Managing Aquatic Invasions: Optimal Locations and Operating Times for Watercraft Inspection Stations

Samuel M. Fischer[1,*], Martina Beck[2], Leif-Matthias Herborg[3], and Mark A. Lewis[1,4]

[1] *Department of Mathematical and Statistical Sciences, University of Alberta, Edmonton, AB.*

[2] *BC Ministry of Environment and Climate Change Strategy, Conservation Science Section, Victoria, BC.*

[3] *Fisheries and Oceans Canada, Institute of Ocean Sciences, Sidney, BC.*

[4] *Department of Biological Sciences, University of Alberta, Edmonton, AB.*

[*] *Department for Mathematical and Statistical Sciences; 632 Central Academic Building; University of Alberta; Edmonton, AB; T6G 2G1; E-Mail: samuel.fischer@ualberta.ca*

Abstract

Aquatic invasive species (AIS) cause significant ecological and economic damages around the world. A major spread mechanism for AIS is traffic of boaters transporting their watercraft from invaded to uninvaded waterbodies. To inhibit the spread of AIS, Canadian provinces and American states often set up watercraft inspection stations at roadsides, where potentially infested boats are screened for AIS and, if necessary, decontaminated. However, since budgets for AIS control are limited, watercraft inspection stations can only be operated at specific locations and daytimes. Though theoretical studies provide managers with general guidelines for AIS management, more specific results are needed to determine when and where watercraft inspections would be most effective. This is the subject of this paper. We show how linear integer programming techniques can be used to optimize watercraft inspection policies under budget constraints. We introduce our approach as a general framework and apply it to the prevention of the spread of zebra and quagga mussels (*Dreissena spp.*) to the Canadian province of British Columbia. We consider multiple scenarios and show how variations in budget constraints, propagule sources, and model uncertainty affect the optimal policy. Based on these results, we identify simple, generally applicable principles for optimal AIS management.

Keywords: aquatic invasive species; linear integer programming; optimal management; spatially explicit; zebra mussel

1 Introduction

Human traffic and trade are major vectors for invasive species (Lockwood et al., 2013). Due to the significant ecological and economic damages invasive species cause (Pimentel et al., 2005), government regulations restrict the import of certain goods and require treatment of potentially infested freight and carriers (Shine et al., 2010; Johnson et al., 2017; Turbelin et al., 2017). While such regulations may be enforced comparatively easily at sea ports, air ports, and border crossings, control of inland traffic is more difficult, as a vast number of routes need to be monitored. This applies, for example, to the spread of zebra and quagga mussels (*Dreissena spp.*) and other aquatic invasive species (AIS), which often spread with watercraft and equipment transported from invaded to uninvaded waterbodies (Johnson et al., 2001). Zebra and quagga mussels are invasive in North America and have negative effects on native species and ecosystems, water quality, tourism, and infrastructure (Rosaen et al., 2012; Karatayev et al., 2015).

To counteract the spread of these AIS, watercraft inspection stations are often set up on roads, where transported watercraft are inspected for AIS and decontaminated if at risk for carrying AIS (Mangin, 2011; Alberta Environment and Parks Fish and Wildlife Policy, 2015; Inter-Ministry Invasive Species Working Group, 2015). However, since budgets for inspections are limited, not all pathways can be monitored around the clock and managers need to prioritize certain locations and daytimes. Though several theoretical studies provide managers with helpful guidelines for their work (Leung et al., 2002; Potapov and Lewis, 2008; Potapov et al., 2008; Vander Zanden and Olden, 2008; Finnoff et al., 2010; Hyytiäinen et al., 2013), more specific results are needed in practice to determine the locations and times for where and when control is most effective. To date it has been difficult to tackle these questions rigorously, as comprehensive models for road traffic of potential vectors have been missing. Therefore, AIS managers have relied on past watercraft inspection data, shared experience between jurisdictions, and iterative improvements of control policies. Recent modelling advances (Fischer et al., 2020), however, now permit the application of quantitative methods to optimize control measures in road networks and to evaluate their effectiveness. This is the subject of this paper.

Our goal is to minimize the number of boaters reaching uninvaded waterbodies without being inspected for AIS. We assume that a fixed budget is available for AIS control. This problem setup differs from scenarios considered in other studies, where the budget is optimized along with the control actions so as to minimize the joint costs of invasive species damages and management (Hastings et al., 2006; Potapov and Lewis, 2008; Potapov et al., 2008; Finnoff et al., 2010; Epanchin-Niell and Wilen, 2012; Hall et al., 2018). This approach, however, requires that the damage caused by invasive species is quantified in monetary terms. This is an often difficult and politically sensitive task. Furthermore, the budget available for AIS control may be subject to political and social influences and determined on a different decision hierarchy than management actions. Therefore, AIS managers may seek to spend a fixed yearly budget optimally rather than to determine the theoretically best control budget. The presence of fixed budget constraints also reduces the need to consider the invasion as a dynamic process.

Identifying the locations where a maximal number of boaters could be screened for AIS is similar to the problem of finding optimal locations for road-side infrastructure (Trullols et al., 2010). A well-known technique to solve such problems is linear integer programming (Conforti et al., 2014). The idea is to model the optimization problem with functions that are linear in terms of the decision variables. Though solving linear integer programs is a computationally difficult task in general, good approximate solutions can often be determined and a variety of software tools are available to compute solutions. Unsurprisingly then, linear integer programming has also been used in the context of invasive species management (Epanchin-Niell and Wilen, 2012; Kıbış and Büyüktahtakın, 2017).

A crucial step in linear integer optimization is to find a problem formulation that facilitates good approximations (Ageev and Sviridenko, 2004). In this paper, we provide such a formulation to optimize locations and operating times of watercraft inspection stations. This problem differs from comparable resource allocation problems (Surkov et al., 2008; Trullols et al., 2010), as we need to account for temporal variations in traffic. These variations are key when considering the trade-off between operating a few inspection stations intensely, e.g. around the clock, and distributing resources over many locations operated at peak traffic times only.

We demonstrate the potential of our approach by applying it to optimize watercraft inspection

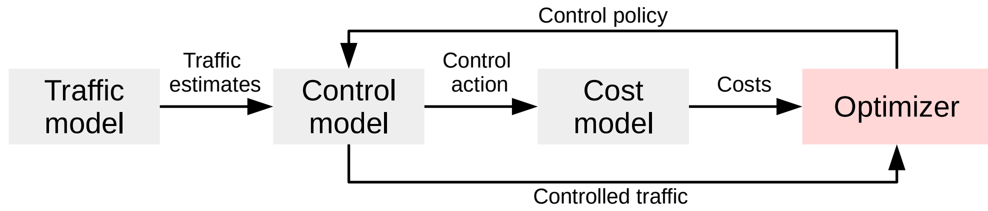

Figure 1: Components of our approach. The control model determines how the traffic estimated by the traffic model changes under a given control policy. The cost model yields the costs for control actions. The optimizer maximizes the controlled traffic subject to a cost constraint.

policies for the Canadian province of British Columbia (BC). We show how uncertainty, different cost constraints, and additional propagule sources impact the optimal policy. Based on this, we identify control principles applicable beyond the specific case study.

