

Optimal boarding method for airline passengers

Jason H. Steffen ¹

*MS 127, PO Box 500
Batavia, IL 60510*

Abstract

Using a Markov Chain Monte Carlo optimization algorithm and a computer simulation, I find the passenger ordering which minimizes the time required to board the passengers onto an airplane. The model that I employ assumes that the time that a passenger requires to load his or her luggage is the dominant contribution to the time needed to completely fill the aircraft. The optimal boarding strategy may reduce the time required to board an airplane by over a factor of four and possibly more depending upon the dimensions of the aircraft. I explore some features of the optimal boarding method and discuss practical modifications to the optimal. Finally, I mention some of the benefits that could come from implementing an improved passenger boarding scheme.

1 Introduction

Several passenger boarding schemes are used by the airline industry in effort to quickly load passengers and their luggage onto the airplane. Since the passenger boarding time often takes longer than refueling and restocking the airplane its reduction could constitute a significant savings to a particular carrier, especially for airplanes which make several trips in a day.

Conventional wisdom would suggest that boarding from the front to the back is the worst case but that boarding from the back to the front is optimal or nearly so. Indeed, this is the strategy that is often employed, boarding passengers in blocks from the rear of the plane to the front. In this case, conventional wisdom only provides an answer that is half right. The worst

Email address: `jsteffen_AT_fnal.gov` (Jason H. Steffen).

¹ Brinson Postdoctoral Fellow, Fermilab Center for Particle Astrophysics

boarding method is, indeed, to board the plane from front to back. As I will show, however, boarding the airplane from the back to the front is very likely the *second worst* method.

Generally, the result of boarding passengers from the back to the front simply shifts the line of passengers into the airplane but does not, thereafter, allow the airplane to load significantly faster. Other boarding schemes, such as boarding the window seats, then the middle seats, then the aisle seats, can reduce the board time by a significant fraction—better than half of the worst case. The reason that these schemes work better is subtle and therefore not obvious; though their improvement over the worst case is significant. The question remains however, “what method gives the fastest boarding time?”

The answer to this question depends upon the layout of the airplane, upon human nature in general, and upon the particular set of passengers on a given flight. The first of these items can be modeled exactly, or nearly so. The second can be approximated with some assumptions and calibrated with data. The third item lies outside any real modeling, but statistical methods can be used to predict how often a particular set of passengers might significantly affect the board time.

In this work I infer from the results of a computer simulation and an associated optimization algorithm the optimal passenger boarding method. The fundamental assumption that I make is that the bulk of the time to load the airplane is consumed by time that it takes the passengers to load their luggage. All other effects, such as the time required to climb over the person sitting in the aisle seat or the time used by passengers standing up to retrieve an item, are treated as negligible. This assumption is based upon my personal experience as well as the experiences of several interviewees. A more sophisticated model could include such detail. Or, many of their effects could be folded into the distribution of the luggage loading times of the passengers. However, I do not attempt to include them in this study.

With the stated assumption, I find that the boarding time for the optimal scheme can be significantly faster than the boarding time of the worst case—between a factor of 4 and 10 faster depending upon the length of the airplane and other model parameters. In this article I will describe the techniques that I use to find the optimal method, I interpret the results and use that interpretation to discuss the merits of some schemes that are employed by the industry, finally I give some concluding remarks and suggestions for refinements of this model. Note that I generally use the term “boarding” to refer to the boarding process itself and the term “loading” to refer to the passengers loading their luggage. Thus, boarding time and loading time are the times required to fill the aircraft and the time required to load one’s luggage respectively.

2 Analysis Approach

2.1 Airplane and Passenger Models

The nominal airplane model that I use seats 120 passengers with six passengers per row and 20 rows. My focus is on the general boarding procedures. Consequently, I do not include a first-class cabin nor do I have any priority seating; and each flight is completely full. I will discuss the effects of changes to this airplane model in section 3.

The passengers are each assigned a seat number and a time to load their luggage, a random number between 0 and 100 time steps unless otherwise stated. Other, human nature assumptions include: 1) that a person will not move unless there is enough space between them and the person in front of them—two steps in this case, 2) that they will occupy any empty space in front of them and then stop (that is, they will bunch up again as they wait), 3) that they require one space either in front of or behind them in order to load their luggage, and 4) that they only load their luggage into the bins above their assigned row. In section 3 I discuss the effects of changing any of these parameters or the distribution from which the luggage loading times are assigned.

