



Contents lists available at ScienceDirect

Journal of Air Transport Management

journal homepage: www.elsevier.com/locate/jairtraman

Optimal boarding method for airline passengers

Jason H. Steffen

Fermilab Center for Particle Astrophysics, MS 127, PO Box 500, Batavia, IL 60510, USA

ARTICLE INFO

Keywords:

Passenger boarding
Optimization
Operations

ABSTRACT

Using a Markov Chain Monte Carlo optimization algorithm and a computer simulation, the passenger ordering that minimizes the time required to board an airplane is found. The model assumes that the time a passenger requires to load luggage makes 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 a factor of four or more depending on the dimensions of the aircraft. Some of the salient features of the optimal boarding method are discussed along with practical modifications to the optimum.

© 2008 Elsevier Ltd. All rights reserved.

1. Introduction

Several boarding schemes are used by airlines in an effort to seat passengers quickly. Since the overall passenger boarding time often takes longer than refueling and restocking its reduction could constitute a significant cost 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 boarding method is, indeed, to board the plane from front to back. However, boarding the airplane from the back to the front is very likely the second-worst method. Previous studies of airplane boarding use a variety of approaches. Pioneering work by Van Landeghem and Beuselinck (2002) used computer simulations of several boarding schemes and found there is room for improvement over traditional back-to-front boarding. In particular, they point out that random boarding is superior to many traditional methods. A later computer simulation study (van den Briel et al., 2005) confirmed that traditional methods are not optimal, even in light of different assumptions regarding the primary cause of delay in the boarding process. Work by Bachmat et al. (2006), that models optimal airplane boarding as an extremal path in a two-dimensional, Lorentzian geometry, broadly confirmed the findings of Van Landeghem and Beuselinck. They were able to interpret their findings to provide an explanation for why the different boarding strategies perform as they do.

Each of these studies, particularly the computational works, focus on different components of delay. Consequently, while all found traditional methods lacking, none agree on the best overall approach. A portion of their disagreement arises from the differences in assumptions. Another point of discrepancy stems from the fact that their focus was largely on finding the best practical approach or the best of a set of approaches. This limitation on the scope of previous works meant that they did not always study the same boarding methods.

Thus, the question remains as to what method, practical or otherwise, gives the fastest boarding time. The answer, whether directly applicable or not, is valuable as it can inform decisions regarding the value of changing an existing policy by indicating how much room there is for improvement. Moreover, once the optimal boarding method is identified one can explore the reasons why it is successful and incorporate these characteristics into a more practical scheme.

Here inference is made from the results of a computer simulation and an associated algorithm that optimizes passenger boarding. The fundamental assumption is that the bulk of the time to load the airplane is consumed by the time it takes passengers to load their luggage. All other effects, such as the time used by passengers standing up to retrieve an item, who sit in the wrong seat, or who must pass by someone that is already seated in their row are initially treated as negligible. A more sophisticated model could include such detail; or, many of their effects could be folded into the model via some of the available parameters or components (such as the distribution from which the passenger luggage loading times are selected). With these assumptions, 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.

E-mail address: jsteffen@fnal.gov

2. The approach

2.1. Airplane and passenger models

The nominal airplane model takes a 120 seat passenger aircraft with six passengers per row in 20 rows. Since the focus is on the general boarding procedures, there is no first-class cabin, no priority seating, and each flight is completely full. The passengers are each assigned a seat and the number of time steps that they need to load their luggage, a random number between 0 and 100 unless otherwise stated. Other, human nature assumptions include: 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), that if they are moving, then they will occupy any empty space in front of them prior to stopping (this results in passengers bunching-up as they come to a halt), that they require one space either in front of or behind them to load their luggage (this is related to the congestion parameter k introduced by Bachmat et al.), and that they only load their luggage into the bins above their assigned row.

This model does not include the effects of aisle versus window seats, the clustering of passengers into companions or families, and other effects of human nature. While adding these features might improve the accuracy of the results, they are not likely to be the primary issue and consequently should not be of fundamental concern when finding the general strategy for a passenger-boarding scheme. Moreover, many of their effects can be accounted for once the optimal-boarding method that is based upon the stated assumptions is identified.