2 Method

2.1 Model

We assume that the control strategy counteracting the introduction of AIS can be adjusted in two aspects: the locations and operating times of watercraft inspection stations. As traffic typically follows cyclic patterns, we consider one such cycle as the time horizon for the control optimization. To find an optimal inspection policy, we need three models (Figure 1): (1) a model for boater traffic, (2) a model for control, and (3) a model for control costs. The traffic model gives us estimates of when, where, and along which routes boaters travel. The control model shows us when and where inspections could be conducted and how effective they are. Lastly, the cost model measures the costs for inspections. The information from the three models serve as input for a control optimizer that determines a good – or, if possible, the best – watercraft inspection strategy. Below, we describe each of the models in greater detail before introducing suitable optimization routines in the next section.

2.1.1 Traffic model

The traffic model provides estimates of when, where, and along which routes boaters drive. Knowledge about routes is key to understanding whether boaters passing one control location have already been inspected at another location. For each considered route, the traffic model provides us with a traffic estimate. In this study, we use a hybrid gravity and route choice model (Fischer et al., 2020) to estimate traffic. The model includes components accounting for boaters' travel incentive, their route choice, the timing of traffic, and boaters' compliance with inspections.

In practice, it is rarely feasible to consider all routes that boaters could possibly take, so we need to focus on some set of "reasonable" routes (Bovy, 2009; Fischer, 2020). As a consequence, there may be some boaters travelling along unexpected routes. When boaters travelling along such routes arrive at inspection locations, we do not know whether their watercraft have been inspected earlier. This makes it difficult to optimize inspection strategies. Nonetheless, we may want to account for these boaters by introducing a "noise" term to our model. To that end, we assume that a fraction of the travelling boaters could be observed at any inspection location with a small probability (see Fischer et al., 2020).

As road traffic is rarely uniform over time, we furthermore need a submodel predicting how traffic varies with time. While it may be comparatively easy to determine the temporal distribution of traffic at a specific location, it can be difficult to identify the temporal relationship between traffic at two locations on the same route. For example, boaters passing one location in the morning may not be able to reach another location before the afternoon. Modelling such relationships is particularly difficult for locations far from each other, as boaters may have different travel speeds. We therefore apply a simplification and assume all boaters travelling along a route have the same speed.

2.1.2 Control model

We assume that there is a specific set of locations where watercraft inspections could be conducted. For example, these locations could be pullouts large enough to provide a safe environment for inspections. We suppose that compliant boaters stop for an inspection whenever they pass an operated inspection station. Conversely, uncompliant boaters are assumed to bypass every inspection station on their route. Consequently, we seek to maximize the number of boaters that pass at least one operated watercraft inspection.

We further assume that there are specific time intervals when inspections can be conducted. The available time intervals may be determined by safety concerns or practical considerations and can be location dependent. As staff cannot move between distant inspection locations easily and the working hours of inspection staff are subject to legal and practical constraints, we further assume that every inspection station can be operated in shifts of predetermined lengths only.

2.1.3 Cost model

Control costs may be split into two classes: infrastructure costs that apply once for each chosen inspection location and operational costs, which depend on when and for how long an inspection station is operated. The operational costs may also account for ongoing equipment maintenance costs and training of staff. The control costs may be location and time dependent. For example, it may be expensive to conduct inspections at remote locations if staff must travel long distances to their work place. Furthermore, some locations will require significantly more infrastructure costs (e.g. lighting and washrooms) in order to operate overnight shifts. In addition, wages are often higher in overnight shifts.

2.2 Optimizing control locations

Given the submodels described in the previous section, we proceed with describing how to optimize the inspection strategy. Optimizing both locations and operating times of watercraft inspection stations at the same time is conceptually and computationally challenging. To ease the introduction of our approach, we first consider a scenario in which inspection stations are operational around the clock. In this case, we can ignore the temporal variations of traffic and focus on only choosing optimal control locations (cf. Trullols et al., 2010).

We formulate the corresponding optimization problem as a linear integer problem. To that end, we let L be the set of all admissible inspection locations and introduce for each location $l \in L$

a binary variable x_l that assumes the value 1 if and only if an inspection station is set up at l. Let R be the set of potential routes that boaters may choose, n_r the expected number of complying boaters travelling along route $r \in R$, and $L_r \subseteq L$ the set of locations where boaters travelling on route r could be inspected.

As noted earlier, one inspection station suffices to control all complying boaters driving along a route r. Consequently, boaters travelling on route r will be controlled if and only if

$$\sum_{l \in L_r} x_l \geq 1. \tag{1}$$

Therefore, we can express the total number of inspected boaters by

$$F_{\text{loc}}(\boldsymbol{x}) := \sum_{r \in R} \min \left\{ 1, \sum_{l \in L_r} x_l \right\} n_r. \tag{2}$$

To formulate the cost constraint, let c_l be the cost for operating control location $l \in L$ and B the available budget. As we assume that all inspection stations are operated for the same time, we do not need to distinguish between infrastructure and operation costs. Hence, we can write the cost constraint as

$$\sum_{l \in L} c_l x_l \leq B. \tag{3}$$

The optimal placement policy can be identified by maximizing $F_{\text{loc}}(\boldsymbol{x})$ over all $\boldsymbol{x} \in \{0,1\}^{|L|}$ subject to constraint (3). Though F_{loc} contains a "minimum" function, F_{loc} can be easily transformed to a linear function by introducing further variables and linear inequality constraints (see e.g. Ageev and Sviridenko, 1999). Since the left hand side of the cost constraint (3) is linear in $\boldsymbol{x}$ as well, and $\boldsymbol{x}$ is constrained to be a vector of integers, the considered optimization problem is a linear integer problem. This can be solved with a general linear integer programming solver (see Conforti et al., 2014) or a heuristic method (Ageev and Sviridenko, 2004). We discuss possible optimization routines in section 2.5.

2.3 Optimizing control locations and timing

We now extend our approach to permit free choice of inspection station operating times. In this generalized case, we need to balance the trade-off between operating a few highly frequented inspection stations around the clock and distributing efforts over many locations operated at peak traffic times only. This trade-off makes combined optimization of location choice and timing more challenging than separate optimization of location choice and timing (cf. Epanchin-Niell and Wilen, 2012).

While location choice is a discrete optimization problem – each potential inspection location is either chosen or not – optimization of operating times is a continuous problem, since inspections could in theory be started at any time. To exploit the toolset of discrete optimization, however, we simplify our problem by discretizing time. More specifically, we split boater departure times into disjoint time intervals and choose the timing of inspection shifts accordingly.

Let T be a set of disjoint time intervals covering the complete time span of interest. We write n_{rt} for the expected number of boaters willing to comply with inspections who travel on route $r \in R$ and start their journey in time interval $t \in T$. Furthermore, let S_l be the set of admissible inspection shifts for location $l \in L$. Each shift corresponds to a time interval in which the inspection station is operated. Since shift lengths are given, S_l can be fully characterized by the shifts' start times.

