in such a way as to even out congestion to the point where total costs are reduced as far as possible.

Before discussing these costs in more detail it should be said first that we are concerned here with what the economist would call the "short-run" problem of distributing classification work. That is, we have a given set of yards of fixed sizes and capabilities. It isn't open to us, say, to change a flat yard into a hump yard, or to add new tracks or retarders to existing yards. We restrict our discussion to a period that is too short for changes in these items, and we then ask the question of how we should act in this short period. In any case, this short-run problem has to be solved before much progress can be made with longer-run ones. Secondly, we suppose that the origin to destination traffi flows are nearly constant; a little bunching or dispersal is not ruled out, but no traffic is growing or declining steadily. This is fairly realistic for the short run. We also assume here that the routes each component of traffic takes are fixed from beginning to end. In other words, this is not a routing problem.

The problem is to find, for each yard of the system, three things:

(1) The points to be classified, that is to say, the grouping, or blocking of cars into groups,

(2) The make-up of trains, that is, what groups to into what trains,

(3) The scheduling of trains. This may be in the form of a departure time strictly adhered to, or perhaps a rule that a particular train will be dispatched when filled, or possibly some compromise between these two.

We shall now discuss how, for a single yard, these three "policy var-
iables" together with the characteristics of incoming traffic and the phys-
ical aspects of the yard combine to determine the various costs of classifi-
cation operations. The discussion in many respects will be incomplete; for
instance we shall take the third variable, scheduling, to be fixed through
most of the discussion.

Two major categories of cost are of interest to us -- money costs, rep-
resented by switch-engine cost and wages, and time costs, consisting of delays
of various sorts. In both cases averages are used, ideally a separate average
for each component of traffic.

It is convenient to define four or five different kinds of delay (see
Figure 1).

handling —— switching —— sorting
trimming

miscellaneous — icing
inspection
lubrication
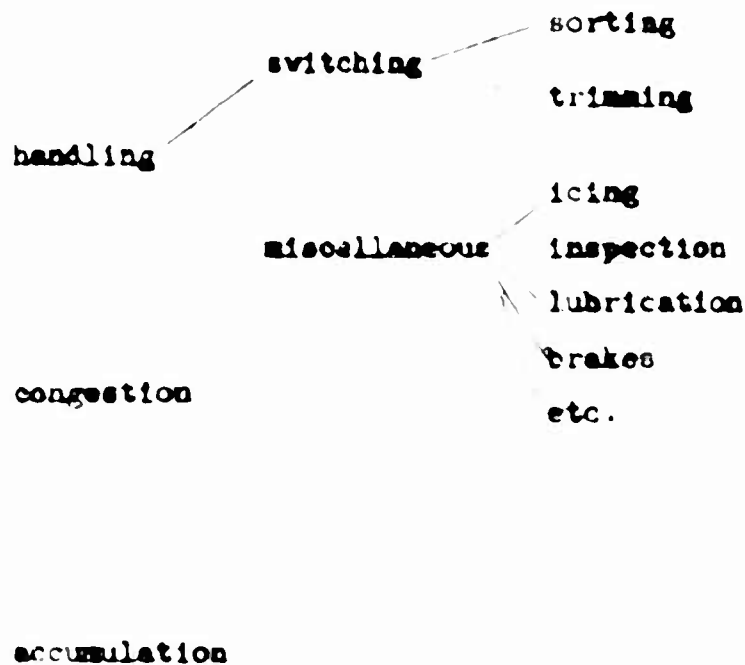brakes
etc.

congestion

accumulation

Figure 1: Sources of Delay

First is accumulation delay about which something has already been said. For
the traffic to be dispatched on any one train it is found in the following way:
Look at the times of arrival in the receiving yard for all the cars going out

on this train. For each of these cars, find the difference between its
arrival time and the arrival time of the very last such car to come in.
Take the average. The resulting average accumulation delay figure repre-
sents the average time spent by a car in waiting for the rest of its train
to come in. These delays will be large if inputs to the yard are dispersed
over time but outputs bunched. They will not be affected by changes in
grouping alone, since the composition of outbound trains remains the same.
They will however, be very seriously affected by changes in train make-up.

The next type of delay is switching delay, consisting of sorting and
trimming delay. The word sorting is used here to cover humping and the
corresponding operation in a flat yard -- the sorting of cars from incoming
trains into the classification tracks. The time consumed per car in sorting
probably depends for the most part on the average number of cars per cut.
This last item is particularly important in a flat yard, and probably not
so important in a hump yard. In fact, doesn't this difference in sorting
abilities of the two kinds of yards constitute one of the main operational
distinctions between them? If the first policy variable, grouping, is changed
so that the number of classifications is increased, that is, the grouping is
made finer, then the number of cars per cut will decrease and the time consumed
in the sorting operation will increase, a little bit in a hump yard and quite
a lot in a flat yard. On the other hand, if incoming traffic has already been
grouped to some degree by earlier yards, then the number of cars per cut will
be greater and the sorting operation will require a little less time in a hump
yard and a lot less time in a flat yard. Analysis of sorting delay and the
way it enters into the problem of selecting the best grouping policy for a
yard requires two things. First, time studies must be made to determine how
the sorting time per car depends on cars per cut. Such studies would reveal

how this relation varied from yard to yard, and most interestingly, from flat yard to hump yard. The second thing needed is a statistical method of estimating how cars per cut varies when policy changes, especially in grouping, are made. Some progress has been made along these lines, what has to be done now is to compare some of these numbers from statistical formulas with numbers from actual inbound train consists to see if the formulas are realistic.

