



HAL
open science

Optimal control subsumes harmonic control

Amine Boumaza, Bruno Scherrer

► **To cite this version:**

Amine Boumaza, Bruno Scherrer. Optimal control subsumes harmonic control. [Research Report] 2006, pp.8. inria-00119243

HAL Id: inria-00119243

<https://inria.hal.science/inria-00119243>

Submitted on 8 Dec 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Optimal control subsumes harmonic control

Amine Boumaza

LORIA Campus Scientifique B.P. 239
54506 Vandoeuvre-lès-Nancy, France
Email: amine.boumaza@loria.fr

Bruno Scherrer

LORIA Campus Scientifique B.P. 239
54506 Vandoeuvre-lès-Nancy, France
Email: bruno.scherrer@loria.fr

Abstract— We consider trajectory planning within the frameworks of optimal control and harmonic control. We present a formal evidence, in continuous domain and in a standard discretization, that harmonic control is the limit case of a some optimal control problem in which we make the noise level tend to infinity. In other words we show that optimal control subsumes harmonic control. We discuss properties of both paradigms and present simulations that illustrate this relationship.

I. INTRODUCTION

The use of harmonic functions, that is solutions to Laplace equation, as potential fields for trajectory planning was independently proposed by Connolly et al. [1] and Akishita et al. [2] in 1990 and was motivated by the fact that harmonic functions do not have local extrema (unlike other potential based methods like those described by Khatib [3]). In such an approach, obstacles in the configuration space correspond to maxima of the potential, while goals correspond to minima ; control algorithms then reduce to *locally* descend the potential until they reach the *global* minimum. The harmonic approach for control (we will, in the remaining, use the terms “harmonic control”) has had some impact on the robotics community (see e.g. [4]–[11]). However, the use of potential functions derived from partial differential equations (PDE), does not date from these seminal works on harmonic functions. Optimal control and reinforcement learning [12] are based on the computation of a potential: the value function. In a continuous domain, the theory of dynamic programming, which was pioneered in the 1950s by Bellman [13], implies that the value function satisfies a PDE, the Hamilton-Jacobi-Bellman equation. Similarly to the harmonic case, the *globally* optimal control is *locally* derived from the value function.

The motivation of this paper is to emphasize a strong relationship between harmonic and optimal control. Our contribution is to argue that the former is a special case of the latter: we formally show, in the continuous domain as well as in a standard finite difference discretization counterpart, that harmonic control is the limit case of trajectory planning with optimal control in a space with isotropic noise¹, when the noise level tends to infinity. In other words, we insist on the fact that our aim is not so much to propose some new control algorithms but to *explain* that many practitioners of harmonic control do optimal control without necessarily

knowing it. To our knowledge, this is the first time that such a close relation is established between these two paradigms. Among the articles that discuss both harmonic and optimal control, we did not find any that highlights such a specific-to-general relation: most articles rather consider them as complementary alternatives. Before going into further details, we review some of these related studies.

We begin by briefly discussing some articles that present harmonic and optimal control as unrelated but complementary paradigms. Trajectories derived from harmonic functions may be controlled through two parameters: the boundary constraints (Dirichlet or Neumann) and the conductance properties of the domain (solutions of discretized Laplace equations can be related to electrical potentials on a resistive network). Singh et al. [14] consider two variants of control with harmonic functions, one with Dirichlet conditions and the other with Neumann conditions, and use a reinforcement learning algorithm to combine them in order to shorten the trajectories. Coelho et al. [15] introduce a gradient-based algorithm for updating the inside conductance values so that the resulting trajectories become “closer to” shortest path trajectories. In both the above cases, the idea is to slightly modify the harmonic approach so that it becomes closer to the solution of the shortest path optimal control problem. In a different spirit, Perkins et al. [16] present harmonic functions as a possible Lyapunov function for constraining the set of actions taken by a reinforcement learning agent; this ensures that a goal state is reached at every exploration trial. In these works, harmonic control and optimal control are used as complementary techniques although, as we will show, they are strongly related.

There exist studies where even though the connection between harmonic control and optimal control seems close, it is not stated by the authors. Connolly et al. [17] propose a way to incorporate non-holonomic constraints² in harmonic control, where the idea is to use Neumann conditions that constrain the system to satisfy its real degrees of freedom. Though this is what is implicitly done in optimal control through the state dynamics’ function (see section II), no link with optimal control is made explicit. Masoud et al. [4] suggest a way to enforce directional constraints in some parts of the state space by introducing an anisotropic harmonic

¹Isotropic noise means that the noise is the same in every direction.

²Non-holonomic means that the controllable degrees of freedom are fewer than the number of state space degrees of freedom.

potential. This is strictly equivalent to having a constant drift (in the dynamics function) and a non trivial diffusion matrix in optimal control. Connolly [18] also provides an analysis of harmonic functions in terms of collision probabilities. He shows that, on a grid, the discrete harmonic function at one point is related to the probability of colliding with an obstacle given that the process follows a random walk from that point. He explains that a reinforcement learning algorithm, TD(0), can calculate such a harmonic function (in this case, the harmonic function is indeed equal to the value function of the random uniform policy). Though close, the link with optimal control is missed: 1) This analysis does not clarify to which extent “descending the gradient of the value function of the random uniform policy” makes sense. 2) The relaxation technique for computing the harmonic function is presented as a new rapid technique for computing the value function even though it is strictly equivalent to the (standard) Value Iteration algorithm. 3) Last but not least, only a discrete version of the problem is studied and the continuous setting, from which the harmonic function comes, is not addressed. The work presented in this paper addresses these issues: we show that a harmonic controller is the optimal controller of some navigation problem on a (finite difference) discretized version, as well as on the (original) continuous domain.

The organization of the rest of the paper is as follows. Section II shows how to do trajectory planning with isotropic noise within the framework of optimal control. Section III begins by formulating our main contribution: control with harmonic function corresponds to the case where, in this planning approach, we make the noise level tend to infinity. It goes on by describing some simulations and comments that intuitively explains what happens when the noise tends to infinity. Finally, section IV provides a discussion that clarifies some properties that are shared by the specific (harmonic) case and the general (optimal control) case and others that are not (interesting particular properties of the specific case over the general case).