As we assume that all boaters travelling along a route have the same speed, we can determine the set $S_{lrt} \subseteq S_l$ of control shifts during which boaters who started their journey in time interval $t \in T$ arrive at location $l \in L$ via route $r \in R$. Under reasonable error allowance, it is possible to construct the sets S_{lrt} in a way that each shift covers the departure time intervals either completely or not at all, respectively. This setup prevents issues arising if some intervals overlap only partially.

To formulate our optimization problem as a linear integer problem, we describe the control policy again with binary variables $x_{ls} \in \{0, 1\}$. Here, x_{ls} is 1 if and only if an inspection station at location $l \in L$ is operated in shift $s \in S_{lrs}$. Boaters travelling on route $r \in R$ who departed in time interval $t \in T$ are controlled if and only if

$$\sum_{l\in L_r}\sum_{s\in S_{lrt}} x_{ls} \quad \geq \quad 1. \tag{4}$$

Consequently, the total controlled boater flow is given by

$$F_{\text{full}}(\boldsymbol{x}) \quad := \quad \sum_{r\in R}\sum_{t\in T} \min\left\{1, \sum_{l\in L_r}\sum_{s\in S_{lrt}} x_{ls}\right\} n_{rt}. \tag{5}$$

To derive the cost constraint, recall that we distinguish between infrastructure costs c_l^{loc} for using location l and operating costs c_{ls}^{shift} payable per inspection shift s conducted at l. Consequently, the total costs for control at l are given by

$$\sum_{s\in S_l} c_{ls}^{\text{shift}} x_{ls} + c_l^{\text{loc}} \max_{r\in R,\, t\in T}\left(\sum_{s\in S_{lrt}} x_{ls}\right), \tag{6}$$

and the cost constraint reads

$$\sum_{l\in L}\left(\sum_{s\in S_l} c_{ls}^{\text{shift}} x_{ls} + c_l^{\text{loc}} \max_{r\in R,\, t\in T}\left(\sum_{s\in S_{lrt}} x_{ls}\right)\right) \quad \leq \quad B. \tag{7}$$

As in the previous section, B denotes the available budget. Optimizing F_{full} subject to (7) is a linear integer problem, since the "minimum" term in (5) and the "maximum" terms in (7) can be replaced by introducing constrained auxiliary variables.

2.4 Boaters travelling on unknown routes

Even if the traffic model accounts for most routes boaters use, some boaters may travel along unexpected routes. It is difficult to optimize inspection station operation with regards to these boaters, as we do not know which inspection stations cover the same routes. Nonetheless, it can be desirable to account for boaters travelling on unknown routes, since the level of uncertainty may affect the optimal inspection policy.

In the absence of a mechanistic model for traffic along unknown routes, we assume that boaters travelling on unexpected routes are passing any inspection location with a small probability η_o and choose their passing time randomly. Let τ_{sl} be the probability that a boater passing location l on an unexpected route does so during shift s. Then, the expected number of inspected boaters travelling along unknown routes is given by

$$F_{\text{noise}} = \left(1 - \prod_{l \in L}\left(1 - \eta_o \sum_{s \in S_l} x_{ls}\tau_{sl}\right)\right) n_{\text{noise}}. \tag{8}$$

Here, n_{noise} denotes the expected number of boaters travelling on unknown routes.

Adding the non-convex function F_{noise} to the objective function would make optimization difficult, but after introducing additional variables and constraints, F_{noise} could be replaced with an equivalent linear term (Morton et al., 2007; O'Hanley et al., 2013). Alternatively, we may linearize F_{noise} via a Taylor expansion (see Supplementary Appendix A): since η_0 is typically small, equation (8) is well approximated by

$$\hat{F}_{\text{noise}} = \eta_o n_{\text{noise}} \sum_{l \in L} \sum_{s \in S_l} x_{ls}\tau_{sl}, \tag{9}$$

which can be incorporated in a linear integer problem. In this study, we apply approximation (9), which is most precise if $x_{ls} = 0$ for most l and s. If the budget is high enough to operate many inspection stations for long times, the noise term may be overestimated. However, if n_{noise} is small compared to the total boater traffic, the approximation error is unlikely to alter the overall optimization results significantly.

2.5 Solving the optimization problems

The inspection station placement problem described in section 2.2 is equivalent to the budgeted maximum coverage problem (Khuller et al., 1999), also called maximum coverage problem with knapsack constraint (Ageev and Sviridenko, 2004). This problem is well studied in computing science and difficult to solve exactly in general (Feige, 1998). Nonetheless, good approximate or even optimal solutions can often be obtained in practical applications.

We solve the full optimization problem using a branch and bound algorithm (see Conforti et al., 2014), in which the distance between upper and lower bounds on the optimal objective is reduced until the gap is satisfactorily small. Branch and bound algorithms are commonly used to solve linear integer programming problems such as ours, and efficient implementations are included in many optimization packages. Branch and bound algorithms are most effective if they can improve upon a good initial starting solution. To find such an initial starting solution, we implement a "greedy" rounding algorithm that (1) repeatedly solves a relaxed version of our optimization problem, permitting non-integer decision variables $x_{ls} \in [0, 1]$, and (2) rounds up the respective largest non-integer decision variable x_{ls} that can be set to 1 without violating the cost constraint. In Supplementary Appendix B, we describe the procedure in greater detail.

3 Application

To demonstrate the potential of our approach, we applied it to optimize watercraft inspections in the Canadian province of British Columbia (BC). Below we provide an overview of the case-study-specific submodels we used. Furthermore, we briefly describe our implementation of the presented approach.

3.1 Case-study-specific submodels

3.1.1 Traffic model

To model boater traffic, we used the hierarchical gravity and route choice model for boater traffic presented in Fischer et al. (2020). The model was fitted using data collected at British Columbian watercraft inspection stations in the years 2015 and 2016. At the time this study was conducted, dreissenid mussels were not known to be established anywhere in BC. As sources of potentially infested boaters, we therefore considered the Canadian provinces and American states that (1) were known to be invaded by dreissenid mussels or (2) had connected waterway to an infested jurisdiction and no coordinated mussel detection program in place at the time the data were collected. As sinks we identified 5981 potentially boater accessible lakes in BC.

To estimate the boater traffic between an origin and destination, the model considered characteristics of the donor jurisdiction, the recipient lake, and the distance between the two. Major sources of high-risk boaters were characterized by high population counts. Furthermore, Canadian provinces were found to have higher boater traffic to BC than American states. Attractiveness of destination lakes increased with their surface area, the population counts of surrounding towns and cities, and the availability of close-by touristic facilities, such as campgrounds. Lastly, the boater flow was estimated to decay in cubic order of the distance between an origin and a destination. For a detailed description of the model along with precise parameter estimates, refer to Fischer et al. (2020).

To identify potential boater pathways, we computed locally optimal routes (Fischer, 2020) between the considered origins and destinations. These routes arise if routing decisions on local scales are rational and based on simple criteria (here: minimizing travel time) whereas unknown factors may affect routing decisions on larger scales. Consequently, the model accounts for routes arising from a multitude of mechanisms. The attractiveness of the routes was computed based on their length measured in travel time, because boaters have a strong incentive to carry larger trailered watercraft on short and well-developed routes. Again, a more in-depth description of the model and the fitted parameter values can be found in Fischer et al. (2020).

