Proceedings of PRADS2016
4^{th}-8^{th} September, 2016
Copenhagen, Denmark

A ship egress analysis method using spectral Markov decision processes

Austin A. Kana[1] and David J. Singer[2]

[1]Delft University of Technology, the Netherlands
[2]University of Michigan, USA

Abstract

This paper introduces a means of performing a ship egress analysis by applying eigenvalue analysis to the ship-centric Markov decision process (SC-MDP) framework. This method focuses on how people egress, the decisions they make under uncertainty, and the interaction between the individuals and the layout of the vessel. The objective is to understand the implications of uncertain decision making of people on general arrangement design. One metric is introduced defined as the ratio between the largest eigenvalue and the second largest. This decision metric is used to identify and quantify changes in decisions, as well as to help identify system attributes driving those changes in decisions. A case study is presented showing the utility of this method on a ship egress problem. Sensitivity studies are performed examining the affect of uncertainty and rewards on individuals' decision making behavior.

Keywords

Ship design; decision making; egress analysis; Markov decision process; eigenvalue analysis

Introduction

Understanding evacuation patterns and egress routes is arguably one of the most important aspects of ship design affecting safety of those on board (Guarin et al., 2014). Emergency situations such as fires, flooding, damage due to collisions, or even ballistic damage, happen on vessels. Understanding how individuals on the vessel react and move about to safety zones, muster points, or life boats during these situations is important to save lives and minimize injuries. The IMO has recognized evacuation and egress analysis as an important aspect that needs to be regulated for the safety of mariners (IMO, 2007).

Two types of methods have been developed to approach this problem: those focused on analyzing the physical layout of the vessel (Andrews et al., 2012; Rigterink et al., 2014; van Bruinessen et al., 2011), and those focused on evacuation patterns of the individuals on board (Guarin et al., 2014; Qiao et al., 2014; Vanem and Skjong, 2006). Methods focused on the layout are defined as solution-centric by this paper, meaning they are focused on the physical layout of the rooms themselves. Methods focused on evacuation routes, however, are typically computationally expensive, and can involve multi-agent simulations (Guarin et al., 2014), or optimization methods (Qiao et al., 2014).

Both type of analyses require the full vessel layout as well as the physical distribution of the crew throughout the vessel in order to run detailed discrete event simulations to study evacuation routes (Rigterink et al., 2014). Two major problems arise with this. First, during early stages of design, little information is known about the details of the vessel layout, and second, a new discrete event simulation is required for each simulation involving a different distribution of individuals. This makes the problem nearly intractable.

The complexity of this problem grows with vessel size and complexity. The problem is made difficult by various passenger populations (such as able bodied seamen, children, or handicapped individuals), the number of decks and passageways, and the multitude of ways emergency situations may start or percolate through the vessel. Also, paid crew who know the layout well will likely react differently compared to passengers on a cruise ship who are new to the vessel. For example, in 1915, the passenger vessel, the S.S. Eastland, capsized in the Great Lakes when the passengers all rushed to one side, causing the deaths of nearly 850 individuals (Eastland Memorial Society, 2015). Thus, for a full detailed analysis, advanced methods are necessary.

While some propose using more advanced, computationally expensive evacuation models in earlier stages of design (Vanem and Skjong, 2006), the IMO (2007) has recognized the importance of using simplified methods in the concept stage. This has led to the increased prominence of analyses aimed at evacuation routes and egress patterns in preliminary stages of design (Casarosa, 2011; Guarin et al., 2014). However, finding the proper balance between computational expense, analysis time, and model fidelity remains difficult.

This paper proposes a method that is focused on understanding the decision making of individuals on the vessel as they evacuate. In early stage design this method will prove to be more tractable than having to fully enumerate all layouts and all crew combinations to understand egress patterns. This method is focused on understanding the impact that uncertainty and pain (possibly from the inhalation of smoke) may have on the decisions individuals make. To do this, eigenvalue analysis applied to the ship-centric Markov decision process (SC-MDP) framework is proposed to understand the implications of evacuation decision making as it pertains to general arrangements design.

