A Sequential Stochastic Multilevel Passenger Screening Problem for Aviation Security

Laura A. McLay
Department of Statistical Sciences & Operations Research
Virginia Commonwealth University
lamclay@vcu.edu

Sheldon H. Jacobson
Simulation and Optimization Laboratory
Department of Computer Science
University of Illinois at Urbana-Champaign
shj@uiuc.edu

Adrian J. Lee
Department of Mechanical Sciences and Engineering
University of Illinois at Urbana-Champaign
ajlee4@uiuc.edu

John E. Kobza
Department of Industrial Engineering
Texas Tech University john.kobza@ttu.edu

December 5, 2008

Abstract

Designing effective aviation security systems has become a problem of national concern. Passenger screening is an important component of aviation security. Developing real-time passenger screening strategies can be quite challenging. This paper identifies a methodology that can be used in real-time to optimally assign passengers to aviation security resources for the Sequential Stochastic Multilevel Passenger Screening Problem (SSMPSP). In SSMPSP, there is a set of classes available for screening passengers, each of which corresponds to several device types for passenger screening. Passengers arrive sequentially, and a prescreening system determines the passengers' perceived risk levels, which becomes known upon check-in. The objective of SSMPSP is to use the passengers' perceived risk levels to determine the optimal policy for screening passengers that maximize the expected total security, subject to capacity and assignment constraints. SSMPSP is formulated as a Markov Decision Process, and an optimal policy is found using dynamic programming. The Sequential Stochastic Assignment Heuristic (SSAH) is introduced to assign passengers to classes in real-time. A condition is provided under which SSAH yields the optimal policy. The model is illustrated with an example that incorporates flight schedule and passenger volume data extracted from the Official Airline Guide. Analysis of the example shows that SSAH is robust in thoroughly screening the high-risk passengers. The models and analysis provided in this paper provide insight into real-time operations of passenger screening systems.

1 Introduction

In the aftermath of September 11, 2001, widespread aviation security policy and operational changes were made throughout the nation's airports. The piecemeal and reactive nature of many of these changes has resulted in large increases in costs and inconvenience to travelers, without a corresponding increase in security. The numerous changes in passenger screening policy following the August 2006 arrest in London of several suspected terrorists plotting to blow up ten US-bound transatlantic flights, serve to further illustrate this point. These experiences suggest the need for new aviation security paradigms that bring together the key security system factors (namely, technology, intelligence, and procedures) needed to achieve screening operations that are, efficient, cost-effective, and nonintrusive.

There are two basic approaches to aviation security screening: *uniform screening* and *selective screening*. Uniform screening subjects every passenger and their baggage to identical security screening procedures. The argument for uniform screening is that anyone could pose a risk, and hence, all passengers should be screened using the most effective technology and procedures available. The 100% baggage screening mandate, which requires all checked baggage to be screened by a federally certified explosive detection technology (effective December 31, 2002), is a move towards uniform screening (US GAO 2006). One disadvantage of uniform screening is that it can be very costly to apply expensive new technologies to every passenger or all baggage. Butler and Poole (2002) and Poole and Passantino (2003) argue that 100% checked baggage screening is not cost-effective, and suggest that creating multiple levels of security for screening passengers may be more effective than treating all passengers the same.

Selective screening selectively applies security technologies and procedures on a targeted subset of passengers, since most passengers do not pose a risk. Selective screening subjects passengers perceived as high-risk to closer scrutiny by screening them and their baggage with more sensitive and accurate technologies and procedures, while passengers perceived as low-risk are subjected to lower levels of scrutiny. This approach requires that a prescreening system perform a risk assessment of each passenger prior to the passenger's arrival. A weakness of selective screening is that it can assign an incorrect degree of risk to a passenger, either by error or through "gaming" of the system by terrorists. McLay et al. (2006) conclude that using expensive and accurate baggage screening technologies on a small proportion of passengers is warranted only if there is an accurate prescreening system in place.

The terrorist events on September 11, 2001 resulted in the implementation of additional passenger screening procedures (Mead 2002, 2003, US GAO 2005, 2006). One policy change following September 11, 2001 is the development of *Secure Flight*. Secure Flight performs a check on passenger information against a consolidated Federal terrorist watch lists (US TSA 2006). Passengers who are permitted to fly by Secure Flight are then prescreened by CAPPS, the Computer-Aided Passenger Prescreening System. CAPPS performs a risk assessment on each passenger and partitions passengers into two classes: nonselectees and selectees, where nonselectees are passengers who have been cleared of posing a risk, and selectees are passengers who have not been cleared, based on limited information known about them. A prescreening system such as CAPPS can be used to assign an *assessed threat value* to each passenger, a scalar value that quantifies the risk associated with the characteristics of the passenger. CAPPS distinguishes selectees and nonselectees by requiring additional screening (such as hand searches) for selectees and their baggage. The selectee and nonselectee classes are each defined by a preassigned subset of devices and a procedure through which passengers are processed prior to boarding an aircraft. CAPPS was refined after September 11, 2001, resulting in a new prescreening system called CAPPS II. However, CAPPS II was eventually dismantled over privacy concerns. At present, an updated version of CAPPS is being used to prescreen passengers at the nation's airports.

There are several devices available for screening passengers, where each device is an aviation security technology or procedure used to identify a threat. A device screens passengers in one of three ways: by screening checked baggage, carry-on baggage, or passengers. At present, all checked baggage is screened for explosives either by an explosive detection system (EDS) or an explosive trace detection device (ETD). All passengers are screened with a magnometer and their carry-on baggage is screened with an X-ray machine. Each device has an associated capacity, the upper bound on the number of passengers or bags that a device can screen in a given amount of time. Selectees and their carry-on baggage are differentiated from nonselectees by undergoing hand searches by airport screening personnel. In some airports, selectees are screened with hand wands or trace portals and their carry-on baggage is screened by ETDs.

A weakness of any prescreening system, including CAPPS, is that such systems can be gamed through extensive trial and error sampling by a variety of passengers passing through the system (Barnett 2001; Chakrabarti and Strauss 2002). Martonosi (2005) and Martonosi and Barnett (2006) report that the underlying screening process has a larger impact on reducing successful attacks than an effective prescreening system. Barnett (2004) suggests that CAPPS II may only improve aviation security under a particular set of circumstances and recommends that prescreening systems be transitioned from a security centerpiece to one of many components in future aviation security strategies. The TSA describes prescreening systems as critical components in a layered system for aviation security, including reinforced cockpit doors, bomb-sniffing dogs, and deploying Federal air marshals on numerous flights (US TSA 2004).

Integer programming and discrete optimization models have been used to formulate several aviation security problems when a system such as CAPPS prescreens passengers. Jacobson et al. (2001) provide a framework for measuring the effectiveness of baggage screening security device deployments for screening selectee baggage at a particular airport station. Jacobson et al. (2003) introduce three performance measures for baggage screening security systems and introduce models to assess the security effect for single or multiple airport stations. Jacobson et al. (2005a) formulate problems that model multiple sets of flights originating from multiple airport stations subject to a finite amount of security resources; these problems consider the three performance measures introduced in Jacobson et al. (2003). Examples are presented to illustrate strategies that may provide more robust device allocations across all these performance measures. Jacobson et al. (2005b) construct integer programming models for problems that consider multiple sets of flights originating from multiple airports. Virta et al. (2002) consider the impact of originating and transferring passengers on the effectiveness of baggage screening security systems. Both these papers consider classifying selectees into two types; those at their point of origin and those connecting through a hub airport. This is noteworthy since at least two of the hijackers on September 11, 2001 were connecting passengers. Babu et al. (2006) use linear programming models to investigate the benefit acquired from using multiple risk groups for screening passengers. They conclude that using multiple risk groups is beneficial for security, even when a prescreening system is not used to differentiate passenger risk. Nie et al. (2006) extend this model to consider passenger risk levels, as determined by a passenger prescreening system, and formulate the resulting model as a mixed integer program. They find that using passenger risk levels results in a more efficient security system.

