Northwestern University
The Institute for the Learning Sciences

Qualitative Spatial Reasoning: The CLOCK Project

Technical Report # 9 • February, 1991

Kenneth D. Forbus
Paul Nielsen
Boi Faltings

1851

Established in 1989 with the support of The Arthur Andersen Worldwide Organization

Qualitative Spatial Reasoning: The CLOCK Project

Kenneth D. Forbus
Paul Nielsen
Boi Faltings

February, 1991

The Institute for the Learning Sciences
Northwestern University
Evanston, IL 60201

This research was supported by the Office of Naval Research, Contract No. N00014-85-K-0225. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of The Arthur Andersen Worldwide Organization. The Institute receives additional funding from Ameritech (an Institute Partner) and IBM.

1 Introduction

Spatial reasoning is a central problem in formalizing commonsense knowledge. People routinely solve sophisticated problems involving space, such as designing buildings, navigation, assembling a car from parts, and packing suitcases. An important subclass of spatial reasoning involves reasoning about motion. Human engineers design sophisticated mechanisms, many sports involve projectile motion, and everyone has some rough idea about whether a falling cup of coffee requires action on their part, based on a quick glance at its motion. Much of the reasoning we do about motion does not appear to use the formalisms of mechanics we are taught in school. For example, someone who has never taken a physics course knows that two balls thrown into a well can collide, but if one ball is always in the well and the other always outside they cannot.

Understanding spatial reasoning brings us closer to an explicit understanding of commonsense knowledge, which is important both for understanding how people work and for making our machines smarter. The qualitative understanding of a problem invariably precedes a quantitative understanding and serves as an engineer's intuition. Indeed, often it is at this level that most of the difficult engineering work is done. Designs are proposed and eliminated based on crude sketches. Only when an design seems reasonable are mathematical tools employed to verify and refine it.

Qualitative physics is the area of AI which focuses on formalizing and using commonsense knowledge about the physical world. Most work in qualitative physics has focused on *qualitative dynamics*, the representation and organization of qualitative time-varying differential equations (cf. [5,18,40,66]). By contrast, qualitative spatial reasoning has received much less attention. Examples of spatial reasoning problems central to qualitative physics are reasoning about motion, the geometry of liquid flow, and the shape of charge distributions. Problems which overlap with the concerns of qualitative physics, but involve significant content from other domains, include planning assemblies and navigation.

This paper presents a theoretical framework for qualitative spatial reasoning. The framework is organized around three ideas:

1. *The poverty conjecture*: We claim there is no purely qualitative, general-purpose, representation of spatial properties. That is, while qualitative descriptions are useful in spatial reasoning, they are not sufficient to describe a situation in a task-independent and problem-independent fashion.

2. The *MD/PV model*: Qualitative spatial reasoning requires two representations: a *metric diagram*, a mixed symbolic-quantitative representation, which serves as an oracle for a class of spatial queries, and a ***place vocabulary*** which provides relevant quantizations of shape and space according to the particular task. Importantly, the place vocabulary is computed from the metric diagram. This overcomes the limitation suggested by the poverty conjecture, and allows the metric diagram to serve as a communication media between qualitative and quantitative representations of space and shape.

3. *The Connectivity/Shape Hypothesis*: The appropriate notion of state for qualitative kinematic reasoning concerns ***connectivity*** and ***shape***, since changes in connection usually determine when forces change, while shape (along with position) determines connectivity.

The next section (Section 2) develops these ideas in detail. This paper focuses on *qualitative kinematics* (QK), the geometric aspects of representing and reasoning about motion. However, we argue below that the same concerns are relevant to all forms of qualitative spatial reasoning. This section also describes a set of ***basic inferences*** for qualitative spatial reasoning which serve as a basis for organizing theories and algorithms. Section 3 then illustrates the utility of the MD/PV model by describing `CLOCK`, a program which reasons about mechanisms, including mechanical clocks (c.f. Figure 1). Section 4 analyzes other relevant research in terms of this framework. Finally, Section 5 outlines some questions this framework suggests for further research.

Figure 1: The QRG clock.

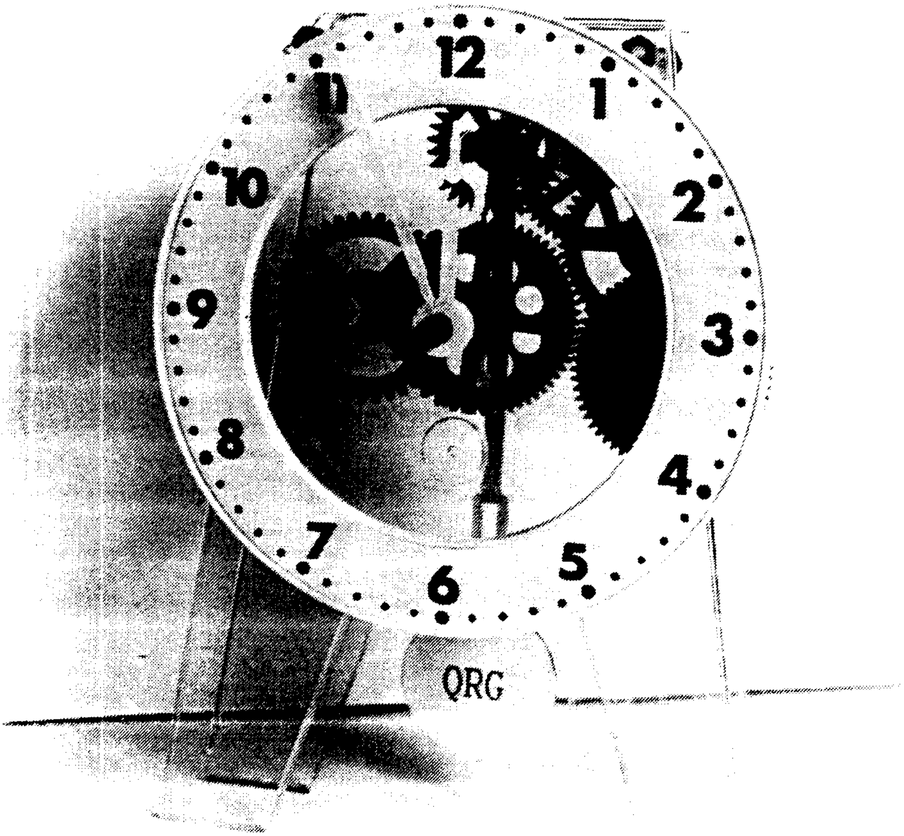

2 A framework for qualitative spatial reasoning

Motion pervades the physical world — things roll, swing, fly, gyrate, spin, slide, push, and collide. The breadth of the phenomena and wide variation in the kinds of answers we desire argues against a single representation for all of qualitative kinematics, much less all of qualitative spatial reasoning. Still, we believe there are important underlying constraints which place important limits on the search for solutions to specific spatial reasoning problems.

This section begins by laying out some specific desiderata that any account of qualitative spatial reasoning should be measured against. Then we argue for the *poverty conjecture*, which concerns the kinds of qualitative representations there can be and fundamental limits on them. Our answer to the negative claim of this conjecture is the *Metric Diagram/Place Vocabulary* model of spatial reasoning, which uses a mixture of qualitative and quantitative representations to overcome the limitations of each. A specific hypothesis about the nature of qualitative kinematics, the *connectivity/shape hypothesis*, is made next. Finally, the section closes with a description of a set of basic inferences for qualitative kinematics.

2.1 Desiderata

2.1.1 Minimality

A spatial representation should require no more information than is strictly necessary. Quantitative simulation may provide sufficient information to predict the behavior of objects with known mass, speed, direction, and location, but often we are faced with situations where these quantities are incompletely known. Our programs should be capable of making intelligent decisions even in the absence of this knowledge. For example, a robot should not have to determine the actual weight of a cup of coffee in order to predict that if dropped it will fall.

Even if good information were available, time and memory constraints may make the computation of simple answers by numerical techniques intractable. For example, to determine whether a sheet of paper can rest stably on a desktop potentially requires taking into account complex shape deformations and its interaction with hundreds of other papers, books, magazines, and office supplies which also occupy the desktop.

2.1.2 Composability

A spatial representation should be able to handle novel shapes and interactions. While there are advantages to maintaining libraries of common objects, such libraries are not sufficient for dealing with the infinite variations found in the real world. Consider a CAD system which used a fixed, finite set of symbols to describe the standardized types and sizes of nuts and bolts, say, along with rules that state which combinations of these parts can be threaded onto each other. This representation only captures a small subset of how these parts can interact. For instance, one bolt could be used as a spacer to prevent a nut on another bolt from being tightened all the way down. Moreover, this representation is useless for understanding how these parts would interact with new components: a washer, say, or figuring out what stable resting positions exist for a bolt-nut combination on a flat or corrugated surface. All such representation schemes can only describe a limited set of shapes, and overly restricts the kinds of motions that can be considered. Labeling an object as a gear, for instance, is a clear violation of the *No function in structure* principle [6], since a gear-shaped object can be used in many ways. For instance, gear-like objects are used as water wheels, circular saws, paddles for boats, or scape wheels in clocks. In addition, the sides of a gear might serve as a spacer to keep objects a fixed distance apart.

In this paper we show that qualitative abstractions of shapes appropriate for reasoning about motion can be extracted automatically from a generative language of input shapes.

2.1.3 Explanation

The spatial representation must support explanation generation. Explanations describe how a property arises due to other factors in the situation. Thus explanations can be used for credit assignment in design and diagnosis. Explanations are often causal, but they need not be. A constraint argument describing why two blocks do not move due to mutual contact, for example, still provides the information needed to change the situation to make them move. Traditional numerical simulation systems do not provide such abilities. Their output is a set of state parameters varying over time, perhaps displayed graphically, which must then be interpreted by some other system (or person). Worse yet, the dependence of different aspects of a behavior on particular features of a situation is totally opaque. By contrast, a qualitative description of behavior should identify behavioral regimes which represent important classes of system behavior, and identify which aspects of the situation they depend on.

In addition to being able to efficiently tease apart the rationale for conclusions, it is also desirable to be able to support "natural" explanations. That is, given the necessity of explaining its results to human beings, it should be possible to use the representation to generate intuitively plausible explanations.

In this paper we show that the coarse granularity of qualitative descriptions allows more explicit reasoning about motion, and hence provides the ability to make better explanations.

2.1.4 Integration

The spatial representation must facilitate the integration of different kinds of knowledge. Reasoning about motion, for example, requires combining information about forces with shape. Many functional constraints are often expressed spatially, such as the range over which a hinge must swing.

In this paper we show how a simple qualitative vector algebra can be used to represent this and other information.

2.2 The poverty conjecture

In retrospect, the success of qualitative dynamics is surprising. To perform a numerical simulation requires identifying a large set of precise numerical values and equations. While numerical simulations often give precise answers, modeling complex systems often requires prodigious amounts of computation (e.g., consider the current interest in supercomputers). By contrast, qualitative dynamics can often provide insights into the kinds of behavior possible to a system with only a smattering of information about inequalities. Can similarly powerful representations for qualitative spatial reasoning be found?

We claim the answer is no. Specifically, we make the following conjecture:

> *poverty conjecture: There is no purely qualitative, general-purpose representation of spatial properties.*

The poverty conjecture is subtle and requires some elaboration. We are not claiming that no useful qualitative representation of space exists. Indeed, the opposite is true. We believe qualitative representations are fundamental to capturing the flexibility of commonsense spatial reasoning. But such representations ***by themselves*** are inadequate: They must reference quantitative information for many classes of predictions. Furthermore, we claim that qualitative representations of space and shape must be computed in a task-specific, and sometimes a problem-specific, manner.

A specific application of the poverty conjecture is the problem of finding a qualitative representation for the shape of an object which supports predicting how it will behave when it is installed as part of a mechanism. We claim that this cannot be done independently of the specifics of the mechanism. The existence of problem-independent representations of shape which include qualitative information along with quantitative information is well-known (e.g., CAD descriptions). Indeed as shown below, such representations provide a basis for solving this problem. However, we claim that the quantitative component is a necessity; less information will not do.

To see this, consider the ***rolling problem***: Given two objects, can one smoothly roll across the other? For prototypical cases little information is needed: A ball can roll across a table, and if two meshing gears are aligned properly then one can roll across the other. But a general-purpose reasoning system cannot rely solely on prototypes. To provide generativity it must at least have the ability to compose prototypes, and preferably provide the ability to generate new shapes from surface or volume primitives. And here is where purely qualitative representations fail. Without some metric information as to the relative sizes and positions of the parts of a compound surface, the rolling problem cannot be solved. Consider for example two wheels, one with a bump on it and the other with a notch carved out of it. Without more details one cannot say how smoothly they will travel across each other: Both perturbations of the shape could be trivial, or the notch might include sharp corners that cause the bump to catch. Figure 2 illustrates. A seductive approach is to state that the shapes are complementary and their sizes are identical, but this solves only a single class of cases.

The problem becomes even worse when one considers objects which are not always in contact. To understand the behavior of a mechanical clock, for example, requires knowing how contact relationships between the parts of its escapement can change. Without knowing the relative sizes and placements of every part of each object which may come into contact, one has no hope of answering this question. But without an initial quantitative representation, there is insufficient information to compute relative sizes and positions.

The centrality of connectivity in spatial reasoning is why we believe that simple qualitative representations of an object's shape, generated in isolation from the specific context in which it is used, can never be adequate. It is the interactions between objects which determines their kinematic behavior, not just the objects themselves. For example, gears only exhibit rotation transfer when their teeth mesh properly with certain other gears. The same gear, on a flat surface, would behave as a wheel.

What about using crude approximations of sizes, for instance by describing some holes a priori as being "small"? Again, this judgment cannot be safely made in isolation from the other objects it will be in contact with. For example, a small hole in an otherwise flat surface will typically be insignificant, but it may interact with a peg. (Anyone who has tried to take apart a computer or mechanical object and was temporarily flummoxed by a small flange catching has experienced this problem.) Since detailed distinctions between shapes may be crucial, an a priori qualitative rendering of shape cannot be generally adequate.

Finally, even if we restrict ourselves to reasoning about a small primitive class of objects, we still cannot escape the need for detailed shape representations. There are arbitrary possible perturbations in the separation between objects' surfaces that will affect their behavior. For example, some gears will jam if they are placed too close together even if the teeth are aligned and they mesh properly.

Figure 2: Some examples of the rolling problem

Consider two wheels, one with a bump and the other with a notch. Can they roll smoothly along each other? The answer depends on detailed metrical information, and thus cannot be answered if the initial representation of each individual shape is purely qualitative. For each row, suppose that the object in column A can roll smoothly when properly aligned with the corresponding object in column B. In both cases, A and C still will not roll smoothly. In the first row, even though the notch in C is larger than the notch in B, its shape is such that the bump in A won't fit. In the second row, the notch in C is exactly the same as the notch in B, but the diameter of C is larger. This means that even if the bump and notch are initially aligned, the next time the bump contacts C it will not hit the notch.

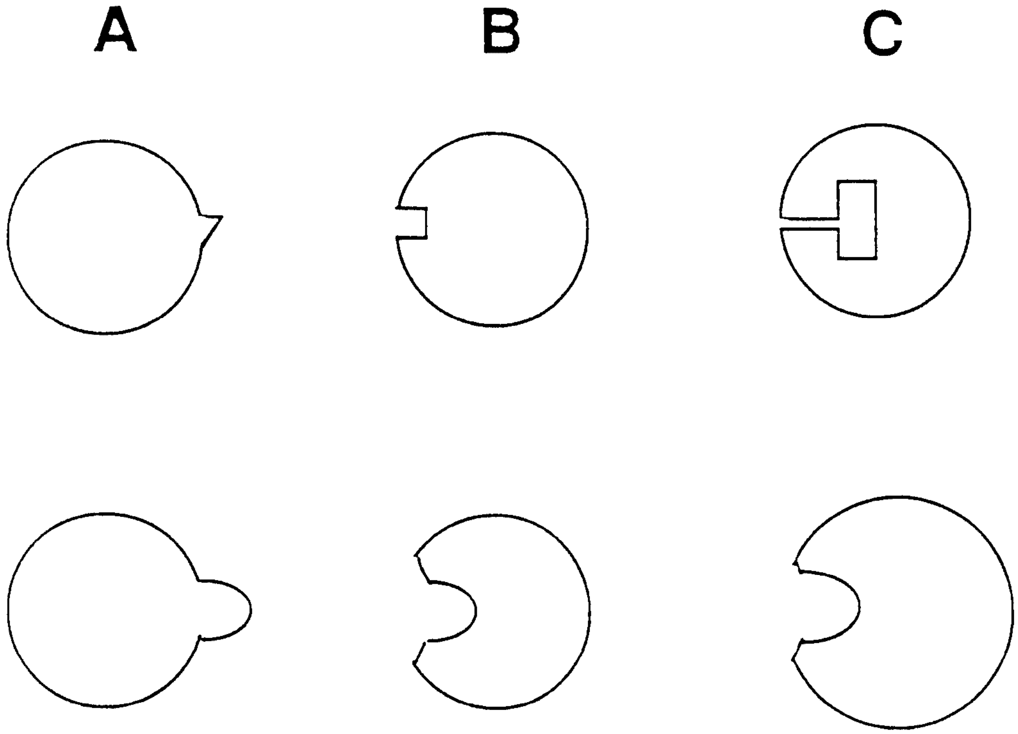

It is difficult to make "purely qualitative" precise because there is a spectrum of possible representations. Clearly a representation which includes elements of $\Re$ as constituents is not purely qualitative. Analytic functions are also not permitted in definitions of qualitative values, since they do not correspond to observable "qualities" of the modeled system. On the other hand, representing the relationship between gears by a `Mesh` predicate, or representing a $2D$ boundary by a list of segments described only as `concave`, `convex`, or `straight` are exactly the kind of temptingly deceptive representations we argue against. In these purely qualitative shape representations, symbols represent entire classes of shapes, and numeric dimensions and relationships are abstracted away. But in spatial reasoning, very fine distinctions between the shapes are often crucial.

