

A Model for Integrated Qualitative Spatial and Dynamic Reasoning about Physical Systems *

Raman Rajagopalan
Artificial Intelligence Laboratory
Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712
raman@cs.utexas.edu

Abstract

Qualitative spatial reasoning has many applications in such diverse areas as natural language understanding, cognitive mapping, and reasoning about the physical world. We address problems whose solutions require integrated spatial and dynamic reasoning.

In this paper, we present our spatial representation, based on the extremal points of objects, and show that this representation is useful for modeling the spatial extent, relative positions, and orientation of objects, and in reasoning about changes in spatial relations and orientation due to the translational and rotational motion of objects. Our theory has been implemented to support a magnetic fields problem solving application using the QPC and QSIM systems for qualitative modeling. The issues encountered in integrating spatial and dynamic reasoning in the context of these systems are also discussed.

Introduction

Spatial reasoning has been studied in Artificial Intelligence from many perspectives, including natural language understanding [Freksa 92, Mukerjee and Joe 90, Retz-Schmidt 88], cognitive mapping [Kuipers and Levitt 88], and qualitative reasoning about physical systems [Forbus, et. al. 91, Joskowicz and Sacks 91, Nielsen 88, Weinberg, Uckun, and Biswas 92]. A fundamental starting point for spatial reasoning is the representation of the spatial extent of an object. Problems which require knowledge of the exact shapes of objects, such as determining whether two gears will mesh together, will require a numerically precise spatial representation, such as the configuration space [Forbus, et. al. 91, Joskowicz and Sacks 91].

For many problems, knowledge of the *approximate extent* of objects is sufficient, or that could be only information available. Several spatial reasoning methods, based on the use of simplifying abstractions to

approximate the actual extents of objects [Abella and Kender 93, Cui, Cohn, and Randell 92, Mukerjee and Joe 90], have been developed to address these cases. None are sufficiently rich to simultaneously model all of the following properties: the spatial extents, relative positions, and orientations of objects, and the effects of translational or rotational motion on the spatial state.

For example, Cui, Cohn, and Randell [1992] use convex hulls as an abstraction method. They present techniques for reasoning about the relative positions of objects, but do not discuss the orientation of objects. They also do not provide a *qualitative* method to describe the convex hull itself, or how to compute its changing location or orientation if the underlying object translates or rotates. They assume that an external function will provide this information.

We consider problems whose solutions require integrated spatial and dynamic reasoning. Specifically, we are implementing a magnetic fields problem solver, which, given a diagram to describe the initial spatial state and text to describe any dynamic changes taking place, solves the problem through qualitative simulation. An input processing specialist, The Figure Understander [Rajagopalan and Kuipers 94], is used to integrate the information in the text and diagram input into a unified model of the initial state. A library of model fragments, for spatial reasoning and for the magnetic fields domain, is used by the QPC [Crawford, et. al. 90, Farquhar 94] qualitative modeling system to detect spatial configurations that could lead to dynamic processes. The effects of any active processes are determined using the QSIM [Kuipers 86] qualitative simulation system.

Our problems cover the operation of such commonly used devices as motors, generators, and transformers. Relative position and orientation information is used to study the effects of changing the magnetic flux passing through a closed conducting object. The change in flux can be due to a time-varying magnetic field (transformer), or to the translational or rotational motion of the conducting object (motor, generator).

Forbus, Neilsen, and Faltings [1991] and Joskowicz and Sacks [1991] also address the issue of integrated

*This work has taken place in the Qualitative Reasoning Group at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Qualitative Reasoning Group is supported in part by NSF grants IRI-8904454, IRI-9017047, and IRI-9216584, and by NASA contracts NCC 2-760 and NAG 9-665.