The Ship-Centric Markov Decision Process

The SC-MDP framework is defined as applying Markov decision processes (MDPs) to ship design and decision making. MDPs were first applied to ship design by Niese and Singer (2013) as a means to generate and analyze predictive time domain design data. MDPs are a mathematical model developed in the 1950's to solve dynamic decision-making problems under uncertainty (Puterman, 2005). They are a state-based model that represent uncertain systems, can differentiate various decision making scenarios, and can handle temporal system variations. MDPs have been used across a wide variety of other disciplines, including: robot navigation (Russell and Norvig, 2003), financial, economic, or portfolio management (Sheskin, 2011), inventory management, scheduling of industrial system maintenance and replacement, and even behavioral ecology (Puterman, 2005).

From a ship design perspective, MDPs allow for analysis of the physical product, the sequential life-cycle decisions associated with that product, and the projected life-cycle expected utility of the products and decisions. From a decision space perspective, the advantages include an explicit model of the system uncertainties and environmental risks present throughout the system's life-cycle, the ability to analyze dynamic operating profiles and external environments, and the ability to enable active management and decision making. Previous research using the SC-MDP framework include analysis of ballast water treatment methods (Niese and Singer, 2013, 2014), designing for the Energy Efficiency Design Index (Niese et al., 2015), and design for evolving Emission Control Area regulations (Kana et al., 2015).

Despite these advantages, the current SC-MDP framework has several limitations. First, the size and scope of the results can be large, and in some cases can become overwhelming for the decision maker. Second, MDPs can only be used to identify changes in decision behavior through the use of simulations and extensive sensitivity studies (Niese et al., 2015; Tan and Hartman, 2011). The decision maker must manually backtrack through the model to find where these decisions change. There is also no way to quantify how much a single change in the decisions has on the process as a whole. This research proposes to address these limitations by introducing eigenvalue analysis techniques to the framework itself.

Eigenvalue Spectral Analysis

Spectral analysis applied to the SC-MDP framework is introduced as a means to elicit decision making insight by quantifying the impact of changes in decisions. The applicability of spectral methods to analyze ship design and decision making in the SC-MDP framework is vast. It can be used to help identify the magnitude of the relationships and inter-dependencies between the various attributes of the system. It can also be used to help identify those secondary, tertiary, and and weaker inter-dependencies in the system that may not be noticeable by other methods. Similar eigenvalue techniques have been used in autoregressive models to study principal oscillation patters, excitation, and damping times (Neumaier and Schneider, 2001). This type of analysis is beneficial to designers by allowing them to focus their efforts systematically on the important factors on the design. While no single framework can capture all aspects of design decision making (Reich, 1995; Seram, 2013), the objective here is to provide a unique perspective on engineering decision making to help improve understanding and design.

Methods

The Markov Decision Process

The classic MDP has four parts:

1. a set of states, S, where an agent can exist
2. a set of actions, A, available to the agent
3. a set of probabilities, T, of transitioning from one state to another after following a given action
4. a set of rewards, R, received after performing a given action, a, and transitioning to a new state

The objective is to identify the sequence of actions that maximizes the cumulative, long term expected utility of the system. This sequence of actions identifies the set of decisions the agent should take during each decision epoch (Puterman, 2005). An example three state, two action MDP is given in Table 1 and Table 2. Table 1 displays the transition matrices for the various possible actions, where the probability of transitioning from state s_i to state s'_j is given. Note the transition probabilities and rewards may vary between different actions.

Table 1. Sample transition matrices for a 3 state, 2 action MDP. Each action is represented by its own transition matrix.

	Action 1				Action 2		
	s'_1	s'_2	s'_3		s'_1	s'_2	s'_3
s_1	0	0.7	0.3	s_1	0	0	1
s_2	0.2	0.8	0	s_2	0	1	0
s_3	1	0	0	s_3	0.1	0.1	0.8

Table 2 is the associated reward matrix, which outlines the reward obtained after taking a given action, a_i, and landing in a new state, s'_j. Essentially, MDPs can be thought of as a series of action dependent Markov chains with rewards (Sheskin, 2011), where each action can be represented by its own transition matrix.