Considered in the light of qualitative dynamics, these features of spatial reasoning may seem surprising. Suppose the poverty conjecture were true of qualitative dynamics. To generate the input to a qualitative simulator, we would first have to include numerical parameters for the properties of the objects in a system and a wealth of detail about the relationships between them. The fact that we can calculate the classes of behaviors of a system in qualitative dynamics with far less information indicates that it is in some ways a fundamentally simpler problem than qualitative spatial reasoning.

The poverty conjecture does not rule out the existence of useful problem-specific qualitative representations of spatial properties. For instance, one might summarize a piece of a mechanism by stating

$$\mathtt{Gear(x)} \land \mathtt{Gear(y)} \land \mathtt{Mesh(x,y)}$$

However, such representations are extremely limited. One often sees predicates like this in initial attempts to formalize reasoning about mechanisms, but rarely are their consequences explored in detail. For example, suppose we are given two new gears. What can we say about them which enables us to conclude whether or not they mesh? Are tooth spacing and pitch irrelevant? Is the `Mesh` relation transitive?[1] Such high-level descriptions are seductive, but it is difficult to actually infer much from them beyond what was initially stated. (Empirically, this is a good indication that the "No function in structure" principle has been violated.)

One motivation for the poverty conjecture is the simple fact that we (and others) have tried for years to find satisfactory, purely qualitative representations for shape and space and have failed. This by itself of course means nothing, since one might be discovered tomorrow. Given the argument above, we believe this is extremely unlikely. Furthermore, there are two additional arguments for the poverty conjecture. First, people appear to require more than qualitative information in spatial reasoning, and second, the combinatorics of connectivity in higher dimensions suggests that simple, local representations do not provide enough constraint to support powerful reasoning. We examine each argument in turn.

2.2.1 The argument from human performance

Much of the motivation for tackling information-processing problems comes from reflection on human skills and capabilities. While it certainly might be useful for a machine to be able to see through walls, for instance, since people do not have this ability we know it isn't necessary to be intelligent. And while some solutions exist for seeing through walls (e.g., X-rays), we have ample evidence that people do not commonly have this ability. Similarly, even though it would be nice to have "stand-alone" qualitative spatial representations, psychological evidence suggests that people do not have them, and hence one can do qualitative spatial reasoning without them.

There is ample evidence that people resort to diagrams or models for all but the simplest spatial problems [30,38]. This includes situations where the initial description of the problem is quite vague on its spatial aspects (such as a textbook physics problem [52]). Generating a diagram from such descriptions involves choosing particular values for unknown geometric parameters which satisfy

---

[1] `Mesh` is not transitive for real gears, unless the tooth width and spacing are equal. Consider a sequence of gears, each with a tooth size just slightly smaller than the next gear. If the gear sequence is long enough and there is enough play in each pair of gears, then each pair of gears in the sequence will mesh, while the first and last gear in the sequence will not.

all the significant geometric constraints on the system. This is not always easy, as anyone who has ever tried to draw a detailed map can attest. The fact that people are willing to go to such trouble to generate diagrams suggests that they are getting some strong benefit from them. Thus it seems likely that our fluency in spatial reasoning does not spring from a set of very clever axioms for handling purely relational descriptions of space.

On the other hand, people use a rich vocabulary of qualitative descriptions, along with diagrams, when describing mechanisms and other spatial situations (as can be ascertained by examining any engineering text). This is one of the motivations for the MD/PV model, described below.

2.2.2 The argument from mathematics

The power of a qualitative representation is the ability to combine weak relationships between its elements to draw interesting conclusions. In qualitative dynamics weak representations of time-varying differential equations suffice for a broad spectrum of inferences. The secret lies in the fact that numbers have a total ordering. This limits the amount of information needed to constrain a quantity's value. With a small amount of inequality information we can provide discrete ranges of values that serve to mark important behavioral differences, and can predict state transitions and provide useful notions of continuity [6,18,40,66]. Allen's temporal logic [1] is another example of a system of relationships which individually are weak but together provide enormous constraint. Unfortunately, similarly weak qualitative spatial representations are virtually useless.

Both Allen's temporal logic and quantity spaces crucially rely on transitivity for their inferential power, which in turn relies on the existence of an underlying total order. It is instructive to try and develop a transitivity table for spatial relations between two dimensional figures, by analogy with Allen's transitivity table for temporal relations. One might for example imagine a vocabulary that included `Equal`, `Inside`, `Abut` (two figures share a side), and `Overlap` (two figures have a two-dimensional intersection). A few minutes of exploration suffices to confirm that the only entries in this table which provide significant constraint are those which impose a partial order (e.g., `Equal` and `Inside`). Combining `Abut` and `Overlap` relations, for example, yields almost no information.

For some tasks, one-dimensional parameterizations can be found which provide useful power in spatial domains. For example, Mukerjee and Joe [53] impose axes on a two-dimensional map, and define an Allen-style transitivity table for spatial relationships by exploiting this reduction in dimensionality. However, it seems unlikely that such inference schemes will be useful for tasks which require full higher-dimensional manipulations. Suppose we have a purely qualitative representation which can solve the rolling problem defined above. Could this representation support qualitative kinematic reasoning? To do so, it must be support at least the following kind of inference: Given an object $\mathcal{OB}$, predict which combinations of contact between it and other objects are possible. In spaces with more than a single dimension, an object can be in simultaneous contact with any number of other objects. This can lead to complex freedoms of motion. For example, a plane polygon which has two simultaneous points of contact with other polygons can slide along a 4th-degree algebraic curve while maintaining these contacts. In fact, motion constraints of arbitrarily high algebraic degree can result from combinations of objects which are in simultaneous contact among each other.

One implication of this fact is that our hypothesized purely qualitative representation cannot consist solely of descriptions of pairwise contacts, such as between a peg and a hole or between wheels which roll against each other. If it did, there would be no way to figure out which combinations of pairwise contacts involving $\mathcal{OB}$ and two different objects were mutually consistent. Instead, the qualitative representation of $\mathcal{OB}$'s shape would have to include qualitative models of contact for all *combinations* of objects in the world model. But this violates the composability desideratum, since we have encoded our desired conclusion in the initial representation. Next, we show how the MD/PV model overcomes this problem by ***computing*** a tailor-made qualitative model for the particular problem, using the precise quantitative information about the shapes of objects provided by the metric diagram.

2.3 The MD/PV model

We believe that qualitative spatial reasoning requires a combination of quantitative and qualitative representations. We call this the *MD/PV* model because it has two parts:

- A ***metric diagram:*** a combination of symbolic and quantitative information used as an oracle for a class of spatial questions.
- A ***place vocabulary:*** a purely symbolic description of shape and space, grounded in the metric diagram.

A reasoner starts with a metric diagram, which is intended to serve the same role that diagrams and models play for people. The metric diagram is used to ***compute*** the place vocabulary, thus ensuring the qualitative representation is relevant to the desired reasoning.

The particular form of these representations varies with the class of problem and architecture. The quantitative component of the metric diagram could be floating point numbers, algebraic expressions, or bitmaps. The place vocabulary can be regions of free space, configuration space, or something else entirely. The key features are that (a) the place vocabulary exists and (b) it is computed from a metric representation. These features mean that we can still draw some conclusions even when little information is known (by using the place vocabulary as a substrate for qualitative spatial reasoning) and that we can assimilate new quantitative information (such as numerical simulations or perception) into the qualitative representation. (The `CLOCK` system provides one example; others are discussed in Section 4.)

Let us consider human spatial reasoning for a moment. Why are diagrams useful? The marks that represent the geometric aspects of a problem in a diagram have a fixed location and size. Their arrangement on paper models the spatial relations between the things they represent. This property allows our visual apparatus to interpret these relationships as we would those of the real objects. We do not yet understand the complexities of human vision, but there are other ways to encode the spatial structure of a diagram for use by a program.

What does one do with a diagram? Usually we decompose space and shapes into distinct regions, according to variations in the kinds of interesting things that can happen in each piece. There are several different classes of problems that we consider to involve spatial reasoning, including navigation, knot tying, and motion problems. We claim that the most important factor these problems share is a notion of ***place***. By place, we mean a piece of space (point, line, surface, region, etc.) such that all parts of it share some common property. Qualitative reasoning about space involves the use of a vocabulary of places, whose interconnections and relationships are specified symbolically.

2.4 The Connectivity/Shape Hypothesis

We claim that the notion of state in qualitative kinematics is organized around ***connectivity*** and ***shape***. Connectivity is important because contact (of some kind) is required for one object to affect another's behavior. Shape is important because it determines the possible connections.

The kinematic state of a system is primarily the collection of connectivity relationships that hold between its parts. Changes in connectivity signal changes in state. For example, a ratchet is clearly in a different state when the pin is on top of a tooth than when jammed in a corner.

A system's ***mechanical state*** is the union of its kinematic and dynamic state. The dynamical component can be represented in many ways, including qualitative state vectors [15,16,55] and Qualitative Process (QP) theory [18]. The particular vocabularies for connectivity and shape will of course be domain-dependent. The notion of mechanical state and transitions involving mechanical state in CLOCK are detailed in Section 3.5.2.

2.5 Basic Inferences in Qualitative Kinematics

A particularly important problem in spatial reasoning is reasoning about motion. Any system which reasons about motion must be able to determine what motions are possible and what effects these motions lead to. The key to progress in qualitative dynamics was finding appropriate notions of state and state transitions. The use of connectivity for kinematic state suggests a similar set of basic inferences for qualitative kinematics which can be combined for more complex reasoning.

1. *Finding potential connectivity relationships:* Computing the place vocabulary from the metric diagram must yield the connectivity relationships that will be the primary constituents of kinematic state. In the CLOCK system, for instance, this corresponds to finding consistent pairwise contacts.

2. *Finding kinematic states:* The constituent connectivity relationships must be consistently combined to form full kinematic states. Although typically quantitative information will still be required (being able to calculate relative positions and sizes is essential), we claim the resulting symbolic description can suffice for the remaining inferences.

3. *Finding mechanical states:* By imposing dynamical information (i.e., forces and motions) complete mechanical states are formed. The key to this inference is identifying qualitative reference frames and the ways in which objects are free to move.

4. *Finding state transitions:* Motion can eventually lead to change in connectivity, providing kinematic state transitions. Dynamical state transitions are also possible (pendulums exhausting their kinetic energy, for instance) as well as combinations of kinematic and dynamical transitions.

These operations are analogous to the basic dynamical inferences of Qualitative Process theory. In particular, the QP analog to finding potential connectivity relationships is finding potential process and view instances. The QP analog to finding kinematic states is finding process and view structures, and the analog to finding mechanical states is resolving influences. Finally, the QP analog to finding state transitions is limit analysis.

These operations can be exploited in a variety of ways. For example, Section 3 uses them to produce envisionments. One can imagine others. For example, these operations could be used to generate histories which would then be subjected to a comparative analysis [70], to determine how to adjust a pendulum in order to speed up a tardy clock.

3 Qualitative Reasoning about Mechanisms

The task of understanding the behavior of mechanisms involves a variety of hard spatial reasoning problems. One might expect that, given the long history of engineering interest in mechanisms, that numerical simulation of such systems would be more or less a solved problem. Unfortunately, this is not the case. The numerical simulation of complex mechanical systems is still an area of active research. Existing engineering CAD programs only analyze systems with pre-calculated kinematics. They are unable to reason with incomplete information, have no notion of functionality, and lack understanding of basic physical principles [49]. Mechanical systems where connections change over time (such as mechanical clocks) have proven particularly hard to model. Even if general-purpose, robust numerical simulators for arbitrary mechanisms are perfected, many engineering tasks still require qualitative representations. Numerical simulators only produce a single path through the space of behaviors from a given, precisely described starting configuration. This makes them unsuitable for conceptual design, where such details have not yet been worked out. Even when applied to later stages of analysis, numerical simulations can overlook potentially interesting (and problematic) behaviors, such as gears jamming, if the choice of initial parameters isn't appropriate. Sometimes such behaviors may be found by searching through the space of numerical parameters, of course. But this is likely to be much less efficient than searching through the highly abstracted description of behaviors provided by qualitative representations, which capture entire classes of specific behaviors within a much smaller symbolic vocabulary.

While we have not solved the general problem of mechanism understanding, we do have an account which is capable of predicting the possible behaviors of fixed-axis mechanisms which can be decomposed into two-dimensional interactions. This class of mechanisms is very broad, including for example most clockwork mechanisms. Basically, we use the MD/PV model to reason about mechanism behavior. The metric diagram is used to compute a place vocabulary consisting of regions of *configuration space*, the space spanned by the set of possible motion parameters of the parts. These places provide the basis for reasoning about dynamics. They furnish a substrate for the propagation of forces and provide connectivity information required to calculate the kinematic aspects of state transitions. We have tested these ideas by implementing them in a computer program, called CLOCK. CLOCK has been successfully tested on a variety of examples, including the QRG clock in Figure 1.

This section describes our account of reasoning about mechanisms, and shows how these ideas are implemented in CLOCK. Section 3.1 begins with an overview of CLOCK, describing its inputs, outputs, and overall operation. Section 3.2 describes CLOCK's metric diagram. Section 3.3 summarizes the theory of qualitative kinematics used in CLOCK. (The details of this theory can be found in [9].) Empirically, we found an important problem was controlling the complexity of the place vocabulary, and developed abstraction techniques to make the resulting place vocabulary manageable. These techniques are described in Section 3.4. Section 3.5 outlines Nielsen's theory of *qualitative mechanics*, which describes qualitative representations of vectors, mechanical constraint, and motion. We also note the simplifying assumptions regarding dynamics used in CLOCK. Section 3.6 sketches CLOCK's algorithms, while Section 3.7 illustrates its operation. Finally, Section 3.8 discusses some specific lessons learned from building CLOCK.

3.1 Overview of CLOCK

CLOCK takes as input a specification of the parts of a mechanism, described quantitatively using a CAD-like representation language. In addition, each part is annotated with a description of what external forces may be acting on it. For example, gravity on a pendulum is represented by an external force which pulls clockwise for some angles and counterclockwise for others. (Section 3.7 provides more details about `CLOCK`'s input.) `CLOCK` produced as output a total envisionment, describing all the possible states of the mechanism relative to the information given, and all the possible transitions between them. This envisionment is useful because several global features of the mechanism's behavior can be detected by further calculations on it (c.f. Section 3.7.1).

`CLOCK`'s computations are organized using the basic inferences of Section 2.5. Roughly, `CLOCK` works like this. Potential connectivity relationships are found by first computing a configuration space for each kinematic pair in the mechanism's parts. The free regions of each configuration space are quantized to form its place vocabulary, abstracting as necessary to keep the size of the vocabulary manageable. Kinematic states are then described as consistent combinations of these pairwise places. The possible mechanical states are computed by analyzing the dynamical properties of each kinematic state, ascertaining what external forces may be acting and propagating them. Finally, state transitions are computed for each mechanical state by detecting changes in motion due to dynamics and by using the place vocabulary to ascertain how connectivity can change due to motion. The information flow and representations used are illustrated in Figure 3.

Figure 3: Information flow of CLOCK
This diagram illustrates the representations and information flow used in CLOCK.

