

Label-Setting Methods for Multimode Stochastic Shortest Path Problems on Graphs

Alexander Vladimirovsky

Department of Mathematics, Cornell University, Ithaca, New York 14853,
vlad@math.cornell.edu, <http://www.math.cornell.edu/~vlad/>

Stochastic shortest path (SSP) problems arise in a variety of discrete stochastic control contexts. An optimal solution to such a problem is typically computed using the value function, which can be found by solving the corresponding dynamic programming equations. In the deterministic case, these equations can be often solved by highly efficient label-setting methods (such as Dijkstra’s and Dial’s algorithms). In this paper we define and study a class of multimode stochastic shortest path (MSSP) problems and develop sufficient conditions for the applicability of label-setting methods. We illustrate our approach in a number of discrete stochastic control examples. We also discuss the relationship of SSPs with discretizations of static Hamilton-Jacobi equations and provide an alternative derivation for several fast (noniterative) numerical methods for these partial differential equations (PDEs).

Key words: stochastic shortest path; dynamic programming; label-setting; Dijkstra’s method; Dial’s method; optimal control; Hamilton-Jacobi PDEs; fast marching method

MSC2000 subject classification: Primary: 90C39, 90C40, 49L20, 65N22, 35B37; secondary: 93E20, 65K05, 49L25

OR/MS subject classification: Primary: dynamic programming; secondary: Markov finite state

History: Received June 30, 2007; revised February 28, 2008. Published online in *Articles in Advance* October 17, 2008.

1. Introduction. Stochastic shortest path (SSP) problems constitute a large class of Markov decision processes. Their accurate and efficient solution is important for numerous applications, including mathematical finance, optimal resource allocation, design of discrete-time risk-sensitive controls, and controlled queuing in communication networks. Our goal in this paper is to study the conditions under which an important subclass of SSPs can be solved by efficient label-setting methods.

In SSP, the current state of the system at the k th stage is \mathbf{y}_k , an element in a finite state space $X = \{\mathbf{x}_1, \dots, \mathbf{x}_M, \mathbf{t} = \mathbf{x}_{M+1}\}$. At the next stage, \mathbf{y}_{k+1} is a random variable whose probability distribution on X depends on \mathbf{y}_k and on the decision made (control value chosen) at the previous stage. The process terminates on reaching a special target state \mathbf{t} . At each stage, our choice of control determines the incurred cost, and the overall goal is to minimize the *value function* (i.e., the expected value of the total accumulated cost up to termination). We provide a formal description of SSP in §2; here we simply note that the dynamic programming approach yields a system of M coupled nonlinear equations for the value function. Under mild technical assumptions this system has a solution, which can be found by *value iteration*. However, since these iterations are performed in \mathbf{R}^M , this can be quite costly, especially considering that infinitely many steps are generally needed for convergence.

On the other hand, SSPs can be considered as a generalization of classical deterministic shortest path (SP) problems on directed graphs, for which there is a variety of well-understood efficient algorithms. In particular, noniterative *label-setting* methods are applicable provided the transition costs in the graph are nonnegative. If a constant $\kappa \ll M$ is an upper bound on outdegrees, Dijkstra’s method [12] and Dial’s method [11] solve the deterministic dynamic programming equations in $O(M \log M)$ and $O(M)$ operations, respectively. We provide a brief overview of these methods in §2.1; here we simply note that both methods hinge on the *absolute causality* present in a deterministic problem—the fact that the value function is decreasing along every optimal path to \mathbf{t} .

Thus, to build similar methods for SSPs, one needs to find similar causal properties in the stochastic problem. In fact, Bertsekas showed that a Dijkstra-like method will correctly compute the value function of an SSP if there exists a *consistently improving optimal policy* (Bertsekas [4, Vol. II, p. 98]). In §2.2, we define a similar notion of a *consistently δ -improving optimal policy*, which guarantees the applicability of a Dial-like method. Unfortunately, both of these criteria are implicit because the existence of such optimal policies is generally not known a priori.

The main contribution of this paper is the development of explicit conditions on transition cost function(s), which guarantee the existence of consistently improving and/or consistently δ -improving optimal policies for a large class of “*Multimode SSPs*.” The exact class of SSPs that we consider is formally defined in §3, but generally our criteria apply provided

- (i) each state $\mathbf{x} \in X$ has a collection of “modes” $m_1(\mathbf{x}), \dots, m_r(\mathbf{x})$ —(possibly overlapping) subsets of X ;
- (ii) each individual control is restricted to one of the modes (i.e., has nonzero transition probabilities only into the states available in that mode);

(iii) there exists an available control corresponding to each possible probability distribution over the states in each mode;

(iv) the control cost is defined for each mode separately as a continuous function of the corresponding probability distribution over the states in that mode.

In this setting, it is natural to interpret the decision made at each stage as a deterministic choice among the modes of \mathbf{y}_k plus the choice of a desirable probability distribution for the transition to one of the states in that mode.

This class obviously includes the classical SP problem (when each mode contains only one possible successor state). More interestingly, it includes the problem of selecting optimal randomized or mixed controls for deterministic SP problems, when such randomized or mixed controls might be available at a discount. In §3.1, we consider several representative examples and discuss the differences between *explicitly causal* problems (where the causality stems from a particularly simple structure of transition probabilities) and *absolutely causal* problems (where the applicability of label-setting methods stems from certain properties of transition costs, as derived in §3.3).

The MSSPs also include (but are not limited to) the Markov chain approximations of deterministic continuous optimal trajectory problems. (E.g., consider a vehicle starting at some point \mathbf{x} inside the domain $\Omega \subset \mathbf{R}^n$, which is controlled to minimize the time needed to reach the boundary $\partial\Omega$.) The value function for such problems is typically found as a viscosity solution of a static first-order Hamilton-Jacobi partial differential equation (PDE). It is well known that semi-Lagrangian discretizations of that PDE (similar to those in Falcone [13] and Gonzales and Rofman [14]) can also be obtained from controlled Markov processes on the underlying grids. This approach was pioneered by Kushner and Dupuis [18] to design approximation schemes for deterministic and stochastic continuously controlled processes. Recent extensions include higher-order approximations (Szpiro and Dupuis [31]) and methods for stochastic differential games (Kushner [17]). The resulting systems of equations are typically treated iteratively, but relatively recently provably convergent label-setting algorithms were introduced for several important subclasses. For the isotropic case (when the vehicle's speed depends only on its current position in Ω and is independent of the chosen direction of motion), the corresponding PDE is Eikonal. In 1994, Tsitsiklis introduced the first Dijkstra-like and Dial-like methods for semi-Lagrangian discretizations of this PDE on a uniform Cartesian grid (Tsitsiklis [32, 33]). The family of Dijkstra-like Fast Marching Methods, introduced by Sethian in [23] and extended by Sethian and coauthors in Sethian [25], Kimmel and Sethian [16], Sethian and Vladimirsky [26], was developed for Eulerian upwind discretizations of the Eikonal PDE in the context of isotropic front propagation problems. A detailed discussion of similarities and differences of these approaches can be found in Sethian and Vladimirsky [29]. More recently, another Dial-like method for the Eikonal PDE on a uniform grid was introduced by Kim [15]. For the anisotropic case, the resulting semi-Lagrangian discretization typically does not have that causal property and the label-setting methods are not directly applicable. The label-setting Ordered Upwind Methods (Sethian and Vladimirsky [27, 29]) circumvent this difficulty; the key idea behind them can be interpreted as “modifying the computational stencil on-the-fly to ensure the causality.” In the appendix of Sethian and Vladimirsky [29], we also demonstrated that the causality is present for the first-order semi-Lagrangian discretizations of the Eikonal PDE on arbitrary acute meshes.

In all of the above cases, the proofs of causality heavily relied on a geometric interpretation of the problem (e.g., a discretization of a particular PDE on a specific grid or mesh in \mathbf{R}^n). In contrast, we first demonstrate that the applicability of label-setting methods to MSSPs can be proven even if no geometric interpretation is available (§3). We then show that the absolute causality of prior numerical methods for the Eikonal PDE can be easily rederived from the more general criteria introduced in here. In addition, our formulation yields the following two new results for deterministic continuous optimal trajectory problems (§4):

- a formula for the bucket width in a Dial-like method for Eikonal PDEs on acute meshes;
- an applicability criterion for the label-setting techniques in anisotropic optimal control problems. Finally, in §5, we discuss the limitations of our approach and list several related open problems.

2. SSP problem. Typically, SSP is described on a directed graph with nodes $X = \{\mathbf{x}_1, \dots, \mathbf{x}_M, \mathbf{t} = \mathbf{x}_{M+1}\}$. Our exposition here closely follows the standard setting described in Bertsekas [4].

For each \mathbf{x}_i , the problem specifies a compact set of allowable controls $A_i = A(\mathbf{x}_i)$. If \mathbf{x}_i is the current state of the process (i.e., if $\mathbf{y}_k = \mathbf{x}_i$), then our choice of a control value $\mathbf{a} \in A_i$ determines the cost of the next transition $C(\mathbf{x}_i, \mathbf{a})$ as well as the probability of transition into each node \mathbf{x}_j ,

$$p(\mathbf{x}_i, \mathbf{x}_j, \mathbf{a}) = p_{ij}(\mathbf{a}) = P(\mathbf{y}_{k+1} = \mathbf{x}_j \mid \mathbf{y}_k = \mathbf{x}_i, \text{ and the chosen control is } \mathbf{a} \in A_i).$$

A class of problems where the transition cost $C(\mathbf{x}_i, \mathbf{a}, \mathbf{x}_j)$ also depends on the resulting successor node \mathbf{x}_j can also be handled in the same framework by defining $C(\mathbf{x}_i, \mathbf{a}) = \sum_{j=1}^{M+1} p_{ij}(\mathbf{a})C(\mathbf{x}_i, \mathbf{a}, \mathbf{x}_j)$. It is assumed that the cost is accumulated until we reach the *absorbing target* \mathbf{t} , i.e., $p_{it}(\mathbf{a}) = 1$ and $C(\mathbf{t}, \mathbf{a}) = 0, \forall \mathbf{a} \in A_t$.

Consider the class of control mappings $\mu: X \mapsto (\bigcup_{i=1}^M A_i)$ such that $\mu(x_i) \in A_i$ for all $\mathbf{x}_i \in X$. A *policy* is an infinite sequence of such mappings $\pi = (\mu_0, \mu_1, \dots)$. A *stationary policy* is a policy of the form (μ, μ, \dots) and for the sake of brevity, we will also refer to it as “the stationary policy μ .”

If the process starts at $\mathbf{x} \in X$ (i.e., $\mathbf{y}_0 = \mathbf{x}$), the expected cost of using a policy $\pi = (\mu_0, \mu_1, \dots)$ is defined as

$$\mathcal{J}(\mathbf{x}, \pi) = E\left(\sum_{k=0}^{\infty} C(\mathbf{y}_k, \mu_k(\mathbf{y}_k))\right).$$

The value function is then defined as usual, $U_i = U(\mathbf{x}_i) = \inf_{\pi} \mathcal{J}(\mathbf{x}_i, \pi)$, and a policy π^* is called *optimal* provided $U(\mathbf{x}_i) = \mathcal{J}(\mathbf{x}_i, \pi^*)$ for all $\mathbf{x}_i \in X$.