Table 2. Sample reward matrix for a 3 state, 2 action MDP. The entries indicate the reward received for taking action, a_i, and landing in state, s'_j.

	a_1	a_2
s'_1	5	-1
s'_2	1	3
s'_3	0	6

The expected utility of the MDP can be obtained via Equation 1, known as the Bellman equation, where U is the expected utility, γ is the discount factor, and the other variables are defined previously. The set of decisions, π, is found by Equation 2 (Russell and Norvig, 2003).

$$U(s) = R(s) + \gamma \max_a \sum_{s'} T(s,a,s')U(s') \qquad (1)$$

$$\pi(s) = \arg\max_a \sum_{s'} T(s,a,s')U(s') \qquad (2)$$

The common output for displaying the decisions is with a decision matrix which provides the preferred action for each state for each decision epoch. An epoch is defined as an instance when the agent must make a decision. Epochs can represent any such decision making period, such as a time step, or physical movement by an individual. The decision matrix for the sample system above is given in Table 3. To read the decision matrix, the decision maker identifies the preferred action in the table that corresponds to the system state and epoch. For instance, Action 2 is preferred if the system were in State 2 during Epoch 2. Note the decisions may change through time, as is the case between Epoch 4 and Epoch 5 for State 1.

Table 3. The decision matrix for the system above. For example, Action 2 is best for State 2 and Epoch 2.

	State 1	State 2	State 3
Epoch 1	Action 1	Action 2	Action 2
Epoch 2	Action 1	Action 2	Action 2
Epoch 3	Action 1	Action 2	Action 2
Epoch 4	Action 1	Action 2	Action 2
Epoch 5	Action 2	Action 2	Action 2

The decision matrix can be a beneficial way of displaying information. First, it provides a road-map for the decision maker by displaying the optimal decision for each state throughout the process lifecycle. Second, it can be used to simulate different decision scenarios to discern differences in seemingly similar situations (Niese and Singer, 2013).

In order to overcome the limitations outlined in the *Introduction*, this research proposes eigenvalue spectral analysis. To perform this, a different formulation of the output is necessary; one that involves using the decision matrix to develop a series of transition matrices to represent the dynamics of the system. Eigenvalue spectral analysis is then performed in these transition matrices.

The representative transition matrix

Instead of displaying the decisions using the decision matrix, this methodology uses a series of transition matrices, $\mathbf{M}$, to represent each decision epoch. Using this method has several advantages. First, it preserves state transition probabilities, as opposed to simply stating the optimal action. Second, and more importantly, formation of a series of transition matrices enables the ability to perform eigenvalue spectral analysis.

These transition matrices are developed from the decision matrix (Sheskin, 2011). This is done by selecting the state transitions for each state from its respective optimal action and placing it in its respective row in the representative transition matrix. For example, if Action 1 is optimal for State 1 according to the decision matrix, then the first row for the representative transition matrix, $\mathbf{M}$, is identical to the first row of the Action 1 transition matrix. Likewise for State 2, if Action 2 is optimal, then the second row of $\mathbf{M}$ will be identical to the second row of the Action 2 transition matrix. This logic is followed for all states. This new representative transition matrix is square stochastic. This matrix is able to represent the various optimal actions for all states. It is essentially an amalgamation of the set of action transition matrices displaying only the optimal action for each state. The result is a different representative transition matrix for each decision epoch.

An example of the formation of the representative transition matrix is given in Tables 4 through 6. The action transition matrices from the sample system above are presented again. The shading denotes the specific transition probabilities that are transferred to the representative transition matrix according to the decision matrix. The decision matrix for the first decision epoch is presented in Table 5, and the resulting representative transition matrix is developed in Table 6. Eigenvalue analysis is then performed on this transition matrix given in Table 6.

Table 4. Sample transition matrices for 3 state, 2 action MDP. The colors denote the specific transition probabilities that are transferred to the representative transition matrix, $\mathbf{M}$, according to the decision matrix.

Action 1

	s'_1	s'_2	s'_3
s_1	0	0.7	0.3
s_2	0.2	0.8	0
s_3	1	0	0

Action 2

	s'_1	s'_2	s'_3
s_1	0	0	1
s_2	0	1	0
s_3	0.1	0.1	0.8

Table 5. Sample one epoch decision matrix. The preferred actions have been selected via Equation 2.