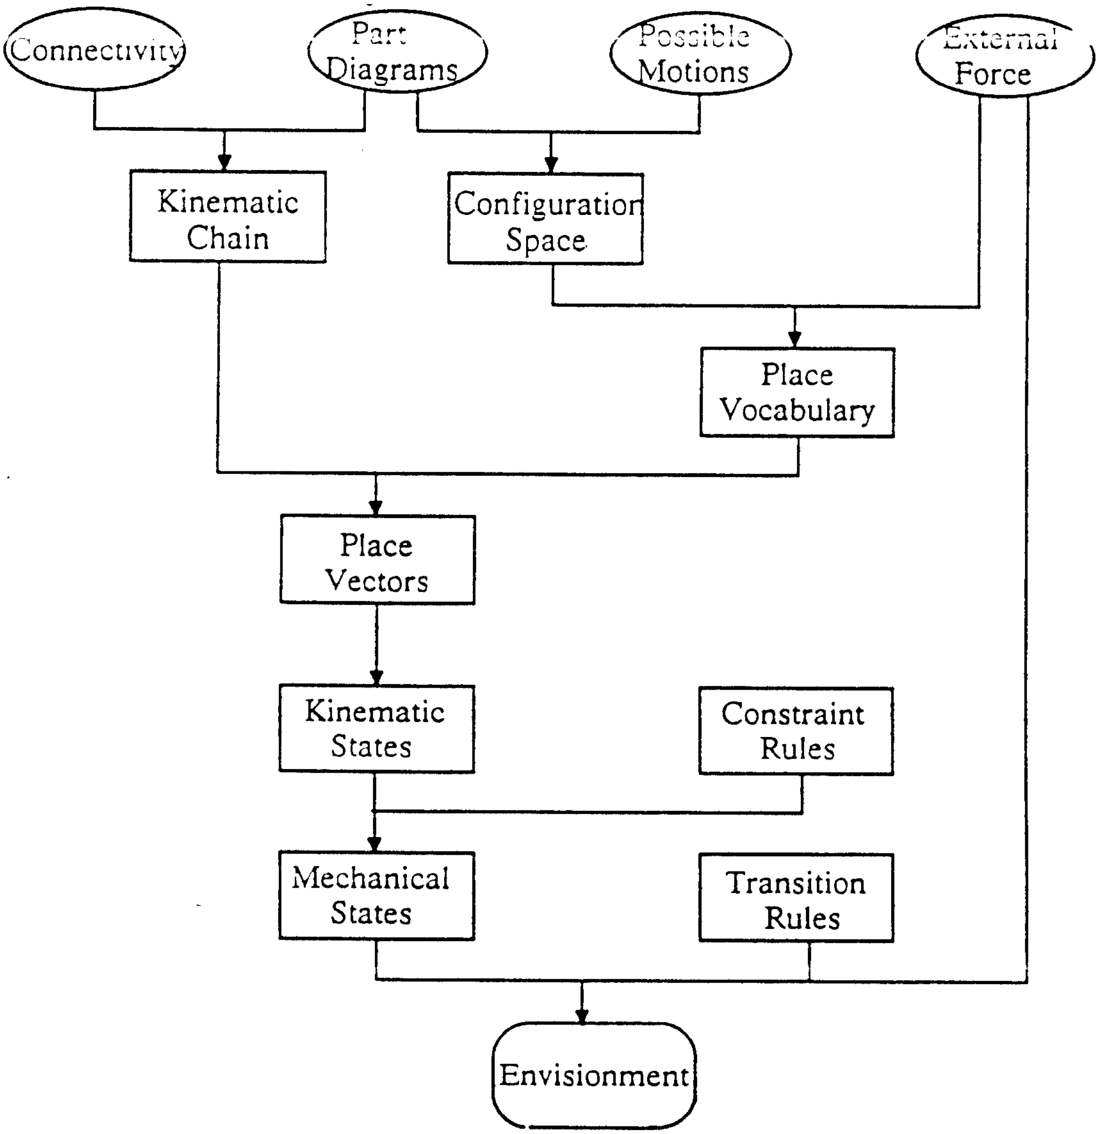

3.2 `CLOCK`'s Metric Diagram

In `CLOCK`, the metric diagram is a boundary-based representation of the geometries of the parts and their arrangement in space. Each boundary consists of a set of vertices connected by smooth edges. The edges can be either segments of straight lines or arcs of circles. An example of a metric diagram for a ratchet is shown in Figure 4.

This particular choice of representation satisfies two of our desiderata. It is minimal, in that we leave out irrelevant properties of the parts (such as their colors). It provides composability, in that we can describe a wide variety of shapes by combinations of these primitive boundary elements.

The metric diagrams fed to `CLOCK` were generated by the following procedure. First, close-up photographs were taken of the parts of the mechanism. These photographs were digitized by hand to the input format demonstrated in Figure 4. Periodically repeating parts of surfaces were noted as part of the input description. The placements of parts within a global coordinate frame was accomplished by careful measurements of the relationships between the parts in the mechanism itself. As Section 3.8 notes, accurate measurements turned out to be very important for subsequent analysis.

Figure 4: Example of a metric diagram for a ratchet.
The wheel and the lever are described in terms of line segements, specified as floating-point coordinates in a local coordinate frame.

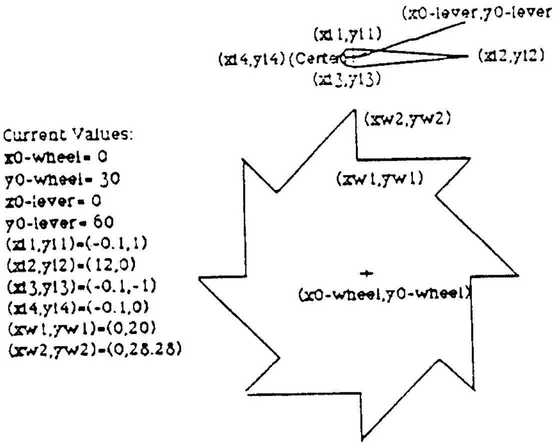

```
Current Values:
x0-wheel= 0
y0-wheel= 30
x0-lever= 0
y0-lever= 60
(xl1,yl1)=(-0.1,1)
(xl2,yl2)=(12,0)
(xl3,yl3)=(-0.1,-1)
(xl4,yl4)=(-0.1,0)
(xw1,yw1)=(0,20)
(xw2,yw2)=(0,28.28)
```

```
((:ROT-REF :ROT-REF)
 ((30 30) (30 60))
 ("Wheel" "Lever")
 (((0 20)
   (:PERIOD 8
    ((0 20) :LINE (0 28.284271d0)
     :LINE (-14.142135d0 14.142138d0))
    (:LINE (0 28.284271d0) :LINE (-14.142135d0 14.142138d0) :LINE (-20 20)
            :LINE (-20 0) :LINE (-28.284271d0 0) :LINE (-14.142135d0 -14.142139d0)
            :LINE (-20 -20) :LINE (0 -20) :LINE (0 -28.284275d0)
            :LINE (14.142136d0 -14.142139d0) :LINE (20 -20) :LINE (20 0)
            :LINE (28.284271d0 0) :LINE (14.142136d0 14.142138d0) :LINE (20 20)
     :LINE (0 20))) (0 20)))
 (((-0.1d0 -1) :LINE (12 0) :LINE (-0.1d0 1) (:ARC-S (-0.005d0 0)) (-0.1d0 -1))))
```

3.3 CLOCK's Place Vocabulary

We use regions in configuration space (c-space) to provide a place vocabulary for reasoning about mechanisms. Each degree of freedom of a mechanism is represented by a dimension of c-space. This means every point in c-space corresponds to a particular arrangement of a mechanism's parts. Since solid physical objects cannot overlap, the regions of c-space which correspond to arrangements where any parts overlap are defined as *blocked space*. The regions of c-space which correspond to arrangements where no parts touch are *free space*. The boundary between free and blocked space corresponds to the subset of configurations where parts of the mechanism are touching. Free space and its boundary represent the legal arrangements of a mechanism's parts.

How should configuration space be quantized to form useful places? There are two sources of constraint: The shapes of parts and dynamics. We examine shape first. Given our metric diagram representation, there are only two possible ways that plane objects O_a and O_b can touch:

1. A vertex of O_a touches a boundary segment of O_b, or vice versa. Note that this subsumes the case where a pair of vertices touch each other, since every vertex is part of some boundary segment, and no other condition besides point of contact is imposed on vertex/boundary contacts.
2. A boundary segment of O_a touches a boundary segment of O_b. Unlike the first case, this case requires that the tangents of the boundaries be parallel.

Each possible instance of these cases defines a *c-space constraint* between O_a and O_b. We call the first kind *vertex constraints* and the second kind *boundary constraints.* A c-space constraint is an algebraic curve containing all the configurations where the defining condition of touch is satisfied. Consequently, *Blocked space* is given as the space enclosed by the envelope (or envelopes) of all configuration space constraints. In our work the place vocabulary is only concerned with actually achievable configurations, hence we focus on representing *free space*, the complement of blocked space. (Quantizing blocked space might be useful for explaining why a proposed mechanism could not work, but that task has not concerned us.)

The idea of configuration space was developed in the last century by Heinrich Hertz for the purpose of formalizing mechanism kinematics [29]. Its use as a computational device is due to Tomas Lozano-Perez [46], who used it for the problem of robot motion planning. In his work, objects are approximated by polyhedra, and thus requires only vertex constraints. However, in kinematics an approximation of a curve by line segments results in spurious discontinuities in behavior. For example, a wheel that is approximated by a polygon will not run smoothly, contrary to its actual behavior. To avoid such anomalous qualitative descriptions, we have extended his work to allow arcs of circles in boundaries as well.

The configuration space is defined by the motions we need to consider. In general, a mechanism has an underlying configuration space of a finite, but possibly enormous number of dimensions, describing the simultaneous locations of each of its parts. A key issue in our approach is how to avoid the computational complexity inherent in directly manipulating high-dimensional spaces.

One source of constraint is our assumption of fixed-axis mechanisms. In such mechanisms each part has only one degree of freedom. This means one can reduce the dimensionality by considering a number of two-dimensional subsets of configuration space, each representing the interaction of a pair of parts, and compose these subspaces as needed to generate descriptions of the whole mechanism. Thus by careful decomposition, the problem can remain tractable no matter how large the complete mechanism is. For the rest of this section we assume that we are dealing with only

Figure 5: Configuration space for the QRG Clock's ratchet

Areas marked as *blocked* indicate impossible configurations which would require some part of the objects be inside each other. Areas marked *open* indicate configurations where the objects are not in contact. The divisions between free and blocked regions indicate configurations where the objects are touching. Because both of the parts have a rotational degree of freedom the configuration space of the ratchet wraps around so that the top and bottom join as well as the left and right sides.

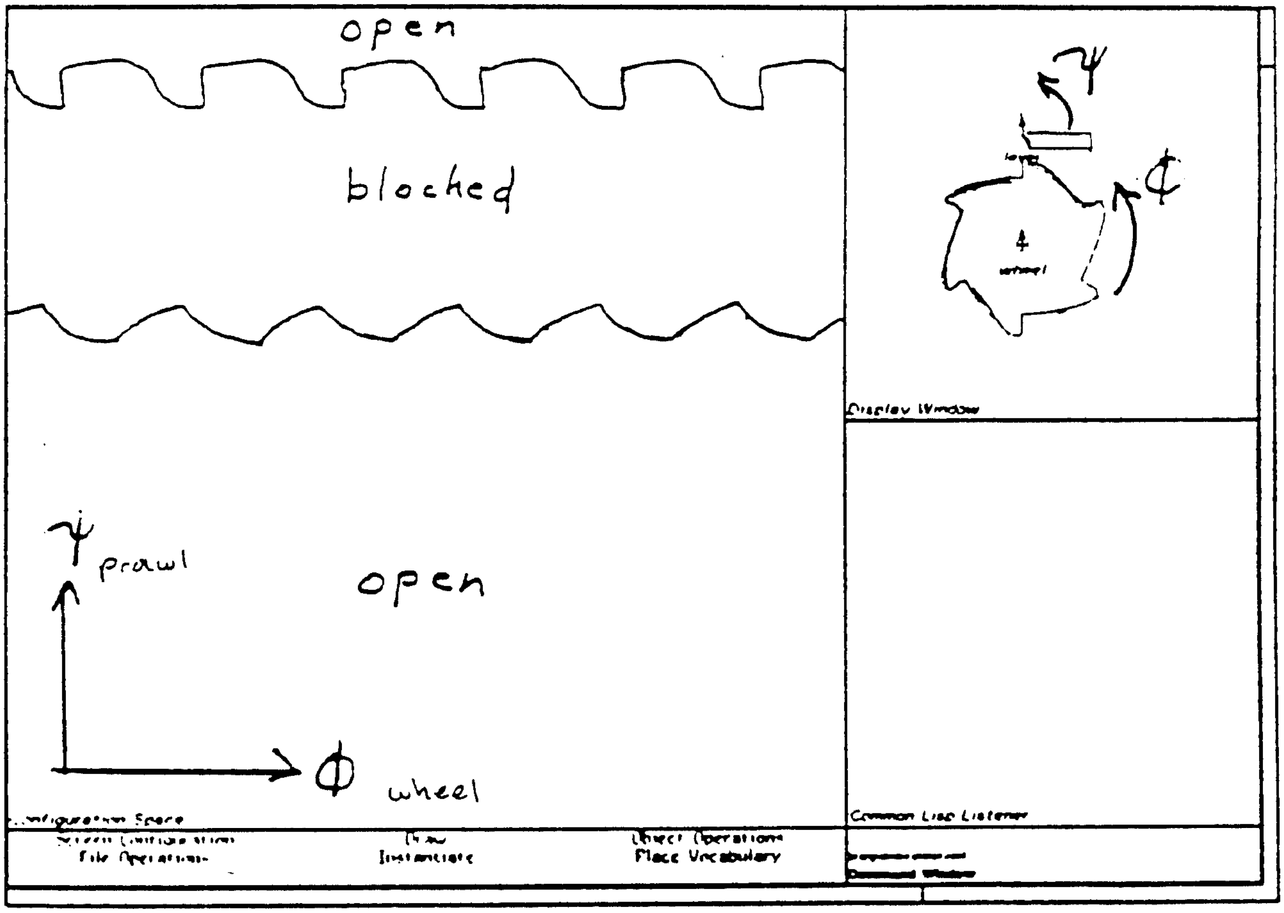

two interacting parts, and in Section 3.3.2 we describe how to compose these pairwise vocabularies to describe the whole mechanism.

Mechanism parts can only interact when they are in contact, i.e. when the device is in a configuration which falls on a c-space constraint. Consequently, the place vocabulary must distinguish configurations where parts touch from those where they do not. The place vocabulary thus becomes a *cell complex* of the following four kinds of elements:

FULL-FACEs: Subsets of free space whose dimensionality is the same as that of the configuration space.

CSEGs: Places with one contact point, bounding the FULL-FACEs, whose dimensionality is one less than that of the configuration space. (i.e., "Constraint SEGments")

JOINs: Intersections points between places with one contact.

FSDs: Divisions of free space which mark important distinctions. (i.e., "Free Space Divisions")

In FULL-FACE configurations, no mechanical interaction of the mechanism's parts occurs. These regions correspond to "play" between the parts of the mechanism. To minimize ambiguity in the representation of connectivity, we create distinct CSEGs for each monotone subsegment of the constraints which define a FULL-FACE's border. JOINs represent corners between CSEGs. FSDs are defined below. Topologically, the JOINs link together the CSEGs and FSDs in a graph. The CSEGs and FSDs in turn enclose the FULL-FACEs, forming a complete graph which represents the cell complex.

Each place in the graph is labeled with the set of qualitative directions of motions possible in it (using the qualitative vector representation of Section 3.5.1). This information is calculated from the type of the place. Each place is also annotated with the place (or places) which can be reached from it by moving in each legal direction. This information provides a concise representation of the potential changes in connectivity wrought by motion.

3.3.1 The effect of dynamics on the place vocabulary

In addition to representing the possible positions of parts, the place vocabulary must also provide a substrate for reasoning about forces and motion. This requirement adds additional restrictions to the decomposition of c-space which forms the place vocabulary. These additional distinctions are called free-space divisions (FSDs). Free space divisions serve two purposes. First, they reduce ambiguity in state transitions. Second, free-space divisions provide a geometric reflection of the effects of external forces.

The excess spatial ambiguity problem arises because some FULL-FACEs can have many neighbors in a particular direction. The more neighbors there are, the greater the ambiguity there is with regard to how connectivity can change, since each neighbor in a given direction represents a possible kinematic state achievable by moving in that direction. Introducing FSDs can subdivide the FULL-FACE, and thus minimize the ambiguity. In CLOCK we introduce FSDs which run parallel to the free directions of each object and distinguish where monotonic changes in the constraint curves occur. (See figure 6.) These directions were chosen because mechanism components can only move along their free directions. These divisions will unambiguously distinguish the next possible contact if the motion of a single object were to continue while the other object remained stationary.

Free-space divisions are also used to mark the geometric aspects of external forces. Recall that when describing a mechanism, external forces may be associated with an object (for example, a spring will always push a gear in one direction) or with objects in various positions (for example, in certain positions a pendulum is pulled clockwise by gravity, in others counterclockwise). CLOCK models this by associating forces with sets of places in the place vocabulary. For example, gravity will pull the pawl in Figure 4 clockwise in configurations where the pawl is to the left of its pivot, counterclockwise when it is to the right, and have no effect at the instant when it is perfectly upright. While sometimes information about external forces lines up with divisions introduced for motion, they often don't. Thus new FSDs are introduced to further decompose places when the set of forces acting changes.