The fraction η_c of boaters travelling on routes not covered by our traffic model was estimated as 4.9%. However, this number is not estimable from survey data obtained at watercraft inspection stations alone because it is negatively correlated with the parameter η_o (section 2.4), which denotes the probability to observe a boater travelling on an unknown route at an arbitrary inspection location. Therefore, Fischer et al. (2020) introduced an additional model assumption bounding η_c below 5%. Note that due to the dependency of η_o on η_c, the estimability issue has little effect on how many boaters driving on unknown routes are observed at watercraft inspection stations. Hence, the estimability problem does not affect the inspection policy significantly. Based on a noise level of $\eta_c = 4.9\%$, Fischer et al. (2020) estimated $\eta_o = 0.06$.

The temporal distribution of traffic was modelled with a von Mises distribution. This is a unimodal circular distribution often used in models (Lee, 2010). The temporal pattern was assumed to have a period of one day. The traffic high was estimated to be at 2 PM, and the estimated peak traffic was 15 times higher than the estimated traffic volume at night. As traffic data were available for specific inspection locations only, the temporal traffic distribution was assumed to be the same for all locations, thereby neglecting the time boaters need to travel between sites. This limitation can be addressed by shifting the obtained operating times according to local traffic patterns after the optimization.

3.1.2 Control model

As described in section 2.1.2, we assume that every complying boater passing an operated inspection location is inspected for invasive mussels. The compliance rate across all inspection stations was estimated to be 80% (Fischer et al., 2020). To find potentially suitable locations for inspections, we identified pullouts across BC. We reduced the number of possible options by disregarding some pullouts in close proximity to others. We included currently used locations as candidates without giving them any preference over other locations. In total, we considered 249 candidate locations.

Due to the large number of candidate locations, we did not conduct a detailed evaluation of the operational suitability of all considered locations (e.g. pullout size, signage, and safety). Instead, we consulted with the BC Invasive Mussel Defence Program to gauge the general suitability of the locations suggested by the optimizer. If a suggested location seemed unsuitable, we removed it from the candidate set and repeated the optimization procedure. Despite this superficial suitability check, a more detailed analysis would be necessary to account for all potential practical constraints. These must be considered independent of the model before an inspection station can be potentially located.

For each location, we assumed that 8 hours long inspection shifts could be started at each full hour of the day. Note that "shift" here refers to the time inspections are conducted and does not include time required for staff to access or set up an inspection station. The work time of staff will therefore be longer in practice. The assumed length of the inspection shifts aligns with average operation patterns of watercraft inspection stations in BC and divides each day into three equally sized shifts, which simplifies the model. Though the effective operation time (limited by access time of staff) is lower at remote locations, our time model provides a good first approximation.

3.1.3 Cost model

We determined the inspection costs based on correspondence with the BC Invasive Mussel Defence Program. Since the considered optimization problem is often easier to solve if costs are rounded to well aligned cost units, we expressed costs in a custom monetary unit, with one cost unit representing the infrastructure costs for setting up an inspection station. In terms of this cost unit, the costs per conducted inspection shift are 3.5 units during daytime and 5.5 units between 9 PM and 5 AM. These operational costs include salary, training, and equipment for inspection staff. In 2017, the BC Invasive Mussel Defence Program was operating on a budget of approximately 80 cost units.

As an in-depth location-specific cost analysis would have been difficult, we assumed that the inspection costs are equal for all considered locations. Note, however, that site specific costs can vary significantly and may be a limitation when assessing a location for overnight operations.

3.2 Implementation

As we considered approximately $300,000$ origin-destination pairs connected by 6.7 routes on average, considering all boater pathways individually would be difficult. Therefore, we merged traffic of boaters passing the same sets of potential inspection locations. As a result, the number of distinct boater flows reduced to 2026.

We determined optimal inspection locations and operating times under different budget constraints. This allowed us to determine the budget required to minimize the fraction of uninspected high-risk boaters to a desired level. We also varied the fraction of boaters travelling on unknown routes to test how inspection strategies change under increased uncertainty. To see how new infestations in close-by jurisdictions change the inspection policy, we furthermore considered a scenario in which the American states Idaho, Wyoming, and Oregon are invaded. Note that boaters from Montana were considered high-risk in the base scenario already. Finally, we determined how the US-Canadian border closure in response to the global COVID-19 pandemic in 2020 affected the optimal inspection policy.

We implemented the model in the high-level programming language Python version 3.7. To formulate the linear integer problem, we used the modelling software CVXPY version 1.0.25 with added support for initial starting solutions. To solve the linear integer problem, we used the commercial solver MOSEK. We computed initial starting solutions with the greedy rounding procedure described in Supplementary Appendix B. We let the solver terminate if a solution was found that was guaranteed to be less than 0.5% below the optimum or if 50 minutes had passed. We conducted the computations on a Linux server with a 20 core Intel Xeon 640 E5-2689 CPU (3.1GHz per CPU) and with 512GB RAM. The computer code can be retrieved from vemomoto.github.io.

4 Results

In 72% of the considered scenarios, we were able to identify a solution with the desired accuracy. In the remaining cases, the relative optimality gap never exceeded 8.7%; i.e., the optimal solution was guaranteed to be no more than 8.7% above the solution we found. In scenarios with budgets $B \geq 25$ units, we could always identify solutions with optimality gap below 2%. The initial starting solutions computed via the greedy algorithm reached the desired final accuracy in 58% of the considered cases and never exceeded an optimality gap of 11%.

Figure 2 displays the optimized locations and operating times for watercraft inspection stations in the considered case study. We depict the optimal policy for scenarios with a budget of 25, 50, and 100 cost units respectively. The optimal locations for inspections are located close to border crossings if suitable locations are available. However, where the traffic through many border crossings merges on a major highway (e.g. in the Vancouver metropolitan area), it is optimal to place the inspection stations farther inland.

Figure 3 depicts characteristics of the optimal inspection stations for various model scenarios. The optimized operating times are correlated with the expected traffic volume at inspection stations. Additional budget is preferably spent on additional inspection locations rather than longer operating hours. However, larger portions of the additional budget are spent on additional locations (see also Figure 2).

