Applications of Harmonic Functions to Robotics *

Christopher I. Connolly
Laboratory for Perceptual Robotics,
Computer and Information Science Department,
University of Massachusetts at Amherst †

May 26, 1992

Abstract

Harmonic functions are solutions to Laplace's Equation. As noted in a previous paper, they can be used to advantage for potential-field path planning, since they do not exhibit spurious local minima. In this paper, harmonic functions are shown to have a number of other properties (including completeness) which are essential to robotics applications. These properties strongly recommend harmonic functions as a mechanism for robot control.

1 Introduction

Harmonic functions were proposed in Connolly et al. [1] as a means for generating paths in a robot configuration space. These functions are solutions to Laplace's Equation. One important feature of harmonic functions is that they can be used to generate smooth, collision-free paths without the threat of spurious local minima. In this paper, harmonic functions are shown to provide a global, complete path-planning scheme. As noted by Tarassenko and Blake [2], a resistive grid can compute such functions very quickly, further motivating their use.

It is important to note that nearly all the potential field models in use prior to [1] were not based on charge distribution. If they had been, then the resulting functions would have been

*IN PROCEEDINGS OF THE 1992 INTERNATIONAL SYMPOSIUM ON INTELLIGENT CONTROL.

†Portions of this work were funded by the National Science Foundation under grants CDA-8921080 and CDA-8922572.

harmonic, hence free of spurious local minima. The possible exception to this is the work of Rimon and Koditschek [3], whose functions also appear to be free of spurious local minima. The main motivations for the introduction of harmonic functions in [1] were the elimination of local minima and the use of a true analogy to phenomena such as charge distribution and fluid flow. In fact, the mathematical literature often refers to Laplace's Equation as the "potential equation" (see, for example, Courant and Hilbert [4]).

There is also the possibility that an analogue to the resistive grid exists within the brain [5]. The basal ganglia are subcortical nuclei within the brain which have been implicated in motor control. Recent research in neuroscience and neurology has uncovered properties of the basal ganglia which are very suggestive of a state-space planning scheme similar to that described here. A detailed discussion of these properties is beyond the scope of this paper. The interested reader is referred to Connolly and Burns [5] for further details.

In this paper, harmonic functions are shown to have several more useful properties which make them well suited for robotics applications:

- Fast surface normal computation
- Completeness (up to discretization error in the environment model)
- Ability to exhibit different modes of behavior (grazing vs. avoidance)
- Robust control in the presence of unanticipated obstacles and errors

2 Harmonic functions

A harmonic function on a domain $\Omega \subset \mathbf{R}^n$ is a function which satisfies Laplace's equation:

$$\nabla^2 \phi = \sum_{i=1}^{n} \frac{\partial^2 \phi}{\partial x_i^2} = 0$$

In the case of robot path construction, the boundary of Ω ($\partial\Omega$) consists of the boundaries of all obstacles and goals in a configuration space representation. Harmonic functions satisfy the min-max principle (see [4, 1]): Spontaneous creation of local mimima within the region is impossible if Laplace's equation is imposed as a constraint on the functions used. In the context of this paper, harmonic functions are computed over a regular grid of points which is a discrete representation of the robot configuration space. A finite difference scheme is often used to compute harmonic functions (see [1] for details), and one can also employ a resistive grid [2].

In [1], Dirichlet-style boundary conditions were used. With Dirichlet boundary conditions, obstacles are raised to a constant high potential, while goal regions are kept at a low potential. The resulting potential in free space is then constrained by $\nabla^2\phi = 0$. When Dirichlet boundary conditions are used in this manner, the gradient of ϕ is aligned with the surface normals of the obstacles. If this function is used for path generation, it will tend to *repel* the robot away from obstacles.

In [2], a Neumann-style boundary condition is proposed. Here, rather than keep obstacles at a constant potential, the derivative of ϕ is constrained so that $n \cdot \nabla\phi = 0$, where n is the surface normal function for the obstacles[1]. This causes $\nabla\phi$ to be tangential to obstacle surfaces. Here, the paths generated will *graze* any obstacles they encounter.

As noted in [1] and [2], the Dirichlet solution tends to have precision problems. Flat regions can develop resulting in very small (but of necessity nonzero) gradients, requiring higher precision in generating the solution. On the other hand, the grazing behavior exhibited by a Neumann solution can result in a tendency for a robot to stay close to obstacle surfaces.

3 Properties

The original motivation for the use of harmonic functions arose from the prevention of spurious local minima. However, there are several other useful properties than can be exploited. The following sections elaborate on these properties.

3.1 Surface Normals

In subsequent sections, some use will be made of the fact that the gradient of a constant-value Dirichlet solution to Laplace's Equation will be normal to the surfaces of boundaries. The Dirichlet solution for $\nabla^2\phi = 0$ (as described in section 2) causes any obstacle boundary to become an equipotential surface. Therefore, the gradient $\nabla\phi_D$ must be orthogonal to this surface [6]. Hence $\nabla\phi_D$ must be aligned with object surface normals. Several algorithms in computer vision rely on the computation of surface normals ([7], [8]). In combination with a resistive grid [2], harmonic functions provide a very fast method for surface normal computation on discrete solid models (e.g. bitmap or octree models of objects – see [9] for an example of how such data could be acquired).

[1] In general, Neumann boundary conditions are described by $n \cdot \nabla\phi(\mathbf{x}) = f(\mathbf{x})$, where $\mathbf{x} \in \mathbf{R}^n$. To get the gradient to flow *around* obstacles, one can simply set $f \equiv 0$.

3.2 Completeness

Under certain assumptions regarding the boundary conditions, it is possible to show that a path planning scheme using harmonic functions is complete (up to the approximation of the environment). We assume the following:

- Ω is a compact subset of $\mathbf{R}^n$.
- ϕ is a solution to Laplace's Equation on Ω with Dirichlet boundary conditions.
- Streamlines for ϕ are computed via gradient descent.
- All boundaries $\partial\Omega$ are held at the obstacle potential $\phi = 1$.
- All goal regions are held at the goal potential $\phi = 0$.

As noted previously, every streamline of a harmonic function must reach a goal point.

Conjecture 1 *Let* $\mathbf{x} \in \mathbf{R}^n$ *be the manipulator starting configuration, and* $\mathbf{y} \in \mathbf{R}^n$ *be any manipulator goal configuration. There is no path between* $\mathbf{x}$ *and* $\mathbf{y}$ *iff.* $\nabla\phi(\mathbf{x}) = 0$.

Proof: Let $S \subset \Omega$ denote the connected component of freespace which contains $\mathbf{x}$.

$\Rightarrow$ If there is no path from $\mathbf{x}$ to $\mathbf{y}$, it must be the case that $\mathbf{x} \in S$ and $\mathbf{y} \notin S$. Therefore, the only potentials within S must be the high potentials ($\phi = 1$) on ∂S. By the min-max property of harmonic functions, both the minimum and the maximum potentials in any neighborhood of $\mathbf{x}$, and since the neighborhood is arbitrary, any point in S, must be the same, i.e. $\phi(\mathbf{x}) = 1$. Therefore, ϕ is constant and $\nabla\phi$ at any point in S must be zero.

$\Leftarrow$ If $\nabla\phi(\mathbf{x}) = 0$ in S, then the derivatives of ϕ must be zero and ϕ must be a constant function. Consider a point $\mathbf{z} \in S - \partial S$. The min-max property requires that the minimum and maximum values of ϕ in any neighborhood of $\mathbf{z}$ must be attained on the boundary of the neighborhood. Let a given neighborhood of $\mathbf{z}$ contain points in ∂S. By the min-max property, since ϕ is constant, it must also be true that $\phi(\mathbf{z}) = 1$ (the value at ∂S). Since the neighborhood is arbitrary, $\phi(\mathbf{z}) = 1$ must hold for any $\mathbf{z} \in S$. For $\mathbf{y}$ to be in S then would be a violation of the min-max property, since $\phi(\mathbf{y}) = 0$. Hence $\mathbf{y}$ cannot be in S and there is no path between the effector and the goal.

