Robot Planning and Control Via Potential Functions

A paper ~~submitted to~~ to appear in the
International Journal of Robotics Review

Daniel E. Koditschek [1]

Center for Systems Science
Yale University, Department of Electrical Engineering

June 29, 1988

[1] This work was supported in part by the National Science Foundation under grant no. DMC-8505160

1 Introduction

There mingle in the contemporary field of robotics a great many disparate currents of thought from a large variety of disciplines. Nevertheless, a largely unspoken understanding seems to prevail in the field to the effect that certain topics are conceptually distinct. In general, methods of task planning are held to be unrelated to methods of control. The former belong to the realm of geometry and logic whereas the latter inhabit the the earthier domain of engineering analysis; geometry is usually associated with off-line computation whereas everyone knows that control must be accomplished in real-time; the one is a "high level" activity whereas the other is at a "low level". This article concerns one circle of ideas that, in contrast, intrinsically binds action and intention together in the description of the robot's task. From the perspective of task planning, this point of view seeks to represent abstract goals via a geometric formalism which is guaranteed to furnish a correct control law as well. From the perspective of control theory, the methodology substitutes reference dynamical systems for reference trajectories. From the point of view of computation, less is required off-line, while more is demanded of the real-time controller. From the historical perspective, these techniques represent the effort of engineers to avail themselves of natural physical phenomena in the sythesis of unnatural machines.

This article reviews the historical context and some of the more important contemporary practitioners of the potential field based methodology for robot planning and control. In Section 2, the example of PD feedback for a one degree of freedom mechanical system is used to illlustrate the basic idea in a trivial context. Namely, we interpret the proportional gain as representative of a one dimensional "planning system" based upon the geometry of cost functions. The derivative gain allows us to "embed" the limiting behavior of the planning system in the two dimensional physical plant. In Section 3 we sketch the history of the control theoretic aspect of this methodology — the progress of total energy conceived as a Lyapunov function — focusing upon the important work of Arimoto and colleagues. Finally, in Section 4 we suggest the planning capabilities of this point of view by focusing on the pioneering work of Khatib in the domain of obstacle avoidance, and touching upon Hogan's re-interpretation in the domain of tasks requiring contact with the environment.

2 Example: An End-Point Task for a Simple Robot

Consider the "one-degree-of-freedom" robot — a single (revolute or prismatic) joint whose position, q, and velocity, $\dot{q}$, are measured exactly and instantaneously by a perfect sensor, actuated by a motor which delivers exact torque, τ, according to the user's instantaneous command — and its dynamical equation, given by Newton's second law as

$$M\ddot{q} = \tau, \tag{1}$$

where M is the mass (or moment of inertia in the revolute case) of the robot link. Suppose the robot has been given the task of moving to a point, q^*, and remaining there. One might imagine splitting the task up into a "high level" geometric problem — find a curve in the jointspace, c(t), which ends up at q^* — and a "low level" control problem — find a contol law, $\tau(t)$, which "forces" the robot to "track" the commanded behavior, $q(t) \to c(t)$. In much of the robotics literature, these two problems are solved independently. In the methodology under consideration, they are solved at the same time.

2.1 Geometry: Hill Climbing

The geometric aspect of this task may be represented by the following optimization problem. Conceive of a "cost" function on the jointspace, φ, which assigns a scalar value to every position, vanishing uniquely at the "target", $\varphi(q^*) = 0$, and growing larger farther away. For example, the quadratic function,

$$\varphi_{HL} \triangleq \frac{1}{2} K_P \left[q - q^*\right]^2,$$

would do nicely for all positive values of the scalar K_P. If φ is continuously differentiable then it has a well defined negative gradient system,

$$\dot{q} = -grad\, \varphi\,. \tag{2}$$

Solutions of this differential equation follow the "fall line" of the "hill" defined by the cost function — i.e. at every position, the velocity of any solution curve is specified by the directional derivative of φ. If q^* is a local minimum of φ it follows that all solutions of system (2) which originate in some neighborhood of that point, define curves which lead to that point. If, in addition, q^* is the *only* extremum of φ, then *every* solution curve leads there. For example, in the case of the quadratic cost function, φ_{HL}, the gradient system works out to be the scalar linear time invariant differential equation,

$$\dot{q} = -K_P \left[q - q^*\right].$$

Since q^* is a minimum, and the only extremum of φ_{HL} to boot, it follows that this gradient system generates a solution to the geometric "find path" problem from *any* starting position.