II. OPTIMAL CONTROL: A GENERAL FRAMEWORK AND A SIMPLE INSTANCE

We begin by a brief introduction to optimal control theory, after which we will describe the instance that can be linked to harmonic control. In optimal control, one considers a system defined at time t by its state $x(t) \in \bar{\Omega}$ (the state space) where $\bar{\Omega} \subset \mathbb{R}^n$ is the closure of an open set Ω and $\partial\Omega$ is its boundary ($\bar{\Omega} = \Omega \cup \partial\Omega$). This system is controlled by $u(t) \in U$ where U is a compact set (the control space). The dynamics of such a system is governed by a stochastic differential equation (SDE):

$$dx = f(x(t), u(t)) dt + \sigma(x(t), u(t)) dw, \quad (1)$$

where w corresponds to a m -dimensional Wiener process or Brownian motion and σ is a $n \times m$ diffusion matrix. We consider the case of infinite time horizon. For any initial state x_0 , any control law $u(\cdot)$ and trajectory $x(\cdot)$, we note τ the exit time of $x(\cdot)$ from Ω , with the convention that $\tau = \infty$ when the trajectory stays infinitely within Ω . We define the discounted cost functional J as the expected cost over all

possible trajectories:

$$J(x, u(\cdot)) = E \left[\int_0^\tau \gamma^t c(x(t), u(t)) dt + \gamma^\tau C(x(\tau)) \right], \quad (2)$$

where c is the instantaneous (or running) cost at time t , C is the terminal cost at time τ and $\gamma \in (0, 1)$ is a discount factor that guarantees that J is bounded. We define the optimal value function J^* , function of the initial state x , as the minimum of the cost functional J for all control laws $u(\cdot)$:

$$J^*(x) = \min_{u(\cdot)} J(x, u(\cdot)).$$

We let $a = \sigma \cdot \sigma^T$ and note ∇J^* the gradient of J^* . Under reasonable assumptions³, it can be proved [19] that J^* is \mathcal{C}^2 and satisfies the following PDE known as the Hamilton-Jacobi-Bellman (HJB) equation:

$$J^*(x) \ln(\gamma) + \min_{u \in U} \left\{ c(x, u) + \nabla J^*(x) \cdot f(x, u) + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \frac{\partial^2 J^*}{\partial x_i \partial x_j} (x) \right\} = 0, \quad (3)$$

for $x \in \Omega$, with boundary conditions $\forall x \in \partial\Omega, J(x) = C(x)$. The optimal value function is particularly interesting since it enables to compute a deterministic optimal controller. For every state x the optimal control $u^*(x)$ is the argument u for which the min in equation 3 is attained, which is also known as the greedy controller with respect to J^* :

$$u^*(x) = \arg \min_{u \in U} \left\{ c(x, u) + \nabla J^*(x) \cdot f(x, u) + \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n a_{ij} \frac{\partial^2 J^*}{\partial x_i \partial x_j} (x) \right\}. \quad (4)$$

In order to make the connection with harmonic control, we now focus on a simple instance of optimal control problem. Consider that the state space $\bar{\Omega}$ is an environment in which an agent navigates. For simplicity of exposition, we consider the 2-dimensional case $\bar{\Omega} \subset \mathbb{R}^2$; note however that the reasoning generalizes directly to any dimension n . The boundary of the environment is decomposed into 2 sets $\partial\Omega = \mathcal{O} \cup \mathcal{G}$: \mathcal{O} is the set of obstacles and \mathcal{G} the set of goals. At each time $t > 0$, the agent is characterized by its coordinates $x(t) = (x_1(t), x_2(t))$ and its dynamics is related to a command direction $\theta(t)$ according to the following SDE:

$$dx = \vec{u}(\theta(t)) dt + \sigma dw.$$

where $\vec{u}(\theta)$ is a unit vector in the direction of $\theta \in [0, 2\pi]$, and σ is a positive constant. In other words, we consider an optimal control problem with a unit speed control ($u = \theta$ and $f(x, u) = \vec{u}(\theta)$) and a constant isotropic noise ($\sigma \in \mathbb{R}$). In this case it amounts to replace the matrix a by a real number σ^2 . Furthermore, we consider that there are no intermediate

³Sufficient conditions are: 1) the matrix $a = \sigma \cdot \sigma^T$ satisfies a ‘uniform parabolicity assumption’: there exists a $c > 0$ such that $\forall x \in \bar{\Omega}, \forall u \in U$ and $\forall y \in \mathbb{R}^n, \sum_{i=1}^n \sum_{j=1}^n a_{ij}(x, u) y_i y_j \geq c \|y\|^2$; 2) f, σ and the boundary $\partial\Omega$ are of class \mathcal{C}^2 , c and C are Lipschitzian.

costs ($\forall x \in \Omega, c(x) = 0$), only terminal costs C on the boundary: $C(x) = 1$ on obstacle boundaries ($\forall x \in \mathcal{O}$) and $C(x) = -1$ on goal boundaries ($\forall x \in \mathcal{G}$). In this specific case, the HJB equation (eq. 3) reduces to:

$$J^*(x) \ln(\gamma) + \min_{\theta \in [0, 2\pi]} \{ \nabla J^*(x) \cdot \vec{u}(\theta) \} + \frac{\sigma^2}{2} \Delta J^*(x) = 0 \quad (5)$$

for $x \in \Omega$ with boundary conditions $\forall x \in \partial\Omega, J^*(x) = C(x)$, where ΔJ^* denotes the Laplacian of J^* . Under this model, it is easy to see that the optimal control $\theta^*(x)$ in state x is such that $\vec{u} = \frac{-\nabla J^*(x)}{\|\nabla J^*(x)\|}$: indeed, in equation (5), the only term involved in the min is $\nabla J^*(x) \cdot \vec{u}(\theta)$, which in turn is minimal when θ is in the opposite direction of the gradient $\nabla J^*(x)$. Furthermore, for every $\gamma \in (0, 1)$, the value function of the controller $\theta(\cdot)$ (eq. 2) reduces to:

$$J(x, \theta(\cdot)) = E[\gamma^\tau | x(\tau) \in \mathcal{O}] Pr[x(\tau) \in \mathcal{O}] - E[\gamma^\tau | x(\tau) \in \mathcal{G}] Pr[x(\tau) \in \mathcal{G}]. \quad (6)$$

Note that the integral term in equation (2) vanishes because there are no instantaneous costs $c(\cdot)$. Let us now interpret what it means to minimize this specific value function: in order to minimize $J(x, \theta(\cdot))$, the optimal controller both tries to minimize $E[\gamma^\tau | x(\tau) \in \mathcal{O}] Pr[x(\tau) \in \mathcal{O}]$ and to maximize $E[\gamma^\tau | x(\tau) \in \mathcal{G}] Pr[x(\tau) \in \mathcal{G}]$. Since $\tau \mapsto \gamma^\tau$ is a decreasing function of τ , this intuitively means that the optimal controller both tries to maximize the time of hitting an obstacle and to minimize the time of reaching the goal. In the simple deterministic case⁴ when $\sigma = 0$, and for any value of $\gamma \in (0, 1)$, the optimal controller which consists in following the gradient of the optimal value function, is the one which minimizes the path length (and the time) to reach the goal.