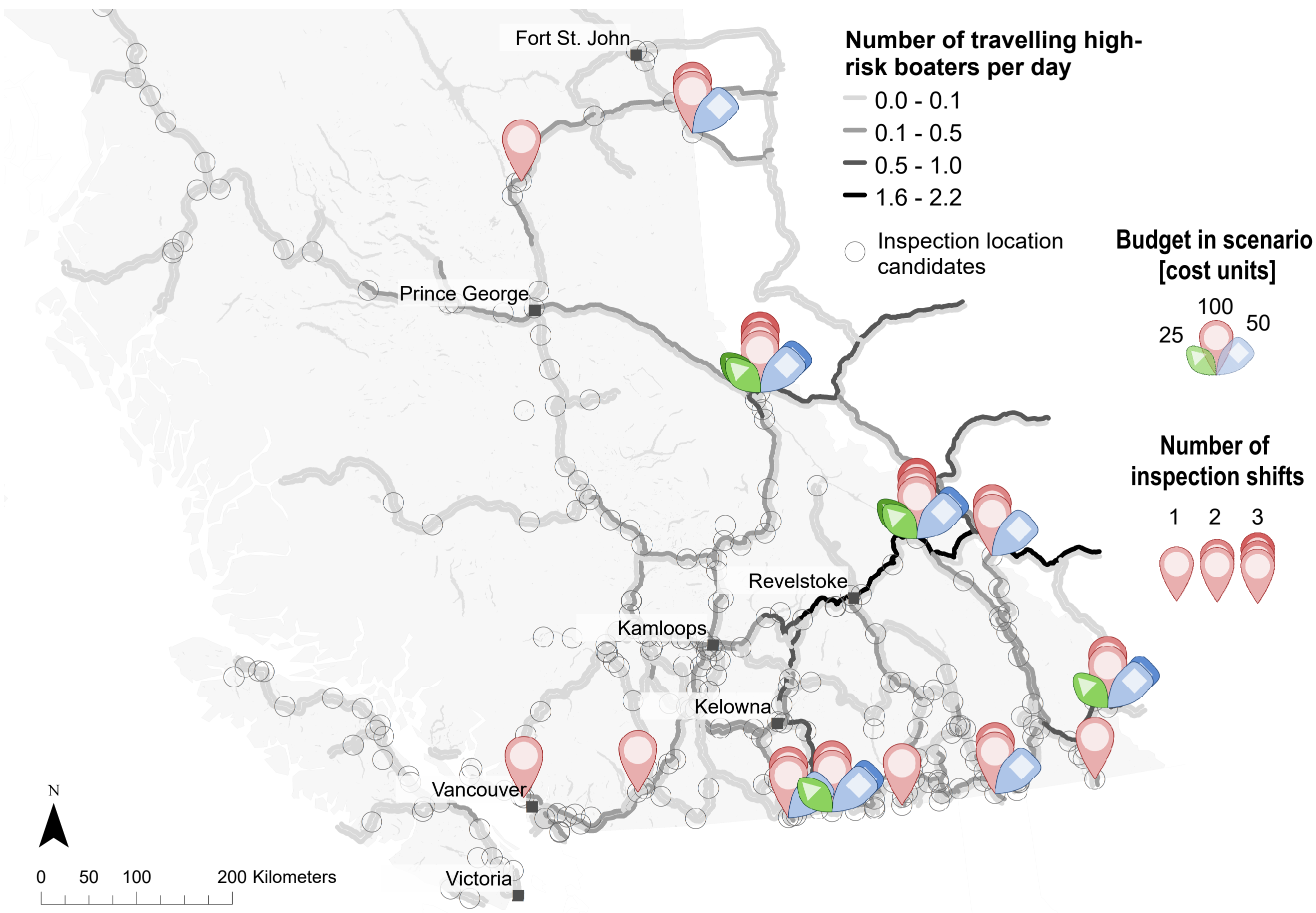

Figure 2: Optimal locations and operation shifts for three different budget scenarios. Most inspection stations are placed close to the British Columbian border. The markers depict the optimal inspection locations for each scenario. Green (triangle): optimal locations with a budget of 25 units; blue (square) 50 unit budget; red (circle) 100 unit budget. The number of markers stacked on top of each other corresponds to the optimal numbers of inspection shifts. The darkness of the roads show the estimated boater traffic volume. The hollow circles depict the considered candidates for inspection locations.

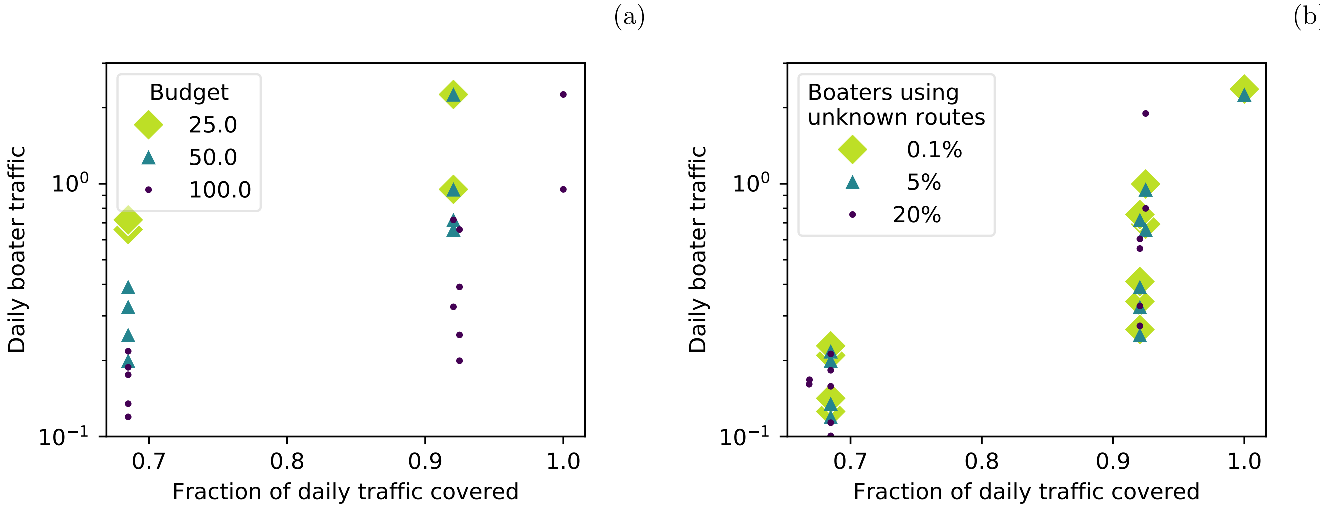

Figure 3: Characteristics of the optimized inspection stations in scenarios with (a) different budget constraints, and (b) different fractions of boaters travelling on unknown routes (cost bound fixed at 80 units). Each marker corresponds to an inspection station. The position of the marker depicts the daily traffic volume expected at the location and the fraction of daily traffic covered under the optimal operation policy (compliance supposed). Note that the traffic volume at the inspection locations changes with the fraction of boaters driving on unknown routes.

If the uncertainty in the traffic predictions increases significantly, more inspection stations are set up at the cost of shorter operations. Little changes of the fraction of boaters travelling on unknown routes (e.g. from 0.001 to 0.05) have no effect, whereas increasing this fraction from 0.05 to 0.2 results in two additional inspection locations being used (+18%) at the expense of reduced shifts at two inspection stations and slightly altered operating times at two other stations.

Optimizing inspection station operation under a range of different budget allowances showed that a moderate inspection budget, corresponding to about half the 2017 BC inspection budget, suffices to inspect half of the incoming high-risk boaters (Figure 4). However, the resources required for inspections increase quickly if more boaters shall be controlled. The fraction of inspected boaters is limited by boaters' compliance with inspections.