Similar arguments may be applied to the case of Neumann boundary conditions. The key to these arguments is simply that the function in any connected component without goals will converge to a constant in that region (see Courant and Hilbert [4] for a detailed treatment of this property).

3.3 Behavior

Solutions to Laplace's Equation obey the principle of superposition. By superposing the Dirichlet and Neumann solutions, one can obtain a harmonic function that exhibits a behavior somewhere between the two original solutions. Let ϕ_D be a Dirichlet solution on a domain Ω, and let ϕ_N be a Neumann solution on Ω. A new harmonic function can be constructed by taking a linear combination of the two:

$$\phi = k\phi_D + (1-k)\phi_N$$

where $k \in [0,1]$. By varying k, the new function can be made to exhibit some mixture of the "repel" and "graze" behaviors inherent in the original solutions. It is easy to envision situations where either behavior might be more desirable than the other. If k is kept away from zero, the precision problems inherent in the Dirichlet solution are avoided. The resulting ϕ is harmonic, has no local minima and, provided that $k \in [0,1]$, guarantees collision-free paths.

To see why the collision-free property is maintained, consider what happens to harmonic functions at the boundaries of obstacles. The gradient of a Dirichlet solution is outward normal to the surface at these points, while the gradient of a Neumann solution is parallel to the surface. Any convex combination of the two solutions will therefore yield a gradient which is bounded by the Dirichlet and the Neumann gradient. This gradient must have a component which is outward normal to the surface.

The variability of robot behavior is especially useful in two ways: Repulsion from obstacles is desirable in fragile environments, where the environment model is accurate and where contact is to be avoided if at all possible. It is also desirable in the case where the robot system itself is fragile or can be damaged by contact with the environment. On the other hand, the grazing behavior of the Neumann solution is useful as an exploratory strategy, in cases where contact with obstacles is acceptable and even desirable. Force sensing can then be used to update the environment model.

To illustrate the effect of this style of superposition, consider the configuration space of a two-jointed robot where each link is 1 unit long. A rectangular obstacle is placed near the robot, and a smaller rectangular goal region is placed just above the obstacle (figure 1). Figures 2, 3 and 4 show the configuration space that results, and the streamlines obtained from the superposed harmonic functions for $k = 0$, $k = 1$, and $k = 0.5$ respectively. In each figure, the right side depicts the robot configuration space, while the left side shows its cartesian workspace. The lines in cartesian space show the paths of the end effector corresponding to the streamlines in configuration space. Lightly shaded regions are obstacles, and black regions are goal points.

Note that in figure 3, many paths graze the configuration space obstacles (the black region). In contrast, figure 2 shows streamlines which tend to stay well clear of obstacles. Finally, 4 shows the streamlines exhibiting a behavior somewhere between the two.

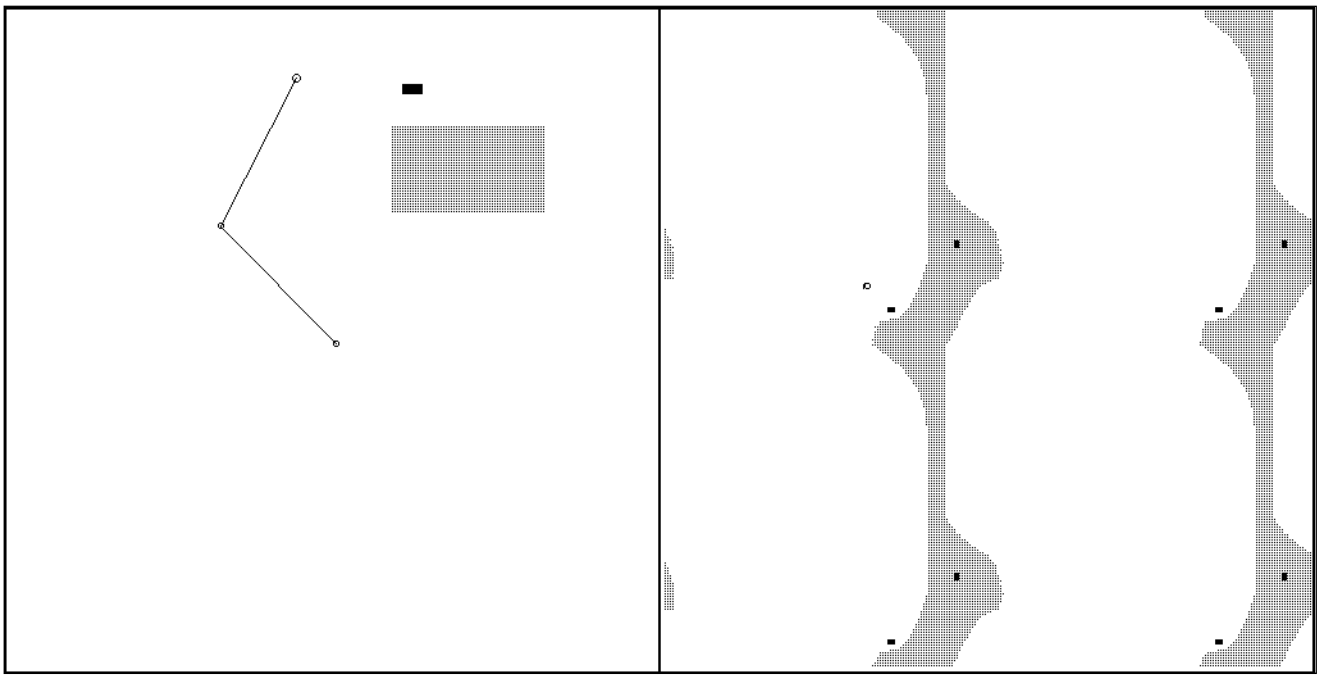

Figure 1: Cartesian space (left) and configuration space (right) for Problem A.

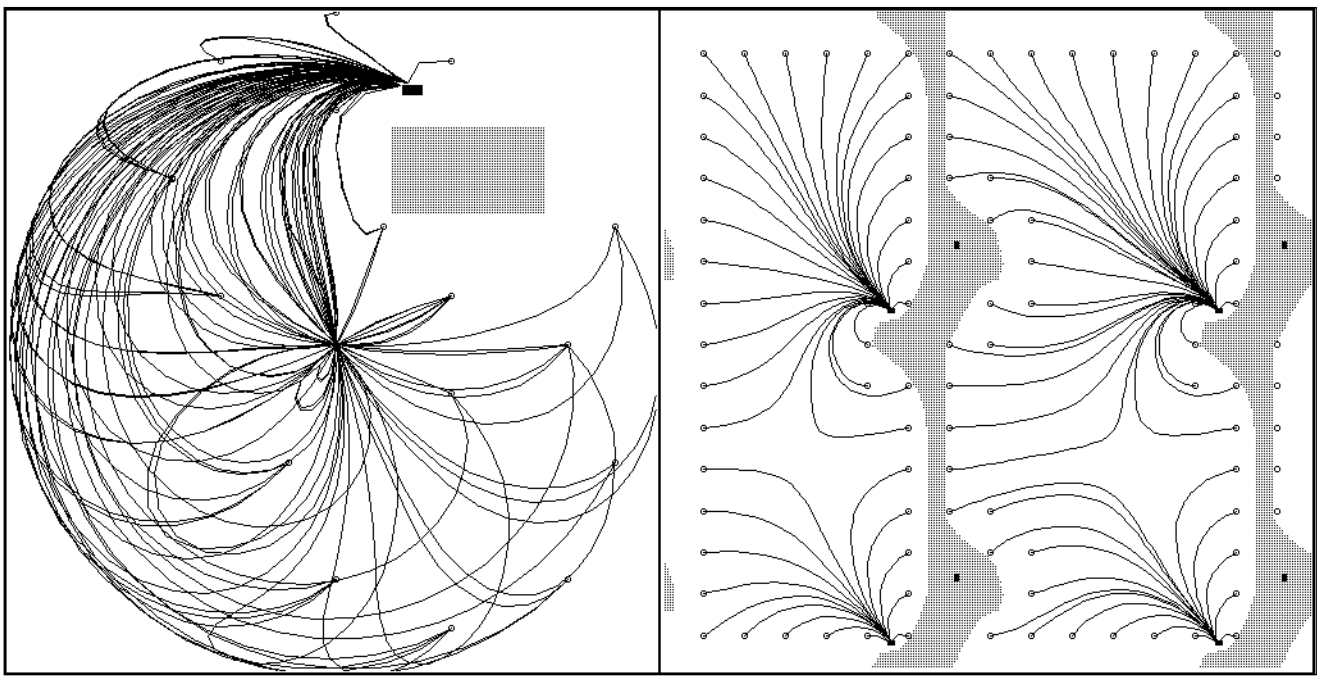

Figure 2: Streamlines for the Dirichlet solution of Problem A.