In practice, one cannot compute exact analytical solutions to equation (5) and a usual approach, which we follow here, consists in building a finite difference scheme. This is what we describe now. To do so, we follow the lines of [21] and build a discrete-time discrete-space controlled Markov chain. Given a resolution $\delta > 0$, we build a grid Σ^δ and its border $\partial\Sigma^\delta$ on the domain of the problem. For simplicity we assume that any point of the border of the discretized domain belongs to the border of the initial domain: $\partial\Sigma^\delta \subset \partial\Omega$. Given a grid resolution δ , the function J is approximated by a function J^δ defined on $\Sigma^\delta \cup \partial\Sigma^\delta$. We define \cos^+ , \cos^- , \sin^+ and \sin^- as the positive and negative parts of \cos and \sin : $\cos^\pm(\theta) = \max(\pm \cos \theta, 0)$ and $\sin^\pm(\theta) = \max(\pm \sin \theta, 0)$. Notice that $\cos \theta = (\cos^+ \theta - \cos^- \theta)$ and $|\cos \theta| = (\cos^+ \theta + \cos^- \theta)$. Furthermore, we define the following transitions probability from grid point (x, y) to its four neighbors when going in direction of θ :

$$\begin{cases} p_\theta[(x, y), (x \pm \delta, y)] = \frac{1}{N_\theta} \left[\frac{\sigma^2}{2} + \delta \cos^\pm \theta \right] \\ p_\theta[(x, y), (x, y \pm \delta)] = \frac{1}{N_\theta} \left[\frac{\sigma^2}{2} + \delta \sin^\pm \theta \right] \end{cases} \quad (7)$$

where $N_\theta = \delta(\cos^+ \theta + \cos^- \theta + \sin^+ \theta + \sin^- \theta) + 4\sigma^2/2 = \delta(|\cos \theta| + |\sin \theta|) + 2\sigma^2$ can be regarded as a normalizing factor that ensures that the transition probabilities

$p_\theta((x, y), (x', y'))$ sum to 1. These transition probabilities on the discretized grid have a natural geometric interpretation: all of them have the same noise component ($\frac{\sigma^2}{2N_\theta}$), and two of them (the ones that are in the control direction θ) have non-zero weights (see figure 1(a)). We finally write $\tau(\theta) = \frac{\delta^2}{N_\theta}$,

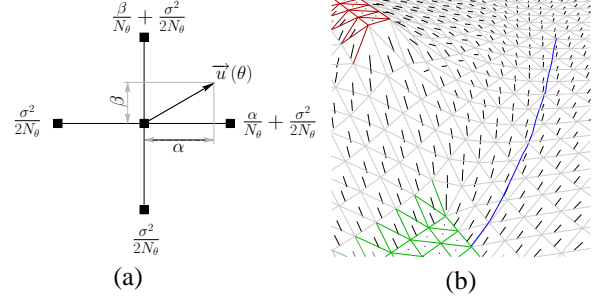


Fig. 1. (a) A geometric interpretation of the finite difference scheme illustrating the transition probabilities $p_\theta[(x, y), (x', y')]$: neighbors in the direction of the control have weights proportional to the coordinates of the direction $\vec{u}(\theta)$. (b) A discrete representation of J^δ , illustrating the orientation of the gradients, and a trajectory computed by interpolation of grid values.

which is a quantity that can be interpreted as the time needed to go from one grid point to another when following direction θ . Then, an approximation of J is obtained by computing the unknown J^δ in the following system:

$$J^\delta(x, y) = \gamma^{\tau(\theta_{x,y}^{J^\delta})} \sum_{(x', y') \in \Sigma^\delta} p_{\theta_{x,y}^{J^\delta}}[(x, y), (x', y')] J^\delta(x', y'), \quad (8)$$

for $(x, y) \in \Sigma^\delta$ and $J^\delta = C$ on $\partial\Sigma^\delta$, where $\theta_{x,y}^{J^\delta}$ is the angle that corresponds to the steepest slope direction at (x, y) when considering a piecewise linear interpolation of J^δ (see figure 1(b)). The notation $\theta_{x,y}^{J^\delta}$ may look a bit heavy, but it is important to remember that the optimal angle depends on the coordinate (x, y) and the value function J^δ . Introducing the following operator B^δ on Σ^δ :

$$B^\delta[W](x, y) = \gamma^{\tau(\theta_{x,y}^W)} \sum_{(x', y') \in \Sigma^\delta} p_{\theta_{x,y}^W}[(x, y), (x', y')] W(x', y')$$

for $(x, y) \in \Sigma^\delta$ and $B^\delta[W] = C$ on $\partial\Sigma^\delta$, equation (8) becomes $J^\delta = B^\delta[J^\delta]$. Since $\gamma < 1$ and $\tau(\cdot) > \frac{\delta}{\delta + 2\sigma^2} > 0$, the operator B^δ satisfies a contraction property (with contraction factor at least $\gamma^{\frac{\delta}{\delta + 2\sigma^2}}$). Therefore J^δ is unique, and it can be computed by *relaxation*, that is as the limit of the sequence $J_{n+1}^\delta \leftarrow B^\delta[J_n^\delta]$ when n tends to infinity. Once J^δ is computed, the (approximate) optimal control is the direction $\theta_{x,y}^{J^\delta}$ that descends the gradient of (the linear interpolation of) J^δ . It can be proved (see [21] for instance) that such a finite difference scheme is convergent: J^δ uniformly tends to J^* and the approximate optimal controller ∇J^δ tends to the optimal controller ∇J^* when the discretization resolution δ tends to 0.

III. HARMONIC CONTROL AS A SPECIFIC LIMIT CASE OF OPTIMAL CONTROL

Now that we have introduced an instance of optimal control problem for trajectory planning, we turn back to the main claim of this study: we demonstrate that when the noise level σ tends to infinity, the optimal control of our simple model tends to harmonic control.

⁴In this case, the analysis of the HJB is a bit more complex, involving viscosity solutions (see [20]) which we do not address in details here.

As we are interested in the dependence on σ , we denote J_σ the value function J^* of eq. 5. In order to study the behavior of J_σ when $\sigma \rightarrow \infty$, we regard it as a function of the Sobolev space $H^1(\Omega)$ with its norm $\|f\|_{H^1(\Omega)} = (\int_\Omega f^2 + \int_\Omega (\nabla f)^2)^{1/2}$ [22]. When $\sigma \rightarrow \infty$, we show that J_σ tends to a harmonic function, that is a function Φ that satisfies Laplace equation $\Delta\Phi = 0$.