2.2 Control: Energy Dissipation

Faced with the particular problem of navigating from some initial position, q_0, to the target, q^*, one might now define a reference trajectory, $c(t)$, by solving the gradient differential equation, (2), for the initial condition, q_0, and then attempt a tracking control. Instead, we will re-interpret the cost function, φ, as a *potential function*, and introduce a control law which achieves the desired result with no recourse to explicit solutions of the original gradient system, (2).

If we are to interpret φ as a potential function, we must form the total energy by taking its sum with the kinetic energy,

$$\eta \triangleq \frac{1}{2}\dot{q}^{\mathrm{T}} M \dot{q} + \varphi,$$

and then apply the Lagrangian formalism,

$$\left[\frac{d}{dt}D_{\dot{q}} - D_q\right]\eta = f_{ext},$$

(where f_{ext} denotes all non-conservative forces) to obtain the Newtonian law of motion,

$$M\ddot{q} + grad\,\varphi = f_{ext}.$$

If f_{ext} represents the effect of a dissipative force, say a Rayleigh damper, $f_{ext} = -K_D\dot{q}$, where K_D is positive, then the total energy must decrease:

$$\dot{\eta} = -K_D\dot{q}^2.$$

In consequence, it seems intuitively plausible, and will be made rigorously clear below, that $(q, \dot{q})$ converges to $(q^*, 0)$ from some neighborhood of that point in phase space — the space of positions and velocities — as long as q^* is a minimum of φ.

The final equation of motion resulting from this formulation is

$$M\ddot{q} + K_D\dot{q} + grad\,\varphi = 0. \tag{3}$$

It is clear that (1) may be made to look like (3) by assignment of the control law,

$$\tau \triangleq -K_D\dot{q} - grad\,\varphi\,.$$

Thus, if q^* is a local minimum of φ, then, under the influence of this control law, our robot is guaranteed to approach $(q^*, 0)$ asymptotically from any initial state, $(q_0, \dot{q}_0)$, which is sufficiently close to $(q^*, 0)$. In the particular case of a Hook's Law spring potential, φ_{HL}, this control strategy corresponds exactly to the time honored "proportional-derivative" feedback control strategy of linear systems theory,

$$\tau \triangleq -K_D\dot{q} - K_P\,[q - q^*]\,,$$

with the familiar closed loop dynamics,

$$M\ddot{q} + K_D\dot{q} + K_P\,[q - q^*] = 0.$$

3 Control: Dissipative Mechanical Systems

The modern history of this idea might be said to begin with the discovery by the American engineering community of the work of Lyapunov [24]. The specific application to the class of mechanical systems seems to have been rediscovered on several different occasions by a variety of researchers in a variety of subdisciplines: Section 3.1 offers an illustrative (and certainly not exhaustive) sketch of this development. [1] Credit for first introducing these results into the robotics literature is due Arimoto and colleagues: Section 3.2 presents an overview of their central contributions.

3.1 Lyapunov Theory and Total Energy

Two hundred years ago, Lagrange showed that a conservative physical system has stable behavior with respect to any minimum of its potential energy. Roughly one century later, Lord Kelvin argued that the addition of a dissipative field would induce asymptotic stability of any minimum state. Since our understanding of the relation between gradient planning systems and robot dynamics stems entirely from these ideas, it is worth spending a little time to sketch a brief history of their subsequent development and discovery by engineers.

3.1.1 Lyapunov's Direct Method

In his doctoral dissertation of 1892, Lyapunov unified the asymptotic analysis of linear systems based upon eigenvalues with the energy-based analysis of mechanical (generally speaking, nonlinear) systems discovered by Lagrange and Lord Kelvin. The so-called "direct method" provides a remarkably simple test for stability of an equilibrium state, x^*, of a general dynamical system,

$$\dot{x} = f(x), \qquad f(x^*) = 0. \tag{4}$$

Namely, one studies the derivative of a scalar valued function, $V(x)$, along trajectories of (4) in the neighborhood of x^*. Notice that

$$\dot{V}[x(t)] = grad\,V \;\cdot \dot{x} = grad\,V \;\cdot f(x),$$

hence no explicit knowledge of the trajectories of (4) (which are, in general, unobtainable for nonlinear systems) is required in order to apply the theory. Lyapunov showed that if V is ***positive definite*** at x^* — that is, takes positive values on a neighborhood, vanishing only at that point — and if $-\dot{V}$ is positive definite as well, then x^* is asymptotically stable.

While this powerful technique was developed and refined by a variety of Soviet mathematicians over the following five decades [5,25,?], it remained virtually ignored in the west. In the nineteen fifties, Lyapunov's method gained the attention of the American mathematician

[1] Of course, the central idea of energy dissipation is to be found in physics texts as well [1][Prop. 3.7.17].

of ϵ relative to α as long as R, P are positive definite square matrices. This last theorem is used in conjunction with the Lyapunov arguments to determine the asymptotic behavior of his connected systems — considered as linked rigid bodies by passing to the limiting case of increasingly large potential forces — in the presence of passive dampers.