Other research has focused on the experimental and statistical analysis of risk and security procedures on aircraft. Barnett et al. (2001) report the results of a large-scale, two-week experiment at several airports to identify costs and disruptions that would arise from using positive passenger baggage matching, an aviation security procedure where passengers' checked baggage is removed from a flight if the passengers do not board the aircraft. They conclude that using positive passenger baggage matching results in an average delay of one minute per flight, and its implementation costs an additional forty cents per passenger. Barnett et al. (1979) and Barnett and Higgins (1989) study mortality rates on passenger aircraft and perform a statistical analysis on this data. Czerwinski and Barnett (2006) analyze differences in airlines protecting passengers from death and recovering from emergencies that have occurred. They find no evidence that established airlines are safer than new-entrant airlines.

This paper introduces the Sequential Stochastic Multilevel Passenger Screening Problem (SSMPSP) that models multilevel passenger screening strategies using Markov Decision Processes and discrete optimization models. Multilevel screening considers an arbitrary number of classes to screen passengers, as opposed to the binary system of CAPPS. Multilevel screening problems are motivated by McLay et al. (2006a,b), who introduce the Multilevel Allocation Problem (MAP) and the Multilevel Passenger Screening Problem (MPSP). In these problems, the set of passengers to be screened at a particular station in an airport in a given period of time is assumed to be known, and hence, the assessed threat values are assumed to be known a priori. This assumption is relaxed by SSMPSP, in which passengers check in sequentially, and each passenger's assessed threat value becomes known upon check-in. This necessitates a change in the solution methodology since passenger screening decisions are made simultaneously in MAP and MPSP, whereas a series of passenger screening decisions are made sequentially for SSMPSP. The Sequential Stochastic Passenger Screening Problem (SSPSP) is a particular case of SSMPSP in which there are only two classes available for screening passengers, namely selectees and nonselectees (McLay and Jacobson 2007). Nikolaev et al. (2007) propose a two-stage model for the sequential stochastic security design problem (SSSDP). The first stage analyzes the purchase of security devices, while the second stage determines the screening assignments of sequentially arriving passengers. However, SSSDP is transformed into a deterministic integer program rather than modeled using Markov Decision Processes.

The primary contribution of this paper is to identify a real-time methodology for screening passengers in a multilevel screening paradigm and to show how the methodology can be used to provide insights into the operation and performance of such real-time systems. Note that this research assumes that a prescreening system such as CAPPS has been implemented and is effective in identifying passenger risk (i.e., the assessed threat values accurately quantify passenger risk).

The paper is organized as follows. Section 2 introduces SSMPSP as a Markov decision process, and shows how the optimal policy for SSMPSP can be obtained by dynamic programming. Section 3 describes the Sequential Stochastic Assignment Heuristic (SSAH) for SSMPSP. SSAH defines a policy for assigning passengers to security resources. A property is provided under which SSAH obtains the optimal policy. Section 4 provides examples using passenger volume and flight data extracted from the Official Airline Guide and discusses computational results for several scenarios based on these examples. The policy provided by SSAH is analyzed to provide insight into the operation of real-time passenger screening systems. Section 5 provides concluding comments and directions for future research.

2 Optimization Models

SSMPSP is an extension of the Multilevel Passenger Screening Problem (MPSP), a general framework for multilevel security screening strategies (McLay et al. 2006a). In SSMPSP, passengers check-in sequentially, and the decisions to assign the passengers to classes are made as passengers check in for their flights. This section first formulates MPSP as a discrete optimization problem and then formulates SSMPSP as a stochastic optimization model and Markov decision process.

2.1 Multilevel Passenger Screening Problem

MPSP considers the problem of optimally screening a set of passengers arriving at a security station in a given time interval. The objective of MPSP is to assign N passengers to M security classes such that the total security is maximized, subject to assignment and device capacity constraints, where each security class is defined in terms of the set of devices it uses (McLay et al. 2006a). Passengers and their baggage are screened by a sequence of devices associated with the class to which they have been assigned. This sequence of devices is assumed to be fixed in advance, based on the procedures set by the TSA. Although MPSP assigns passengers to classes, it implicitly determines which devices are critical for passenger screening operations (i.e., which devices operate at capacity) as well as which classes should (and should not) be used.

Note that each device gives one of two possible responses: alarm or clear. In addition, the system gives one of two possible outcomes: alarm or clear, which is a function of the device outcomes and can be defined in several ways (Kobza and Jacobson 1996, 1997). Each passenger is either a threat or a nonthreat. Ideally, the system yields a clear response for all of the nonthreat passengers and yields an alarm response for all of the threat passengers. Although it is not known in advance whether a given passenger will yield an alarm response, it is assumed that there are procedures in place for resolving alarms and that adequate resources are available to resolve all alarms given by the system.

MPSP is formally stated as a discrete optimization problem.

The Multilevel Passenger Screening Problem (MPSP)
Given:

A set of N passengers, each characterized by an assessed threat value $\alpha(1), \alpha(2), \ldots, \alpha(N)$, $0 < \alpha(j) \leq 1,\ j = 1, 2, \ldots, N$,

a set of M classes,

a set of V devices types, where device type k has capacity $c_k,\ k = 1, 2, \ldots, V$,

a subset of device types associated with each class, with $d_{ik} = 1(0)$ if class $i = 1, 2, ..., M$ uses (does not use) device type $k = 1, 2, ..., V$,

the *security level* of each class, $L_i,\ 0 \leq L_i \leq 1,\ i = 1, 2, \ldots, M$.

Objective: Denote passenger assignments for the N passengers to the M classes by $A_1, A_2, \ldots, A_M$, where $A_i \subseteq \{1, 2, \ldots, N\}$ represents the subset of passengers assigned to class i, and define the *risk level* R_i of class $i = 1, 2, \ldots, M$ as the proportion of assessed threat values of the passengers assigned to class i. Therefore,

$$R_i = \frac{\sum_{j \in A_i} \alpha(j)}{\sum_{j=1}^{N} \alpha(j)}, \; i = 1, 2, \ldots, M.$$

Find passenger assignments $A_1, A_2, \ldots, A_M$ where $\bigcup_{i=1}^{M} A_i = \{1, 2, \ldots, N\}$, $A_i \cap A_j = \emptyset$ for $i = 1, 2, \ldots, M$, $j = 1, 2, \ldots, M$, $i \neq j$, such that each device type is operating within its capacity (i.e., $\sum_{i=1}^{M} d_{ik} |A_i| \leq c_k, k = 1, 2, \ldots, V$) and the total security

$$z^{OPT} = \sum_{i=1}^{M} L_i R_i \tag{1}$$

is maximized.

The assessed threat values are determined by a prescreening system such as CAPPS. The details of CAPPS are classified, and the assumptions of how it works are based on information disseminated in the public domain. It is also assumed that each assessed threat value is determined independently and in real-time, since the prescreening system is a computer program that computes deterministic assessed threat values based on information that passengers provide at the point of ticket purchase.

The security levels and the risk levels can be obtained using information and data available from the TSA. The security level of each class (scaled between zero and one) is based on the security procedures of each device used to screen passengers in that class. In this case, the security level for class i is interpreted as the conditional true alarm rate, the probability that a passenger who is a threat is detected given that they are assigned to class i. Likewise, the risk level for class i is interpreted as the conditional probability that class i contains a passenger who is a threat given that the passenger population contains a passenger who is a threat. It is assumed that the assessed threat values are defined in such a way that the risk level functions can be interpreted as this conditional probability (McLay et al. 2006a,b).

Security devices may require significant amounts of space in airport lobbies or terminals. MPSP does not explicitly incorporate space requirements of screening devices. However, since it is assumed that the device capacities are based on the necessary number of screening devices allowed by the physical space available, then space requirements are implicitly captured by MPSP.

