Toward a Path Co-Processor for Automated Vehicle Control *

Roderic A. Grupen, Computer Science
Chris I. Connolly, SRI International
Kamal X. Souccar, Computer Science
Wayne P. Burleson, Elec. and Comp. Eng

University of Massachusetts, Amherst, MA 01003
Phone (413) 545-3280, Fax (413) 545-1249, grupen@cs.umass.edu

Abstract

The paper describes an embedded controller as the basis for high performance navigation tasks. The approach solves Laplace's equation on a discrete grid representing the roadway map. The resulting harmonic potential is free of local minima and a greedy descent minimizes the probability of collision with posted obstacles. In this paper, we will (1) introduce the important properties of harmonic potential functions for vehicle control applications, (2) describe a numerical procedure to compute harmonic functions on a grid, (3) express nonholonomic constraints as connectivity in the grid, (4) formulate an energy-referenced control scheme to manage vehicle dynamics, and (5) provide estimates for the performance on a four processor DSP architecture.

1 Introduction

Automated vehicle controllers for general road navigation tasks introduce technical problems beyond the scope of most path planning techniques designed for robot control. For instance, many robot navigation algorithms exploit synchrodrive mechanisms to simplify nonholonomic constraints. Steering and drive actions are treated sequentially. Moreover, mobile robots rarely incorporate dynamic properties of the platform because they operate in a regime where dynamics are negligible. Automated controllers for passenger vehicles will have to contend with both of these independently difficult problems. For instance, in highway driving, dynamic constraints dominate (passing and emergency maneuvers, entering and exiting platoons) while urban driving conditions will depend critically on precise treatment of nonholonomic constraints (precision docking, backup maneuvers). These systems must be responsive to changing roadway conditions and capable of incremental policy modifications. Moreover, we argue that to justify general purpose vehicle controllers, we must demonstrate control throughput at least two orders of magnitude greater than their human counterparts. In this paper, we will attempt to formulate a plausible control design for incorporating nonholonomic and dynamic constraints and for achieving the required throughput in navigation systems for automated vehicle control.

Nonholonomic systems have received some attention in the recent robotics literature (for a survey, see [4]). This is a natural consequence of the interest in controlling the motion of wheeled vehicles, since rolling contact is properly treated as a nonholonomic constraint. In this paper, we show that nonholonomic constraints can be expressed by the connectivity in a grid-based map of the roadway. In this model, there are no paths in the grid that are inconsistent with nonholonomic constraints.

To select a path in the constrained grid, we are exploring artificial harmonic potentials to plan and control the trajectories of the vehicle. This approach addresses on-line performance since the result of planning is a path encoded as a *controller* rather than a geometric path. In conjunction with a total energy formulation, bounded-torque "safe" controllers can be derived as long as the artificial potential satisfies certain constraints. Vehicle dynamics can be regulated to avoid collisions by using an energy-referenced feedback compensator based on an artificial Hamiltonian function for the system. This approach manipulates the kinetic energy of the vehicle in response to the artificial potential. The energy referenced controller can be designed so that the vehicle may only approach geometric obstacles (maximum potential energy) by removing all kinetic energy — or more generally,

*This work is supported in part by the National Science Foundation under grants CDA-8922572, IRI-9116297 and IRI-9208920.

by eliminating relative approach velocity.

2 Harmonic Functions

Harmonic functions are solutions to Laplace's equation.

$$\nabla^2 \phi = \frac{d^2\phi}{dx_1^2} + \ldots + \frac{d^2\phi}{dx_n^2} = 0$$

Harmonic functions provide a basis for constructing numerical approximations of globally admissible potential functions [6].

2.1 Properties of the Harmonic Function

Harmonic functions exhibit properties which are not found in conventional path planning algorithms. If a path to the goal exists, these properties guarantee that the system will achieve the specified goal with a smooth, obstacle-avoiding trajectory.

Min-Max Property

Harmonic functions satisfy the min-max principle: *in any compact neighborhood of freespace the minimum and maximum of the function over this domain must occur on the boundary of that neighborhood.* A corollary of this property is that local minima within the solution region can not exist [9, 10].

