

A decision-making framework for planning lifecycle ballast water treatment compliance

Kana, Austin; Brefort, D.C.; Seyffert, H.C.; Singer, D.J.

Publication date

2016

Document Version

Final published version

Published in

Proceedings of the 13th International Symposium on Practical Design of Ships and Other Floating Structures (PRADS2016)

Citation (APA)

Kana, A., Brefort, D. C., Seyffert, H. C., & Singer, D. J. (2016). A decision-making framework for planning lifecycle ballast water treatment compliance. In U. D. Nielsen, & J. J. Jensen et al (Eds.), *Proceedings of the 13th International Symposium on Practical Design of Ships and Other Floating Structures (PRADS2016)* PRADS Organising Committee.

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A decision-making framework for planning lifecycle ballast water treatment compliance

Austin A. Kana¹), Dorian C. Brefort²), Harleigh C. Seyffert²), and David J. Singer²)

¹Delft University of Technology, the Netherlands

²University of Michigan, USA

Abstract

This paper introduces a novel decision-making framework for planning lifecycle compliance of ballast water treatment by applying eigenvalue spectral analysis to the ship-centric Markov decision process (SC-MDP) framework. This method focuses on identifying the relationships of various decision making scenarios, and how those relationships change through time. The objective is to understand both these relationships and the impact of initial technology selection on lifecycle ballast water compliance. Two metrics are used. First, the optimal lifecycle strategy is presented for technology selection. Second, the set of dominant eigenvalues is used as a metric to identify the number of unique, initial condition dependent design absorbing paths the process may converge to. Sensitivity studies are performed examining the affect of policy strength on preferred compliance strategy.

Keywords

Ship design; decision making; ballast water compliance; Markov decision process; eigenvalue analysis

Introduction

Ballast water treatment has become compulsory due to various levels of environmental regulations. Regulating the discharge of ballast water has been recognized as an important part in the fight against invasive species. To combat this, many governing bodies have put into place strict guidelines that specify the quality of the water that is being discharged. Unfortunately for ship owners and operators, decision making and lifecycle planning for ballast water treatment compliance remains difficult due to the interplay of various factors, including: stochastic degradation, technology development, and multiple levels of environmental policy-making.

Ballast water is regulated by multiple governing bodies, including, the IMO, the U.S. Coast Guard, various states and local governing bodies, and the European Union. In 2004, the IMO adopted the Ballast Water

Management Convention designed to regulate global discharge of ballast water. These regulations will apply to all vessels required to carry ballast, including: submersibles, floating craft, floating platforms, floating storage units and floating production storage and offloading vessels. The ramifications for violating the regulations are significant, ranging from monetary fines to criminal sanctions for willful noncompliance (Davis and Levy, 2012).

While the IMO Ballast Water Convention was held in 2004, by 2016 it has not come into effect. The regulation will go into effect 12 months after it has been ratified by 30 member States representing 35% of the world merchant shipping tonnage (IMO, 2016). Despite the fact the IMO regulations are not in force, there are still significant reasons to study their potential impact on ship design and decision making. These regulations will likely come into force soon, creating a necessity to have a strategic plan for them now. Also, vessels already have to consider national and regional regulations that are in force.

Acceptable ballast water technologies are dependent on the size of the ballast capacity and the year the ship was constructed. The type of approved technology changes in 2016. This date was selected to give engineers time to develop applicable technologies. Even though the IMO regulation has yet to come into force, technology developers and vessel owners have had to prepare for this upcoming change well in advance, despite the uncertainty surrounding the enforcement date.

Many technologies already exist that meet some of the regulations, and others are still in development to meet the most stringent of the policies. Ballast water exchange systems will no longer be allowed once the regulation goes into force. The other option is ballast water treatment, which tries to kill the bacteria and living organisms in the ballast water. Currently there are over 70 different manufacturers of ballast water treatment technologies across a range of various treatment options (Lloyd's Register, 2012). This makes the planning and selection of proper technologies complicated and difficult.