McLay et al. (2006a) formulate MPSP as an integer program (IP) and provide its linear programming relaxation. They provide several theoretical results for MPSP. First, MPSP is NP-hard. Second, they show that the optimal passenger assignments can be determined by finding the optimal number of passengers to assign to each class. Third, the linear programming relaxation of MPSP has few fractional variables, given that the number of passengers is much larger than the number of devices and classes. Due to the structure of the linear programming relaxation, conditions exist under which the solution to the linear programming relaxation is integer for all device capacities, assessed threat values, and security levels. Therefore, in practice, it is not difficult to obtain optimal integer solutions to MPSP.

2.2 The Sequential Stochastic Multilevel Passenger Screening Problem

The Sequential Stochastic Multilevel Passenger Screening Problem (SSMPSP) generalizes MPSP to consider stochastic passenger arrivals. Both MPSP and SSMPSP consider a time interval in which passengers arrive at a security station. In MPSP, the passenger set is known whereas in SSMPSP, the passenger set is not known a priori—passengers check-in sequentially over the time interval, which is broken into T stages. Each stage represents a small period of time during which at most one passenger checks in. This assumption is reasonable if each stage represents a small period of time. There is a probability p of a passenger arriving during stage $t = 1, 2, \ldots, T$, with passengers arriving independently. In SSMPSP, the number of passengers N who arrive at the airport to be screened in a given period of time is a binomial random variable with parameters T and p, and with $E[N] = pT$.

The primary difference between MPSP and SSMPSP is that MPSP assigns N passengers to M classes simultaneously prior to the passengers checking in, while for SSMPSP, passengers check in sequentially, and each passenger is assigned to a class upon check-in (i.e., before the next passenger checks in) with the total number of passengers unknown. Note that the passengers may check in several hours before they arrive at the airport. Therefore, the time period when passengers are assigned to classes may be different than the time period when passengers arrive at the airport. Note that a prescreening system such as CAPPS determines which passengers are selectees using information provided by passengers at the point of ticket purchase, and hence, assigning passengers to classes when they check in is a reasonable assumption (TSA 2004). To differentiate the passengers in the static MPSP and stochastic SSMPSP models, the N passengers are indexed as $j = 1, 2, \ldots, N$ in MPSP (when passenger order is not a factor), and the time stages are indexed as $t = 1, 2, \ldots, T$ in SSMPSP (when passenger order is critical).

Several assumptions are needed to define SSMPSP. First, the passengers arrive at the airport to undergo screening during a given time period. The number of passengers as well as their assessed threat values are unknown prior to check-in. Each passenger has exactly one checked and one carry-on bag. The device capacities represent the number of passengers or bags that a device can screen in the time period. It is assumed that each passenger's assessed threat value becomes known upon check-in. Let $\boldsymbol{\alpha}(t)$ represent the random variable of the assessed threat value of passenger t, which is unknown prior to the passenger's check in. The probability density function of the assessed threat values, $f_\alpha(\alpha)$, is identical for all passengers $t = 1, 2, \ldots, T$. The assessed threat value of passenger t becomes known upon the passenger's check-in, taking on value $\alpha(t)$, $t = 1, 2, \ldots, T$. No passenger checks in during stage t with probability $1 - p$, resulting in $\alpha(t) = 0$. Therefore, SSMPSP can be modeled as if T passengers always arrive at the security station. It is assumed that there is a class with sufficient capacity to screen all passengers (call this class 1). In practice, a class would always be available with unlimited capacity and with minimal screening similar to the nonselectee class.

SSMPSP is formulated as a stochastic optimization problem, where the objective is to determine the optimal policy for assigning passengers to classes as they check in for their flights. A policy π defines a rule for assigning each passenger to a class, which may change after each passenger assignment is made. A policy may be deterministic (i.e., the policy always assigns a passenger to the same class given an identical time and state) or random (i.e., the policy may assign a passenger to different classes given an identical time and state). It may also be Markovian (i.e., the policy only depends on the current passenger and current state) or history-dependent (i.e., the policy depends on the passenger assignments of the previous passengers).

The Sequential Stochastic Multilevel Passenger Screening Problem (SSMPSP)
Given:

M, V, $c_1, c_2, \ldots, c_V$, d_{ik}, $i = 1, 2, \ldots, M, j = 1, 2, \ldots, V$, $L_1, L_2, \ldots, L_M$ as given by MPSP,

T stages,

p, the probability that a passenger checks-in during stage $t = 1, 2, ..., T$,

$f_\alpha(\alpha)$, a probability density function that describes the distribution of the passengers' assessed threat values, $0 < \alpha \leq 1$,

Objective: Denote passenger assignments for the N passengers to the M classes by $x(1), x(2), \ldots, x(T)$, where $x(t) \in \{1, 2, \ldots, M\}$, $t = 1, 2, \ldots, T$, represents the class to which passenger t is assigned. Find the policy π^* that determines the passenger assignments $x^{\pi*}(1), x^{\pi*}(2), \ldots, x^{\pi*}(T)$ such that each device type is operating within its capacity (i.e., $\sum_{i=1}^{M} d_{ik} |\{t : x^{\pi*}(t) = i\}| \leq c_k,\ k = 1, 2, \ldots, V$) and the expected total security (i.e., $E\{\sum_{t=1}^{T} L_i \alpha(t) | x^{\pi*}(t) = i\}$) is maximized.

Since SSMPSP has a sequential structure, it is natural to formulate SSMPSP as a Markov decision process (MDP). This MDP formulation also illustrates how the optimal policy is found. SSMPSP can be formulated as an MDP with $T+1$ stages using post-decision state variables, with the state in stage $t = 1, \ldots, T$ describing the system after the first t passengers have been assigned to classes, and $t = 0$ corresponds to the initial stage (Powell 2004).

Let S denote the set of states. Let state $\mathbf{s}(t) \in S$ represent the vector of remaining capacities after the first $t = 0, 1, \ldots, T$ passengers have been assigned to classes, where $\mathbf{s}(t) = (\bar{c}_1^t, \bar{c}_2^t, \ldots, \bar{c}_V^t)$, with $\bar{c}_k^t$ denoting the remaining capacity of device $k = 1, 2, \ldots, V$ after stage t. The initial state corresponds to the initial capacity, $\mathbf{s}(\mathbf{0}) = (c_1, c_2, \ldots, c_V)$. Therefore, $|S| = (c_1 + 1) \times (c_2 + 1) \times \cdots \times (c_V + 1)$. Moreover, define a vector $\mathbf{d}^i = (d_{i,1}, d_{i,2}, \ldots, d_{i,V})$ associated with class $i = 1, 2, \ldots, M$, that represents the vector of devices used by each class.

The set of actions available for assigning passenger t to a class is given by the subset of classes to which a passenger can be assigned. Given state $\mathbf{s}(t-1)$, the set of available classes in stage t is given by $X_t(\mathbf{s}(t-1)) \subseteq \{1, 2, \ldots, M\}$, $t = 1, 2, \ldots, T$, for all $\mathbf{s}(t-1) \in S$, with $i \in X_t(\mathbf{s}(t-1))$ if $\bar{c}_k^t - d_{ik} \geq 0,\ k = 1, 2, \ldots, V$. Let $x(t) \in X_t(\mathbf{s}(t-1))$ denote the security class to which passenger t is assigned. The transition probabilities $p(\mathbf{s}(t)|\mathbf{s}(t-1), x(t)),\ t = 1, 2, \ldots, N$, determine state $\mathbf{s}(t)$ given $\mathbf{s}(t-1)$ and $x(t)$. To define these probabilities for a deterministic policy, as each passenger is assigned to a class, the remaining capacities of the devices associated with that respective class decrease by one, with transition probability

$$p(\mathbf{s}(t)|\mathbf{s}(t-1), x(t)) = \begin{cases} 1, \text{ if } \mathbf{s}(t) = \mathbf{s}(t-1) - \mathbf{d}^{x(t)} \\ 0 \text{ otherwise} \end{cases}$$

for $t = 1, 2, \ldots, T$.

