

General rules for managing and surveying networks of pests, diseases, and endangered species

Iadine Chadès^{a,1}, Tara G. Martin^a, Samuel Nicol^b, Mark A. Burgman^c, Hugh P. Possingham^b, and Yvonne M. Buckley^{a,b}

^aCSIRO Ecosystem Sciences, Ecosciences Precinct, Dutton Park, QLD 4102, Australia; ^bUniversity of Queensland, School of Biological Sciences, Queensland, 4072, Australia; and ^cSchool of Botany, University of Melbourne, Parkville, VIC 3052, Australia

Edited by Simon A. Levin, Princeton University, Princeton, NJ, and approved March 31, 2011 (received for review November 11, 2010)

The efficient management of diseases, pests, or endangered species is an important global issue faced by agencies constrained by limited resources. The management challenge is even greater when organisms are difficult to detect. We show how to prioritize management and survey effort across time and space for networks of susceptible–infected–susceptible subpopulations. We present simple and robust rules of thumb for protecting desirable, or eradicating undesirable, subpopulations connected in typical network patterns (motifs). We further demonstrate that these rules can be generalized to larger networks when motifs are combined in more complex formations. Results show that the best location to manage or survey a pest or a disease on a network is also the best location to protect or survey an endangered species. The optimal starting point in a network is the fastest motif to manage, where line, star, island, and cluster motifs range from fast to slow. Managing the most connected node at the right time and maintaining the same management direction provide advantages over previously recommended outside–in strategies. When a species or disease is not detected and our belief in persistence decreases, our results recommend shifting resources toward management or surveillance of the most connected nodes. Our analytic approximation provides guidance on how long we should manage or survey networks for hard-to-detect organisms. Our rules take into account management success, dispersal, economic cost, and imperfect detection and offer managers a practical basis for managing networks relevant to many significant environmental, biosecurity, and human health issues.

conservation planning | decision theory | metapopulation | optimization | Markov decision process

Infectious diseases, invasive pests, and other threats to species persistence have profound impacts on human health, agriculture, and biodiversity (1–3). Many threatened or invasive species are difficult to detect and their presence in an area can be uncertain due to the imperfect nature of most detection methods (4, 5). Even large charismatic mammals such as the Sumatran tiger (*Panthera tigris sumatrae*) or the Sumatran rhinoceros (*Dicerorhinus sumatrensis*) can be surprisingly hard to detect. It is possible that some areas are being managed while the invasive pests, diseases, or threatened species have already disappeared. It is also likely that managers, in the absence of sighting, might stop managing too early and give up too soon on a species or disease (6, 7). Managers need to know when to stop managing or surveying for species in areas of particular interest.

At present the epidemiology, ecology, and conservation literature provides little guidance on how to approach such a problem. The problem of how to allocate management and surveillance effort has been studied for a single cryptic population (8–10). For example, Regan et al. (8) determined when to stop monitoring an invasive plant with a long-lived seed bank. Chadès et al. (9) determined when to stop managing or surveying cryptic threatened species. However, none of these works have accounted for multiple populations, spatial dispersal, or risk of reinfestation over time. When managing a cryptic invasive plant over time,

Regan et al. (11) showed that accounting for imperfect detection but ignoring the risk of reinfestation led to a low probability of eradication, defeating the purpose of an optimal management strategy. Other studies on managing metapopulations have accounted for spatial connectivity but have not included imperfect detection in a systematic manner (12–17).

Metapopulations can be modeled as susceptible–infected–susceptible (SIS) networks (18). In epidemiology and ecology, SIS infections are challenging to manage because a treated node (person or location) can be cured and reinfected (we use the term “infection” to refer to infectious disease incidence, infestation by invasive pests, and colonization by endangered species). SIS network models can represent sexually transmitted diseases (gonorrhea) (19), other infectious diseases (meningitis, plague, malaria, and sleeping sickness) (20), metapopulations of invasive or threatened (21) species, influence in a social network (22, 23), and computer viruses across physical networks (24, 25). Here, we build on lessons from previous empirical studies and theoretical frameworks for optimal decision making (26–28) to develop coherent guidance for allocating resources across time and space between three activities: managing, surveying, and doing nothing for networks of cryptic diseases, pests, and threatened species. In our case managing a disease or a pest means decreasing its presence, whereas managing a threatened species means increasing its presence. We represent the spatial structure of pest invasions, infectious diseases, and threatened metapopulations (29) as “infected” and “susceptible” nodes in networks (18, 30, 31). The pest, disease, or threatened species can spread from an infected node to a susceptible node with which it shares a direct connection (32). We tackle our problem by first assuming a perfect detection process. We frame our optimization problem as a finite Markov decision process (33) and apply the solution to network motifs (34) (lines, islands, and stars) and combinations of motifs. Network motifs are small subgraphs that capture specific patterns of interactions. Studying optimal strategies for network motifs provides important insight into the management of more complex networks. We include imperfect detection into our model and analyze how the management strategy changes as our belief in the presence of the species or disease decreases. Our simple, but robust, rules of thumb for efficient management take into account management success, dispersal, economic cost, and imperfect detection. Contrary to previous high-profile work on large-scale networks (31, 35–37), we find that an “outside–in” control strategy for preventing the spread of infection is in-

Author contributions: I.C., T.G.M., M.A.B., H.P.P., and Y.M.B. designed research; I.C. and S.N. performed research; I.C. and S.N. contributed new reagents/analytic tools; I.C., T.G.M., S.N., and Y.M.B. analyzed data; and I.C., T.G.M., S.N., M.A.B., H.P.P., and Y.M.B. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

Freely available online through the PNAS open access option.