Theorem 3.1: Let Φ be the unique function satisfying $\Delta\Phi = 0$ on Ω and $\Phi = C$ on $\partial\Omega$. Then the family of solutions $(J_\sigma)_{\sigma>0}$ of equation 5 satisfies

$$\lim_{\sigma \rightarrow \infty} \|J_\sigma - \Phi\|_{H^1(\Omega)} = \lim_{\sigma \rightarrow \infty} \left[\int_\Omega (J_\sigma - \Phi)^2 dx + \int_\Omega (\nabla J_\sigma - \nabla\Phi)^2 dx \right]^{1/2} = 0$$

Notice that this convergence is strong in the sense that 1) J_σ tends to Φ and that 2) its gradient ∇J_σ tends to $\nabla\Phi$.

We can prove a similar result for the standard finite difference approximations of J_σ and Φ . For a resolution $\delta > 0$ and the grid $\Sigma^\delta \cup \partial\Sigma^\delta$ introduced in section II, let Φ^δ be the standard 5-point finite difference approximation of Φ , that is the unique solution to:

$$\Phi^\delta(x, y) = \frac{1}{4} [\Phi^\delta(x + \delta, y) + \Phi^\delta(x - \delta, y) + \Phi^\delta(x, y + \delta) + \Phi^\delta(x, y - \delta)]$$

for all $(x, y) \in \Omega$ and $\Phi^\delta = C$ on $\partial\Omega$. When the noise level σ tends to infinity, the approximate value function J_σ^δ tends to the approximate harmonic function Φ^δ :

Theorem 3.2: Fix a discretization step $\delta > 0$. Let Φ^δ be the standard 5-point finite difference approximation of the harmonic function Φ . Then the family $(J_\sigma^\delta)_{\sigma>0}$ of discretizations of the value function defined in section II satisfies

$$\lim_{\sigma \rightarrow \infty} J_\sigma^\delta = \Phi^\delta \text{ uniformly on } \Sigma^\delta \cup \partial\Sigma^\delta.$$

The proofs of theorems 3.1 and 3.2 are in appendices A and B. On one hand, the proof of theorem 3.2 is rather straightforward (notice that it is specific to the finite difference schemes used to discretize both problems). On the other hand, if the result of theorem 3.1 seems intuitive (in eq. 5, we expect that when the term $\frac{\sigma^2}{2}$ before ΔJ^* tends to ∞ , the other terms get negligible), the proof we make relies on a few technical manipulations of eq. 5 and non-trivial properties of the Sobolev space $H^1(\Omega)$. In fact, the proof of theorem 3.1 we give in appendix A is more general than what we really need here: we show that the value function J_σ and its gradient ∇J_σ tend to the harmonic function Φ and its gradient $\nabla\Phi$ for any bounded control function f and any bounded cost c (recall the general optimal control equations 1 and 3); in the simple instance model, we take $u = \theta$, $f(x, u) = \vec{u}(\theta)$ and $c = 0$ so that the optimal controller matches the opposite direction of the gradient ∇J_σ .

Overall, we can summarize *our* convergence results (when $\sigma \rightarrow \infty$) and the standard numerical schemes convergences (when $\delta \rightarrow 0$) in the following diagram:

$$\begin{array}{ccc} (J_\sigma, \nabla J_\sigma) & \xrightarrow{\sigma \rightarrow \infty} & (\Phi, \nabla\Phi) \\ \delta \rightarrow 0 \uparrow & & \uparrow \delta \rightarrow 0 \\ (J_\sigma^\delta, \nabla J_\sigma^\delta) & \xrightarrow{\sigma \rightarrow \infty} & (\Phi^\delta, \nabla\Phi^\delta) \end{array}$$

When the noise level σ tends to infinity, the value function becomes harmonic, and the corresponding optimal control therefore consists in descending the gradient of this harmonic function. This limit case of optimal control is strictly equivalent to what Connolly [1] and Akishita [2] proposed in 1990: path planning is done by gradient descent of the function Φ . Furthermore, our formal analysis gives some insight on the (discrete) analysis of harmonic control by Connolly [18] discussed in the introduction. If harmonic control happens to be equivalent to “descending the gradient of the value function of a random uniform policy”, this is intuitively because the agent’s actions become negligible when the noise level gets infinitely large: at the limit, its movements are mostly due to the isotropic noise. Also, as we said that theorem 3.1 is general (it is true for any control function f and any cost c), the connection we make has a corollary that may be of interest to harmonic control practitioners: in the limit case, and for general $f(x, u)$ and $c(x, u)$, the optimal control theory prescribes to take the controller given by equation 4, which reduces here to:

$$u^*(x) = \arg \min_u \{c(x, u) + \nabla\Phi(x) \cdot f(x, u)\},$$

(and this is only equal in general to the gradient direction when $u = \theta$, $f(x, u) = \vec{u}(\theta)$ and $c = 0$). Recall that $f(x, u)$ describes the real dynamics for action u in state x while $c(x, u)$ can be seen as a bias to prevent from taking action u in state x . Though such a general control law might generate some complications (the controller may in general get stuck⁵), it constitutes a natural candidate for incorporating kinematical and dynamical constraints such as non-holonomicity and/or a notion of cost within the harmonic control framework.

We now present some simulations in order to illustrate what happens when the noise level tends to infinity. We have considered a maze environment containing two goals (G_1 and G_2 in figure 2). Dark lines represent obstacles, white represents free space and gray lines represent trajectories computed along the value function’s gradient. Goal G_1 at the top left of the environment is accessible through a narrow corridor, and goal G_2 is located in a fairly uncluttered area. Figure 2 illustrates, for five increasing values of the noise level σ , the trajectories computed starting from equally spread initial positions.

In the case where the dynamics is deterministic ($\sigma = 0$), the trajectories optimize the length to the goal; most of the trajectories starting in the upper part of the environment go toward goal G_1 whereas all others go toward goal G_2 . When we increase the noise level ($\sigma = 0.4$), the trajectories stray away from the walls and start to become smoother. With little more noise ($\sigma = 0.5$), the corridor leading to goal G_1 becomes hazardous and most of the trajectories go toward goal G_2 . At $\sigma = 1$, trajectories are smoother, and we begin to see some of them stopping before reaching a goal

⁵A non-zero c can block the trajectory away from the goal (since in practice the gradient of harmonic functions can be very small). If $c = 0$, then a sufficient condition for always getting to a goal is that f allows to move in any direction (possibly with different speeds) in the state space.

IV. DISCUSSION




Fig. 2. Evolution of trajectories with respect to the noise level σ . Each figure illustrates different trajectories leading to goals G_1 and G_2 , and starting at equally distributed grid points in the environment. All trajectories were computed using optimal control (with J_σ^δ) except the last one which is derived from harmonic control (with Φ^δ).