2.2 Optimization Algorithm

The algorithm that I use to find the optimal loading order is based upon a Markov Chain Monte Carlo (MCMC) algorithm and is similar to the METROPOLIS algorithm (Metropolis et al., 1953). Starting with an initial passenger order I load the airplane and record the loading time. Then, I take that initial order, switch the positions of two random passengers, and load the airplane again. If the airplane loads as fast or faster than the previous iteration, then I accept the current passenger order, swap the positions of two additional random passengers, and repeat the process. If the current configuration loads more slowly than the previous, then I reject the change, return to the original, and repeat the process beginning there. I stop after $\sim 10,000$ iterations since adding additional steps does not significantly change the results.

Unlike a traditional MCMC, I do not allow any configuration which loads more slowly to be accepted. When I include this aspect the primary effect is to increase the convergence time while the results remain essentially unchanged. Moreover, several trials of the above algorithm produce indistinguishable results. Thus, the choice to neglect that aspect of an MCMC analysis should not affect the results stated here.

That being said, it is possible for many configurations to load in the same time or to be near enough that the differences in loading time are not important. This shows that a class of configurations that are effectively equivalent is more important than a single, optimal order. For example, there is no difference in the loading time between switching the two aisle-seat passengers in the same row, and there is little difference in switching two random passengers. To identify the class of optimal configurations I tabulate the differences in seat number between adjacent passengers. It is this distribution of seat number differences that remains effectively constant from one optimization run to the next. I call this distribution the “seating distribution”. This representation illustrates the fact that the important aspect of passenger loading is not so much where two adjacent passengers sit on the airplane, but rather how far apart they sit from each other.

Note that I use a single coordinate for the seating such that back row of seats are numbered 1 through 6, the second row from the back are seats 7 through 12, and so forth. Also, for most experiments I collect data from 100 realizations of luggage loading times. That is, I re-assign the loading time for each passenger 100 times and rerun the optimization algorithm for each new assignment. I’ll note exceptions to this when necessary.

3 Results

The results of applying the above analysis gives the seating distribution shown in Figure 1. The largest feature is the peak near a seat difference of 12. That difference corresponds to two rows, or the distance that I assume neighboring passengers require to load their luggage. Other features of the peak, aside from its location, are its shape, its height relative to the rest of the distribution and its width. All four of these aspects depend upon the passenger and airplane model parameters and each of them could be calibrated with data.

As mentioned, the location of the peak corresponds to the distance required by a passenger to load his luggage. If I allow passengers from adjacent rows to load their luggage simultaneously, then the location of the peak would shift to a value of 6 or one row of difference. If passengers require 2 rows of loading space, then the peak shifts to 18. This effect can be seen in figure 2 where I make these changes while leaving all other model parameters fixed.

The peak of the seating distribution is symmetrical near its apex. This is because a passenger can load his luggage using either the space in front of them or the space behind them. Thu, two people can simultaneously load when their seats are in adjacent rows provided that there is space on either side—allowing the peak its symmetry. If I require the passengers to have a

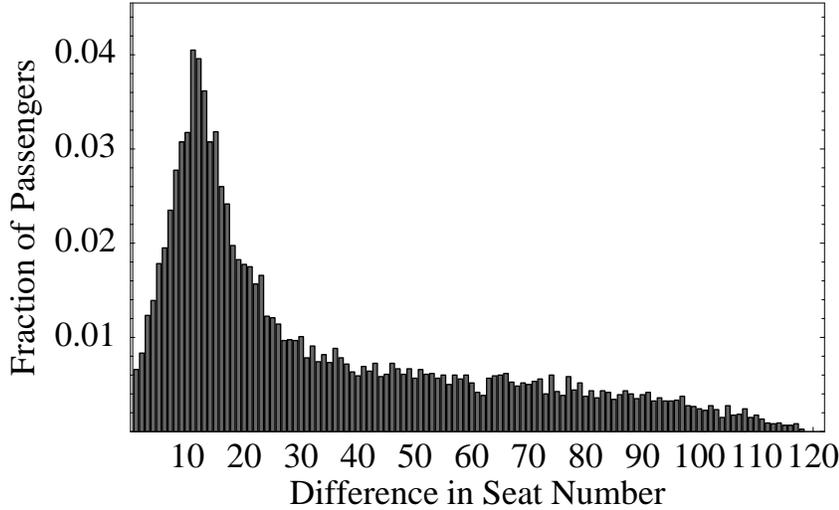


Fig. 1. Example of the resulting seating distribution obtained from 100 realizations of the luggage loading time distribution.