2.2. Optimization algorithm

The algorithm used 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 the airplane is filled and the boarding time recorded. Then, starting with that initial order, the positions of two random passengers are exchanged and the airplane is loaded again. If the airplane boards as fast or faster than the previous iteration, then one accepts the current passenger order, swaps the positions of two additional random passengers, and repeats the process. If the current configuration loads more slowly than the previous one, then the change is rejected, and one returns to the previous configuration, and repeats the process beginning there. After about 10,000 iterations the procedure is stopped because adding additional steps does not significantly change the results.

Unlike a traditional MCMC, I do not allow any configuration that loads more slowly to be accepted, meaning that the stated algorithm only finds a local minimum. However, when this aspect is included the results are unchanged while the time needed to converge increases. Thus, the choice to neglect that aspect of an MCMC analysis should not affect the results.

It is possible for many configurations to board in the same time or to be near enough that the differences in boarding time are not important. This shows that a class of configurations which are effectively equivalent is more important than a single, optimum order. For example, there is no difference in the boarding time if the two aisle-seat passengers in the same row are swapped, and there is little difference in swapping two random passengers. To identify the class of optimal configurations the differences in seat number between adjacent passengers is tabulated. It is this distribution of seat number differences, the seating distribution, that effectively remains constant from one optimization run to the next. This illustrates the important point that the seat numbers of

two adjacent passengers is less important than how far apart they sit from each other.

3. Results

Applying the stated analysis gives the seating distribution shown in Fig. 1. These results are from 100 realizations of the boarding scenario where the passengers are reassigned their luggage loading time for each realization. The main feature is the peak near a seat difference of 12 that corresponds to two rows, or the distance that neighboring passengers are assumed to require to load their luggage. Other features of the peak, aside from its location, are its symmetrical shape, its height relative to the rest of the distribution, and its width. All of these depend on the passenger and airplane model parameters that should ultimately be calibrated with data.

The location of the peak corresponds to the distance required by a passenger to load his luggage. If passengers need only the space that corresponds to their assigned row, then the location of the peak would shift to a value of 6 or one row of difference. If passengers require three rows of loading space (including their assigned row), then the peak shifts to 18. This effect can be seen in Fig. 2, where these changes are made while leaving all other parameters fixed. The peak of the seating distribution is symmetrical near its apex, because passengers can load their luggage using either the space in front of them or the space behind them. 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 narrower. If there are eight seats per row (still with only one aisle), then it is broader.

The height of the peak depends upon the time that passengers take to load their luggage. If passengers load their luggage instantaneously, then the peak disappears altogether. As the luggage loading time increases the penalty for having someone out of order increases and the algorithm forces more passengers to be separated by amounts nearer the minimum row separation required to load their luggage (between 7 and 18 seats). The height of the peak ultimately saturates when the average luggage loading time approaches the time required to walk the length of the airplane. This means that if it takes significantly longer for a typical passenger to walk the length of the airplane than it does for him to load his luggage, then the passenger ordering is much less important.

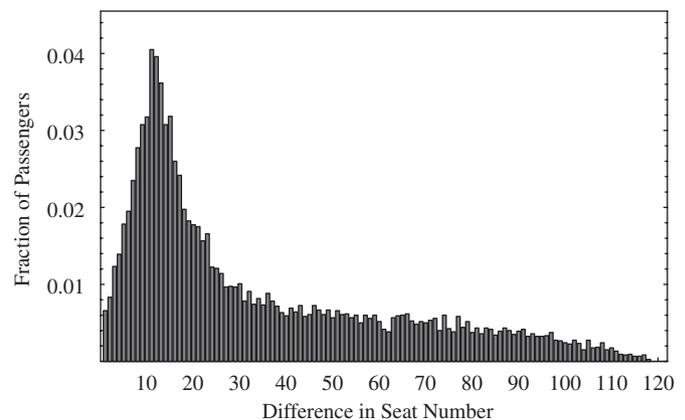


Fig. 1. Example of the seating distribution obtained from 100 realizations of the luggage loading time distribution.