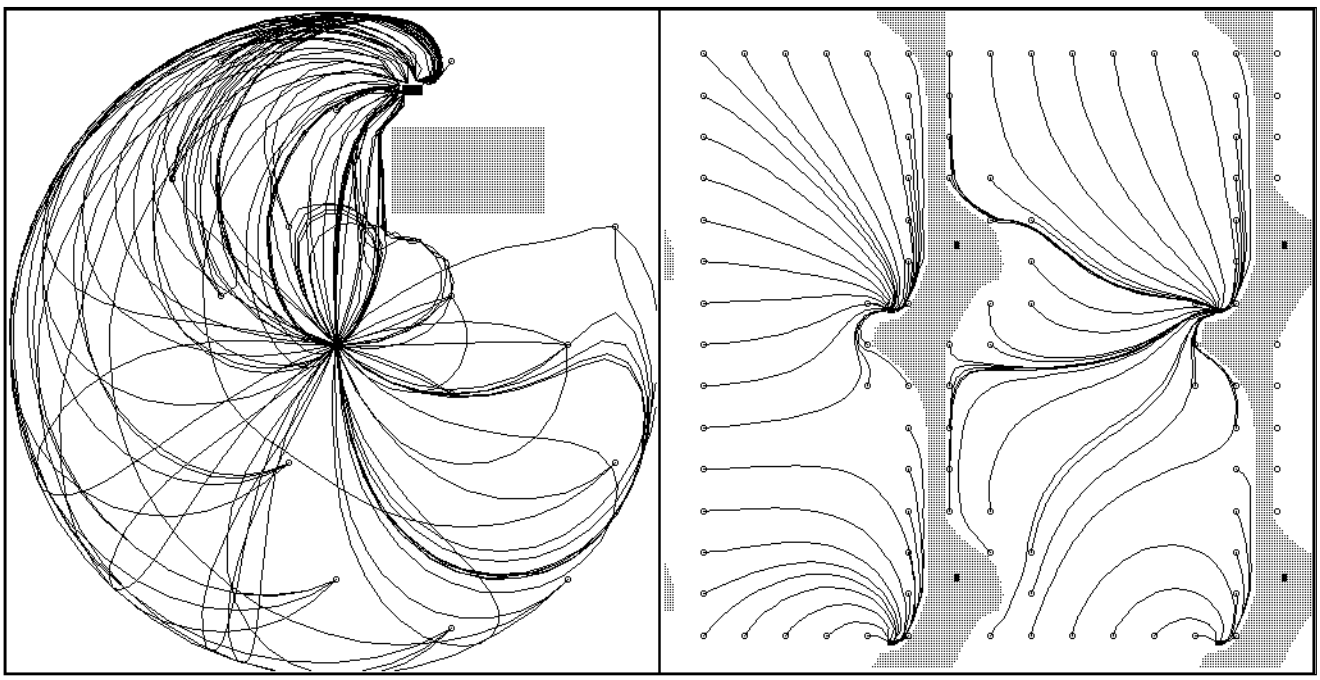

Figure 3: Streamlines for the Neumann solution of Problem A.

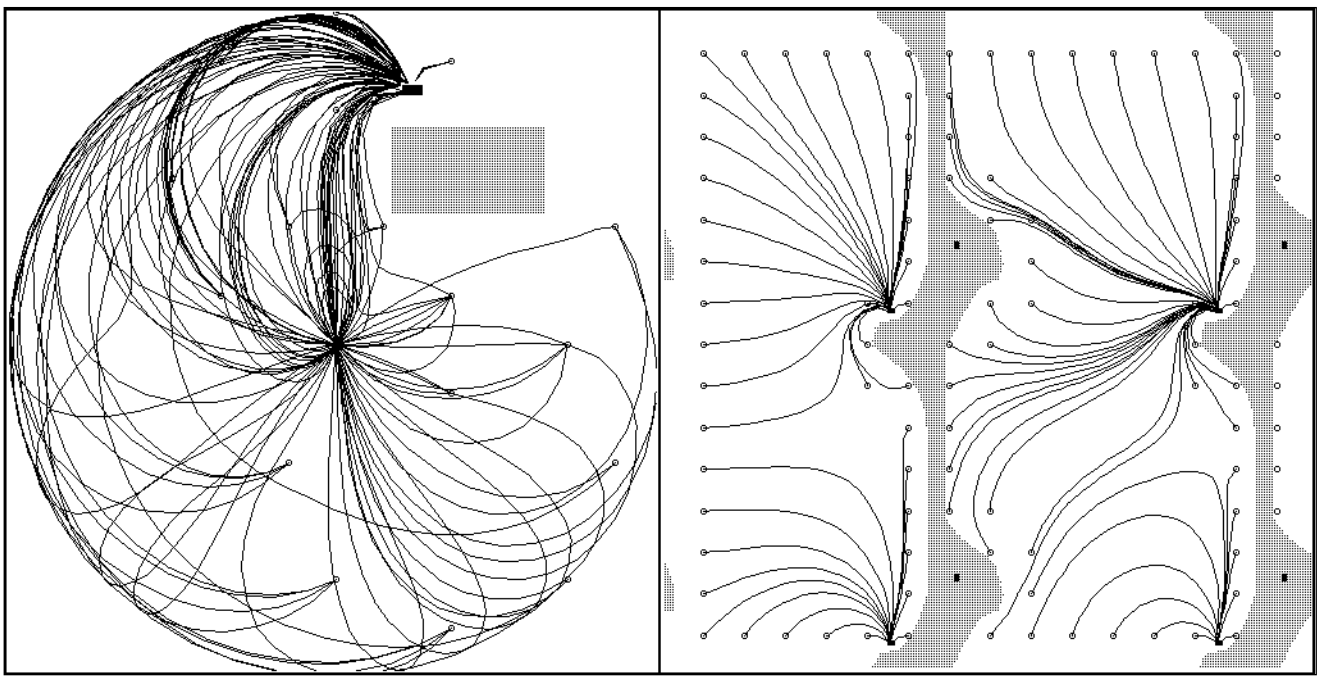

Figure 4: Streamlines for Problem A with a superposition constant of 0.5.

3.4 Velocity Control

Velocity control can be used to great advantage with harmonic functions. The gradient of a harmonic function can be used to supply velocity commands at every point in the robot workspace. In the absence of external forces, a robot so commanded will follow a streamline[2] in configuration space, which will always lead to a goal configuration [10]. Deflections in the streamline which result from external forces will simply place the robot on another streamline.

For example, admittance control [11, 12, 13] is ideally suited for use with harmonic functions. This control law is as follows:

$$\mathbf{V} = \mathbf{V_0} + \mathcal{A}\mathbf{F}$$

where $\mathcal{A}$ is an admittance matrix that alters the nominal command velocities $\mathbf{V_0}$ in response to external forces and torques $\mathbf{F}$. The scaled gradient of a harmonic function can be used as the command velocity $\mathbf{V_0}$.

A robot which uses harmonic functions to provide velocity control in this manner is robust in the presence of uncertainty in the robot's model of its environment. Errors in the model of the environment can result in contact between the robot and obstacles. This causes the robot to be deflected from its path. Since every streamline of a harmonic function *must* terminate at a goal point, the system can simply sample the harmonic function gradient at the current robot configuration and follow it. The application of harmonic functions to velocity control is described in more detail in [14].