The objective function value of SSMPSP is determined by accruing a reward after each stage in the MDP. Define $r(\mathbf{s}(t-1), \boldsymbol{\alpha}(t), x(t))$ as the reward for assigning passenger t to class $x(t)$ given state $\mathbf{s}(t-1) \in S$, $t = 1, 2, \ldots, T$,

$$r(\mathbf{s}(t-1), \boldsymbol{\alpha}(t), x(t)) = L_{x(t)}\boldsymbol{\alpha}(t),\ t = 1, 2, \ldots, T.$$

Denote for passenger t a realized assessed threat value of $\alpha(t)$. Then the reward for assigning passenger t to class $x(t)$ becomes $L_{x(t)}\alpha(t)$ after passenger t checks in. The reward in stage t corresponds to the amount of security obtained from screening passenger $t = 1, 2, \ldots, T$. Note that if no passenger arrives at stage t, then $\alpha(t) = 0$ and $r(\mathbf{s}(t-1), 0, x(t)) = 0$.

The *expected total security* for SSMPSP is determined by policy π, which describes the decision rule for selecting an action (i.e., the class to which each passenger is assigned) in each state and at each stage. Let $S(t-1)$ represent the random variable corresponding to the state after $t-1$ passengers have been assigned to classes, and let $X_t^\pi(S(t-1))$ represent the random variable corresponding to the class to which passenger t is assigned given policy π. Given that the system is initialized in state $\mathbf{s}(0)$, then the expected total security for SSMPSP is defined as

$$E^\pi(\mathbf{s}(0)) = E^\pi \left\{ \sum_{t=1}^{T} r(S(t-1), \boldsymbol{\alpha}(t), X_t^\pi(S(t-1)))\, |S(0) = \mathbf{s}(0) \right\},$$

where $r(S(t-1), \boldsymbol{\alpha}(t), X_t^\pi(S(t-1)))$ is the random variable corresponding to the reward obtained in time period t in state $S(t-1)$ with decision $X_t^\pi(S(t-1))$ based on policy π. The objective of SSMPSP is to find the optimal policy π^* such that $E^{\pi^*}(\mathbf{s}(0)) = \max_\pi\{E^\pi(\mathbf{s}(0))\}$. This optimal policy is a deterministic Markov policy since the number of states is finite (Puterman 1994).

Define the value function in stage $t = 1, 2, \ldots, T$ as the optimal expected total security for assigning the remaining $T-t$ passengers,

$$V_{t-1}(\mathbf{s}(t-1)) = \mathrm{E}\left[\max_{x(t) \in X_t(\mathbf{s}(t-1))} \{ r(\mathbf{s}(t-1), \boldsymbol{\alpha}(t), x(t)) + V_t(S(t)|x(t)) \} \right],$$

for $t = 1, 2, \ldots, T$, where $S(t)|x(t)$ denotes the state after passenger t has been assigned to class $x(t)$. For SSMPSP, this is rewritten as

$$V_{t-1}(\mathbf{s}(t-1)) = \mathrm{E}\left[\max_{x(t) \in X_t(\mathbf{s}(t-1))} \left\{ L_{x(t)}\boldsymbol{\alpha}(t) + V_t(\mathbf{s}(t-1) - \mathbf{d}^{x(t)}) \right\} \right], \tag{2}$$

for $t = 1, 2, \ldots, T$, with boundary conditions

$$V_T(\mathbf{s}(\mathbf{T})) = 0,\ \mathbf{s}(\mathbf{T}) \in S, \tag{3}$$

which are also known as the *optimality equations.* The optimal policy π^*, which solves the optimality equations, can be found using dynamic programming.

The total security of a SSMPSP instance given realized assessed threat values $\alpha(1)$,$\alpha(2)$, $\ldots$,$\alpha(T)$ and their assignments $x(1), x(2), \ldots, x(T)$, is $\sum_{t=1}^{T} \alpha(t) L_{x(t)}$. The total security can be rescaled to be between zero and one, with

$$z^{SS} = \frac{\sum_{t=1}^{T} \alpha(t) L_{x(t)}}{\sum_{t=1}^{T} \alpha(t)},$$

which can be directly compared to (1), the total security of MPSP. Note that the total security for MPSP is normalized by a factor of $\sum_{j=1}^{N} \alpha(j)$, a constant parameter. However, this value is not known for SSMPSP until after the passengers have arrived.

3 The Sequential Stochastic Assignment Heuristic

Finding the optimal solution to the SSMPSP optimality equations (2) and (3) using dynamic programming is computationally intractable. The optimality equations can only be solved in a reasonable amount of computation time for small instances of SSMPSP. In order to solve large instances of SSMPSP in real-time, the Sequential Stochastic Assignment Heuristic (SSAH) is presented to efficiently obtain approximate solutions to SSMPSP.

In SSA, there are n people available for jobs, with the values of the people given by $v_1, v_2, ..., v_n$, which are treated as known constants. The n jobs arrive sequentially, with the values of the jobs being independently and identically distributed random variables $X_1, X_2, \ldots, X_n$ with cumulative density function $F(X)$. Therefore, there are n sequential stages in SSA, where in stage t, a job arrives and its value x_t becomes known, $t = 1, 2, ..., n$. The job is matched to a person who has not yet been assigned to an arriving job. Therefore, job t is assigned to person i_t, $t = 1, 2, ..., n$. The objective function is to maximize the expected total reward, given by the expectation of sum of the products of the values of the job and the person to whom it is matched, $E[\sum_{t=1}^{n} X_t v_{i_t}]$. Derman et al. (1972) provide the optimal policy for SSA.

SSA can also be trivially generalized to consider the number of jobs being different than the number of people. If the number of jobs m is less than the number of people n, then the m people with the highest values are assigned to a job. If $m \geq n$, then $m - n$ pseudo-people are created with value zero. If the probability that a job arrives at each stage is less than one, then the optimal policy can be trivially modified by modeling the event that a job does not arrive as a job arriving with value zero, which can be obtained by modifying $F(X)$. This generalization of SSA to consider a random number of jobs arriving is referred to as the Generalized SSA (GSSA).

SSMPSP is similar to GSSA. In this case, the optimal policy can be trivially modified by modeling the event that a job does not arrive as a job arriving with value zero, which can be obtained by considering a modified cumulative density function,

$$F^{GSSA}(X) = (1-p) + pF(X). \tag{4}$$

The policy defined by SSA assigns passengers to classes by using the main result from Derman et al. (1972). Applying this result requires that the classes be sorted such that $L_1 \leq L_2 \leq \ldots \leq L_M$. In SSMPSP, each passenger's assessed threat value becomes known upon check-in, with each assessed threat value taking on a value between zero and one. SSA breaks the assessed threat value range into a series of intervals at each stage. Passenger are assigned to classes based on the intervals in which their assessed threat values lie.

In order to use SSAH, the number of passengers assigned to each class must be determined before passengers check-in. Given that the number of passengers assigned to each class is known, $n_1, n_2, \ldots, n_M$, then $\sum_{i=1}^{M} n_i = T$ since SSMPSP is interpreted such that T passengers check-in. SSMPSP can be formulated as a particular case of GSSA. First create a set of passengers, with n_i passengers having value L_i, $i = 1, 2, \ldots, M$, corresponding to spaces available in the classes. The job arriving at each stage has value zero with probability $1 - p$, corresponding to no passenger checking in during stage t. With probability p, a job arrives, whose value follows the cumulative density function $F(X)$, corresponding to the assessed threat value distribution $f_\alpha(\alpha)$. Defined in this way, the objective function of GSSA maximizes the expected total security. The optimal policy can be generalized to solve GSSA by considering $F^{GSSA}(X)$ instead of $F(X)$.