If the value function $U(\mathbf{x}_i)$ is finite, it should satisfy the dynamic programming equations: $U(\mathbf{t}) = 0$ and

$$U_i = \inf_{\mathbf{a} \in A_i} \left\{ C(\mathbf{x}_i, \mathbf{a}) + \sum_{j=1}^{M+1} p_{ij}(\mathbf{a})U_j \right\}, \quad \forall \mathbf{x}_i \in X \setminus \{\mathbf{t}\}. \tag{1}$$

An operator T is defined on \mathbf{R}^M componentwise by applying the right-hand side of Equation (1); i.e., for any $W \in \mathbf{R}^M$,

$$(TW)_i = \inf_{\mathbf{a} \in A_i} \left\{ C(\mathbf{x}_i, \mathbf{a}) + \sum_{j=1}^{M+1} p_{ij}(\mathbf{a})W_j \right\}. \tag{2}$$

Clearly, $U = \begin{bmatrix} U_1 \\ \vdots \\ U_M \end{bmatrix}$ is a fixed point of T and one hopes to recover U by *value iteration*:

$$W^{n+1} := TW^n \quad \text{starting from an initial guess } W^0 \in \mathbf{R}^M. \tag{3}$$

However, T generally is not a contraction unless all stationary policies are known to be proper (Bertsekas and Tsitsiklis [6]).

Bertsekas and Tsitsiklis [5] demonstrated the existence of a stationary optimal policy, the uniqueness of the fixed point of T , and that $W^n \rightarrow U$ for arbitrary $W^0 \in \mathbf{R}^M$ under the following assumptions:

- **(A0)** All $C(\mathbf{x}_i, \mathbf{a})$ are lower-semicontinuous and all $p_{ij}(\mathbf{a})$ are continuous functions of controls \mathbf{a} .
- **(A1)** There exists at least one *proper policy* (i.e., a policy, which reaches the target \mathbf{t} with probability 1 regardless of the initial state $\mathbf{x} \in X$).
- **(A2)** Every improper policy π will have cost $\mathcal{J}(\mathbf{x}, \pi) = +\infty$ for at least one node $\mathbf{x} \in X$.

(A0) and the compactness of control sets A_i allow us to replace “inf” with “min” in formulas (1) and (2). **(A1)** corresponds to a graph connectivity assumption in the deterministic case, while **(A2)** is a stochastic analog of the assumption that all cycles have positive cumulative penalty. **(A2)** also follows automatically if

$$\underline{C} = \min_{\mathbf{x} \in X \setminus \{\mathbf{t}\}, \mathbf{a} \in A(\mathbf{x})} C(\mathbf{x}, \mathbf{a}) > 0.$$

The convergence of value iteration provides a method to compute U , but generally that convergence does not occur after any finite number of iterations (for a simple example, see Figure 1). Some error bounds are available, but typically in an implicit form only (Bertsekas [4, Vol. I, §7.2]). A recent work by Bonet [7] provides a polynomial upper bound on the number of value iterations required to achieve a prescribed accuracy for the case when the ratio $(\|U\|_{\infty}/\underline{C})$ is a priori known to be polynomially bounded.

2.1. Label-setting methods: The deterministic case. Fast methods for deterministic discrete control problems (e.g., searching for an SP in a graph or a network) can be found in all standard references (e.g., Ahuja et al. [1], Bertsekas [3]) and we provide a brief overview just for the sake of completeness. The dynamic programming equations are much simpler in this case: $U(\mathbf{t}) = 0$ and

$$U_i = \min_{\mathbf{x}_j \in N(\mathbf{x}_i)} \{C_{ij} + U_j\}, \quad \forall \mathbf{x}_i \in X \setminus \{\mathbf{t}\}, \tag{4}$$




FIGURE 1. A simple example with only one available control at every node; the transition probabilities are indicated above and the control cost is $C > 0$. By the symmetry, it is clear that $U_1 = U_2 = u$ and $u = C + \frac{1}{2}(u + 0)$; thus $u = 2C$. At the same time, the generic value iteration described by formula (3) will not converge after any finite number of steps unless $W^0 = U$.

where $U(\mathbf{x}_i) = U_i$ is the min-cost-to-exit starting from \mathbf{x}_i ; $N(\mathbf{x}_i)$ is the set of nodes to which \mathbf{x}_i is connected; and $C_{ij} = C(\mathbf{x}_i, \mathbf{x}_j)$ is the cost of traversing the corresponding link. In the absence of negative cost cycles and if every \mathbf{x}_i is connected by some path to \mathbf{t} , the value function is finite and well defined. Value iteration (3) will converge to U after at most M iterations resulting in an $O(M^2)$ computational cost.

Label-setting methods provide a better alternative if a suitable lower bound on the transition costs is available. These methods reorder the iterations over the nodes to guarantee that each U_i is recomputed at most κ times, where the constant upper bound on outdegrees κ is assumed to be much smaller than the total number of nodes M . For example, Dijkstra's classical algorithm [12] is a label-setting method for the deterministic case provided all $C_{ij} \geq 0$. The idea is based on the causality of the system (4):

$$U_i \text{ may depend on } U_j \quad \text{only if } U_i \geq U_j. \quad (5)$$

Such an ordering is not known in advance and has to be obtained at runtime. The method subdivides X into two classes: *permanently labeled nodes* P and *tentatively labeled nodes* L and the values for \mathbf{x}_i 's in L are successively reevaluated using only the adjacent values already in P :

$$U(\mathbf{x}_i) := \min_{\mathbf{x}_j \in \tilde{N}(\mathbf{x}_i)} \{C_{ij} + U_j\}, \quad \forall \mathbf{x}_i \in L, \quad (6)$$

where $\tilde{N}(\mathbf{x}_i) = N(\mathbf{x}_i) \cap P$. The algorithm is initialized by placing all nodes into L and setting $U(\mathbf{t}) = 0$ and $U(\mathbf{x}_i) = +\infty$ for $i = 1, \dots, M$. At each stage, the algorithm chooses the smallest of tentative labels $U(\bar{\mathbf{x}})$, "accepts" $\bar{\mathbf{x}}$ (i.e., moves $\bar{\mathbf{x}}$ from L to P), and reevaluates U_i for each $\mathbf{x}_i \in L$ such that $\bar{\mathbf{x}} \in N(\mathbf{x}_i)$. Because $\bar{\mathbf{x}}$ is the only new element in $\tilde{N}(\mathbf{x}_i)$, that reevaluation can be more efficiently performed by setting

$$U(\mathbf{x}_i) := \min\{U(\mathbf{x}_i), (C(\mathbf{x}_i, \bar{\mathbf{x}}) + U(\bar{\mathbf{x}}))\}. \quad (7)$$

The algorithm terminates once the list L is empty, at which point the vector $U \in \mathbf{R}^M$ satisfies the system of Equations (4). The necessity to sort all (finite) temporary labels dictates the use of heap-sort data structures (Ahuja et al. [1]), usually resulting in the overall computational complexity of $O(M \log M)$.

In addition, if all $C_{ij} \geq \Delta > 0$, then Dial's label-setting method [11] is also applicable. The idea is to avoid sorting temporarily labeled nodes and instead place them into *buckets* of width Δ based on their tentative values. If $U(\bar{\mathbf{x}})$ is the *smallest* of tentative labels and $U(\mathbf{x})$ is currently in the same bucket, then even after $\bar{\mathbf{x}}$ is permanently labeled, it cannot affect $U(\mathbf{x})$ since

$$U(\mathbf{x}) < U(\bar{\mathbf{x}}) + \Delta \leq U(\bar{\mathbf{x}}) + C(\mathbf{x}, \bar{\mathbf{x}}).$$

Thus, a typical stage of Dial's method [11] consists of "accepting" (or declaring labels to be permanent) for everything in the current bucket, recomputing all nodes in L adjacent to those newly labeled permanent, switching them to other buckets if warranted by the new tentative labels, and then moving on to the next bucket. Because inserting and deleting from a bucket can be performed in $O(1)$ time, the overall computational complexity of Dial's method [11] becomes $O(M)$. In addition, while Dijkstra's approach [12] is inherently sequential, Dial's method [11] is naturally parallelizable. Many other enhancements of the above label-setting methods are available in the literature (e.g., see Bertsekas [3] and references therein). Most of those enhancements can also be used with the label-setting of SSP—provided the basic versions of the above algorithms are applicable.

2.2. Label-setting methods: The general SSP. Given a stationary policy μ for a general SSP, we can construct its directed *dependency graph* G_μ using the nodes $X \setminus \{\mathbf{t}\}$ and connecting \mathbf{x}_i to \mathbf{x}_j if $p_{ij}(\mu(x_i)) > 0$. Assuming **(A0)**, **(A1)**, and **(A2)**, it is easy to show that the value iteration for this problem converges after at most M iterations provided there exists an optimal stationary policy μ^* such that G_{μ^*} is acyclic. (See Bertsekas [4, Vol. II, §2.2.1].) We will refer to such SSPs as *causal*.

REMARK 2.1. This condition seems to forbid any self-transitions (e.g., $p_{ii}(\mu^*(x_i)) > 0, \forall \mathbf{x}_i \neq \mathbf{t}$), but an SSP with self-transitions can be converted into an SSP without them by setting

$$\tilde{C}(\mathbf{x}_i, \mathbf{a}) = \frac{C(\mathbf{x}_i, \mathbf{a})}{1 - p_{ii}(\mathbf{a})}; \quad \text{and} \quad \tilde{p}_{ij}(\mathbf{a}) = \begin{cases} 0 & \text{if } i = j, \\ \frac{p_{ij}(\mathbf{a})}{1 - p_{ii}(\mathbf{a})} & \text{if } i \neq j; \end{cases} \quad \text{for } \begin{matrix} i = 1, \dots, M; \\ j = 1, \dots, M + 1. \end{matrix}$$

REMARK 2.2. One obvious set of causal SSPs consists of all problems where the dependency graph is acyclic for every stationary policy. The SSP belongs to this class if and only if, $\forall \mathbf{x}_i, \mathbf{x}_j \in X \setminus \{\mathbf{t}\}$,

$$\exists \mu_1 \text{ s.t. there is a path from } \mathbf{x}_i \text{ to } \mathbf{x}_j \text{ in } G_{\mu_1} \implies \nexists \mu_2 \text{ s.t. there is a path from } \mathbf{x}_j \text{ to } \mathbf{x}_i \text{ in } G_{\mu_2}.$$

We will refer to such problems as *explicitly causal* (see Examples 3.1 and 3.5 in later sections). Such explicit causality is independent of the cost function and can be determined based on the available controls and the transition probabilities alone. The above property imposes a partial ordering on X ; using that partial ordering to go through the nodes, we will clearly have the value function computed correctly in a single sweep, yielding the computational complexity of $O(M)$. Thus the applicability of label-setting methods described below is only important for SSPs that are causal, but not explicitly causal. This is similar to the fact that the original Dijkstra's method [12] is not needed to solve the deterministic SP problem on any acyclic digraph.

According to the definition introduced by Bertsekas in [4], an optimal stationary policy μ^* is *consistently improving* if

$$p_{ij}(\mu^*(\mathbf{x}_i)) > 0 \implies U_i > U_j. \tag{8}$$

This is a stochastic equivalent of the causality condition (5). Thus the existence of such μ^* not only guarantees that G_{μ^*} is acyclic, but also allows us to avoid the value iteration process altogether since U_i 's can be computed by a noniterative Dijkstra-like method instead.