To help ship designers with understanding how to comply with these regulations, the authors propose applying eigenvalue analysis to the ship-centric Markov decision process (SC-MDP) framework. The SC-MDP is selected because of its ability to handle uncertain, temporal deci-

sion making problems in the maritime domain. Using a stochastic framework to model the uncertainty in lifecycle analysis is necessary because deterministic methods may provide incomplete and sometimes conflicting information for the decision maker (Zayed et al., 2002). Eigenvalue analysis is chosen to gain an understanding of the relationships and implications of decisions throughout the vessel's life. This is the first application of applying eigenvalue analysis to the SC-MDP framework on a non-stationary temporal problem. An extended background on eigenvalue analysis applied to the SC-MDP framework can be found in Kana and Singer (2016).

Previous work on applying the SC-MDP framework to the ballast water treatment problem has been performed by Niese and Singer (2013, 2014). These previous studies performed simulations to capture the interplay between internal and external forces. This was done in an effort to develop both a design and a lifecycle decision making strategy that minimizes life cycle cost while maintaining compliance and performance. This paper extends this work by introducing temporal eigenvalue spectral methods to gain a deeper understanding of the driving forces behind the different decision making scenarios, as well as quantifying their differences. The authors use eigenvalue analysis to help identify and examine interdependencies between decision paths and projected design scenarios.

Specifically this paper discusses the concept of design absorbing paths. An absorbing path represents the long term behavior of a non-stationary decision process. Some processes may have more than one absorbing path for the whole decision process, each one being dependent on the initial conditions of the system (Gebali, 2008). Niese et al. (2015) discussed the importance of identifying the presence of multiple absorbing paths. They discussed that differing absorbing paths may mean that differing decision sequences may be viewed as only locally optimal. They were able to identify the multiple paths through the use of simulation. This paper, on the other hand, claims that these differing paths are in fact dependent on the initial conditions of the system.

A case study is presented showing the utility of eigenvalue analysis for a non-stationary temporal decision process. Metrics that handle repeated eigenvalues are used to study initial condition dependence of design convergence paths. The focus is on identifying relationships and interdependencies in the decision process and their behavior through time.

Methods

The following procedure outlines how to apply eigenvalue analysis to the SC-MDP framework. A full detailed description of the procedure can be found in Kana and Singer (2016). While each step can be found in the literature, it is the combination of steps, and the applications presented here, that makes this research unique.

1. *Obtain the set of decisions and expected utilities by solving the ship-centric Markov decision process* (Puterman, 2005). The objective is to determine the

decisions that maximize the lifecycle expected utility of the system. An MDP has four attributes: 1) a set of states, S , describing the environment, 2) a set of actions, A , the decision maker can take, 3) a set of probabilities, T , of transitioning from one state, s , to a new state, s' , after taking a given action, a , and 4) a set of rewards, R , received by executing a given action and transitioning to a new state. The expected utility is found using Equation 1, while the decisions, π , are found by Equation 2 (Russell and Norvig, 2003), where U is the expected utility and γ is the discount factor.

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

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

2. *From the set of decisions, develop a series of representative transition matrices, M , for each decision epoch* (Sheskin, 2011). The transition matrices are generated using the decision matrix (Sheskin, 2011). To do this, select the state transitions from the optimal actions for each state and insert them into the associated row in the representative transition matrix. This is done for all states for every time step. The result is one representative transition matrix for each time step. These new transition matrices are by definition square stochastic (Anton and Rorres, 2005). Kana and Singer (2016) present an explicit example of how to form these representative transition matrices.
3. *Perform eigenvalue spectral analysis on the transition matrices to generate the spectrum of the MDP* (Caswell, 2001). Equation 3 gives the eigenvalues, λ_i , and eigenvectors, w_i of M . Cressie and Wikle (2011) define the spectrum of a Markov process as the set of its eigenvalues. For this research, the eigenvalues were found using a built-in MATLAB function.

$$w_i M = \lambda_i w_i \quad (3)$$

Composite Reducible Markov Processes

This paper distinguishes itself from the methods presented in Kana and Singer (2016) because this paper handles situations where the dominant eigenvalue may be repeated. This may be especially helpful for analyzing temporal systems. For stochastic matrices, such as M , the dominant eigenvalue equals one, that is $\lambda_1 = 1$ always (Kirkland, 2009). Two types of behavior may occur when λ_1 is repeated. First, the process may fail to converge to a single steady state distribution. This may happen if the process oscillates between more than one steady state, or if multiple steady states exist simultaneously. Second, the long term convergence may be initial

condition dependent (Kirkland, 2009). This means the system will converge based on both the set of decisions and on where the system starts. In these situations the designer needs to be very careful in how they select their starting state.