¹To whom correspondence should be addressed. E-mail: iadine.chades@csiro.au.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1016846108/-DCSupplemental.

efficient for network motifs. This is a key finding because eradication or effective control is generally possible only where a pest or a disease occupies a small number of nodes (38). Likewise, managing rare or threatened metapopulations is by definition a small network problem.

The first step in solving a resource allocation problem is defining the objective we seek to achieve. When managing pests and diseases, our objective is to minimize the number of infected nodes by taking actions to increase the probability of extinction of an infected node. When managing threatened species, our objective is to maximize the number of infected (occupied) nodes by taking actions that decrease the probability of extinction of an infected node. These equivalent problems have the same solution. The best local population to manage for an infestation is also the best place to manage for a threatened metapopulation.

At each time step, we assume the infection state of each node i ($s_i = \{\text{susceptible, infected}\}$) is observed. Once managed, a node can remain infected with probability $(1 - p_m)$ and infect susceptible nodes to which it is connected with a probability p_d . When management occurs, the node may become susceptible and can be reinfected with probability $1 - (1 - p_d)^\delta$, where δ is the number of infected nodes directly connected to the susceptible node. We assume that each infected node has a probability of becoming susceptible without management (p_e), to account for the natural, stochastic decline of the infection. We assign an economic cost, c , every year that results from a node being infected. Managing a node incurs a cost c_m . We consider that only one node can be managed at a time. This method is the simplest way to prioritize the order in which nodes are managed so that our objective is achieved in the most efficient way. We also demonstrate that our management rules are robust when several nodes can be managed at a time (*Results* and *Methods*). We assume uniform dispersal, management success, benefit, and cost across all nodes (*Methods* and *Table S1*). However, our method can accommodate changes and nonuniform assumptions in any of the following parameters: the probability of a node being infected if one neighbor is infected (p_d), the probability of a node becoming susceptible if infected and not managed (p_e), the probability of a node becoming susceptible if it has been managed (p_m), the cost of a node being infected (c), and the cost of managing one node (c_m).

Results

When two infected nodes have the same probability of infecting each other, the optimal strategy is to manage either one. When several nodes are linked by a directed probability of dispersal, it is optimal to manage the upstream node before managing the downstream nodes (Fig. 1A). In the case of a line network with bidirectional infection risk, the nodes at the extremity of the infection process have a lower probability of being reinfected than nodes with two neighbors. The optimal strategy is to start from either extremity of the network and then manage the nearest infected nodes sequentially (Fig. 1B). In the case of an island network, all nodes have an equal probability of being reinfected once managed; therefore, any node can be the starting point (Fig. 1C). Star networks are composed of a central node connected to satellite nodes. Infected satellite nodes are managed until the number of susceptible satellites is equal to or larger than the remaining number of infected satellites; it is then optimal to manage the central node before managing the remaining infected satellites (Fig. 1D). When managing star networks with lines of nodes of different lengths connected to a central node, the optimal strategy has three phases: (i) manage the longest lines from an outside-in direction until just over half of the nodes connected to the central node remain infected, (ii) manage the shortest infected line and the central node following an outside-in direction, and (iii) manage the remaining

Network motifs	Optimal rules of thumb
A Directed lines 	Manage upstream nodes then manage downstream nodes.
B Lines 	Manage one of the extreme nodes, then manage the nearest infected nodes.
C Islands 	Manage any node, then manage one of the nearest infected nodes.
D Stars 	Manage half of the satellite nodes (rounded up for odd numbers), then manage the central node c , finally manage the remaining infected satellite nodes.

Fig. 1. (A–D) Optimal rules of thumb to manage network motifs of a disease or an invasive or threatened species.

infected lines starting with the shortest lines following an inside-out direction.