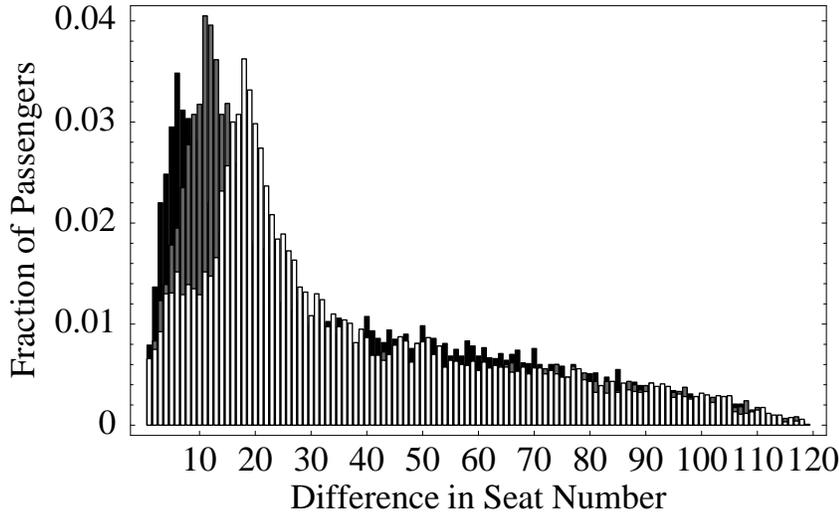


Fig. 2. Changes in the seating distribution as a function of the required “personal space” of the passengers. This shows the distribution if no space is required (black), a single row is required (gray), and two rows are required (white). Each distribution is calculated from 100 realizations of the luggage loading times.

space in front of them only, then there is a larger penalty for being next to someone who is assigned a seat in the row in front of you. Consequently, the peak skews toward larger distances; being more shallow before a separation of 12 and more broad past that separation. Figure 3 compares the peak shape when passengers require a space in front of their assigned row only or if they require either a space in front or a space behind their assigned row.

The width of the peak is related to the number of seats per row. If there are only four seats per row, then the width of the peak is more narrow. If there are eight seats per row, then it is more broad. The position of the peak also

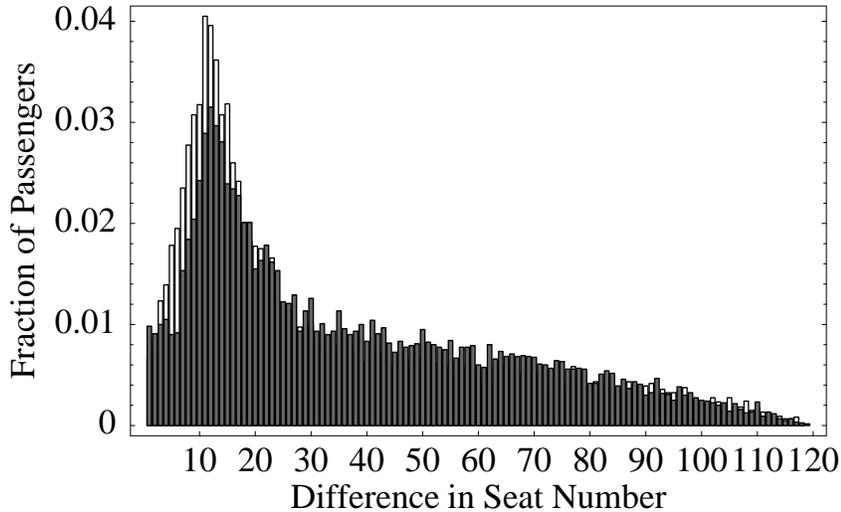


Fig. 3. Changes in the shape of the peak in the seating distribution that arises if a passenger is able to load their luggage with the required space either in front or behind them (white) or if the required space must be in front of them (black). This explains why the peak is symmetric about the two-row separation in the fiducial model.