Mean-Value Property

Up to truncation error, the value of the harmonic potential at a point in a lattice is the average of the values at the points in the (manhattan) neighborhood.

$$\begin{aligned}\nabla^2\phi &\approx \phi(x+1,y) + \phi(x-1,y) + \\ &\quad \phi(x,y+1) + \phi(x,y-1) - 4\phi(x) \\ &= 0\end{aligned}$$

This applies in both the discrete and continuous cases and for any number of dimensions [1]. To compute harmonic functions on a lattice, one may exploit the mean-value property to solve a system of equations using iterative methods such as successive overrelaxation (SOR) [2].

Hitting Probabilities

If we denote $p(x)$ at any state x as the probability that the process, starting from x, will reach an obstacle before it reaches a goal, then $p(x)$ is known as a *hitting probability.* Harmonic functions on the configuration space of the vehicle are equivalent to collision probabilities. This interpretation of harmonic functions is described in more detail in [4], the result is that gradient descent on a harmonic function minimizes the probability of collisions with known obstacles.

3 Nonholonomic Constraints

Kane [5] defines simple nonholonomic constraints on the configuration q as:

$$\sum_{i=1}^{n} A_{ij}\dot{q}_i + B_j = 0 \qquad (j = 1, \ldots, m) \qquad (1)$$

where n is the dimension of the configuration space, and m is the number of constraints imposed. Equation 1 is a linear constraint on the configuration space velocity of the mechanical system. Simple nonholonomic constraints restrict the instantaneous velocity of the system to be within some linear subspace of configuration space.

3.1 Resistive Networks

Resistive grids are direct physical instances of harmonic functions. This perspective permits us to model nonholonomic constraints as nonlinearities in an otherwise homogeneous resistive medium and it leads directly to useful models for analog path planning. In this model, nodes corresponding to obstacle configurations are pulled high (1 *volt*), while goal nodes are pulled down (0 *volts*). The voltage gradient at free nodes corresponds to the gradient of the harmonic function. In practice, paths are generated by interpolating the function in the grid.

For a holonomic system, the grid nodes are connected to their manhattan neighbors. The point representing the current configuration is allowed to move in any direction in the grid. Figure 1 shows the connectivity of such a grid for a simple x, y, θ cspace.

To represent nonholonomic constraints, the appropriate subset of the 8 nearest neighbors in each of the θ-planes is connected by a uniform conductance. A grid connectivity that properly captures the nonholonomic constraints for each value of θ is illustrated in Figure 2. Every node is connected to its neighbors in the adjacent constant-θ planes. The grid is still fully connected so that there is a path between any two points. In contrast to more conventional nonholonomic planning methods, this technique allows continuous and simultaneous changes in the constraining degree of free-

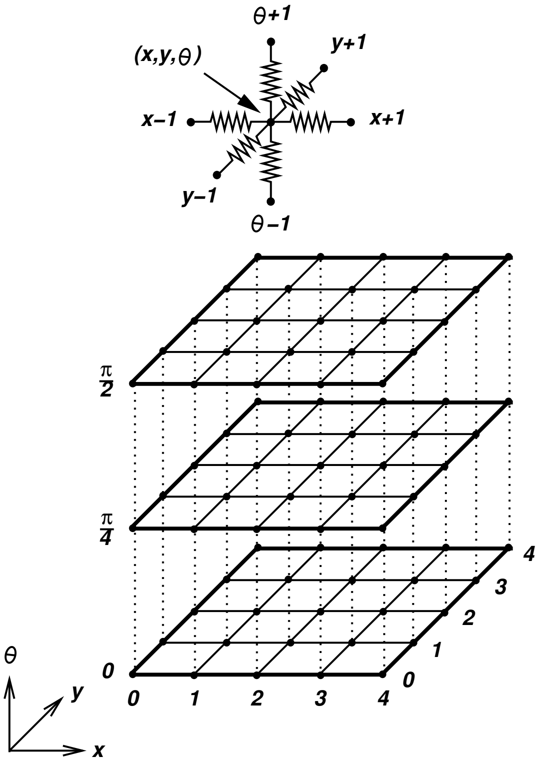

Figure 1: The holonomic 3D grid: lines denote resistances, and every interior grid node is connected to its manhattan neighbors.

