

How I Learned to Stop Worrying and Find the Bomb

Tara Martin
Gautam Thatte
Michael Vrable
Harvey Mudd College
Claremont, CA

Advisor: Hank Krieger

Summary

We develop a queueing system model to determine the optimal number of Explosive Detection System (EDS) and Explosive Trace Detection (ETD) machines to implement 100% baggage screening for airports A and B. We test the model with data from United Airlines at Denver International Airport.

The particular queue system implementation does not affect queue length but can affect the quantity of late bags and length of delay. Our two-queue system model is 92% as efficient as an optimal priority queue, so a complex queueing system is not required. If the system can handle peak-hour volumes, there will be no delays during the rest of the day.

We also compare three flight-scheduling algorithms for peak-hour flight departures and create flight schedules for airports A and B. Optimal scheduling of peak-hour flights does not significantly change the number of machines needed, although use of a greedy algorithm reduces late bags.

To meet the 100% baggage screening requirement using EDSs, we recommend 10 for airport A, 11 for B, and 48 for United Airlines at Denver. These conservative estimates account for breakdowns and a safety margin. To replace EDSs, four times as many ETDs are needed.

Initial cost of implementation at airports A and B is \$22.9 million. This cost could be lowered by speeding the approval of cheaper and faster technologies such as dual-energy X-ray, multiview tomography, and quadrupole resonance.

The UMAP Journal 24 (2) (2003) 123–140. ©Copyright 2003 by COMAP, Inc. All rights reserved. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice. Abstracting with credit is permitted, but copyrights for components of this work owned by others than COMAP must be honored. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior permission from COMAP.

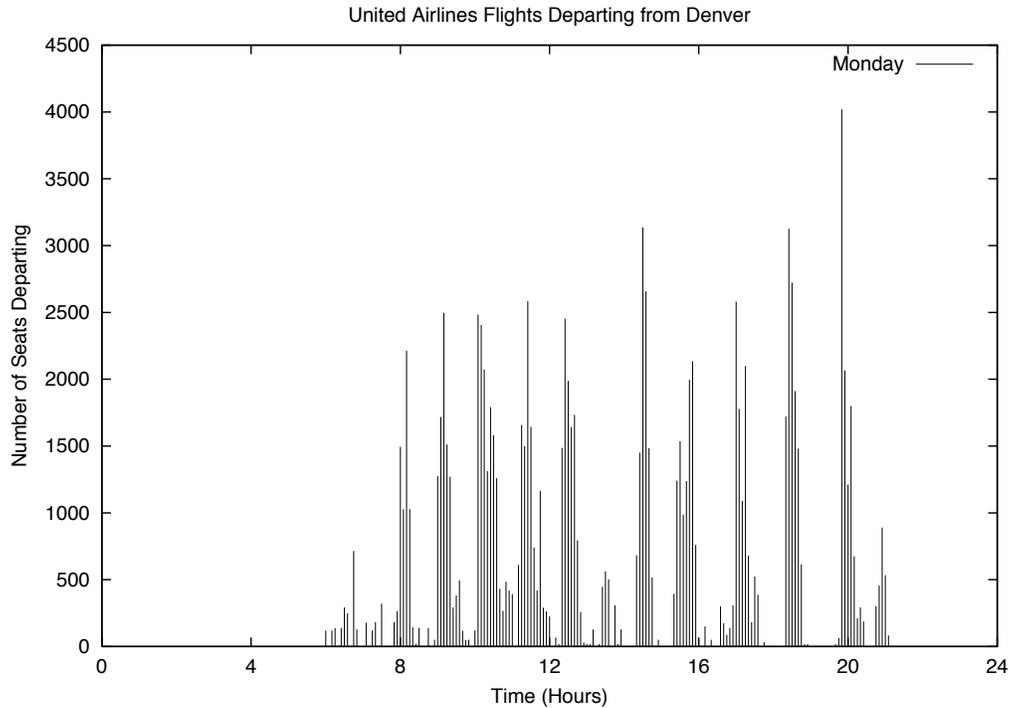


Figure 1. Flight departures from Denver by United Airlines during a single day.

Baggage Screening Queueing Models

We construct a queueing model of the screening baggage for explosives and test it with many more bags than it was designed to handle. Sample loads include peak-hour traffic at airport A and at airport B and a flight schedule modeled after traffic patterns at Denver International Airport. The raw data for the Denver simulation, summarized in **Figure 1**, consists of 991 nonstop flights on a typical Monday, as taken from a United Airlines timetable [United Airlines 2003].

Terminology

Queueing System. A system for storing bags that arrive before a screening machine can take them. The order in which the bags are removed depends on the type of queueing system. Queueing systems are described by their input, queue discipline, and service mechanism.

Queue. A system for storing bags which is first-in, first-out—that is, bags that arrive first are the first to be screened. A single queueing system might be composed of multiple queues.

Input. The *input* describes how the bags enter the system. In our model, the rate at which bags arrive varies during the day.

Queue discipline. The *queue discipline* describes how arriving bags are served, such as first-in, first-out.

Service mechanism. The *service mechanism* tells how the bags are assigned to *servers* (screening machines) as they leave the queue. Our model allows for many servers; in the case of multiple queues, the service mechanism specifies how machines are matched up with queues to process bags.