Fig. 2 shows simple rules of thumb for more complex networks made up of multiple motifs. When an island is connected to a line, the risk of reinfection is minimized by starting management at the extreme node of the line and continuing to the connecting node, before managing the island. Once the connecting node is

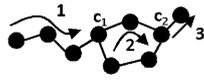
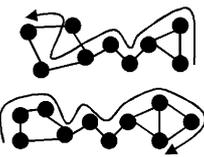
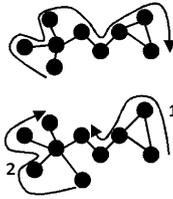
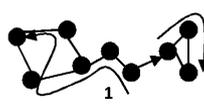
Complex networks	Optimal rules of thumb
A Islands with lines 	Manage the extreme node from the longest line, manage the nearest infected nodes until the connection node c_1 is managed (1). Manage the nodes of the island that are nearest to the connection node c_1 (2). Manage the connection node c_2 and the remaining infected nodes (3).
B Connected clusters 	Manage the nodes in the cluster with the least number of nodes first. When the clusters have equal numbers of nodes manage the nodes from the least connected cluster first.
C Connected island and star 	Manage the star first, when the number of satellite nodes of the star is less than or equal to the number of nodes of the island (in the first case, both equal 3). Otherwise, manage the island first.
D Directed links 	Directed links change the rules of thumb, manage upstream of a directed link before managing downstream.

Fig. 2. (A–D) Optimal rules of thumb to manage small complex networks of a disease or an invasive or threatened species.

susceptible, it is optimal to manage the infected nodes closest to the connecting node. A more complex configuration is an island connected to two lines (Fig. 2A). Maintaining the same management direction reduces the risk of reinfection and is therefore more efficient than managing the “outside” nodes first.

We define a cluster as an island with interconnected nodes. When managing connected clusters (Fig. 2B), it is optimal to start from the smallest cluster or, for the same number of nodes, the least connected cluster. For a connected island and star (Fig. 2C), the star is faster to manage when the number of satellites is equal to or smaller than the number of nodes of the island. In contrast, when the number of star satellites is greater than the number of island nodes, managing the island first reduces the chance of reinfection. In both cases, the rules of thumb derived from individual motifs can be applied. For a star connecting four substars (Fig. 3), the optimal rule of thumb for managing single star motifs holds.

We compared our rules of thumb to an outside-in control strategy by simulating the management of 1, 2, and 5 nodes at a time on large star networks of 101 nodes (Figs. S1–S3). Our optimal strategy outperforms the outside-in strategy, reducing the time to eradication and the management required by 31.2% (90.7 treatments, managing 1 node), 10.7% (20.4 treatments, managing 2 nodes), and 3.3% (5.6 treatments, managing 5 nodes). Our rules are also robust to nonuniform dispersal probabilities (p_d) and management success (p_m) and consistently outperform the traditional outside-in strategy (Table S2 and SI Methods). We found that the relative benefits of our rules of thumb are greater when fewer patches are managed in each time step. Our rules of thumb outperform the outside-in strategy because reinfections are less likely to occur using our priority management order. The longer the time it takes to eradicate the infection, the more opportunity there is for reinfections to take place, and the greater the benefit of our rules of thumb over the outside-in strategy. By managing multiple patches in a time step, the time to eradication is reduced, making the performance of the two strategies more similar.

For all of the complex networks we examined, management starts with the fastest motif to manage, following the priority order line, star, island, and cluster. Management of the most connected node at the right time and maintaining the same management direction minimize the reinfection risk, providing potentially important practical advantages over outside-in strategies, particularly in the case of limited management resources. To determine the management starting point more easily we asked if a network centrality measure could substitute for our rules of thumb. Centrality measures are able to rank nodes in a network according to the topological features that they account for (see refs. 39–41 for more details). We calculated centrality

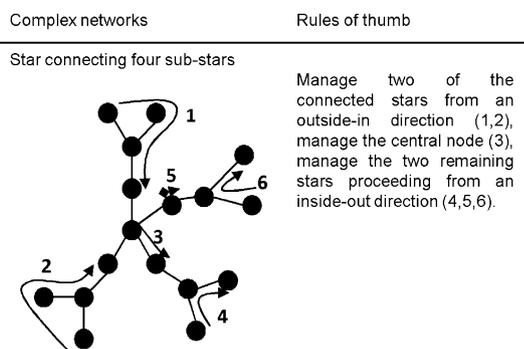


Fig. 3. Optimal management of a complex star network of a disease or an invasive or threatened species.

measures for nine small networks (Table S3). We found that the closeness centrality measure correctly predicted the most appropriate node to start from for seven cases of nine but predicted incorrect starting nodes for the two networks characterized by connected clusters. Betweenness centrality, eigenvector centrality, and subgraph centrality (42) performed inadequately. Although closeness centrality often predicted the correct starting node, it did not predict an optimal priority management order, highlighting the continuing need for our optimization framework.

Most pests, diseases, and threatened species cannot be detected perfectly (9). Apparently susceptible nodes may, in fact, be infected. We extend our framework to take into account our inability to detect subpopulations perfectly, using the partially observable Markov decision process (POMDP) (Methods). In this case, we consider three possible management actions: Manage a node, survey for the infection, or do nothing. We include in our optimization model the probability of imperfectly detecting infected nodes, d_n , when doing nothing or managing. Probability of detection increases when we survey a node, $d_s > d_n$. The cost of survey, c_s , is set to half the cost of management. Using a star network of four nodes, we illustrate the importance of accounting for imperfect detection where we assume the infection is not observed after each action (see the parameters used in Table S1).