Figure 6: Place vocabulary for the ratchet

This diagram illustrates one periodic section of the place vocabulary for the configuration space in Figure 5. Each line, point, and region corresponds to a place. Each place is labelled with an integer for easy reference. Free space divisions (shown dashed) distinguish regions of space where a qualitative change in contact will occur under movement of a single part. The interpretation of three sample places (102, 17, and 26) are shown below.

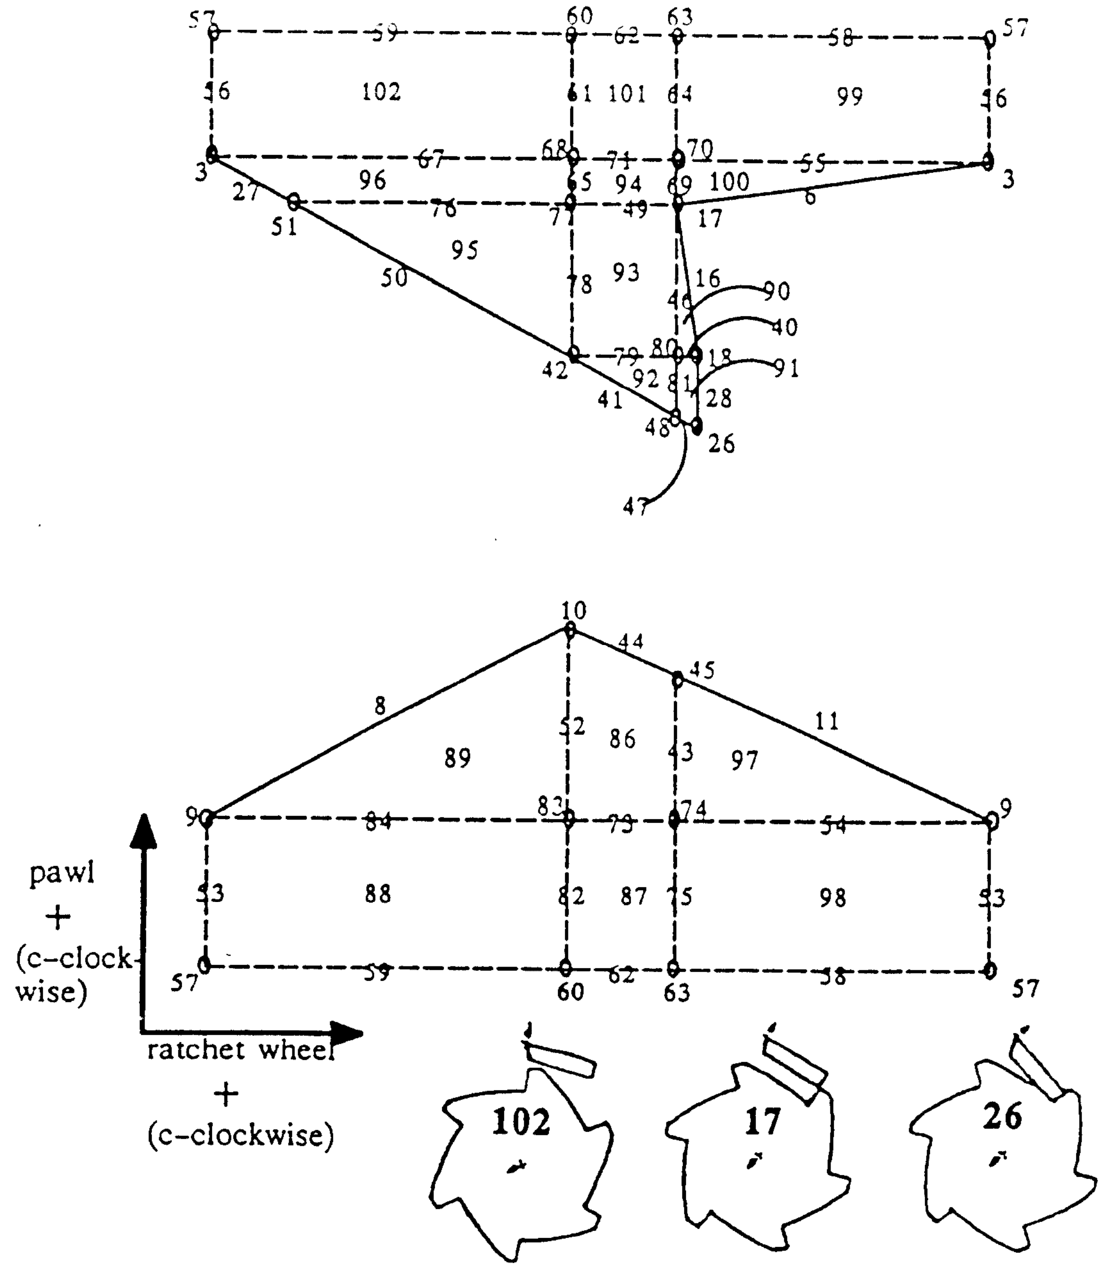

3.3.2 Place vectors

As described so far, the place vocabularies represent all possible pairwise interactions of a mechanism's component parts. While a higher order configuration space could theoretically represent all possible interactions of all the components in a mechanism, explicitly constructing the higher-order configuration space as a separate entity would be extremely inefficient. Furthermore it is unnecessary, since each part of a mechanism typically only interacts with a few others. Instead, we specify the kinematic state for an entire mechanism as a vector of places from each place vocabulary describing the interactions of pairs of parts. We call this description the ***place vector*** of the mechanism.

The number of different place vectors is bounded by the product of the number of different places in the constituent place vocabularies. Typically, only a small subset of these vectors actually correspond to legal kinematic states. A compatibility constraint is imposed to define these legal states. In particular, if two place vocabularies involve a particular part, the only combinations of places which can be part of a legal place vector are those for which the allowable interval for the common part intersects [9]. We call this the ***kinematic compatibility*** constraint. (Interestingly, the place vocabulary can be organized so that these constraints can be computed directly in it, without reference to the metric diagram.) This set can still be very large, but Section 3.4 shows how this size may be reduced through abstraction to make analyzing complex mechanisms tractable.

Information can be propagated across a kinematic chain by establishing ***correlations*** between the place vocabularies of kinematic pairs. Correlations provide additional symbolic information relating places from different place vocabularies. For example, if two parts are connected to the same shaft, all places with the same rotational positions are constrained to be identical. Transmission ratios are related in a similar manner: the fact that one object was rotating twice as slow as another can be represented as a correlation between places in the first object and places offset by 0° and 180° for the second object in the metric diagram.

3.4 Abstracting the place vocabulary

While the subset of legal place vectors is generally much smaller than the product of the sizes of the constituent configuration spaces, the number of legal place vectors can still be extremely large. A complex shape may need to be decomposed into many pieces, to capture the subtleties of its interactions. Small imperfections in the shapes of the objects themselves may be relevant for ascertaining the possibility of binding, but are often irrelevant for figuring out the global character of the part's motion. Empirically, we discovered it was necessary to introduce abstraction techniques to keep the size of place vocabularies tractable.

What should the abstraction be performed over? One possibility is abstracting the shapes of the objects themselves. In some cases it seems plausible that a coarse description of shape could be used to provide an initial topology of configuration space, and refined as necessary using finer-detailed shape descriptions (see Section 4). But, as noted in Section 2.2, this approach will often fail. In analyzing gear chains, for instance, a pair of smoothed wheels may either be inconsistent due to overlap, or seem to slide smoothly, while a closer examination of their shape will reveal they will bind. Smoothing only makes sense when one knows which differences are relevant, and that information is not apparent until the place vocabulary is computed.

This argument suggests that abstracting at the level of the place vocabulary should provide better results, and in our experience it does. We used three abstraction techniques in `CLOCK`. First, adjacent `CSEG`s with qualitatively equivalent surface normals (i.e., those `CSEG`s whose surface normals have identical qualitative vectors) were merged into a single segment. This loses information about exactly which boundary element of one part is in contact with which boundary element of the other, but this information is unnecessary for predicting behaviors. (Backpointers to the original c-space can be kept for explanation generation if desired.) Second, places which are smaller than some grain size δ are eliminated. This operation has the effect of ignoring interactions involving small surface imperfections. Letting δ be too large can introduce erroneous results, of course. Finally, if the width of a `FULL-FACE` is below some ϵ, it can be replaced by a special, two-sided `CSEG`. This operation has the effect of ignoring the play between closely mated parts.

The effect of these three operations can best be illustrated by examining their effect on the possible kinematic states for the QRG clock. Its constituent place vocabularies can be divided into two parts, the place vocabulary for the escapement and the vocabularies for the gear pairs. The initial vocabulary for the escapement consisted of over 1,300 places, and the initial vocabulary for the gear pairs ranged from roughly 6,000 to 60,000 places. However, these huge numbers are due to the periodicity of the gears, and thus are artificially high[2]. Collapsing equivalent configurations reduces the number of places for the escapement to 96 and roughly 16 for the typical gear pair. But there are six gear pairs, and, ignoring kinematic compatibility constraints, the size of the kinematic state space is the product of the sizes of the constituent vocabularies. Thus the worst-case number of states is 1.6×10^9. Merging adjacent `CSEG`s with qualitatively equivalent surface normals reduces the escapement to 80 and the gear pairs to 12, thus reducing the worst case size to 2.4×10^8. Eliminating small places (i.e., ignoring small irregularities) reduces the size of the escapement's vocabulary to 58 and the size of a gear pair's vocabulary to 3. This dramatically reduces the worst-case number of kinematic states to 36,450. Finally, replacing narrow `FULL-FACE`s with two-sided `CSEG`s (i.e., eliminating small gaps) reduces the number of places in each gear pair to one, hence reducing the size of the kinematic state space to simply that of the escapement's vocabulary. Thus in the end the size of the place vocabulary for the QRG clock is 58 places.

Some of the drama of this example, to be sure, is due to ignoring compatibility constraints, which leads to a gross overestimate about the size of the total state space. But applying those constraints is not free, and these abstraction operations allow us to greatly minimize them, replacing pairwise interval consistency tests with simple local operations. Furthermore, it demonstrates that intuitive notions about possible motions can arise as a combination of qualitative reasoning with abstraction operations. For instance, every engineer knows that, to a first approximation, in a well-designed gear chain there is only one degree of freedom no matter how long the chain is. One can imagine encoding such a principle explicitly as an axiom via traditional knowledge-engineering techniques. But we achieve the same results via geometric and qualitative reasoning, which makes explicit the conditions under which this principle is true (i.e., that δ and ϵ aren't "too large"). Thus the resolution can be tuned according to the needs of the task, rather than relying on an a priori categorization of the mechanism's parts.

3.5 Qualitative Mechanics

So far we have defined the kinematic state of a mechanism. To define the full mechanical state of a mechanism requires imposing dynamical information. This dynamical information concerns forces and motions. Nielsen's account of *qualitative mechanics* (QM) provides this. Here we summarize the highlights of qualitative mechanics necessary to understand `CLOCK`; for a complete exposition see [55]. Section 3.5.1 begins by introducing a qualitative representation for vectors, and discuss how they are used to represent motion and shapes. It outlines qualitative laws of motion, including

---

[2] Recurring surface patterns are identified by the user when creating the original part drawings.

Figure 7: Qualitative translational vectors

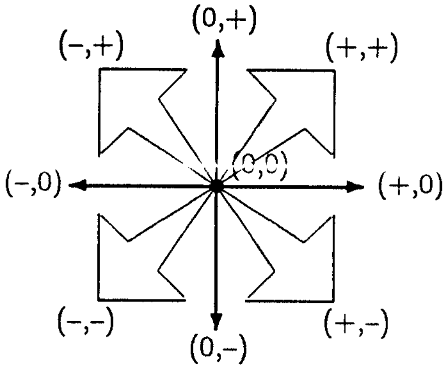

a notion of mechanical constraint that describes how forces are transmitted and how objects can move. Section 3.5.2 defines mechanical states and transitions between them.

3.5.1 Qualitative vectors

Mechanisms are usually subject to external forces (e.g., gravity, springs, someone pushing on them, etc.), and forces are transferred between components depending on the connectivity of its parts. Representing such forces requires a qualitative vector algebra. Traditionally, vectors are also used to represent surfaces (through normals and tangents) as well as other properties, such as velocity and momentum. In its current form qualitative mechanics uses vectors for describing positions, motions, forces, and surface normals and tangents.

The basic idea of qualitative vectors is to extend the standard sign algebra [6] to multiple dimensions. That is, given a set of n orthogonal reference directions, we define a qualitative vector as an ordered tuple of length n whose elements are drawn from { `+`, `0`, `-` }. In the $2D$ case, by convention we align these reference frames with the traditional reference frame, hence "left and up" would be `(+,+)` (see Figure 7). Rotational directions are represented similarly, with counterclockwise corresponding to `+` and clockwise corresponding to `-`. A set of laws defining dot products and various aspects of motion using this notation can be found in [55].

One important use of qualitative vectors in `CLOCK` is decomposing the shapes of objects. Consider a book lying flat on a table. What happens when one pushes the book perpendicular to its spine depends on where along the spine you push it. Pushing it on an end leads to rotation, while pushing it exactly in the middle leads to a pure translational motion. These distinctions must be captured in the qualitative representation of the initial shapes of objects. In particular, the elements of a boundary are individuated by having a unique surface normal and a unique direction to the object's center of rotation, both expressed as qualitative vectors. At first glance this can lead to some counter-intuitive decompositions (c.f. Figure 8). However, the dependence of the qualitative representation of the shape on the choice of center of rotation is a natural consequence of the laws of motion. The situation the object is a part of helps define its center of rotation, and only with this information can the appropriate qualitative representation be constructed. Thus for this class of problem a task-independent qualitative representation of shape is ruled out for all but the simplest

Figure 8: Qualitative decomposition of a $2D$ book

This picture illustrates the qualitative description of the surfaces of a $2D$ book, assuming the center of rotation is the center of mass. Notice that imposing a different center of rotation, by for instance applying downward pressure on some other part of the book, would imply a different qualitative description of the shape.

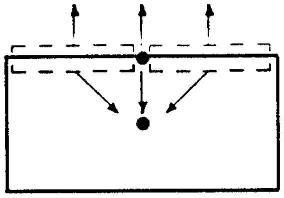

cases. Fortunately, in fixed-axis mechanism domain of `CLOCK`, the center of rotation is always fixed by the construction of the device, so a single qualitative surface description can be computed for each part.

Qualitative vectors are also used in defining mechanical constraint. The laws of mechanical constraint answer several kinds of questions. Given an object $\mathcal{OB}_1$ in contact with another object $\mathcal{OB}_2$, mechanical constraint determines (1) how $\mathcal{OB}_1$ may move if $\mathcal{OB}_2$ is fixed and (2) how $\mathcal{OB}_2$ should be constrained in order to prevent $\mathcal{OB}_1$ from moving in some particular direction. Similarly, if we know $\mathcal{OB}_2$ is moving, then the laws of mechanical constraint can determine (1) how the motion of $\mathcal{OB}_2$ can push $\mathcal{OB}_1$ and (2) how can $\mathcal{OB}_1$ and $\mathcal{OB}_2$ be moving consistently with respect to each other. The essence of these laws is that one object cannot move into the half-plane defined by the surface normal of the other, but can move in any other direction. At corners (i.e., `JOIN`s in the place vocabulary) two cases arise:

- If the free space forms a concave angle, the objects are constrained from moving in directions defined by the union of the constraints imposed by the adjacent surfaces.
- If the free space forms a convex angle, the objects are constrained from moving in directions defined by the intersection of the constraints imposed by the adjacent surfaces.

These intuitions and their implications are formally represented in [55].

The notion of mechanical constraint also defines the way parts interact by propagating forces. When two objects are in contact (i.e., their configuration lies along a `CSEG`) they are constrained from moving in any direction which is into the half plane defined by the opposite of the `CSEG`'s surface normal. If a force is propagated through one of the objects in any of these directions it will produce a resultant propagated force in the second object with direction along the reverse of the contact surface normal.

An important problem in any propagation algorithm concerns what happens at confluences of propagation, where values computed along two different paths collide. Consider for example figure 9. Assume **A** can only move vertically and **B** can only move horizontally. If an external force is applied to **A** in the downward direction (i.e., `(0,-)`), the laws of qualitative mechanics allow the inference that a force is transmitted to **B** at surface **S** in the `(+,-)` direction – that is, to the

Figure 9: Clashing forces example

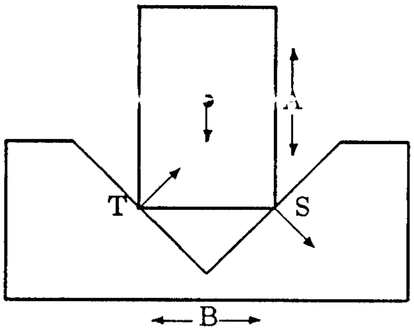

right and down. This force in turn causes a force to be transmitted back to **A** on surface **T** in the (+,+) direction (i.e., to the right and up), which pushes **A** upward. If we attempt to compute the net force on **A** by a qualitative sum, we get an ambiguous result. Without knowing the relative magnitudes of the forces we might hypothesize that pushing down on **A** may cause it to move up!