Formulation of the Model

We compute a schedule for the arrival of passengers and baggage. This baggage arrival schedule is left fixed irrespective of changes we make to the baggage queueing system to determine whether bags make it to flights on time. Our goal is a model that determines how long each bag is delayed and hence suggests an appropriate number of machines required for a specified load.

We make a number of simplifying assumptions:

- **The time required to screen a bag is short.** Any delay in delivery of the bag is due entirely to waiting for screening, not to the screening itself. This assumption allows us to disregard many distinctions among different screening machines; only the rate of screening is important.
- **Discretizing time does not introduce a large error.** Our simulation proceeds in small discrete time steps. This time step, denoted T (usually 2 min), is small in comparison to the time available for screening a bag, so rounding times to the nearest multiple of the time-step does not cause a large error.
- **Screening of a bag must be completed by some fixed time before its flight departs;** we use 10 min. A bag that does not meet this deadline is *late*.
- **Baggage screening, not check-in or other processes, is the only bottleneck.** Passengers do not encounter another bottleneck before baggage screening, such as a long line to check in, that affects the flow of bags into the screening system. This assumption allows us to consider the worst-case scenario of unlimited baggage inflow and to isolate the effects of the screening system from other airport influences.
- **It is not necessary to consider multiple separate screening systems at an airport;** if all are independent and approximately equally loaded, then the system behaves as a single system.
- **Baggage is processed at a constant rate.** We do not allow for oversized baggage or other variations that affect processing time of bags but assume these are included in the averages.

Our model is a *queueing system* [Prabhu 1997]. The input is a list of bags that arrive at each time step; the bags are grouped according to how much time they are allowed before they must be finished with the screening process. A

fixed number of servers each can process a fixed number of bags in any time step.

General Analysis

Our queueing model can be described by several parameters:

- The *service rate* S (bags/time-step) is the rate at which machines can process bags at full efficiency.
- The *input rate* $\lambda(t)$ (bags/time step) is the number of bags added to the queueing system at time t .

Regardless of implementation, the number of bags in the queueing system at any time is determined only by S and $\lambda(t)$. The implementation of the queueing system can affect the order in which bags are removed from the queueing system, not the number in it.

The total number of bags in the queueing system at time t , denoted $Q(t)$, is determined by

$$Q(t + T) = \max\{0, Q(t) + \lambda(t) - S\}.$$

If $\lambda(t) > S$, the number of bags in the queueing system increases; if $\lambda(t) < S$, the number of bags shrinks. **Figure 2** shows the bag input rate $\lambda(t)$ at the Denver airport in our model. The dashed horizontal line shows the service rate S for 36 EDS machines operating at 180 bags/h. The solid line shows the number of bags in the queueing system $Q(t)$, which increases when $\lambda(t) > S$. Approximately 52 EDS machines would be required to prevent a backlog of bags from ever building up.

Queue Disciplines and Service Mechanisms

We analyze several mechanisms for controlling how bags are stored in the queueing system before screening and later removed from it. These mechanisms have a large impact on the timely screening of bags, so choosing an appropriate mechanism is important.

Naive Model

We first develop a simple model to give an upper bound estimate on the number of EDS machines we need, using the assumptions:

- The hour before the peak-hour has significant traffic.
- The minimum number of machines is the number to ensure that no flight is delayed.

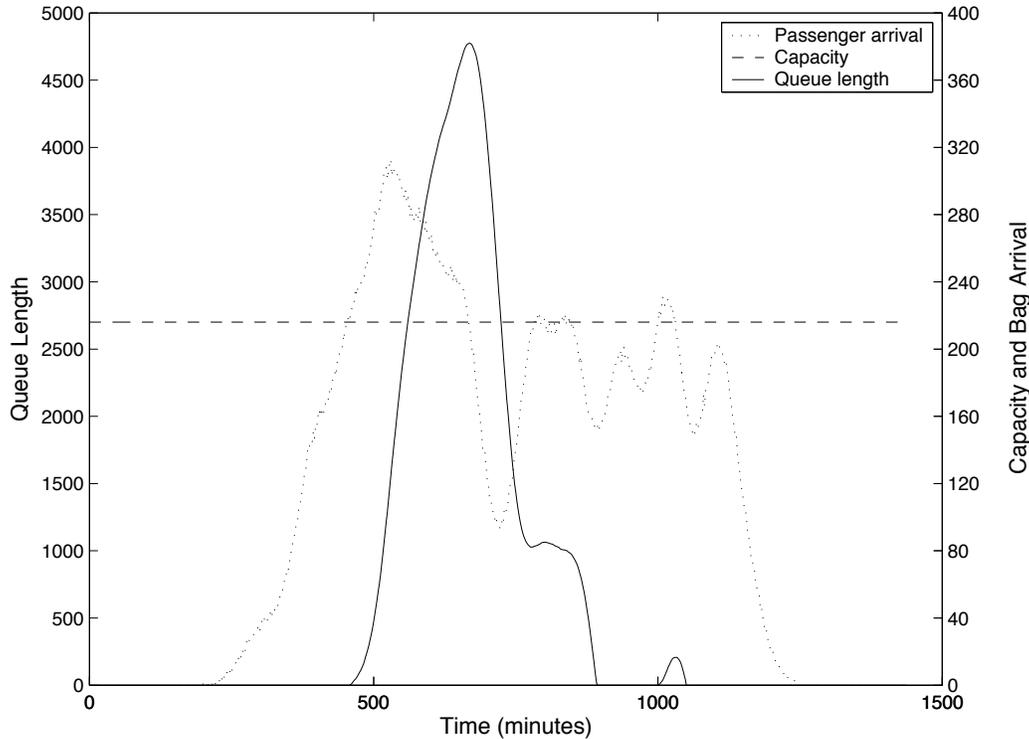


Figure 2. Bag arrival rate and number of bags in queueing system at Denver airport, assuming 36 EDS machines processing 180 bags/h.