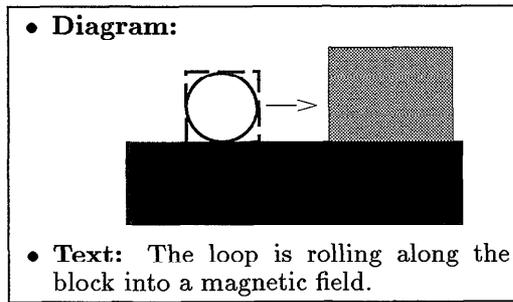


Figure 1: A magnetic fields problem involving translational motion: What happens as the conducting loop (white) rolls across the magnetic field (gray)? The rectangular bounding box abstraction shown for the loop is not included in the input to the problem solver.

spatial and dynamic reasoning, but the characteristics of our problems are not well suited for the configuration space-based approach they have adopted. They work with domains where objects cannot share the same physical space, and objects cannot be dynamically created or destroyed. In such cases, the configuration space must be computed only once, and the boundaries of the forbidden regions can be used to detect the points at which two objects are in contact. In our problem set, the most interesting spatial relation is overlap between objects, and since magnetic fields can be created simply by establishing a current flow in a wire, we must allow for the dynamic creation and destruction of spatial objects.

In the following sections, we describe our spatial representation and its use in reasoning about the relative positions between objects and the orientation of convex objects, particularly in dynamically changing worlds. We will also discuss the issues encountered in integrating spatial and dynamic reasoning in the context of the QPC and QSIM systems.

Example Magnetic Fields Problems

Consider the scenarios in Figures 1 and 2. In Figure 1, a conducting loop is rolling into a magnetic field. In Figure 2, a conducting loop is rotating inside a magnetic field. The problem is to describe what happens as the specified motions take place.

We have to recognize changes in the *relative position* (figure 1) or *orientation* (figure 2) of the loop relative to the field to determine when the magnetic flux through the loop is changing. During those periods, an induced emf will be established in the loop such that the resulting current flow will produce a magnetic field that opposes the change in flux through the loop. Magnetic forces can then act on the loop, and orientation information can be used to determine the net magnetic force/torque and its effects (e.g., the loop in Figure 1 will slow down as it enters and exits the field).

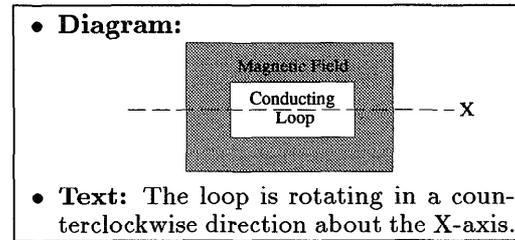


Figure 2: A magnetic fields problem involving rotational motion: A rudimentary generator.

The Figure Understander

The Figure Understander works with text input and a numerical description of a diagram, a PostScript file produced by the InterViews drawing editor, to output a qualitative, constraint-based description of the initial state. A picture description language is used to define a semantics for diagram objects, greatly simplifying the task of processing the input. For example, for Figure 2, given that white objects are conducting loops, The Figure Understander easily associates the direction of rotational motion with the smaller rectangular object.

The diagram-based input method serves two purposes. First, it is a more natural and less error-prone method of entering the initial spatial state than using a special purpose text form. Second, it allows numerically-precise spatial data to be extracted for solving problems where such information is necessary. For example, The Figure Understander includes heuristics for selecting frames of reference, and when appropriate, numerically rotates the diagram before extracting the initial spatial state. The Figure Understander corrects minor human errors in drawing the diagram by assuming that coordinate values within a given epsilon measure are equal.

The Spatial Representation

When an abstraction is used to *approximate* the actual spatial extent of an object, there can exist real distinctions that can no longer be reasoned about without additional case-specific information. If a maximally covering abstraction is used, such as a convex hull, one can *guarantee* non-intersection conclusions based on examining only the approximate shapes, but not intersection conclusions. For example, if the convex hulls of two objects are not intersecting, then the underlying objects cannot be intersecting. However, it's possible that the underlying objects do not intersect even if the convex hulls intersect. Similarly, if an abstract shape is fully contained within the actual shape, then one can guarantee intersection conclusions, but not non-intersection conclusions.

Our problems require reasoning about what can happen *if* two objects are intersecting or *if* one object is fully enclosed inside another. We enforce the con-

straint that a shape abstraction should fully enclose an underlying object at all times, to guarantee when it is at all possible for two objects to be intersecting.