If a consistently improving optimal policy is known to exist, the new equivalent of the causal update Equation (6) for each $\mathbf{x}_i \in L$ is now

$$U(\mathbf{x}_i) := \min_{\mathbf{a} \in \tilde{A}(\mathbf{x}_i)} \left\{ C(\mathbf{x}_i, \mathbf{a}) + \sum_{j=1}^{M+1} p_{ij}(\mathbf{a}) U_j \right\}, \tag{9}$$

where $\tilde{A}(\mathbf{x}_i)$ is the set of controls, which make transition possible to permanently labeled nodes only, i.e., $\tilde{A}(\mathbf{x}_i) = \{\mathbf{a} \in A(\mathbf{x}_i) \mid p_{ij}(\mathbf{a}) = 0 \text{ for all } \mathbf{x}_j \notin P\}$. Once \bar{x} is moved from L to P , each $\mathbf{x}_i \in L$ needs to be updated only if the set

$$\tilde{A}(\mathbf{x}_i, \bar{x}) = \{\mathbf{a} \in \tilde{A}(\mathbf{x}_i) \mid p(\mathbf{x}_i, \bar{x}, \mathbf{a}) > 0\}$$

is not empty. Finally, the new equivalent of the efficient update formula (7) is now

$$U(\mathbf{x}_i) := \min \left\{ U(\mathbf{x}_i), \min_{\mathbf{a} \in \tilde{A}(\mathbf{x}_i, \bar{x})} \left\{ C(\mathbf{x}_i, \mathbf{a}) + \sum_{j=1}^{M+1} p_{ij}(\mathbf{a}) U_j \right\} \right\}. \tag{10}$$

If a constant $\kappa \ll M$ is an upper bound on stochastic outdegrees (i.e., if for $i = 1, \dots, M$, we have $\kappa \geq$ the total number of nodes \mathbf{x}_j for which $\exists \mathbf{a} \in A(\mathbf{x}_i)$ such that $p_{ij}(\mathbf{a}) > 0$), then a Dijkstra-like algorithm described above will have a computational cost of $O(M \log M)$. Upon termination, the resulting $U \in R^M$ will satisfy the system of Equations (1). The proof of this is straightforward and is listed as one of the exercises in Bertsekas [4].

Here, we introduce a similar definition:

Given $\delta \geq 0$, an optimal stationary policy μ^* is *consistently δ -improving* if

$$p_{ij}(\mu^*(\mathbf{x}_i)) > 0 \implies U_i > U_j + \delta. \tag{11}$$

When $\delta > 0$, it is similarly easy to show that the existence of a consistently δ -improving optimal policy guarantees the convergence of a Dial-like method with buckets of width δ to the value function of the SSP. As in

the deterministic case, for $\kappa \ll M$, the resulting cost will be $O(M)$. Every consistently δ -improving policy is obviously also consistently improving; when $\delta = 0$, this reduces to the previous definition.

Unfortunately, conditions (8) and (11) are implicit because no optimal policy is a priori known. Thus, for a general SSP, applicability of label-setting methods is hard to check in advance. It is preferable (and more practical) to develop conditions based on functions $C(\mathbf{x}_i, \mathbf{a})$ and $p(\mathbf{x}_i, \mathbf{x}_j, \mathbf{a})$, which would guarantee that every optimal policy is consistently improving (or δ -improving). We will refer to such SSPs as *absolutely δ -causal* (or simply *absolutely causal* when $\delta = 0$). Before developing such explicit conditions for a particular class of MSSPs in §3, we make several remarks about the general case.

REMARK 2.3 (CONSISTENTLY ALMOST IMPROVING POLICIES). Comparing the deterministic causality condition (5) with the condition (8), it might seem that a Dijkstra-like method should work whenever there exists a “consistently almost improving optimal policy,” i.e., an optimal μ^* such that

$$p_{ij}(\mu^*(\mathbf{x}_i)) > 0 \implies U_i \geq U_j.$$

A simple example in Figure 1 demonstrates that this is false. Indeed, for this example, a Dijkstra-like method would terminate with $U_1 = U_2 = +\infty$ even though the optimal policy is consistently almost improving and the correct value function is $U_1 = U_2 = 2C$.

REMARK 2.4 (LOWER BOUNDS ON CONTROL COST). If μ^* is an optimal policy and $\mathbf{a}^* = \mu^*(\mathbf{x}_i)$, then

$$C(\mathbf{x}_i, \mathbf{a}^*) = U_i - \sum_{j=1}^{M+1} p_{ij}(\mathbf{a}^*)U_j = \sum_{j=1}^{M+1} p_{ij}(\mathbf{a}^*)(U_i - U_j).$$

This means that $C(\mathbf{x}_i, \mathbf{a}^*) > \delta \geq 0$ when μ^* is δ -improving. Thus, when building label-setting methods, the natural class of problems to focus on is an SSP with

$$(A2') \quad C(\mathbf{x}_i, \mathbf{a}) > 0, \quad \text{for all } \mathbf{x}_i \in X \setminus \{\mathbf{t}\} \text{ and } \forall \mathbf{a} \in A(\mathbf{x}_i).$$

We note that (A2') and the compactness of all A_i 's imply (A2).

REMARK 2.5 (LABEL-SETTING ON A REACHABLE SUBGRAPH). Consider a reachable set X_c consisting of all nodes $\mathbf{x} \in X$ such that there exists a policy π leading from \mathbf{x} to \mathbf{t} with probability 1. Assumption (A1) states that $X = X_c$. If this is not the case, but the condition (A2') holds, then $U(\mathbf{x}_i) = +\infty$ for all $\mathbf{x}_i \notin X_c$. If a stationary policy μ^* is optimal, then (A2') implies $p_{ij}(\mu^*(\mathbf{x}_i)) = 0$ whenever $\mathbf{x}_i \in X_c$ and $\mathbf{x}_j \notin X_c$. If μ^* also satisfies (8) on X_c , then a Dijkstra-like method is still applicable. Upon its termination, the value function will be computed correctly on X_c and we will have $U(\mathbf{x}) = +\infty$ for all $\mathbf{x} \notin X_c$. (This is analogous to using the original Dijkstra's method [12] on a digraph that does not contain directed paths to \mathbf{t} from every $\mathbf{x} \in X$.) Of course, an efficient implementation will terminate the method as soon as all nodes remaining in L have a label of $+\infty$.

REMARK 2.6 (LABEL-SETTING FOR SSP: PRIOR WORK). It is natural to look for classes of SSPs for which either (8) or (11) is automatically satisfied by every optimal policy. One simple example is the deterministic case: If for every $\mathbf{x}_i \in X \setminus \{\mathbf{t}\}$ and $\forall \mathbf{a} \in A(\mathbf{x}_i)$, there exists $\mathbf{x}_j \in X$ such that $p_{ij}(\mathbf{a}) = 1$, then every optimal policy is consistently improving because of (A2'). Tsitsiklis [32, 33] was the first to prove causality of two truly stochastic SSPs, which he used to develop Dijkstra-like and Dial-like methods for two special discretizations of the Eikonal PDE on a uniform Cartesian grid. For Eikonal PDEs discretized on arbitrary acute meshes, the equivalent of property (8) for all optimal controls was proven in Sethian and Vladimirsky [29, Appendix]. Another implementation of a Dial-like method for the Eikonal PDE was introduced in Kim [15]. For the optimal control of hybrid systems, a similar property was used to build Dijkstra-like methods in Branicky and Hebbbar [9] and Sethian and Vladimirsky [28]. It is interesting to note that of all these papers, only Tsitsiklis's work mentions the SSP interpretation of the discretizations, but even in Tsitsiklis [33] the proof of causality is very problem specific and relies on the properties of the PDE and on a particular choice of the computational stencil. In §4, we use MSSPs to provide convergence criteria for Dijkstra's method in the above cases as well as the bucket width for Dial's method whenever it applies.

REMARK 2.7 (LABEL-CORRECTING METHODS FOR SSP). Whenever the value iteration converges after finitely many steps, *label-correcting methods* become another viable alternative. Their implementation for the deterministic case can be found in standard references (e.g., Ahuja et al. [1], Bertsekas [3]). Two such methods were introduced in Polymenakos et al. [21] for the SSP considered in Tsitsiklis [33]. In a more recent work (Bornemann and Rasch [8]), a similar method was applied to a finite element discretization of the Hamilton-Jacobi-Bellman PDE. In the latter case, the label-setting is used to obtain convergence-up-to-specified-tolerance even though the equivalent of condition (8) is not satisfied. Label-setting methods have an optimal worst-case computational cost; however, in practice, label-correcting methods can outperform them on many problems. The exact conditions under which this happens are still a matter of debate even in deterministic problems. While clearly interesting, the comparison of their performance on SSPs is outside the scope of this paper.

3. MSSP problems. We will use Ξ_n to denote the set of possible barycentric coordinates in \mathbf{R}^n , i.e.,

$$\Xi_n = \{\xi = (\xi_1, \dots, \xi_n) \mid \xi_1 + \dots + \xi_n = 1 \text{ and } \forall \xi_j \geq 0\}.$$

We will further define $I(\xi) = \{i \mid \xi_i > 0\}$ and use $\{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ to denote the standard canonical basis in \mathbf{R}^n . Finally, we will use $\mathbf{R}_{+,0}^n$ to denote the nonnegative orthant in \mathbf{R}^n , i.e., $\mathbf{R}_{+,0}^n = \{(x_1, \dots, x_n) \mid \forall x_j \geq 0\}$.

We will assume the following:

- (i) For every node $\mathbf{x}_i \in X \setminus \{\mathbf{t}\}$, we are given a list of “modes” $\mathcal{M}_i = \mathcal{M}(\mathbf{x}_i) = \{m_1, \dots, m_{r_i}\}$, where each mode $m \in \mathcal{M}_i$ is a nonempty subset of $X \setminus \{\mathbf{x}_i\}$ and $r_i = r(\mathbf{x}_i) = |\mathcal{M}_i| \geq 1$.
- (ii) The nodes within each mode are ordered; i.e., $m = (\mathbf{z}_1^m, \dots, \mathbf{z}_{|m|}^m)$, where $\mathbf{z}_j^m \neq \mathbf{z}_k^m$ if $j \neq k$.
- (iii) All controls have a special structure $\mathbf{a} = (m, \xi \in \Xi_{|m|})$ and there exists an available control $(m, \xi) \in A(\mathbf{x}_i)$ for all $m \in \mathcal{M}_i$ and all $\xi \in \Xi_{|m|}$.
- (iv) The corresponding transition probability is

$$p(\mathbf{x}_i, \mathbf{x}, (m, \xi)) = \begin{cases} \xi_j, & \text{if } \mathbf{x} = \mathbf{z}_j^m \text{ for some } j \in \{1, \dots, |m|\}; \\ 0, & \text{otherwise.} \end{cases}$$

- (v) The transition costs are defined for each mode separately, i.e., $C(\mathbf{x}_i, (m, \xi)) = C^m(\mathbf{x}_i, \xi)$.
- (vi) For all $\mathbf{x}_i \in X \setminus \{\mathbf{t}\}$ and for all $m \in \mathcal{M}_i$, the function $C^m(\mathbf{x}_i, \xi)$ is a positive continuous function of ξ .
- (vii) There exists a constant upper bound κ on stochastic outdegrees; i.e., $(\sum_{m \in \mathcal{M}_i} |m|) \leq \kappa$ for $i = 1, \dots, M$.

For these MSSPs, it is natural to interpret the decision made at each stage as a deterministic choice of a mode m plus the choice of a desirable probability distribution for the transition to one of the successor nodes in m . We note that the above framework is sufficiently flexible: each node can have its own number of modes, each mode can have its own number of successor nodes, and different modes can have overlaps (e.g., $\mathbf{z}_j^{m_1}$ can be the same as $\mathbf{z}_k^{m_2}$). The fully deterministic case is conveniently included when $|m| = 1$ for each mode m .