(trajectories starting on the right side): paths through narrow corridors becoming more dangerous, the controller favors to maintain the agent in safer areas away from obstacles. This behavior intensifies at $\sigma = 50$. A new phase begins at $\sigma = 500$: the noise being elevated, strategies consisting in staying in safe areas are not efficient anymore (the noise inevitably leads towards an obstacle). In this case, as a kamikaze that would know he is going to die anyway, it becomes more interesting to start moving again towards a goal. This behavior is more noticeable when we go on increasing the noise level. Eventually observe that the last simulation ($\sigma = 10^6$) experimentally confirms our theoretical result: the trajectories are close to the trajectories computed using a harmonic control.

We believe that harmonic control has had a significant impact in the literature because as a paradigm, it has several nice properties: *completeness*, *incrementality*, *flexibility*, and *parallel implementation*; we discuss them in this section. Moreover, as we showed that harmonic control is a special case of optimal control, it is interesting to check which of these properties are also shared by the general framework of optimal control. Conversely, we will highlight specific advantages of harmonic control over optimal control.

A fundamental property of harmonic functions is the absence of local optima. Whatever the starting position, planned trajectories always lead to a goal; we qualify such trajectories as complete. This property makes harmonic function particularly interesting compared to local potential field navigation [3] which might exhibit local optima. In general, the value function of an optimal control problem can have local optima. In the model we described in this paper, the underlying control law is complete in 2 cases: in the deterministic case ($\sigma = 0$) and, obviously, in the limit (harmonic) case ($\sigma = \infty$). In an environment with a large amount of noise, the control law may not attain a goal if this proves riskier (in the sense of the criterion defined in equation 6) than moving towards the goal.

Computation of harmonic functions using relaxation methods is performed iteratively, which allows incremental updates of the environment. Newly detected obstacles may be integrated in the model as new boundary conditions during computation. This property permits to use them in dynamic environments or in environments explored on-line, as for instance is explained in [5]–[7], [9]. Value functions, which can also be computed using relaxation, also permit incremental updating of the environment’s model: new obstacles and goals may be integrated during computation. More generally, updating incomplete or dynamic models of the environment within the framework of optimal control is known as “*indirect reinforcement learning*”. Moreover, research on “*direct reinforcement learning*” shows that incremental planning can be done even without maintaining an explicit model of the environment (see [12] for a general introduction, and [23] for a study of the continuous time case).

Trajectories generated using harmonic control can be quite flexible. Using Dirichlet boundary conditions, one can generate safe trajectories that have the tendency to stray away from obstacles. This is due to the fact that the potential flow is orthogonal to the boundaries. Using Neumann conditions, the potential flow is tangential to the boundary which permits to have riskier trajectories. It is also possible to combine both boundary conditions [14], [24], which allows to have different intermediate behaviors. In optimal control, trajectories may also be refined in numerous ways. Contrary to the harmonic case, we can express precisely, via the cost functions c and C , the relative severity of hitting certain obstacles compared to others, the relative importance of attaining certain goals compared to others and the relative importance of attaining goals rather than hitting obstacles.

In the same way, the laws of dynamics f , and the noise parameter σ , are parameters that may further explicitly influence the nature of the produced trajectories.

Another interesting property of harmonic functions that has been advocated is their inherent distribution of computation. The relaxation methods used to solve Laplace's equation are naturally distributable; computing grid values rely only on the local information at neighboring cells. Low level hardware implementation were proposed by Trassenko and Blake [25] implemented as resistive grid arrays. Furthermore parallelization makes this model a potential candidate to explain the computations of control undergone in the brain. For instance, Connolly and Burns [26] argue that the basal ganglia could compute harmonic function for motor control. The computation of the value function is also a natural candidate for parallelization. Iterative methods such as Gauss-Seidel or Jacobi, which solve the discrete version of the problem, may be implemented completely asynchronously using massively parallel architectures [27]. Though it is not clear whether these algorithms can be implemented in *very low level* architectures such as resistive grids, it is possible to implement them on parallel processors with communication delays and no synchronization [28]. Finally, as it is the case for harmonic functions, such parallel implementations motivated biological analogies with what may happen within the brain [29].

As we have just seen, most of the interesting properties of harmonic control are shared with optimal control, except the possibility of *very low level* implementation such as resistive grids. There is, however, another important property which is characteristic of harmonic control: Laplace equation (which characterizes the harmonic function) is a linear PDE for which one can derive analytical fundamental solutions *in any dimension* whereas this is generally not the case for the Hamilton-Jacobi-Bellman PDE. This does not make a real difference when one uses a relaxation technique for iteratively computing the potential. However, the linearity and the existence of such fundamental solution was exploited by Viéville et al. [30] to tackle the so-called curse of dimensionality: the harmonic function is decomposed in a sum of fundamental (possibly high-dimensional) harmonic functions. With such an approach, the authors can compute trajectories for a 10-degrees of freedom virtual 2D robotic arm in the presence of obstacles. Extending this idea to optimal control in general does not seem obvious and constitutes future research.

V. CONCLUSION

This paper presents a link between harmonic control and optimal control: we argue that optimal control subsumes harmonic control. The results presented show formally that for a well chosen optimal control problem, the value function becomes harmonic when the noise level tends to infinity and that in both cases, control consists in following the gradient of these functions. The convergence is showed in the continuous domain and in a standard finite difference scheme. We believe that such an analysis provides more

insight on harmonic control. We discussed different properties of both theories, and showed that many interesting properties of harmonic control are shared by optimal control. However, a specific power of harmonic controllers lies in the ability to implement them on very low level hardware as resistive arrays, and this does not seem to be the case for the Hamilton-Jacobi-Bellman equation. It would be interesting to investigate further the possibility for such hardware implementation of optimal control. Also, harmonic control relies on the computation of a simple linear PDE, which has some analytical fundamental solutions; this has recently been exploited to tackle the curse of dimensionality [30]. The transposition of this work to the optimal control framework is currently under investigation.