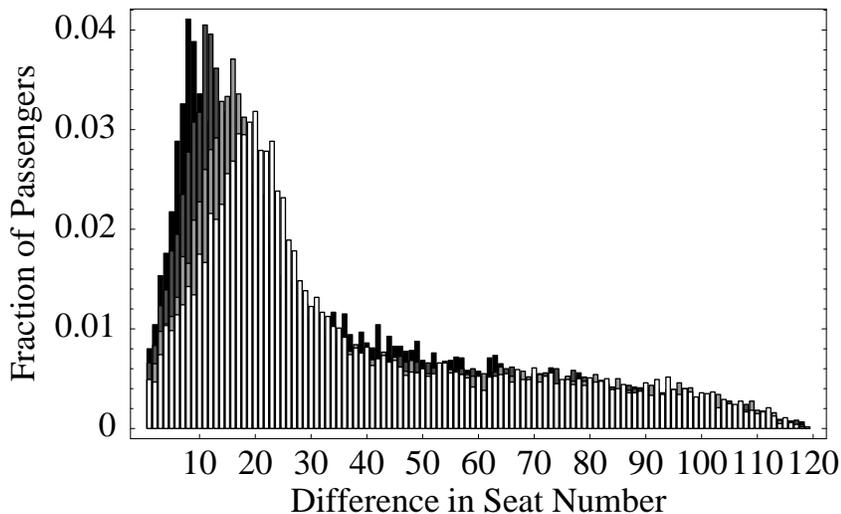


Fig. 4. The changes in the seating distribution that result from changes in the width of the airplane. These peaks correspond to an airplane width of 4 (black), 6 (dark gray), 8 (light gray), and 10 (white) seats. I always assume that there is one aisle.

shifts such that it is located at a separation of two rows. This effect is shown in figure 4 which displays the distribution that arises on airplanes with 4, 6, 8, and 10 seats per row (always with one aisle).

The height of the peak depends upon the time that the passengers take to load their luggage. If passengers load their luggage instantaneously, then the peak would disappear altogether. As the luggage loading time increases the penalty for having someone out of order increases and the algorithm forces

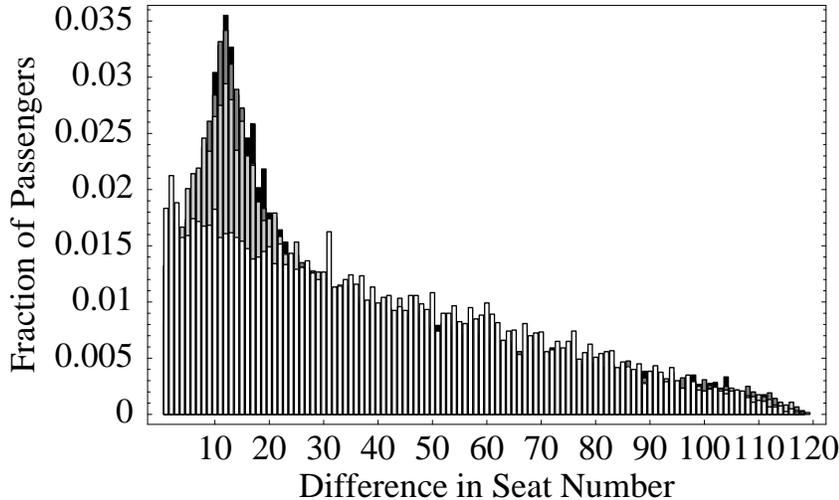


Fig. 5. Changes to the seating distribution that arise from changing the mean time to load one’s luggage. The time to walk the length of the airplane is 20 counts in this scenario. The mean luggage loading times that correspond to these distributions are 0 (white), 6 (light gray), 12 (dark gray), and 18 (black) time steps. A distribution with a loading time of 24 steps is virtually indistinguishable from that with the a loading time of 18 and is not shown.

more passengers to be separated by amounts nearer the minimum row separation required to load their luggage (here between 7 and 18 seats). The height of the peak ultimately saturates when the mean luggage loading time approaches the time to walk the length of the airplane. As the loading time becomes longer than this time, the height of the peak remains constant. If the airplane is longer, then the peak saturates at longer loading times. This effect can be seen in figure 5 where I show the distributions that arise from several values of the mean luggage loading time.