There are two phases to SSAH: the preprocessing phase and the assignment phase (see Algorithm 1 pseudocode). In the preprocessing phase, the number of passengers to assign to each class is determined for the passengers who are expected to arrive in the time period. A series of intervals are created to construct the expected value of the assessed threat values of the T passengers arriving in the time period as follows, given the value of T and the probability density function of the assessed threat values $f_\alpha(\alpha)$, with cumulative density function $F_\alpha(a) = P(\boldsymbol{\alpha}(t) \leq a)$. Since a passenger may not arrive in each stage, then the resulting assessed threat distribution has cumulative density function $F_\alpha^{GSSA}(a) = (1-p) + pF_\alpha(a)$. A series of $t' = T - t + 1$ assignment intervals are constructed for passenger $t = 1, 2, \ldots, T$, where t' represents the number of passengers that have not yet been assigned to a class. Interval $j = 1, 2, \ldots, t'$ is defined by the boundaries $J_{t',j-1}$ and $J_{t',j}$, with $J_{t',0} \leq J_{t',1} \leq \ldots \leq J_{t',t'}$. These intervals can be determined by the recursion

$$J_{t'+1,j} = J_{t',j-1}F_\alpha^{GSSA}(J_{t',j-1}) + J_{t',j}[1 - F_\alpha^{GSSA}(J_{t',j})] + \int_{J_{t',j-1}}^{J_{t',j}} y \, dF_\alpha^{GSSA}(y), \quad (5)$$

with boundary condition $J_{t',0} = 0$ and $J_{t',t'} = 1$, $t' = 1, 2, \ldots, T$.

Define $E[\alpha^j]$ as the expected value of the jth smallest assessed threat value, given that there are T stages remaining, $j = 1, 2, ..., T$, as

$$E[\alpha^j] = J_{T,j-1}F_\alpha^{GSSA}(J_{T,j-1}) + J_{T,j}[1 - F_\alpha^{GSSA}(J_{T,j})] + \int_{J_{T,j-1}}^{J_{T,j}} y \, dF_\alpha^{GSSA}(y),$$

(Derman et al. 1972). The remaining parameters of the SSMPSP instance and $E[\alpha^1]$, $E[\alpha^2], \ldots, E[\alpha^T]$ define an instance of MPSP. The IP formulation of this MPSP instance is solved to give the *passenger partition*, the number of passengers assigned to each class, $n_1, n_2, \ldots, n_M$. Note that McLay et al. (2006a) indicate that the IP can be solved with little computational effort.

In the assignment phase, SSAH uses the intervals (5) to define a policy for assigning each passenger to a class upon check-in, given the current state. Passenger $t = 1, 2, ..., T$ falls into one of the t' intervals, and this interval is matched with one of the M classes, to which passenger t is assigned. Define $\tilde{n}_i = \sum_{j=1}^{i} n_j$ as the number of passengers assigned to class $i = 1, 2, \ldots, M$ and less secure classes, with $\tilde{n}_0 = 0$. When the first passenger arrives, $\alpha(1)$ becomes known and the passenger is assigned to position j_1^* (such that $J_{T,j_1^*-1} < \alpha(1) \leq J_{T,j_1^*}$) and to class x_1 (such that $\tilde{n}_{x_1-1} < j_1^* \leq \tilde{n}_{x_1}$). The objective function value z^{SSA}, n_{x_1}, and $\tilde{n}_{x_1}, \tilde{n}_{x_1+1}, \tilde{n}_M$ are updated. This procedure is repeated for the remaining $T-1$ passengers.

SSAH determines $M-1$ breakpoints at each stage in SSMPSP. Each passenger assignment is determined by the passenger's assessed threat value and the two breakpoints between which it lies. The breakpoints are dynamic, changing at each stage based on the previous passengers. The breakpoints for passenger t are $J_{t',\tilde{n}_1}, J_{t',\tilde{n}_2}, \ldots, J_{t',\tilde{n}_M}$. Passenger t is assigned to class 1 if the assessed threat value is between 0 and $J_{t',\tilde{n}_1}$ and class M if the assessed threat value is between $J_{t',\tilde{n}_M}$ and 1.

Given $J_{t',0}, J_{t',1}, \ldots, J_{t',t'}$ and $n_1, n_2, \ldots, n_M$, SSAH requires $\mathrm{O}(N(\log N + M))$ time in order to find position j_t^* for passenger $t = 1, 2, \ldots, N$ and to update $\tilde{n}_1, \tilde{n}_2, \ldots, \tilde{n}_M$ based on the passenger assignments. The time complexity of computing the intervals $J_{1,0}, J_{1,1}, \ldots, J_{N,N}$ depends on the method used for computing the integrals in (5). However, the intervals do not need to be recomputed each time SSAH is executed if $f_\alpha(\boldsymbol{\alpha})$ remains constant—the intervals can be saved for future access. Solving the IP of the corresponding MPSP instance is computationally intractable. However, optimal solutions to MPSP instances have been observed to be obtained efficiently in practice (McLay et al. 2006a).

Theorem 1 outlines the condition under which SSAH provides the optimal policy. Note that the condition in Theorem 1 can only be checked after all of the passengers have checked in and their assessed threat values have become known. Therefore, the condition given in Theorem 1 only has value retrospectively.

Theorem 1 *If the passenger partition $n_1, n_2, \ldots, n_M$ given $E[\alpha^1], E[\alpha^2], \ldots, E[\alpha^T]$ is identical to the SSMPSP passenger partition when the true assessed threat values $\alpha(1)$, $\alpha(2)$, $\ldots, \alpha(T)$ are assumed to be known a priori, then the policy defined by SSAH is the optimal policy for SSMPSP.*

Proof: The proof is a direct application of Theorem 1 in Derman et al. (1972). □

4 Computational Results

This section reports computational results for an instance of SSMPSP that incorporates flight schedule and passenger volume data extracted from the Official Airline Guide (OAG) for the domestic flights of a single airline carrier at particular airport stations in the United States. The passenger set is based on data extracted from the OAG, which includes the set of flights, the number of available seats on each flight, and the departure time of each flight. The flights departing from a single terminal at a hub airport are considered. Since some flights depart from the terminal infrequently, and including

Algorithm 1 The Sequential Stochastic Assignment Heuristic

Comment: Preprocessing Phase
Compute $J_{t',0}, J_{t',1}, \ldots, J_{t',t'}$ for $t' = 1, 2, \ldots, T$ based on (5)
Initialize objective function value $z^{SSA} \leftarrow 0$
Solve an IP instance of MPSP with $E[\alpha^1], E[\alpha^2], \ldots, E[\alpha^T]$ to obtain the number of passengers to assign to each class $n_1, n_2, \ldots, n_M$
Compute $\tilde{n}_i = \sum_{j=1}^{i} n_j$, $i = 1, 2, \ldots, M$
Comment: Assignment Phase
for $t \leftarrow 1$ to T **do**
  $t' \leftarrow T - t + 1$
  $\alpha(t)$ becomes known as passenger t checks in
  Determine position $j_t^* \in \{1, 2, \ldots, t'\}$ such that $J_{t',j_t^*-1} < \alpha(t) \leq J_{t',j_t^*}$
  Determine class $x(t) \in \{1, 2, \ldots, M\}$ such that $\tilde{n}_{x(t)-1} < j_t^* \leq \tilde{n}_{x(t)}$
  $\tilde{n}_i \leftarrow \tilde{n}_i - 1$ for $i = x(t), x(t)+1, \ldots, M$.
  $z^{SSA} \leftarrow z^{SSA} + L_{x(t)}\alpha(t)$
end for
$z^{SSA} \leftarrow z^{SSA} / (\sum_{t=1}^{T} \alpha(t))$
return z^{SSA}, $x(1), x(2), \ldots, x(T)$

these flights would inaccurately increase the peak number of passengers, only flights with a frequency of at least three flights per week (twelve per month) are considered.

To find the peak number of passengers in the time window, passengers are assumed to arrive randomly according to a triangular distribution between 40 and 120 minutes prior to the departure time of each flight. This arrival interval is recommended by airlines for domestic flights, with 30 minutes prior to departure being the latest check-in time (e.g., see www.nwa.com, www.ual.com, www.aa.com, or www.delta.com). Moreover, each flight is assumed to have an enplanement rate of 80% (i.e., the number of passengers divided by the number of available seats). The data set is chosen based on the largest expected number of arriving passengers in a one hour minute window, resulting in $N = 882$ passengers.