3.2 Arimoto's Use of Total Energy for Robot Systems

The first clear exposition of the use of total energy in robotics was contributed by Arimoto. Over the better part of a decade, he and his colleagues have offered numerous extensions and refinements of these ideas. An increasing number of researchers (this author among them) have had the experience of reporting a "new" result in the application of Lyapunov theory to robotics and subsequently discovering the same or similar ideas in papers written several years earlier by Arimoto. This section attempts to present an overview of his work in the area by concentrating on two important papers.

3.2.1 The Central Idea

In constrast to the situation two decades ago treated by Pringle who analyzed passive arrangements of dampers and springs, the advance of technology in robotics lends the designer complete freedom (at least in theory) to choose an arbitrary force law at each degree of freedom. Arimoto's contribution was to argue that computationally very simple force laws based upon the principles of energy dissipation discussed so far could accomplish rather sophisticated end-point tasks. The first exposition of this point of view seems to be his 1981 paper with Takegaki [36].

The authors observe that existing potential fields may be exactly cancelled and replaced by artificial fields in a manipulator whose joints are all actuated by user commanded torques. They next introduce a Rayleigh damping term and argue, as in our simple example of Section 2.2, that the total energy — the robot's physical kinetic energy added to the artificial potential function — must decrease. By invoking LaSalle's invariance principle, they obtain their desired asymptotic stability result. Since the Lagrangian formulation of dynamics obtains from variational principles, the authors are able to assert the optimality of the Rayleigh damping scheme with respect to a quadratic performance index based upon it. They next choose a quadratic artificial potential function: since the stability result requires only that the potential function be positive definite at the desired rest point, they point out that simple decoupled PD feedback will suffice.

The paper now shifts to applications involving workspace based planning. A quadratic artificial potential function in work space leads to a feedback law whose proportional term involves the workspace errors multiplied by the transposed jacobian of the forward kinematics. A certain amount of unnecessary confusion is added by sticking with the Hamiltonian rather than Lagrangian formulation of the dynamics: the former requires analysis involving the inverse jacobian whereas the latter would not. Holonomic constraints in the workspace are handled by

Solomon Lefschetz [23]. His colleague, Lasalle, discovered a fundamental relationship between Lyapunov functions and limit sets of dynamical systems which proves to be of primary importance for mechanical systems.

A glance back at the time derivative, $\dot{\eta}$, in Section 2.2 shows that the total energy is not a *strict* Lyapunov function: this means that its derivative vanishes (in this case, at every zero velocity state) even though the trajectory has not arrived at the desired equilibrium state. In general, if V is positive definite and $-\dot{V}$ is merely non-negative, then one may deduce stability (solutions originating near the equilibrium, x^* remain near that point) but not necessarily asymptotic stability (nearby solutions not only remain near, but eventually end up at x^*). Thus, the total energy appears to have a significant flaw for purposes of deducing asymptotic attraction to a nearby minimum of the potential energy: indeed, Lord Kelvin's original argument [38][§345] ignores this flaw. In his simple but illuminating book with Lefschetz, Lasalle [22] presented his "Invariance Principle", which demonstrated in the context described above that all solutions eventually end up in an *invariant* subset of the set $\dot{V} \equiv 0$. An invariant set of (4) is characterized by the property that all solutions originating there, remain there for all time. Thus, with a little more effort, Lyapunov theory may be used to make rigorous the intuitive arguments of Lord Kelvin.

3.1.2 Application to Problems in Mechanical Engineering

In conjunction with the advocacy of these ideas by Lefschetz came the the awareness of their relevance to engineering problems. The weakness of Lyapunov's method is that it provides information only if a "candidate function", V, is available: the construction of such functions remains something of a black art in general. In particular, however, when the dynamical system (4) has a specific structure, then it is often possible to say a great deal. This is the situation for linear systems — Lyapunov had already proven the existence of quadratic functions, V, in that case — as well as for mechanical systems where, as we have seen, the total energy provides the obvious choice. Kalman and Bertram published an influential article in 1960 describing the relevance of Lyapunov functions for control applications [11], leading to an avalanche of interest of which the present application are represents just one stream. Credit for first making use of total energy as a Lyapunov function for a specific engineering application involving mutually constrained bodies would appear to be due Pringle.

