Supplementary Appendices for "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.*

[*] *samuel.fischer@ualberta.ca*

A Linearization of the term accounting for boaters travelling on unexpected routes

In the main text, we use a linear approximation of

$$F_{\text{noise}} = \left(1 - \prod_{l\in L}\left(1 - \eta_o \sum_{s\in S_l} x_{ls}\tau_{sl}\right)\right) n_{\text{noise}}. \tag{A1}$$

to facilitate faster maximization of the objective function. Here, we show how this approximation can be computed. We start by noting that the parameter η_o is small. Hence, F_{noise} can be approximated via a Taylor expansion in η_o about $\eta_o = 0$. Let us write $F_{\text{noise}} = F_{\text{noise}}(\eta_o)$ as a function in η_o. Then,

$$F_{\text{noise}}(\eta_o) \approx F_{\text{noise}}(0) + F'_{\text{noise}}(0) \cdot \eta_o. \tag{A2}$$

With $f_l(\eta_o) := 1 - \eta_o \sum_{s \in S_l} x_{ls}\tau_{sl}$, we obtain by the product rule

$$\frac{d}{d\eta_o} \prod_{l \in L} f_l(\eta_o) = \sum_{l \in L} f'_l(\eta_o) \prod_{\tilde{l} \in L \setminus \{l\}} f_{\tilde{l}}(\eta_o) \tag{A3}$$

and thus

$$F'_{\text{noise}}(0) = n_{\text{noise}} \sum_{l \in L} \sum_{s \in S_l} x_{ls}\tau_{sl}. \tag{A4}$$

Inserting (A4) into (A2) yields

$$F_{\text{noise}}(\eta_o) \approx \eta_o n_{\text{noise}} \sum_{l \in L} \sum_{s \in S_l} x_{ls}\tau_{sl} =: \hat{F}_{\text{noise}}, \tag{A5}$$

which is the approximation given in equation (9) in the main text.

B Greedy rounding algorithm

In this Appendix, we describe the greedy rounding algorithm we applied to obtain initial starting solutions for the general branch and bound solvers. We start by introducing some helpful notation. Let P be the linear integer problem that we desire to solve and P_{cont} its continuous relaxation, in which decision variables x_{ls} may attain fractional values. In contrast to the integer problem P, which is NP-hard, the relaxed problem P_{cont} can be solved efficiently with linear programming techniques. We write $\boldsymbol{x}$ for the N-dimensional vector of decision variables, indexed by $(l, s) \in L \times S$. Let $\boldsymbol{e}_{ls}$ be a unit vector that is 0 everywhere except for the component corresponding to the index (l, s). Suppose that $C(\boldsymbol{x})$ denotes the cost for implementing a policy given by $\boldsymbol{x}$. We provide pseudo code for the greedy rounding algorithm in Algorithm 1.

The algorithm repeatedly solves the relaxed problem P_{cont} with different constraints fixing some decision variables to integer values. The algorithm proceeds in two phases. In the first phase, the maximal non-integral decision variable that can be rounded up without violating the budget constraint is determined. With this variable fixed, problem P_{cont} is solved again. When no additional component can be rounded up without violating the cost constraint, all previous constraints are removed, and the set of utilized locations is fixed instead. The algorithm sets a flag *locked* to `True` to show that the second phase of the algorithm has started.

In the second phase, components of $\boldsymbol{x}$ are still rounded up if possible. However, now we do not round up the largest non-integral component of $\boldsymbol{x}$. Instead, we determine for some location $l \in L$ with non-integral operation (i.e. $\exists \tilde{s} \in S_l \,:\, x_{l\tilde{s}} \notin \{0,1\}$) the first time interval

$$t \;:=\; \underset{t\in T}{\text{minargmax}} \left\{ \sum_{s\in S_{lt}} x_{ls} \,\middle|\, x_{ls} < 1 \forall s \in S_{lt} \right\} \tag{A6}$$

that is operated strongest at this location. Here, minargmax $\{\cdot\}$ refers to the minimal admissible value for argmax $\{\cdot\}$ if the maximum is not unique. Then, we round up the latest affordable shift $s \in S_l$ that covers the time interval t and add $x_{ls} = 1$ to the set of constraints. If no additional shift can be operated at location l, we add a constraint fixing the usage of this location: $x_{ls} = \lfloor x_{ls} \rfloor$ for all $s \in S_l$.