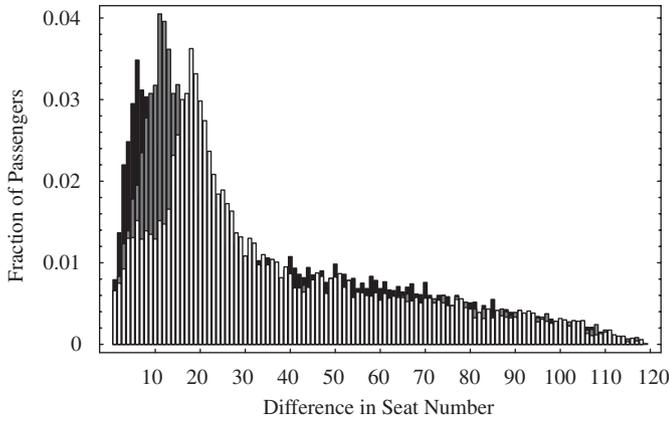


Fig. 2. Changes in the seating distribution as a function of the required personal space of the passengers. *Note:* This shows the distribution if no space is required (black), if a single row is required (gray), or if two rows are required (white). Each distribution is calculated from 100 realizations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

3.1. The optimal loading method

The optimal boarding scheme is found by extrapolating from the previous results, using the insight they provide. The reason that these seating 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 during the boarding process. That is the case where all adjacent passengers are assigned seats that are separated by exactly two rows such that the person at the front of the line is assigned a seat on the back row. The second-order effect of windows versus aisle seats can be incorporated by having the passengers in the window seats enter the airplane first, then the middle seats, and then aisle seats. The resulting seat ordering could be one of those shown in Fig. 3 although there are many equivalent orderings. These ordering provides nearly a fivefold reduction in the time that it takes over the worst case, an improvement that gets larger with a larger aircraft.

A look at the optimal boarding method shows why loading from the back of the plane to the front does not provide any benefit. If the back two rows of passengers were to board the airplane first, they would occupy roughly 12 rows of the aisle. All but the first few would be putting their luggage away while the others waited their turn—the passengers load their luggage serially. The optimal boarding strategy uses this aisle space more efficiently because each member of the first group of passengers who enter the airplane (10 passengers in the fiducial model) can put their luggage away—they load their luggage in parallel. In this manner the aisle is not used as a passive extension of the waiting area, but rather as a place for passengers to actively situate themselves. Ideally, the passengers inside the aircraft should either be seated or be loading their luggage with none waiting.

One issue that arises from this is whether or not it is practical to implement the optimal boarding scheme whereby each passenger enters the airplane in a particular order. Such a scenario may well be possible since Southwest Airlines has recently implemented a similar policy, at least to some extent. Given that, however, there will always be some fraction of passengers who are out of order; there will always be families or other groups who board together regardless of their assignments. These same issues will affect other boarding methods in similar ways.

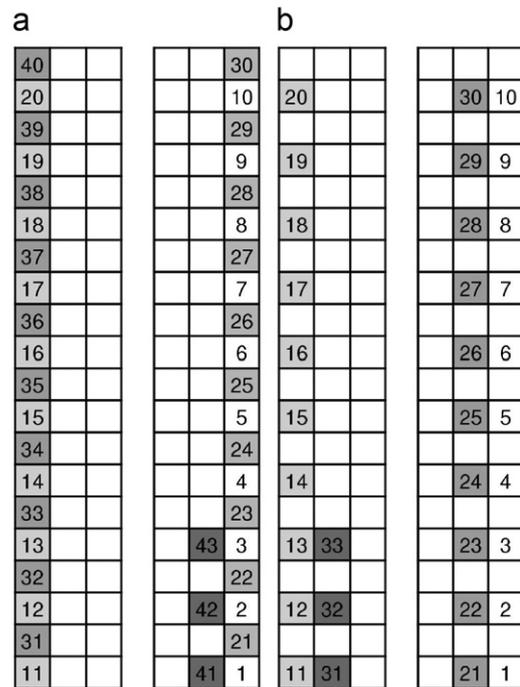


Fig. 3. Examples of the optimal passenger ordering. *Note:* There are other permutations that would give identical results. Seating would proceed following the patterns illustrated. The shading only indicates the passengers who would be inside the plane at the same time.