When the dominant eigenvalue is repeated, different analysis techniques are necessary. One analysis technique is presented in this paper that involves the idea of composite reducible Markov processes to group specific aspects of the design and decision process that align with each other (Gebali, 2008). Reducible Markov processes are those in which not every state is reachable from every other state. Composite reducible Markov processes are those in which there is more than one set of grouped states. That is, by starting in a specific set of states, it is not possible to reach certain other states.

This technique helps identify relationships and interdependencies between specific decisions within the whole decision process. From a ship design perspective, this is similar to deciphering the dependencies between various decisions a single designer must make or between various design teams working on a single project. Likewise, from a model perspective, the physical meaning of this is that this technique decomposes the one decision process into multiple independent decision processes. This aspect is important because it highlights the initial condition dependence of reducible Markov processes. Also, it may inform the decision maker that certain design or decision paths may not be reachable given a specific initial condition. It is possible to determine the presence of composite reducible Markov processes by examining the set of dominant eigenvalues Gebali (2008). To show this, this paper uses the number of repeated dominant eigenvalues to identify the number of unique, initial condition absorbing paths on a case study involving life-cycle planning for ballast water treatment compliance.

Case Study: Lifecycle Planning for Ballast Water Treatment Compliance

A notional 150,000 deadweight tonnage containership with a 30,000 metric ton ballast water capacity routed along the trans-pacific route is used for this study. The ballast water treatment system must have a capacity of at least $10,000m^3/h$. The vessel has a 20 year lifespan and is put in service sometime before the 2004 IMO Ballast Water Management Convention. Ten ballast water systems, labeled 1-10, are considered. System 1 is a commercially available ballast water exchange system. Systems 2-10 represent ballast water treatment systems that become commercially available at some time during the lifespan of the vessel. Specifics of the ballast systems, including performance, capital costs, operating expenditures, availability and approval have been derived from Lloyd's Register (2007, 2010); California State Lands Commission (2010). These systems represent various treatment technologies, such as: filtration, electrochlorination, cavitation, radiation, and de-oxygenation. The original case setup, including inputs, stochastic variables,

and economic parameters have all been tested and validated against historical data by Niese and Singer (2013, 2014).

Markov decision process framework

This section details how the MDP states, actions, transition probabilities, and rewards are defined for this study.

States

The states are defined by the individual ballast water systems, their commercial availability and regulatory approval, and their deterioration level. For each ballast system, there are six availability states: unavailable, commercially available, basic approval, final approval-Tier 1, final approval-Tier 2, and final approval-Tier 3. Tier-2 regulation is roughly 10x more stringent than Tier-1, and the Tier-3 is roughly 100x more stringent than Tier-1 regulation. Each system also has four deterioration levels. The deterioration level is defined as a percentage of total deterioration. There are 240 states, accounting for ten systems, six approval states, and four deterioration levels.

Starting State

This analysis assumes that the initial state is unknown. That is, there is equal probability of being located in any of the states at the start of the model. Niese et al. (2015) discussed the problems associated with, and importance of, selecting the correct start state, and its implications on future decision making opportunities. Designs may be dependent on the initial conditions, and thus selecting differing starting states may lead the design down a different path. By assuming equal starting probabilities, this analysis aims to find the most natural path the design would take as opposed to pre-determining its trajectory.

Actions

There are twelve actions available to the decision maker:

1. *No Action*: The system continues to deteriorate, yet no action is necessary to maintain it.
2. *Maintain*: Maintenance is performed and the ballast water system is restored to a less deteriorated state.
3. *Replace System (1-10)*: The ballast system is replaced with one of the 10 possible systems. The system installed is identified by the index 1-10. A system can only be installed after it becomes commercially available and meets regulatory requirements.