When all of the nodes in a star are initially infected (Fig. 4), the order of management follows the rules of thumb we described under perfect detection. After this first round of management, if the infection is not detected in any node, a stochastic dynamic programming solution recommends doing nothing; however, when accounting for imperfect detection, our results recommend managing all of the nodes a second time for this star network. The number of times a node must be managed depends on the POMDP parameters used. As our belief in eradication increases, our solution shifts to surveying the most connected node. Finally, when our belief that the infection has been eradicated from the network has reached a threshold (0.82 in the case of Fig. 4), and in the absence of observing an infection, “do nothing” becomes the optimal action. When we observe an infection to be present in a node(s), the belief of eradication falls to zero and “manage” is recommended.

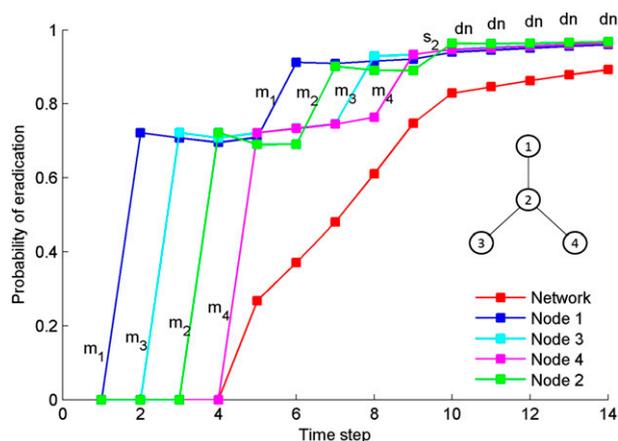


Fig. 4. Resource allocation recommended over time for the management of a hypothetical invasive species when infection may be undetected on a star network (Inset). m_i , manage node i ; s_i , survey node i ; and dn , do nothing. The red line represents the probability (belief) of eradication of the infection from the network. The other lines represent the belief in eradication of the individual nodes. The probability of eradication starts at zero for each node because initially, the network is assumed to be infected.

When we do not have any prior information about the infection of the network, every infection state in every node is equally probable (Fig. S4). In this case, our solution recommends first managing all of the nodes at least once and more than once if the node is highly connected (e.g., node 2, Fig. S4). The action of surveillance then becomes optimal. In the absence of detection of the infection under surveillance, our belief that the infection is eradicated reaches a high value (e.g., probability of eradication in the network >0.82). It is then recommended to do nothing.

Under imperfect detection, as our belief in eradication increases, our control strategy evolves through three phases: Manage, survey, and do nothing. As demonstrated in Chadès et al. (9), the problem of when to stop managing and surveying a single population of cryptic species is dependent on the relative cost of survey and management, the cost of the infection, the benefits of eradication, and the efficiency of management and surveillance determine the switching points (9). To avoid re-running the POMDP algorithms for every network studied, we build on this work and derive an analytic approximation by transforming our network POMDP problem into a single POMDP. To make the transformation, we assume the dispersal probability between nodes is a less important factor than the number of subpopulations (*Methods*). Although this assumption may not always hold, this approximation is the best indicator of how long we can afford to manage derived thus far and forms the basis for future theoretical work. We find that, in the absence of sighting, as the number of subpopulations increases, the time we should spend managing increases substantially irrespective of the economic costs (Fig. S5).

Discussion

We identify simple rules for the management of small networks under perfect and imperfect detection of susceptible–infected–susceptible network invasions, infections, and threatened metapopulations. Our findings provide guidance for managers in determining (i) where to prioritize management over time and space to achieve the most efficient control strategy and (ii) when the species or the disease is difficult to detect, how long they should invest in active management or in surveying and when it is reasonable to stop investing limited resources. Our simulation results demonstrate the improvement in performance when following our rules of thumb over traditional outside-in management approaches on large motif networks. Our rules are also robust to nonuniform distribution of dispersal and management efficiency. When deciding where we should prioritize management, a key finding is to start with the fastest motifs to manage and keep the same management direction through the network. Centrality measures are not good proxies for determining the priority order of management of nodes. Finally we derived an analytic approximation providing guidance on the time we should spend managing and surveying across a network in the absence of detection.

Details of the number of nodes, the costs and benefits of managing nodes, and the dispersal probability drive the optimal management strategies. In ecology and economics, evaluating the loss of a species or the impact of a pest or disease is a challenging and often highly subjective undertaking (43). We find that for a pest or disease with high management costs combined with a high dispersal probability a do nothing strategy can be optimal when control becomes unachievable under budget constraints. In contrast, high economic value of a threatened species and low detection probability result in management and surveying for longer periods of time being optimal (ref. 9 and Fig. S5).