APPENDIX A

PROOF OF THEOREM 3.1: CONVERGENCE IN THE CONTINUOUS SETTING

We consider the family of functions J_σ that satisfy

$$\forall x \in \Omega, -\lambda J_\sigma(x) + \min_u [c(x, u) + \nabla J_\sigma(x) \cdot f(x, u)] + \frac{\sigma^2}{2} \Delta J_\sigma(x) = 0 \quad (9)$$

$$\forall x \in \partial\Omega, J_\sigma(x) = C(x)$$

and the harmonic function Φ with the same boundary conditions:

$$\begin{aligned} \forall x \in \Omega, \Delta \Phi(x) &= 0 \\ \forall x \in \partial\Omega, \Phi(x) &= C(x) \end{aligned}$$

where

- Ω is a bounded open subset of R^n .
- $\partial\Omega$, the boundary of the closure $\bar{\Omega}$ of Ω , and $C(\cdot)$ are "smooth" enough so that J_σ and Φ exist, are unique, and belong to the Sobolev space $H^1(\Omega)$ (see [22]).
- $\lambda \geq 0$ (when $\lambda = -\log \gamma$, $0 < \gamma \leq 1$ implies $\lambda \geq 0$)
- $\|f\|_\infty \leq M_f < \infty$ (when $f(x, \theta) = \vec{u}(\theta)$, we have $\|f\|_\infty = 1$)
- $\|c\|_\infty \leq M_c < \infty$.

The goal here is to show that J_σ (resp. ∇J_σ) converges to Φ (resp. $\nabla \Phi$) for the L_2 norm when σ tends to ∞ . It is equivalent to show that $U_\sigma = J_\sigma - \Phi$ and its gradient ∇U_σ tend to 0 for the L_2 norm. Replacing J_σ by $U_\sigma + \Phi$ in equation 9, we see that U_σ satisfies:

$$\begin{aligned} \forall x \in \Omega, -\frac{\sigma^2}{2} \Delta U_\sigma(x) + \lambda U_\sigma(x) = \\ -\lambda \Phi(x) + \min_u [c(x, u) + \nabla(U_\sigma + \Phi) \cdot f(x, u)] \quad (10) \\ \forall x \in \partial\Omega, U_\sigma(x) = 0 \end{aligned}$$

Note that since U_σ is null on the boundary $\partial\Omega$, U_σ belongs to $H_0^1(\Omega)$ ([22], Theorem 2 p. 259). From this point it is implied that U_σ and Φ are functions of x , and c and f are functions of (x, u) . Multiplying equation 10 by U_σ and integrating over Ω , we get

$$\begin{aligned} -\frac{\sigma^2}{2} \int_\Omega U_\sigma \Delta U_\sigma dx + \lambda \int_\Omega U_\sigma^2 dx = \\ -\lambda \int_\Omega U_\sigma \Phi dx + \int_\Omega U_\sigma \min_u [c + \nabla(U_\sigma + \Phi) \cdot f] dx. \quad (11) \end{aligned}$$

When integrating by parts, we have:

$$\begin{aligned} & -\frac{\sigma^2}{2} \int_{\Omega} U_{\sigma} \Delta U_{\sigma} dx \\ & = \frac{\sigma^2}{2} \int_{\Omega} \nabla U_{\sigma} \nabla U_{\sigma} dx - \frac{\sigma^2}{2} \int_{\partial\Omega} \nabla U_{\sigma} U_{\sigma} \vec{n} ds \\ & = \frac{\sigma^2}{2} \int_{\Omega} (\nabla U_{\sigma})^2 dx, \end{aligned}$$

where the second term (in which \vec{n} is the usual outward normal vector) vanishes since $U_{\sigma}|_{\partial\Omega} = 0$. Therefore, equation 11 becomes

$$\begin{aligned} & \frac{\sigma^2}{2} \int_{\Omega} (\nabla U_{\sigma})^2 dx + \lambda \int_{\Omega} U_{\sigma}^2 dx \\ & = -\lambda \int_{\Omega} U_{\sigma} \Phi dx + \int_{\Omega} U_{\sigma} \min_u [c + \nabla(U_{\sigma} + \Phi) \cdot f] dx. \end{aligned}$$

Noticing that all terms on the left are non-negative, and using the triangle inequality, we have:

$$\begin{aligned} & \frac{\sigma^2}{2} \int_{\Omega} (\nabla U_{\sigma})^2 dx + \lambda \int_{\Omega} U_{\sigma}^2 dx \\ & = \left| -\lambda \int_{\Omega} U_{\sigma} \Phi dx + \int_{\Omega} U_{\sigma} \min_u [c + \nabla(U_{\sigma} + \Phi) \cdot f] dx \right| \\ & \leq \left| -\lambda \int_{\Omega} U_{\sigma} \Phi dx \right| + \left| \int_{\Omega} U_{\sigma} \min_u [c + \nabla(U_{\sigma} + \Phi) \cdot f] dx \right| \end{aligned} \quad (12)$$

We now consider the right terms of equation 12:

- Using the triangle inequality, and the inequality $ab \leq \frac{a^2}{2} + \frac{b^2}{2}$, we see that

$$\begin{aligned} \left| -\lambda \int_{\Omega} U_{\sigma} \Phi dx \right| & \leq \lambda \left| \int_{\Omega} U_{\sigma} \Phi dx \right| \\ & \leq \lambda \int_{\Omega} |U_{\sigma}| |\Phi| dx \\ & \leq \frac{\lambda}{2} \int_{\Omega} U_{\sigma}^2 dx + \frac{\lambda}{2} \int_{\Omega} \Phi^2 dx \end{aligned}$$

- Using the triangle inequality, the fact that $|\min_u g(u)| \leq \sup_u |g(u)|$, the fact that for any vectors a and b , $|a \cdot b| \leq \|a\| \cdot \|b\|$, and the inequality $ab \leq \frac{a^2}{2} + \frac{b^2}{2}$, we see that

$$\begin{aligned} & \left| \int_{\Omega} U_{\sigma} \min_u [c(x, u) + \nabla(U_{\sigma} + \Phi) \cdot f(x, u)] dx \right| \\ & \leq \int_{\Omega} |U_{\sigma}| \left| \min_u [c(x, u) + \nabla(U_{\sigma} + \Phi) \cdot f(x, u)] \right| dx \\ & \leq \int_{\Omega} |U_{\sigma}| \sup_u |[c(x, u) + \nabla(U_{\sigma} + \Phi) \cdot f(x, u)]| dx \\ & \leq \int_{\Omega} |U_{\sigma}| \sup_u [|c(x, u)| + \|\nabla U_{\sigma} + \nabla \Phi\| \cdot \|f(x, u)\|] dx \\ & \leq \int_{\Omega} |U_{\sigma}| \left[\sup_u |c(x, u)| + \|\nabla U_{\sigma} + \nabla \Phi\| \sup_u \|f(x, u)\| \right] dx \\ & = \int_{\Omega} |U_{\sigma}| [M_c + \|\nabla U_{\sigma} + \nabla \Phi\| M_f] dx \end{aligned}$$