The device types are chosen based on current devices and procedures used to screen passengers in commercial airports in the United States, with each class using multiple devices. Five devices are considered; two for screening checked baggage, and three for screening passengers and their carry-on baggage. The three devices for screening passengers and carry-on baggage are assumed to be a combination of metal detectors and X-ray machines, labeled D1, trace portals, labeled D2, and detailed hand searches, labeled D3. The two devices for screening checked baggage are explosive detection systems (EDSs), labeled D4, and explosive trace devices, labeled D5. It is assumed that all passengers and carry-on baggage are screened by metal detectors and X-ray devices (D1) and all checked baggage is screened by EDSs (D4).

The combination of metal detectors and X-ray machines is considered as one device for two reasons. First, passengers and carry-on baggage are screened together, with passengers always in the presence of their carry-on baggage. For this reason, either the metal detector or the X-ray machine is the screening bottleneck, as passengers must either wait for their carry-on baggage screening to be completed or carry-on baggage must wait for passenger screening to be completed. Therefore, the capacity (i.e., throughput) of these two devices is the same. The second reason is that these devices work together to remove threat items. Once the screening procedures have been completed, threat passengers have access to threat items in carry-on baggage after screening, whereas they do not have access to threat items in checked baggage.

Table 1 summarizes the false clear rates and capacities associated with the devices. The device values are estimated using information available in the public domain (see Butler and Poole 2002; Virta et al. 2003; McLay et al. 2006b). It is assumed that the devices for passengers operate independently of the devices for checked-baggage screening. This is reasonable since the screening for these two classes of devices is often done in different areas in airports, and the system response for checked baggage (i.e., either alarm or clear) does not depend on the system response to passenger and carry-on baggage.

Table 1: Security Device Data

Label	Device Type	False Clear	Units/hour	Units/10 minutes
D1	Metal Detector/X-ray Machine	0.20	120	20
D2	Trace Portal	0.15	90	15
D3	Detailed Hand Search	0.1	30	5
D4	EDS	0.12	150	25
D5	Explosive Trace Devices	0.15	30	5

Using these devices, six classes can be defined as follows. First, passengers can be screened in three ways: by D1, by D1 and D2, or by D1 and D3. Likewise, checked baggage can be screened in two ways: by D4 or by both D4 and D5. The resulting six classes enumerate all the possible ways of how passengers and baggage can be screened. The security level of the classes are the average of the true alarm rates associated with passenger screening and checked-baggage screening. Note that a different metric could be used to determine the security level of a class based on the false clear rates associated with both types of devices. Without loss of generality, assume that a passenger (checked baggage) is screened by D1 before D2 or D3 (D4 before D5) when screened by two devices.

When passengers or checked baggage are screened by both devices, it is assumed that the system gives an alarm response if the passenger or bag triggers an alarm at either device in the system. Therefore, the system gives a clear response only if both devices give a clear response. First, define T as the event that a passenger is a threat, and Ci as the event that a passenger is cleared by device Di, $i = 1, 2, 3, 4$. Kobza and Jacobson (1996) define the dependence between two screening devices for passengers as

$$\varepsilon^T_{C2|C1} = P(C2|C1 \cap T) - P(C2|T), \tag{6}$$

with the lower and upper bounds for $\varepsilon^T_{C2|C1}$ being

$$-P(C2|T) \leq \varepsilon^T_{C2|C1} \leq \min\left(P(C2|C1 \cap T),\ \frac{P(C2|T)P(\bar{C}1|T)}{P(C1|T)}\right). \tag{7}$$

Then, the false clear rate when being screened by both devices is

$$P^{D1,D2}_{FC} = P^{D1}_{FC}(P^{D2}_{FC} + \varepsilon^T_{C2|C1}). \tag{8}$$

Devices D1 and D2 operate independently if $\varepsilon^T_{C2|C1} = 0$. If knowing that a threat is cleared by device D1 increases (decreases) the probability of a false clear by device D2, then there is a positive (negative) dependence and $\varepsilon^T_{C2|C1} > (<) 0$. The dependence between D1 and D3 can be computed by substituting C3 for C2 and D3 for D2 in (6), (7), and (8).

The dependence between the two screening devices for checked baggage $\varepsilon^T_{C4|C3}$ is defined by substituting D4 for D1 and D5 for D2 in (6), (7), and (8). Based on the bounds in (7), $\varepsilon^T_{C2|C1} = \varepsilon^T_{C4|C3} = 0.1$ for the scenarios considered. This results in a conservative relationship between the screening devices in which the false clear rate associated with the combination of device types is greater than when the device types are assumed to operate independently. The security level is computed as the true alarm rate, $P^{D1,D2}_{TA} = 1 - P^{D1,D2}_{FC}$.

Table 2 contains class information, the devices used by each class, and the security level associated with each class. If device type k is (not) used by class i, then a one (zero) is listed in the Class i row and the Devices Used column k in Table 2. The classes are sorted such that $L_1 < L_2 < \ldots < L_6$. Since all passengers are screened by D1 and D4, these two devices are omitted from Table 2.

Table 2: Class data

	Devices Used (d_{ik})			
Class (i)	D2 ($k = 2$)	D3 ($k = 3$)	D5 ($k = 5$)	Security Level (L_i)
1	0	0	0	0.84
2	0	0	1	0.885
3	1	0	0	0.915
4	0	1	0	0.92
5	1	0	1	0.96
6	0	1	1	0.965

An example with a total of eight scenarios is considered by constructing eight device capacities. There are two levels of available capacities for each of the device types, resulting in eight device capacities when all combinations of available capacities are considered across the four device types. Table 3 summarizes the eight device capacity levels available in the time window. Since all passengers are screened by devices D1 and D4, it is assumed that the capacities of these devices are 3600.

Table 3: Device Capacity Levels

Capacity Level	D2	D3	D5
1	90	120	150
2	90	120	300
3	90	360	150
4	90	360	300
5	360	120	150
6	360	120	300
7	360	360	150
8	360	360	300

The assessed threat value distribution assumes that most passengers can be cleared of being a threat and hence, have low assessed threat values. This is consistent with what is known in the public domain about how Secure Flight works (US GAO 2005). Recall that $\boldsymbol{\alpha}(t)$ denotes the random variable that represents the assessed threat value of passenger $t = 1, 2, \ldots, T$ prior to passenger t checking in. The assessed threat values only become known upon passenger check-in. Note that for all three types of distributions, the same distribution is used for all values of t. The assessed threat distribution is an exponential random variate with mean 1/16, where values greater than one are discarded when assessed threat values are randomly generated. A randomly generated assessed threat value is greater than one with a probability of approximately 10^{-7}. This distribution results in approximately 80% of the assessed threat values being less than 0.1.

The optimal policy for SSMPSP that maximizes the expected total security can be found by solving the optimality equations (2) and (3). However, finding the optimal policy in this manner is too computationally intensive for the scenarios considered. SSAH provides the optimal policy under the condition in Theorem 1. However, this condition can only be verified retrospectively, after all passengers have checked in.

SSAH found solutions for thirty replications of each of the eight scenarios, where new assessed threat values were randomly generated for each of the thirty replications. SSAH has two phases: the preprocessing phase and the assignment phase. Since the preprocessing phase does not depend on the realized assessed threat values, the preprocessing phase did not execute for each replication. The SSAH intervals (5) were computed once (using Matlab) for each assessed threat distribution (i.e., a total of three sets of SSAH intervals were computed for the 24 scenarios). The three sets of SSAH intervals were used to create 24 IPs, one corresponding to each scenario. The 24 IPs were solved in the preprocessing phase using CPLEX 9.0 on a Sun Blade 1500 with a 1.5 GHz UltraSPARC IIIi processor. Each IP was solved in fewer than 2 CPU seconds. For each of the eight scenarios, thirty replications of the assignment phase was executed using Matlab. The Matlab computations were performed on a 3.19 GHz Pentium IV processor with 3.5 GB of RAM. The assignment phase required no more than 42 seconds for SSAH for all thirty replications of each scenario.