	State 1	State 2	State 3
Epoch 1	Action 1	Action 2	Action 2

Table 6. Representative transition matrix, M, developed from the action transition matrices (Table 4) and the decision matrix (Table 5).

	s'_1	s'_2	s'_3	
s_1	0	0.7	0.3	(from Action 1)
s_2	0	1	0	(from Action 2)
s_3	0.1	0.1	0.8	(from Action 2)

Eigenvalue spectral analysis

Once the transition matrices are formed for each decision epoch, eigenvalue spectral analysis can be performed. The eigenvalues, λ_i, and eigenvectors, $\mathbf{w}_i$, are defined according to Equation 3. The set of eigenvalues define the spectrum of the Markov process (Cressie and Wikle, 2011). For this research, Equation 3 was solved numerically using a built-in MATLAB function.

$$\mathbf{w}_i\mathbf{M} = \lambda_i\mathbf{w}_i \tag{3}$$

The eigenvalues are key to understanding the underlying dynamics of the system (Salzman, 2007). They represent the analytic solution to the linear system, and can be used to examine the system attributes driving the behavior through time. Identification of oscillatory patterns, system stability, and bifurcation regions (Cressie and Wikle, 2011) in the decision pathways is possible with eigenvalue analysis. Applying eigenvalue spectral analysis on transition matrices developed from a Markov decision process to analyze design decision making behavior is unique to this research. Presented below is one derived metric, designed to identify and quantify changes in decision making behavior.

The decision metric

The decision metric is defined in Equation 4. λ_1 is the largest eigenvalue, and $|\lambda_2|$ is the magnitude of the second largest. For Markov matrices, such as M, λ_1, is always one. This equation has previously been used as a damping ratio metric for linear models to study the transient behavior of a system (Caswell, 2001).

$$\rho = \frac{\lambda_1}{|\lambda_2|} \tag{4}$$

For this paper, the decision metric will be used in two unique ways. First, it will identify and quantify changes in decisions. Changes in this metric are coincident with changes in the transition matrix, and thus represent changes in the decisions. When the decision of an individual state changes, that state's row in the representative transition matrix changes as well. The row consisting of the transition probabilities from the action transition matrix from the previous set of decisions is replaced by the row of the transition probabilities of the new action transition matrix associated with the new decisions. This process is shown in Figure 1.

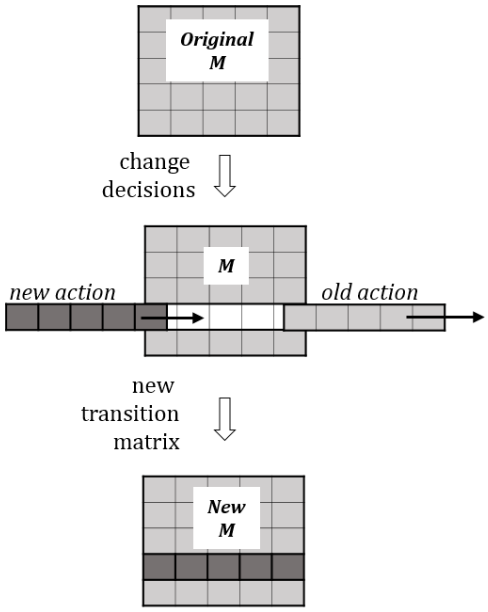

Figure 1. Visualization of how a change in decisions affects the representative transition matrix, M.

Since eigenvalue analysis is performed on these transition matrices, any changes will result in changes in the eigenvalue spectrum. When λ_2 is affected, it is concluded that a major change in the system has occurred because the primary spectral mode has changed. When the decision metric is unaffected by changes in decisions, there is no significant change to the system. This is similar to only affecting the minor spectral modes of the system.

The second new application for the decision metric is that it can help identify the significant state/action combinations that affect the process as a whole. By associating the changes in the sets of decisions with the changes in the decision metric, the decision maker can identify important state/action combinations. A case study is presented showing the utility of this method and metric on a ship egress problem.

Case Study: Ship Egress Analysis and General Arrangements Design

This case study is designed as one application of how this technique can be applied in the ship design context. The case involves studying personnel movement inside a ship. The assumption is that a fire has broken out in one of the rooms, and that individuals need to find the exit. Their movement is probabilistic to simulate the confusion of an individual's location due to smoke that may be percolating throughout the ship or uncertainty with blocked passageways. A utility function is used to simulate the pain experienced while in the ship looking for the exit. The rooms have not been designated for a specific use for this analysis, as it is the attempt of this study to understand their interaction prior to designating their use. That is, these rooms have not yet been classed as galley, engine room, etc.