Distinguishing between the two phases of the algorithm yields optimized operating times. Suppose we are in phase 2, and consider the example depicted in Figure A2. The solution to the relaxed problem P_{cont} suggests that 3 inspection shifts s_1, s_2, and s_3 are conducted fractionally at the considered location l, with s_2 overlapping with s_1 and s_3. The respective operation intensities are $x_{ls_1} = x_{ls_3} = 0.8$ and $x_{ls_2} = 0.2$. The budget assigned to this location does not suffice to operate both s_1 and s_3 completely. Hence, only one shift can be operated at l. Naive greedy rounding would suggest to operate shift s_1, as it is the earliest shift with the maximal fractional operation. However, in the optimal solution, the time interval between 8 AM and 4 PM should be operated strongest. Therefore, shift s_2 would be the optimal choice.

In its second phase, the suggested algorithm rounds up shifts based on the maximal *cumulative* operation rather than choosing the shift with the highest operation variable. Nonetheless, it would

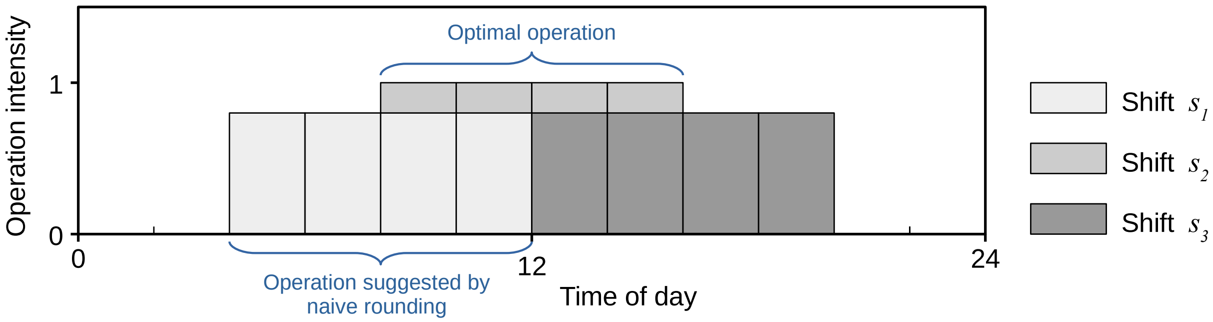

Figure A1: Motivation for the changed rounding procedure in phase 2 of the greedy rounding algorithm. The operation intensity is depicted as a function of time for some inspection location. The intervals on the time axis depict the discretization of the day time. The grey boxes show the extent to which the inspection station would be operated in the respective time intervals if fractional operation would be allowed. The boxes' colours correspond to the respective operation shifts. Naive greedy rounding would suggest to operate shift s_1. Improved rounding, however, would prefer the time interval in which the cumulative operation is maximal (shift s_2).

be of disadvantage to apply this rounding scheme in phase 1 of the algorithm, in which the set of used locations is not fixed. In this case, shifts in the middle of the day would always be chosen with preference, which make operation of *two* shifts on a day less efficient. In the second phase, it is typically known how many shifts should be operated at each location.

Slight improvements to the suggested algorithm are possible. For example, we added constraints in phase 1 to suppress fractional operation of shifts that would not be affordable completely under the costs of the already constrained variables. However, this improvement is unlikely to have a major effect on the results.

C Optimal inspection policy for the invasion scenario and the border closure scenario

To assess how the optimal inspection policy changes in response to altered external conditions, we considered (1) a scenario in which the invasion has progressed to states bordering BC and (2) a scenario in which the US-Canadian border is closed, such as experienced in the summer of 2020 in response to the global COVID-19 pandemic.

In the first scenario, we considered boaters from Idaho, Oregon, and Wyoming as high-risk

Algorithm 1: Greedy rounding algorithm.