This simple example illustrates a more general problem. A force may be transferred back to its source through a complex chain of other objects. In the mechanism domain, for example, three gears connected to each other produce such clashes. To solve this problem, every external force represented in `CLOCK` is given a unique name. This name is propagated along with the magnitude of its contribution, to provide a causal attribution to each contribution to the net force on an object. This additional information is used when comparing forces with opposite signs. That is, we presume that opposing forces with the same root cause will cancel. In this case, since the cause of the push upward on **A** is the same as the cause of the push downward on **A**, we conclude there is no net vertical force on **A**. (This is only true for rigid bodies.)

Consider now Figure 10(a), which illustrates a form of feed-forward. Assume **A** can only move vertically and **B** can only move horizontally. Further, assume the actual orientation of the slopes is ambiguous, but somewhere in the (-,-) (i.e., left, down) to (+,+) (i.e., right, up) region. A downward external force applied to **A** pushes **B** to the right at point **S**. In turn, **B** moving to the right will push **A** downward at point **T**. The qualitative sum of the two forces suggests that **A** moves downward. Yet if the slope at **S** is less than the slope at **T**, the mechanism will jam. (By contrast, consider figure 10(b). Assume part **A** can only move vertically, and part **B** can only move horizontally. A downward external force applied to **A** pushes **B** in the left, down direction at points **S** and **T**. In this case the external force does not get propagated back to **A**, since both contact points push **B** to the left.)

In general, a force with the same cause transmitted by two different paths causes the mechanism to jam when it acts in the same direction and the magnitude of the transmitted force is not equal to the original force. However, if the magnitudes of the forces are equal the mechanism will move freely. (Mathematically the only way the motion of object $\mathcal{OB}_m$ can also be $\mathcal{OB}_{\alpha m}$ is when $\alpha = 1$ or $m = 0$.) Even when detailed surface information were available, this ambiguity arises through

Figure 10: Positive feedback example

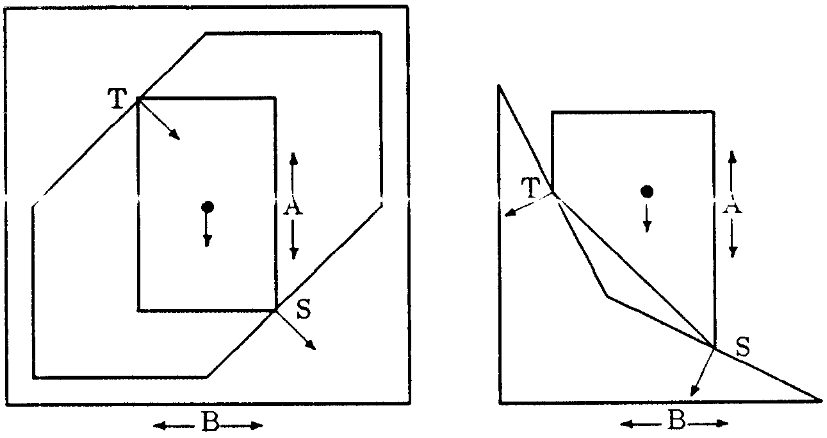

inequalities of transmitted forces. For example, an even number of gears connected in a circle will jam unless the torques balance at each contact for every gear. Even if detailed surface information were available, one cannot tell if the magnitudes of forces are equal, except in the restricted case where all elements in intermediate chains are identical. Consequently, in CLOCK we presume that jamming does not occur, but all occurrences of positive feedback are noted for possible external verification.

3.5.2 Mechanical states and transitions

The set of mechanical states represents all possible individual behaviors of the mechanism. Mechanical states are formed out of consistent combinations of where things are (kinematic state) with how they are moving (dynamic state). The kinematic component consists of a place vector, while the dynamic component is a consistent set of vectors describing the motion of each part. (Notice that since each part of a mechanism can have only one degree of freedom, we need only a single dimension to describe its motion. That is, a part's `Motion` will be interpreted as translational or rotational according to the input specification of how it was free to move.) A set of motion vectors is consistent exactly when (a) each individual motion is legal, given the object's location and (b) each object is given a unique motion vector. Examples of inconsistent dynamical states include requiring a gear to move both clockwise and counterclockwise, or requiring the teeth of one gear to pass through the teeth of another. Figure 11 illustrates the contents of a complete mechanical state for the QRG clock.

Figure 11: Sample complete mechanical state

Motion(Pallets) = 0
Motion(Scape) = -
Motion(G1) = -
Motion(G2) = +
Motion(G3) = +
Motion(G4) = -
Motion(G5) = -
Motion(G6) = +
Motion(G7) = -
Motion(G8) = -
Motion(G9) = +
Motion(G10) = +
Motion(G11) = -
Loc(Scape, Pallets) = SP-PL-22
Loc(G1, G2) = G1-2-PL-1
Loc(G3, G4) = G3-4-PL-1
Loc(G5, G6) = G5-6-PL-1
Loc(G6, G7) = G6-7-PL-1
Loc(G8, G9) = G8-9-PL-1
Loc(G10, G11) = G10-11-PL-1

Just as mechanical states are a combination of dynamic and kinematic information, state transitions are the combination of the changes in the dynamic component and the kinematic component of the current state. Here we consider state transitions in detail. As in qualitative dynamics, we distinguish changes which occur instantly from those which require an interval of time to occur. As usual, transitions which occur in an instant are presumed to occur before those which require an interval of time. We focus on dynamical considerations first, then focus on kinematics. As before, we start by considering the motions of individual objects (for dynamics) and pairs of objects (for kinematics) and then describe how to combine them to compute global state transitions.

Because forces are associated with places, the active external forces in a mechanical state may be determined by inspecting the place vocabulary. These external forces propagate through the kinematic chain either by sharing a common component or through contact. The simplest case is when an object touches nothing else. This corresponds to the set of places describing its location all being either `FULL-FACE`s or `FSD`s. In this case any motion it has will persist unless acted upon by an external force. Table 1 shows how to predict the possible next motions for each force and current motion. An *ambiguous* value is represented by "?". An ambiguous change indicates a branch in the predicted next motion. Those transitions which occur in an instant are distinguished by the subscript "i". All other transitions require some interval of time, and hence their current values may persist in the next state. When an object is in contact with other objects, there can be

		Force Component			
		+	0	−	?
	+	+	+	0	$+ \vee 0$
Motion Component	0	$+_i$	0	$-_i$	?
	−	0	−	−	$- \vee 0$
	?	$+ \vee 0$	?	$- \vee 0$	?

Table 1: Effects of force on next possible motion

propagated forces as well as external forces. Table 1 also applies to the net force in these situations, once cancellations due to identical underlying causes have been performed.

In addition to these dynamical transitions, kinematic transitions can occur due to motion. These transitions are detected by examining the connectivity information stored in the place vocabularies. Consider a pair of objects which might interact, that is, whose place vocabulary is non-trivial. If either object is moving, the connectivity information in their place vocabulary can be used to determine potential changes in connectivity. That is, the set of places which can be found in the direction of motion from the current place can be retrieved and each such element hypothesized as the next potential relationship between the objects. The duration of kinematic transitions depends on their dimensionality. If the current place has no dimensionality in the direction of motion, the kinematic change occurs in an instant. If the current place has dimensionality in the direction of motion, then the transition requires some interval of time to occur. Suppose a pair of objects both have `(0,0)` as their motion. Then any state transition will be due to dynamics (i.e., the introduction of motion through an external or transmitted force) and will occur in an instant. If the current place has no dimension in the direction of imparted motion it will also transition in an instant.

The transitions hypothesized for a given mechanical state are found by combining dynamical and kinematic transitions. If a given mechanical state has instantaneous kinematic transitions then the place vector hypothesized for the next state consists of the effects of these instantaneous changes and any resulting instantaneous place transitions. Otherwise, every combination of transitions which yields a consistent place vector constitutes a valid hypothesis about the next kinematic state. (We include, of course, the null transition, to provide a kinematic state for the case where a dynamical transition occurs.) The dynamical transitions were described previously. The cross-product of these new dynamical and kinematic states defines the set of initially hypothesized transitions. Legal transitions are a subset of this cross-product which satisfy certain additional criteria.

Two criteria always apply. First, the resulting state must be a legal mechanical state of the system. That is, its place vector does not violate the kinematic compatibility constraint, and the set of motion vectors when combined with the place vector does not violate the laws of mechanical constraint. The second criteria is that the new state not be identical to the old one (i.e., the *no change* filter of [41]).

Additional constraints are applied if a collision occurs. A collision is detected when the new place vector includes a contact relationship not found in the previous one. Modeling collisions can be difficult, so we make the following three simplifying assumptions: (1) The collision forces are great enough to overwhelm the external forces. (This is a good assumption because the collision forces between perfectly rigid bodies are infinite.) (2) No two collisions will occur at exactly the same time. and (3) All collisions are inelastic.

During a rigid body collision, a change in the motion of the two bodies may instantaneously occur in order to satisfy the constraints imposed by the new contact. Notice that the way we have defined state transitions enforces continuity in other circumstances. The duration of the states both before and after the collision may persist for some interval. Because we do not know the relative masses of the objects, the resultant motion is highly ambiguous. By assuming the collisions are inelastic, the resultant mass is always greater than either of the original masses. This reduces the ambiguity, since then the resultant motion of the coupled objects is the qualitative vector sum of the directions of motion of the original masses.

Any change in motion which causes objects to move inside each other is disallowed by the rigid body assumption. Thus motions of any objects which were indirectly in contact with the objects involved in the collision may also have their motions affected in the same instant. Unfortunately, the best way to describe the collision constraints is procedurally [55]. After determining the direction of the objects after the collision, the resultant motion must be compared with the motion of adjacent objects along the kinematic chain. If the newly determined motions of the colliding objects and the other objects in the chain violate the rigid body assumption, the motion of the object farthest from the collision is modified until this conflict is resolved. This change ripples along the chain as long as any modification of motion has been made.

Figure 12: CLOCK's algorithm

```
1. Compute kinematic component of state.
   1.1 For each pair of interacting objects,
       1.1.1 Compute their configuration space.
       1.1.2 Apply abstraction operators.
       1.1.3 Generate place vocabulary.
   1.2 Establish kinematic chains by clustering
       interacting place vocabularies.
   1.3 For each kinematic chain, compute consistent place vectors.
2. Compute mechanical states by finding all motions consistent with each state vector
3. Find state transitions.
   3.1 For each mechanical state,
       3.1.1 Using motion and places, find kinematic transitions.
       3.1.2 Propagate forces.
       3.1.3 Compute net force on each object.
       3.1.4 Use forces to compute dynamical transitions.
       3.1.5 Combine kinematic and dynamical transitions to generate hypothetical next states.
       3.1.6 Record each combined transition and its result.
```

3.6 Algorithms

CLOCK operates as a total envisioner. That is, it creates a description of every possible qualitative state of a mechanism and all potential transitions between them, independently of any particular choice of initial state. The overall algorithm is outlined in Figure 12. A detailed description of the algorithm is beyond the scope of this paper; see [9] for the details of place vocabulary construction, and [55] for the rest. Our description here is intended to show the reader how CLOCK instantiates the framework of Section 2.

CLOCK is structured as a set of justify/assume/interpret cycles [21], using an assumption-based truth maintenance system to efficiently share constraints between multiple states. This radically simplifies certain operations. For instance, by the time transition checking occurs, every legal mechanical state has been generated. Thus the constraint of state consistency is trivially enforced by seeing if the result of a transition corresponds to an already-known mechanical state. The drawback of envisioning is that for a particular question one may only need a small subset of the possible behaviors. However, since our main goal was to explore the sufficiency of the representations, generating envisionments makes a great deal of sense. By exploring the entire space of behaviors, we can better understand what our descriptions actually allow.

Figure 13: The kinematic chains of the QRG clock.

Below is a two dimensional representation made by digitizing photos of the components of the clock shown in figure 1. Adjacent components, which are not shown interacting, share a common axis. The hands share a common axis with gears 7 and 11. Gear 6 is attached to a spring which is attached to the ratchet.

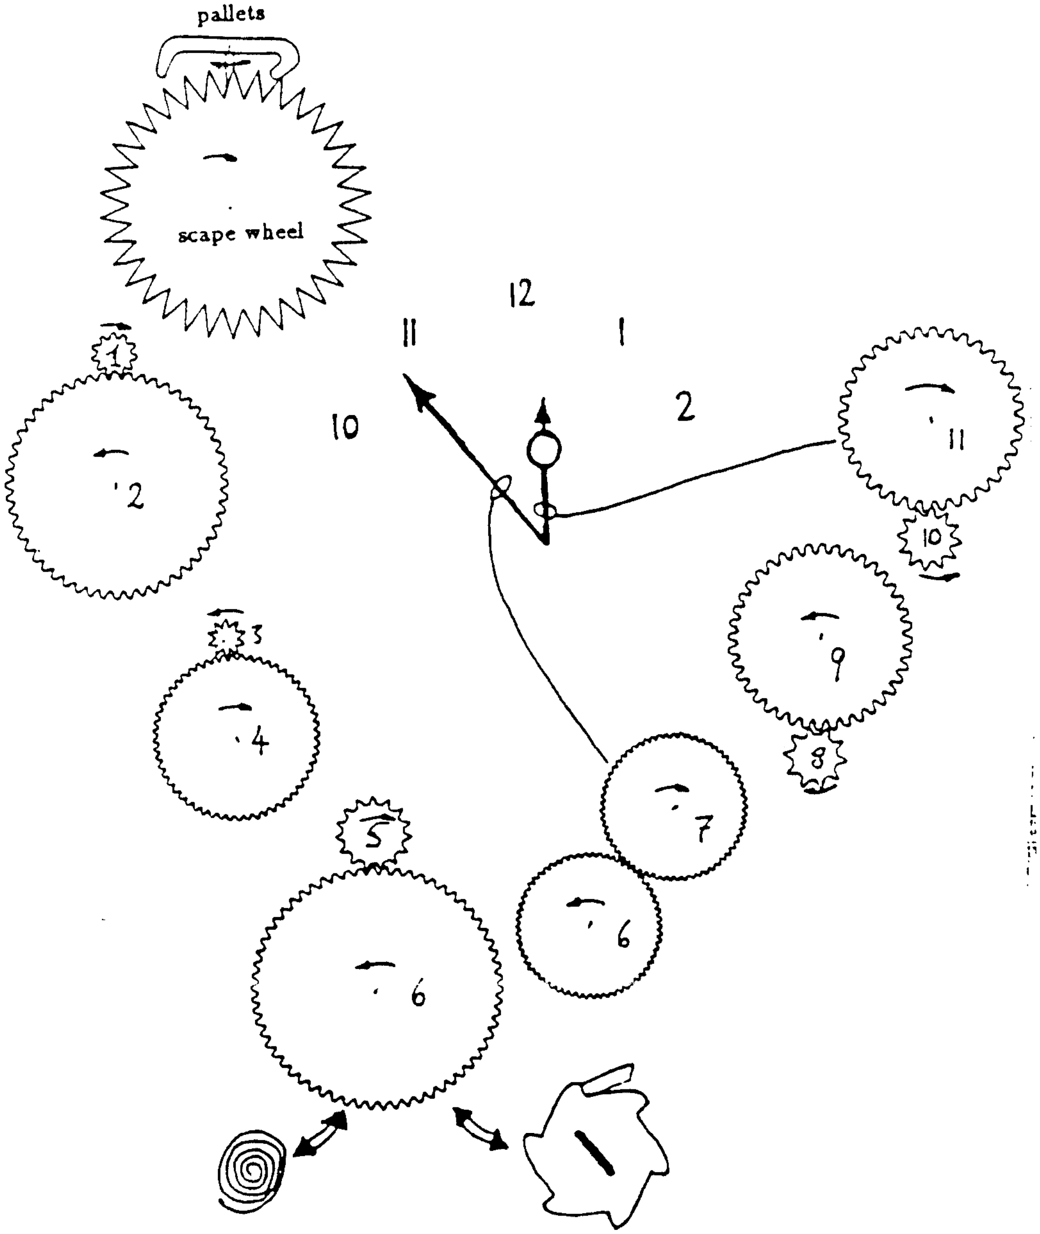

3.7 Example: Analysis of the QRG Clock

Consider again the mechanical clock shown in Figure 1. The QRG clock contains an escapement, eleven gears, a spring, and a ratchet. While it is obvious from the photograph that the clock is a three-dimensional object, we approximate it as a two-dimensional object by pretending all the interacting parts are in a single plane (see Figure 13).