Transition Probabilities

The probability of transitioning between states is defined as follows:

- Transitioning between ballast water systems is determined based on the best action selected by the MDP and its availability. A system can only be selected once it is available and approved.

- Transitioning between approval states is based on the regulatory environment and the commercial availability. The following schedule is used to model various regulatory scenarios.

1. The ballast water convention is held which outlines the expected strength of the legislation, as well as expected date of enforcement.
2. Laboratory testing procedures specific to ballast water treatment efficacy are available.
3. The legislation is ratified by member States.
4. The legislation enters force.

The implementation schedule is defined as the number of years following the convention a policy trigger occurs. For instance, the 1-4-4-9 schedule simulates a convention being held one year after the ship enters service. Four years later testing procedures become available and the legislation is ratified. Nine years after the convention the legislation enters force. Prior to the convention, there is little demand for development of the treatment technologies, and thus it is assumed the treatment technologies will not become available until after the convention is held. Each individual technology will meet a different threshold of regulatory compliance and will become available at different times. The schedule outlining the expected year each technology will be available is given in Table 1. The table outlines the number of years following the convention that the technologies are expected to be available, and their expected regulatory compliance level. This data has been based on actual dates when the technologies became available, while the deviation has been included to simulate uncertainty in the commercialization process (Niese and Singer, 2013).

Table 1. Ballast water technology availability schedule and compliance level. The mean availability gives the number of years after convention the technology is expected to be commercially and regulatory compliant.

| Ballast System | Mean Availability (years) | Deviation σ (years) | Compliance Level |
|----------------|---------------------------|----------------------------|------------------|
| 1 | - | - | Exchange |
| 2 | 3 | 0.5 | Tier 3 |
| 3 | 2 | 0.4 | Tier 3 |
| 4 | 7 | 1.0 | Tier 1 |
| 5 | 3 | 0.5 | Tier 1 |
| 6 | 5 | 0.75 | Tier 2 |
| 7 | 7 | 1.0 | Tier 3 |
| 8 | 5 | 0.75 | Tier 3 |
| 9 | 4 | 0.6 | Tier 2 |
| 10 | 3 | 0.5 | Tier 3 |

- Transitioning between deterioration levels is modeled by the exponential distribution (Equation 4).

$$f_j(x) = \gamma_j e^{-\gamma_j x} \quad (4)$$

Deterioration happens independently, and follows an exponential distribution for γ (Equation 5). λ_j is a function of the system's treatment method. This is due to ballast water treatment systems using filtration, electrochlorination, cavitation, radiation, deoxygenation, and/or ozone-generation degrade differently (Niese and Singer, 2013). A full description of this model can be found in van Noortwijk (2007).

$$\gamma_j = a_j e^{-b_j} + c_j \quad (5)$$

Figure 1 shows the availability calculated by the Markov process of the various systems according to commercial availability and regulatory compliance for the 1-4-4-9 regulatory implementation schedule. For visualization purposes the 240 states have been condensed to 60 representative states. To do this, the four deterioration levels for each ballast water system and for each approval status have been added together. This creates a single representative state that accounts for all four deterioration levels (Niese and Singer, 2014).




Figure 1. Ballast water system commercial availability and regulatory compliance for the 1-4-4-9 schedule. Shading represents the percent likelihood a given system will be located in that state.

Rewards

The rewards are based on the system capital costs, installation costs, and operating and maintenance costs (Table

2) (Lloyd's Register, 2007, 2010; California State Lands Commission, 2010; Rigby and Taylor, 2001). The cost function is given in Equation 6.

$$\text{cost} = \min(\text{capital} + \dots \\ \text{install} + \text{operating} + \text{maintenance}) \quad (6)$$

The capital costs are dependent on whether the system meets basic approval. Capital costs tend to increase after achieving basic approval because the approval status may warrant a cost increase, or supply and demand economics may dictate it (Niese and Singer, 2013). Installation costs vary depending on whether it is during vessel new construction or a retrofit. In cases of a retrofit, it is assumed there is sufficient space.