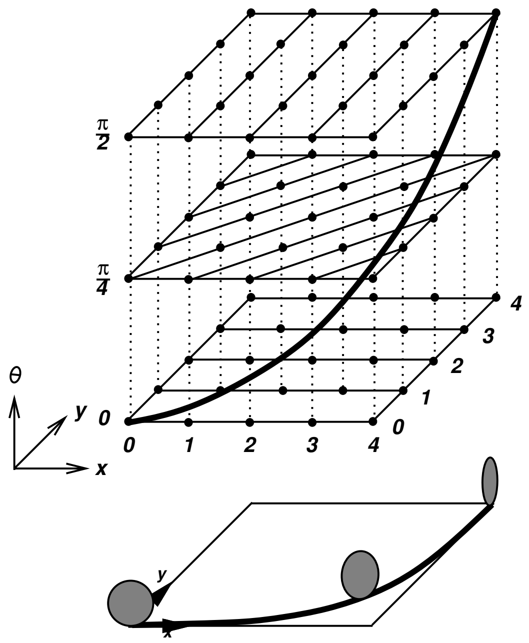

Figure 2: The nonholonomic 3D grid: each constant-θ plane has restricted connections in order to constrain the flow to be in the direction of θ.

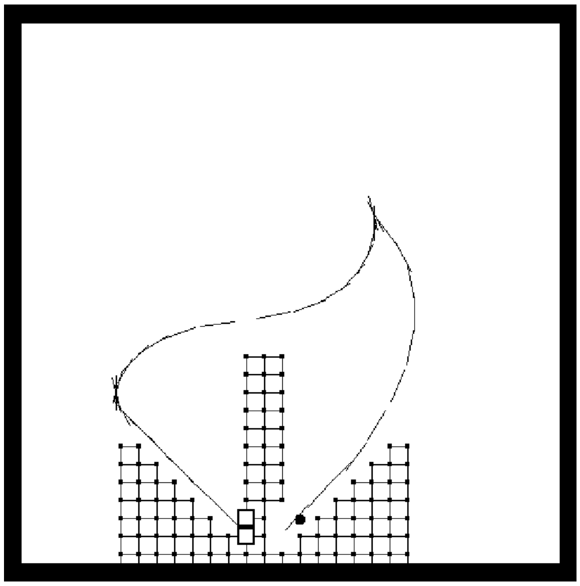

Figure 3: Example of a more complicated nonholonomic trajectory. The wheel starts in the right half of the Y-shaped corridor at $\theta = \frac{3\pi}{4}$ and proceeds to the goal at $\theta = \frac{\pi}{4}$.

dom (e.g., the steering angle), rather than a sequence of discrete infinitesimal reorientations.

3.2 Example Nonholonomic Control Tasks

Figure 3 illustrates an example of this approach applied to a single wheel. Four θ directions ($\theta \in \{0, \frac{\pi}{4}, \frac{\pi}{2}, \frac{3\pi}{4}\}$) are explicitly represented in the grid. The θ dimension wraps around so that $\theta = \pi$ maps to the $\theta = 0$ plane. This results in a configuration space which represents all unique headings of the wheel with no preferred forward direction. In each control cycle, the simulated wheel advances by δs and is simultaneously rotated by $\delta\theta$. The start configuration is at $\theta = \frac{3\pi}{4}$, and the goal orientation is $\theta = \frac{\pi}{4}$. The wheel must exit the channel on the right, and re-enter on the left. There is a cusp in the trajectory where the wheel reverses direction.

The unicycle (x, y, θ) is extended to the bicycle (x, y, θ, ϕ) to accommodate the trailing rear wheel.