3.1 *The Optimal Loading Method*

The optimal boarding scheme is found by extrapolating from the results given above, using the insight that they provide. The reason that these distributions load faster than the worst case is that they allow multiple passengers to load their luggage at the same time. The peak occurs at a distance that corresponds exactly to the space needed by adjacent passengers to do so. Taking this to an extreme, we wish to find the configuration that allows the maximum number of passengers to load their luggage at all times. That is the case where all adjacent passengers are assigned seats that are separated by exactly two rows.

For the fiducial airplane model stated above there can be a maximum of ten passengers loading their luggage at once. After they finish, the next ten passengers are sent in, again to sit in every other row, and so forth. This

ordering scheme provides nearly a five-fold reduction in the time that it takes over the worst case (for the fiducial model). This also shows why simply loading from the back of the plane to the front does not provide any benefit. While there may be many passengers inside the aircraft when you load from back to front, very few will actually be loading their luggage. Ideally, you want all of the passengers that are inside the aircraft either to be seated or to be loading their luggage with none of them waiting.

One question that arises from this is whether or not it is practical to implement the optimal boarding scheme, where each passenger enters the airplane in a particular order. Such a scenario may well be possible since Southwest Air has recently implemented a similar policy, at least to some extent. Given that, however, there will always be some fraction of the passengers who are out of order; there will always be families or other groups who board together regardless of their assignments. These scenarios lie outside the scope of this model, though they could be incorporated in a number of ways; for example, by using a multimodal distribution for the loading times. These same issues will affect other boarding methods in similar ways. In the next section I test the robustness of the optimal boarding method. The section following is a comparison of a few practical boarding methods where the passengers board in groups but are ordered randomly within those groups.

Finally, I note and explain why the MCMC algorithm did not produce the optimal ordering scheme on its own. Search algorithms that are based upon random changes are not likely to produce a highly ordered result—such as having each passenger exactly 12 rows from the two adjacent passengers. This would be, in essence, a violation of the second law of thermodynamics which states that a randomly evolving system of objects will tend toward a disordered state instead of an ordered one. Since the total number of seating configurations is $100!$ and the class of optimal loading schemes has, at most, of order $12! \times 10!$ (corresponding to the number of seats per row and the number of starting rows that are separated by two), the probability of a random search, with a modern computer, stumbling across one of these solutions within the age of the universe is less than 10^{-100} . Thus, there is a need to make the intuitive leap to arrive at the optimal solution.

4 Robustness of the Optimal

To test the robustness of the optimal boarding scheme I conducted two experiments. The first is to change the distribution from which I select the passenger's loading time. The second test is to make random changes to the passenger ordering. These changes include swapping the locations of several random pairs of passengers and shifting the entire line by some random num-

ber (moving people at the end of the line to the front).

4.1 Changes to Loading Time Distribution

To test the effect of a different distribution of luggage loading time, I ran my minimization software on 100 realizations of each of several distributions. These distributions include a uniform distribution with a given mean, a normal distribution with the same mean and with a variance equal to that mean (essentially a Poisson distribution), and an exponential distribution with the same mean. For each of these cases the resulting seat distribution were statistically indistinguishable as shown by a Kolmogorov-Smirnov test (Press et al., 2002). Moreover, the time required to load the entire plane is not affected by these different distributions; it depends primarily upon the mean luggage loading time.

4.2 Random Shifts and Swaps

The effect of randomly shifting the line does not affect the boarding time in any significant way. This is because it only changes the starting point of the boarding process. All of the passengers keep their 12 seat spacing from their neighbors and so the advantage of the optimal boarding scheme is preserved.

If pairs of passengers are swapped, which effectively randomizes portions of the line, then the time to board the airplane can change significantly. Indeed, a 20% increase in the boarding time results from randomly swapping only 10% of the passengers—that is 6 pairs for the case of 120 passengers. However, there is an upper limit to the effect of swapping passengers since once they are completely randomized additional swaps maintain the random nature of the passenger order and the ordering doesn't get any worse. To completely randomize the optimal ordering one must swap at least 60 pairs of passengers. It is unlikely, even in the worst of cases, that so many of the passengers would be out of order. Interestingly, random boarding takes much less than half the time of the worst case boarding; indicating that randomization is not catastrophic. Indeed, we will see in the next section that random boarding compares favorably with traditional boarding techniques.