Table 2. Ballast water technology costs. The Capex ## corresponds to costs before/after basic approval. The Install ## corresponds to costs for newbuild/retrofit.

| System Reference | Capex (2,000m ³ /h) | Install (2,000m ³ /h) | O&M (\$/m ³ /h) |
|------------------|-----------------------------------|-------------------------------------|-------------------------------|
| 1 | 50/50 | 0/0 | 0.06 |
| 2 | 800/820 | 40/55 | 0.08 |
| 3 | 950/1,200 | 5/15 | 0.07 |
| 4 | 950/1,500 | 50/65 | 0.06 |
| 5 | 690/670 | 60/60 | 0.13 |
| 6 | 800/450 | 80/100 | 0.32 |
| 7 | 500/975 | 65/125 | 0.013 |
| 8 | 1,600/1,600 | 5/15 | 0.06 |
| 9 | 559/600 | 100/150 | 0.03 |
| 10 | 1,800/1,200 | 25/40 | 0.01 |

As equipment deteriorates, it becomes less efficient and more costly. Equation 7 and 8 model the increasing operating costs as a function of deterioration. For this study $g = 0.01$, and $x = [1, 2, 3, 4]$ depending on the deterioration level. $\lambda = [0.72, 0.78]$ and is a function of the system installed (Niese and Singer, 2013). A full description of this deterioration cost function model can be found in (Nguyen et al., 2010).

$$\text{O\&M cost} = \\ \text{Annual trips} \times \text{required ballast} \times \phi(x) \quad (7)$$

$$\phi(x) = \phi_0 + ge^{\gamma_j x} \quad (8)$$

Results

Three sets of results are examined. First, the optimal states are analyzed to see the impact a given regulatory strength and schedule has on the ballast water system of choice. This is done without the use of spectral methods. Second, spectral methods are used to examine interdependencies of the decision process and how those dependencies change through time. This is done through analysis of the set of dominant eigenvalues and using them as

a metric for identifying independent decision absorbing paths. Third, the number of absorbing paths identified by the eigenvalues is validated by varying the initial conditions. Again, an absorbing path describes the long term behavior of the decision process.

Optimal States Accessed

The model was run to determine what the best decisions are, when they should be made, and what ballast system is best to install under given conditions. Figure 2 gives the optimal states plot that displays the preferred ballast system that should be installed at any given time. This plot accounts for uncertainty in technology availability, thus there is no uncertainty between making the choice to install a particular ballast water system and actually having it installed. When a given system is selected in the optimal states plot, it is assumed that the optimal action is to select that particular system. This metric was used extensively by (Kana et al., 2015) to study temporal decision making behavior in the face of evolving Emission Control Area regulations.


A sensitivity study was performed on the strength of the regulation as to its affect on the preferred ballast system. As shown in Figure 2, for the 1-4-4-9 regulatory schedule, with a Tier 1 regulation strength, the best choice is to install ballast System 9 after year 9. System 1 becomes unavailable due to regulatory requirements at year 9, thus necessitating a change. System 9 is selected as the best option, which meets Tier 2 requirements, despite the regulation only requiring Tier 1 compliance. When the strength of the regulation is increased to requiring Tier 3 compliance, ballast System 2 becomes preferred after year 9. Only 5 of the original 10 systems meet Tier 3 standards, and System 2 was selected due to it lower lifecycle costs.

While this metric shows what the best decisions are, it does not show which other technologies may also be desirable or how the initial conditions may be affecting future decision opportunities. Also, this analysis is only able to display one particular absorbing path. The following study on the set of dominant eigenvalues aims to address these limitations.