We also wish to minimize the need for quantitative information in describing the initial spatial state and to maintain the spatial state as objects translate and rotate. The goal is that a qualitative solution should apply to any physical scenario that fits the initial state description. The generality of the solution will degrade with the requirement for precise numerical information in describing the initial spatial state or in determining subsequent spatial states.

Modeling Spatial Extent

We describe the spatial extent of an object qualitatively in terms of its **extremal points**: the **topmost** (tm), **rightmost** (rm), **bottommost** (bm), and **leftmost** (lm) points [Rajagopalan 93]. To reason about three dimensional objects, we add the frontmost and rearmost points.

We use two abstractions, a **rectangular bounding box** drawn around the extremal points of an object, and a **bounding circle**, as described below, to approximate the region occupied by the object. For problems in three-dimensional space, we use a **bounding cube** and a **bounding sphere** respectively. In Figure 3, we illustrate our abstraction method on two polygons. We draw a bounding circle around object A and a rectangular bounding box around object B. The radius of the bounding circle is the maximum distance from the center of gravity of an object to any point on its perimeter.

The bounding box abstraction is sufficient for reasoning about problems involving static worlds or only translational motion. The bounding circle/sphere abstraction is used for problems involving rotation about the center of gravity. The latter abstraction provides the guarantee that an underlying rotating object will always remain within the bounding abstraction, and that the abstractions will have a fixed shape.

The bounding circle/sphere abstraction is, in general, weaker than the bounding box abstraction. Consider the real possibilities that a rotating object could come into contact with another object, and that rotation could remove an existing contact. If two bounding spheres are not intersecting, we can conclude with certainty that the underlying objects cannot come into contact after any rotation of either object about its center of gravity. If two bounding spheres intersect, they will continue to intersect after any rotation of the underlying objects. Additional information, such as the distances between the objects, their sizes, and relative orientations, will be required in that case to determine with certainty if rotation can create or remove contacts between the underlying objects.

Modeling Relative Positions/Orientation

We describe the relative positions of objects through *inequality relations* between their extremal points. If

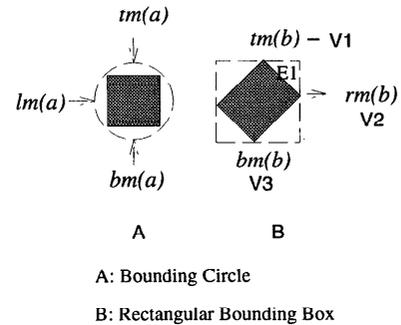


Figure 3: Examples of the rectangular bounding box and bounding circle as methods of shape abstraction for two polygons.

the rectangular/cubic bounding box abstraction is used, the extremal points of the shape abstraction and the underlying object are the same. If the bounding circle/sphere abstraction is used, the extremal points used are those of the bounding abstraction, and not the underlying object. For example, for the objects in Figure 3, the inequality relation $rm(A) < lm(B)$, represents the fact that A is to the left of B.

Although we can conclude that objects are intersecting with certainty in some special cases (such as overlap across an extremal edge), in general, for relations involving intersections between objects, we can prove only the weaker relation that the bounding abstractions of the objects satisfy the desired property. For example, object A **bounding-abstraction-encloses** object B if $(lm(A) \leq lm(B)) \wedge (rm(A) \geq rm(B)) \wedge (tm(A) \geq tm(B)) \wedge (bm(A) \leq bm(B))$.

We define the **orientation** of a surface of an object with respect to a global Cartesian frame of reference through its surface normal direction, as used by Nielsen [1988]. In our representation, this may be obtained through knowledge of the identities of the extremal points between which a surface lies. For example, the orientation vector of edge E1 for object B in Figure 1 has positive X and Y components since it lies between Vertices V1 and V2, the topmost and rightmost points.

Rectangular bounding boxes and bounding circles are powerful abstractions for qualitative spatial reasoning: numerically precise information is not required, even in a dynamically changing world, to maintain the bounding boxes and bounding circles. All we need are the inequality relationships between the extremal points of the bounding abstractions that are used. The use of circular/spherical abstractions for rotating objects ensures that only translational motion can change the position of the extremal points of the shape abstractions and thus, the relative positions of objects (assuming that objects have fixed shapes).