The considered change in the invasion state of three American states had only a moderate impact on inspection policy. As the additional propagule sources were located south of BC, the inspection effort increased at the southern border under the optimal policy. Furthermore, the optimal policy contained no overnight inspection at the eastern border of BC and distributed resources more evenly across inspection stations.

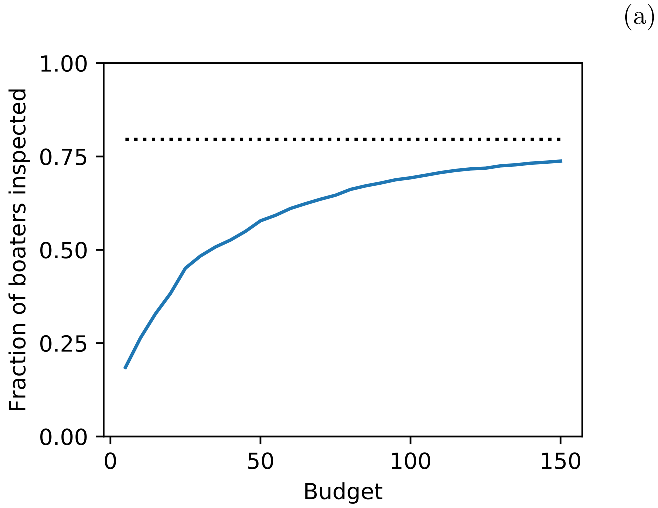

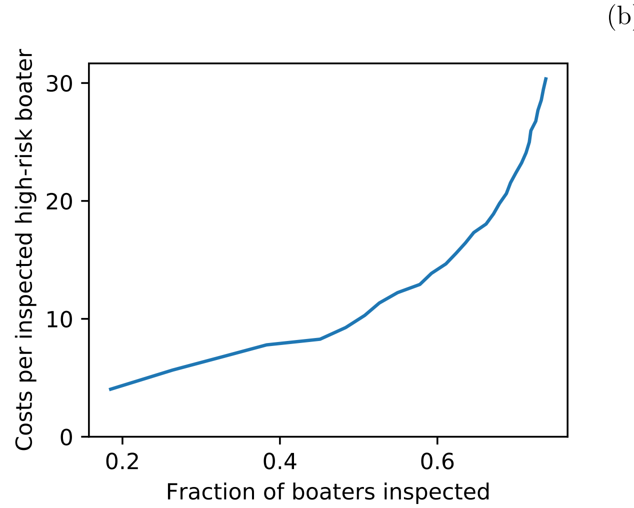

Figure 4: Inspection effectiveness dependent on the budget constraint (a) and price per inspected high-risk boater dependent on the proportion of inspected boaters (b). While a large fraction of high-risk boaters can be covered with moderate effort, inspecting all complying boaters is costly. Panel (a) shows the expected fraction of incoming high-risk boaters that can be inspected under the optimal policy. The dotted line shows the level of complying boaters, which is the maximal fraction of boaters that can be inspected.

If the US-Canadian border is closed, all boater traffic concentrates on the eastern provincial border crossings. Consequently, the inspection effort increases at these locations. If a high budget is available (e.g. 100 units), all major boater routes can be covered, and the distribution of the excess budget depends strongly on the submodel for boaters travelling on unknown routes. Figures depicting the results obtained for the invasion and border closure scenarios can be found in Supplementary Appendix C.

5 Discussion

We present a method to optimize placement and operating times of watercraft inspection stations. The approach is suited to model management scenarios on a detailed level and gives specific advice for management actions. We applied our approach to invasive mussel management in BC and investigated the impact of budget constraints, model uncertainty, potential future invasions, and border closures on management actions and efficiency. Though our model did not account for all critical operational factors, such as site safety and jurisdiction, the presented results provide valuable insights into optimal management of AIS when combined with critical operational factors.

Most of our results agree with common sense. In general, it is optimal to inspect boaters as soon as they enter the managed region. That way, waterbodies close to the border can be protected. If multiple routes via different border crossings merge close to the border, it can be optimal to inspect boaters after this merging point. Inspection stations should operate longer at locations with high traffic volume. Furthermore, uncertainty in traffic predictions increases the benefit of spreading the inspection efforts over many locations. Driven by these simple principles, our results were remarkably robust throughout considered scenarios and agree well with the watercraft inspection policy currently implemented in BC.

While these qualitative principles may seem obvious, it can be challenging to identify quantitative definitions of terms like "close to the border" and "longer". The difficulty in optimizing management policies is in balancing trade-offs, such as between leaving some waterbodies close to the border unprotected and maximizing the overall number of inspected boaters, or between long-time operation of few highly frequented inspection stations and distribution of resources over many locations. As the approach proposed in this paper is suited to account for these trade-offs, it is a valuable extension to earlier more theoretical results on AIS management (Potapov and Lewis, 2008; Potapov et al., 2008; Finnoff et al., 2010; Warziniack et al., 2011), which typically do not include enough detail to answer specific operational questions. Since the presented approach is easily adjusted to new circumstances, it can furthermore help managers to respond rapidly and effectively to short-term changes in traffic patterns, such as experienced due to southern border closures in the global COVID-19 pandemic.

Considering scenarios with different budget constraints allowed us to investigate the trade-off between resources invested in AIS control and the number of inspected high-risk boaters. In combination with the expected monetary damage caused by the arrival of an uncontrolled boater at an uninvaded lake, this trade-off curve can be used to identify the optimal budget for inspections. Since both invasion risk and damages due to invasions are difficult to quantify, a rigorous computation of the optimal inspection budget may not always be feasible in practice. Nonetheless, the cost-effectiveness curve provides an estimate of the efficacy of control efforts and shows which budget is required to achieve a certain management goal.

In the case of AIS control in BC, a moderate budget suffices to inspect a significant portion of the incoming high-risk boaters. This is because boater traffic in BC concentrates on a small number of major highways. Nevertheless, inspecting all high-risk boaters would be very costly, as many minor roads would have to be considered as well. It could, therefore, be more cost-effective to implement measures to increase the compliance of boaters at existing inspection locations. These measures can include additional / improved road signs, public outreach and education, and increased presence of enforcement officers at inspection locations. In BC, it is mandatory for boaters to stop at open inspection stations along their travel routes, and tickets and fines are issued to uncompliant boaters. The compliance rates are tracked continuously, and with the implementation of the above measures, the BC program has seen an increase in compliance to 83% in 2019. Furthermore there is an added layer of protection provided by inspection programs operating in other western Canadian provinces and US states.

We see a particular contribution of our approach in supporting and evaluating concrete management decisions. Often, significant funds are allocated for invasive species management with mostly qualitative and heuristic research insights at hand. Quantitative evaluation of the implemented policy is thus key to ensure that the resources are used effectively. This, in turn, requires a spatially explicit approach incorporating enough details to account for the trade-offs between available options. Our approach adds on to previously developed methods by incorporating these details. Our results confirming the effectiveness of the British Columbian watercraft inspection policy are a valuable contribution to justify the implemented policy to stakeholders and funders.