To set up the problem, individuals are located somewhere in the eleven state environment shown in Table 7, where the entries show the labeling convention of the accessible rooms. The solid black state (state (2,2)) represents an inaccessible area, such as a column or unusable elevator shaft. The top right room (state (3,4)) represents the exit, while the room adjacent below (state (2,4)) is the room with the fire. The objective is to minimize the pain from smoke inhalation while heading towards the exit.

Table 7. Diagram of the eleven room arrangement. Entries show the labeling convention of the states.

(3,1)	(3,2)	(3,3)	(3,4)
(2,1)		(2,3)	(2,4)
(1,1)	(1,2)	(1,3)	(1,4)

The SC-MDP framework is used to determine the best sequence of decisions and the expected utility. The states are defined as the individual rooms in the ship. The individuals may move one step at each decision epoch; either up, down, left, or right. There is uncertainty in the individual's movement, and the probability of moving in the desired direction is only 0.8. There is a 0.1 probability of moving in either direction laterally. The transition probability in this case is defined as $p = 0.8$. If the decision is made to step directly into a wall, the individual remains in the same location. A reward is received for landing in a given room (Table 8). The fire is modeled with a $r = -1$ reward, while the exit has a $r = +1$ reward. The $r = -0.04$ rewards are varied in the following analyses, while the $r = +1$ and $r = -1$ rewards remain fixed.

Table 8. The rewards received for landing in a given state. The following analyses vary the -0.04 rewards, while the $+1$ and -1 rewards remain fixed.

-0.04	-0.04	-0.04	+1
-0.04		-0.04	-1
-0.04	-0.04	-0.04	-0.04

For this study, the transition probabilities and rewards do not change with time. The MDP is run for 30 decision epochs, allowing individuals to take up to 30 steps. The expected utility and best decision paths are given in Table 9. The decision paths display the best action a person should take while in each state. These results are unique only to the given rewards and uncertainty, as changes in either will affect both the expected utility and decisions (Russell and Norvig, 2003).

Table 9. Expected utility and decision paths for $p = 0.8$ and $r = -0.04$.

Expected utility

0.81	0.87	0.92	+1
0.76		0.66	-1
0.71	0.66	0.61	0.39

Decision Paths

→	→	→	+1
↑		↑	-1
↑	←	←	←

Eigenvalue spectral methods are used to gain a deeper understanding into these variations in the policy, and the driving system characteristics that may be causing them. The decision metric as defined previously is examined to elicit this information.

Results

The representative transition matrix was formed after the MDP was run for 30 decision epochs. This process is shown in Figure 2. The eigenvalues were generated from this representative transition matrix, and the decision metric was calculated. As such, these results represent one single time step. The objective is to obtain the relationship between the decision metric and the system for a single time step before examining its behavior through time.

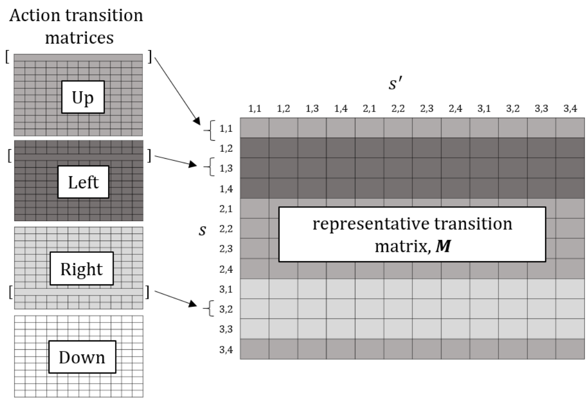

Figure 2. Building the representative transition matrix, M, for $p = 0.8$ and $r = -0.04$. Shading represents specific actions. Note the action "down" is never optimal for any of the states.

A sweep of the uncertainty and rewards was performed to examine the behavior of the decision metric across a broad environment. The uncertainty was varied from complete uncertainty of $p = 0$ to complete certainty of $p = 1$, while the rewards were varied between $-2.0 < r < 0.5$ (Figure 3).