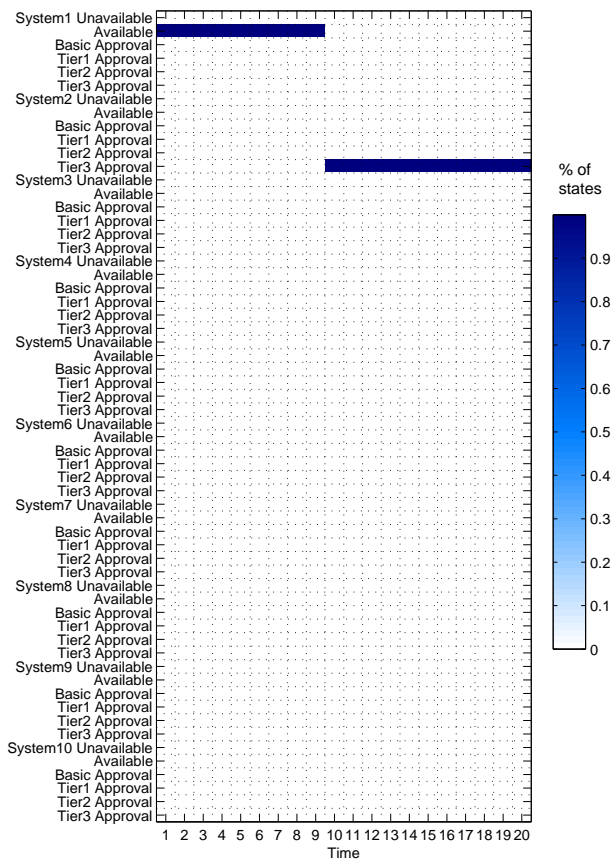
Eigenvalue Analysis

A metric is presented that uses the number of dominant eigenvalues to show the number of possible absorbing paths and how they may evolve through time. This was done to show how the structure of decision process evolves through time. As discussed in Gebali (2008), the number of dominant eigenvalues, $\lambda_i = 1$, is equal to the number of unique absorbing paths of the decision process. In a sense, the number of unique dominant eigenvalues signify that the decision process is not a single connected process, but rather a collection of independent decision processes.

Figures 3 and 4 show the number of unique absorbing paths for Tier 1 and Tier 3 regulatory strength respectively. Up to year 4 there is only one possible path, mean-



(a) Regulation: Tier 1



(b) Regulation: Tier 3

Figure 2. Optimal states accessed for 1-4-4-9 regulatory schedule and two treatment strengths.

ing the process will always converge to a single set of states. Beginning at year 5, when testing becomes available and when the regulation is ratified by the member States, multiple paths become possible. The increasing number of absorbing paths with time is representative of the number of ballast water systems that may be installed in the long term. At year 10 the regulation enters force, thus removing ballast System 1 from compliance. This explains the drop in both figures at year 10. After year 10, only those technologies that meet the regulation can become a possible absorbing path. Thus, the number of unique paths for the Tier 3 schedule is only five (Figure 4), while there are nine unique paths for the Tier 1 regulation (Figure 3).

The number of absorbing paths represents more than just technology availability and compliance. It is essentially a synthesis of technology availability, compliance, their uncertainty, as well as lifecycle costs. For instance, four different ballast systems are potentially available at year 4 (see Figure 1), and yet there is only one absorbing path. This is because, while those systems may be technically feasible, there is no decision path that will select them in the long run. All ballast systems become available by year 9, yet it is not until year 12 that the number




Figure 3. The number of initial condition dependent absorbing paths: 1-4-4-9 schedule and Tier 1 strength.




Figure 4. The number of initial condition dependent absorbing paths: 1-4-4-9 schedule and Tier 3 strength.

of absorbing paths becomes steady. Thus, lifecycle time also affects which systems may be selected. This is a subtle but valuable contribution of this method.

Showing initial condition dependence

A study was performed to show how varying the initial conditions of the system can identify what those specific absorbing paths may be. To show this, the initial conditions were changed so that at a given year there was equal probability of landing in any state. The model was then run to see how the process evolves through time given this new set of conditions. Year 8 was chosen for this validation study. Thus, at year 8, the system is run assuming that the prior year there is equal probability of being in any state. This is different from the original analysis where the process was started at year 1.