Function `lock_location`($\tilde{\boldsymbol{x}}$, l, Θ)**:**
    foreach $s \in S_l$ **do**
        $\Theta := \Theta \cup \{x_{ls} = \tilde{x}_{ls}\}$;

$locked$:= `False`; $\Theta := \emptyset$;
while `True` **do**
    $\boldsymbol{x}$:= solution to P_{cont} subject to additional constraints in Θ;
    if $\boldsymbol{x} \in \mathbb{Z}^N$ **then**
        return $\boldsymbol{x}$;
    $\tilde{\boldsymbol{x}} := \lfloor \boldsymbol{x} \rfloor$;
    $\Omega := \{(l,s) \in L \times S \,|\, 0 < \boldsymbol{x}_{ls} < 1;\, C(\tilde{\boldsymbol{x}} + \boldsymbol{e}_{ls}) \leq B\}$;
    if $\Omega = \emptyset$ **then**
        if not $locked$ **then**
            $locked$:= `True`; $\Theta := \emptyset$;
            foreach $l \in L$ **with** $\max_{s \in S_l} x_{ls} = 1$ **do**
                $\Theta := \Theta \cup \{\max_{s \in S_l} x_{ls} = 1\}$;
        else
            l := some location with $0 < x_{ls} < 1$ for some $s \in S_l$;
            `lock_location`($\tilde{\boldsymbol{x}}$, l, Θ);
    else
        $(l,s) := \operatorname{minargmax}_{(l,s) \in \Omega} x_{ls}$;
        if $locked$ **then**
            $t := \operatorname{minargmax}_{t \in T} \left\{ \sum_{s \in S_{lt}} x_{ls} \middle| x_{ls} < 1 \forall s \in S_{lt} \right\}$;
            $\Psi := \left\{ s \in S_{lt} \middle| C(\tilde{\boldsymbol{x}} + \boldsymbol{e}_{ls}) \leq B \right\}$;
            if $\Psi = \emptyset$ **then**
                `lock_location`($\tilde{\boldsymbol{x}}$, l, Θ);
                continue;
            else
                $s := \max S_{lt}$;
        $\Theta := \Theta \cup \{x_{ls} = 1\}$;