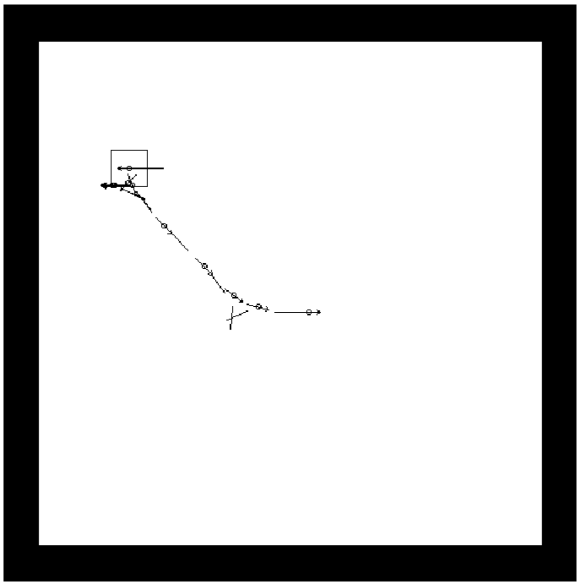

Figure 4: Example of a linear, nonholonomic bicycle trajectory from the middle, rightfacing configuration to the upper left, leftfacing configuration.

ϕ is vehicle heading and θ is the steering angle relative to ϕ. Figure 4 illustrates the front wheel of the bicycle as the arrow with a circle to designated the (x, y) control reference. The circle is drawn at the point at which the steering wheel contacts the ground, and is also the point at which the bicycle is articulated.

Figure 4 shows a trajectory where the vehicle backs up from the initial horizontal, right facing configuration to the x, y, ϕ goal and then rotates the steering angle θ into its goal configuration. There are no constraints on the relative steering angle in this example. It is a relatively simple matter to attach steering angle constraints as boundary conditions in the grid.

4 Energy-Reference Control

We consider the class of controllers consisting of a model-based controller whose command accelerations are obtained from the gradient of an artificial potential. We assume that the unit mass system has the following artificial Hamiltonian function:

$$H(q,p) = \frac{1}{2}||p||^2 + \phi(q) \tag{2}$$

where ϕ is an artificial potential, and q and p are configuration and momentum, respectively. This results in the following phase space trajectory:

$$\begin{aligned} \dot{q} = & \quad \tfrac{\partial H}{\partial p} & = p \\ \dot{p} = & -\tfrac{\partial H}{\partial q} & = -\nabla \phi \end{aligned} \tag{3}$$

An energy-referenced damper is used to regulate system energy. It is derived by computing the change in momentum required to achieve the energy reference.

$$\Delta p = [2(H_{ref} - \phi(q))]^{1/2} - p. \tag{4}$$

where q and p are the configuration and momentum, respectively, and Δp is the momentum change required to satisfy the energy economy. For a given reference energy, H_{ref} in Figure 5, this compensator is $\psi(\phi(q), p)$ defined as follows:

$$\begin{aligned} \psi(\phi(q), p) &= [2(H_{ref} - \phi(q))]^{1/2} - p \\ &\quad \textit{when} \;\; ((H_{ref} - \phi(q)) > 0) \\ &= -p \qquad \textit{otherwise} \end{aligned} \tag{5}$$

In the first case, $H_{ref} > \phi(q)$, and enough energy is introduced in the form of momentum to make up for the system energy deficit. In the second case, the system potential, $\phi(q)$, is too high and momentum is dissipated. In either case, kinetic energy is added or removed only along the current trajectory which is consistent with the nonholonomic constraints.

Obstacle potentials can be fixed at a uniform value, ϕ_{obs}, which is the maximum potential energy in the system. As long as $H_{ref} < \phi_{obs}$, the system can not encounter a configuration space obstacle with non-zero velocity. In general, equipotential sets of ϕ at $\phi = H_{ref}$ will bound the motion of the system. To adjust vehicle velocity, a single parameter is required to scale kinetic energy onto the range of potential energy.