- All bags arriving for a peak hour flight are processed in one hour.
- Bags for a flight are computed using parameters in the problem statement.

Bags arriving for peak-hour flights must be processed within a 1-h time period. Our model suggests that 34 and 37 EDS machines are required for airports A and B, respectively, and 55 for Denver International Airport. We believe that these are upper bounds. Any optimization in the passenger arrival model or the organization of people at the airport would probably achieve the same 100% success rate but with fewer machines.

Optimal Queueing

We develop an optimal queueing system that bounds the performance of any queueing system and compare various queueing models to this optimum.

We minimize the total amount of time by which bags are late.

Our optimal queueing system uses a priority queue: As bags arrive, they are added to a pile. When a bag is to be processed, we pick the bag that needs to be finished soonest.

In the Denver simulation with an optimal queue, 35 EDSs operating at 180 bags/min each are sufficient to process all bags before their deadlines. The

queue fills with up to nearly 5,500 bags at one point (26 min of uninterrupted processing is required to screen all of these).

Practical implementation of such an optimal priority queue at an airport would be too complicated. Thus, we look at other less-complex queueing systems.

Single Queue

In a single first-in, first-out queue, bags that arrive earlier are screened earlier. This scheme could be implemented with a single conveyor belt carrying bags from check-in to machines.

As long as bags can be screened quickly enough that a significant line never develops, this scheme works well. We find that 47 EDS machines at Denver suffice to deliver all bags on time; this is 34% more than required by the optimal solution.

If bags must be finished screening at least 10 min before departure, then to guarantee that all bags arriving at least 30 min before the flight are processed in time, the wait must never grow to more than 20 min. In the Denver simulation, this can be done with 38 EDS machines; approximately 0.75% of all bags arrive within 30 min of departure and are delivered late.

This single-queue system does not perform very well under load. As the queue increases in length, the chance of processing a bag late rises quickly. Although most bags arrive with more than an hour that they could wait, the few bags with less time available forces the queue length to be kept small at all times. Many bags are processed much more quickly than necessary so that the few bags that need rapid processing are not late. This situation is not optimal, and it is improved by our next queueing model.

Double Queue Model

Giving preferential treatment to some bags can produce a better queueing system. In particular, bags that arrive late should be processed more quickly. We propose a two-queue system consisting of two first-in, first-out queues for bags of different priority: A normal queue is used for bags that arrive sufficiently early and a rush queue for bags that do not arrive as early.

The total throughput of the system is not increased, but bags are much more likely to be processed before their deadline. In effect, time is borrowed from bags that have it (by placing them in a slower queue) and given to those that need it (by allowing them to jump ahead of bags in the normal queue), approximating the optimal queue discipline. The number of machines can be decreased, resulting in longer lines but without causing bags to be processed late, and also in significant cost savings.

The double queue model requires several implementation decisions:

The method for sorting bags into the two queues (the queue discipline). The cutoff may be fixed (e.g., all bags with less than 40 min to departure go into

the rush queue) or vary with the lengths of the queues.

The service mechanism. At each time step, the number of bags to remove from each queue must be determined. A fixed number of machines can be assigned to each queue; but if one queue empties, this leaves machines idle. It is better to adjust dynamically the number of machines processing bags from each queue. We suggest increasing the number of machines processing the rush queue as the rush queue increases in size.

In the Denver simulation, 42 EDS machines are sufficient to get all bags delivered on time. With 38 EDS machines, the only late bags are those that arrive late (only 0.05% of bags). This system requires 9% more machines than the optimal solution.

Evaluation

Adding more queues allows for more flexible scheduling of bag processing, which may help keep more bags from being late. However, more queues mean more parameters in the queue discipline and service mechanism, and a poor choice may harm performance. Additionally, adding queues adds complexity, with more potential for failures and higher labor cost. We believe that the benefits of a many-queue system are not worth the complexity incurred.

A double-queue system provides a performance competitive with the optimal system; we recommend its use. With only a few more machines than the 35 of the optimal queue, only a few bags are delivered late; and with only 20% more machines, no bags are late.

Validation of the Model

We account for

- unfilled seats (ranging from 0% to 50% and partially depending on the size of the flight),
- some of the passengers transfer from another flight and do not have bags rescreened (35%), and
- distribution of checked bags from 0 to 2 per passenger.

In the Denver simulation, a total of 82,500 bags are screened in a day.