Orientation information, as required to determine the direction of current flow in an edge of a conducting loop, may be obtained if the connectivity of an object

and its extremal points are known. Only rotational motion can change the identities of the vertices and edges that form the extremal points of an object, and thus, its orientation.

Translational Motion

Our model of relative positions, based on inequality relations between the extremal points of objects, allows us to recognize that a qualitatively interesting event will occur when the inequality relation changes (e.g., when $rm(A) = lm(B)$ in Figure 3). This can be used to study the effects of translational motion.

The values of the rightmost and leftmost points of an object will change over time if the X-velocity of the object is non-zero, and the values of the topmost and bottommost points will be affected by the Y-velocity. In Figure 3, if the X-velocity of object A is positive, the coordinate value of its rightmost point will increase over time, and we can recognize that a qualitatively interesting state will occur once $rm(A) = lm(B)$.

Rotational Motion

To maintain the orientation of a rotating object, we must determine the changes in the identities of the vertices and edges that form its extremal points as it rotates. Currently, we only consider the rotation of two dimensional objects about an axis that passes through their center of gravity. Rotation can also change the projected area of an object onto another object in a different plane, a property of great interest for the generator problem described in Figure 2.

The projected area of an object rotating about the X or Y axis onto another object in the XY plane is a function of the area of overlap between the two objects (as if both were in the XY plane) and the cosine of the angle between the plane of the rotating object and the XY plane. We independently model the effects of rotation about the X and Y axes, and use qualitative addition to determine the net change in the projected area. We model the extent of rotation about the X and Y axes, respectively, through two angles measured counterclockwise from the XY plane to the plane of the object: θ_Y , the angle between the plane of the object and the Y axis, and θ_X , the angle between the plane of the object and the X axis.

When θ_X or θ_Y is equivalent to 90 or 270 degrees, the projected area of an object onto the XY plane is zero. These angles also mark the point at which the orientation of the object changes. During counterclockwise rotation about the X-axis, when θ_Y is in the interval 90 degrees to 270 degrees, the identities of the topmost and bottommost points will be reversed. During rotation about the Y-axis, the identities of the leftmost and rightmost points can be reversed. In Figure 3, considering object B, if the object has a value for θ_Y in the interval 90 to 270 degrees, then the orientation of edge E1 would have a positive X-component and a negative

Y-component since vertex V1 would be a bottommost point instead of a topmost point.

Rotation in the XY Plane Rotation about the Z-axis can only change the orientation of an object, and not the projected area onto the XY plane. Given, as in Figure 3, that vertex V1 is the topmost point of an object rotating in the counterclockwise direction, and that vertex V2 follows V1, V2 will be the next topmost point. The difficult problem is to determine which extremal point will change first. Given that V2 is the rightmost point and that V3 follows V2, the question now becomes will V2 become the topmost point before V3 becomes the rightmost point? Since the bottommost and leftmost points may also change, there is always a four-way race condition.

For a general polygon, we would require quantitative knowledge (e.g., the internal angles of the polygon) to avoid intractable branching during a qualitative simulation. However, for such commonly encountered special shapes as circles, rectangles and parallelograms, qualitative solutions do exist. For example, rectangles have the property that all four extremal points change simultaneously, eliminating the race condition.

Integrating Spatial and Dynamic Reasoning

We have implemented our spatial reasoning methods using the QPC [Crawford, et. al 90, Farquhar 94] qualitative modeling system, which in turn uses the QSIM [Kuipers 86] qualitative simulation system. We encountered two difficulties in integrating spatial and dynamic reasoning in the context of these systems - the need to model piecewise continuous variables, and the need to model variables with circular quantity spaces.

Modeling Piecewise Continuous Variables

Several qualitative modeling tools, such as QPC and QPE [Forbus 90], require model variables to be continuous. In problems involving both spatial and dynamic reasoning, it's possible for certain model variables to be piecewise-continuous. Figure 1 shows a conducting loop, whose abstracted shape is a rectangle, moving into a rectangular field. As the loop enters, the derivative of the area of overlap between the two objects is positive and is proportional to the X-velocity of the loop. Once the loop is enclosed in the field, the derivative of the area of overlap drops discontinuously to zero. This derivative is a *piecewise-continuous* variable - one whose value is continuous within a given world model, but which may change discontinuously at a transition point, when a new world model is computed to reflect a qualitatively significant change in the world state.