In his 1966 paper [29], Pringle first presents a succint historical review and then defines his notion of a "connected system." This is a single rigid body together with a set of point masses interconnected via holonomic constraints. There are three theorems: the first expresses Lyapunov's direct method; the second presents the application to mechanical systems taken from the Soviet literature cited above, together with the appropriate appeal to Lasalle's invariance principle; the final theorem states that the block matrix

$$\begin{bmatrix} R & \frac{1}{2}\epsilon C \\ \frac{1}{2}\epsilon C^{\mathrm{T}} & \alpha P \end{bmatrix}$$

may be made positive definite, regardless of the properties of C, by a sufficiently small choice a local technique based upon the implicit function theorem which yields tangent information to the jointspace based torque inputs. This necessitates a different feedback structure for each region of the task giving rise to the use of time varying potentials for which no stability proof is offered.

Subsequent to the publication of Arimoto's first paper, independent derivations of the same result were reported by Van der Schaft [32] and this author [20]. The latter development proceeded from the Lagrangian rather than Hamiltonian formulation of robot arm dynamics, thus the author was led to discover a structural aspect of the Coriolis terms which had not been reported in Arimoto's work. This "skew-symmetric" structure was later clarified and exploited by Slotine and Li [33] in the construction of the first globally stable adaptive controller for robot arms.

3.2.2 Extensions and Problems

The important 1984 paper by Arimoto and Miyazaki [35], receives an independent review in the same issue of this journal. For the purposes of the present article, that paper provides a convenient point of departure for discussion of extensions and open problems in this domain. In the effort to gain a stability result for the traditional PID control algorithm from linear systems theory, the authors are forced to confront and "fix" the weakness of Lord Kelvin's original result: namely, the lack of a negative definite derivative which necessitates appeal to LaSalle's Invariance Principle.

They modify the total energy function

$$V_E(x) = x^{\mathrm{T}} \begin{bmatrix} K_P & 0 \\ 0 & M \end{bmatrix} x$$

by adding a cross term which is bilinear in position and velocity,

$$V_{Ari}(x) = x^{\mathrm{T}} \begin{bmatrix} K_P & \alpha M \\ \alpha M & M \end{bmatrix} x$$

resulting a locally strict Lyapunov Function,

$$\begin{aligned} \dot{V} &= x^{\mathrm{T}} \begin{bmatrix} K_P & \alpha M \\ \alpha M & M \end{bmatrix} \begin{bmatrix} 0 & I \\ -M^{-1}K_P & -M^{-1}K_D \end{bmatrix} x \\ &= -x^{\mathrm{T}} \begin{bmatrix} \alpha K_P & \frac{1}{2}\alpha K_D \\ \frac{1}{2}\alpha K_D & K_D - \alpha M \end{bmatrix} x. \end{aligned}$$

For the purposes of this illustrative linear example, both V and $\dot{V}$ may be made sign definite via small enough choice of the parameter, α: a global stability argument results. In the general situation, the presence of the Coriolis forces results in a cubic term in $\dot{V}$ (it is linear in the position error and quadratic in the velocity error) which, though necessarily sign indefinite, is dominated by $K_D - \alpha M$ for small enough position error magnitudes. Thus V_{Ari} yields only a local stability argument: the designer must insure that the initial conditions originate in a neighborhood of the desired state small enough to guarantee that decreasing V implies the error remains less than a magnitude proportional to $\|K_D - \alpha M\|$. The alternative — enlarging the domain of attraction — may be accomplished only by recourse to larger derivative gains, K_D.

Tracking. The notion of tracking appears to be out of context in a review article concerning robot task encoding via reference dynamics. Actually, the use of a "moving Hook's Law Potential" — that is, a PD-based tracking method — serves as the first instance of generalized moving potential fields. For example, it would be very nice to use the obstacle avoidance schemes discussed below to navigate amidst moving environments. Two years ago, this author [19] and Bayard and Wen [39] announced independent work leading to the proof of exponential stability of PD compensated nonlinear mechanical systems. This construction afforded a proof of global boundedness of PD-based tracking schemes, and reasonable estimates of the rate of convergence to the bound. Both results were based upon a modified version of total energy which yielded a locally strict Lyapunov function. At the following year's IEEE Conference on Robotics and Automation, Arimoto quietly pointed out that he had obtained the same result several years earlier [3]. The three independent papers offer strikingly similar variants of the total energy — all with the same local limitations discussed above. In fact, they are all similar to the original modification Arimoto had introduced in 1984.

Very recently, this author has developed a further modification of the total energy which yields a strict Lyapunov function whose extent is global[15]. Moreover, unlike the earlier modifications, which assume a Hook's Law Spring Potential, the new variant is defined for a large class of potential laws. This opens the desirable possibility of combining time variation with the complex artificial potentials to be discussed in the last Section of this article.