We validate our model by comparing its predictions with numbers for EDS machines at actual airports. There are no statistics for the number of machines at Denver, but Dallas/Fort Worth processes 55,000 bags/day with 60 EDS machines [Douglas 2002]. If scaled to the same numbers of bags processed by Denver in our model, Dallas/Fort Worth would use 90 EDS machines. This is larger than the number we predict is necessary. However, on initial testing, EDS machines were less than half as fast as predicted (72 bags/min vs. 180 bags/min) [Clark County Department of Aviation 2002]; combined with a safety margin, our results are in agreement with the Dallas/Fort Worth figure.

Extensions to the Model

We present extensions to account for various modifications of the problem, with each change considered in isolation, not in combination with other extensions.

Accounting for Error Rates

EDS machines have a false positive rate of 30% [Butler and Poole 2002]. The result of a false positive is that the bag must be more closely examined, causing delay for that bag and some bags to be late that otherwise would not be—so more machines may be needed. We incorporate this false positive rate into our model by randomly adding a fixed time (6 min) to 30% of the bags.

At the Denver airport, the effect is to slow down screening enough that 40 EDS machines (instead of 38) are required to process all but late bags in a timely fashion. Doing so for all bags becomes nearly impossible, since some bags arrive with less than 16 min to departure.

Incorporating ETD Machines

Although we developed our model for EDS machines, it is generic enough to study other devices. We identify ways to incorporate ETD machines:

- **ETD Machines in series with EDS machines.** The problem statement relates that up to 20% of passengers may need to have bags screened through both an EDS and an ETD machine. We can account for this by giving 20% of bags an extra delay of 4 min.

We assume that there is no queue between EDS and the ETD machines following them—appropriately many ETD machines are purchased to match the processing speed of EDS machines. In the Denver simulation, an increase to 39 EDS machines, instead of 38, allows all but late-arriving bags to be processed on time.

- **ETD machines replace EDS machines.** We can calculate the number of ETD machines necessary to obtain the same service rate as for EDS machines and compare the costs. Any mixture of the two machine types with the same overall service rate will behave the same in our model; but since the cost varies linearly as machines of one type are replaced with the other, the most cost-effective operation will occur at one of the extremes, either all EDS or all ETD machines.

Assuming a rate of 45 bags/min for an ETD machine (one-fourth the rate of an EDS machine), four times as many ETDs will be needed. According to Butler and Poole [2002], ETD machines cost less than one fifth the amount of EDS machines to operate.

Strengths and Weaknesses of the Model

Our model succeeds in capturing the essence of the problem and allows for good predictions, as a result of its many strengths:

- **Our model is based on real-world data.** Use of data from Denver makes it much more likely that the results from our model are realistic and not artifacts of an artificial flight distribution, such as the isolated peak hour of flight at airports A and B. Additionally, agreement with figures for EDS and ETD machines currently installed at airports gives us confidence that our model is accurate.
- **Our model is flexible enough to handle other types of screening machines, passenger arrival schedules, etc.** Our model's parameters can be varied to account for changes in screening machinery, training of baggage screening personnel, and so on. Since our queueing simulation takes as input merely a list of arrival times for bags, it is also very easy to study airline flight schedules at any other location, or to modify the arrival behavior of passengers.
- **Our model can predict the screening capacity needed as well as predict how the system will fail.** Our model goes beyond merely predicting the number of baggage screening machines needed to process all bags on time to give a complete model for the flow of bags through the system. The model can thus be used to see exactly how the baggage screening system will begin to break down as it is pushed past its limits. This information will help airports evaluate what margin of safety they require.

At the same time, there are aspects of our model that could be improved:

- **More detailed data for machine operation could be incorporated.** Our model is rather simplistic in that all behavior is based only on the waiting time to process bags. Including the actual time to scan a bag (not just the queue wait time) may be better, especially for systems that are slower to screen bags.
- **Queue scheduling could be optimized further.** Our proposed two-queue system generally performs well, but we have not completed a detailed analysis of it nor systematically determined optimal values for its parameters.

Recommendations

Based on simulations and an analysis of our model, we are able to make a number of recommendations:

- **A safety margin can make a significant difference.**

The loss of just a small percentage of the capacity of the system can make the difference between no late bags and a significant fraction of late bags.

This is shown in **Figures 3a** and **3b**. After the number of machines in use drops by about 10%, the number of late bags rises dramatically, regardless of queueing algorithm.

Seemingly paradoxically, the optimal queueing algorithm has the highest fraction of late bags for some values; this is because it sacrifices the percentage of bags on time for decreasing the average amount by which bags are late.

Since unpredictable slowdowns or large arrivals should be anticipated, planning to handle a larger than expected number of bags is necessary to avoid breakdown of the system. With EDS machines operational 92% of the time, at least 8% more EDS machines should be installed than predicted as necessary by our model. We recommend a further margin of safety, perhaps 10%, to account for any other unexpected circumstance, such as unusually high traffic.

Based on these considerations, we recommend 48 EDS machines for Denver, 10 for airport A, and 11 for airport B.