Gear wheel 6 is attached to a spring which drives it to the left. This spring is the source of energy for the clock. It is wound up using a handle connected to a ratchet, which prevents the handle from turning backwards and relaxing the spring. The spring and ratchet are indicated at the bottom of Figure 13. Because there is no rigid connection between gear 6 and the ratchet we consider the gear train and the ratchet to be separate kinematic chains which may be analyzed independently. Furthermore, we identify the pairs of objects which can interact by hand. (In principle CLOCK has the ability to do this computation for itself, but as discussed below, the code for computing place vocabularies was extremely inefficient and thus we wanted to avoid the n^2 work of finding empty place vocabularies.) This consists of six pairs of gears, the scape wheel and pallets, and the ratchet. To model the effects of coupling, parts which share the same axes are constrained to have the same forces and motion.

Recall that the abstraction operations described in Section 3.4 required two numerical parameters, ϵ and δ. Empirically, we found that a good value for both parameters were $\frac{1}{100}$th of the size of the configuration space, and that the results were not very sensitive to this particular choice of value. Much larger values tended to merge opposite CSEGs of the escapement, while much smaller values left extraneous places in the place vocabularies.

After using the abstraction operations, there were 58 distinct place vectors describing the kinematic states of the QRG clock. Combining these place vectors with consistent motions of each part produced an envisionment of 462 mechanical states. Reprinting the whole envisionment here would be excessive detail; readers wanting the complete envisionment should see [55]. Instead, we summarize some results obtained by analyzing it. Figure 14 and Figure 15 contain a subset relevant to these analyses and the next section.

Perhaps the key question is whether or not the envisionment is correct. This question is difficult, as Kuipers [41] points out. Our approach is to examine four questions:

1. Are the individual states physically reasonable?
2. Are the individual state transitions physically reasonable?
3. Does the envisionment include all reasonable behaviors of the mechanism?
4. Does the envisionment include only reasonable behaviors?

We (Nielsen) found that the answers to Questions 1 and 2 are yes, by examining every state and transition to ensure that the clock could behave that way.

Figure 14: Place vocabulary for scape wheel and pallets of the QRG clock.

A section of the configuration space for the scape wheel and pallets has been enlarged here to show detail. It depicts the results of the abstraction operations, includes FSDs, and shows the symbolic names of the spatial regions which will be used as the place vocabulary. Representative configurations for some of these regions are given to either side.

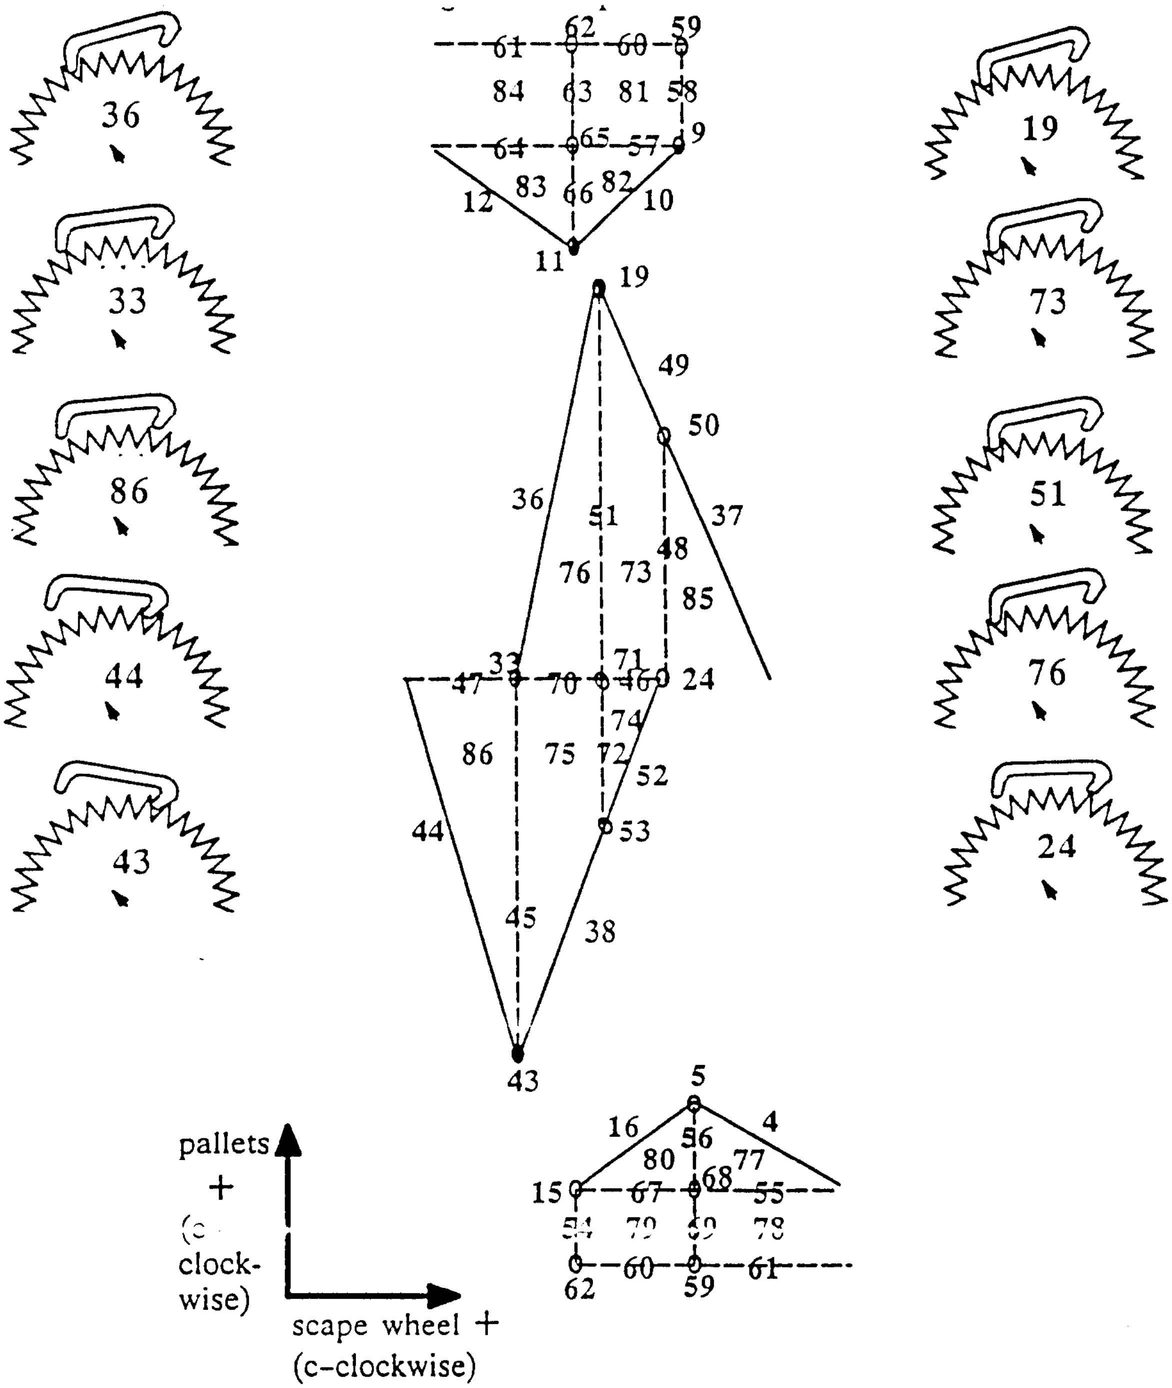

Figure 15: Subset of the envisionment for the QRG clock

The table below shows a portion of the possible state transitions for the QRG clock. It indexes the next possible location and motion of the parts by current position (left hand column) and current part motions (topmost row). Counterclockwise motion is denoted "+". Multiple entries indicate ambiguous transitions.

Current Location	Current Direction of Motion (scape wheel / pallets)								
	- -	- 0	- +	0 -	0 0	0 +	+ -	+ 0	+ +
19	76 - - 36 - -			51 0 -	19 - -		73 + - 49 + -		
24	74 - - 52 - -	46 - 0	73 - +		24 - 0	48 0 +	86 + - 44 + -	47 + 0	85 + +
33	86 - -	47 - 0	85 - + 37 - +	45 0 -	33 - 0		75 + -	70 + 0	76 + + 36 + +
36	33 - - 76 - -			76 0 -	36 - -		76 + -	76 + 0	76 + + 19 + - 19 0 0 36 0 0 36 + 0
43			86 - + 44 - +		43 - +	45 0 +			75 + + 38 + +
44			86 - + 24 - +		44 - +	86 0 +	43 + + 43 0 0 86 + - 44 0 0 44 + 0	86 + 0	86 + +
51	76 - -	76 - 0	76 - +	71 0 - 51 - -	51 - -	19 0 0 51 - +	73 + -	73 + 0	73 + +
73	71 - - 51 - - 46 - -	51 - 0 73 - -	19 0 0 19 - - 51 - + 49 - + 73 - 0	46 0 - 73 - -	73 - -	49 - + 73 - +	50 + - 49 + - 24 + - 48 + - 46 + - 73 0 -	50 + - 49 + - 48 + 0 73 + -	50 + - 50 0 0 50 - + 49 + - 49 0 0 49 - + 48 + + 73 0 + 73 0 0 73 + 0
76	33 - - 70 - - 36 - -	36 - - 76 - -	36 + + 36 0 0 36 - - 76 - 0	70 0 - 76 - -	76 - -	36 + + 76 - +	71 + - 70 + - 51 + - 76 0 -	51 + 0 76 + -	19 + - 19 0 0 36 + + 51 + + 76 0 + 76 0 0 76 + 0
86	44 + - 44 0 0 44 - + 86 - 0	44 - + 86 - +	24 - + 44 - + 47 - +	44 + - 86 - -	86 - +	47 0 + 86 - +	43 + + 43 0 0 44 + - 45 + - 86 0 - 86 0 0 86 + 0	45 + 0 86 + +	33 + + 47 + + 45 + + 86 0 +

Question 3 only makes sense with respect to the background assumptions under which the envisionment was made. If we violate those assumptions, say by bashing the clock with a hammer, one should no longer expect the envisionment to correspond to the new artifact's behavior. Consequently, the only changes that make sense are changes in the relative "play" between parts. Even for these, most of the behavioral changes that can be introduced by pushing or pulling the gears cannot be modeled, since values of δ and ϵ were large enough to reduce the gear train to a single place[3]. The only degree of freedom left is the play between the scape wheel and the pallets. On closer examination, we found that the envisionment can be divided into four components:

1. The normal behavior of the clock.
2. Pallets in normal position, but with enough clearance so that the gears can turn freely.
3. The pallets are inverted with pendulum balancing upright, and gears turning freely.
4. Inverted pallets, but still interacting with the scape wheel.

These do a reasonable job of capturing the degree of variability in behavior. We have not figured out all possible ways the QRG clock can go amiss, but those we have thought of (barring hammers) can all be found in the envisionment. For example, the envisionment indicates that there are possible configurations where the gears spin without being regulated by the scape wheel. Another phenomena concerns decaying motion. Clocks are built such that each swing of the pendulum advances the scape wheel by one tooth during normal operation. As the clock runs down, the force of the swing is no longer great enough to advance the scape wheel and the pallets may hit multiple times between the same teeth before advancing. Such behavior can be found in `CLOCK`'s envisionment, for example, by transitions from state (36, –, 0) to (76, –, –) to (70, –, –) to (75, –, –) to (38, –, 0) then (38, –, +) to (75, –, +) to (70, –, +) to (76, –, +) back to (36, –, 0), where the cycle may repeat.

Question 4 is a trick question, in that the answer is no but we shouldn't worry about it. Recall that de Kleer and Brown [6] originally proposed a very strong restriction on envisionments: Every physically possible behavior should be represented by a path through the envisionment, and every possible path through the envisionment should correspond to some physically realizable behavior. Kuipers [41] showed that, for `QSIM`, the former held but the latter didn't, and argued that requiring every path to be physically realizable was too strong a restriction. So the question really is, does our account fall prey to the same kind of problem? That is, are there paths through the envisionment which do not correspond to physically realizable behaviors?

The answer is yes: The envisionments produced by `CLOCK` include paths that do not correspond to physically realizable behaviors. For example, consider the instant of contact between the pallets and the scape wheel. On each cycle one of three outcomes is possible: (1) The pallets could smash into the scape wheel, forcing it backwards[4], (2) the momentum of pallets and scape wheel could exactly balance, or (3) the momentum of the scape wheel is high enough to throw the pallets immediately backwards. Now suppose we attempt to generate a history by walking through the envisionment, selecting an arbitrary outcome at each ambiguity (for under the restrictive notion of envisionments, each such path must be physically possible). The first time we reach the state of contact between the pallets and the scape wheel, any of the three choices is possible. But the physical implications of each choice are such that the only consistent possibility the next time

---

[3] These behaviors could be captured for pairs of gears in isolation by envisioning them without applying the abstraction operators, of course.

[4] This is the actual behavior of a recoil escapement under normal operating conditions.

around the cycle is the same choice. But by assumption each choice is independent, so our simple algorithm can generate paths which do not correspond to physically realizable behaviors.

This is not just an isolated example, and the problem is not confined only to cycles in the envisionment. Here is another specific behavior sequence which CLOCK's envisionment would view as legal behavior, but which real clocks fail to exhibit. During recoil the scape wheel is driven backward by the pallets by a small amount. (This is necessary for restoring energy lost by the pendulum.) Without knowing the magnitude of this collision, CLOCK would find nothing wrong with the hypothesis that the recoil force persists and drives the clock backward for an arbitrary length of time.

This of course does not mean that the envisionment is not useful. Given that it contains a description of a desired behavior, further knowledge can be used to ascertain if that behavior is in fact possible[5]. We consider one such task next.

3.7.1 A task: Understanding an explanation

One of the possible uses of a qualitative mechanical envisionment is to establish the plausibility of a design. This includes determining whether or not a desired behavior can occur. In this case we want to know if a path exists through the envisionment which exhibits "clock behavior". To do this we need a sample of the envisionment. Figure 15 describes a relevant subset of the envisionment, where the places are those referred to in Figure 14. We demonstrate that such a path exists through this sample of the envisionment, by following a chain of events described for a typical escapement by de Carle [4].

Here is the explanation given by de Carle of a similar escapement, interleaved with the traversal of the envisionment which corresponds to this description.

> "The escape wheel is rotating in a clockwise direction. As the pendulum swings to the left it allows the tooth *A* of the escape wheel to slide along the impulse face *B* of the pallet pad."

The contact relationship is satisfied by place-36, and to get these motions we must have both the scape wheel and the pallets moving clockwise, making the state (36, –, –).

> "Eventually the tooth *A* drops off the pallet pad."

This corresponds to taking the transition to state (33, –, –) and then the transition to (86,–, –).

> "The tooth *C* drops on to the pallet pad *D* and as the pendulum continues to swing to the left this locking becomes deeper and by reason of the curve of the pad the escape wheel is made to recoil."

While there are three possible transitions, the only one which includes the escape wheel recoiling is (44, +, –).

> "When the pendulum has reached the end of its journey and starts to return the escape tooth *C* will then give impulse to the pallet and so to the pendulum."

---

[5] Given a better physical theory, it may well be possible to develop a correct algorithm for generating only paths corresponding to physically possible behaviors. Consider: If we have full quantitative information we could always calculate the next state. Therefore such an algorithm does exist. The important open question is, what is the minimum information required to do it?

There are transitions to states (44, 0, 0) or (43, 0, 0) and then (43, –, +) which leads to (44, –, +).

"This cycle is repeated on the pallet pad *B*."

Briefly, the transitions are from (44, –, +) to (24, –, +) to (73, –, +) to (51, –, +) to (76, –, +). From this state there is another collision to (36, +, +). There are transitions to states (36, 0, 0) or (19, 0, 0) and then (19, –, –) which leads to (36, –, –), the starting state. This sequence of states is illustrated in Figure 16.

Applying a force in the (0 +) direction from place 19 demonstrates the necessity of feedback. Though the pallets cannot move any further counter-clockwise the transmission of force from the pallets to the scape wheel causes a reflective force from the scape wheel back to the pallets, pushing them clockwise. Without feedback information the net force is ambiguous so we might conclude that the pallets move clockwise, but by analyzing the feedback we determine that these forces exactly cancel.

Figure 16: Sequence of states corresponding to de Carle's explanation
A subset of the possible states and transitions of the QRG clock is shown here. States are depicted as representative configurations of a region given in the PV as well as annotations for motion. Symbolic names and motions correspond to the notation given in Figures 14 and 15.