4 Examples

A system has been implemented which demonstrates the use of harmonic functions for control of a revolute robot arm. The system simulates a two-link planar robot which uses harmonic functions for generating paths. In the experiments shown here, the left side of each figure depicts the cartesian space grid, while the right side shows the corresponding configuration space grid. As before, lightly shaded regions denote obstacles, while solid black regions are goals. Each harmonic function is computed over the configuration space grid. Not shown is a "virtual" obstacle which in each case is placed at singular points of the manipulator Jacobian matrix (diagonal lines in cspace).

Figures 5, 6, and 7 all show paths generated from the same starting configuration, but with different superposition constants. The letter "S" denotes the start configuration. Note the tendency of the end effector to roam closer to the obstacles as the constant approaches zero.

[2] The integral curve resulting from the gradient descent of the harmonic function.

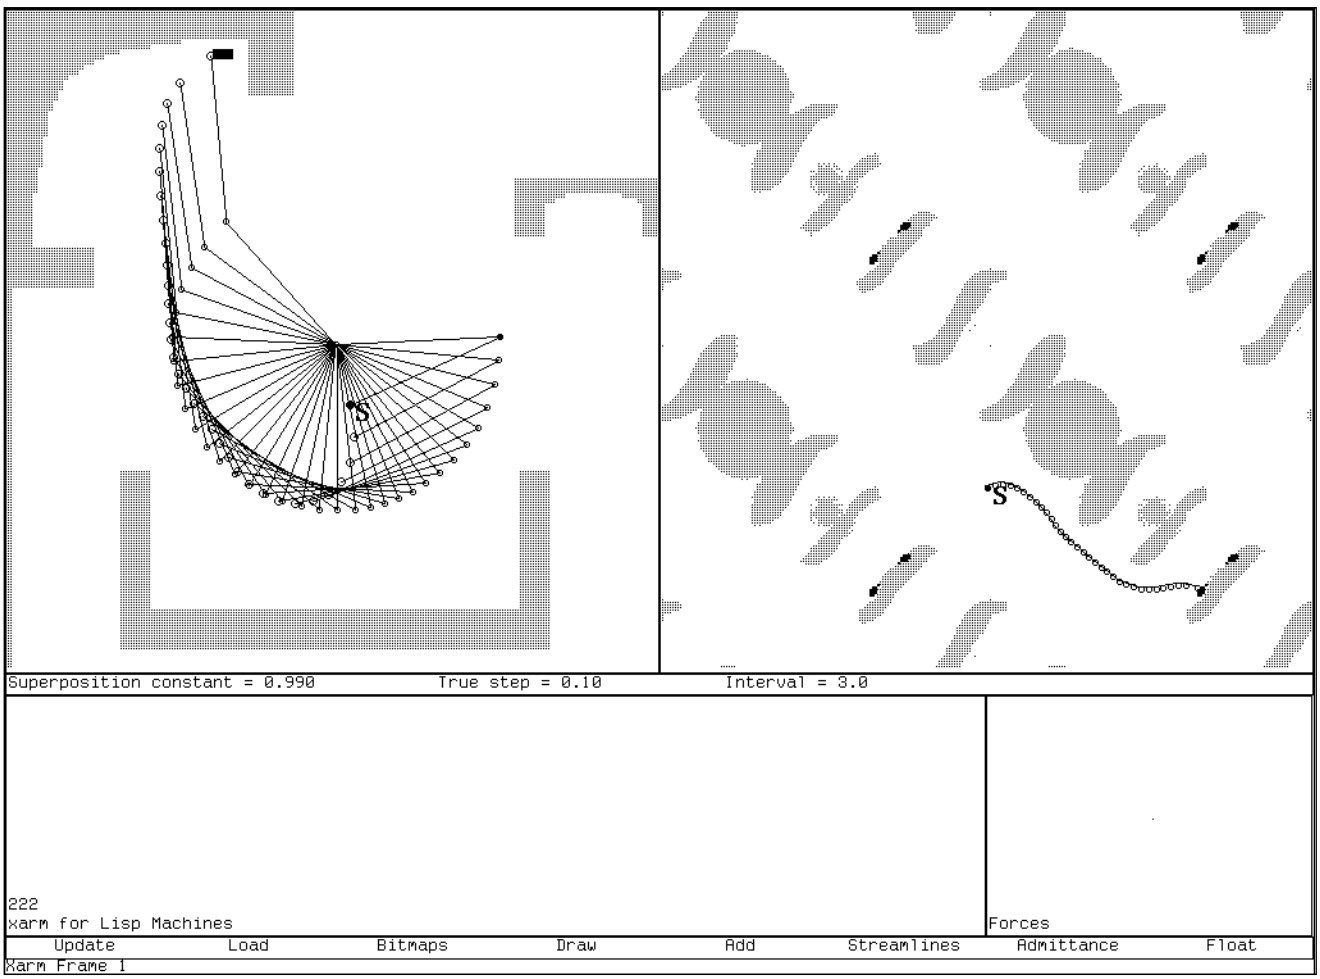

Figure 5: Path using $k = 0.99$ (Dirichlet).