QPC inherits the values of model variables from the previous model when determining the new model after a transition point. Discontinuous changes in the values of variables could lead to contradictions in the

new model. In our implementation, we allow the QPC model builder to explicitly declare that a variable is piecewise-continuous within a QPC model fragment description, and prevent their values and any inequality relations involving such variables from being inherited at a transition point. This technique may be used for any dependent variable, since the new value for the variable may be recomputed in a subsequent model, or for any independent variable whose new value is explicitly given in the subsequent model. Variables that are dependent on piecewise continuous variables must also be modeled as piecewise continuous since their values could also change discontinuously.

Modeling Circular Quantity Spaces

A second modeling restriction imposed by the existing qualitative reasoning systems is that the range of values for a model variable is given through a *linear* quantity space. For modeling such phenomena as rotation, it is beneficial to model some variables through a circular quantity space to directly model the fact that the behavior of the system is cyclic.

Consider the variable θ_Y for the example of Figure 2. This measures the extent of rotation around the X-axis. The qualitatively significant values for θ_Y include 0, 90, 180, 270, and 360 degrees. We would like to directly state that the value of 360 degrees is functionally equivalent to zero degrees (resulting in a circular quantity space), instead of having to model, in a linear quantity space, that 450, 540, etc., are also significant values. We cannot model this fact in a linear quantity space since we would have that $(0 < 270)$, $(270 < 360)$, and $(0 = 360)$, which leads to a contradiction.

We use the circular-quantity-space declaration -

```
(circular-qspace-quantity
  theta-y convex-2-d-objects
  (theta-y-0 theta-y-90 theta-y-180 theta-y-270))
```

to define the variable theta-y as a quantity for the domain class convex-2-d-objects. It has a circular quantity space with four qualitatively significant values represented by theta-y-0, theta-y-90, theta-y-180, and theta-y-270. After theta-y-270, the quantity space wraps back to theta-y-0. Note that each of these values actually represents a set. We maintain the inequality relationships between θ_Y and each of θ_{Y0} , θ_{Y90} , θ_{Y180} , and θ_{Y270} to model the relative value of θ_Y . The same inequality relations, $(\theta_Y > \theta_{Y0})$ and $(\theta_Y < \theta_{Y90})$, are used to capture the fact that θ_Y is between 0 and 90 degrees or 360 and 450 degrees.

We avoid contradictions in the model by asserting only a linear subset of the circular quantity space into each QPC model. Assuming that each of the qualitatively significant values in a circular quantity space is associated with a model transition condition, in any QPC model, we only have to consider the relationship between the circular-valued variable (θ_Y) and three of those values. For example, in the initial state of the

problem described in Figure 2, we have that θ_Y is equal to θ_{Y0} , is less than θ_{Y90} , and is greater than θ_{Y270} . We may model θ_Y to be greater or less than θ_{Y180} , and not both, to linearize the circular quantity space since we know that θ_Y cannot equal θ_{Y180} without first crossing through θ_{Y90} or θ_{Y270} , when a model transition is required to occur.

We use the piecewise-continuous variable feature to model the fact that the inequality relationships between θ_Y and any of θ_{Y0} , θ_{Y90} , θ_{Y180} , and θ_{Y270} can change discontinuously at a transition point. The appropriate inequality relations are automatically asserted into the subsequent model. For example, in Figure 2, the loop is rotating in a counterclockwise direction about the X-axis and the next qualitatively significant change occurs when θ_Y equals θ_{Y90} . Then, we ignore all previous inequality relations between θ_Y and its possible values, and insert into the model that $(\theta_Y = \theta_{Y90})$, $(\theta_Y < \theta_{Y180})$, and $(\theta_Y > \theta_{Y0})$. The relationship between θ_Y and θ_{Y270} is insignificant, and may be modeled as ' $<$ ' or ' $>$ ' since the relationship in the previous model ($>$) will not have been inherited.