- **Backlogs should be avoided except at peak times.** The processing capacity (bags/min) should be set higher than the arrival rate of bags at all but peak times. Further, peak times must be fairly well isolated (to an hour or so), or queue lengths will grow quickly to unmanageable levels. When a line develops for scanning, it can then take a good deal of time to get back to a no-wait situation. While our model shows that a persistently long queue can sometimes be handled as long as it does not continue to grow, a long backlog of bags is unstable—any event that causes the queue to grow in length quickly causes many late bags. Thus, a persistently long queue indicates insufficient screening capacity safety margin.
- **Set stricter deadlines for passenger arrival before flights.** We assume that airlines are fairly lenient about accepting bags from passengers up to the final deadline for placing them on a flight. An airline could establish a policy wherein bags that are not checked by a certain time before the flight—say, 30 min—are not guaranteed to make the flight. With such a policy, any time we have identified a strategy as handling all but “late bags,” all bags would be handled in time—the “late bags” would have been rejected by the airline outright and would not delay the flight.
- **Plan for future growth in aircraft travel.** Historical data shows a growth of about 6% per year in the number of airline passengers [Metropolitan Airports Commission 2003]. Since a screening system is a large investment, an airport should plan with an eye to future capacity. The dip in traffic since 2001 may be only temporary and airline traffic may return to its normal growth curve, with a corresponding larger-than-usual increase in traffic in the next year or two.

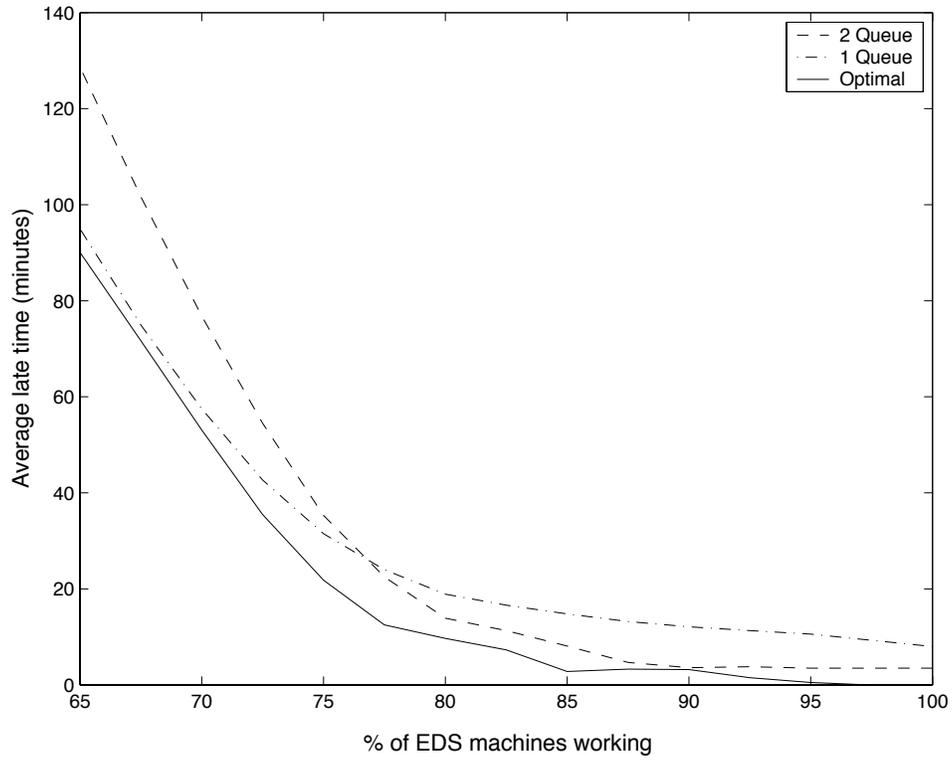


Figure 3a. Effects of the removal of baggage-screening capacity on the percentage of bags screened on schedule, for various queueing algorithms.