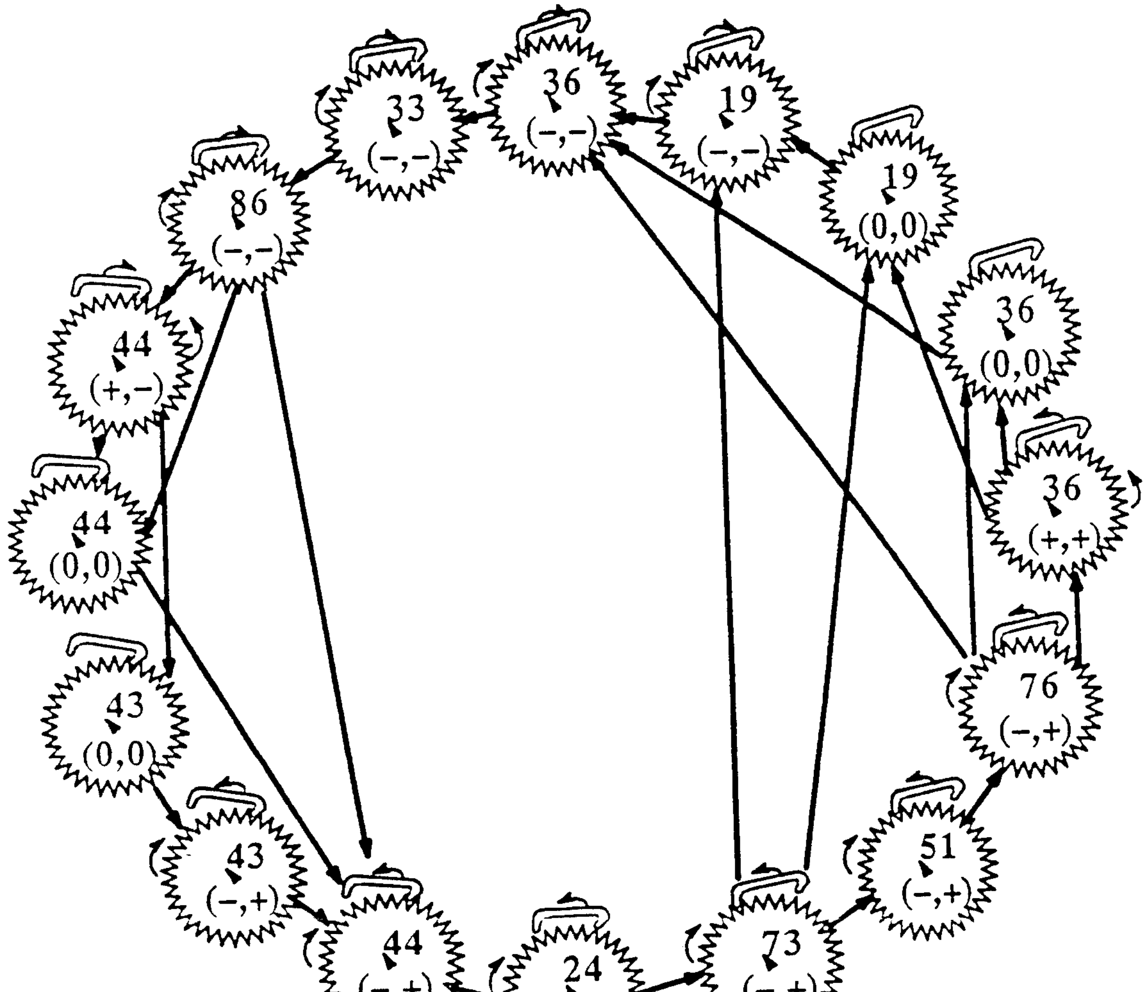

3.8 Lessons learned by building CLOCK

CLOCK and its constituent programs have been tested on a large number of examples, as illustrated in [10] and [55]. Using our qualitative spatial reasoning framework as a guide, we have been able to develop accounts which provide the basis for programs which can analyze mechanisms whose intricate kinematic interactions are well beyond the capabilities of earlier programs. This section documents some of the things we learned while building CLOCK.

We begin by examining how well the representations used in CLOCK satisfy the desiderata in Section 2.1. To the extent that the poverty conjecture is true for reasoning about mechanisms, CLOCK does reasonably well at satisfying the minimality desideratum. While precise shape information is required as part of its input, numerical values for masses and forces are not. Furthermore, the place vocabulary provides a useful quantization of kinematic state which elides irrelevant distinctions. CLOCK's methods also do a good job at providing composability, assuming that the vocabulary of line segments and arcs of circles provides an adequate approximation to the shapes being modeled. (Recall that the abstraction operations significantly reduce the complexity imposed by complex shapes; a very detailed approximation could be used to ascertain that each pair of objects can move smoothly together, and then the resulting place vocabulary smoothed down for tractability.)

There are ways in which CLOCK satisfies the explanation desideratum well, and ways in which it does not. The less interesting aspect of this is, perhaps surprisingly, the one most commonly used to evaluate systems, namely the ability to represent the rationale underlying its conclusions. The reason this is not very interesting is that any system which either uses some form of dependency records or can reconstruct dependencies has this capability. The way CLOCK happens to be implemented using an ATMS makes some queries very easy to answer. For instance, figuring out what the contributions to a net force are is quite simple. Unfortunately, explanations for some other kinds of queries would require additional computation. For instance, describing why a particular state transition is inconsistent would involve making explicit the closed-world assumptions underlying the interpretation construction steps used to build the set of mechanical states. Again, this is a consequence of how CLOCK was implemented, as opposed to a limitation of the spatial and physical representations it uses. Every implementation choice makes some things easy and some hard; if our goal is to generate great explanations, either the underlying style of implementation could be changed or a separate explanation generation system could be added that reconstructs the rationale from the existing data structures.

The most important way CLOCK satisfies the explanation desideratum is naturalness. We showed that the descriptions CLOCK produces are very natural, in the sense that they easily map into the intuitive explanations given by engineers for how mechanisms work. This is a knowledge-level property, and hence of more lasting importance.

Finally, CLOCK satisfies the integration desideratum through its use of qualitative vectors to represent directions, forces, and motions. Of critical importance is the metric diagram, which provides the substrate for computing the appropriate qualitative representations of shape and space according to the demands of the task.

While we view CLOCK as a success, we also discovered some very interesting problems in the course of building it. First, it was very difficult to get accurate descriptions of the part geometries. While we obtained some quite good approximations from photographs, these had to be corrected by trial and error using the results of the configuration space computation. For example, in the case of the escapement the distance between the escape wheel and the pallet has to be accurate to 0.1 % in order for the pair to function properly. This problem is inherent in the domain: The behavior of a mechanism is often quite sensitive to small perturbations in the shape of its parts. For example, the actual QRG clock contains a special lever in its frame which is used to adjust the spacing between the pallets and scape wheel with great accuracy. While the gear train can have a great deal of slop and the clock will still run, the adjustment of this lever must be made with great delicacy.

Floating-point truncation errors posed the second major problem. The computer used for the implementation (Symbolics) did not have a very refined floating-point arithmetic, and the program thus uses a rather large tolerance (10^{-7}) to allow for truncation error. However, it turned out (in the case of the ratchet) that the example required finer distinctions. We avoided the problem in this particular case by a very slight variation of the dimensions of the parts, but this is of course an unsatisfactory solution.

The third problem is that the algorithms we used for the computation of configuration spaces was very slow. With the exception of the ratchet, the computation for all the kinematics pairs had to be run overnight. More sophisticated algorithms, and faster (and more robust) floating point performance would solve this problem[6]. By contrast, building place vocabularies and adding dynamic information to the configuration space representations required less than five minutes per pair. Generating complete envisionments of the clock from the symbolic descriptions of each pair also required less than five minutes. Given recent advances in computing technology, we suspect that the envisionment of even very complex mechanisms is quite practical.

Aside from these problems, there are many incremental improvements one could make to `CLOCK`. For instance, the place vector representation offers a further advantage over a higher order configuration space. Given that the computation of configuration space for kinematic pairs is expensive, it makes sense to cache the configuration spaces for common kinematic pairs, perhaps in parameterized form (e.g., different combinations of gear ratios). Then only the interaction of unusual pairs would require new configuration space computations.

`CLOCK` embodies several assumptions that restrict the domain of mechanisms it can handle. First, the restriction to fixed-axis mechanisms means `CLOCK` cannot reason about systems such as linkages, which are found in many common machines. (This is discussed further in Section 4.) Second, the dynamics used in `CLOCK` is extremely simple, and does not support reasoning about energy. This could be fixed by integrating a more powerful dynamics, such as Qualitative Process theory [18]. Third, the restriction of metric diagram elements to be composed of line segments and arcs of circles means that more subtle mechanisms, whose purpose is to achieve motion described by a complicated analytic function as an output, cannot be completely verified. Overcoming this limitation would be substantially harder, since it requires a concomitant increase in the complexity of the symbolic algebra capabilities in the place vocabulary computation [9]. Finally, there are many situations where providing even the degree of specificity of the current metric diagram is asking too much, such as evaluating a "back of the envelope" sketch. Falting's notion of *kinematic topology* [14] provides a useful technique for such circumstances, but it needs to be integrated with richer descriptions to allow the results of earlier analyses to be used as a consistency check on later analyses of more fully specified designs.

---

[6] Such algorithms have been developed by Faltings' group, but are still undergoing experimental tests

4 Analysis of other spatial reasoning systems

Here we examine other spatial reasoning efforts and relate them to our framework.

4.1 Naive Physics

Some of our framework has its roots in the work of Pat Hayes on Naive Physics [26,27]. His seminal concept of histories was a key inspiration for our work, and his arguments about the locality of histories (i.e. things don't interact if they don't touch) indirectly suggest the connectivity aspect of the Connectivity/Shape hypothesis. We differ in our view of how rich and varied the spatio-temporal representations underlying histories must be, and see no clues in Hayes' work pointing to the poverty conjecture.

4.2 Representing space via axioms for topology

There have been several attempts to axiomatize spatial reasoning in a more or less topological fashion, but these have achieved very little. Shoham [62] developed a set of axioms for point contact involving completely constrained objects. While interesting from a standpoint of axiom-writing, such restrictions mean these axioms cover little of the phenomena. By contrast, Nielsen's formulation (Section 3.5, see also[54]) handles surface contact and partially constrained objects, and is powerful enough to support envisioning complex devices, as shown here. Davis [3] has developed axioms for special cases of motion, and has made an excellent case for the addition of non-differential, conservation-style arguments to qualitative physics. Like us, he argues that purely qualitative representations are "too weak" to support reasoning about motion involving solid objects. However, like Shoham, he has focused on developing an axiomatic formalism for a very small piece of the phenomena, rather than attempting to model a class of systems like mechanisms. We suspect the lack of breadth of coverage in these systems is due in part to their lack of any representational facility which plays the role of a metric diagram.

The most promising attempt to formalize topology to date is the work of Randell and Cohn [59], who are developing a general-purpose formal representation of space and connectivity. Their theory is already powerful enough to model certain spatial aspects of things that can occur in force pumps. Although much remains to be done, their system is an interesting start. Importantly, they are focusing on systems where the contact relationships are fairly constant (e.g., flows and changes in obstruction in fluid paths), which is exactly where a topological representation will have the most leverage. The poverty conjecture makes us suspect that their axioms cannot be extended to handle mechanisms without the incorporation of the equivalent of a metric diagram.

Similarly, Leyton's work on process grammar [44,45] provides a qualitative formulation of shape tuned to express the relationship between the spatial boundaries of an object and the physics which formed it. His theory is powerful enough to predict the changes in the shape of a deformable "stuff" if you squeeze it or push it in various ways. Recently Hayes and Leyton extended this grammar to incorporate discontinuities [28]. This work illustrates the power of purely qualitative representations in dealing with isolated shapes. But consider the following problem: We take two round clumps of stuff, and squeeze them in various ways. The process grammar can predict what their shapes will be, but this description of shape will not suffice to solve the rolling problem for them. Establishing the possible patterns of interactions between shapes requires more precise information.

4.3 FROB

The MD/PV model was first used in FROB [15,16], a program which reasoned about the motion of balls in a $2D$ world. This problem is an important subset of the spatial reasoning problems faced by any robot operating in a world of moving objects. For example, one should be able to quickly figure out that two balls thrown into the same well might collide, while two balls thrown into different wells cannot, unless at least one of them escapes. To answer such questions, FROB used both qualitative and quantitative information to provide as detailed an answer as possible given its state of knowledge. For example, collisions can be ruled out if the envisionments of the moving objects reveal that they are never in the same qualitative place. Providing more detailed information can lead to better answers. For example, if qualitative information alone indicated that two balls might collide, FROB could use constraint-based numerical simulation to demonstrate that they in fact do not.

Like CLOCK, FROB's metric diagram was based on a mixed symbolic/numerical model of the shapes of surfaces, but was much simpler, restricting balls to points and surfaces to straight lines. FROB's place vocabulary was a quantization of physical space, capturing distinctions imposed by surfaces and gravity. Since contacts were infrequent (i.e., collisions) and kinematically simple due to modeling balls as point masses, this quantization of physical space was in fact the configuration space for this domain.

However, FROB demonstrated some ideas that CLOCK does not. The reason is that FROB incorporated quantitative knowledge of its domain as well (i.e., the laws of projectile motion). In fact, FROB was the first program to demonstrate that a diagram could be used as a communication medium between qualitative and quantitative representations. The simplest such computation was FROB's use of its metric diagram to evaluate the spatial boundary conditions for quantitative simulation. FROB incrementally built a constraint network to represent the history of a ball. Extending the history of a ball in flight requires figuring out what, if anything, it will hit. FROB performed this computation by drawing a tentative trajectory in the diagram and calculating which surface (if any) the ball would hit first.

FROB's metric diagram also played a central role in more subtle combined qualitative and quantitative inferences. For instance, some spatial concepts, such as wells, can be represented by combinations of simple places. The ability to use such abstractions is crucial in predicting global behavior, namely bounding the future location of a ball. FROB also used spatial reasoning in performing energy analyses, casting quantitative estimates of the maximum energy of a moving object as a geometric constraint on the maximum height the ball could reach. FROB used this information to rule out the ball being in higher places, which in turn could rule out yet more potential locations. Such conclusions were used in reasoning about possible collisions between moving balls, allowing FROB to rule out possible collisions with a minimum of numerical simulation (often none).

4.4 Other spatial reasoning efforts in qualitative physics

The research of Joskowicz [32] shares many of our goals and fits well within the framework described here. He also uses configuration space to formulate a place vocabulary (he calls it a *region graph*). He has made several interesting proposals about how place vocabularies can be used. In [33] he proposed to analyze single interactions in a mechanism by recognition, describing kinematic pairs by patterns in configuration space. In [34] he proposed several abstraction techniques, some of which correspond to those already implemented in CLOCK [56]. His focus has been more on proving useful theorems about configuration space and algorithms for it, rather than exploring these ideas experimentally. We suspect that current versions of his formulation have certain limitations which will make them unsuitable for qualitative simulations of complex devices such as clockworks when they are tested. For example, his place vocabulary is annotated with labels corresponding to possible motions that do not include the directionality of mechanical contact. Without this information it is hard to see how they can be used to predict the propagation of forces through a mechanism.

Another line of research on reasoning about mechanisms which is accurately described by our framework is that of Gelsey [25]. His metric diagram is a constructive solid geometry CAD system, and his place vocabulary is the set of motion envelopes and kinematic pairs computed from this representation. His program uses a combination of heuristic kinematic analysis (to identify kinematic pairs) and numerical simulation to derive a qualitative description of particular behaviors of the device. The reliance on numerical parameters allows his system to always compute unambiguous descriptions of behavior. However, unlike CLOCK's qualitative simulation, quantitative simulation of mechanisms is a complex undertaking. It requires significantly more information on input (e.g., the weight of every component, numerical estimates of friction, initial velocities, etc.), substantially more computation (e.g., days of CPU time on a microvax to simulate just an escapement), and produces results which are held hostage to numerical errors and instabilities. By contrast, our qualitative mechanics formalism produces stable (although ambiguous) answers, far more quickly, and with less initial data. To be sure, similar requirements for exact shape descriptions and numerical issues affect CLOCK in the construction of the place vocabulary from the metric diagram. However, the abstraction operations simplify the description in ways that reduce the sensitivity to noise. Moreover, the coarse distinctions of the qualitative representations eliminate the need for detailed dynamical information. One potential use of the qualitative simulation in fact would be to guide a more detailed numerical simulation. For example, it would be interesting to build a FROB-like system, which combined the abilities of CLOCK to quickly analyze possible behaviors, and then using a simulation like Gelsey's to constrain its envisionment.

An alternate approach for modeling mechanical devices is to represent them by a fixed vocabulary of known interaction models. This is the approach adopted by Craig Stanfill [65] for the domain of pistons and cylinders. A typical example of how this approach can be applied to mechanisms is the work of Pu [58]. While the fixed vocabulary greatly simplifies modeling, it lacks the generativity of a first-principles geometric analysis like ours.

Some evidence that our framework applies to spatial reasoning not involving kinematics can be found in the work of Simmons [63] on qualitative and quantitative reasoning about geological processes. The problem he was solving was to test if a proposed sequence of geological events (such as uplift, deposition, and erosion) could account for a rock formation, as represented by a diagram indicating the positions and boundaries of various strata. The qualitative sequence of events was used to derive a series of transformations on elements of a diagram representing primordial rock, and the results compared against the diagram representing the measured state.