The results of the tests outlined in this section indicate that the optimal boarding method is not largely affected by small changes to the passengers or their order. This serves to validate the use of the optimal method since there is not a significant penalty in boarding time when small deviations from the optimal are present. While there is a steep initial loss of time due to passengers being out of order, the effect is bounded by the complete randomization of the

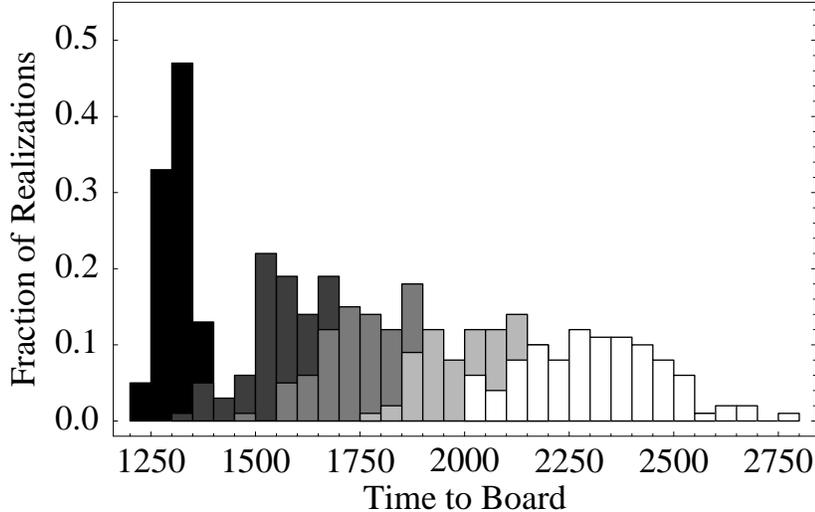


Fig. 6. Boarding times for 100 realizations of passengers when pairs of passengers are swapped. The black is optimal with no swapping, then there are the histograms which correspond to 10% (dark gray), 20% (medium gray), 40% (light gray), and 60% (white) swapping. 10% means that 6 pairs of passengers are exchanged out of 120 passengers. The mean boarding times for these scenarios are 1312, 1585, 1795, 2084, and 2311 counts respectively.

passengers. Since that random boarding is still a marked improvement over the worst case, it is not grounds for dismissing the approach.

5 Practical Comparison

While the optimal scheme would produce the fastest boarding times, there are issues of practicality to consider. It may not be possible to arrange all of the passengers in the proper order or near enough to that order to capitalize on the available gains. Though, as mentioned, Southwest Airlines has implemented such a policy—the only difference for this model is that the boarding order is chosen based upon seat assignment. Regardless, most airlines board the airplane in groups, presumably out of convenience and in effort to reduce confusion. In this section I introduce a few practical modifications to the optimal boarding scheme and choose one to compare with existing methods.

5.1 Practical Modification to the Optimal

The advantage that the optimal comes from the fact that neighboring passengers do not sit near each other and consequently can load their luggage simultaneously. One way to accomplish a similar effect while allowing blocks

of passengers to board is to have each block contain passengers from widely separated rows. For example, passengers from every fifth row; which gives five boarding groups and provides significant space between each of the allowed rows, thus reducing the number of passenger collisions.

After trying several possibilities, I instead settled on having blocks of three consecutive seats separated by 12. This scenario has four boarding groups and is equivalent to calling all passengers that are from one side of the airplane and that sit in every other row. The three remaining groups are for the other side of the airplane in the same row, then the two sides of the next row. This approach gave the overall best loading times of the variations that I tried. The loading time that results from this scheme is not as fast as the optimal, indeed it took about twice as long to load, but it was more than a factor of two faster than the worst case. I call this boarding method the “modified optimal” method.

5.2 Comparison with Other Methods

Initially, I selected three different group boarding strategies to compare with the optimal, the modified optimal, the worst case, and the second worst case (boarding from the back to the front with the passengers in order). These are: 1) ordered blocks, where a fourth of the cabin is loaded at a time starting in the back and moving to the front—this is the most common loading strategy that I’ve experienced, 2) unordered blocks where I load the back fourth of the airplane, then the second fourth, then the third, then the first, and 3) a scheme where I load the windows, middle, then aisle seats. Within each of these groups the travelers were randomly distributed. I ultimately dropped the unordered block scheme since it gave similar results to the ordered blocks.