Two major trends are apparent. First, there is a significant drop around where the reward transitions from neg-

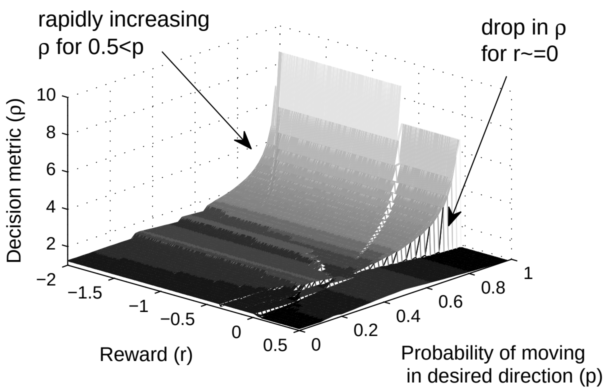

Figure 3. Decision metric, ρ, for a range of rewards and uncertainties. Note the drop in ρ near $r = 0$, and the rapidly increasing ρ for $0.5 < p$.

ative to positive. As the rewards move close to and into the positive region, the set of decisions changes drastically, and may even become non-unique for areas where the rewards are fully positive. Second, as p increases, the decision metric grows rapidly for $p > 0.5$. For the region $p < 0.5$ where the uncertainty is so great that the individuals take a misstep more than 50% of the time, the decision metric is roughly consistent across all rewards. Figure 3 shows only the high level trends. A more detailed analysis of the underlying shape and behavior follows.

To examine the underlying shape of Figure 3, a closer examination of the data was performed with the uncertainty fixed at $p = 0.8$, while the rewards were varied from $-2.0 < r < 0.5$. The results are shown in Figure 4. Five step changes are apparent, located at points b, e, f, h, and i, which are associated with $r = -1.65$, $r = -0.45$, $r = -0.09$, $r = -0.03$, and $r = -0.02$ respectively, with the major drop occurring at $r = -0.03$.

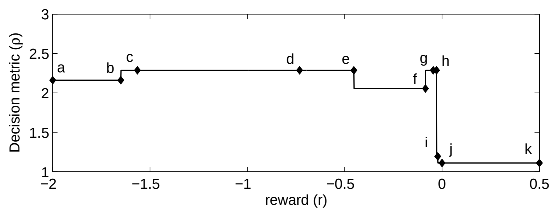

Figure 4. Decision metric, ρ, over a range of rewards for $p = 0.8$. Points indicate transition regions in the decision paths. Note the step function behavior and the drop at $r = -0.03$ (between h and i).

The step changes in Figure 4 occur when there are changes in the decisions. However, not all transitions in the decisions are identified as changes in the decision metric. There are ten different decision paths for an uncertainty of $p = 0.8$. These are given in Table 10, and are arranged moving from left to right in the decision metric plot. The change in the decisions is identified by a double arrow in the state that was affected. Those changes in the decisions that occurred simultaneously with changes in the decision metric are denoted with both double arrows as well as highlighted in gray.

Table 10. Variations in the decision paths for $p = 0.8$. Moving from lowest rewards to highest, the state that changed actions is identified by a double arrow. States that changed actions with a change in the decision metric have a double arrow and a gray background.

$r \leq -1.65$
(a-b)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\rightarrow$	-1
$\rightarrow$	$\rightarrow$	$\rightarrow$	$\uparrow$

$-1.65 < r \leq -1.57$
(b-c)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\Uparrow$	-1
$\rightarrow$	$\rightarrow$	$\rightarrow$	$\uparrow$

$-1.57 < r \leq -0.73$
(c-d)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\uparrow$	-1
$\rightarrow$	$\rightarrow$	$\Uparrow$	$\uparrow$

$-0.73 < r \leq -0.45$
(d-e)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\uparrow$	-1
$\Uparrow$	$\rightarrow$	$\uparrow$	$\uparrow$

$-0.45 < r \leq -0.09$
(e-f)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\uparrow$	-1
$\uparrow$	$\rightarrow$	$\uparrow$	$\Leftarrow$

$-0.09 < r \leq -0.05$
(f-g)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\uparrow$	-1
$\uparrow$	$\Leftarrow$	$\uparrow$	$\leftarrow$

$-0.05 < r \leq -0.03$
(g-h)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\uparrow$	-1
$\uparrow$	$\leftarrow$	$\Leftarrow$	$\leftarrow$

$-0.03 < r \leq -0.02$
(h-i)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\Leftarrow$	-1
$\uparrow$	$\leftarrow$	$\leftarrow$	$\leftarrow$