Figure 6 illustrates the use of energy-referenced control for a roadway map. Two lane change scenarios are illustrated differentiated by the placement of the goal in either slow or fast lane. The map consists of Dirichlet boundary conditions representing road boundaries and a Neumann boundary condition representing the lane marker on a two lane highway. For a discussion of these boundary constraints, see [3]. A slower moving vehicle (white arrow in the black rectangle) is marked as an obstacle by placing the system of boundary conditions within the dotted lines in the roadway map. Within this obstacle template, the vehicle is represented as a Dirichlet B.C. and the lane marker is removed in the neighborhood of the obstacle. Control is expressed in terms of relative

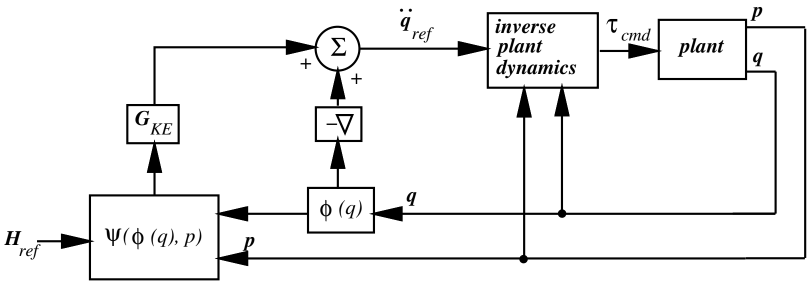

Figure 5: Energy-reference controller.

velocity. For example, while the vehicle and the obstacle may be traveling 65 *mph* and 60 *mph* respectively, this presentation assumes the obstacle is stationary and the vehicle approaches at a relative velocity of 5 *mph*. In practice, the obstacle B.C. template is dilated an amount proportional to approach velocity, and control is derived by scaling steering commands proportional to absolute velocity.

5 Hardware for Real-Time Motion Control

Automated vehicle path control must be responsive to changing road conditions. The local character of harmonic relaxations are well suited to parallel implementations.

Analog VLSI solutions for path control have been demonstrated [8, 7]. However, complex path planning problems lead to digital approaches which show more promise for larger grids, higher dimensions, programmability, testing and reliability.

5.1 Performance Requirements for Motion Control

The complexity of the numerical harmonic function solution scales super-linearly in the number of nodes in the grid (or roughly, exponential in the dimension of the grid). For this reason, it is clearly advantageous in terms of throughput, to specify coarse grids when possible. Several tradeoffs must be considered when designing the discrete grid for a particular application. The algorithm is complete up to the resolution of configuration space. Coarse grids may be overly conservative. Control actions are derived from gradients in the harmonic potential that are computed by finite difference approximation on interpolated potentials. The coarseness of this approximation dictates how precisely the system will trace the true path of steepest descent.

We are interested in recomputing the artificial potential field at a rate proportional to the highest frequency disturbance expected in the application. "Disturbances," in this context, relate to the changing structure of goals and obstacles in the environment.

5.2 Parallelism

There is a powerful **functional parallelism** in the two major steps in harmonic control: (1) SOR relaxation on a coarse grid with fixed-point arithmetic, and (2) gradient descent with floating-point interpolation on a fine grid. If the configuration memory is duplicated, the gradient descent for the previous step and the harmonic function for the present step can be computed simultaneously. With appropriate parallel hardware architectures, this leads to expected speed-up of approximately **2x** by exploiting functional parallelism.

There is also an inherent **data parallelism** in the numerical relaxation used to compute the harmonic function. The potential value at each grid point depends only on the values of its nearest neighbors and this implies that with a parallel hardware architecture an almost linear speed-up is achievable.

Harmonic function computation can be considered as consisting of two phases. A coarse first phase *propagates* boundary conditions into the free configuration space using fixed point precision to gain performance. To converge to the final artificial potential function, and subsequently to compute the gradient, floating-point precision is required. By adopting multiple grid resolutions