The above assumptions imply **(A0)** and **(A2')**; hence the value iteration converges at least on the reachable subgraph X_c (see Remarks 2.4 and 2.5).

The dynamic programming Equations (1) can be now rewritten as

$$U(\mathbf{x}) = \min_{m \in \mathcal{M}(\mathbf{x})} \{V^m(\mathbf{x})\}, \tag{12}$$

$$V^m(\mathbf{x}) = \min_{\xi \in \Xi_{|m|}} \left\{ C^m(\mathbf{x}, \xi) + \sum_{j=1}^{|m|} \xi_j U(\mathbf{z}_j^m) \right\}. \tag{13}$$

Before developing criteria for solvability of the above equations by label-setting methods (§§3.2 and 3.3), we provide a number of representative examples to illustrate the MSSP framework.

3.1. MSSPs and modeling. In this subsection, we list several examples of discrete stochastic control problems, which are naturally modeled in the MSSP framework. Our goal is twofold: (i) to explore the type of stochasticity present in MSSPs, and (ii) to understand which types of MSSPs make the development of label-setting methods worthwhile.

We begin by considering two very simple MSSPs, which illustrate the difference and relationship between explicit and absolute causalities.

EXAMPLE 3.1. For $M = 3$, suppose that each node has only one mode, and nodes $\mathbf{x}_1, \mathbf{x}_3$ have only one node \mathbf{t} in their modes. I.e., the transition to \mathbf{t} is deterministic and costs $C_{it} > 0$ for $i = 1, 3$. The \mathbf{x}_2 's only mode is $m = \{\mathbf{x}_1, \mathbf{x}_3\}$. (See Figure 2A.) Because the problem is so simple, it is clear that

$$U_1 = C_{1t}; \quad U_3 = C_{3t}; \quad U_2 = \min_{\xi \in \Xi_2} \{C^m(\mathbf{x}_2, \xi) + (\xi_1 U_1 + \xi_2 U_3)\}.$$

This SSP is obviously *explicitly* causal: U_2 will be computed correctly, provided it is computed after U_1 and U_3 (see Remark 2.2).

However, whether this SSP is *absolutely* causal depends on the cost function.

Suppose ξ^* is the unique minimizer of the above and C^m is such that

$$U_1 < U_2 < C^m\left(\mathbf{x}_2, \begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) + U_1 < U_3.$$




FIGURE 2. Two simple examples of MSSP. In both cases, starting from x_i , one needs to select an optimal probability distribution over two successor nodes (dashed and dotted lines) or to opt for the deterministic transition to t (priced at $C_{it} > 0$ and shown by solid lines wherever available).

If $\xi_2^* > 0$, it is clear that the Dijkstra-like method of §2.2 would compute U_2 incorrectly (since x_2 would be moved from L to P before x_3). If label-setting methods were to be used here, we would need to find conditions on $C^m(x_2, \xi)$, which make the above scenario impossible for any choice of positive C_{1t} and C_{3t} .

EXAMPLE 3.2. For a somewhat more interesting example, consider a circular doubly linked list of M nodes. (See Figure 2B for the case $M = 6$.) Each x_i has two modes: $m = (x_{\text{prev}}, x_{\text{next}})$ and $m' = \{t\}$.

The applicability of label-setting methods seems harder to judge in this case, but it is clear that we do not want x_i moved from L to P before its neighbors if the optimal choice at x_i involves a possible (nondeterministic) transition to one of them. For instance,

$$U_2 = \min \left\{ C_{2t}, \min_{\xi \in \Xi_2} \{ C^m(x_2, \xi) + (\xi_1 U_1 + \xi_2 U_3) \} \right\},$$

and we note that this equation is provably absolutely causal if in Example 3.1, a Dijkstra-like method produces correct U_2 for all allowable C_{1t} and C_{3t} . In fact, if the same $C^m(x_i, \xi)$ is used for each x_i , the above condition is sufficient to show the absolute causality of the full problem. This idea is generalized in §3.2.

In practical terms, whether or not the MSSP in Example 3.1 is absolutely causal is irrelevant because the value function can be easily computed directly (see Remark 2.2). On the other hand, Example 3.2 can be viewed as a variant of an optimal stopping problem, whose absolute causality would yield a more efficient alternative to the basic value iteration when M is large. We continue by considering a number of interesting single-mode-for-each-node examples.

In the opening act of Tom Stoppard’s [30] famous play, the title characters engage in statistical experimentation with supposedly fair coins. The fairness of their coins is highly suspect because they are observing a very long and uninterrupted run of “heads.” Rosencrantz (Ros) is bored by the game and would be glad to stop playing, but Guildenstern (Guil) insists on continuing. The following two examples are inspired by the above.

EXAMPLE 3.3. Suppose Guil will agree to stop only after observing K “heads” in a row. Ros has to pay some fee for every toss of a coin and is interested in minimizing his expected total cost up to the termination. Moreover, suppose that for each toss, Ros can request a coin with any probability distribution $(p, (1 - p))$ on possible outcomes (“heads” versus “tails”), but Guil intends to charge him a different fee $C(p)$ based on his request. The problem is to find an optimal $p_i^* \in [0, 1]$ that Ros should request after observing i “heads” in a row (i.e., in the state x_i).

Figure 3 (left) shows the graph representation of the game for $K = 3$. Denoting $x_K = t$, we set $U_K = 0$. Since there is exactly one mode per node, and two successor-nodes only ($\xi \in \Xi_2$; $\xi_1 = p, \xi_2 = (1 - p)$), the Dynamic Programming equations of this game can be rewritten as

$$U_i = \min_{\xi \in \Xi_2} \{ C(x_i, \xi) + \xi_1 U_{i+1} + \xi_2 U_0 \} = \min_{p \in [0, 1]} \{ C(p) + p U_{i+1} + (1 - p) U_0 \} \quad \text{for } i = 0, \dots, K - 1.$$

We note that the self-transition in the node x_0 can be dealt with in the spirit of Remark 2.1; see Figure 3 (right). This results in a deterministic transition to x_1 :

$$U_0 = C_{01} + U_1, \quad \text{where } C_{01} = \min_{p \in (0, 1]} \frac{C(p)}{p} = \frac{C(p_0^*)}{p_0^*}.$$

After this simplification, the example satisfies all the assumptions listed for MSSPs; therefore the applicability of label-setting methods can be determined by checking if $C(\xi)$ satisfies any of the criteria developed in §3.3.




FIGURE 3. The first Guildenstern and Rosencrantz game for $K = 3$ (left). After i “heads” in a row the **game-state** is \mathbf{x}_i . Transitions corresponding to “heads” and “tails” are shown by dashed and dotted lines, respectively. The self-transition in \mathbf{x}_0 can be removed and replaced by a deterministic transition (solid line) with the optimal cost C_{01} (right).

REMARK 3.1. Even though this MSSP is not explicitly causal, a simple structure of the graph makes it almost trivial for our purposes:

- (i) Every (eventually terminating) path from \mathbf{x}_i leads through \mathbf{x}_{i+1} , which implies $U_i > U_{i+1}$. Thus, label-setting methods are only applicable if $p_i^* = 1$ for all i .
- (ii) On the other hand, recursive relations similar to the one used above can be repeatedly employed to make this into a deterministic problem. For example,

$$\begin{aligned} U_i &= \min_{p \in [0,1]} \{C(p) + pU_2 + (1-p)U_0\} = \min_{p \in [0,1]} \{C(p) + pU_2 + (1-p)(C_{01} + U_1)\} \\ &= \min_{p \in [0,1]} \{[C(p) + (1-p)C_{01}] + pU_2 + (1-p)U_1\} = C_{12} + U_2, \end{aligned}$$

where

$$C_{12} = \min_{p \in (0,1)} \frac{C(p) + (1-p)C_{01}}{p} = \frac{C(p_1^*) + (1-p_1^*)C_{01}}{p_1^*}.$$

By repeating this procedure, we can compute the value function in $O(K)$ steps (counting the above minimization as a single operation) even if some of the p_i^* s are less than one (in which case the value iteration would not converge in a finite number of steps).

EXAMPLE 3.4. Now, suppose that Guil will agree to stop only after observing an uninterrupted run of K_h “heads” or K_t “tails”; see Figure 4.

Identifying $\mathbf{t} = \mathbf{x}_{K_h}^h = \mathbf{x}_{K_t}^t$ and $\mathbf{x}_0 = \mathbf{x}_0^h = \mathbf{x}_0^t$, we can rewrite the dynamic programming equations as

$$\begin{aligned} U(\mathbf{t}) &= 0; \\ U(\mathbf{x}_i^h) &= \min_{p \in [0,1]} \{C(p) + pU(\mathbf{x}_{i+1}^h) + (1-p)U(\mathbf{x}_i^t)\}; \quad \text{for } i = 0, \dots, K_h - 1; \\ U(\mathbf{x}_i^t) &= \min_{p \in [0,1]} \{C(p) + pU(\mathbf{x}_i^h) + (1-p)U(\mathbf{x}_{i+1}^t)\}; \quad \text{for } i = 0, \dots, K_t - 1. \end{aligned}$$

Because Remark 3.1 does not apply, in this case, it is possible to have a nontrivial optimal strategy (i.e., $p^* \in (0, 1)$), which might be computable by the label-setting methods. Their applicability can be guaranteed by certain properties of the cost function as will be shown by theorems in §3.3. For example, this MSSP is absolutely causal (and thus efficiently computable using a Dijkstra-like method regardless of specific values of K_h and K_t) for

$$C_1(p) = 3 + 2p - p^4 - (1-p)^2; \quad \text{or} \quad C_2(p) = \sqrt{p^2 + (1-p)^2}; \quad \text{or} \quad C_3(p) = 4 + (p - 0.5)^3.$$




FIGURE 4. The second Guildenstern and Rosencrantz game for $K_h = 4$ and $K_t = 3$. After i “heads” or “tails” in a row the **game-state** is \mathbf{x}_i^h or \mathbf{x}_i^t , respectively.

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.




FIGURE 5. The multitasking problem for $K_A = 3$ and $K_B = 2$. Each of the nodes \mathbf{x}_{i, K_B} or $\mathbf{x}_{K_A, j}$ has a deterministic transition to \mathbf{t} only. All other nodes $\mathbf{x}_{i, j}$ have a single mode $(\mathbf{x}_{i, j+1}, \mathbf{x}_{i+1, j})$. The node \mathbf{x}_{K_A, K_B} is not needed because the process always terminates before reaching it.

(Recalling that $\xi_1 = p$ and $\xi_2 = (1 - p)$, it will be easy to check that Theorems 3.1, 3.2, and 3.3 apply to C_1 , C_2 , and C_3 , respectively.)

A similar analysis works when Guil is allowed to use different prices, depending on the current state of the game (i.e., with $C(\mathbf{x}_i, p)$ instead of $C(p)$) or when the number of possible outcomes is higher (e.g., dice instead of coins, Ξ_6 instead of Ξ_2).

EXAMPLE 3.5. Suppose a person is engaged in multitasking, dividing her attention between activities A and B. This allocation of resources is described by $\xi = (\xi_A, \xi_B) \in \Xi_2$. We assume that

- per every time unit she reaches a new milestone in *exactly one* of these activities;
- the probability of a milestone reached in A or B is proportional to the fraction of her attention invested in that activity (ξ_A or ξ_B) during that time unit;
- the current state of the process $\mathbf{x}_{i, j}$ reflects the number of milestones reached in both activities;
- the cost (per time unit) of all possible resource allocations is specified by $C(\mathbf{x}_{i, j}, \xi)$;
- the process terminates after at least K_A milestones are reached in A or at least K_B milestones reached in B;
- the goal is to minimize the total expected cost up to a termination.