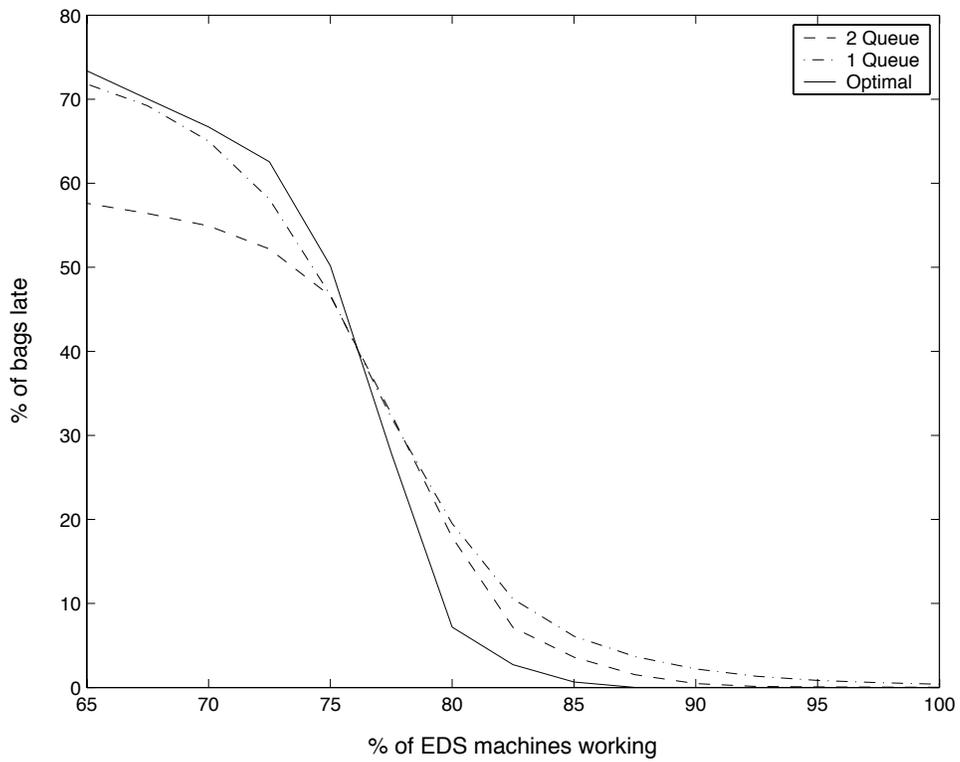


Figure 3b. Effects of the removal of baggage-screening capacity on the average delay for a bag, for various queueing algorithms.

- **Install a baggage screening system early, and ramp up use.** Unexpected difficulties may arise with a new screening system. In addition, machine operators need to become proficient. If an airport installs a baggage screening system in advance of the federally mandated deadline, screening can begin below 100% and increase to 100% by the deadline as problems are dealt with.

Optimal Peak-Hour Scheduling

We develop three passenger arrival models to schedule flights during the peak hour, with three distinct passenger arrival profiles and two arrival concentration distributions. The following assumptions simplify the model without reducing the validity of the simulations.

Assumptions

- **On average, passengers arrive 1.5 h before departure.** The problem statement says “between forty-five minutes and two hours”; although 1.5 h is not the middle of that range, it is close and makes for easier modeling.
- **Passengers arrive according to a Gaussian distribution.** We adopt a Gaussian arrival model from Clark County Department of Aviation [2002]; such a distribution encompasses realistic features, such as a peak in arrivals considerably before flight departure. We chose a mean of 90 min and tried standard deviations of 15 min and 30 min, implying that respectively 95% and 70% of passengers arrive between 2 h and 1 hour before their flights.
- **Flights scheduled to leave during the peak hour are uniformly spaced.** This assumption accommodates a generic runway structure.

Passenger Arrival Models

We apply three passenger arrival models to airports A and B. The peak-hour data given in the problem statement were processed both in isolation (no other flights during the day) and as part of a busier schedule that affects peak-hour departures.

Random Placement Algorithm

A random placement of flights within the peak hour, according to a uniform distribution, makes different parts of the hour look approximately the same. We regard this algorithm as a baseline.

Bimodal Distribution Algorithm

A bimodal distribution schedules the largest flights at the beginning and at the end of the hour in an attempt to reduce the peak in passenger arrival. This method is useful only when the standard deviation of arrival distributions is low (such as $\sigma = 15$ min). At higher standard deviations (such as $\sigma = 30$ min), the bimodal distribution converges to the distribution obtained with the greedy algorithm below.

Greedy Algorithm

A greedy algorithm always makes the optimal local choice in the hope that the final solution will be globally optimal [Cormen et al. 2001]. Our greedy algorithm attempts to minimize the peak in the arrival distribution and thus reduce a major peak in passenger arrival for peak hour flights. The following methodology is used:

- We consider the flights sequentially from largest flight to smallest.
- At each step, the center of the passenger arrival Gaussian distribution being considered is assigned to the minimum value among the possible centers of the distributions.
- Each center cannot be used for more than one distribution.

Simulation Results

We ran each of the passenger models through the optimal baggage screening model to determine which would be best suited for airports A and B. The $\sigma = 30$ cases outperformed the $\sigma = 15$ cases for all arrival distributions, which implies that having nearly all the passengers arrive for peak hour flights at the same time backs up the queue significantly.

The procedure used to combine the given peak-hour data and the Denver data involved:

- The peak-hour of the Denver data was identified as 10 A.M. to 11 A.M., with a maximum rate of baggage arrival of 160 bags/min.
- The peak-hour data for airports A and B were scaled up by a factor of 3.5 to approximate better the volume at Denver.
- The peak hour in the Denver data was entirely replaced by the airport A and B data in their respective simulation.

Both the embedded and the isolated peak data were processed using the optimal baggage screening algorithm.

The greedy algorithm creates a schedule that performs up to 50% better (in terms of total late time for bags) compared to the random schedule, when