Internal Model Principle. Arimoto's extension of Lord Kelvin's ideas in the 1984 paper expand the analogy between linear second order systems and nonlinear mechanical systems from PD to PID techniques. In classical linear systems theory, the introduction of integral feedback action is associated with tracking a reference (or rejecting a disturbance) which is known structurally, but only up to an unknown set of parameters. For example, the addition of an integrator is dictatce the presence of a step input; two cascaded integrators are required in the presence of a ramp input; and so on. These ideas from the classical theory were formalized by Wonham and colleagues as "The Internal Model Principle" [8].

Since the most popular reference trajectories in the robotics literature are generated by polynomial spline techniques — that is, they may be modeled as outputs of unforced linear time invariant dynamical reference systems whose dimension is specified by the degree of the polynomial — the internal model principle is immediately applicable to Cartesian robot arms. Its extension to nonlinear mechanical systems would be of great significance for practical robot control.

Transient Analysis. Even though the desired asympotic behavior is assured, it has been pointed out on many occasions, originally in Arimoto's 1981 paper [36], that the transients resulting from potential field control of nonlinear mechanical systems may be poorly behaved. More generally, given a kinetic energy, it seems very important to develop a theory that matches the dissipative term to the particular form of the potential, φ, in order to assure that the closed loop mechanical system evince a "sufficiently accurate" copy of the gradient "planning system" trajectories. In the simple example of Section 2 we already have an intuitive idea of what this means. The "best" match between (2) and (3) arguably obtains from the assignment of gains for the latter which produces a critically damped system whose natural frequency is the time constant of the former. It seems important to generalize this notion in the nonlinear context.

the level lines of the original "objective" function [9]. Thus, a careful construction methodology for potential functions brings the promise of good transient shaping in the planning system. Of course, our true interest is in an on-line implementation of the gradient algorithm implicit in the motion of the robot arm itself. Work of this author has established the global similarity in limiting behavior between the planning system and the ultimate closed loop mechanical system [20,15]. Yet, as mentioned at the end of the previous section, it remains to determine the relationship between the transients.

4.2 Khatib's Method of Obstacle Avoidance

There are certain ideas which are so natural and compelling that it becomes hard to distinguish a particular individual responsible for their introduction to a community researchers. Happily, in the case of artificial potential fields — surely the exemplar of such a compelling idea — it seems clear that credit for its introduction to the robotics community is due Khatib.

4.2.1 FIRAS

In his original 1978 paper with Le Maitre [12], Khatib suggested the desirability of bringing more "adaptibility" to the control level, while maintaining a relatively lean (hence, rapidly executable), yet dynamically meaningful maneuvering algorithm, for robots in cluttered environments. He proposed adopting the general paradigm of Renaud, who had suggested the relationship between cost function gradients and mechanical systems in a doctoral dissertation two years earlier [?], to the specific problem of obstacle avoidance.

Assuming that each obstacle is described as the zero level surface of a known scalar valued analytic function, $f(x, y, z) = 0$, Khatib formed a local inverse square potential law: it goes to infinity as the inverse square of f near the obstacle, and gets cut off at zero at some positive level surface, $f(x, y, z) = f_0$, presumably "far enough" away from the obstacle, as determined by the designer. A particle moving according to Newton's Laws in such a potential field would clearly never hit this obstacle. Khatib further observed that the sum of the gradients is the gradient of the sum: thus adding up the potential laws for many obstacles would result in a single function under whose influence the particle could not hit any obstacle.

Since the objects to be maneuvered are not point masses, but rigid links of a kinematic chain, Khatib proposed identifying a number of distinguished points on the manipulator body, and subjecting each of them to the potential field of the obstacles. This induces a potential function on the joint space: the techniques of stable hill climbing via Lagrangian dynamics described above are immediately applicable. To complete the methodology, Khatib suggested adding obstacle potentials representing the joint limits to the induced workspace function. Thus was born FIRAS — "force inducing an artificial repulsion from the surface" — in 1978.

4 Planning: Gradient Dynamical Systems

The idea of using "potential functions" for the specification of robot tasks was pioneered by Khatib [?] in the context of obstacle avoidance, and further advanced by fundamental work of Hogan [10] in the context of force control. The methodology was developed independently by Arimoto in Japan [26], and by Soviet investigators as well [28]. Of course, the possibility of solving complex problems by resort to analog methods of computation has a much older history which we touch upon briefly in Section ?? before turning to Khatib's contributions in Section ??.