A particular instance of this problem is illustrated in Figure 5.

The above MSSP is obviously explicitly causal because the number of milestones achieved in each activity can only increase as time goes on regardless of the chosen policy. As usual with explicitly causal SSPs, the causal ordering of the nodes is a priori known regardless of the cost functions and the label-setting methods are really not needed. However, the following slight variation of the above is already computationally challenging.

EXAMPLE 3.6. Suppose the same person also dedicates a part of her attention to some distraction D and her resource allocation is now $\xi = (\xi_A, \xi_B, \xi_D) \in \Xi_3$, where ξ_D is the probability of getting completely distracted and inadvertently *resetting* the process (i.e., transition into $\mathbf{x}_{0,0}$).

If the diversion is appealing (i.e., if $C(\mathbf{x}_{i, j}, \xi)$ is a decreasing function of ξ_D), this problem is not explicitly causal and the applicability of label-setting methods becomes relevant. The possibility of self-transition in $\mathbf{x}_{0,0}$ is again dealt with in the spirit of Remark 2.1 and theorems from §3.3 can be then used to test for the absolute causality. Generalizations of this example (to an arbitrary number of activities and/or partial resets because of a diversion) can be handled similarly.

We note that the MSSPs occupy a niche in between purely deterministic and generally stochastic shortest path problems. It is easy to see that in all of the above examples, the stochastic aspect of the model is not due to some uncontrollable event (after all, the deterministic or pure controls are always available in MSSPs), but rather due to our belief that a randomized or mixed control might carry a lower cost.

REMARK 3.2 (RANDOMIZED OR MIXED CONTROLS AND DETERMINISTIC SP). In most deterministic discrete control problems, mixed policies are considered unnecessary. But this is mainly because of the fact that the cost of implementing such mixed or randomized controls is usually modeled by a linear function, i.e., $C^m(\mathbf{x}, \xi) = \sum_{j=1}^{|\mathbf{m}|} \xi_j C_j$. (More generally, Theorem 3.1 will show that an optimal control can be found among the pure controls $\{\mathbf{e}_j\}$ for any concave cost function.) However, if the cost is nonconcave, i.e., if $C^m(\mathbf{x}, \xi) < \sum_{j=1}^{|\mathbf{m}|} \xi_j C^m(\mathbf{x}, \mathbf{e}_j)$ for at least some $\xi \in \Xi_{|\mathbf{m}|}$, then a mixed strategy is available *at a discount* and might be advantageous.

The methods developed in this paper are therefore most useful for MSSPs that

- are not explicitly causal (otherwise direct methods are more efficient);
- but are absolutely causal because of (possibly nonconcave) costs satisfying criteria in §3.3.

Additional examples (stemming from discretizations of continuous optimal control problems) are discussed in §4.

3.2. Causality of MSSP and single-mode auxiliary problems. Checking whether a given MSSP is absolutely causal can be hard, although sufficient conditions can be developed hierarchically. This approach was already used in §3.1 to show the relationship between Examples 3.1 and 3.2.

In the general case, if μ^* is an optimal policy and $\mu^*(\mathbf{x}) = (m^*, \xi^*)$, then the formulas (12 and 13) imply

$$U(\mathbf{x}) = V^{m^*}(\mathbf{x}) = C^{m^*}(\mathbf{x}, \xi^*) + \sum_{j=1}^{|\mathbf{m}^*|} \xi_j^* U(\mathbf{z}_j^{m^*}).$$

If μ^* is consistently δ -improving, we have $(\xi_j^* > 0) \implies V^{m^*}(\mathbf{x}) > U(\mathbf{z}_j^{m^*}) + \delta$.

OBSERVATION 3.1. For each mode m , let $\Xi^* \subset \Xi_{|m|}$ be a set of all minimizers in formula (13). If

$$\begin{aligned} (\xi_j^* > 0) \implies V^m(\mathbf{x}) > U(\mathbf{z}_j^m) + \delta, & \quad \forall \mathbf{x} \in X \setminus \{\mathbf{t}\}; \quad \forall m \in \mathcal{M}(\mathbf{x}); \\ & \quad \forall \xi^* \in \Xi^*; \quad j = 1, \dots, |m|, \end{aligned} \tag{14}$$

then every optimal policy is consistently δ -improving (and this MSSP is absolutely δ -causal).

As a result, we can develop label-setting applicability conditions on a mode-per-mode basis. In the following, we will focus on one $\mathbf{x} \in X \setminus \{\mathbf{t}\}$ and one mode $m \in \mathcal{M}(\mathbf{x})$ to develop conditions on $C^m(\mathbf{x}, \cdot)$ that guarantee causality for all possible values of $U(\mathbf{z}_j^m)$ s. Since \mathbf{x} and m are fixed, we will simplify the notation by using

$$V = V^m(\mathbf{x}); \quad C(\cdot) = C^m(\mathbf{x}, \cdot); \quad W_j = U(\mathbf{z}_j^m); \quad n = |m|.$$

Furthermore, interpreting ξ and W as column vectors in \mathbf{R}^n , we define $F: \Xi_n \times \mathbf{R}^n \mapsto \mathbf{R}$ as follows:

$$F(\xi, W) = C(\mathbf{x}, \xi) + \xi^T W.$$

The dynamic programming Equation (13) can now be rewritten as

$$V = \min_{\xi \in \Xi_n} \left\{ C(\xi) + \sum_{j=1}^n \xi_j W_j \right\} = \min_{\xi \in \Xi_n} F(\xi, W). \tag{15}$$

Once the vector W is specified, this also determines the set of minimizers $\Xi^*(W) = \arg \min_{\xi \in \Xi_n} F(\xi, W)$.

DEFINITION 3.1. The mode m is *absolutely δ -causal* if

$$(\xi_j^* > 0) \implies V > W_j + \delta, \quad \forall W \in \mathbf{R}_{+,0}^n; \quad \forall \xi^* \in \Xi^*(W); \quad j = 1, \dots, n.$$

We will also refer to a mode as *absolutely causal* if the above holds at least with $\delta = 0$.

A simple way to interpret this definition is by considering an auxiliary single-mode MSSP on the nodes $\{\mathbf{x}, \mathbf{z}_1^m, \dots, \mathbf{z}_n^m, \mathbf{t}\}$ with a single mode for each node (see Figure 6). Let the transition from each \mathbf{z}_j^m to \mathbf{t} be deterministic with cost $C_{j\mathbf{t}} = W_j \geq 0$, and for \mathbf{x} , let the mode be $m = (\mathbf{z}_1^m, \dots, \mathbf{z}_n^m)$ using the transition cost $C^m(\mathbf{x}, \cdot)$ from the original problem.

The mode m is absolutely causal if a Dijkstra-like method solves the auxiliary single-mode problem correctly for every vector $W \in \mathbf{R}_{+,0}^n$. The mode is absolutely δ -causal if the same is true for a Dial-like method with buckets of width δ . In fact, Example 3.1 can be viewed as such an auxiliary problem for the mode $m \in \mathcal{M}(\mathbf{x}_2)$ of Example 3.2. We emphasize that the absolute causality of auxiliary problems is desirable not because we intend to use label-setting on any of them (after all, each auxiliary problem is explicitly causal, and a direct computation is efficient; see Remark 2.2), but because the label-setting methods might be advantageous on the original MSSP.

The conditions on the mode m in Definition 3.1 are more restrictive than those in (14) since in the latter case, the δ -causality is needed for only one (albeit unknown) vector W . Thus, Observation 3.1 yields the following sufficient condition for applicability of label-setting methods to MSSPs.




FIGURE 6. An auxiliary single-mode problem for $m \in \mathcal{M}(\mathbf{x})$. Deterministic transitions are shown by solid lines; $n = |m| = 4$.

COROLLARY 3.1. For a general MSSP, if every mode of every node is absolutely causal, then the MSSP is also absolutely causal and a Dijkstra-like method is applicable. If each mode m is absolutely δ_m -causal, then the MSSP is absolutely Δ -causal with

$$\Delta = \left(\min_{\mathbf{x} \in X \setminus \{t\}, m \in \mathcal{M}(\mathbf{x})} \{\delta_m\} \right),$$

and a Dial-like method is applicable if $\Delta > 0$.

We note that it is possible to have an absolutely causal MSSP, some of whose modes are not absolutely causal. This is reminiscent of the fact that the original Dijkstra’s method [12] might be converging correctly even for some deterministic problems with negative transition penalties.

3.3. Criteria on cost and absolute causality. Consider a mode $m \in \mathcal{M}(\mathbf{x})$ such that $n = |m| > 1$. In view of Corollary 3.1, it is important to find additional conditions on the transition cost function $C(\cdot) = C^m(\mathbf{x}, \cdot)$ that guarantee m ’s absolute δ -causality.

One obvious example is $C(\xi) = \sum_{j=1}^n \xi_j C_j$, where all C_j ’s are positive constants. In that case, F is linear and Equation (15) reduces to

$$V = \min_{\xi \in \Xi_n} F(\xi, W) = \min_{\xi \in \Xi_n} \left\{ \sum_{j=1}^n \xi_j (C_j + W_j) \right\} = \min_j \{(C_j + W_j)\},$$

which is not different from the deterministic SP Equation (4). The same principle works for arbitrary concave costs.

THEOREM 3.1. Suppose

$$(A3) \quad C: \mathbf{R}^n \mapsto \mathbf{R}_+ \text{ is concave.}$$

Then, the mode m is absolutely δ -causal with $\delta = \min_j C(\mathbf{e}_j)$. Moreover, V can be more efficiently evaluated as $V = \min_j \{(C(\mathbf{e}_j) + W_j)\}$.

PROOF. Since $F(\xi, W) = C(\xi) + \xi^T W$, we know that the function $F(\xi, W)$ is concave on Ξ_n . Thus, if $\xi^* \in \Xi^*(W) = \arg \min_{\xi \in \Xi_n} F(\xi, W)$, then

$$(\xi_j^* > 0) \implies F(\mathbf{e}_j, W) = F(\xi^*, W) \implies U - W_j = C(\mathbf{e}_j) \geq \delta > 0;$$

hence the mode m is absolutely δ -causal. \square

Homogeneous cost functions naturally arise in many SSPs. We recall that a function $L(\mathbf{y})$ is absolutely homogeneous of degree d if $L(a\mathbf{y}) = |a|^d L(\mathbf{y})$ for all $\mathbf{y} \in \mathbf{R}^n, a \in \mathbf{R}$. If L is also smooth, by Euler’s Homogeneous Function Theorem, it satisfies the equation $\mathbf{y}^T \nabla L(\mathbf{y}) = dL(\mathbf{y})$.

LEMMA 3.1. Suppose the cost

$$(A4) \quad C: \mathbf{R}_+^n \mapsto \mathbf{R}_+ \text{ is continuously differentiable and absolutely homogeneous of degree } d.$$

INFORMS holds copyright to this article and distributed this copy as a courtesy to the author(s). Additional information, including rights and permission policies, is available at http://journals.informs.org/.