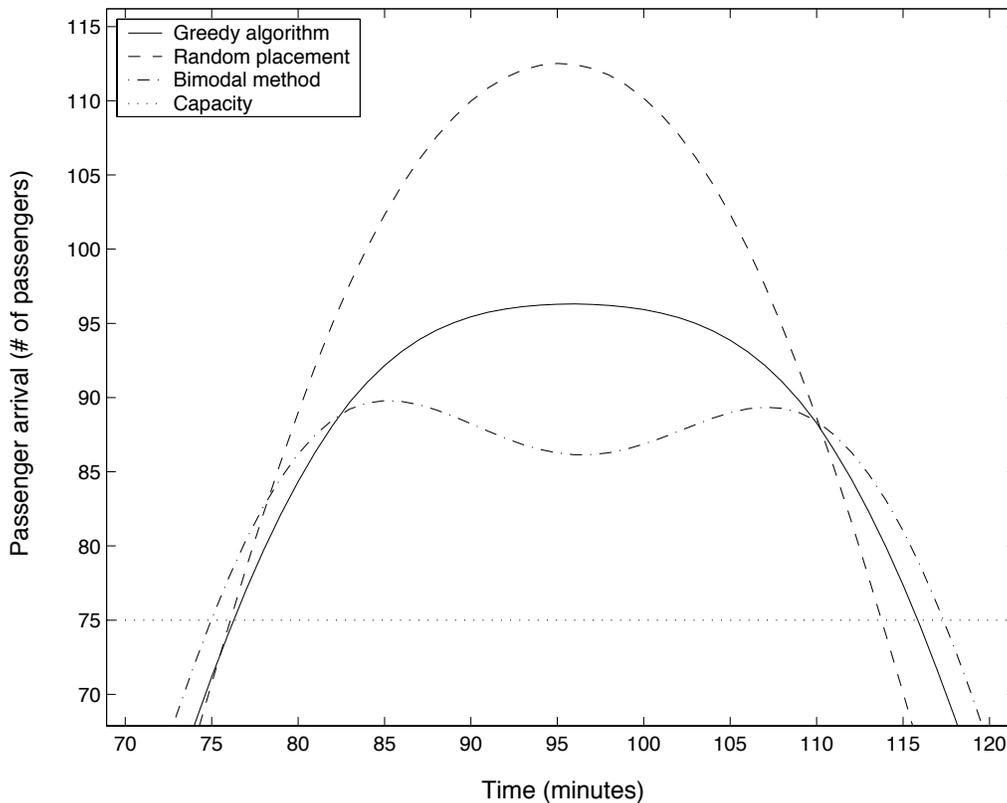


Figure 4. Comparison of peaks of passenger arrival profiles illustrating the superiority of the greedy schedule over the random and bi-modal distributions.

the peak hour is embedded in the relatively busy day; in isolation, a greedy algorithm schedule was about 30% superior. The bimodal algorithm produced a schedule that was worse than the baseline; we therefore eliminate it.

At a load below capacity of the machines, any scheduling algorithm will do. Above capacity, some methods perform better than others. The efficiency of a scheduling algorithm may be gauged by how long its operating capacity is exceeded and how backed up the queue becomes.

In **Figure 5**, notice that the bimodal profile exceeds its capacity first and continues to operate above capacity for the longest time. Even the intermediate decrease in queue backup is not enough to allow the bags to be processed faster than either the random or the greedy profiles. On the other hand, although the random placement profile exceeds its capacity latest and again drops below capacity earliest, its high peak leads to a significant queue backup that cannot be cleared as quickly as in the greedy profile. This latter profile balances both factors, giving the best result.

We used the two better algorithms to develop schedules for airports A and B. [EDITOR'S NOTE: We omit the details of the schedules.] For both airports, the greedy algorithm generated a better schedule. Both methods resulted in the use of the same number of EDS machines at the airports, although the greedy

schedule results in fewer late bags. Airports A and B require 8 and 9 EDS machines, respectively, for 100% baggage screening and no delays due to the screening process.

Recommendations

With normal or above-normal traffic during pre-peak hours, the scheduling of flights during the peak hours does not matter much, because passenger arrivals are spread out over 3 hours, reducing the impact of changes within the peak hour.

If the peak hour has significantly more traffic than pre-peak hours, then the greedy algorithm is better than either the random or the bimodal distributions.

Review of Future Technologies

Current technology approved by the FAA is highly limited and extremely expensive.

EDS machines produce a three-dimensional image of the contents of a bag, allowing observation of hidden materials, zoom, and rotation of perspective to focus on suspicious objects. Unfortunately, EDSs use a powerful X-ray that requires screening to protect operators, is very expensive, and—due to the high sensor rotation rate required to resolve images—is limited in speed.

ETD machines use mass spectrometry to detect trace levels of explosives. The sample collection takes much longer and has much higher labor requirements than the EDS, with a critically high false-negative rate of 30% for a surface sample and 15% for an open-bag sample. This poor detection rate is due to the uneven concentration of explosive residues within a bag [Butler and Poole 2002].

Few alternatives have been developed as fully as EDSs and ETDs, but some appear very promising:

- **Coherent scatter** is slower than EDS (60–240 bags/h), but with a near perfect detection rate and an order of magnitude fewer false-positives, it is still relatively efficient [Butler and Poole 2002].
- **Dual-energy X-ray** has a high false alarm rate of 20% [Singh and Singh 2003] but can process 1,500 bags/h. These systems are being installed in London and other European airports and await certification in the U.S. [Butler and Poole 2002].
- **Stereoscopic tomography**, slightly different from the computed Tomography used in EDSs, scans 1,200–1,800 bags/h and is being tested for accuracy and false alarm rates [Singh and Singh 2003].