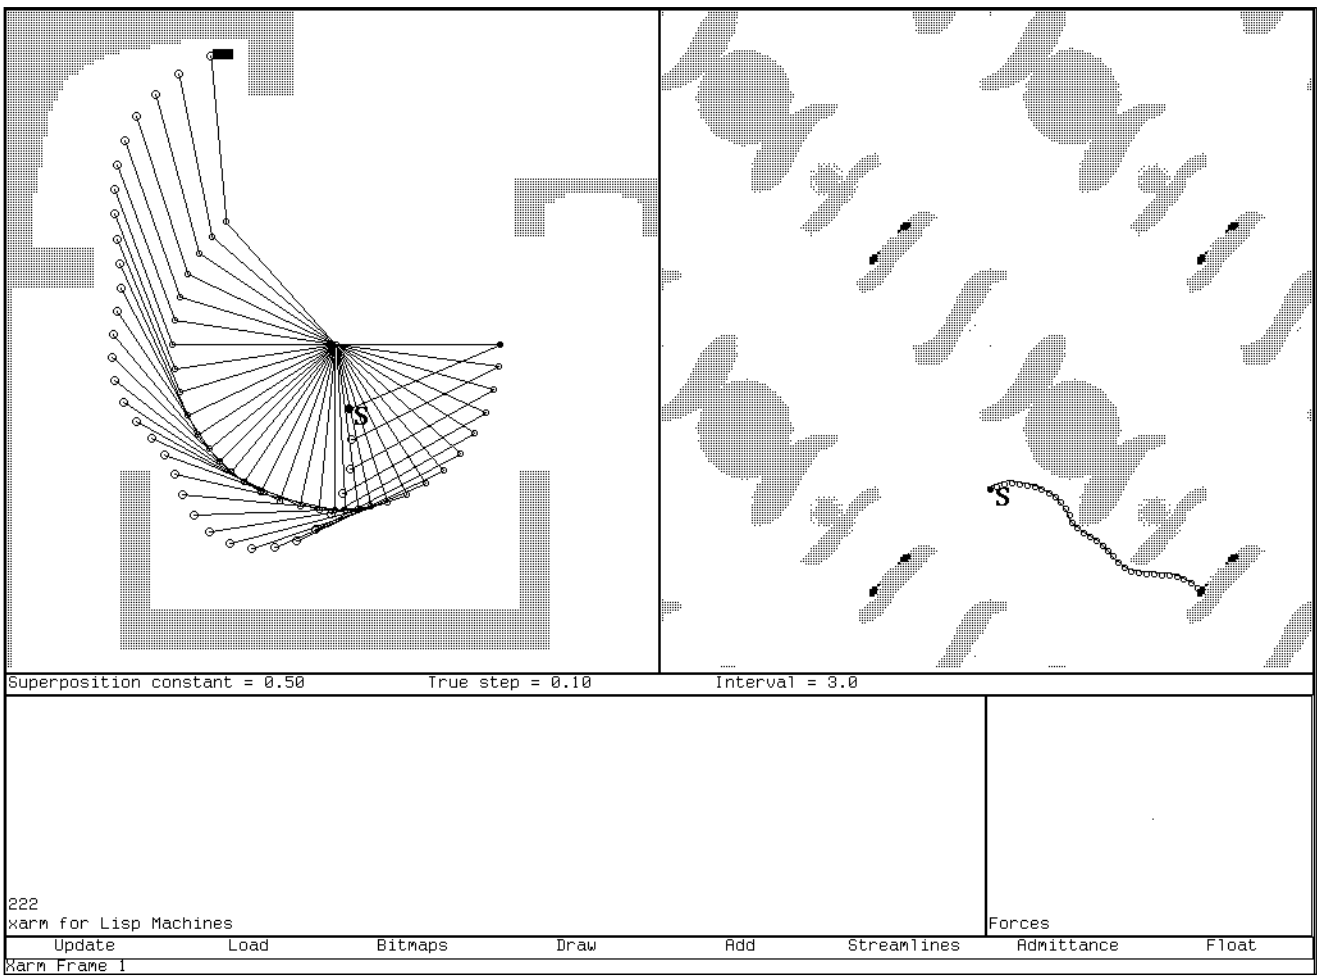

Figure 6: Path using $k = 0.5$ (blend).

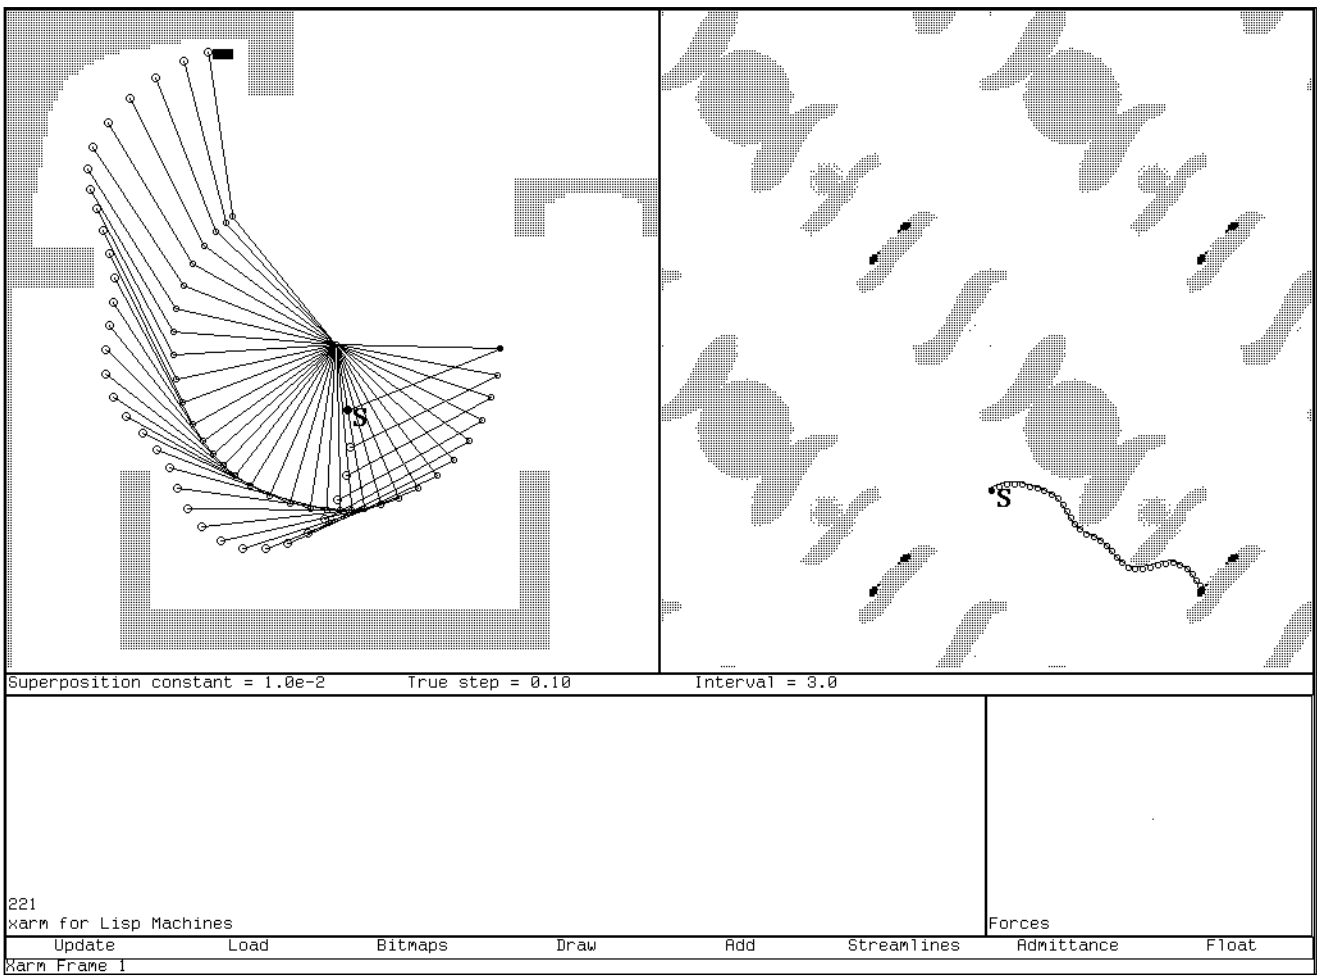

Figure 7: Path using $k = 0.01$ (Neumann).

In figures 8, 9, and 10, the ability of harmonic functions to use more than one goal region is demonstrated. These examples also help to show that harmonic functions are capable of dealing with a variety of environments.

Finally, figures 11, 12 and 13 show the ability of harmonic functions to detect that no path can be obtained from start to goal. This is a two-goal problem in which certain configurations are isolated from both goals. In this case, it is sufficient to test that the gradient is zero and that the harmonic function value is 1. Note that in figure 11, a path is found from start to goal. However, in figure 12, there is no path, and the program terminates and notifies the user that this is the case.

5 Conclusion

In conclusion, harmonic functions exhibit several properties which are useful for robot control:

- lack of spurious minima
- completeness up to discretization
- behavioral variability through superposition
- robust control in the presence of environmental model error
- computation of surface normals as a by-product

In conjunction with the results of [2], this demonstrates that harmonic functions provide a very attractive mechanism for robot control. It is also important to note that the method and properties described here may have a close biological analogue [5].

Even in the case where fast computation of harmonic functions is not possible, one can still drive the manipulator in parallel with the harmonic function calculation. Since new obstacles essentially place a high potentials at points in the configuration space grid, these potentials will result in wavefronts which will propagate outward from those points until convergence is again achieved.

The author would like to thank Kamal Souccar, who implemented the 2-link robot simulator.

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*, pages 2102–2106, May 1990.