4. Robustness of the optimal

To test the robustness of the optimal boarding scheme two experiments are conducted. The first is to change the distribution from which passenger's loading times are selected. The second is to make random changes to the passenger ordering including swapping the locations of several random pairs of passengers and shifting the entire line by some random number (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, the minimization software was run 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 passenger seating distributions 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.

The fact that the results do not depend upon the distribution from which the luggage loading times are chosen indicates that the optimal boarding method does not depend upon the presence of non-Gaussian outlier disturbances (as it should not). Thus, if a particular passenger takes an unusually long time to load his luggage, the optimal method of passenger boarding is still, on average, optimal. Moreover, since one cannot generally identify which passenger, if it is a passenger, will be the cause of delay and since that passenger's seat will be at a random location within the aircraft, the best that one can robustly expect to do is to board optimally both before and after the disturbance.

4.2. Random shifts and swaps

Randomly shifting the line does not significantly affect the boarding time. This is because it only changes the starting point of the boarding process. All 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 six pairs for the case of 120 passengers. However, there is an upper limit to the adverse results of swapping passengers. Once they are completely randomized, additional swaps preserve the random nature of the passenger ordering and it does not get any worse. Interestingly, random boarding takes much less than half the time of the worst case boarding indicating that randomization is not catastrophic; a finding also identified by Van Landeghem and Beuselinck and Bachmat et al. and these methods also prove successful in the respective studies.

5. Comparisons

While the optimal scheme would produce the fastest boarding times, there are issues of practicality to consider. It may be challenging to arrange all of the passengers in the proper order, although, as mentioned, at least one airline has implemented this policy. Regardless, most airlines board the airplane in groups, presumably out of convenience and in an effort to reduce confusion.

5.1. Practical modification to the optimal

The advantage of the optimal boarding method 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. After trying several possibilities, as a modification to the optimal method, blocks of three consecutive seats separated by 12 was selected.

This scenario has four boarding groups and is equivalent to calling all passengers that are in even rows and from one side of the airplane. The three remaining groups are for the other side of the airplane in the same row, then the two sides of the odd rows. The loading time that results from this is not as fast as the optimum, indeed it took about twice as long to board, but it was more than a factor of two faster than the worst case. I call this boarding method the modified optimal method. Similar boarding methods, arrived at by very different means, are suggested by Van Landeghem and Beuselinck and Bachmat et al. and these methods also prove successful in the respective studies.

5.2. Comparison with other methods

Two group-boarding strategies are selected 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 ordered blocks, where a fourth of the cabin is loaded at a time starting in the back and moving to the front and a scheme where the windows are boarded first, then the middle, then the aisle seats. Within each of these groups the travelers were randomly distributed. Tests that used unordered blocks (blocks of five rows but not in order from back to front) gave results similar to the ordered blocks and were therefore not included.

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 at most three passengers out of 30 will be simultaneously loading their luggage.

The windows–middle–aisle approach reduced the boarding time to 43% of the worse 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 of 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 of getting past the person in the aisle is eliminated. One drawback of this approach is that many people travel in groups and would likely not adhere strictly to this boarding policy. In general, however, this approach did very well (as also found by van den Briel et al. (2005) though for different reasons).

The modified optimal approach performed slightly better than 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. An advantage it has over boarding windows–middle–aisle is that it allows passengers who are traveling together and sitting side by side to board at the same time without boarding out of order.

Random boarding, where the passengers positions in line are completely uncorrelated with their seat assignment, shares the advantage of the windows–middle–aisle method 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 slightly worse than the windows–middle–aisle and the modified optimal methods. 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 most efficient of the methods that employ boarding groups.