To derive our rules of thumb we have assumed uniform dispersal across all nodes; however, infection/colonization processes can be driven by directed factors such as transport vectors (wind, water, or road dispersal). In this case, directed dispersal changes

the management rules (Fig. 2D). Although our optimization tool can accommodate differences in parameters across a network, such information is often not available. In the absence of explicit information, expert elicitation can be used and these qualitative values can be incorporated into our framework (44).

We were able to derive management strategies for many different kinds of small networks (<25 nodes). Deriving a priority order of which nodes to manage is useful only for networks that exhibit a clear structure. Highly connected networks will not benefit from our approach as the amount of resources invested to manage the threat will be more important than the order of nodes managed. In other words networks need to be sparsely connected to benefit from our approach. If a structured network is found but the number of nodes is too large to be solved following our approach, we recommend using an abstraction process or decomposition methods to extract small subnetworks (45–47). An alternative approach is to use a model that does not include network structure at all, such as a diffusion process (48).

Managers of diseases might find our rules counterintuitive as we do not recommend managing the most connected nodes first. There are several plausible explanations for this. First, our rules provide a priority order of management from the initial condition of all nodes being infected. This order might not correspond to the starting management point of some diseases, e.g., when only a few nodes in a wider known network are infected. Although our optimal strategy defines an optimal action for any infection stage, it is impossible to derive a unique priority order that applies to every infection stage. Second, although common, the assumption that a node can represent the state of a subpopulation might not reflect epidemic dynamics; however, we can consider a node as a representation of an individual and our rules continue to apply to a network of infected individuals. Third, we assume the structure of the network is known. Therefore, the infection cannot spread to unknown connected nodes and managing the most connected node does not bring any additional advantage. Finally, the susceptible–infected–susceptible network model might not capture the complexity of some diseases where vaccinated or recovered and immune states are possible; other models such as susceptible–infected–recovered (SIR) might be more appropriate.

We present general solutions for the management of connected subpopulations of pests, diseases, and threatened species arranged in a SIS network. The efficient management of diseases, pests, or threatened species is a globally important issue, especially in agencies constrained by limited resources. The challenge of managing new incursions is even greater when the species or disease is cryptic and impossible to detect perfectly. To date, most of the literature in network theory has focused on making predictions about large complex systems such as food webs, epidemics, and gene regulation (49, 50). We provide general rules of thumb for setting management priorities in small networks, where our chances of eradication or recovery of a species are maximized (51). Our results provide a practical basis to manage SIS networks relevant to many significant biosecurity, human health, and environmental issues.

Methods

When there is perfect detection, finding the optimal management strategy for an infected network can be solved using Markov decision process (MDP) and stochastic dynamic programming (SDP) (33). In our case MDP is inconvenient for two reasons: The computational complexity grows exponentially with the number of nodes and SDP methods mostly provide a single optimal solution. Determining alternative optimal solutions is essential when deriving rules of thumb. We overcome these disadvantages by modeling our SIS network management problem as a factored Markov decision process (FMDP) (26), taking advantage of the structure of our problem and fast optimization methods from artificial intelligence.

A factored MDP is a MDP that explicitly represents the independent relationships between variables of the system (conditional independence).

A factored MDP is defined by four elements: a factored set of states (S), a finite set of actions (A), a transition probability function (P), and a reward function (R). Let s in S be the state of the network at any given time t and s represent the nodes infected or the susceptible state: $s = (s_1, s_2, \dots, s_n)$, where the state variable $s_i \in S_i = \{\text{infected, susceptible}\}$ represents the state of the node i . The system has 2^n possible states. Let a in A be the decision the manager can make at any given time t and a represent the nodes managed by the decision makers: $a = (a_1, a_2, \dots, a_n)$, where the components $a_i \in A_i = \{\text{manage, do nothing}\}$, the set of decisions available in node i . Because we are interested in the prioritization of management, we assume only one node can be managed at each time step. Our simulation results have shown that when we increase the number of nodes we can manage at a time, our rules continuously outperform the outside-in strategy (*Results and SI Methods*). As currently defined, our model and approach will still hold when more resources are available, allowing the management of several nodes at a time. However, the computational complexity is dependent on the number of actions we consider. If we were to consider managing k nodes among n at a time, the size of the action set grows with k , $|A| = n!/(n-k)!$, limiting the use of our methods to very small networks.

The dynamics of the system are captured by a matrix P that contains the transition probabilities of moving from any state to any other state of the system under different actions following the SIS process (Introduction and Fig. S6). In classical MDP solution methods, all transition probabilities must be represented and stored using $|A|$ matrices of $|S|$ by $|S|$ dimensions. Each element of these matrices represents a transition probability: $p(s_{t+1}|s_t, a) = \prod_{i=1, \dots, n} p(s_{i,t+1}|s_{i,t}, \dots, s_{n,t}, a)$.