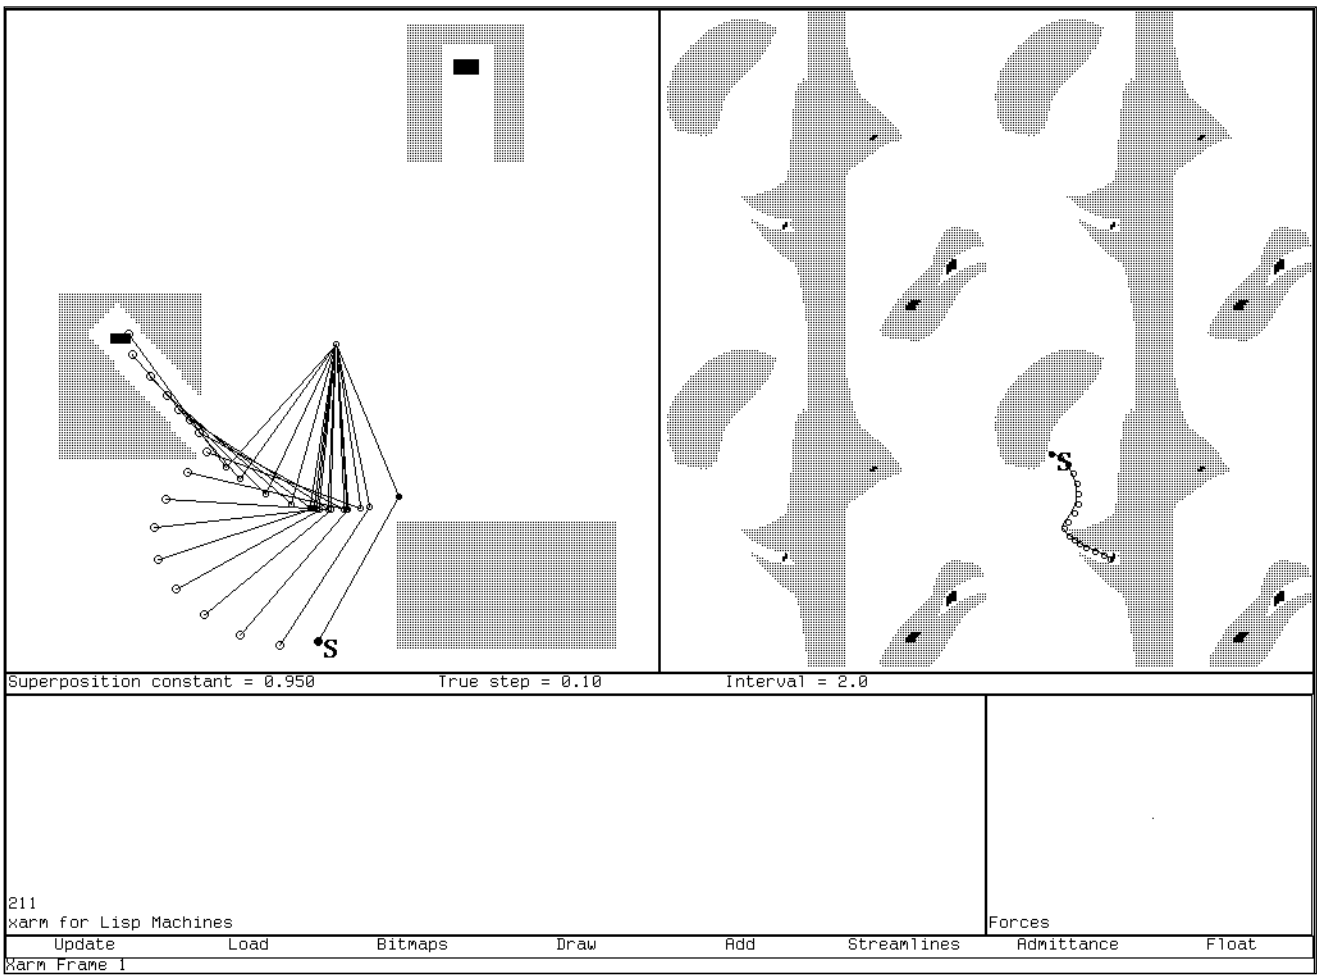

Figure 8: Two-goal problem: first path.

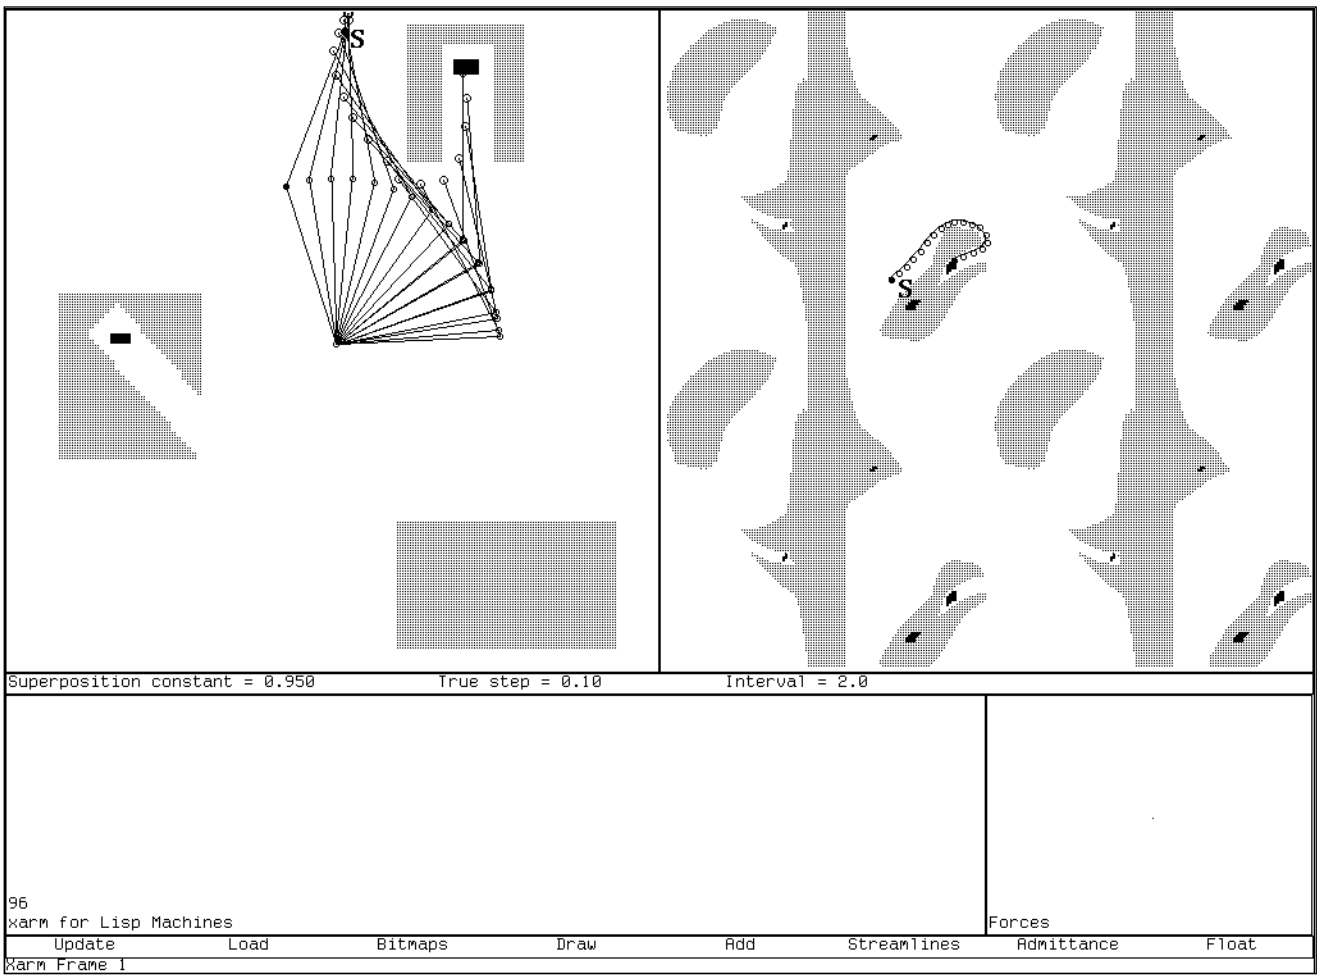

Figure 9: Two-goal problem: second path.