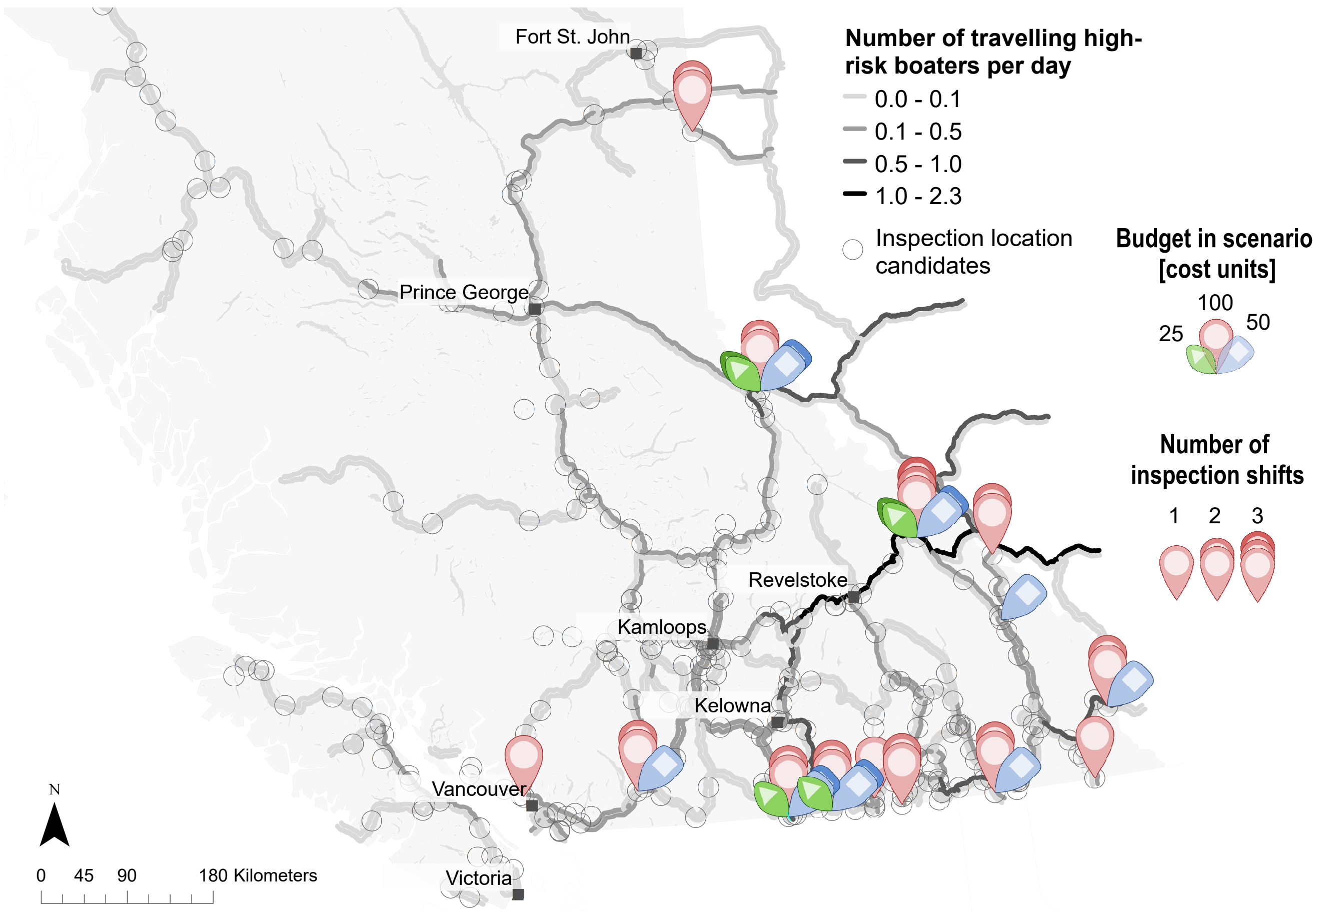

Figure A2: Optimal locations and operation shifts assuming that Idaho, Oregon, and Wyoming are mussel invaded. Compared to the base scenario with fewer infested States south of BC, more inspections are conducted at the southern border of the province. The symbols have the same meaning as in Figure 2 (main text).

boaters. Boaters from Montana were considered high-risk in the base scenario already. The results are depicted in Figure A2. As more high-risk boaters enter BC via the southern border, inspection efforts at this border are increased. The required resources are freed by operating fewer inspection stations over night and by abandoning inspection locations in the north. Nonetheless, the overall changes are moderate, because even in the changed invasion scenario, most high-risk boaters are expected to enter the province via the eastern border.

In the border closure scenario, boaters from the US were disregarded and routes to BC via the US were excluded. The latter led to a slight increase of boater traffic on some of the remaining routes. The optimized inspection locations and operating times are depicted in Figure A3. All resources are allocated to sites close to the eastern provincial border. In the Fort St. John area there are three eastern border crossings but location candidates were not identified near the border for all three routes. Hence, a location further inland, north of Prince George, had to be chosen to cover boaters travelling westbound from all three routes.

If a high budget is available, not all resources can be used in a meaningful way, and the distribution of the excess funds is dominated by model artefacts. Specifically, our model for traffic along unexpected routes does not incorporate the road topology. Hence, some traffic is predicted on inland roads even if all border crossings are covered with inspection stations. This biases the results to suggest inspections at inland locations, such as the location in northern BC (Figure A3). This issue can be resolved by disregarding traffic along unexpected routes when optimizing inspections in high-budget scenarios (Figure A4).

D Flexible and location-specific compliance rates

It may be more cost-effective to implement measures enforcing boaters' compliance than to operate many inspection stations for long hours. Furthermore, compliance of boaters may be higher or enforced more easily at some specific locations. In this appendix, we show how the approach presented in this paper can be adjusted to take these factors into account.

D.1 Location-specific compliance rates

We start by considering the case of non-uniform compliance rates. To that end, we split the boater flows based on the compliance of the boaters. Let C be the set of possible compliance rates, $c_l \in C$ the expected compliance rate of boaters at location $l \in L$, and L_c the set of locations with compliance rate $\tilde{c} \geq c$. For a route $r \in R$ and a time interval $t \in T$ Let n_{rtc} be the expected number of boaters who travel along route $r \in R$, started their journey in time interval $t \in T$, and comply at all inspection locations l with $l_c \geq c$ but not at inspection locations with $l_c < c$. These boaters will be inspected if and only if

$$\sum_{l \in L_r \cap L_c} \sum_{s \in S_{lrt}} x_{ls} \;\; \geq \;\; 1. \tag{A7}$$