A factored MDP takes into account the independence between state variables explicitly, using a dynamic Bayesian network implemented as algebraic decision diagrams (ADD) (Fig. S6). For example, let us consider a line network with three nodes. The classic MDP needs to store the probabilities of transition for every state given all possible states of nodes 1, 2, and 3, even though node 1 does not affect the state of node 3 and vice versa. In this way the classic algorithms are inefficient because they store information that is not needed to find the optimal solution. A better way is to use the structure of the network to store information for the nodes that directly affect the state of nodes 3 and 1; e.g., $p(s_{3,t+1}|s_{3,t}, s_{2,t}, s_{1,t}, a) = p(s_{3,t+1}|s_{3,t}, s_{2,t}, a)$ and $p(s_{1,t+1}|s_{3,t}, s_{2,t}, s_{1,t}, a) = p(s_{1,t+1}|s_{2,t}, s_{1,t}, a)$ —this simplification is factored MDP. If an optimization problem has many independent variables, we can solve larger-size MDPs using factored MDP because we have fewer state interactions to consider and store. In our case, as the network studied becomes larger, the savings gained by representing the network dynamic increase exponentially.

When solving the corresponding optimization problem, we seek to determine an optimal strategy $\pi: S \rightarrow A$ that minimizes the expected number of nodes infected over an infinite time horizon. The optimal strategy matches an optimal action to each possible state of the system. For any network thus specified, we first determine the conditional dependence/independence between nodes over one time step. We automatically generate the corresponding FMDP model and solve the optimization process using the fast algorithm *stochastic planning using decision diagrams* (SPUDD) (27). We then analyze the solutions and derive clear management rules.

In reality many pests, diseases, and threatened species cannot be detected perfectly. A factored partially observable Markov decision process (FPOMDP) model is an augmented FMDP with a factored set of observations (Z) and an observation function (O). Similar to a factored MDP, a factored POMDP is a POMDP that explicitly represents the independent relationships between variables of the system. In our case the FPOMDP tackles the problem of deciding the most appropriate management action given the difficulty of detecting an infected node. We define the finite set of local observations for each node at time t , $z_i = \{\text{Absent, Present}\}$ and their corresponding observation function o_i that maps to each state-action pair a probability distribution over z_i . In other words, the probability of detection of the infection given that the node is infected and that the previous decision is to do

nothing is defined by $o_i(\text{Present} | \text{Infected}, \text{Do nothing})$. Let z in Z be an observation of the whole system $z = \{z_1, \dots, z_n\}$. We define the observation function of the system as the joint probability of the local observation function: $O(z | s, a) = o_1(z_1 | s_1, a_1) \times \dots \times o_n(z_n | s_n, a_n)$.

As it is neither practical nor tractable to use the history of the action-observation trajectory to compute or represent an optimal solution, belief states are preferred to summarize and overcome the difficulties of incomplete detection. Indeed, Astrom (52) has shown that belief states are sufficient statistical tools to summarize all of the observable history of a POMDP without loss of optimality. A POMDP can be cast into a framework of a fully observable Markov decision process where belief states represent the continuous but fully observable state space. Here, a belief state b is defined as a distribution probability over states in S . In our case, solving a FPOMDP is finding a strategy $\pi: B \times T \rightarrow A$, mapping an allocation of resource $a \in A$ given a current belief state $b \in B$ and a time step $t \in T$. An optimal strategy minimizes the expected sum of costs or rewards (R) over a finite time horizon, T . This expected summation is also referred to as the value function (53). A value function essentially ranks strategies by assigning a real value to each b .

Whereas various algorithms from the operations research and artificial intelligence literature have been developed over the past years, the computational complexity of exact algorithms remains intractable for most problems: Finite horizon POMDPs are PSPACE-complete (54) and infinite-horizon POMDPs are undecidable (55). In the last few years, approximate methods have been developed to solve POMDPs (28, 56–59). Among these approaches we chose *symbolic perseus*, a point-based approach for factored POMDP. Point-based approaches approximate the value function by updating it only for some selected belief states. The point-based methods sample belief states by simulating some random interactions in the environment and then update the value function and its gradient over those sampled belief states. The symbolic perseus method (28) benefits from both the factored representations of FPOMDP using ADD and a fast point-based method using *Perseus* (58).

In summary, being able to model our optimization problems as a factored MDP (under perfect detection) or factored POMDP (under imperfect detection) is innovative and constitutes a key point of our work. Factored MDP and POMDP allow us to (i) compactly represent the optimization problems using trees and dynamic Bayesian networks, (ii) implement efficient computation using algebraic decision diagrams, and (iii) identify management rules by using decision trees to represent all of the optimal solutions in a structured way.

We derived an analytic solution for a population of invasive species following previous work by Chadès et al. (9) (*SI Methods*). We also derived an analytic approximation for n independent populations by transforming an n -population POMDP into a single POMDP (*SI Methods* and Fig. S5). Due to the high level of complexity of an n -population POMDP, we were unable to include interactions between populations. In most cases, interactions between populations are likely to increase the local persistence of each population. As a consequence, our approximation provides a lower time limit for which it is beneficial to manage or survey. This approximation constitutes a first step toward further theoretical development.