- **X-ray diffraction** uses unique diffraction patterns of scanned materials to determine their chemical composition. Current experiments show a nearly perfect detection rate and extremely small false-alarm rate [Singh and Singh 2003]. Throughput rates and cost will likely be similar to that of normal X-ray scanners, making this a promising technology.
- **Neutron-based detection** is used in several developing techniques:
 - Thermal neutron analysis (TNA)** can detect nitrogen levels particular to many plastic explosives but has limited sensitivity, a high false-alarm rate due to background nitrogen levels, and is at least as expensive as an EDS, making it a less promising candidate.
 - Fast neutron analysis (FNA)** is similar to TNA except that it can also detect oxygen, carbon, and hydrogen levels, allowing greater sensitivity and accuracy. However, the high-energy neutrons used create large amounts of noise, making information difficult to detect.
 - Pulsed fast neutron analysis (PFNA)** solves the noise problem but requires a collimated, pulsed energetic neutron beam, which is hard to make and tends to be unsafe and expensive.
 - Pulsed fast thermal neutron analysis (PFTNA)** uses a shorter pulse. It measures both thermal and fast neutron information. Portable models for landmine, unexploded ordinance, and narcotic detection have very high accuracy levels [Singh and Singh 2003].
- **Quadrupole resonance** uses magnetic resonance techniques to identify the composition of the scanned object. Every material releases a unique signal; those corresponding to explosive compounds can be isolated and identified. Machines using this technique are under construction; the manufacturer predicts that this technology will be faster (300 bags/h) and more accurate than both EDS and ETD [Quantum Magnetism 2002].
- **Millimeter wave imaging** is a noninvasive technique that detects short wavelength electromagnetic radiation from scanned objects. While this appears to work well for locating weapons concealed about a person, it does not seem able to distinguish explosive materials from inert ones and is thus not useful for baggage scanning [Homeland Security Research 2002]. Microwave imaging is similar to millimeter wave imaging.

Conclusion

Frankly, we've tried everything else We've put up more metal detectors, searched carry-on luggage, and prohibited passengers from traveling with sharp objects. Yet passengers still somehow continue to find ways to breach security. Clearly, the passengers have to go.

—*The Onion* (16 October 2002)

Since excluding passengers is unrealistic, we study the more practical technique of scanning baggage. Our results are:

- We develop a model that predicts the behavior of a queueing system for baggage in an airport security screening system and allows prediction of delays caused by the system. This model is then expanded to include multiple types of screening machines and false-positive results.
- We evaluate our model against real-world data for Denver International Airport and for the data given for airports A and B.
- Using our model, we predict the optimal number of Explosive Detection System (EDS) or Explosive Trace Detection (ETD) machines to use at several different airports and provide other recommendations for the implementation of a security screening system. For Denver, we recommend 48 EDS machines; at airports A and B we recommend 10 and 11 machines, respectively. We also compare these figures to actual figures for EDS use at the Dallas/Fort Worth Airport.
- We study how the distribution of flights during the peak hour of the day affects the efficiency of the system. We propose a greedy algorithm for optimally scheduling flights.
- We review promising technologies for future security screening machines.

Our evaluation of the requirements for 100% baggage screening suggests that such high security goals are cost-ineffective, so research into alternative technologies and screening systems is needed.

References

- 107th Congress. 2001. Aviation and Transportation Security Act. http://www.tsa.dot.gov/interweb/assetlibrary/Aviation_and_Transportation_Security_Act_ATSA_Public_Law_107_1771.pdf.
- Butler, Viggio, and Robert W. Poole, Jr. 2002. Re-thinking checked-baggage screening. July, 2002. Los Angeles, CA: Reason Public Policy Institute. www.rppi.org/baggagescreening.html.
- Clark County Department of Aviation. 2002. EDS bag screening analysis. Presentation posted online. <http://www.aci-na.org/docs/EDS%20Bag%20Screening%20Analysis.ppt>. Revised 17 January 2002.
- Cormen, Thomas H., Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. 2001. *Introduction to Algorithms*. 2nd ed. Cambridge, MA: MIT Press.
- Devore, Jay L. 2000. *Probability and Statistics for Engineering and the Sciences*. 5th ed. Belmont, CA: Duxbury.

- Douglas, Jim. 2002. House votes down Jan. 1 airport-security deadline. http://www.wfaa.com/jdouglas/stories/wfaa020726_am_dfwdelays.2b8745ed.html.
- FAA considering passenger ban. 2002. *The Onion* (16 October 2002); http://www.theonion.com/onion3838/faa_passenger_ban.html.
- Homeland Security Research. 2002. Homeland Security Analyst—May 2002—Technology Focus. http://www.hsrb.biz.newsletter/May_2002/newsletter_May_2002_tech_focus.htm.
- Metropolitan Airports Commission. 2003. *2002 MAC Annual Report*. <http://www.msairport.com/MAC>.
- Prabhu, N.U. 1997. *Foundations of Queueing Theory*. Boston, MA: Kluwer.
- Quantum Magnetics. 2002. Quadrupole resonance. http://www.qm.com/core_technology/quadrupole_resonance_body.htm.
- Singh, Sameer, and Maneesha Singh. 2003. Explosives detection systems (EDS) for aviation security. *Signal Processing* 83: 31–55; http://www.dcs.ex.ac.uk/research/pann/pdf/pann_SS_084.PDF.
- United Airlines. 2003. United.com Electronic Timetable. <http://www.ual.com/page/middlepage/0,1454,1891,00.html>.