Then, for every $W \in \mathbf{R}_{+,0}^n$, $\xi^* \in \Xi^*(W)$, we have

$$(\xi_j^* > 0) \implies V - W_j = \frac{\partial C}{\partial \xi_j}(\xi^*) - (d-1)C(\xi^*).$$

PROOF. For all $j \in I(\xi^*)$, the Kuhn-Tucker optimality conditions state that

$$\lambda = W_j + \frac{\partial C}{\partial \xi_j}(\xi^*), \tag{16}$$

where λ is a Lagrange multiplier. We recall that $\sum_{j \in I(\xi^*)} \xi_j^* = 1$. Multiplying (16) by ξ_j^* and summing over all $j \in I(\xi^*)$, we obtain

$$\lambda = \sum_{j \in I(\xi^*)} \xi_j^* \lambda = \sum_{j \in I(\xi^*)} \xi_j^* \left(W_j + \frac{\partial C}{\partial \xi_j}(\xi^*) \right) = \sum_{i=1}^n \xi_i^* W_i + \left(\sum_{i=1}^n \xi_i^* \frac{\partial C}{\partial \xi_i}(\xi^*) \right). \tag{17}$$

Thus, by Euler's Homogeneous Function Theorem,

$$\lambda = \sum_{i=1}^n \xi_i^* W_i + dC(\xi^*) = F(\xi^*, W) + (d-1)C(\xi^*). \tag{18}$$

Since ξ^* is a minimizer, $V = F(\xi^*, W) = \lambda - (d-1)C(\xi^*)$ and it follows from (16) that

$$V - W_j = \frac{\partial C}{\partial \xi_j}(\xi^*) - (d-1)C(\xi^*) \quad \text{for all } W \in \mathbf{R}_{+,0}^n, \quad \xi^* \in \Xi^*(W), \quad j \in I(\xi^*). \quad \square$$

THEOREM 3.2. If C satisfies (A4) and

$$(A5) \quad \frac{\partial C}{\partial \xi_j}(\xi) - (d-1)C(\xi) > \delta \geq 0, \quad \forall \xi \in \Xi_n, \quad \forall j \in I(\xi),$$

then the mode is absolutely δ -causal.

PROOF. If $\xi^* \in \Xi^*(W)$ and $\xi_j^* > 0$, condition (A5) and Lemma 3.1 imply that $V - W_j > \delta \geq 0$. \square

REMARK 3.3. The case most frequently encountered in applications of SSPs is the homogeneity of degree one. When $d = 1$, Equation (18) states that $\lambda = V$ and condition (A5) becomes even simpler:

$$(A5') \quad \frac{\partial C}{\partial \xi_j}(\xi) > \delta \geq 0, \quad \forall \xi \in \Xi_n, \quad \forall j \in I(\xi).$$

Lemma 3.1 and Theorem 3.2 can be viewed as generalizations of the key idea in proofs of causality in Tsitsiklis [33] and Sethian and Vladimirsky [29, Appendix].

REMARK 3.4. If (A4) holds and C is strictly convex, then (A5) is a necessary condition for the absolute δ -causality of the mode. Indeed, suppose (A5) is violated for some $\bar{\xi} \in \Xi_n$, $j \in I(\bar{\xi})$ and let $K = 1 + \max_i (\partial C / \partial \xi_i)(\bar{\xi})$. If for each $i = 1, \dots, n$, we choose $W_i = K - (\partial C / \partial \xi_i)(\bar{\xi})$, this ensures that $W \in \mathbf{R}_+^n$, $K = \lambda$, and $\Xi^*(W) = \{\bar{\xi}\}$, which implies $V \leq W_j + \delta$ even though $\bar{\xi}_j > 0$.

LEMMA 3.2. Suppose the cost

$$(A6) \quad C: \mathbf{R}_+^n \mapsto \mathbf{R}_+ \text{ is twice continuously differentiable.}$$

Then, for every $W \in \mathbf{R}_{+,0}^n$, $\xi^* \in \Xi^*(W)$, $j \in I(\xi^*)$, there exists a point $\hat{\xi}$ on the straight line segment $[\mathbf{e}_j, \xi^*]$ such that

$$V - W_j = C(\mathbf{e}_j) - \frac{1}{2}(\mathbf{e}_j - \xi^*)^T H(\hat{\xi})(\mathbf{e}_j - \xi^*),$$

where H is the Hessian matrix of $C(\xi)$.

PROOF. If $\xi^* \in \Xi^*(W)$ and $j \in I(\xi^*)$, then the Kuhn-Tucker optimality conditions yield two different formulas (16) and (17) for the Lagrange coefficient λ . By combining these, we see that

$$V - W_j = (V - \lambda) + (\lambda - W_j) = \left(C(\xi^*) - \sum_{i=1}^n \xi_i^* \frac{\partial C}{\partial \xi_i}(\xi^*) \right) + \frac{\partial C}{\partial \xi_j}(\xi^*) = C(\xi^*) + (\mathbf{e}_j - \xi^*)^T \nabla C(\xi^*).$$

By Taylor's theorem, there exists a point $\hat{\xi} \in [\mathbf{e}_j, \xi^*] \subset \Xi_n$ such that

$$C(\mathbf{e}_j) = C(\xi^*) + (\mathbf{e}_j - \xi^*)^T \nabla C(\xi^*) + \frac{1}{2}(\mathbf{e}_j - \xi^*)^T H(\hat{\xi})(\mathbf{e}_j - \xi^*);$$

thus $V - W_j = C(\mathbf{e}_j) - \frac{1}{2}(\mathbf{e}_j - \xi^*)^T H(\hat{\xi})(\mathbf{e}_j - \xi^*)$. \square

THEOREM 3.3. Consider an n by $(n - 1)$ matrix B , whose columns form an orthonormal basis for the subspace orthogonal to $[1, \dots, 1]^T \in \mathbf{R}^n$. Suppose the cost C satisfies **(A6)** and $\hat{H}(\xi) = B^T H(\xi) B$ is its projected Hessian. If $\Lambda(\hat{H}(\xi))$ is the maximum eigenvalue of $\hat{H}(\xi)$ and

$$(A7) \quad \min_i C(\mathbf{e}_i) > \delta + \max \left\{ 0, \max_{\xi \in \Xi_n} \Lambda(\hat{H}(\xi)) \right\},$$

then the mode is absolutely δ -causal.

PROOF. First, we assume that $\max_{\xi \in \Xi_n} \Lambda(\hat{H}(\xi)) > 0$ (the other case is already covered by Theorem 3.1). If $\xi^* \in \Xi^*(W)$ and $j \in I(\xi^*)$, then there exists $\beta \in \mathbf{R}^{n-1}$ such that $(\mathbf{e}_j - \xi^*) = B\beta$. We note that $\|\beta\| = \|\mathbf{e}_j - \xi^*\| \leq \sqrt{2}$. Since Lemma 3.2 applies,

$$V - W_j = C(\mathbf{e}_j) - \frac{1}{2} \beta^T \hat{H}(\hat{\xi}) \beta \geq C(\mathbf{e}_j) - \frac{1}{2} \|\beta\|^2 \Lambda(\hat{H}(\hat{\xi})) \geq \min_i C(\mathbf{e}_i) - \max_{\xi \in \Xi_n} \Lambda(\hat{H}(\xi)) > \delta. \quad \square$$

REMARK 3.5. Because the cost function is always evaluated on Ξ_n , condition **(A4)** is somewhat awkward: The cost can *always* be considered absolute homogeneous of degree one since $C(\xi)$ can be replaced by $\tilde{C}(\xi) = \|\xi\|_1 C(\xi/\|\xi\|_1)$, which has the same values as C on Ξ_n . A more meaningful question is: Assuming that C is smooth and homogeneous of degree one, what additional conditions on C and its directional derivatives inside Ξ_n are sufficient to guarantee **(A5')**? It is easy to see that **(A7)** is an answer to that question since $\max_{\xi \in \Xi_n} \Lambda(\hat{H}(\xi))$ is the upper bound on the second derivative of C restricted to any straight line in Ξ_n .

4. MSSPs approximating continuous deterministic problems. As already mentioned, MSSPs naturally arise in approximations of deterministic continuous optimal control problems. To illustrate this, we consider a class of time-optimal trajectory problems. Many variants of these problems are studied in robotic navigation, optimal control, and front propagation literature; a detailed discussion of the version presented here can be found in Sethian and Vladimirsky [29].

Suppose $\mathbf{y}(t) \in \mathbf{R}^2$ is the vehicle's position at the time t and the vehicle starts at $\mathbf{y}(0) = \mathbf{x}$ inside the domain Ω . We are free to choose any direction of motion (any vector in $S_1 = \{\mathbf{a} \in \mathbf{R}^2 \mid \|\mathbf{a}\| = 1\}$), but the speed will depend on the chosen direction and the current position of the vehicle. The vehicle's dynamics is governed by $\mathbf{y}'(t) = f(\mathbf{y}(t), \mathbf{a}(t))\mathbf{a}(t)$, where $f: \mathbf{R}^2 \times S_1 \mapsto \mathbf{R}$ is a Lipschitz-continuous speed function, satisfying $0 < F_1 \leq f(\mathbf{x}, \mathbf{a}) \leq F_2$ for all \mathbf{x} and \mathbf{a} . Additional exit-time-penalty q is incurred at the boundary; we will assume that $q: \partial\Omega \mapsto \mathbf{R}$ is nonnegative and Lipschitz continuous.

The goal is to cross the boundary $\partial\Omega$ as quickly as possible.

The value function of this problem is $u(\mathbf{x})$ (the minimal time to exit after starting from \mathbf{x}). It is well known that $u(\mathbf{x})$ is the unique viscosity solution (Crandall and Lions [10]) of the following static Hamilton-Jacobi-Bellman PDE:

$$\begin{aligned} \max_{\mathbf{a} \in S_1} \{(\nabla u(\mathbf{x}) \cdot (-\mathbf{a}))f(\mathbf{x}, \mathbf{a})\} &= 1, \quad \mathbf{x} \in \Omega \subset \mathbf{R}^2, \\ u(\mathbf{x}) &= q(\mathbf{x}), \quad \mathbf{x} \in \partial\Omega. \end{aligned} \quad (19)$$

The optimal trajectories coincide with the characteristic curves of this PDE. If the problem is *isotropic* (i.e., if $f(\mathbf{x}, \mathbf{a}) = f(\mathbf{x})$), the above PDE is equivalent to the usual Eikonal equation $\|\nabla u(\mathbf{x})\|f(\mathbf{x}) = 1$ and the optimal trajectories coincide with the gradient lines of $u(\mathbf{x})$.

For simplicity, we will first assume that the domain Ω is rectangular and that X is a uniform Cartesian grid on $\bar{\Omega}$. Concentrating on one particular gridpoint $\mathbf{x} \in X \cap \Omega$, we will number all of its neighbors as in Figure 7. Suppose that the optimal initial direction of motion \mathbf{a} lies in the first quadrant and assume that the corresponding optimal trajectory remains a straight line until intersecting the segment $\mathbf{x}_1\mathbf{x}_3$ at some point $\tilde{\mathbf{x}}$ (see Figure 7A). Then, it follows that

$$u(\mathbf{x}) = \frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{f(\mathbf{x}, \mathbf{a})} + u(\tilde{\mathbf{x}}).$$

Let $\tilde{\mathbf{x}} = \xi_1\mathbf{x}_1 + \xi_2\mathbf{x}_3$; a linear approximation yields

$$u(\mathbf{x}) \approx \frac{\|\tilde{\mathbf{x}} - \mathbf{x}\|}{f(\mathbf{x}, \mathbf{a})} + \xi_1 u(\mathbf{x}_1) + \xi_2 u(\mathbf{x}_3).$$

Of course, since $\tilde{\mathbf{x}}$ is not a priori known, we would have to minimize over all possible intersection points and all four quadrants.




FIGURE 7. Two simple stencils using four nearest neighbors (A) or eight nearest neighbors (B) on a uniform Cartesian grid.