Figure 5 shows the results for Tier 1 regulatory strength. For this case there are six different absorbing paths identified by the set of dominant eigenvalues. When the initial conditions for year 8 are changed so that there is equal probability of being in each state there are six paths identified using the state vector. The probability of landing in one absorbing path over another is not equal. For example, it is more likely that System 9 will be the preferred choice over System 2, 3, 5, or 10. System 1 appears as a long term absorbing path even though System 1 is not viable for the whole lifespan of the vessel.

Discussion

The results presented in this paper are significant for ship designers and decision makers for several reasons. First, the spectral techniques presented gives a unique perspective into the structure of the decision process. Understanding the interdependencies of the decision making process and how those dependencies may change and evolve throughout the lifecycle of the vessel provides ship designers great power as they aim to understand the impact of their decisions. The number of dominant eigenvalues clearly displays the evolution of these relationships and dependencies. Second, the spectral methods are inherently a leading indicator highlighting the impact of decision making. Spectral analysis has represented the long term absorbing paths the design may follow without the need for simulation. The focus was less on what the final design is, but instead this analysis has focused on why that final design was selected, the paths that lead the decision process to that point, and the underlying structure of the entire process.

Conclusion

A method for applying eigenvalue spectral analysis to the SC-MDP framework for a temporal, non-stationary problem has been presented. The set of dominant eigenvalues was used to identify and group independent states and processes within the SC-MDP framework. This method benefits ship designers by clearly eliciting the number of




Figure 5. Optimal states accessed for 1-4-4-9 regulatory schedule assuming equal probability of being in any state prior to year 8.

feasible design absorbing paths without the need for simulation and exhaustive perturbation of initial conditions.

Acknowledgement

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

References

Anton, H. and Rorres, C. (2005). *Elementary Linear Algebra: Applications Version*. Wiley, Hoboken, NJ, ninth edition.

California State Lands Commission (2010). 2010 assessment of the efficacy, availability, and environmental impacts of ballast water treatment systems for use in California waters. Technical report, California State Lands Commission.

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.

- Davis, A. N. and Levy, A. D. (2012). *U.S. Coast Guard ballast water discharge rule establishes new requirements and enforcement risks for ship owners and operators*. Shipman and Goodwin LLP.
- Gebali, F. (2008). *Analysis of Computer and Communication Networks*, chapter Reducible Markov Chains. Springer, Victoria, B.C.
- IMO (2016). Ballast water management. 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.
- Kana, A. A. and Singer, D. J. (2016). A ship egress analysis method using spectral Markov decision processes. In *13th International Symposium on Practical Design of Ships and Other Floating Structures (PRADS'2016)*, Copenhagen, Denmark.
- Kirkland, S. (2009). Subdominant eigenvalues for stochastic matrices with given column sums. *Electronic Journal of Linear Algebra*, 18:784–800.
- Lloyd's Register (2007). Ballast water treatment technology. Technical report, Lloyd's Register.
- Lloyd's Register (2010). Ballast water treatment technology. Technical report, Lloyd's Register.
- Lloyd's Register (2012). Ballast water treatment technologies and current system availability. Technical report, Lloyd's Register.
- Nguyen, T. P. K., Castanier, B., and Yeung, T. (2010). Optimal maintenance and replacement decisions under technological change. In *European Safety and Reliability*, Rhodes, Greece.
- Niese, N. D., Kana, A. A., and Singer, D. J. (2015). Ship design evaluation subject to carbon emission policy-making 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.
- Rigby, G. R. and Taylor, A. H. (2001). Progress in the management and treatment of ship's ballast water to minimize the risks of translocating harmful nonindigenous aquatic organisms. *Journal of Marine Environmental Engineering*, 6:153–173.
- Russell, S. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, 2nd edition.
- Sheskin, T. J. (2011). *Markov Chains and Decision Processes for Engineers and Managers*. CRC Press, New York.
- van Noortwijk, J. M. (2007). A survey of the application of gamma processes in maintenance. *Reliability Engineering and System Safety*, 94:2–21.
- Zayed, T. M., Chang, L.-M., and Fricker, J. D. (2002). Life-cycle cost analysis using deterministic and stochastic methods: Conflicting results. *Journal of Performance of Constructed Facilities*, 16(2):63–74.