Besides evaluating existing policy, we see particular use of our approach for optimizing rapid response actions under scenarios of interest. The scenarios considered in this paper show that slight adjustments to the inspection policy may suffice to react on new invasions, whereas stronger adjustments are needed when border closures alter traffic patterns strongly. Our approach could also be used to assess the benefit from cross-border and and cross-agency collaborations, in which inspection efforts are combined to control the boater inflow to a large joint area and at international border crossings. Due to the flexibility of our model, managers can consider a variety of scenarios at little cost.

5.1 Limitations and possible extensions

Our approach does not consider human behavioural responses to the inspection policy. Uncompliant boaters may try to avoid inspections by bypassing inspection stations on alternate routes or by travelling outside the operating hours. In the former case, identifying locations that cannot be bypassed easily becomes more important; additional inspection stations might be necessary to cover alternative routes of uncompliant boaters. If boaters take the inspection policy into account when they evaluate the quality of routes, different routes would need to be computed for each tested set of inspection locations. This is computationally expensive and incompatible with the optimization routines we used. As for our case study in BC, however, the road network is relatively sparse, and by placing inspection stations at key entry points into the province, the detours required to bypass inspection stations typically take longer than watercraft inspections.

The accuracy of our approach is strongly dependent on the accuracy and level of detail of the utilized data and models. Therefore, the results should be combined with expert knowledge and refined iteratively if necessary. For example, our model did not incorporate site-specific costs and operational constraints. In some scenarios, the model thus suggested overnight inspections at remote sites that are lacking the required infrastructure to safely operate at night, e.g. proper road infrastructure (lanes/barriers), lighting, access to safe communication and nearby living accommodations for staff. Such constraints could be incorporated in a more detailed model as well as increased costs at remote locations. A more detailed model could also account for inspection stations operated by other jurisdictions or federal agencies. Nonetheless, the presented model includes major factors affecting inspection station operation and can serve as a helpful resource to inform managers' decisions in parallel with operational constraints.

Limitations exist with respect to the objective function we considered. Though the number of potentially infested watercraft arriving at a waterbody is a valuable proxy for invasion risk, the establishment probability of dreissenid mussels is a more complicated, nonlinear function of the boater inflow (Leung et al., 2004). Minimizing this function is a computationally more difficult optimization problem, however, and minimizing a proxy for invasion risk instead may be the preferred option in many practical applications. If desired, the suitability of the destination waterbodies as habitat for AIS could be considered in our linear model by weighting boater flows corresponding to the AIS establishment risk at the destination waterbodies.

Our model could also be extended to incorporate location-specific and management-dependent compliance rates. At certain sites, such as cross-national border crossings, compliance can be enforced more easily than at other locations. Compliance may furthermore depend on management efforts: it may be possible to increase the compliance rate of boaters at some costs. In Supplementary Appendix D, we show how location-dependent and flexible compliance rates can be considered with small model adjustments. The idea is to consider boaters with different compliance levels separately from one another and to extend the objective function (5) by a multiplier modelling boater compliance dependent on management effort.

Despite the high level of detail incorporated in our model, some factors remain difficult to consider. Since our traffic model does not explicitly account for the time boaters need to travel between locations, the optimized inspection station operating times may have to be adjusted to local traffic patterns. This correction may not fully resolve the issue if boaters pass multiple inspection stations under the optimal policy, but such configurations are rarely optimal in practice. In our case study, for example, inspection stations were optimally operated during traffic peak hours, which indicates that boaters passing multiple inspection stations did not affect the optimal operating times significantly.

Another modelling challenge is to appropriately account for boaters travelling along unexpected routes. As our submodel for this traffic is not based on road topology, some traffic is predicted on inland roads even if all border crossings are covered with inspection stations. This can bias the results in high-budget scenarios, in which boaters can be inspected at most border crossings, thereby also covering traffic along unexpected routes. In such cases, more meaningful results may be obtained by disregarding traffic along unexpected routes (see Supplementary Appendix C). To improve the submodel for traffic along unexpected routes, it could be adjusted to reflect that boater traffic is typically higher on major highways than on minor roads. In most scenarios of our case study, however, the results were not very sensitive to the traffic along unexpected routes, and an extended model may not yield significantly different inspection policies.

The numerical optimization techniques we applied gave good results in this study, but there may be cases in which linear integer solvers fail to provide good solutions. This may happen, for example, if many boaters pass multiple inspection stations under the optimal policy. For a more detailed discussion of such cases and their prevalence in real-world applications refer to Supplementary Appendix E.

5.2 General conclusions for invasive species management

In this paper, we considered specific management scenarios with focus of AIS control in BC. Nonetheless, some common patterns were consistent throughout our results and may thus apply with greater generality. These principles may be used as rules of thumb if no comprehensive modelling and optimization effort is possible. Below we summarize these conclusions.

- Inspection stations should be placed close to the border of the uninfested region. Consequently, cross-border collaborations between uninvaded jurisdictions have a high potential of improving the cost-effectiveness of control.
- If traffic flows merge close to the border, inspections are more cost-effective after the merging point. Hence, identifying such points is crucial for successful management.
- If traffic predictions involve a high level of uncertainty, inspection efforts should be distributed over many locations at the cost of lower inspection effort at each site.
- If a high reduction of the propagule inflow is desired, it may be most cost-effective to implement measures increasing the compliance rate rather than operating more inspection stations for longer hours.

These conclusions agree well with earlier results describing the importance of the geometry of control (Potapov and Lewis, 2008; Epanchin-Niell and Wilen, 2012), the benefits of collaborative management (Epanchin-Niell et al., 2010; Epanchin-Niell and Wilen, 2015), and the effects of uncertainty on the control policy (Finnoff et al., 2010; Hall et al., 2018).

Authors' contributions

All authors conceived the project; SMF conceived the methods jointly with MAL. SMF and MB jointly prepared the data for the analysis. SMF conducted the mathematical analysis, implemented the model, and wrote the manuscript. All authors revised the manuscript.

Acknowledgements

The authors would like to give thanks to the BC Ministry of Environment and Climate Change Strategy staff of the BC Invasive Mussel Defence Program, who conducted the survey this study is based on. Furthermore, the authors thank the members of the Lewis Research Group at the University of Alberta for helpful feedback and discussions. SMF is thankful for the funding received from the Canadian Aquatic Invasive Species Network and the Natural Sciences and Engineering Research Council of Canada (NSERC); MAL gratefully acknowledges an NSERC Discovery Grant and Canada Research Chair.

Competing Interests

The authors declare no competing interests.

References

Ageev, A. and Sviridenko, M. (2004). Pipage rounding: a new method of constructing algorithms with proven performance guarantee. *Journal of Combinatorial Optimization*, 8(3):307–328.

Ageev, A. A. and Sviridenko, M. I. (1999). Approximation algorithms for maximum coverage and max cut with given sizes of parts. In Cornuéjols, G., Burkard, R. E., and Woeginger,

G. J., editors, *Integer Programming and Combinatorial Optimization*, volume 1610, pages 17–30. Springer Berlin Heidelberg, Berlin, Heidelberg.