ACKNOWLEDGMENTS. We thank J. Firn, R. McVinish, P. Pollett, A. Franc, A. Sheppard, and R. McAllister for discussions. We are grateful to P. Poupart and J. Hoey for making available their toolbox. This work was supported by the Australian Centre of Excellence for Risk Analysis (I.C. and M.A.B.), the Applied Environmental Decision Analysis (I.C., H.P.P., and S.N.), the Centre of Excellence for Mathematics and Statistics of Complex Systems (I.C.), Australian Research Council Australian Research Fellowships (to Y.M.B. and H.P.P.), a Julius Career award (to T.G.M.), and a Commonwealth Scientific and Industrial Research Organisation (CSIRO) Office of the Chief Executive postdoctoral fellowship (to I.C.).

- Jones KE, et al. (2008) Global trends in emerging infectious diseases. *Nature* 451: 990–993.
- Pimentel D (2002) *Biological Invasions: Economic and Environmental Cost of Alien Plant, Animal, and Microbe Species* (CRC, Boca Raton, FL), p 384.
- Chapin FS, 3rd, et al. (2000) Consequences of changing biodiversity. *Nature* 405: 234–242.
- MacKenzie DI, Kendall WL (2002) How should detection probability be incorporated into estimates of relative abundance? *Ecology* 83:2387–2393.
- MacKenzie DI, et al. (2002) Estimating site occupancy rates when detection probabilities are less than one. *Ecology* 83:2248–2255.
- Stokstad E (2007) Endangered species. U.S. announces recovery plan for a ghost bird. *Science* 317:1158.
- Fisher DO, Blomberg SP (2011) Correlates of rediscovery and the detectability of extinction in mammals. *Proc R Soc Biol Sci Ser B*, 10.1098/rspb.2010.1579.
- Regan TJ, McCarthy MA, Baxter PWJ, Dane Panetta F, Possingham HP (2006) Optimal eradication: When to stop looking for an invasive plant. *Ecol Lett* 9:759–766.
- Chadès I, et al. (2008) When to stop managing or surveying cryptic threatened species. *Proc Natl Acad Sci USA* 105:13936–13940.
- Hauser CE, McCarthy MA (2009) Streamlining 'search and destroy': Cost-effective surveillance for invasive species management. *Ecol Lett* 12:683–692.
- Regan TJ, Chadès I, Possingham HP (2011) Optimal strategies for managing invasive plants in partially observable systems. *J Appl Ecol* 48:76–85.
- Day JR, Possingham HP (1995) A stochastic metapopulation model with variability in patch size and position. *Theor Popul Biol* 48:333–360.