Qualitative Behavior for the Examples in Figures 1 and 2

To illustrate both solutions, we have artificially combined the problems in Figures 1 and 2 into a single problem: *As in Figure 1, the loop moves to the right until it is entirely enclosed in the magnetic field. At that point, its translational motion stops, and it rotates around the X-axis as in Figure 2.* This was done by treating translational and rotational velocity as piecewise continuous variables, and resetting their values once the loop was enclosed in the field.

During translational motion (denoted by time points t_0 , t_1 , and t_2), as the loop enters the field (t_1 to t_2), the area of overlap, and thus, the magnetic flux through the loop, increases steadily (steady due to the use of the rectangular bounding box abstraction for the conducting loop). During this period, an induced emf is established in the loop, which results in a steady current flow. The current flow is negative to show that it opposes the increase in the flux through the loop. The direction of current flow may be either clockwise or counterclockwise depending on the direction of the magnetic field, a property we ignore for the purposes of this example.

The variables $\text{DIF-}\langle\text{F.LM-XVAL}\rangle\text{-}\langle\text{L.RM-XVAL}\rangle$ and $\text{DIF-}\langle\text{F.LM-XVAL}\rangle\text{-}\langle\text{L.LM-XVAL}\rangle$ are generated by QPC to encode the inequality relations between the rightmost and leftmost points of the loop and the leftmost point of the field. Note that a qualitatively interesting change occurs in the *relative positions* of the loop and the field whenever one of the DIF terms is zero (i.e., an inequality relation changes). When the former DIF variable is zero, the loop begins to *enter* the field. When the latter is zero, the loop is fully *enclosed* in the field. Note also that the value of the derivative of the

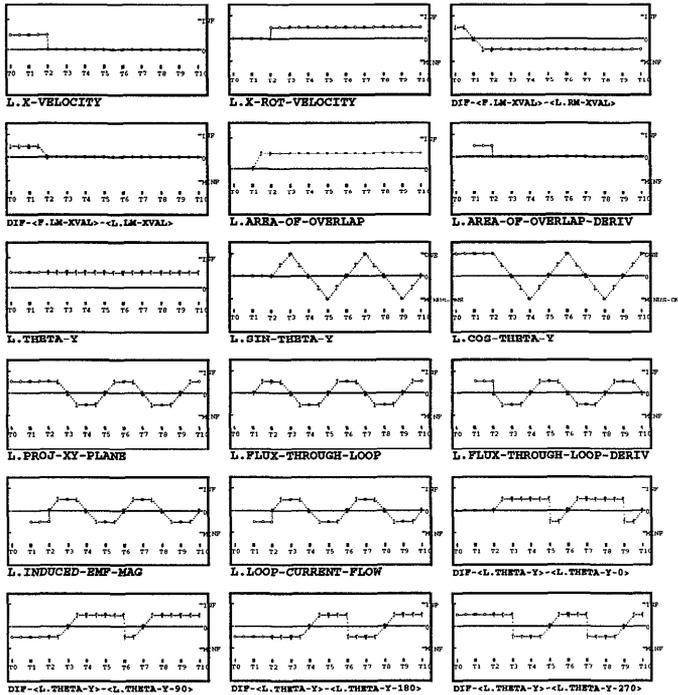


Figure 4: Qualitative behavior produced by QPC/-QSIM for the problems in Figure 1 (up to time t_2) and Figure 2 (after time t_2). The shape abstraction for the loop is assumed to be rectangular in the models for translational motion, and is assumed to be circular in the models for rotational motion. The values of `L.AREA-OF-OVERLAP` (constant after t_2) and `L.AREA-OF-OVERLAP-DERIV` (constant during t_1 - t_2) reflect these assumptions.

area of overlap between the loop and the field changes discontinuously from a positive value to zero at time t_2 . The variables dependent on this derivative, such as the derivative of the flux passing through the loop, the magnitude of the induced emf, and the current flow in the loop, also change discontinuously.