$$\begin{aligned} & \leq M_c \int_{\Omega} |U_{\sigma}| dx + M_f \int_{\Omega} |U_{\sigma}| [\|\nabla U_{\sigma}\| + \|\nabla \Phi\|] dx \\ & = M_c \int_{\Omega} |U_{\sigma}| dx \\ & \quad + M_f \left(\int_{\Omega} |U_{\sigma}| \|\nabla U_{\sigma}\| dx + \int_{\Omega} |U_{\sigma}| \|\nabla \Phi\| dx \right) \\ & \leq M_c \left(\frac{1}{2} \int_{\Omega} U_{\sigma}^2 dx + \frac{1}{2} \right) \\ & \quad + M_f \left(\int_{\Omega} U_{\sigma}^2 dx + \frac{1}{2} \int_{\Omega} (\nabla U_{\sigma})^2 dx + \frac{1}{2} \int_{\Omega} (\nabla \Phi)^2 dx \right) \\ & = \frac{2M_f + M_c + 1}{2} \int_{\Omega} U_{\sigma}^2 dx + \frac{M_f}{2} \left(\int_{\Omega} (\nabla U_{\sigma})^2 dx + \int_{\Omega} (\nabla \Phi)^2 dx \right) \end{aligned}$$

Gathering everything back into 12, we can write:

$$\begin{aligned} & \frac{\sigma^2}{2} \int_{\Omega} (\nabla U_{\sigma})^2 dx + \lambda \int_{\Omega} U_{\sigma}^2 dx \leq \frac{\lambda}{2} \int_{\Omega} U_{\sigma}^2 dx + \frac{\lambda}{2} \int_{\Omega} \Phi^2 dx \\ & + \frac{2M_f + M_c + 1}{2} \int_{\Omega} U_{\sigma}^2 dx + \frac{M_f}{2} \left(\int_{\Omega} (\nabla U_{\sigma})^2 dx + \int_{\Omega} (\nabla \Phi)^2 dx \right) \end{aligned}$$

Putting all terms involving ∇U_{σ} on the left, this can be rewritten as follows:

$$\left(\frac{\sigma^2 - M_f}{2} \right) \int_{\Omega} (\nabla U_{\sigma})^2 dx \leq \frac{2M_f + M_c + 1 - \lambda}{2} \int_{\Omega} U_{\sigma}^2 dx + K$$

where $K = \frac{\lambda}{2} \int_{\Omega} \Phi^2 dx + \frac{M_f}{2} \int_{\Omega} (\nabla \Phi)^2 dx$ is a finite number that does not depend on σ . Since $\lambda \geq 0$, and we also have:

$$\left(\frac{\sigma^2 - M_f}{2} \right) \int_{\Omega} (\nabla U_{\sigma})^2 dx \leq M \int_{\Omega} U_{\sigma}^2 dx + K$$

with $M = \frac{2M_f + M_c + 1}{2}$.

Poincaré's inequality ([22], Theorem 3 p. 265) states that there exists a constant $A < \infty$ such that for all functions $v \in H_0^1(\Omega)$

$$\int_{\Omega} v dx \leq A \int_{\Omega} (\nabla v)^2 dx.$$

Since for all σ , $U_{\sigma} \in H_0^1(\Omega)$, we have

$$\int_{\Omega} U_{\sigma}^2 dx \leq A \int_{\Omega} (\nabla U_{\sigma})^2 dx. \quad (13)$$

Therefore

$$\left(\frac{\sigma^2 - M_f}{2} \right) \int_{\Omega} (\nabla U_{\sigma})^2 dx \leq AM \int_{\Omega} (\nabla U_{\sigma})^2 dx + K$$

and consequently

$$\int_{\Omega} (\nabla U_{\sigma})^2 dx \leq \frac{K}{\frac{\sigma^2 - M_f}{2} - AM} \xrightarrow{\sigma \rightarrow \infty} 0.$$

By equation 13, we also automatically get

$$\int_{\Omega} U_{\sigma}^2 dx \xrightarrow{\sigma \rightarrow \infty} 0.$$

APPENDIX B

PROOF OF THEOREM 3.2: CONVERGENCE IN THE DISCRETE CASE

Here we prove that, for any discretization level δ , the approximate value function J_{σ}^{δ} tends to the approximate harmonic function Φ^{δ} when the noise level σ tends to infinity. For clarity, and because we consider here that δ is a predefined constant, we remove the δ exponent on functions and operators. We know that J_{σ} and Φ satisfy fixed point equations: $J_{\sigma} = \mathcal{B}_{\sigma} J_{\sigma}$ and $\Phi = \mathcal{H}\Phi$. Therefore

$$\begin{aligned}
J_\sigma - \Phi &= \mathcal{B}_\sigma J_\sigma - \mathcal{H}\Phi \\
&= \mathcal{B}_\sigma J_\sigma - \mathcal{H}J_\sigma + \mathcal{H}J_\sigma - \mathcal{H}\Phi \\
&= (\mathcal{B}_\sigma J_\sigma - \mathcal{H}J_\sigma) + \mathcal{H}(J_\sigma - \Phi)
\end{aligned}$$

\mathcal{H} is a linear operator such that $(I - \mathcal{H})$ is invertible (\mathcal{H} is a contraction mapping, which assures that the equation $\Phi = \mathcal{H}\Phi$ has a unique solution), thus:

$$J_\sigma - \Phi = (I - \mathcal{H})^{-1} (\mathcal{B}_\sigma J_\sigma - \mathcal{H}J_\sigma).$$

So it is sufficient to prove that $\mathcal{B}_\sigma J_\sigma - \mathcal{H}J_\sigma$ tends to 0 when $\sigma \rightarrow \infty$. First notice that there exists a constant M_J such that for all σ , $|J_\sigma| \leq M_J$; indeed, the values of J_σ belong to the interval $(\min_{x \in \partial\Omega} (C(x)), \max_{x \in \partial\Omega} (C(x)))$ (in the very model discussed in the paper this interval is $(C(\mathcal{G}), C(\mathcal{O})) = (-1, 1)$). Let $N(x, y)$ be the neighborhood of (x, y) on the grid: $N(x, y) = \{(x + \delta, y), (x - \delta, y), (x, y + \delta), (x, y - \delta)\}$. We note in the sequel $\sum_{(x', y') \in N(x, y)}$ as $\sum_N(x, y)$, then for all $(x, y) \in \Omega$

$$\begin{aligned}
&|\mathcal{B}_\sigma J_\sigma(x, y) - \mathcal{H}J_\sigma(x, y)| \\
&= \left| \gamma^{\tau(\theta_{x,y}^{J_\sigma})} \sum_{N(x,y)} p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] J_\sigma(x', y') \right. \\
&\quad \left. - \frac{1}{4} \sum_{N(x,y)} J_\sigma(x', y') \right| \\
&= \left| (\gamma^{\tau(\theta_{x,y}^{J_\sigma})} - 1) \sum_{N(x,y)} p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] J_\sigma(x', y') \right. \\
&\quad \left. + \sum_{N(x,y)} \left(p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] - \frac{1}{4} \right) J_\sigma(x', y') \right| \\
&\leq |\gamma^{\tau(\theta_{x,y}^{J_\sigma})} - 1| \left| \sum_{N(x,y)} p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] J_\sigma(x', y') \right| \\
&\quad + \left| \sum_{N(x,y)} \left(p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] - \frac{1}{4} \right) J_\sigma(x', y') \right| \\
&\leq |\gamma^{\tau(\theta_{x,y}^{J_\sigma})} - 1| M_J + \sum_{N(x,y)} \left| p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] - \frac{1}{4} \right| M_J
\end{aligned}$$