4.1 Hill Climbing via Analog Computation

Possibly the most common use of gradient vector fields in engineering occurs in the context of root finding: given a function, $g(x)$, find those points which attain a particular value. For example, in robotics, g might likely represent the kinematic transformations. Recent work of Smale [34] and Hirsch and Smale [?] has established variants of the Newton-Raphson technique which succeed for almost all initial points, in the search space. They point out that the gradient algorithm — climbing down the "error hill" — is a much older variant with poorer properties. Yet recourse to hill climbing has two other advantages that make it so attractive in the present context. First, peaks as well as valleys may be encoded: this idea leads to the notion of repelling sets of obstacles. Second, as we have seen in the previous section, the dissipative mechanical systems make natural analog computers for integrating gradient vector fields.

Of course, the appeal to analog computation has an old and established history. Even after the domination of electronic computation by digital technology in the sixties, serious engineering effort has been expended upon analog technology for special purpose problems of data acquisition and manipulation, e.g. for pulse code modulation devices [4], or image sensory processing [7,6]. More general classes of combinatorially complex problems have been approached from this point of view in recent years. Work of Hopfield and Tank [37] has demonstrated the possibility of finding good sub-optimal solutions to the travelling salesman problem (in the class np-complete) by formulating them in terms of a scalar potential function and implementing their solution using "neural nets" consisting of many processing elements endowed with elementary analog capability. Kirkpatrick and colleagues [14] have generated a great deal of interest in a class of objective functions denoted "simulated annealing" (because of their origin in models of physical cooling processes) used, again, to find good sub-optimal solutions to the problem of VLSI circuit device placement (also np-complete). It should be mentioned that the gradient method has a long history in systems science as well as the general world of applied mathematics. For instance, the parameter adjustment algorithms of adaptive control and estimation schemes may be seen as a gradient solution to a set of linear algebraic equations [27].

Unlike most nonlinear differential equations, gradient vector fields have particularly simple dynamical behavior. Systems possessed of isolated equilibrium states come with the guarantee that any bounded solution must asymptotically approach one of them travelling orthogonal to

4.2.2 Problems and Extensions

In the decade since its introduction, the idea of using artificial potential functions for robot task description and control has been adopted or re-introduced independently by a growing number of researchers [26,2,28]. There are several conceptual problems with the potential function methodology, of which the more important include the following.

Terrain Shaping Khatib [13] observed that the use of analytic level surfaces to prescribe obstacle boundaries and neighboring regions becomes difficult when the surface is complex because the relative locus of the concentric level surfaces may be poorly behaved. His alternative — reliance upon distance functions computed online — represents an interesting point of departure but would surely give rise to equally difficult problems in an effort to obtain theoretical understanding of the properties of such trajectories given the documented trickiness of set distances — see the work of Gilbert and Johnson [?], for example.

It seems equally plausible to imagine a more careful examination of the construction of implicit representations of solid body surfaces via scalar valued functions. Of course, the study of representations of solid bodies is a vast discipline of its own. While those researchers seem to favor explicit parametrizations rather than implicit representations [30], it seems certain that increased contact between the two communities would result in enhanced understanding for both. In particular, a careful listing of desiderata for such functions — both with respect to their use in robot navigation as well as in computer models of solid bodies — would seem to be essential before they are rejected out of hand.

Rigid Bodies While it is most natural to construct obstacle avoidance potentials on the workspace, existing methods are actually more suited to configuration space where the exact state of the robot at any instant is represented by a single point. As discussed above, work space boundaries (obstacle surfaces) are plausibly available from a CAD/CAM model: their "inverse images" under the kinematic transformation representing any particular robot are not. Thus practitioners of the methodology must rely upon heuristics. A creditable example is proposed by Khatib [13] who places his "points subjected to potential" forces on a specific locus — a line through the body — of each link in the kinematic chain. While this seems like a reasonable procedure, and the designer is guaranteed that the particular points in question will not collide with obstacles, there is no such guarantee for the entire link itself. Thus, there is a crying need for provably correct procedures which build obstacle avoidance potentials on the configuration space of rigid bodies rather than simply point robots.

Spurious Minima Even given a good description of the configuration space, a central difficulty w ith this technique has been the possibility of extra minima. As has been asserted, gradient vector fields have particularly simple dynamical behavior. Systems possessed of isolated equilibrium states come with the guarantee that any bounded solution must asymptotically approach one of them [9]. However, there is no generally applicable method for determining the

5 Conclusion

This article has attempted a brief review of the potential function methodology for robot task encoding and control. The author hopes that it may simultaneously serve the purpose of an elementary introduction to this point of view, while highlighting the significant contributions of three researchers: Arimoto, Hogan, and Khatib. A more elaborate treatment of these ideas may be found in the author's recent tutorial articles [21,17].

particular equilibrium state achieved from a particular initial condition. All points in some open neighborhood of a local minimum will approach that point, whereas the set of points which approach a saddle has measure zero, and the set of points which approach a local maximum is empty. Thus, if the problem formulation gives rise to an objective function with more than one local minimum a sub-optimal solution is guaranteed for all initial conditions in some set of non-zero measure whose location is generally impossible to characterize analytically.