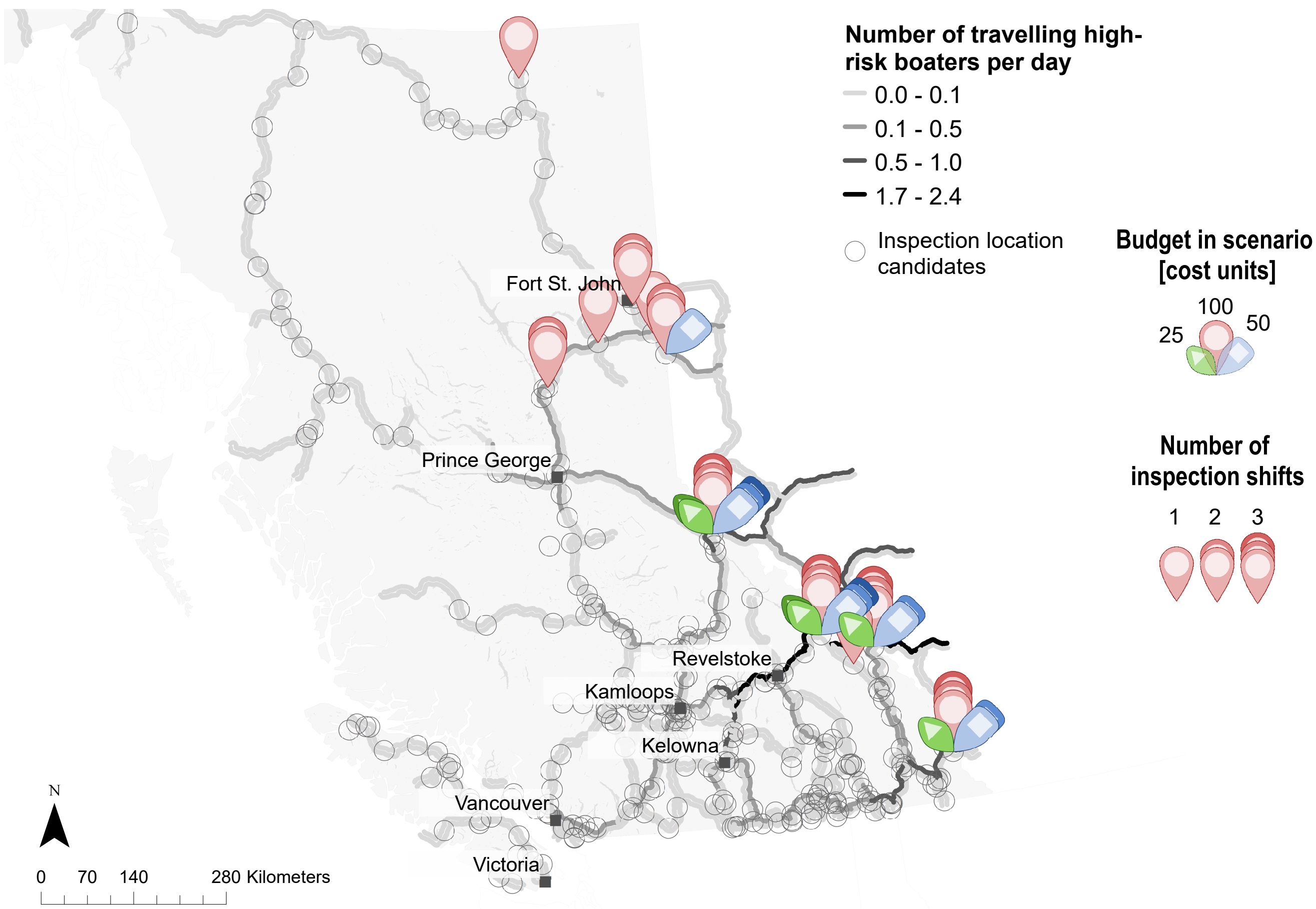

Figure A3: Optimal locations and operation shifts if the US-Canadian border is closed. All inspection efforts are distributed along the eastern provincial border. If the budget is high (100 units), all major traffic flows are blocked, and the allocation of the remaining resources is dominated by model and optimization artefacts, e.g. resulting in the choice of an inland location in northern BC. Note that the locations surrounding Fort St. John cover northbound traffic only; the suggested location north of Prince George is the first location candidate to cover westbound traffic from the border in this area.