We will enumerate all quadrants as follows: $\mathcal{M}(\mathbf{x}) = \{(\mathbf{x}_1, \mathbf{x}_3), (\mathbf{x}_3, \mathbf{x}_5), (\mathbf{x}_5, \mathbf{x}_7), (\mathbf{x}_7, \mathbf{x}_1)\}$. For any $(\mathbf{z}_1^m, \mathbf{z}_2^m) \in \mathcal{M}$ and any $\xi \in \Xi_2$, we can similarly denote

$$\tilde{\mathbf{x}}_\xi = \xi_1 \mathbf{z}_1^m + \xi_2 \mathbf{z}_2^m; \quad \tau(\xi) = \|\tilde{\mathbf{x}}_\xi - \mathbf{x}\|; \quad \mathbf{a}_\xi = (\tilde{\mathbf{x}}_\xi - \mathbf{x})/\tau(\xi).$$

We can now state a semi-Lagrangian discretization of the PDE (19):

$$U(\mathbf{x}) = \min_{(\mathbf{z}_1^m, \mathbf{z}_2^m) \in \mathcal{M}(\mathbf{x})} \min_{\xi \in \Xi_2} \left\{ \frac{\tau(\xi)}{f(\mathbf{x}, \mathbf{a}_\xi)} + (\xi_1 U(\mathbf{z}_1^m) + \xi_2 U(\mathbf{z}_2^m)) \right\}, \quad \forall \mathbf{x} \in X \cap \Omega; \quad (20)$$

$$U(\mathbf{x}) = q(\mathbf{x}), \quad \forall \mathbf{x} \in X \cap \partial\Omega.$$

This (fully deterministic) derivation is similar to the one used by Gonzales and Rofman in [14].

On the other hand, it is easy to see that this system of equations also describes the value function for an MSSP on $X \cup \{\mathbf{t}\}$:

- for the nodes $\mathbf{x} \in X \cap \partial\Omega$, there is a single (deterministic) transition to \mathbf{t} with the cost $q(\mathbf{x})$;
- for the nodes $\mathbf{x} \in X \cap \Omega$, the set of quadrants $\mathcal{M}(\mathbf{x})$ can be interpreted as a set of modes and $C^m(\mathbf{x}, \xi) = \tau(\xi)/f(\mathbf{x}, \mathbf{a}_\xi)$.

This interpretation is in the spirit of Kushner and Dupuis' [18] approach of approximating continuous optimal control by controlled Markov processes.

On a uniform Cartesian grid and the stencil of Figure 7A, we can express $\tau(\xi) = h\sqrt{\xi_1^2 + \xi_2^2}$, where h is the grid size. If the problem is isotropic, the cost function becomes $C(\mathbf{x}, \xi) = (h/f(\mathbf{x}))\sqrt{\xi_1^2 + \xi_2^2}$. A similar construction in \mathbf{R}^n leads to modes containing n neighbor-nodes each and the cost function

$$C(\mathbf{x}, \xi) = \frac{h}{f(\mathbf{x})} \sqrt{\sum_{i=1}^n \xi_i^2}. \quad (21)$$

This function is homogeneous of degree one in terms of ξ . Moreover, $(\partial C/\partial \xi_j)(\mathbf{x}, \xi) = (h^2/f(\mathbf{x}))(\xi_j/\tau(\xi))$, which is positive if and only if $\xi_j > 0$. By Theorem 3.2, each mode is absolutely causal and a Dijkstra-like method can be used to solve the problem. This is, in fact, the first of two methods introduced by Tsitsiklis in [33]. Since this C is also convex in ξ , Remark 3.4 shows that the modes are not absolutely δ -causal for any $\delta > 0$; hence, Dial's method is generally not applicable.

Another obvious computational stencil in \mathbf{R}^2 uses all eight neighboring gridpoints as shown in Figure 7B. Here, the optimal trajectory is still assumed to remain a straight line until the intersection with a segment, but the list of segments is different:

$$\mathcal{M}(\mathbf{x}) = \{(\mathbf{x}_1, \mathbf{x}_2), (\mathbf{x}_3, \mathbf{x}_2), (\mathbf{x}_3, \mathbf{x}_4), (\mathbf{x}_5, \mathbf{x}_4), (\mathbf{x}_5, \mathbf{x}_6), (\mathbf{x}_7, \mathbf{x}_6), (\mathbf{x}_7, \mathbf{x}_8), (\mathbf{x}_1, \mathbf{x}_8)\}.$$

The discretized Equation (20) still holds, but the difference is that

$$\tau(\xi) = \|\tilde{\mathbf{x}}_\xi - \mathbf{x}\| = \|(\xi_1 \mathbf{z}_1^m + \xi_2 \mathbf{z}_2^m) - \mathbf{x}\| = h\sqrt{1 + \xi_2^2}.$$

If the problem is isotropic, the cost function becomes $C(\mathbf{x}, \xi) = (h/f(\mathbf{x}))\sqrt{1 + \xi^2}$. Theorem 3.3 is certainly applicable, but instead we rewrite the above as a function homogeneous of degree one (see Remark 3.5): $C(\mathbf{x}, \xi) = (h/f(\mathbf{x}))\sqrt{(\xi_1 + \xi_2)^2 + \xi_2^2}$. We now note that

$$\frac{\partial C}{\partial \xi_1}(\mathbf{x}, \xi) = \frac{h^2}{f(\mathbf{x})} \frac{1}{\tau(\xi)}; \quad \frac{\partial C}{\partial \xi_2}(\mathbf{x}, \xi) = \frac{h^2}{f(\mathbf{x})} \frac{1 + \xi_2}{\tau(\xi)} \geq \frac{\partial C}{\partial \xi_1}(\mathbf{x}, \xi).$$

Since $\tau(\xi) \leq h\sqrt{2}$, we conclude that

$$\frac{\partial C^m}{\partial \xi_j}(\mathbf{x}, \xi) \geq \frac{h}{F_2\sqrt{2}} = \delta > 0, \quad \forall \mathbf{x} \in X \cap \Omega, \quad \forall m \in \mathcal{M}(\mathbf{x}), \quad \forall \xi \in \Xi_2, \quad j = 1, 2. \quad (22)$$

By the Theorem 3.2, each mode is absolutely δ -causal and a Dial-like method can be used with buckets of width δ (corresponding to the second method introduced in Tsitsiklis [33]).

More generally, suppose that X is a simplicial mesh on $\bar{\Omega} \subset \mathbf{R}^n$ with the minimum edge length of h . Let $S(\mathbf{x})$ be the set of all simplexes in the mesh that use \mathbf{x} as one of the vertices. Each such simplex $s \in S(\mathbf{x})$ corresponds to a single mode $m \in \mathcal{M}(\mathbf{x})$ consisting of all other vertices of s besides \mathbf{x} . For any mode $m = (\mathbf{z}_1^m, \dots, \mathbf{z}_n^m)$ and any $\xi \in \Xi_n$, we can similarly define

$$\tilde{\mathbf{x}}_\xi^m = \sum_{i=1}^n \xi_i \mathbf{z}_i^m; \quad \tau^m(\xi) = \|\tilde{\mathbf{x}}_\xi^m - \mathbf{x}\|; \quad \mathbf{a}_\xi^m = (\tilde{\mathbf{x}}_\xi^m - \mathbf{x})/\tau^m(\xi).$$

Since $\tau^m(\xi) = \sqrt{\sum_{i=1}^n \sum_{k=1}^n \xi_i \xi_k (\mathbf{z}_i^m - \mathbf{x})^T (\mathbf{z}_k^m - \mathbf{x})}$, we see that

$$\frac{\partial \tau^m}{\partial \xi_j}(\xi) = \frac{(\mathbf{z}_j^m - \mathbf{x})^T (\sum_{i=1}^n \xi_i (\mathbf{z}_i^m - \mathbf{x}))}{\tau^m(\xi)} = \frac{(\mathbf{z}_j^m - \mathbf{x})^T (\tilde{\mathbf{x}}_\xi^m - \mathbf{x})}{\tau^m(\xi)} = (\mathbf{z}_j^m - \mathbf{x})^T \mathbf{a}_\xi^m = \|\mathbf{z}_j^m - \mathbf{x}\| \cos \beta_{\xi,j},$$

where $\beta_{\xi,j}$ is the angle between \mathbf{a}_ξ^m and $(\mathbf{z}_j^m - \mathbf{x})$. Suppose $\beta(\mathbf{x}, m)$ is the maximum angle between a pair of vectors $(\mathbf{z}_k^m - \mathbf{x})$ and $(\mathbf{z}_i^m - \mathbf{x})$ maximizing over all $i, k \in \{1, \dots, n\}$. Furthermore, define

$$\beta(\mathbf{x}) = \max_{m \in \mathcal{M}(\mathbf{x})} \beta(\mathbf{x}, m); \quad \beta = \max_{\mathbf{x} \in X \cap \Omega} \beta(\mathbf{x}).$$

Since \mathbf{a}_ξ^m lies in the cone defined by $(\mathbf{z}_1^m - \mathbf{x}), \dots, (\mathbf{z}_n^m - \mathbf{x})$, we know that $\beta_{\xi,j} \leq \beta(\mathbf{x}, m) \leq \beta(\mathbf{x}) \leq \beta$. Therefore

$$\left(\beta < \frac{\pi}{2}\right) \implies \frac{\partial \tau^m}{\partial \xi_j}(\xi) = \|\mathbf{z}_j^m - \mathbf{x}\| \|\mathbf{a}_\xi^m\| \cos \beta_{\xi,j} \geq h \cos \beta > 0.$$

The dynamic programming equations in this case become

$$U(\mathbf{x}) = \min_{m \in \mathcal{M}(\mathbf{x})} \min_{\xi \in \Xi_n} \left\{ \frac{\tau^m(\xi)}{f(\mathbf{x}, \mathbf{a}_\xi^m)} + \left(\sum_{i=1}^n \xi_i U(\mathbf{z}_i^m) \right) \right\}, \quad \forall \mathbf{x} \in X \cap \Omega; \quad (23)$$

$$U(\mathbf{x}) = q(\mathbf{x}), \quad \forall \mathbf{x} \in X \cap \partial\Omega.$$

The cost function $C^m(\mathbf{x}, \xi) = \tau^m(\xi)/f(\mathbf{x}, \mathbf{a}_\xi^m)$ is homogeneous of degree one in ξ . For the isotropic case, we see that

$$\frac{\partial C^m}{\partial \xi_j}(\mathbf{x}, \xi) = \frac{1}{f(\mathbf{x})} \frac{\partial \tau^m}{\partial \xi_j}(\xi) = \frac{\|\mathbf{z}_j^m - \mathbf{x}\| \cos \beta_{\xi,j}}{f(\mathbf{x})}.$$

Thus, for the Eikonal PDE on any acute mesh (i.e., for $\beta < \pi/2$), each mode of the discretization is absolutely causal by Theorem 3.2 and a Dijkstra-like method is applicable (this is a rederivation of the result in Sethian and Vladimirsky [29, Appendix]). Moreover, if $\beta < \pi/2$, then $(\partial C^m / \partial \xi_j)(\mathbf{x}, \xi) \geq h \cos \beta / F_2 = \delta > 0$. This provides the optimal bucket width δ to use in a Dial-like method when solving the Eikonal PDE on any acute mesh. As far as we know, no general formula for δ has been derived elsewhere up till now.