At time point t_2 , translational motion stops and rotational motion begins. We show the output for two complete, counterclockwise rotations of the loop. Each time interval after t_2 corresponds to an additional 90 degree increment in the value of θ_Y . `Theta-Y` has a positive value equivalent to zero in the initial state (I.e., the actual value may be 360, 720, etc.). It is continuously increasing after time t_2 since the rotational velocity of the loop is positive and constant.

The magnetic flux through the loop is proportional to the projected area of the loop onto the magnetic field, which is given by $Proj-Area = (Area-Loop)(\cos \theta_Y)$. The rate of change of flux through the loop, $d\Phi/dt$, is proportional to the rate of change of the projected area of the loop onto the field, which

is given by $dProj-Area = (Area-loop)(-\sin \theta_Y)(d\theta_Y)$. The induced emf in the loop is proportional to $-d\Phi/dt$.

The simulation output shows the sinusoidal behavior of the model variables as the loop rotates around the field. The last four DIF variables encode the inequality relationships between `theta-Y` and its possible values. The qualitatively interesting changes (in orientation) occur whenever one of the DIF terms crosses zero. Note that at time t_3 , when the value of `theta-Y` is equivalent to 90 degrees (the difference term `dif-<l.theta-y>-<l.theta-y-90>` is zero), the projected area of the loop onto the XY plane becomes zero. The projected area is then negative through time point t_5 , when `theta-Y` becomes equivalent to 270 degrees (the difference term `dif-<l.theta-y>-<l.theta-y-270>` is zero). This represents the period when the orientation of the loop is reversed due to the identities of the topmost and bottommost points being reversed.

Note the discontinuous changes in the values of the last four DIF variables. These occur when a newly asserted inequality relation (after a transition) between `theta-Y` and any of θ_{Y0} , θ_{Y90} , θ_{Y180} , and θ_{Y270} is different from that in the previous model. This demonstrates not only the operation of the circular quantity space mechanism, but also the usefulness of the ability to model piecewise continuous variables.

Summary and Conclusions

Forbus, Nielsen, and Faltings [1991] have conjectured that there exists no purely qualitative, general purpose, representation of spatial properties. Their conclusions are based on the need to reason about the *exact shapes* of objects in reasoning about mechanical devices, a spatial property that is indeed difficult to describe in a qualitative fashion.

Many qualitative formalisms have since been developed to support spatial reasoning [Mukerjee and Joe 90, Cui, Cohn, and Randell 92, Abella and Kender 93]. They share the property that *simplifying abstractions* are used to *approximate* the actual shapes of objects. The authors describe how abstractions such as the convex hull, collision parallelograms, or inertia tensor-based bounding boxes may be used for qualitative spatial reasoning, but do not describe how these abstractions can be represented internally in a qualitative fashion, or how to *qualitatively compute* the changes in the positions and orientations of the abstractions as the underlying objects translate or rotate.

We discussed the advantages and fundamental limitations of using shape abstractions for spatial reasoning, and presented an *extremal point-based* spatial representation that requires little information to compute and maintain even as objects translate and rotate. With this method, we can describe the relative positions of objects through inequality relations between their extremal points, and easily determine the effects of translational motion on the spatial state. We showed that the orientation of a convex object with respect to

a global Cartesian frame of reference could be determined given just the connectivity of the object and the identities of its extremal points, and discussed methods for maintaining the orientation if the object rotates about an axis passing through its center of gravity.

We are applying our spatial reasoning methods in a problem solver for the magnetic fields domain. The remaining work in this ongoing project will cover the use of orientation information to compute and reason about the effects of magnetic forces, and will address methods for modeling dynamically created spatial entities, such as magnetic fields due to current flow in a wire. We have also tested the utility of our spatial representation for use in natural language understanding, including reasoning about objects with intrinsic fronts, by considering problems involving the use of spatial reasoning to isolate objects in a complex scene to name them [Rajagopalan 94].

Acknowledgements

I thank Professor Benjamin Kuipers for his many suggestions in developing the ideas presented in this paper, and Bert Kay, James Lester, Toyooki Nishida, Sowmya Ramachandran, and Lalitha Rajagopalan for providing valuable comments on earlier drafts.