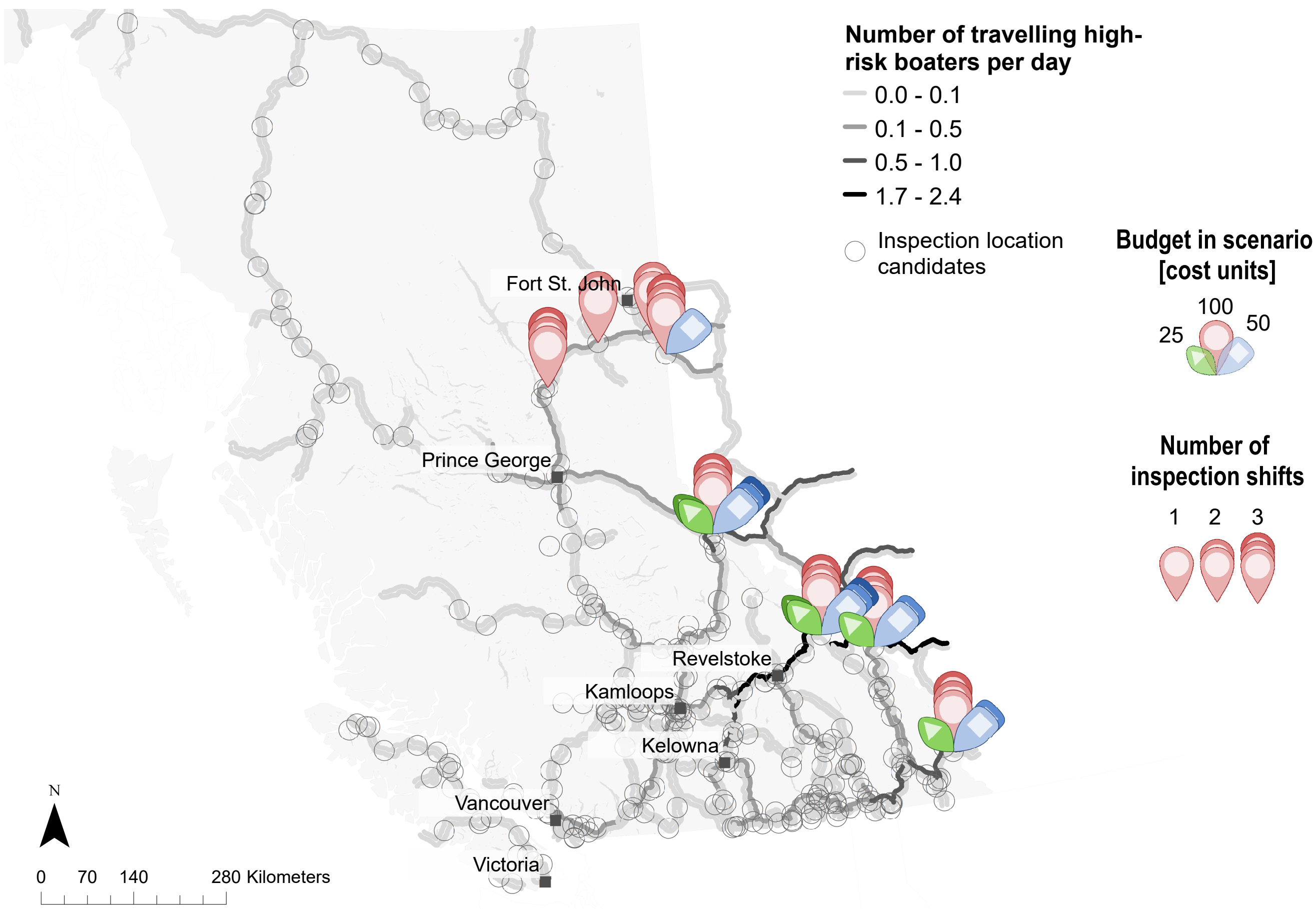

Figure A4: Optimal locations and operation shifts if the US-Canadian border is closed and traffic along unexpected routes is ignored. For the low and medium budget scenario, the chosen locations do not differ from those obtained when traffic along unexpected routes is considered (Figure A4). For the high-budget scenario, however, all resources are distributed to locations in border proximity, and no inland locations are used. Note that the locations surrounding Fort St. John cover northbound traffic only; the suggested location north of Prince George is the first location candidate to cover westbound traffic from the border in this area.