We note that the last result is applicable even in a more general situation when the computational stencil $S(\mathbf{x})$ does not correspond to a set of nonoverlapping simplexes present in the mesh; e.g., in Figure 7B, $\beta = \pi/4$, and this yields the same δ as already derived in (22). That leads to an interesting dilemma: including more nearby nodes into a computational stencil usually decreases β and increases the bucket width, thus reducing the total number of “bucket acceptance” steps until the termination of Dial’s algorithm. On the other hand, a larger $S(\mathbf{x})$ increases both the computational complexity of a single step (more tentative labels to update after each

acceptance) and the discretization error (proportional to h in the above examples). Finding an optimal way for handling this trade-off, could further speedup noniterative methods for Eikonal PDEs on acute meshes. We note that h/F_2 remains the upper bound for δ and corresponds to the situation when the vehicle is allowed to move only along the directions $(\mathbf{z}_j^m - \mathbf{x})$.

A much harder question is the applicability of label-setting methods to semi-Lagrangian discretizations of anisotropic optimal control problems. It is well known that Equations (23) are generally not causal; this issue is discussed in detail in Sethian and Vladimirsky [29] and Vladimirsky [34]. On uniform Cartesian grids, the criteria for applicability of a Dijkstra-like method to anisotropic problems were previously provided in Sethian [24], Osher and Fedkiw [19], and more recently in Alton and Mitchell [2]. All of these criteria are grid-orientation dependent; i.e., given a Hamilton-Jacobi PDE, its semi-Lagrangian or Eulerian discretization may or may not be computed correctly by a Dijkstra-like method, depending on whether the anisotropy in the PDE is aligned with the grid directions. Here, we provide a criterion for applicability of a Dijkstra-like method for discretizations based on arbitrary acute stencils. In the anisotropic case,

$$\frac{\partial C^m}{\partial \xi_j}(\mathbf{x}, \xi) = \frac{f(\mathbf{x}, \mathbf{a}_\xi)(\partial \tau^m / \partial \xi_j)(\xi) - \tau^m(\xi)(\partial f / \partial \xi_j)(\mathbf{x}, \mathbf{a}_\xi)}{f^2(\mathbf{x}, \mathbf{a}_\xi)}.$$

Suppose there exists $\delta \geq 0$ such that $\forall \mathbf{x} \in X \cap \Omega, \forall m \in \mathcal{M}(\mathbf{x}), \forall \xi \in \Xi_n, \forall j \in \{1, \dots, n\}$,

$$(\xi_j > 0) \implies \frac{\partial f}{\partial \xi_j}(\mathbf{x}, \mathbf{a}_\xi) < \frac{f(\mathbf{x}, \mathbf{a}_\xi)[(\partial \tau^m / \partial \xi_j)(\xi) - \delta f(\mathbf{x}, \mathbf{a}_\xi)]}{\tau^m(\xi)}.$$

By Theorem 3.2, this implies that a Dijkstra-like method will be applicable and a Dial-like method will also be applicable if $\delta > 0$. Building label-setting methods based on this sufficient condition could potentially yield algorithms outperforming the Ordered Upwind Methods specially designed to restore the causality of anisotropic problems by dynamically extending the stencil (Sethian and Vladimirsky [27], Vladimirsky [34], Sethian and Vladimirsky [29]). We intend to explore this approach in future work.

5. Conclusions. We defined a large class of MSSP problems and derived a number of sufficient conditions to check the applicability of the label-setting methods. We illustrated the usefulness of our approach to the numerical analysis of first-order nonlinear boundary value problems by reinterpreting previous label-setting methods for the Eikonal PDE on Cartesian grids. For the Eikonal equation on arbitrary meshes, we reinterpreted the prior Dijkstra-like method and derived the new formula of bucket width for Dial-like methods. We also developed a new sufficient condition for the applicability of label-setting methods to anisotropic Hamilton-Jacobi PDEs on arbitrary stencils.

In practice, the applicability of label-setting methods to a particular SSP can be tested directly in $O(M)$ operations: Upon the method's termination, a single value iteration can be applied, and if it results in no changes, the value function was computed correctly. However, the sufficient conditions (presented above for MSSPs) allow us to avoid these additional computations.

Unfortunately, the framework of MSSPs is not flexible enough to express many common discrete stochastic control problems where not all possible probability distributions over successor nodes are available. Nevertheless, we hope that the key idea of our approach (splitting the original MSSP into a number of absolutely causal auxiliary problems) can be generalized to test the applicability of label-setting methods to other SSPs. Because SSPs can be naturally extended to describe stochastic games on graphs (Patek and Bertsekas [20]), we also intend to investigate the applicability of our approach to the latter. If successful, this could yield efficient numerical methods for a wide class of first- and second-order static Hamilton-Jacobi equations.

In Dial-like methods, the bucket width can sometimes be adjusted on the fly based on the not-yet-accepted part of the problem only. We expect such extensions to be advantageous for problems where the cost function C has very different lower bounds for different nodes. Another open question of practical importance is the use of label-setting methods to obtain an *approximation* of the value function for noncausal SSPs. Recently, a numerical method based on a related idea was introduced in Yatziv et al. [35] for Eikonal PDEs: A Dial-like method is used with buckets of width δ for a discretization that is not δ -causal. This introduces additional errors (analyzed in Rasch and Satzger [22]), but decreases the method's running time.

Finally, the performance comparison of label-setting and label-correcting methods on MSSPs is yet another interesting topic for future research.

Acknowledgments. The author thanks B. Van Roy, M. J. Todd, and D. Shmoys. The author's research is supported by National Science Foundation (NSF) Grants DMS-0514487 and CTS-0426787.

References

- [1] Ahuja, R. K., T. L. Magnanti, J. B. Orlin. 1993. *Network Flows*. Prentice-Hall, Upper Saddle River, NJ.
- [2] Alton, K., I. M. Mitchell. 2007. Fast marching methods for a class of anisotropic stationary Hamilton-Jacobi equations. Technical Report TR-2006-27, Department of Computer Science, University of British Columbia.
- [3] Bertsekas, D. P. 1998. *Network Optimization: Continuous & Discrete Models*. Athena Scientific, Boston.
- [4] Bertsekas, D. P. 2001. *Dynamic Programming and Optimal Control*, 2nd ed., Vols. I and II. Athena Scientific, Boston.
- [5] Bertsekas, D. P., J. N. Tsitsiklis. 1991. An analysis of stochastic shortest path problems. *Math. Oper. Res.* **16**(3) 580–595.
- [6] Bertsekas, D. P., J. N. Tsitsiklis. 1996. *Neuro-Dynamic Programming*. Athena Scientific, Boston.
- [7] Bonet, B. 2007. On the speed of convergence of value iteration on stochastic shortest-path problems. *Math. Oper. Res.* **32**(2) 365–373.
- [8] Bornemann, F., C. Rasch. 2006. Finite-element discretization of static Hamilton-Jacobi equations based on a local variational principle. *Comput. Visualization Sci.* **9**(2) 57–69.
- [9] Branicky, M. S., R. Hebban. 1999. Fast marching for hybrid control. *Proc. IEEE Internat. Sympos. Comput.-Aided Control System Design*, Kohala-Kona, HI.
- [10] Crandall, M. G., P.-L. Lions. 1983. Viscosity solutions of Hamilton-Jacobi equations. *Trans. Amer. Math. Soc.* **277**(1) 1–42.
- [11] Dial, R. 1969. Algorithm 360: Shortest path forest with topological ordering. *Comm. ACM* **12**(11) 632–633.
- [12] Dijkstra, E. W. 1959. A note on two problems in connection with graphs. *Numerische Mathematik* **1** 269–271.
- [13] Falcone, M. 1987. A numerical approach to the infinite horizon problem of deterministic control theory. *Appl. Math. Optim.* **15**(1) 1–13; corrigenda 23 (1991) 213–214.
- [14] Gonzales, R., E. Rofman. 1985. On deterministic control problems: An approximate procedure for the optimal cost, part I: The stationary problem. *SIAM J. Control Optim.* **23**(2) 242–266.
- [15] Kim, S. 2001. An $O(N)$ level set method for Eikonal equations. *SIAM J. Sci. Comput.* **22**(6) 2178–2193.
- [16] Kimmel, R., J. A. Sethian. 1998. Fast marching methods on triangulated domains. *Proc. Nat. Acad. Sci.* **95**(15) 8341–8435.
- [17] Kushner, H. J. 2002. Numerical approximations for stochastic differential games. *SIAM J. Control Optim.* **41**(2) 457–486.
- [18] Kushner, H. J., P. G. Dupuis. 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Academic Press, New York.
- [19] Osher, S., R. P. Fedkiw. 2001. *Level Set Methods: An Overview and Some Recent Results*. *J. Comput. Phys.* **169** 463–502.
- [20] Patek, S. D., D. P. Bertsekas. 1999. Stochastic shortest path games. *SIAM J. Control Optim.* **37**(3) 804–824.
- [21] Polymenakos, L. C., D. P. Bertsekas, J. N. Tsitsiklis. 1998. Implementation of efficient algorithms for globally optimal trajectories. *IEEE Trans. Automatic Control* **43**(2) 278–283.
- [22] Rasch, C., T. Satzger. 2007. Remarks on the $O(N)$ implementation of the fast marching method. Preprint, <http://www.citebase.org/abstract?id=oai:arXiv.org:cs/0703082>.
- [23] Sethian, J. A. 1996a. A fast marching level set method for monotonically advancing fronts. *Proc. National Acad. Sci.* **93**(4) 1591–1595.
- [24] Sethian, J. A. 1996b. *Level Set Methods and Fast Marching Methods: Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision and Materials Sciences*. Cambridge University Press, Cambridge, UK.
- [25] Sethian, J. A. 1999. Fast marching methods. *SIAM Rev.* **41**(2) 199–235.
- [26] Sethian, J. A., A. Vladimirsky. 2000. Fast methods for the Eikonal and related Hamilton–Jacobi equations on unstructured meshes. *Proc. National Acad. Sci.* **97**(11) 5699–5703.
- [27] Sethian, J. A., A. Vladimirsky. 2001. Ordered upwind methods for static Hamilton-Jacobi equations. *Proc. National Acad. Sci.* **98**(20) 11069–11074.
- [28] Sethian, J. A., A. Vladimirsky. 2002. Ordered upwind methods for hybrid control. *Proc. (LNCS 2289) 5th Internat. Workshop, HSCC 2002*, Stanford, CA, 393–406.
- [29] Sethian, J. A., A. Vladimirsky. 2003. Ordered upwind methods for static Hamilton-Jacobi equations: Theory and Algorithms. *SIAM J. Numer. Anal.* **41**(1) 325–363.
- [30] Stoppard, T. 1967. *Rosencrantz and Guildenstern Are Dead*. Faber, London, UK.
- [31] Szpiro, A., P. Dupuis. 2002. Second-order numerical methods for first-order Hamilton–Jacobi equations. *SIAM J. Numer. Anal.* **40**(3) 1136–1183.
- [32] Tsitsiklis, J. N. 1994. Efficient algorithms for globally optimal trajectories. *Proc. IEEE 33rd Conf. Decision and Control*, Lake Buena Vista, FL, 1368–1373.
- [33] Tsitsiklis, J. N. 1995. Efficient algorithms for globally optimal trajectories. *IEEE Trans. Automatic Control* **40**(9) 1528–1538.
- [34] Vladimirsky, A. 2001. Fast methods for static Hamilton-Jacobi partial differential equations. Doctoral dissertation, Department of Mathematics, University of California, Berkeley, CA.
- [35] Yatziv, L., A. Bartesaghi, G. Sapiro. 2006. $O(N)$ implementation of the fast marching algorithm. *J. Comput. Phys.* **212**(2) 393–399.