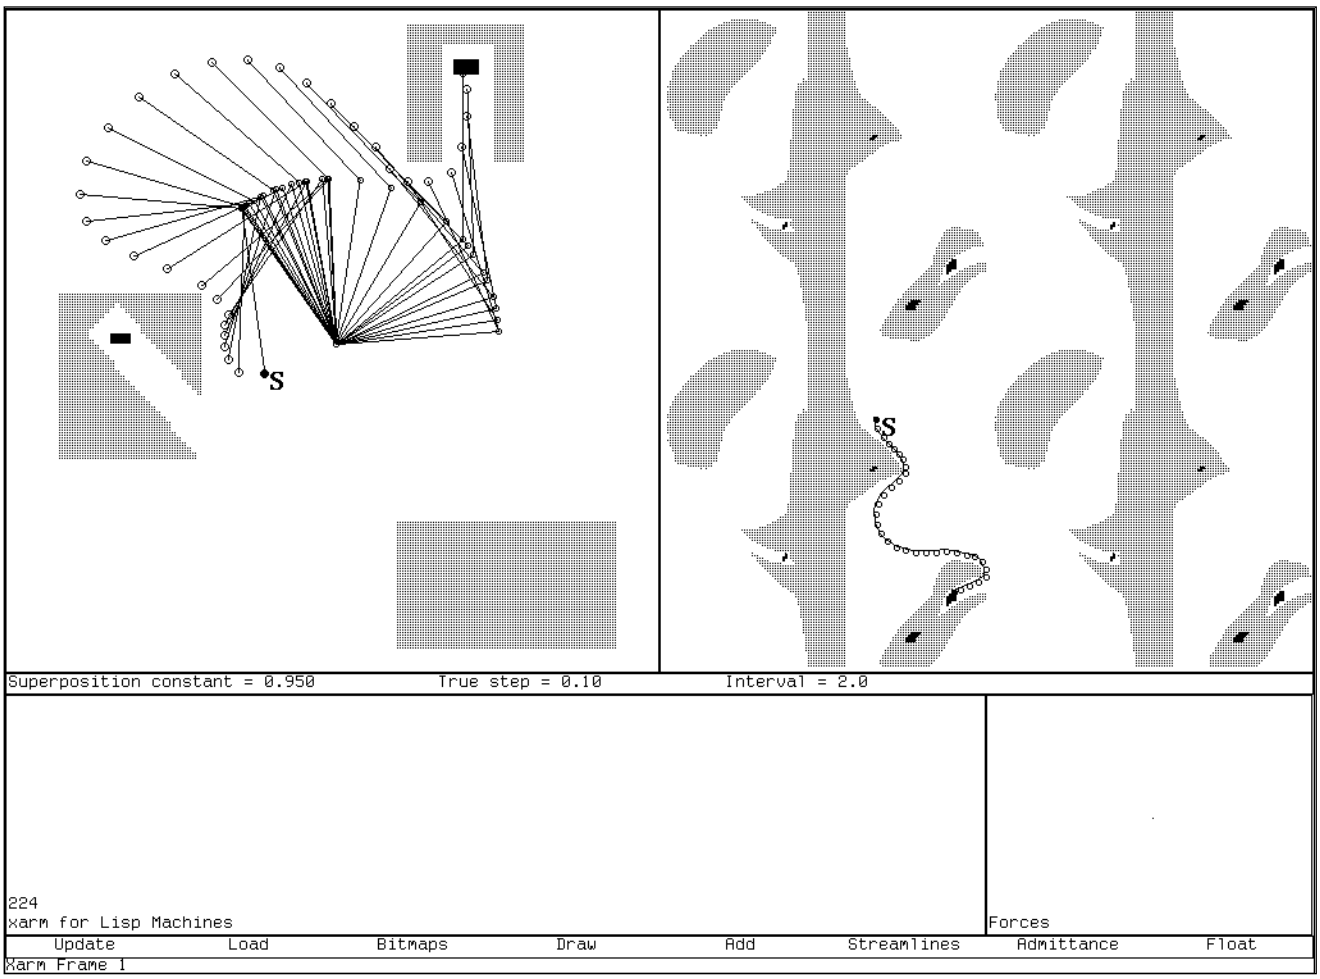

Figure 10: Two-goal problem: third path.

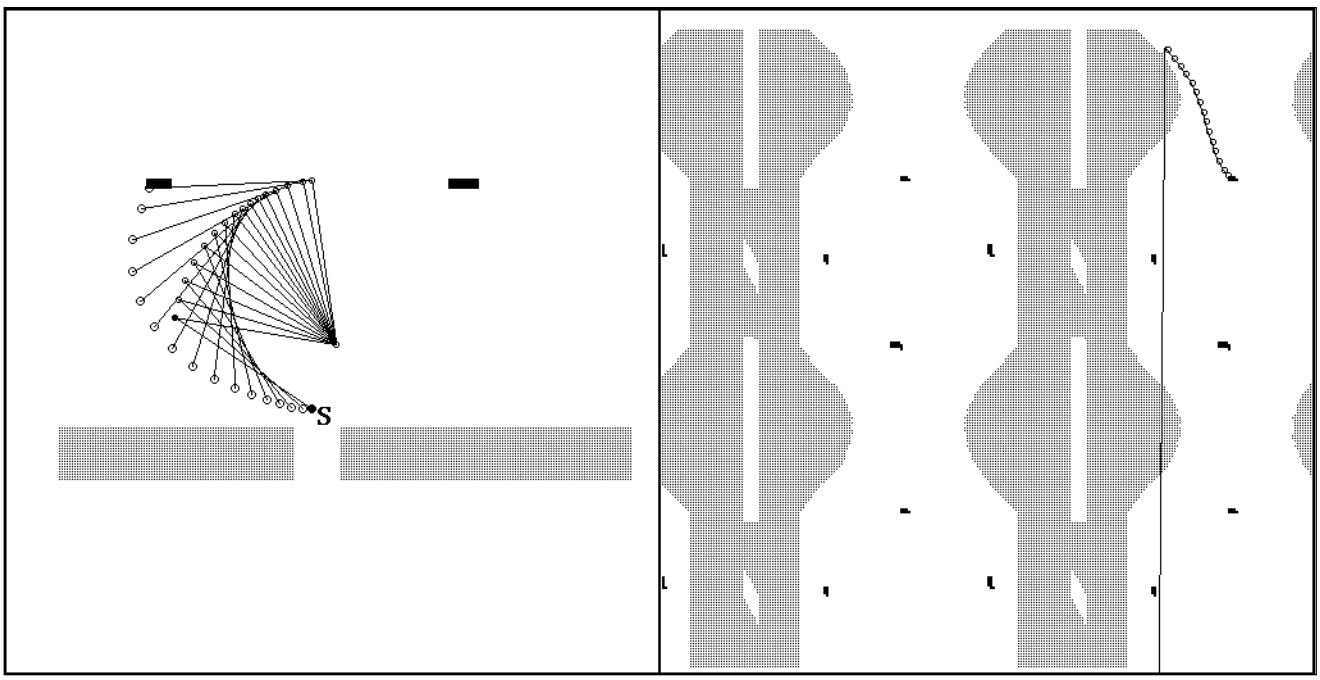

Figure 11: Problem B: A two-goal, two-obstacle workspace.