This has prompted the author to spend a growing amount of attention focused on the problem of how to construct "good" potential functions on manifolds with boundary [18]. We have been led to define the notion of a *navigation function* [31] — a refinement of the potential function which includes the condition of a single minimum (at the desired destination), a uniform finite height on all the boundaries, along with several other technical requirements — and have constructed them on a variety of spaces. Thirty year old topological results of S. Smale assure the existence of smooth navigation functions on any smooth space with boundaries [16]: it is now up to the engineering community to actually construct them.

4.3 Hogan's Impedance Functions

It seems useful to include in this article a brief look at some work which demonstrates that the potential function approach to unified task description and control is not limited to task domains involving a purely geometric environment. Fundamental work by Hogan [10] advances persuasive arguments for encoding general manipulation tasks in the form of "impedances". Impedances and admittances are formal relationships between the force exerted on the world at some cartesian position and the motion variables - displacement, velocity, acceleration, etc. - at that position with respect to some reference point (or "virtual position" in Hogan's terminology [10]). He argues that for purposes of modeling manipulation tasks, the kinematic and dynamical properties of a robots's contacted environment must be understood as admittances - systems for which the relationship operates as a function describing a specified displacement for any input force. Arguing, further, that physical systems may only be coupled via port relationships which match admittances to impedances, and that robots can violate physics no more than any other objects with mass, he arrives at the conclusion that the most general model of manipulation is the specification of an impedance - a system which returns force as a function of motion. By construing motions relative to a virtual position as defining tangent vectors at that position, Hogan notes that an impedance may be defined in terms of a scalar valued function on the cross product of two copies of the tangent space at each virtual position whose gradient co-vector determines the relationship between motion and force. Thus, an impedance may be re-interpreted as the gradient co-vector field of an "objective function", whose fall lines specify the desired dynamical response of the robot end-effector in response to infinitesimal motions imposed by the world. In this context, unlike the other gradient vector field task definitions, it is intended a priori that the dynamics be second order - i.e. define changes of velocity (force) rather than changes of position.

References

[1] Ralph Abraham and Jerrold E. Marsden. *Foundations of Mechanics.* Benjamin/Cummings, Reading, MA, 1978.

[2] J. R. Andrews and N. Hogan. Impedance control as a framework for implementing obstacle avoidance in a manipulator. In David E. Hardt and Wayne J. Book, editors, *Control of Manufacturing and Robotic Systems*, pages 243–251, A.S.M.E., Boston, MA, 1983.

[3] S. Arimoto and F. Miyazaki. Asymptotic stability of feedback controls for robot manipulators. In *Proceedings 1st IFAC Symposium on robot Control*, Barcelona, Spain, 1985.

[4] K. W. Cattermole. *Principles of Pulse Code Modulation.* American Elsevier, NY, 1969.

[5] N. G. Chetaev. *The Stability of Motion.* Pergammon, New York, 1961.

[6] E. R. Fossum. *Charge-Coupled Analog Computer Elements and Their Application to Smart Image Sensors.* PhD thesis, Yale University, New Haven, Ct., 1984.

[7] Eric R. Fossum and Richard C. Barker. A linear and compact charge-coupled charge packet differencer/replicator. *IEEE Transactions on Electron Devices*, ED-31(12):1784–1789, Dec 1984.

[8] B.A. Francis and W. M. Wonham. The internal model principle for multivariable regulators. *Appl. Math. & Optimization*, 2(2):170–194, 1975.

[9] Morris W. Hirsch and Stephen Smale. *Differential Equations, Dynamical Systems, and Linear Algebra.* Academic Press, Inc., Orlando, Fla., 1974.

[10] Neville Hogan. Impedance control: an approach to manipulation. *ASME Journal of Dynamics Systems,Measurement, and Control*, 107:1–7, Mar 1985.

[11] R.E. Kalman and J.E. Bertram. "Control systems analysis and design via the 'second method' of Lyapunov". *Journal of Basic Engineering*, :371–392, June 1960.

[12] O. Khatib and J.-F. Le Maitre. Dynamic control of manipulators operating in a complex environment. In *Proceedings Third International CISM-IFToMM Symposium*, pages 267–282, Udine, Italy, Sep 1978.

[13] Oussama Khatib. Real time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–99, Spring 1986.