References

[Abella and Kender 93] Abella, A. and J. Kender. Qualitatively Describing Objects Using Spatial Prepositions. In: AAAI-93, Washington, DC, 1993.

[Crawford, et. al. 90] Crawford, J., A. Farquhar, and B. Kuipers. QPC : A Compiler from Physical Models into Qualitative Differential Equations. In: AAAI-90, Boston, MA, 1990.

[Crawford and Kuipers 91] Crawford, J. and B. Kuipers. Algernon - A Tractable System for Knowledge-Representation. AAAI Spring Symposium on Implemented Knowledge Representation and Reasoning Systems, Palo Alto, CA 1991.

[Cui, et. al. 92] Cui, Z., A. G. Cohn, and D. Randell. Qualitative Simulation Based on a Logical Formalism of Space and Time. In: AAAI-92, San Jose, CA, 1992.

[Egenhofer and Al-Taha 92] Egenhofer, M. and K. Al-Taha. Reasoning about Gradual Changes in Topological Relationships. In: A. Frank, I. Campari, and U. Formentini, Eds., *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, Springer-Verlag, Berlin, 1992.

[Farquhar 94] Farquhar, A. A Qualitative Process Compiler. In: AAAI-94, Seattle, WA, 1994.

[Forbus, et. al. 91] Forbus, K., P. Nielsen, and B. Faltings. Qualitative Spatial Reasoning: the CLOCK Project. *Artificial Intelligence* 51, 1991.

[Forbus 90] Forbus, K. The Qualitative Process Engine. In Weld, D., and de Kleer, J., eds., *Readings in Qualitative Reasoning about Physical Systems*. Morgan Kaufman, Los Altos, CA, 1990.

[Freksa 92] Freksa, C., Using Orientation Information for Qualitative Spatial Reasoning. In: A. Frank, I. Campari, and U. Formentini, Eds., *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*, Springer-Verlag, Berlin, 1992.

[Galton 93] Galton, A., Towards an Integrated Logic of Space, Time, and Motion. In: IJCAI-93, Chambery, France, 1993.

[Joskowicz and Sacks 91] Joskowicz, L., and E. Sacks. Computational Kinematics. *Artificial Intelligence* 51, 381-416, 1991.

[Kuipers 86] Kuipers, B. Qualitative Simulation. *Artificial Intelligence* 29, 289-338, 1986.

[Kuipers and Levitt 88] Kuipers B., and T. Levitt. Navigation and Mapping in Large-Scale Space. *AI Magazine*, 9(2), 25-43.

[Mukerjee and Joe 90] Mukerjee, A. and G. Joe. A Qualitative Model for Space. In: AAAI-90, Boston, MA, 1990.

[Nielsen 88] Nielsen, P. A Qualitative Approach to Mechanical Constraint. In: AAAI-88, Saint Paul, MN, 1988.

[Rajagopalan 93] Rajagopalan, R. A Model of Spatial Position Based on Extremal Points. In Proceedings: ACM Workshop on Advances in Geographic Information Systems, Arlington, VA, 1993.

[Rajagopalan 94] Rajagopalan, R. Integrating Text and Graphical Input to a Knowledge Base. In Working Notes: AAAI Workshop on Integration of Natural Language and Vision Processing, Seattle, WA, 1994.

[Rajagopalan and Kuipers 94] Rajagopalan, R. and B. Kuipers. The Figure Understander: A System for Integrating Text and Diagram Input to a Knowledge Base. In Proceedings: The Seventh International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems (IEA/AIE-94), Austin, TX, 1994.

[Resnick and Halliday 88] Resnick J. and D. Halliday. *Fundamentals of Physics*, 3rd. Edition, John Wiley and Sons, 1988.

[Retz-Schmidt 88] Retz-Schmidt, G. Various Views on Spatial Prepositions. *AI magazine*, 9(2).

[Weinberg, et. al. 92] Weinberg, J., S. Uckun, G. Biswas, and S. Manganaris. Qualitative Vector Algebra. In: B. Faltings and P. Struss, Eds., *Recent Advances in Qualitative Physics*, pp. 177-192, The MIT Press, Cambridge, MA, 1992.