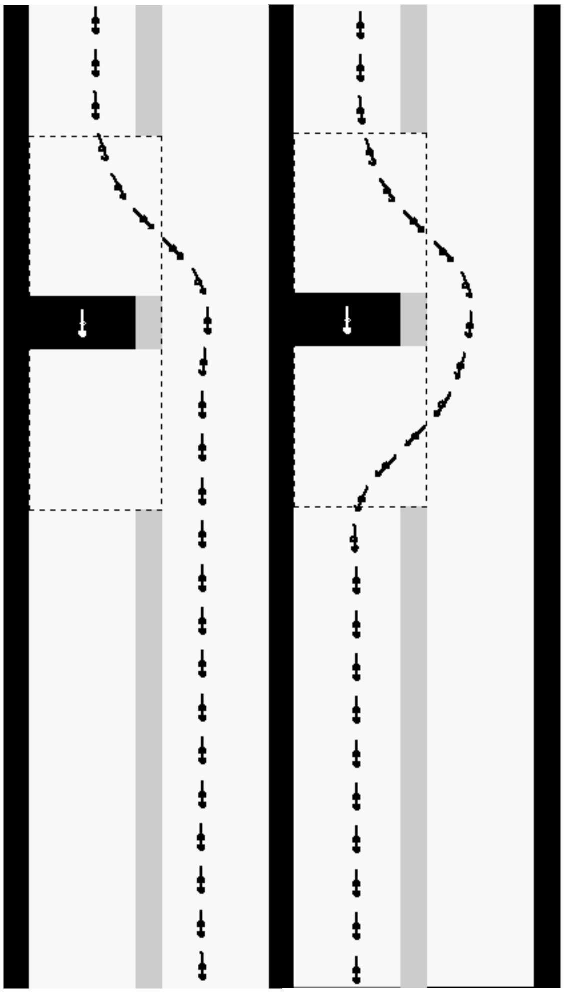

Figure 6: Example energy-referenced, nonholonomic control for highway lane change maneuvers in relative velocity space.

and numerical precisions reflecting these computational phases, we estimate a speed-up of at least **2x**.

In summary, we expect the following speed-up to be possible relative to a standard sequential implementation:

- **2x** by using DSPs who can access two operands from memory in each cycle as opposed to RISCs that have a load/store architecture,
- **4x** by using the 4 ADSPs in parallel,
- **2x** by using 8-bit operations in the first steps, followed by 16-bit and finally followed by 32-bit operations for harmonic function.

The current sequential version of the path control algorithm uses SOR on a 4D, 10,000 node grid and settles on average in about 13 iterations to 10^{-6}. This requires about 100 milliseconds on a SPARC 10. We expect roughly 6.25 milliseconds or 160 Hz on the DSP architecture.

References

[1] S. Axler, P. Bourdon, and W. Ramey. *Harmonic Function Theory*, volume 137 of *Graduate Texts in Mathematics.* Springer-Verlag, New York, 1991.

[2] R. Burden, J. Faires, and A. Reynolds. *Numerical Analysis.* Prindle, Weber and Schmidt, Boston, 1978.

[3] C. Connolly and R. Grupen. On the applications of harmonic functions to robotics. *Journal of Robotics Systems*, 10(7):931–946, 1993.

[4] C. Connolly and R. Grupen. Nonholonomic path planning using harmonic functions. In *Submitted to the 1995 IEEE Int. Conf. Robotics Automat.*, 1995.

[5] T. Kane. *Dynamics.* Holt, Rinehart and Winston, Inc., New York, 1968.

[6] E. Rimon and D. Koditschek. The construction of analytic diffeomorphisms for exact robot navigation on star worlds. In *Proc. 1989 IEEE Int. Conf. Robotics Automat.*, pages 21–26, Scottsdale, AZ, May 1989.

[7] M. Stan, W. Burleson, C. Connolly, and R. Grupen. Analog path computation using vlsi relaxation networks. *Journal of VLSI Signal Processing*, 8(1), Aug 1994.

[8] L. Tarassenko and A. Blake. Analogue computation of collision-free paths. In *International Conference on Robotics and Automation*, pages 540–545. IEEE, April 1991.

[9] R. Weinstock. *Calculus of Variations with Applications to Physics and Engineering.* Dover Publications, Inc., 1974.

[10] E. Zachmanoglou and D. Thoe. *Introduction to Partial Differential Equations with Applications.* Dover Publications, Inc., New York, 1986.