Note that if the passenger partition for a SSMPSP solution is the same as the partition of the corresponding MPSP scenario when the assessed threat values are assumed to be known a priori, then SSAH is the optimal policy. However, the subsets of passengers assigned to classes in the SSMPSP solution may be different than those of the corresponding MPSP solution since a passenger's assignment depends on the order of arrival for check-in in SSMPSP.

Table 4 summarizes the integer programming and SSAH solutions for all eight scenarios. The partitions for all thirty replications of each of the eight scenarios are the same as the corresponding MPSP partitions when the assessed threat values are assumed to be known a priori. Therefore, SSAH is the optimal policy for all eight scenarios considered. The small variation in the SSAH solution values suggests that SSAH is robust in assigning passengers to classes.

Table 4: Integer Programming and SSAH Solutions

	Passenger Partition						Average	Standard
Capacity Level	n_1	n_2	n_3	n_4	n_5	n_6	Value	Deviation
1	3390	0	60	0	30	120	0.906	0.002
2	3300	90	0	0	90	120	0.917	0.002
3	3150	0	90	210	0	150	0.929	0.001
4	3150	0	90	60	0	300	0.939	0.001
5	3120	0	330	0	30	120	0.928	0.002
6	3120	0	180	0	180	120	0.939	0.001
7	2880	0	360	210	0	150	0.938	0.001
8	2880	0	360	60	0	300	0.949	0.001

To quantify the worst-case performance of SSAH, two extreme cases are considered: passengers who check in with increasing assessed threat values and passengers who check in with decreasing assessed threat values. These two cases are compared to the original scenarios when passengers are assumed to arrive at random. The case with increasing (decreasing) assessed threat values represents the scenario when lower-risk passengers arrive first (last), and higher-risk passengers arrive last (first). The case with increasing assessed threat values provides insight on whether a large number of lower-risk passengers arriving early in the time period use scarce screening resources associated with class 6, which are ideally reserved for the higher-risk passengers who arrive later. The case with decreasing assessed threat values provides insight on whether a large number of higher-risk passengers who arrive early in the time period are assigned to less secure classes in order to keep resources associated with class 9 available for higher-risk passengers that are expected to arrive later in the time period.

One goal of multilevel passenger screening systems is to direct more security resources toward passengers perceived as higher-risk (National Research Council 2002). However, the objective function value captures the expected total security, not the accuracy of SSAH in assigning higher-risk passengers to more secure classes. Since there are few higher-risk passengers, their effect on the objective function value may be diluted by the more numerous lower-risk passengers. SSAH indirectly determines $M-1$ breakpoints during each stage $t = 1, 2, \ldots, T$, with passengers assigned to classes based on which breakpoints their assessed threat values lie between. The breakpoints change after every passenger arrival.

The scenario with capacity level 6 and the assessed threat value is used to illustrate how SSA assigns passengers to classes. SSAH assigns 120 passengers to class 6 (see Table 4). Ideally, the 120 passengers with the highest assessed threat values should be assigned to class 6. Note that there are only two unique breakpoints between the classes, since only four of the classes are used in the solutions.

Figure 1 shows the average breakpoints over the thirty replications to the scenario with capacity level 6, when passengers randomly arrive for check-in. The breakpoints are relatively constant for most of the time period. Both breakpoints appear to decrease at the end of the time period, which indicates that passengers who check in at the end of the time period are more likely to be assigned to more secure classes. In all SSAH replications, it was observed that the last passenger to check in is always assigned to class 6, regardless of the passenger's assessed threat value. The SSAH policy appears to save available space in class 6 for high-risk passengers that may check in at the end of the time period. Therefore, SSAH provides a policy that is difficult for high-risk passengers to game—extremely high-risk passengers are always assigned to class 6 in the scenarios considered.

Figure 2 shows the average breakpoints across thirty replications for the cases with increasing and decreasing assessed threat values. The breakpoints increase near the end of the time period for case with increasing assessed threat values. This indicates that some high-risk passengers would be screened with the least secure class, and hence, it is possible for terrorists to game the system by altering when passengers check-in. However, this would be difficult to achieve in practice, particulary at a hub airport with a large number of passengers. SSAH is more robust in assigning high-risk passengers to class 6 for the case with decreasing assessed threat values. This result is counterintuitive since the case with random passenger arrivals suggests that SSAH is most effective in screening high-risk passengers who check in at the end of the time period, not at the beginning, while in the case with decreasing assessed threat values, the high-risk passengers who check in at the beginning of the time period are consistently assigned to class 6.

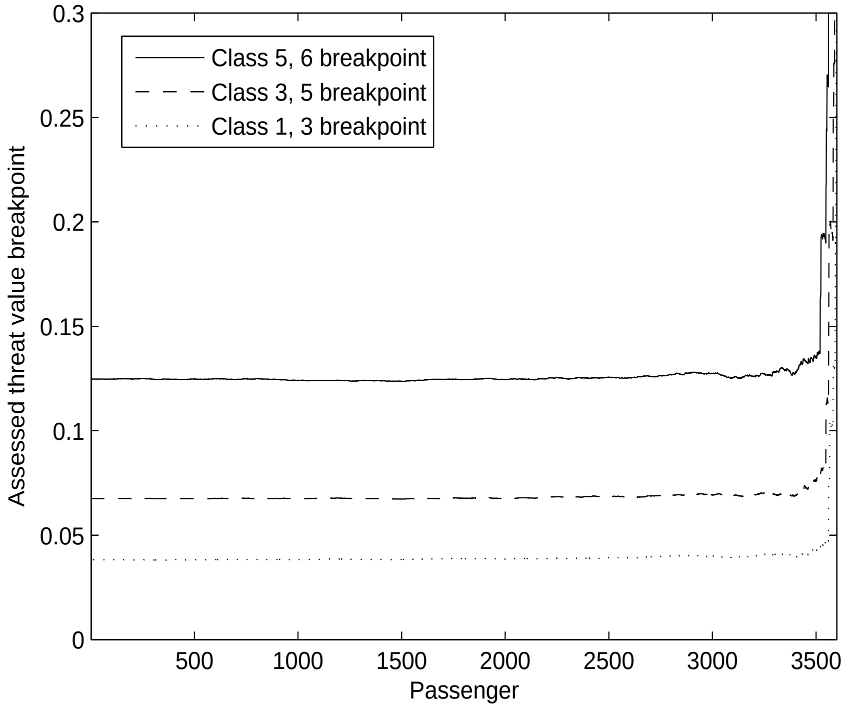

Figure 1: Average Breakpoints with Capacity Level 6 with Random Arrivals

5 Conclusions

Passenger screening is a critical component of any aviation security system operation. This paper introduces SSMPSP, which models stochastic passenger and baggage screening systems. SSMPSP, an extension of MPSP that considers passengers arriving over time, is formulated as a Markov Decision Process, where the optimal policy can be computed by dynamic programming. The SSA heuristic is presented to provide approximate solutions to SSMPSP in real-time. A condition is provided under which SSAH yields the optimal solution. SSAH obtains solutions to SSMPSP in real-time, and this suggests that SSAH could be part of a tool used by the TSA in assigning passengers to classes in real-time.