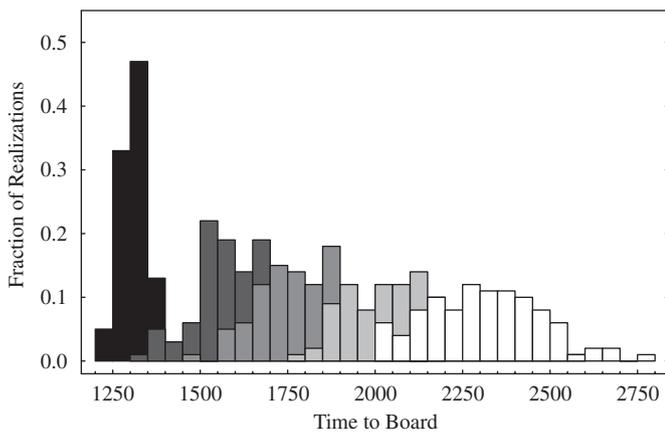


Fig. 4. Boarding times for 100 realizations of passengers when pairs of passengers are swapped. Note: Black is optimal with no swapping, then there are the histograms that correspond to 10% (dark gray), 20% (medium gray), 40% (light gray), and 60% (white) swapping. 10% means that six pairs of passengers are exchanged out of 120 passengers. The mean boarding times for these scenarios are 1312, 1585, 1795, 2084, and 2311 counts.

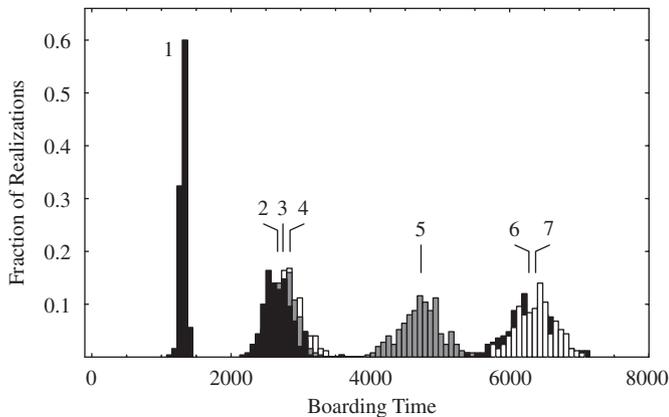


Fig. 5. The loading times for 100 realizations of seven different boarding schemes. Note: The luggage loading time for each realization is drawn from a uniform distribution with a mean of 50 counts. The alternatives are: (1) optimal boarding (mean boarding time: 1312 counts), (2) modified optimal approach (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) worst case, front to back with all passengers in order (6373).

Fig. 5 shows a histogram of the loading times for 100 realizations of seven boarding schemes, each realization having different selections for the time it takes to load the luggage and different passenger ordering where applicable. These include (in order): optimal boarding, the modified optimal approach, the windows–middle–aisle method, random boarding, ordered blocks from back to front, back to front with all passengers in order, and the worse-case (front to back with all passengers in order). We see from the 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 worst case.

6. Conclusion

Based on the assumption that a passenger loading his luggage consumes the bulk of the time that it takes for him to be seated, the primary cause for delay in the boarding process as well as the best means to overcome these delays have been explored. By boarding passengers in a manner allowing several to load their luggage simultaneously, the boarding time can be dramatically reduced. This 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 hardly better than the worse-case scenario. The goal of an optimized boarding strategy should focus on spreading the passengers who are loading their luggage 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 of these methods are likely to be psychological rather practical. Groups of passengers that 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 for passengers to line up in their assigned order were to be adopted, then there is potential for a substantial savings in time that would 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 that are used by a computer model.

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.

In the end, the time that it takes to load passengers into an airplane affects not only the airline company and the airport, but also the passengers. Few enjoy standing in line longer than necessary and still fewer enjoy sitting in an airplane longer than needed. Faster boarding would be a significant improvement for all involved parties.

Acknowledgments

The author acknowledges the generous support of the Brinson Foundation. Fermilab is supported by the US 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

- Bachmat, E., Berend, D., Sapir, L., Skiema, S., Stolyarov, N., 2006. Analysis of airplane boarding via space-time geometry and random matrix theory. *Journal of Physics A* 39, L453–L459.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculation by fast computing machines. *Journal of Chemical Physics* 21, 1087–1092.
- Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., 2002. *Numerical Recipes in C*. Cambridge University Press, Cambridge.
- van den Briel, M.H.L., Villalobos, J.R., Hogg, G.L., Lindemann, T., Mulé, A., 2005. America West Airlines develops efficient boarding strategies. *Interfaces* 35, 191–201.
- Van Landeghem, H., Beuselinck, A., 2002. Reducing passenger boarding time in airplanes: a simulation based approach. *European Journal of Operational Research* 142, 294–308.