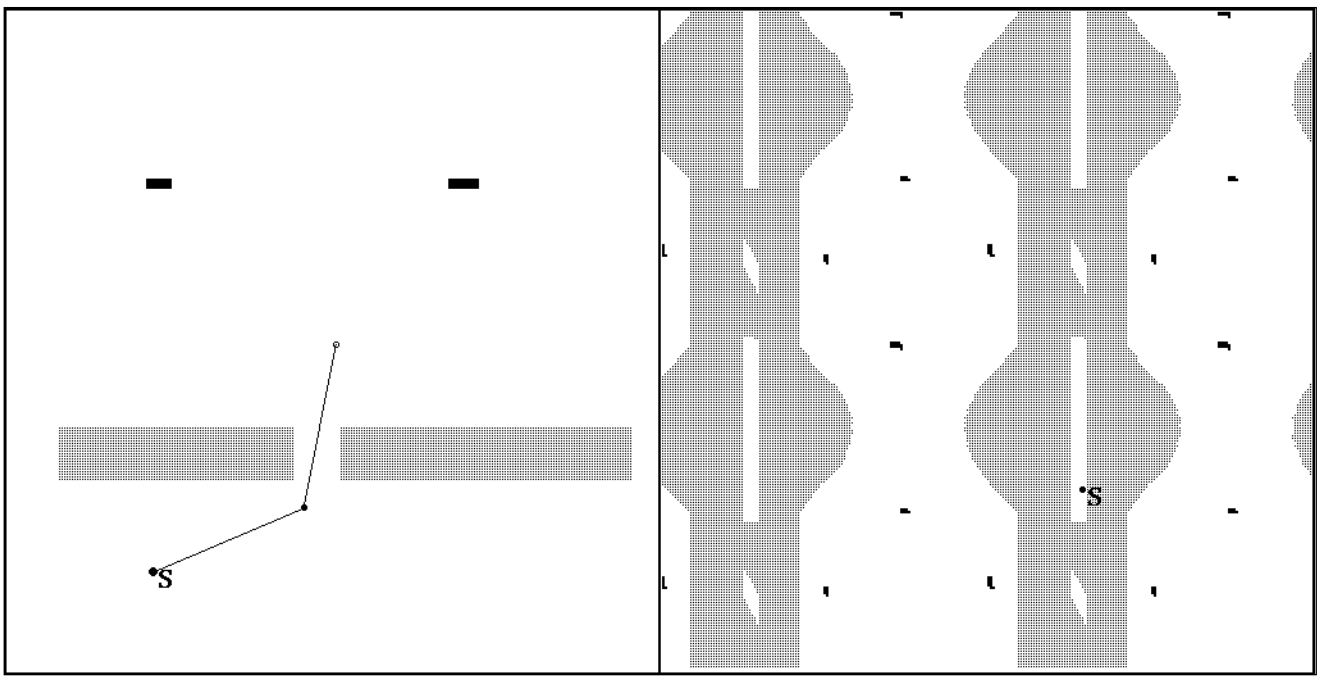

Figure 12: A start configuration for problem B which has no path to a goal.

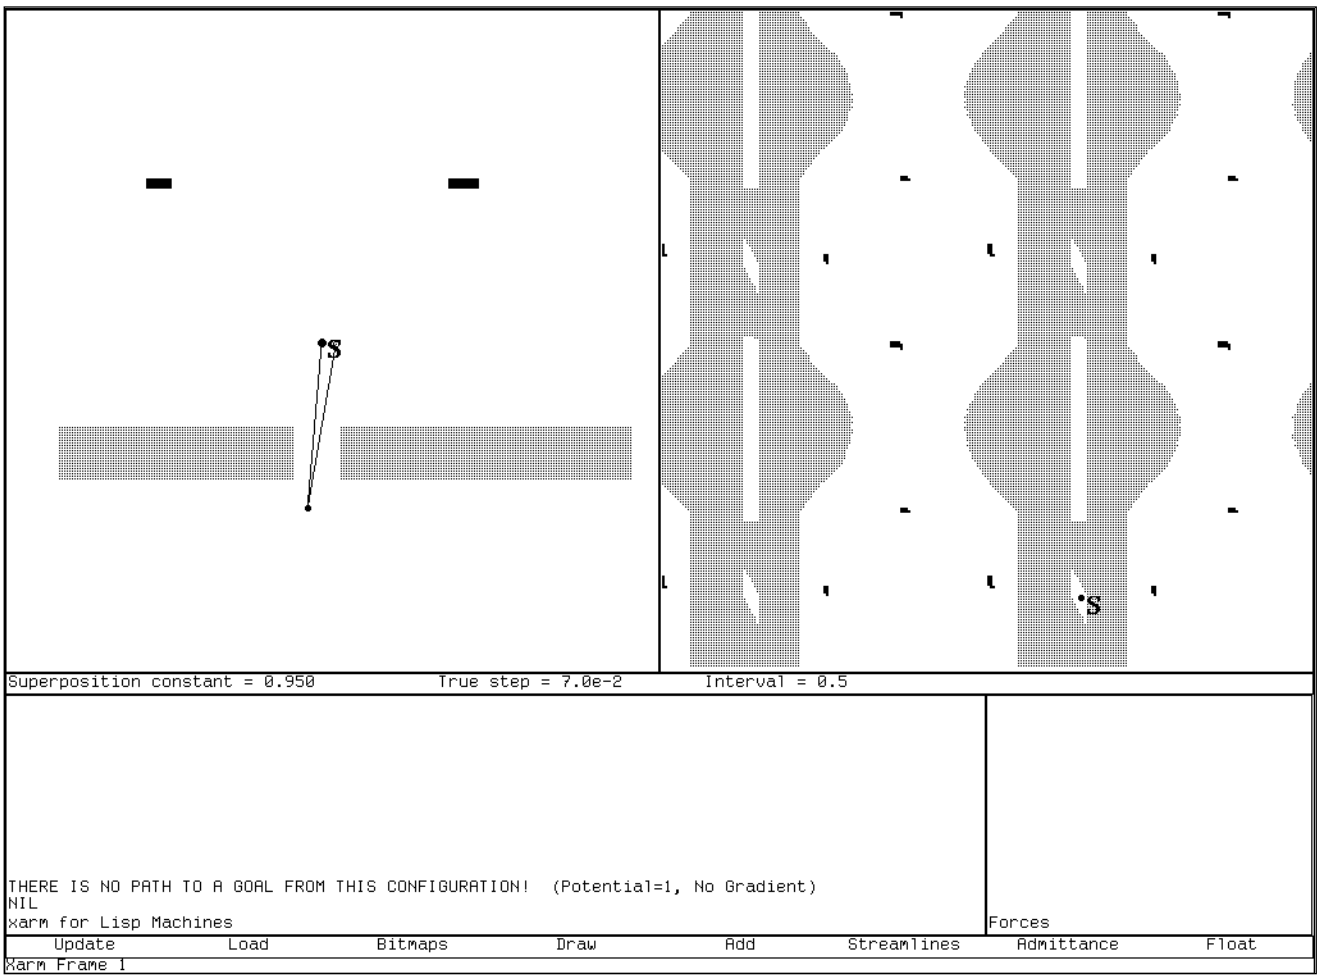

Figure 13: Another start configuration for problem B which has no path to a goal.

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

[3] Elon Rimon and Daniel E. Koditschek. Exact robot navigation using cost functions: The case of distinct spherical boundaries in E^n. In *Proceedings of the 1988 IEEE International Conference on Robotics and Automation*, pages 1791–1796. IEEE, April 1988.

[4] R. Courant and D. Hilbert. *Methods of Mathematical Physics*, volume 2. John Wiley and Sons, New York, 1989.

[5] Christopher I. Connolly and J. Brian Burns. The planning of actions and the basal ganglia. In J. Hendler, editor, *Artificial Intelligence Planning Systems: Proceedings of the First International Conference (AIPS92)*. Morgan Kaufmann, San Mateo, CA, 1992.

[6] Ruel V. Churchill, James W. Brown, and Roger F. Verhey. *Complex Variables and Applications*. McGraw-Hill Book Company, 3rd edition, 1976.

[7] Paul Besl and Ramesh Jain. Intrinsic and extrinsic surface characteristics. In *Proceedings of the 1985 Conference on Computer Vision and Pattern Recognition*, page 226. IEEE, June 1985.

[8] John Ross Stenstrom. Finding points of high curvature and critical curves from a volume model. In *Proceedings of the 1985 Conference on Robotics and Automation*, page 436. IEEE, April 1985.

[9] C. I. Connolly. Cumulative generation of octree models from range data. In *Proceedings of the First International Conference on Robotics and Automation*, page 25. IEEE, March 1984.

[10] Rolf H. Sabersky, Allan J. Acosta, and Edward G. Hauptmann. *Fluid Flow*. MacMillan Publishing Company, 2nd edition, 1971.

[11] Daniel E. Whitney. Force feedback control of manipulator fine motions. *Journal of Dynamic Systems, Measurement, and Control*, pages 91–97, June 1977.

[12] Daniel E. Whitney. Historical perspective and state of the art in robot force control. *International Journal of Robotics Research*, 6(1):3–14, Spring 1987.

[13] Michael A. Peshkin. Programmed compliance for error corrective assembly. *IEEE Transactions on Robotics and Automation*, 6(4):473–482, August 1990.

[14] C. I. Connolly and R. Grupen. Applications of harmonic functions to robotics. Technical Report 92-12, COINS Department, University of Massachusetts, February 1992.