Data extracted from the OAG was used to construct the passenger set for the SSMPSP example with eight capacity levels. It was verified retrospectively that SSAH

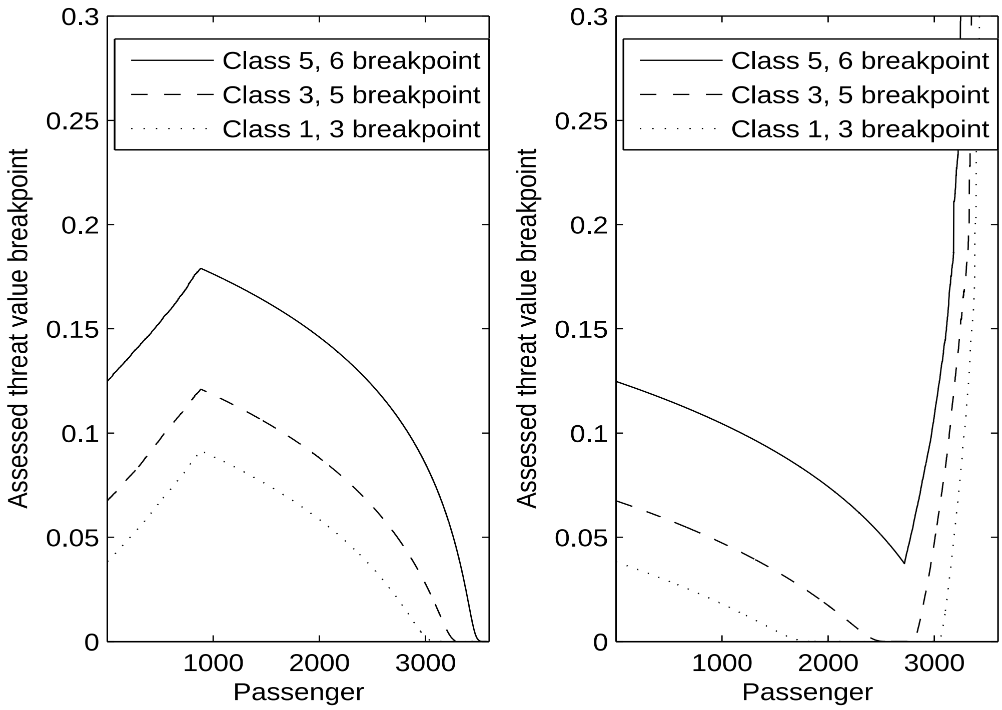

Figure 2: Average Breakpoints with Capacity Level 6 when Passengers Arrive at the Beginning and End of the Time Interval

provided the optimal policy for all eight scenarios. Analysis suggests that SSAH is almost certain to assign extremely high-risk passengers to the most secure class, regardless of when these passengers check in. An analysis of two extreme cases, passengers arriving at the beginning and the end of the time interval, suggests that there may be scenarios in which SSAH can be gamed if the passenger check in pattern is skewed.

There are several possible directions to extend the research results presented. First, incorporating screening costs into the model can give insight into the design of cost-effective screening strategies. Second, understanding how these models can be manipulated by terrorists wishing to be screened by less secure classes is of interest. In particular, modeling how terrorists could take advantage of how a heuristic such as SSAH operates in real-time as well as manipulating the passenger pool (e.g., flooding the airport with decoys) may provide the impetus for designing robust heuristics and algorithms. Work is in progress to address these extensions.

Acknowledgements

This research was supported in part by the National Science Foundation (DMI-0114499), the Air Force Office of Scientific Research (FA9550-07-1-0232), and the U.S. Department of Homeland Security (2008-DN-077-ARI001-02). The computational work was done at Virginia Commonwealth University and in the Simulation and Optimization Laboratory housed within the Department of Computer Science at the University of Illinois. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views or official policies, either expressed or implied, of the National Science Foundation, the Air Force Office of Scientific Research, or the U.S. Department of Homeland Security.

References

Babu, V. L. L., R. Batta, and L. Lin (2006). Passenger grouping under constant threat probability in an airport security system. *European Journal of Operational Research 168*, 633–644.

Barnett, A. (2001). The worst day ever. *OR/MS Today 28*(6), 28–31.

Barnett, A. (2004). CAPPS II: The foundation of aviation security? *Risk Analysis 24*(4), 909–916.

Barnett, A., M. Abraham, and V. Schimmel (1979). Airline safety: Some empirical findings. *Management Science 25*(11), 1045–1056.

Barnett, A. and M. K. Higgins (1989). Airline safety: The last decade. *Management Science 35*(1), 1–21.

Barnett, A., R. Shumsky, M. Hansen, A. Odoni, and G. Gosling (2001). Safe at home? An experiment in domestic airline security. *Operations Research 49*(2), 181–195.

Butler, V. and R. W. Poole, Jr. (2002). Rethinking checked-baggage screening. Technical report, Reason Public Policy Institute, Public Policy No. 297, Los Angeles, CA.

Chakrabarti, S. and A. Strauss (2002). Carnival Booth: An algorithm for defeating the computer-assisted passenger screening system. *First Monday 7*(10). Accessed at www.firstmonday.org on 23 October 2003.

Czerwinski, D. and A. Barnett (2006). Airlines as baseball players: Another approach for evaluating an equal-safety hypothesis. *Management Science 52*(9), 1291–1300.

Jacobson, S. H., J. M. Bowman, and J. E. Kobza (2001). Modeling and analyzing the performance of aviation security systems using baggage value performance measures. *IMA Journal of Management Mathematics 12*(1), 3–22.

Jacobson, S. H., L. A. McLay, J. E. Kobza, and J. M. Bowman (2005). Modeling and analyzing multiple station baggage screening security system performance. *Naval Research Logistics 52*(1), 30–45.

Jacobson, S. H., L. A. McLay, J. L. Virta, and J. E. Kobza (2005). Integer program models for the deployment of airport baggage screening security devices. *Optimization and Engineering 6*(3), 339–359.

Jacobson, S. H., J. E. Virta, J. M. Bowman, J. E. Kobza, and J. J. Nestor (2003). Modeling aviation baggage screening security systems: A case study. *IIE Transactions 35*(3), 259–269.

Kobza, J. E. and S. H. Jacobson (1996). Addressing the dependency problem in access security system architecture design. *Risk Analysis 16*(6), 801–812.

Kobza, J. E. and S. H. Jacobson (1997). Probability models for access security system architectures. *Journal of the Operational Research Society 48*(3), 255–263.

Martonosi, S. E. (2005). *An Operations Research Approach to Aviation Security.* Ph. D. thesis, Massachusetts Institite of Technology, Cambridge, MA.

Martonosi, S. E. and A. Barnett (2006). How effective is security screening of airline passengers? *Interfaces 36*(6), 545–552.

McLay, L. A., S. H. Jacobson, and J. E. Kobza (2006). A multilevel passenger prescreening problem for aviation security. *Naval Research Logistics 53*(3), 183–197.

McLay, L. A., S. H. Jacobson, and J. E. Kobza (2007). Integer programming models and analysis for a multilevel passenger screening problem. *IIE Transactions 39*(1), 73–81.

Mead, K. M. (2002). Key issues concerning implementation of the Aviation and Transportation Security Act. *Office of Inspector General, Department of Transportation, Washington, D.C. Report Number CC-2002-098.*

Mead, K. M. (2003). Aviation security costs, Transportation Security Administration. *Office of Inspector General, Department of Transportation, Washington, D.C. Report Number CC-2003-066.*

Nie, X., R. Batta, C. G. Drury, and L. Lin (2006). Passenger grouping with risk levels in an airport security system. Technical report, University of Buffalo, Buffalo, NY.

Nikolaev, A., S. H. Jacobson, and L. A. McLay (2007). A sequential stochastic security system design problem for aviation security. *Transportation Science 41*(2), 182–194.

Poole, Jr., R. W. and G. Passantino (2003). A risk-based airport security policy. Technical report, Reason Public Policy Institute, Public Policy No. 308, Los Angeles, CA.

Powell, W. B. (2004). Approximate dynamic programming for asset management: Solving the curses of dimensionality. Technical report, Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ.

United States Government Accountability Office (March 2005). Secure Flight development and testing under way, but risks should be managed as system is further developed. Technical Report GAO-05-356, Washington, D.C.

United States Transportation Security Administration (2006). Privacy frequently asked questions. Web site, Department of Homeland Security, Washington, D.C. (available at www.tsa.gov/research/privacy/faqs.shtm, accessed on 11 December 2006).

United States Transportation Security Administration (26 August 2004). “Secure Flight” to be tested before year’s end. Press release, Department of Homeland Security, Washington, D.C.

Virta, J. E., S. H. Jacobson, and J. E. Kobza (2002). Outgoing selectee rates at hub airports. *Reliability Engineering and System Safety 76*(2), 155–165.