As in the main text, x_{ls} is a binary variable denoting whether inspections are conducted at location $l \in L$ in shift $s \in S$. Consequently, the total number of inspected boaters is given by

$$F_{\text{loc-compliance}}(\boldsymbol{x}) \quad := \quad \sum_{c \in C} \sum_{r \in R} \sum_{t \in T} \min \left\{ 1, \sum_{l \in L_r \cap L_c} \sum_{s \in S_{lrt}} x_{ls} \right\} n_{rtc}. \tag{A8}$$

This function can be optimized with the same method discussed in the main text. With a similar approach, time-dependent compliance rates could be incorporated, too.

D.2 Flexible compliance rates

In some applications, the compliance rate may be altered at a specific cost. If these costs can be expressed as a convex function of the achieved compliance rate, a flexible compliance rate can be incorporated in our model easily. Below, we consider for simplicity the base case with a uniform compliance rate at all locations. Allowing location-specific flexible compliance rates can be done by combining the two approaches introduced in this appendix.

Let n_{rt} be the expected number of boaters travelling on route $r \in R$ and who started their journey in time interval $t \in T$. Note that other than in the main text, compliance of these boaters is not supposed. Altering equation (5) from the main text to

$$F_{\text{flex-compliance}}(\boldsymbol{x}) \quad := \quad c \sum_{r \in R} \sum_{t \in T} \min \left\{ 1, \sum_{l \in L_r} \sum_{s \in S_{lrt}} x_{ls} \right\} n_{rt} \tag{A9}$$

accounts for the flexible compliance rate c.

Let us assume assume that the costs for enforcing a specific compliance rate c at a location $l \in L$ and during shift $s \in S$ are given by the linear function

$$\text{cost}_{ls}(c) \quad = \quad \alpha_l \left(c - c_0 \right), \tag{A10}$$

where c_0 is the base compliance rate if no actions are taken to increase compliance. More complex cost functions can be modelled with convex piece-wise linear functions or general convex functions.

Adding these costs to the overall cost function changes the cost constraint to

$$\sum_{l\in L}\left(\sum_{s\in S_l}\left(c_{ls}^{\text{shift}}+\alpha_l\left(c-c_0\right)\right)x_{ls}+c_l^{\text{loc}}\max_{r\in R,\,t\in T}\left(\sum_{s\in S_{lrt}}x_{ls}\right)\right) \;\leq\; B. \tag{A11}$$

In addition to changing the objective function and the cost constraint, we have to introduce one further constraint limiting the compliance rate to the feasible range:

$$c_0 \;\leq\; c \;\leq\; 1.$$

With these changes, the compliance rate can be optimized along with the inspection locations and operating times.

E Difficult inspection optimization scenarios

In many real-world instances, good solutions to the linear integer problems derived in this paper can be identified within reasonable time. Nonetheless there are examples in which the optimization is computationally challenging. In this appendix, we discuss two important mechanism that can make it difficult to find a highly optimal solution in short time. We also provide examples for the discussed mechanisms.

Difficulties can arise (1) if a significant fraction of the budget is unused under the optimal policy and (2) if many boaters pass multiple operated inspection locations under the optimal policy. We start by considering budget-related issues before we discuss problems arising from unfavourable relationships between potential inspection locations. At the end of this appendix we discuss why these challenges are not of major concern in many real-world applications. To simplify explanations, we consider the case of optimizing inspection station placement only. The described mechanisms extend easily to the full problem in which operating times must be optimized as well.

E.1 Difficulties due to cost constraints

Let us first consider a scenario in which a fraction of the given budget remains unused under the optimal policy. For example, suppose that operation of an inspection station costs 5 cost units and that we are given a budget of 9 units. Consequently, 4 cost units of the budget will remain unused. To obtain an approximate solution and obtain an upper bound on the optimal objective value, solvers consider the problem's continuous relaxation, in which partial use of inspection locations (and shifts) is permissible. In this relaxed scenario, all 9 cost units will be spent, which allows the inspection of more boaters than in the realistic scenario with binary choices. Consequently, the upper bound on the solution given by the solution to the relaxed problem may be much higher than the true optimal solution. This makes it difficult to check whether an identified solution is highly optimal and thus increases computation time.