The result comes from the fact that when $\sigma \rightarrow \infty$, $|\gamma^{\tau(\theta_{x,y}^{J_\sigma})} - 1|$ and $p_{\theta_{x,y}^{J_\sigma}} [(x, y), (x', y')] - \frac{1}{4}$ both tend to 0.

REFERENCES

- [1] C. I. Connolly, J. B. Burns, and R. Weiss, "Path planning using laplace's equation," in *Proceedings of the 1990 IEEE International Conference on Robotics and Automation*. IEEE, May 1990, pp. 2102–2106.
- [2] S. Akishita, S. Kawamura, and K. Hayashi, "Laplace potential for moving obstacle avoidance and approach of a mobile robot," in *Japan-USA Symposium on Flexible Automation, A Pacific Rim Conference*, 1990, pp. 139–142.
- [3] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," *International Journal of Robotic Research*, vol. 5, no. 1, pp. 90–98, 1986.
- [4] A. M. S.A. Masoud, "Motion planning in the presence of directional and obstacle avoidance constraints using nonlinear anisotropic, harmonic potential fields: A physical metaphor," *IEEE Transactions on Systems, Man, & Cybernetics, Part A: systems and humans*, vol. 32, no. 6, pp. 705–723, November 2002.
- [5] J. Zelek, "Complete real-time path planning during sensor-based discovery," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1998.

- [6] D. Alvarez, J. Alvarez, and R. Gonzalez, "Online motion planning using laplace potential fields," in *Proceedings of the IEEE International Conference Robotics and Automation*, 2003., vol. 3, 2003, pp. 3347–3352.
- [7] H. Feder and J. Slotine, "Real-time path planning using harmonic potentials in dynamic environments," in *Proceedings of the IEEE International Conference Robotics and Automation*, 1997, vol. 1, 1997, pp. 874–881.
- [8] M. Huber, W. MacDonald, and R. Grupen, "A control basis for multi-legged walking," in *Proceedings of the IEEE International Conference Robotics and Automation*, 1996, vol. 4, 1996, pp. 2988–2993.
- [9] M. Kazemi, M. Mehrandezh, and K. Gupta, "An incremental harmonic function-based probabilistic roadmap approach to robot path planning," in *Proceedings of the IEEE International Conference Robotics and Automation*, 2005, 2005, pp. 2136–2141.
- [10] J. Sweeney, H. Li, R. Grupen, and K. Ramamritham, "Scalability and schedulability in large, coordinated, distributed robot systems," in *Proceedings of the IEEE International Conference Robotics and Automation*, 2003, vol. 3, 2003, pp. 4074–4079.
- [11] Y. Wang and G. Chirikjian, "A new potential field method for robot path planning," in *Proceedings of the IEEE International Conference Robotics and Automation*, 2000, vol. 2, 2000, pp. 977–982.
- [12] R. S. Sutton and A. G. Barto, *Reinforcement Learning : An Introduction*. MIT Press, 1998.
- [13] R. E. Bellman, *Dynamic Programming*. Princeton University Press, 1957.
- [14] S. P. Singh, A. G. Barto, R. A. Grupen, and C. I. Connolly, "Robust reinforcement learning in motion planning," in *Advances in Neural Information Processing Systems Conference*, 1994, pp. 655–662.
- [15] J. A. Coelho Jr., R. Sitaraman, and R. A. Grupen, "Parallel optimization of motion controllers via policy iteration," in *Advances in Neural Information Processing Systems*, vol. 8, 1996, pp. 996–1002. [Online]. Available: citeseer.ist.psu.edu/coelho96parallel.html
- [16] T. Perkins and A. Barto, "Lyapunov-constrained action sets for reinforcement learning," in *Proc. 18th International Conf. on Machine Learning*. Morgan Kaufmann, San Francisco, CA, 2001, pp. 409–416. [Online]. Available: citeseer.ist.psu.edu/perkins01lyapunovconstrained.html
- [17] C. I. Connolly and R. A. Grupen, "Nonholonomic path planning using harmonic functions," Computer Science Department, University of Massachusetts, Tech. Rep. 94-50, May 1994.
- [18] C. Connolly, "Harmonic functions and collision probabilities," *International Journal of Robotics Research*, vol. 16, no. 4, pp. 497–507, Aug 1997.
- [19] N. Krylov, *Controlled Diffusion Processes*. New York: Springer-Verlag, 1980.
- [20] R. Munos, "A study of reinforcement learning in the continuous case by the means of viscosity solutions," *Machine Learning*, vol. 40, no. 3, pp. 265–299, 2000.
- [21] H. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer Verlag, 1992.
- [22] L. Evans, *Partial Differential Equations*. AMS, 1991.
- [23] R. Munos and P. Bourgin, "Reinforcement learning for continuous stochastic control problems," in *Advances in Neural Information Processing Systems Conference*, 1997.
- [24] C. Connolly and R. Grupen, "On the applications of harmonic functions to robotics," *Journal of Robotic and Systems*, vol. 10, no. 7, pp. 931–946, October 1993.
- [25] L. Trassenko and A. Blake, "Analogue computation of collision-free paths," in *International Conference on Robotics and Automation*. IEEE, April 1991, pp. 540–545.
- [26] C. Connolly and J. Burns, "The basal ganglia and the planning of actions," COINS Department, University of Massachusetts, Tech. Rep. 92-18, Feb 1992.
- [27] D. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*. Prentice hall, 1989.
- [28] D. Bertsekas, "Distributed asynchronous computation of fixed points," *Math. Programming*, vol. 27, pp. 107–120, 1983.
- [29] B. Scherrer, "Asynchronous neurocomputing for optimal control and reinforcement learning with large state space," *Neurocomputing*, vol. 63, pp. 229–251, 2005.
- [30] T. Viéville and C. Vadot, "An improved biologically plausible trajectory generator," INRIA, Tech. Rep. RR-4539, September 2002.