The ordered block scenario reduced the boarding time to 74% of the worst case. Since the main contributor to fast boarding is having multiple passengers load their luggage at once, this method suffers from the fact that only passengers within a small portion of the airplane are boarding at a given time. Within the boarding block there is the possibility of multiple people loading their luggage at once, but it is relatively small since the block is only 5 rows deep and passengers require more than one row of space to load their luggage. This would improve with a larger airplane, but is still limited by this effect.

The windows-middle-aisle approach reduced the boarding time to 43% of the worst case. The advantage here is that passengers from anywhere in the cabin are allowed to board. Thus, many people can load their luggage at once and the probability having two passengers from adjacent rows near each other is small. Moreover, this probability decreases as the length of the airplane increases.

Another advantage is that the second-order effect, one that I ignore in my model, of getting past the person in the aisle is eliminated. One drawback of this approach is the fact that many people travel in groups and would likely not adhere to this particular boarding policy. The result would be bottlenecks that occur during the boarding process, most likely towards the end when people from all boarding groups are allowed onto the airplane. In general, however, this approach did very well.

The modified optimal approach performed slightly better than, though almost identically to, the windows-middle-aisle approach. This method does not depend as strongly on the length of the airplane. So, for shorter airplanes it compares better while on longer planes it compares less well to the windows-middle-aisle method. The difference in boarding times between these two schemes is likely to be less than the accuracy of my model (less than 5%) and should be considered accordingly. An advantage that the modified optimal approach has to boarding window-middle-aisle is that it allows passengers who are travelling together and sitting side-by-side to board at the same time without boarding out of order—though the effects are not modeled here.

Random boarding, where the passengers positions in line are completely uncorrelated with their seat assignment, has the same advantage that the windows-middle-aisle method has of spreading passengers throughout the length of the airplane. Moreover, it is not disadvantaged, in implementation at least, by traveling groups. The random method performed the similar to, but slightly worst than, the windows-middle-aisle method. This demonstrates that the optimal approach, even with a significant fraction of people out of order can do at least as well as or better than the common methods that employ boarding groups. That is, there is effectively no penalty in attempting to implement the optimal approach since the worst that it will do is to perform as well as the modified optimal or the windows-middle-aisle methods; the best block-loading methods that I studied.

Figure 7 shows a histogram of the loading times for 100 realizations of seven different boarding schemes, each realization having different selections for the time it takes to load the luggage and different passenger ordering where applicable. These include: 1) optimal boarding, 2) the modified optimal approach, 3) the window-middle-aisle method, 4) random boarding, 5) ordered blocks from back to front, 6) back to front with all passengers in order, and 7) the worst-case (front to back with all passengers in order). We see from this figure that optimal boarding has, by far, the best improvement in loading times, nearly a factor of five faster than the worst case, more than a factor of three better than the ordered blocks, and it is more than a factor of two faster than the modified optimal, random, and windows-middle-aisle methods. This improvement grows with the length of the airplane such that an airplane that seats 240 passengers (40 rows) will board over seven times faster than the