The problem described above becomes even more difficult if control actions with different costs are possible. Suppose that we may operate one of three inspection stations, which are passed by different sets of boaters, respectively. That is, no boater passes two of the potential inspection locations. Assume that per day $n_1 = n_2 = 5$ boaters pass stations l_1 and l_2, respectively, and that $n_3 = 8$ boaters may be inspected at location 3. Suppose we are given a budget of 9 units and that the costs for operating stations l_1 and l_2 are $c_1 = c_2 = 5$ cost units, whereas operation of station l_3 requires $c_3 = 9$ cost units.

Again, optimizers may consider the problem's continuous relaxation to find an approximate solution and a quality estimate. An optimal solution to the relaxed problem is to operate both station 1 and station 2 fractionally with weight $x_1 = x_2 = 0.9$. Then, the total costs $x_1 c_1 + x_2 c_2 = 9$ satisfy the budget constraint and the total number of inspected boaters is given by $x_1 n_1 + x_2 n_2 = 9$. However, in the original integer problem, stations cannot be operated fractionally, and only one station can be chosen. As more boaters pass location 3 than locations l_1 or l_2, it would be optimal to conduct inspections at location l_3, where 8 boaters can be inspected. Applying a greedy rounding algorithm to the relaxed solution, however, would suggest to operate either location l_1 or l_2, where only 5 boaters would be expected.

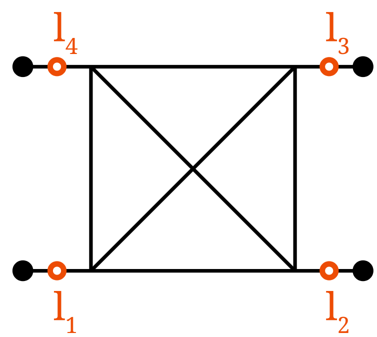

Figure A5: Inspection location setup that leads to a challenging optimization problem. The lines denote roads, the solid black circles origins and destination, and the hollow orange circles potential inspection locations.

E.2 Difficulties due to unfavourable relations between inspection locations

Besides challenges induced by cost constraints, specific relationships between potential inspection locations can make the optimization difficult. Consider the example depicted in Figure A5 and assume that an arbitrary number of boaters may drive from each origin/destination (black circle) to each other origin/destination. Suppose that operating an inspection location at any of the permissible locations has unit cost and that we are provided a budget of 2 cost units. If the relaxed version of the problem is considered and fractional operation of stations is permitted, operating each location with intensity $\frac{1}{2}$ would cover all boater flows and hence be the optimal solution. However, if discrete choices must be made, some boaters will not be inspected. As all locations are operated equally in the optimal solution to the relaxed problem, this solution does not provide any hint towards which of the locations should be operated in the original scenario with binary decisions. Therefore, the problem is difficult to solve.

E.3 Prevalence of difficult scenarios in real-world applications

Any of the challenging scenarios discussed above can occur in real-world problems. However, certain characteristics of real-world scenarios lower the risk of running into optimization issues. In many management scenarios of interest, various inspection stations can be operated. Problems induced by the budget constraint become less significant if a large budget is considered so that a potential remainder of the budget becomes insignificant. For example, in all scenarios with a budget above 30 units considered in this paper, we reached a solution with at least 98% optimality within minutes. Furthermore, issues induced by budget constraints can be mitigated by investigating alternative scenarios with slightly adjusted budgets.

Scenarios with unfavourable relationships between potential inspection locations can be expected in real-world applications. Note that the issue with the setup in Figure A5 persists if the roads connecting the potential inspection locations have a shape different from the road pattern drawn in the figure. Furthermore, the depicted situation may refer to a portion of the road network only, with multiple origins and destinations connected to each of the depicted origin/destination vertices. In fact, situations such as the considered one could appear multiple times in a road network. Therefore, the considered challenges do not only occur in scenarios in which inspections are restricted to locations close to origins and destinations.

Nonetheless, invasion patterns frequent in real-world scenarios reduce the prevalence of such unfavourable inspection station relationships. As short distance dispersal of invasive species is typically more likely than long-distance dispersal, invaded habitat patches form clusters so that the inflow of potentially infested vectors, such as boaters, comes from specific directions only. For example, high-risk boaters enter BC through the southern and eastern border only. Therefore, it is often possible to identify inspection location configurations in which only few high-risk boaters pass multiple operated inspection stations. This simplifies optimization of the inspection policy. Greater optimization challenges can be expected if origins and destinations are intermixed.