$-0.02 < r \leq 0.00$
(i-j)

$\rightarrow$	$\rightarrow$	$\rightarrow$	+1
$\uparrow$		$\leftarrow$	-1
$\uparrow$	$\leftarrow$	$\leftarrow$	$\Downarrow$

$0.00 < r$
(j-k)

$\leftrightarrow\updownarrow$	$\leftrightarrow\updownarrow$	$\Leftarrow$	+1
$\leftrightarrow\updownarrow$		$\leftarrow$	-1
$\leftrightarrow\updownarrow$	$\leftrightarrow\updownarrow$	$\leftrightarrow\updownarrow$	$\downarrow$

Starting from the most negative rewards, the first change in decisions occurs at $r = -1.65$ (point b), where state (2,3) changes from action "right" to "up". At the same time the decision metric increases through a small step change. Due to the highly negative rewards in the environment for $r \leq -1.65$ the best decision is to take the shortest path to the -1 state. This is a situation where it is less painful to be in the room with the fire than outside of it. However, when the penalty is changed to $-1.65 < r$, the best path is no longer to step directly at the -1 state from the (2,3) state. This change in the decisions is considered a major change to the system because of the effect it has on the decision metric.

This is in contrast to the change that occurs at $r = -1.57$ (point c). Here, state (1,3) changes from "right" to "up"; however, there is no change in the decision metric.

This change is considered minor because the individuals are still two steps away from the -1 state. The new decision recognizes the +1 state is preferable, but due to the high painful incremental rewards, may only be slightly more preferable than stepping into the -1 state. A similar trend is apparent at $r = -0.73$ (point d), where state (1,1) changes from "right" to "up" with no change in the decision metric. This change in decisions is minor as the individuals are still 5 steps away from the safe exit state no matter the decision path.

For the other transitions this relationship is consistent, both qualitatively and quantitatively. When the decisions change for a given state, the change is considered significant if there is an associated change in the decision metric. These changes typically affect states adjacent to the room with the fire ($r = -1$ state). State (1,2), while not directly adjacent to the -1 state does have a significant change because the best decision is now to go the long way around state (2,2) as opposed to taking the shorter route. On the other hand, states that are farther away, in general, have less of an effect on the system as changes in their decisions do not change the decision metric.

Of particular note is the significant drop in the decision metric that occurs when state (2,3) changes from "up" to "left", located at $r = -0.03$ (point h). Two changes occur here. First, by deciding to go "left", individuals will never take an uncertain misstep into the room with the fire ($r = -1$ state) from the (2,3) state. Second, this decision effectively blocks the passage between state (1,3) and (3,3). They must now travel clockwise around the inaccessible area (state (2,2)). This result is consistent for all $0.6 < p \leq 1$. The significant drop in the decision metric identified in Figure 3 always occurs when the decision for state (2,3) changes from "up" to "left". This drop happens at various rewards for various uncertainties. For $p = 0.6$ this change happens at $r = -0.07$, while it occurs at $r = -0.003$ for $p = 0.9$. Accordingly, as more uncertainty is added into the system, this change in the decision metric occurs further from 0. For $p \leq 0.5$ there is no change in the decision metric when the decisions changes for this state.

Discussion

The method presented is significant for understanding decision making regarding egress routes and general arrangements design. This paper focused exclusively on the egress problem, defined as how people egress, understanding the decisions they make under uncertainty, and the interaction between the individuals themselves and the layout of the vessel. This is juxtaposed to traditional analyses that focus on the solution, namely the physical layout of the vessel and the distribution of the crew throughout the ship.

The decision metric highlights important transition regions in the decisions, and which state/action combinations are significant. Without the use of the spectral methods, the designer would have to examine all nine transition regions in the decisions; however, the decision metric reduces that number to five areas, with one of significant importance. The room next to the fire (state (2,3)) showed to be an important area, especially when the decision transitions from "up" to "left". Designers need to be careful when allocating this space as a particular room so not as to block this egress route, or they need to ensure that proper smoke ventilation is installed so as not to affect the pain of the individuals, thus impacting their decision making. By highlighting the importance of this room, and the relative insignificance of other rooms (such as state (1,1)), this metric has been able to show which areas of the vessel deserve greater focus and which rewards are likely to cause changes in decision making behavior for those on board.