13. Shea K, Possingham HP (2000) Optimal release strategies for biological control agents: An application of stochastic dynamic programming to population management. *J Appl Ecol* 37:77–86.
14. Hastings A, et al. (2005) The spatial spread of invasions: New developments in theory and evidence. *Ecol Lett* 8:91–101.
15. Tilman D, Kareiva PM (1997) *Spatial Ecology: The Role of Space in Population Dynamics and Interspecific Interactions* (Princeton Univ Press, Princeton, NJ).
16. Blackwood J, Hastings A, Costello C (2010) Cost-effective management of invasive species using linear-quadratic control. *Ecol Econ* 69:519–527.
17. Moody ME, Mack RN (1988) Controlling the spread of plant invasions: The importance of nascent foci. *J Appl Ecol* 25:1009–1021.
18. May RM (2006) Network structure and the biology of populations. *Trends Ecol Evol* 21:394–399.
19. Hethcote HW, Yorke JA (1984) Gonorrhea transmission dynamics and control. *Lect Notes Biomath* 56:1–110.
20. Allen LJS (1994) Some discrete-time SI, SIR, and SIS epidemic models. *Math Biosci* 124: 83–105.
21. Hanski I (1998) Metapopulation dynamics. *Nature* 396:41–49.
22. Wasserman S, Faust K (1994) *Social Network Analysis: Methods and Applications* (Cambridge Univ Press, Cambridge, UK).
23. Saito K, Kimura M, Motoda H (2009) Discovering influential nodes for SIS models in social networks. *Lect Notes Comput Sci* 5808:302–316.
24. Pastor-Satorras R, Vespignani A (2001) Epidemic spreading in scale-free networks. *Phys Rev Lett* 86:3200–3203.
25. Kephart JO, White SR, Chess DM (2002) Computers and epidemiology. *Spectrum IEEE* 30(5):20–26.
26. Boutilier C, Dean T, Hanks S (1999) Decision theoretic planning: Structural assumptions and computational leverage. *J Artif Intell Res* 11:1–94.
27. Hoey J, St-Aubin R, Hu A, Boutilier C (1999) SPUDD: Stochastic planning using decision diagrams. *Fifteenth Conference on Uncertainty in Artificial Intelligence*, eds Laskey K, Prade H (Morgan Kaufmann, Stockholm, Sweden), pp 279–288.
28. Poupard P (2005) Exploiting structure to efficiently solve large scale partially observable Markov decision processes. PhD thesis (Univ of Toronto, Toronto).
29. Levins R (1969) Some demographic and genetic consequences of environmental heterogeneity for biological control. *Bull ESA* 15:237–240.
30. May RM, Nowak MA (1994) Superinfection, metapopulation dynamics, and the evolution of diversity. *J Theor Biol* 170:95–114.
31. Proulx SR, Promislow DEL, Phillips PC (2005) Network thinking in ecology and evolution. *Trends Ecol Evol* 20:345–353.
32. Harris TE (1974) Contact interactions on a lattice. *Ann Probab* 2:969–988.
33. Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (Wiley, New York).
34. Milo R, et al. (2002) Network motifs: Simple building blocks of complex networks. *Science* 298:824–827.
35. Ferguson NM, et al. (2005) Strategies for containing an emerging influenza pandemic in Southeast Asia. *Nature* 437:209–214.
36. Keeling MJ, Woolhouse MEJ, May RM, Davies G, Grenfell BT (2003) Modelling vaccination strategies against foot-and-mouth disease. *Nature* 421:136–142.
37. Ferguson NM, et al. (2006) Strategies for mitigating an influenza pandemic. *Nature* 442:448–452.
38. Simberloff D (2003) How much information on population biology is needed to manage introduced species? *Conserv Biol* 17:83–92.
39. Freeman LC (1977) A set of measures of centrality based on betweenness. *Sociometry* 40:35–41.
40. Newman MEJ (2001) Scientific collaboration networks. II. Shortest paths, weighted networks, and centrality. *Phys Rev E Stat Nonlin Soft Matter Phys* 64:016132.
41. Newman MEJ (2003) The structure and function of complex networks. *SIAM Rev* 45: 167–256.
42. Estrada E, Rodríguez-Velázquez JA (2005) Subgraph centrality in complex networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 71:056103.
43. Pearce DW, Moran D (1994) *The Economic Value of Biodiversity* (Earthscan, London, UK).
44. Kuhnert PM, Martin TG, Griffiths SP (2010) A guide to eliciting and using expert knowledge in Bayesian ecological models. *Ecol Lett* 13:900–914.
45. Palomar DP (2006) A tutorial on decomposition methods for network utility maximization. *IEEE J Sel Areas Comm* 24:1439–1451.
46. Fortunato S, Barthélemy M (2007) Resolution limit in community detection. *Proc Natl Acad Sci USA* 104:36–41.
47. Reichardt J, Bornholdt S (2006) Statistical mechanics of community detection. *Phys Rev E Stat Nonlin Soft Matter Phys* 74:016110.
48. Mollison D (1977) Spatial contact models for ecological and epidemic spread. *J R Stat Soc B* 39:283–326.
49. Montoya JM, Pimm SL, Solé RV (2006) Ecological networks and their fragility. *Nature* 442:259–264.
50. Hufnagel L, Brockmann D, Geisel T (2004) Forecast and control of epidemics in a globalized world. *Proc Natl Acad Sci USA* 101:15124–15129.
51. Epanchin-Niell RS, Hastings A (2010) Controlling established invaders: Integrating economics and spread dynamics to determine optimal management. *Ecol Lett* 13: 528–541.
52. Astrom KJ (1965) Optimal control of Markov decision processes with incomplete state estimation. *J Math Anal Appl* 10:174–205.
53. Cassandra AR, Kaelbling LP, Littman ML (1994) Acting optimally in partially observable stochastic domains. *National Conference on Artificial Intelligence* (American Association for Artificial Intelligence, Seattle, WA), pp 1023–1028.
54. Papadimitriou CH, Tsitsiklis JN (1987) The complexity of Markov decision processes. *Math Oper Res* 12:441–450.
55. Madani O, Hanks S, Condon A (2003) On the undecidability of probabilistic planning and related stochastic optimization problems. *Artif Intell* 147:5–34.
56. Hauskrecht M (2000) Value-function approximations for partially observable Markov decision processes. *J Artif Intell Res* 13:33–94.
57. Pineau J, Gordon G, Thrun S (2003) Point-based value iteration: An anytime algorithm for POMDPs. *International Joint Conference on Artificial Intelligence*, (Lawrence Erlbaum, Acapulco, Mexico), pp 1025–1032.
58. Spaan MTJ, Vlassis N (2005) Perseus: Randomized point-based value iteration for POMDPs. *J Artif Intell Res* 24:195–220.
59. Smith T, Simmons R (2004) Heuristic search value iteration for POMDPs. *Uncertainty in Artificial Intelligence (AUAI)*, Arlington, VA, pp 520–527.