[14] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. *Optimization by Simulated Annealing.* Research Report RC 9355 (#41093), IBM Thomas J. Watson Research Center, Yorktown Heights, NY, Apr 1982.

[15] D. E. Koditschek. *A Strict Global Lyapunov Function for Mechanical Systems.* Technical Report 8707, Center for Systems Science, Yale University, November 1987 .

[16] D. E. Koditschek and E. Rimon. *The Existence of Navigation Functions for Robot Obstacle Avoidance.* Technical Report 8803, Center for Systems Science, Yale University, January 1988 .

[17] Daniel E. Koditschek. Automatic planning and control of robot natural motion via feedback. In Kumpati S. Narendra, editor, *Adaptive and Learning Systems: Theory and Applications*, pages 389–402, Plenum, 1986.

[18] Daniel E. Koditschek. Exact robot navigation by means of potential functions: some topological considerations. In *IEEE International Conference on Robotics and Automation*, pages 1–6, Raleigh, NC, Mar 1987.

[19] Daniel E. Koditschek. High gain feedback and telerobotic tracking. In *Workshop on Space Telerobotics*, page , Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, Jan 1987.

[20] Daniel E. Koditschek. Natural motion for robot arms. In *IEEE Proceedings 23rd Conference on Decision and Control*, pages 733–735, Las Vegas, Dec 1984.

[21] Daniel E. Koditschek. Robot control systems. In Stuart Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 902–923, John Wiley and Sons, Inc., 1987.

[22] J. P. Lasalle and S. Lefschetz. *Stability by Lyapunov's Direct Method with Applications.* Academic, New York, 1961.

[23] Solomon Lefschetz. *Differential Equations: Geometric Theory.* Dover, NY, 1977.

[24] A. M. Lyapunov. *Problème Gén éral de la Stabilité du Mouvement.* Princeton University, Princeton, NJ, 1949.

[25] I. G. Malkin. *The Theory of Stability of Motion.* US Atomic Energy Commission AEC-TR-3352, Wash. DC, English ed. (translation).

[26] Fumio Miyazaki and S. Arimoto. Sensory feedback based on the artificial potential for robots. In *Proceedings 9th IFAC*, Budapest, Hungary, 1984.

[27] K. S. Narendra and L. S. Valavani. Stable adaptive observers and controllers. *Proceedings of the IEEE*, 64(8), 1976.

[28] V. V. Pavlov and A. N. Voronin. The method of potential functions for coding constraints of the external space in an intelligent mobile robot. *Soviet Automatic Control*, (6), 1984.

[29] Ralph Pringle, Jr. On the stability of a body with connected moving parts. *AIAA*, 4(8):1395–1404, Aug 1966.

[30] A.A.G. Requicha and H. B. Voelcker. Solid modeling: current status and reserach directions. *IEEE Computer Graphics and Applications*, 3(10):25–37, Oct 1983.

[31] E. Rimon and D. E. Koditschek. Exact robot navigation using cost functions: the case of spherical boundaries in $\mathbf{E}^n$. In *IEEE International Conference on Robotics and Automation*, page , Philadelphia, PA, Apr 1988.

[32] A. J. Van Der Schaft. *Stabilization of Hamiltonian Systems.* Memo 470, Technische Hogeschool Twente, Twente, Netherlands, Jan 1985.

[33] Jean-Jacques E. Slotine and Weiping Li. On the adaptive control of robot manipulators. In *Proc. ASME Winter Annual Meeting*, page , Anaheim, CA., Dec 1986.

[34] S. Smale. The fundamental theorem of algebra and complexity theory. *Annals of Mathematics*, 4(1):1–36, Bull. AMS.

[35] M. Takegaki and S. Arimoto. Stability and robustness of pid feedback control for robot manipulators of sensory capability. In *Robotics Research, First International Symposium*, MIT Press, 1984.

[36] Morikazu Takegaki and Suguru Arimoto. A new feedback method for dynamic control of manipulators. *ASME Journal of Dynamics Systems,Measurement, and Control*, 102:119–125, 1981.

[37] David W. Tank and John J. Hopfield. Simple "Neural Optimization Networks": an a/d converter, signal decision circuit, and a linear programming circuit. *IEEE Transactions on Circuits and Systems*, CAS-33(5):533–541, May 1986.

[38] Sir W. Thompson and P. G. Tait. *Treatise on Natural Philosophy.* University of Cambridge Press, 1886, Cambridge.

[39] John T. Wen and David S. Bayard. *Robust Control for Robotic Manipulators Part I: Non-Adaptive Case.* Technical Report 347-87-203, Jet Propulsion Laboratory, Pasadena, CA, 1987.