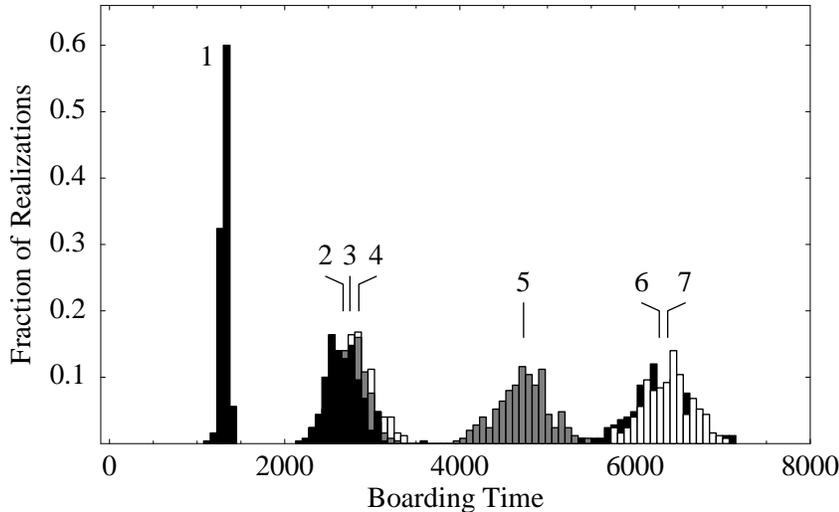


Fig. 7. Histogram of the loading times for 100 realizations of seven different boarding schemes. The luggage loading time for each realization is drawn from a uniform distribution with a mean of 50 counts. These are: 1) optimal boarding (mean boarding time: 1312 counts), 2) the modified optimal approach (2670), 3) the window-middle-aisle method (2750), 4) random boarding (2846), 5) ordered blocks from back to front (4727), 6) back to front with all passengers in order (6276), and 7) the worst-case—front to back with all passengers in order (6373).

worst case!

6 Conclusion

The results of this study, based upon the assumption that a passenger loading his luggage consumes the bulk of the time that it takes for him to be seated, identify the primary cause for delay in the boarding process as well as the best means to overcome these delays. By boarding passengers in a manner that allows several passengers to load their luggage simultaneously the boarding time can be dramatically reduced. This result contradicts conventional wisdom and practice that loads passengers from the back of the airplane to the front. Indeed, it shows that loading from the back to the front is little different from the worst case of loading from the front to the back. The goal of an optimized boarding strategy should focus on spreading the passengers throughout the length of the airplane instead of concentrating them in a particular portion of the cabin.

By boarding in groups where passengers whose seats are separated by a particular number of rows, by boarding from the windows to the aisle, or by allowing passengers to board in random order one can reduce the time to board by better than half of the worst case and by a significant amount over conventional

back-to-front blocks—which, while better than the worst case performed worse than all other block loading schemes. The primary drawbacks for any of these methods is likely to be psychological instead of practical. Groups of passengers who wish to board together would be an issue to investigate from both a customer satisfaction point of view and as a component in a more detailed model.

If a workable method to have passengers line up in an assigned order could be found—and it likely may be employed already, then there is the potential for a substantial savings in time. Such a savings would most likely benefit flights between nearby cities where a particular airplane would make several trips in a given day since it might allow one or two additional flights. Or, it might allow an airline to reduce the number of gates that it requires to meet its obligations since each gate would be cleared more rapidly.

While the generic features of this model are well understood, a real application of it would require some data so that it can be properly calibrated. In particular, the distributions of luggage loading times, the fraction of people traveling in groups, the queueing habits of passengers, and empirical measurements of “personal space” are all pieces of information that are necessary to state these results in terms of actual times and distances instead of arbitrary time steps and lengths.

Regardless of the ultimate application of this technique, the establishment of a firm lower bound can be used to inform a decision maker of the worth of further improvements to a particular boarding strategy. If an improvement could provide only a marginal gain while costing significant amounts money and time to implement then it is not likely to be worth the investment. On the other hand, if a particular strategy is clearly failing to meet the demands of competition and customer satisfaction, then knowing just how much room there is for improvement could expedite changes.

Clearly the model described here can be improved. Including the effects of aisle vs. window seats, the clustering of passengers into companions or families, and other effects of human nature would improve the accuracy of the results. But these effects are not likely to be the primary issue and consequently should not be the fundamental concern when finding the general strategy for a passenger boarding scheme. On the other hand, they should not be neglected outright and indeed should be what drives the refining process needed to settle on a workable solution. In the end, the time that it takes to load passengers into the airplane affect not only the airline company and the airport, it also affects the passengers. Few enjoy standing in line longer than necessary and fewer still enjoy sitting in an airplane longer than needed. Faster boarding would be a significant improvement for all involved parties.

Acknowledgements: J. Steffen acknowledges the generous support of the Brinson Foundation. Fermilab is supported by the U.S. Department of Energy under contract No. DE-AC02-07CH11359. A record of invention of this software and its results are on file at Fermilab.

References

- N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller.,
Journal of Chemical Physics, 21(6):1087-1092, (1953).
W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, “Numerical
Recipes in C”, Cambridge University Press, USA, (2002).