Conclusion

A method for enabling decision making insight has been presented involving applying eigenvalue spectral analysis to the SC-MDP framework. A decision metric for Markov decision processes was introduced, defined as the ratio between the largest eigenvalue and the magnitude of the second largest eigenvalue. This metric was applied in two new ways: first to identify and quantify changes in the sets of decisions, and second, to identify the specific system attributes causing the major changes in decisions. Use of eigenvalue spectral methods will be beneficial for ship design and decision making by eliciting new insight into the design and decision space that may not be possible using traditional methods. Future work will include using these results to help guide the design of the vessel.

Acknowledgement

The authors would like to thank Ms. Kelly Cooper and the U.S. Office of Naval Research for their support.

References

Andrews, D., Duchateau, E., Gillespe, J., Hopman, H., Pawling, R., and Singer, D. (2012). State of the art report: Design for layout. In *11th International Marine Design Conference*, pages 111–137, Glasgow.

Casarosa, L. (2011). *The integration of human factors, operability and personnel movement simulation into the preliminary design of ships utilising the Design Building Block approach*. PhD thesis, University College London, London, UK.

Caswell, H. (2001). *Matrix Population Models: Construction, Analysis, and Interpretation*. Sinauer Associates, Inc., Sunderland, MA, 2nd edition.

Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, NJ.

Eastland Memorial Society (2015). Eastland memorial society. www.eastlandmemorial.org/. April, 2016.

Guarin, L., Hifi, Y., and Vassalos, D. (2014). Passenger ship evacuation - design and verification. In *Virtual, Augmented and Mixed Reality. Applications of Virtual and Augmented Reality*, volume 8526, pages 354–365, Heraklion, Crete, Greece.

IMO (2007). Guidelines for evacuation analysis for new and existing passenger ships. www.imo.org.

Kana, A. A., Knight, J. T., Sypniewski, M. J., and Singer, D. J. (2015). A Markov decision process framework for analyzing LNG as fuel in the face of uncertainty. In *12th International Marine Design Conference*, volume 2, pages 297–308, Tokyo, Japan.

Neumaier, A. and Schneider, T. (2001). Estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Transactions on Mathematical Software*, 27(1):27–57.

Niese, N. D., Kana, A. A., and Singer, D. J. (2015). Ship design evaluation subject to carbon emission policymaking using a Markov decision process framework. *Ocean Engineering*, 106:371–385.

Niese, N. D. and Singer, D. J. (2013). Strategic life cycle decision-making for the management of complex systems subject to uncertain environmental policy. *Ocean Engineering*, 72:365–374.

Niese, N. D. and Singer, D. J. (2014). Assessing changeability under uncertain exogenous disturbance. *Research in Engineering Design*, 25:241–258.

Puterman, M. L. (2005). *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, Hoboken, NJ.

Qiao, Y., Han, D., Shen, J., and Wang, G. (2014). A study on the route selection problem for ship evacuation. In *2014 IEEE International Conference on Systems, Man, and Cybernetics*, pages 1958–1962, San Diego, CA.

Reich, Y. (1995). A critical review of general design theory. *Research in Engineering Design*, 7:1–18.

Rigterink, D., Piks, R., and Singer, D. J. (2014). The use of network theory to model disparate ship design information. *International Journal of Naval Architecture and Ocean Engineering*, 6:484–495.

Russell, S. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, 2nd edition.

Salzman, J. (2007). *Spectral Analysis with Markov Chains*. PhD thesis, Stanford University.

Seram, N. (2013). Decision making in product development - a review of the literature. *International Journal of Engineering and Applied Sciences*, 2(4):1–11.

Sheskin, T. J. (2011). *Markov Chains and Decision Processes for Engineers and Managers*. CRC Press, New York.

Tan, C. H. and Hartman, J. C. (2011). Sensitivity analysis in markov decision processes with uncertain reward parameters. *Journal of Applied Probability*, 48:954–967.

van Bruinessen, T., Hopman, H., DeNucci, T., and van Oers, B. (2011). Generating more valid designs during design exploration. *Journal of Ship Production and Design*, 27(4):153–161.

Vanem, E. and Skjong, R. (2006). Designing for safety in passenger ships utilizing advanced evacuation analyses - a risk based approach. *Safety Science*, 44:111–135.