Trimming delay (per car) is the time taken to pull the classification tracks and make up the outbound train, divided by the number of cars in the train. If the number of switch engines is fixed, this probably depends for the most part on the particular make-up policy decided on. The more groups there are in a train, the greater the number of classification tracks that will have to be pulled, and the greater the time and work involved in assembling the train.

The miscellaneous category is meant to contain all those other sorts of delay that do not depend on the degree of prior classification at different parts of the system. The average time per car for inspection and lubrication, for instance, is probably the same whatever the order of incoming cars. This means that for the problem we are here concerned with, the selection of our policy variables -- grouping, make-up, and scheduling, these delays are not so important. They affect our choice through their influence on congestion delays (see below), but they are not themselves affected by our choice. The same cannot be said of the other delays we have discussed.

Congestion delay is defined here as the average time spent by a car in the receiving yard, not counting those parts of this receiving yard time devoted to inspection and other miscellaneous items. If handling time is very regular and arrivals of trains are very regular, then congestion delays will

be all. They increase with the variability of each. They are affected
by grouping and make-up to the extent that these policy choices affect
the variability of handling delays and arrival times. Since Mr. Crane's
paper deals specifically with congestion delays, nothing more will be said
about them at this point.

Money cost per car can be pretty well categorized in the same way as
delays were. The miscellaneous costs corresponding to lubrication, brakes,
etc., are probably independent of the degree of grouping and make-up. Switch-
engine cost and wages for sorting and trimming probably change in the same
directions as sorting and trimming delays when traffic characteristics and
the policy variables are altered. They money cost side of the picture is a
little simpler than the delay side in that there seem to be no money costs
corresponding to congestion and accumulation delays.

The important question concerning money cost is the following one.
To what extent and at what places can cost be substituted for delay? That
is, can delay be reduced by a more intensive application of switch engines
and labor to the work at hand? One suspects that at some places this is
possible to a large extent, and very little at others. The time it takes
to hump a given train is probably not subject to reduction of this type.
The sorting operation in a flat yard however, can sometimes be speeded up
by cutting a train in two, using two engines instead of one, and feeding
cars into the classification tracks from both ends. Mr. Laku mentioned a
few minutes ago that an extra crew seldom helps in yard work unless the
yard is designed with an eye toward this additional labor.

After this hasty examination of the various costs involved in classi-
fication, let's look at a simple problem and try to see how a proposed
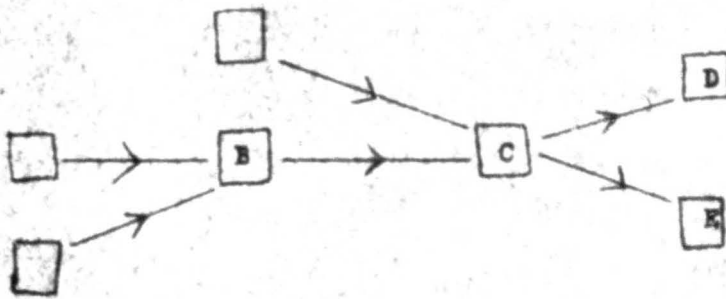prior classification alters costs at different parts of the system.

Figure 2.

Each of the squares in the diagram of Figure 2 is a yard of some kind. Because the tracks leading away from C diverge, it is clear that C must classify D and E. Let's suppose to start with however, that no prior classification is done. What would happen if the grouping at B were changed so that now B also classified D? With this new grouping policy cars per cut at B would of course increase, and the time and cost of B's sorting would go up correspondingly. Trimming costs would also increase because of the extra classification track to be pulled. Costs at C however, would decrease for the number of cars per cut has gone up with B's prior classification work. Now if make-up and scheduling remain unchanged at the two yards, so that accumulation delay is unaffected, and if neither yard is so hard pressed that congestion delays are affected by the change, then the answer we seek may be found in a simple comparison of the cost increase at B with the cost decrease at C. Some problems are undoubtedly just as simple as this. If, for instance, B is a hump yard and C a flat yard, then the cost increase at B is likely to be small, and the cost decrease at C large; comparison then points to prior classification.

Upon closer examination of such simple problems however, many difficulties show up. Delay comparisons for instance only make sense if some scheduling

changes are allowed. Prior classification at B may require that we delay a train's departure from B to enable the extra work to be done. If this time is more than made up (as it will be if prior classification is profitable, time-wise) by the reduction in sorting time at C, then the departure from C may take place earlier. But if we allow scheduling to be changed, then we can no longer ignore changes in accumulation delays. A complete analysis must handle both grouping and scheduling problems simultaneously.

Such considerations lead us straight into the mathematics of the problem. This mathematics is admittedly difficult. Provided, however, our assumptions are sufficiently realistic and comprehensive, these difficulties only correspond to the difficulties of the problem which the operating department of a railroad faces much of the time. The great advantage of an analytic approach is that it can help us to see more clearly the relationships of the factors involved, and to judge more systematically the consequences of various possible operating decisions and procedures.