Alberta Environment and Parks Fish and Wildlife Policy (2015). Alberta Aquatic Invasive Species Program 2015 annual report. Technical report, Edmonton, AB.

Bovy, P. H. L. (2009). On modelling route choice sets in transportation networks: a synthesis. *Transport Reviews*, 29(1):43–68.

Conforti, M., Cornuéjols, G., and Zambelli, G. (2014). *Integer programming*, volume 271 of *Graduate Texts in Mathematics*. Springer International Publishing, Cham.

Epanchin-Niell, R. S., Hufford, M. B., Aslan, C. E., Sexton, J. P., Port, J. D., and Waring, T. M. (2010). Controlling invasive species in complex social landscapes. *Frontiers in Ecology and the Environment*, 8(4):210–216.

Epanchin-Niell, R. S. and Wilen, J. E. (2012). Optimal spatial control of biological invasions. *Journal of Environmental Economics and Management*, 63(2):260–270.

Epanchin-Niell, R. S. and Wilen, J. E. (2015). Individual and cooperative management of invasive species in human-mediated landscapes. *American Journal of Agricultural Economics*, 97(1):180–198.

Feige, U. (1998). A threshold of ln n for approximating set cover. *Journal of the ACM*, 45(4):634–652.

Finnoff, D., Potapov, A., and Lewis, M. A. (2010). Control and the management of a spreading invader. *Resource and Energy Economics*, 32(4):534–550.

Fischer, S. M. (2020). Locally optimal routes for route choice sets. *Transportation Research Part B: Methodological*, 141:240–266.

Fischer, S. M., Beck, M., Herborg, L.-M., and Lewis, M. A. (2020). A hybrid gravity and route choice model to assess vector traffic in large-scale road networks. *Royal Society Open Science*, 7(5):191858.

Hall, K. M., Albers, H. J., Alkaee Taleghan, M., and Dietterich, T. G. (2018). Optimal spatial-dynamic management of stochastic species invasions. *Environmental and Resource Economics*, 70(2):403–427.

Hastings, A., Hall, R. J., and Taylor, C. M. (2006). A simple approach to optimal control of invasive species. *Theoretical Population Biology*, 70(4):431–435.

Hyytiäinen, K., Lehtiniemi, M., Niemi, J. K., and Tikka, K. (2013). An optimization framework for addressing aquatic invasive species. *Ecological Economics*, 91:69–79.

Inter-Ministry Invasive Species Working Group (2015). Zebra and quagga mussel early detection and rapid response plan. Technical report, Victoria, BC.

Johnson, L. E., Ricciardi, A., and Carlton, J. T. (2001). Overland dispersal of aquatic invasive species: a risk assessment of transient recreational boating. *Ecological Applications*, 11(6):1789–1799.

Johnson, R., Crafton, R. E., and Upton, H. F. (2017). Invasive species: major laws and the role of selected federal agencies. Technical report, Congressional Research Service, Washington, DC.

Karatayev, A. Y., Burlakova, L. E., and Padilla, D. K. (2015). Zebra versus quagga mussels: a review of their spread, population dynamics, and ecosystem impacts. *Hydrobiologia*, 746(1):97–112.

Khuller, S., Moss, A., and Naor, J. S. (1999). The budgeted maximum coverage problem. *Information Processing Letters*, 70(1):39–45.

Kıbış, E. Y. and Büyüktahtakın, İ. E. (2017). Optimizing invasive species management: a mixed-integer linear programming approach. *European Journal of Operational Research*, 259(1):308–321.

Lee, A. (2010). Circular data. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(4):477–486.

Leung, B., Drake, J. M., and Lodge, D. M. (2004). Predicting invasions: Propagule pressure and the gravity of Allee effects. *Ecology*, 85(6):1651–1660.

Leung, B., Lodge, D. M., Finnoff, D., Shogren, J. F., Lewis, M. A., and Lamberti, G. (2002). An ounce of prevention or a pound of cure: bioeconomic risk analysis of invasive species. *Proceedings of the Royal Society B: Biological Sciences*, 269(1508):2407–2413.

Lockwood, J. L., Hoopes, M. F., and Marchetti, M. P. (2013). *Invasion ecology*. Wiley-Blackwell, Chichester, West Sussex, UK, 2nd edition.

Mangin, S. (2011). The 100th Meridian Initiative: a strategic approach to prevent the westward spread of zebra mussels and other aquatic nuisance species. Technical Report 152, U.S. Fish and Wildlife Service, Arlington, VA.

Morton, D. P., Pan, F., and Saeger, K. J. (2007). Models for nuclear smuggling interdiction. *IIE Transactions*, 39(1):3–14.

O'Hanley, J. R., Scaparra, M. P., and García, S. (2013). Probability chains: A general linearization technique for modeling reliability in facility location and related problems. *European Journal of Operational Research*, 230(1):63–75.

Pimentel, D., Zuniga, R., and Morrison, D. (2005). Update on the environmental and economic costs associated with alien-invasive species in the United States. *Ecological Economics*, 52(3):273–288.

Potapov, A. B. and Lewis, M. A. (2008). Allee effect and control of lake system invasion. *Bulletin of Mathematical Biology*, 70(5):1371–1397.

Potapov, A. B., Lewis, M. A., and Finnoff, D. C. (2008). Optimal control of biological invasions in lake networks. *Natural Resource Modeling*, 20(3):351–379.

Rosaen, A. L., Grover, E. A., and Spencer, C. W. (2012). The costs of aquatic invasive species to Great Lakes states. Technical report, Anderson Economical Group, East Lansing, MI.

Shine, C., Kettunen, M., Genovesi, P., Essl, F., Gollasch, S., Rabitsch, W., Scalera, R., Starfinger, U., and ten Brink, P. (2010). Assessment to support continued development of the EU Strategy to combat invasive alien species. Final Report for the European Commission, Institute for European Environmental Policy (IEEP), Brussels, Belgium.

Surkov, I. V., Oude Lansink, A. G. J. M., van Kooten, O., and van der Werf, W. (2008). A model of optimal import phytosanitary inspection under capacity constraint. *Agricultural Economics*, 38(3):363–373.

Trullols, O., Fiore, M., Casetti, C., Chiasserini, C., and Barcelo Ordinas, J. (2010). Planning roadside infrastructure for information dissemination in intelligent transportation systems. *Computer Communications*, 33(4):432–442.

Turbelin, A. J., Malamud, B. D., and Francis, R. A. (2017). Mapping the global state of invasive alien species: patterns of invasion and policy responses. *Global Ecology and Biogeography*, 26(1):78–92.

Vander Zanden, M. J. and Olden, J. D. (2008). A management framework for preventing the secondary spread of aquatic invasive species. *Canadian Journal of Fisheries and Aquatic Sciences*, 65(7):1512–1522.

Warziniack, T., Finnoff, D., Bossenbroek, J., Shogren, J. F., and Lodge, D. (2011). Stepping stones for biological invasion: a bioeconomic model of transferable risk. *Environmental and Resource Economics*, 50(4):605–627.