A formalism for modeling *kinematic topology* is described in [14]. While kinematic topology is much weaker than place vocabularies and does not allow envisioning kinematic behavior, it can be computed from a less precise, primitive-based object representation. In making a clear distinction between object model and functional model, this work follows the MD/PV model we argue for in this paper.

The deepest integration so far of symbolic descriptions of space and connectivity with traditional equational models for mechanism simulation has been achieved by Kramer's *degrees of freedom analysis* technique [39]. Roughly, the technique works as follows. Given a description of the parts of a mechanism in terms of mechanical constraints which must hold between them, the first step is to construct a plan for assembling it. This assembly plan is then used to guide the solution of equations governing the mechanism, leading to solutions that are both found more easily and yield better results than traditional numerical techniques. His system, TLA, differs from CLOCK in several respects. TLA is not restricted to fixed-axis mechanisms, as CLOCK is. TLA's goal is to produce a numerical simulation that can support visualization, rather than a qualitative simulation to support explanation. TLA represents objects essentially as mechanical links connected at points, since one of Kramer's goals is the automatic synthesis of linkages. This means it could not analyze the class of systems CLOCK can. In principle envisionments produced by CLOCK could be used to produce numerical simulations of behavior (e.g., by adapting the SIMGEN technique [22]). It would be very interesting to compare the results of such a system with TLA.

While Kramer's approach tightly integrates qualitative and quantitative representations, a purely qualitative theory of linkages has been developed by H. Kim [36]. She has extended Nielsen's qualitative mechanics by adding a richer notion of angle, including relative inclination as well as the quandrant. Her purely qualitative representation is successful on this task for two reasons. First, like Kramer, constraints imposed by the detailed shapes between the parts are ignored, essentially modeling links as one-dimensional objects. Second, she assumes that relative lengths between adjacent links of the mechanism are available. This is a reasonable stipulation given that the parts of a mechanism are known in advance, but it is clearly too restrictive for general reasoning about motion. Incorporating her theory of linkages into the accounts described here would be straightforward, since information about relative lengths can easily be computed from a metric diagram.

In related work, Kim has found a novel use for qualitative mechanics: generating descriptions of streamlines in two-dimensional laminar fluid flow situations [37]. Streamlines are routinely used by engineers to visualize how fluid flows through a system. To capture these intuitive laws, Kim uses an algebra of reflections that describes how fluid interacts with surfaces. These interactions then introduce new places in the space the fluid is flowing in. These laws are a first step towards qualitative reasoning about spatially distributed systems.

4.5 AI studies of physics problem solving

Cognitive science has spawned an entire genre of research on human problem solving in complex domains. Some of this work has focused on solving physics problems, sometimes involving diagrams. In some cases the diagram is not central, but used as a device to communicate results to human users (e.g., [57]). In others, the diagram is modeled mainly as an additional scratchpad to extend human short-term memory, rather than as a computational device to support spatial reasoning (e.g., [52]). None of these studies have used diagrams to support qualitative reasoning, and none have yielded systems capable of performing the kind of reasoning about motion that CLOCK can do.

The approach most like our framework is that of Simon and Larkin [64]. Like us, they view the ability to answer a class of spatial questions easily (e.g., via perception) to be an important advantage of diagrams. They also highlight an advantage we have not considered here, namely that diagrams can be used to group information about a problem and elements of it to reduce search. While our notion of metric diagram is intended to provide the same functional role as perception, the details of how our systems are organized make other means of indexing advantageous. We suspect, however, that there could be circumstances where computer problem-solvers will need to utilize these same advantages.

4.6 Robotics

While the concept of configuration space was originally develop to formalize the kinematic analysis of mechanisms, its development as a computational formalism is due to the robotics work of Lozano-Perez and others [46,47]. This work provided the algebraic formalization of configuration space constraints, which Faltings [9] extended to handle objects with curved boundaries. We expect that progress in robotics will lead to complementary progress in qualitative spatial reasoning.

Like our approach, the motion planning work of Schwartz & Sharir [61] is based on decomposing configuration space into a graph of regions. However, their formulation does not take into account the limits of constraint validity, and their methods result in an unnecessary combinatorial explosion of regions. While this may be acceptable for motion planning, it is not reasonable for a good qualitative description. Another difference is that they do not show any implementable algorithm by which their decomposition could be computed.

4.7 Route-finding

Finding one's way around in the world is an important spatial reasoning problem. Although the problems of route-finding and map-learning differ in many respects from the examples we have used here, it seems our framework captures at least some aspects of successful work in the area. For example, the *fuzzy map* representation of McDermott & Davis [51] can be viewed as a metric diagram, with the hierarchical region representation which structures it can be viewed as a place vocabulary. The work of Kuipers and Byun [42] explicitly separates metrical information from topological information, which corresponds to the metric diagram/place vocabulary split.

The QUALNAV system [43] incorporates a navigation algorithm based on a symbolic representation of space as a graph of *orientation regions*. In our terminology, the graph of orientation regions is a place vocabulary, while the exact location of landmarks which defines them is the metric diagram. Interestingly, in QUALNAV only required parts of the place vocabulary are computed. Furthermore, the metric diagram is only accessed by local observations of the robot. This makes QUALNAV particularly interesting in that it shows how it is possible to compute place vocabularies incrementally given only partial, local quantitative information.

4.8 Vision and Imagery

We suspect the MD/PV model is a good account of human spatial reasoning. There have been strong suggestions and evidence that the computations involved in vision and imagery are tightly linked, even to the extent of "shared hardware". Thus we would hope that further elaboration of the capabilities needed to perform the kinds of tasks we have explored could lead to some information-processing constraints on imagery and high-level vision.

The most direct connection to our work is the *visual routines* idea of Ullman [68]. Ullman proposes that the human visual system contains a *visual routines processor* which operates on the base representations produced bottom-up from lower-level visual processes, executing what are essentially programs to answer a variety of spatial queries relevant to recognition and other tasks. This is exactly the functional role of a metric diagram. The visual tasks he and his collaborators have studied to date have focused on recognition, however, so they have made no corresponding claims about qualitative descriptions of space.

There have been several attempts to produce computational models of visual imagery, some of which have emphasized the use of a diagram as an oracle for spatial questions. One of the original inspirations for FROB was a natural language understanding system developed by Waltz and Boggess [69], which argued for, and demonstrated, the use of something like a metric diagram for understanding sentences like "A fly is on the table" by constructing models. Kosslyn [38] developed a simulation of imagery which used an array scanned by a processing unit which was supposed to be an analog of a retina scanning a visual scene. Their goal was to fit psychological data concerning imagery, and unfortunately it seems no detailed account of how this model would aid spatial reasoning was developed. A similar model by Funt [23] was used to reason about stability of blocks world scenes. Both of these systems can be viewed as a form of metric diagram, although neither of them included aspects which could be interpreted as a place vocabulary.

In [30], Hinton describes some interesting phenomena that array-based models, like those of Kosslyn and Funt, appear to be unable to explain. Hinton's proposed alternative was to use a mixture of propositional and numerical representations, where the propositional description is computed from a metrical description, much as we compute place vocabularies from metric diagrams.

5 Discussion

While qualitative spatial reasoning is clearly a crucial part of qualitative physics, it has received relatively little attention. Part of the problem, we believe, is that it is fundamentally harder. Unlike qualitative dynamics, purely qualitative representations provide very little leverage. This is an unattractive conclusion. We know many people who have spent a great deal of time trying to build powerful, purely qualitative spatial theories. If the poverty conjecture is right most of this effort has been wasted. We ourselves would prefer for the poverty conjecture to be wrong. However, all our experience to date suggests it is correct. By putting this conjecture explicitly before the community, we hope that we can save other researchers time by alerting them to a potential tar pit. (And, perhaps, goad someone into providing us wrong!)

Our claims are not all negative, of course: We described the metric diagram/place vocabulary model, which provides an organization for spatial reasoning. The power of this model was demonstrated via the `CLOCK` program, which performed the first qualitative simulation of a mechanical clock (February, 1988), a milestone in qualitative physics.

`CLOCK` demonstrated that a quantitative spatial representation could be used as a substrate to compute a powerful qualitative representation of space. This point was illustrated earlier in `FROB` (c.f. Section 4.3), but only for a more limited domain. `CLOCK` showed that configuration space could be automatically quantized in a manner to form a useful qualitative description of the potential interactions of shapes. The ability to compute place vocabularies from a metric diagram was crucial, since it allowed the complexity of the place vocabulary to be controlled by the details of the shapes and it allowed the complexity and accuracy of the analysis to be controlled by varying the resolution. `CLOCK` also demonstrated that a powerful *qualitative mechanics* could be developed which captured intuitive notions of force, motion, and stability. We believe Nielsen's qualitative mechanics will find broad use in spatial reasoning tasks, since the notions of mechanical constraint are so central.

Research often raises more questions than it answers. Some ideas for specific extensions to `CLOCK` have already been described in Section 3.8. More importantly, though, the MD/PV model offers a new set of research questions and opportunities. We highlight some of these below.

Form of metric diagram: There is a spectrum of potential representations for metric diagrams. These range from representations based on analytic geometry, to combinations of numerical and algebraic expressions, to arrays. Little is currently known about which representations are useful for what tasks.

Dynamics: When is a qualitative state vector description versus a process-centered description appropriate? Are there other reasonable possibilities? How can the distinctions needed for qualitative spatial reasoning provide a foundation for formalizing spatial derivatives, so that spatially distributed systems can be modeled?

Theory of places: What are the commonalities underlying place vocabularies across various domains? It appears convexity, or at least quasi-convexity, is important. More empirical studies are needed to gain the insight required for a general theory

Links to vision and robotics: We view Ullman's theory of visual routines [68] in part as a theory of human metric diagrams. If Ullman is right, then people have metric diagrams. The obvious question is, do people use place vocabularies? (Some speculations on this question may be found in [17,13].) A better understanding of visual routines should lead to hints at building better metric diagrams. For example, we view Saund's work on scale-space blackboards [60] as an exciting set of ideas for significantly improving our ability to represent and reason about shapes.

6 Acknowledgments

This research was supported by the Office of Naval Research, Contract No. N00014-85-K-0225. The programming and analysis of examples was carried out while the authors were at the Department of Computer Science, University of Illinois at Urbana-Champaign. We thank John Collins, Johan de Kleer, Brian Falkenhainer, John Hogge, and the anonymous reviewers for their useful comments and suggestions.

References

[1] Allen, J. "Maintaining knowledge about temporal intervals", TR-86, Department of Computer Science, University of Rochester, January, 1981

[2] Brooks, R. A. "Solving the find-path problem by good representation of free space", *IEEE Systems, Man and Cybernetics*, SMC-13, 190–197, March 1983; also in *Proceedings of National Conference on Artificial Intelligence,* Pittsburgh, PA, August, 1982

[3] Davis, E. "A logical framework for solid object physics", New York University Computer Science Department Technical Report no. 245, October, 1986

[4] de Carle, D. *Watch and Clock Encyclopedia,* Bonanza Books, New York, 1975

[5] de Kleer, J. "Causal and teleological reasoning in circuit recognition", MIT AI Lab Technical Report no. 529, September, 1979

[6] de Kleer, J. and Brown, J. "A qualitative physics based on confluences", *Artificial Intelligence*, **24**, 1984

[7] de Kleer, J. and Bobrow, D. "Qualitative reasoning with higher order derivatives", Proceedings of the National Conference on Artificial Intelligence, August, 1984

[8] Faltings, B. "Qualitative kinematics and intelligent CAD," Proceedings of the 2nd IFIP WG 5.2 workshop on ICAD, North Holland, 1988

[9] Faltings, B. "Qualitative kinematics in mechanisms", *Artificial Intelligence* **41**(1), June, 1990

[10] Faltings, B. "Qualitative place vocabularies for mechanisms in configuration space", University of Illinois Technical Report No. UIUCDCS-R-87-1360, July, 1987

[11] Faltings, B. "A symbolic approach to qualitative kinematics", Proceedings of the Third International Conference on Fifth Generation Computer Systems, Tokyo, November, 1988

[12] Faltings, B. "A theory of qualitative kinematics in mechanisms", University of Illinois at Urbana-Champaign, Department of Computer Science Technical Report No. UIUCDCS-R-86-1274, May, 1986

[13] Faltings, B. "Visual representation for qualitative kinematics", AAAI Spring Symposium on Physical and Biological Approaches to Computer Vision, Stanford, March, 1988

[14] Faltings, B., Baechler E., and Primus, J. "Reasoning about kinematic topology", Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, August, 1989

[15] Forbus, K. "Spatial and qualitative aspects of reasoning about motion", Proceedings of the National Conference on Artificial Intelligence, Palo Alto, California, August, 1980

[16] Forbus, K. "A study of qualitative and geometric knowledge in reasoning about motion", MIT AI Lab [illegible]

[17] Forbus, K. "Qualitative reasoning about space and motion" in Gentner, D. and Stevens, A. (Eds.), *Mental Models*, LEA associates, inc., 1983

[18] Forbus, K. "Qualitative process theory", *Artificial Intelligence*, **24**, 1984

[19] Forbus, K. "The logic of occurrence", Proceedings of the International Joint Conference on Artificial Intelligence, August, 1987

[20] Forbus, K. "Interpreting observations of physical systems", *IEEE transactions on systems, man, and cybernetics*, Vol. SMC-17, No. 3, May/June, 1987

[21] Forbus, K. "The Qualitative Process Engine" in Weld, D. and de Kleer, J. (Eds.) *Readings in qualitative reasoning about physical systems*, Morgan-Kaufman, 1990.

[22] Forbus, K. and Falkenhainer, B. "Self-explanatory simulations: An integration of qualitative and quantitative knowledge", Proceedings of AAAI-90, August, 1990.

[23] Funt, B. V. "WHISPER: A compu[illegible] implementation using analogues in reasoning", Report 76-09, University of British Columbia, Department of Computer Science, 1976

[24] Gelsey, A. "Automated physical modeling", Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, August, 1989

[25] Gelsey, A. "Automated reasoning about machine geometry and kinematics", Third IEEE conference on AI applications, Orlando, Florida, February, 1987

[26] Hayes, P. "The naive physics manifesto", in *Expert systems in the micro-electronic age*, D. Michie (Ed.), Edinburgh University Press, 1979

[27] Hayes, P. "Naive physics 1: Ontology for liquids", in Hobbs, R., Moore, R. (Eds.), *Formal Theories of the Commonsense World*, Ablex Publishing Corporation, Norwood, New Jersey, 1985

[28] Hayes, P. and Leyton, M. "Processes at discontunities", Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, pp. 1267-1272, August, 1989

[29] H. Hertz, *Die Principien der Mechanik in neuem Zusammenhange dargestellt.* Leipzig, 1894.

[30] Hinton, G. "Some demonstrations of the effects of structural descriptions in mental imagery", *Cognitive Science*, Vol. 3, No. 3, July–September, 1979

[31] Joskowicz, L. " A framework for the kinematic analysis of mechanical devices", TR No. 313, New York University, 1987

[32] Joskowicz, L. "Reasoning about shape and kinematic function in mechanical devices", Ph. D. Dissertation, Courant Institute, New York University, 1988

[33] Joskowicz, L. "Shape and function in mechanical devices", Proceedings of the Sixth National Conference on Artificial Intelligence, August, 1987

[34] Joskowicz, L. "Simplification and abstraction of kinematic behaviors", Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, August, 1989

[35] Joskowicz, L. and Addanki, S. "From kinematics to shape: An innovative approach to design", Proceedings of the National Conference on Artificial Intelligence, August, 1988

[36] "Qualitative kinematics of linkages", Technical Report No. UIUCDCS-R-90-1603, Department of Computer Science, University of Illinois at Urbana-Champaign, May, 1990.

[37] "Qualitative reasoning about the geometry of fluid flow" *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*, Cambridge, MA, July, 1990.

[38] Kosslyn, S. *Image and Mind* Harvard University Press, Cambridge, Massachusetts, 1980

[39] Kramer, G. "Solving geometric constraint systems", Proceedings of AAAI-90, August, 1990.

[40] Kuipers, B. "Common sense causality: Deriving behavior from structure", *Artificial Intelligence*, **24**, 1984

[41] Kuipers, B. "Qualitative simulation" *Artificial Intelligence*, **29**, 1986.

[42] Kuipers, B. and Byun, Y.T. "A robust, qualitative method for robot spatial learning", Proceedings of the Sixth National Conference on Artificial Intelligence, August, 1988

[43] Levitt, T.S., and Lawton, D.T.: "Qualitative Navigation for Mobile Robots," *Artificial Intelligence* **44**(3), pp. 305-360, August 1990.

[44] Leyton, M. "A process grammar for shape", *Artificial Intelligence*, **34**, pp. 213-247, 1988

[45] Leyton, M. "Inferring cognitive history from shape", *Cognitive Science*, Vol. 13, No. 3, pp. 357-387, July–September, 1989

[46] Lozano-Perez, T. and Wesley M. "An algorithm for planning collision-free paths among polyhedral obstacles", Communications of the ACM, Vol. 22 pp. 560-570, 1979

[47] Lozano-Perez, T., Mason, M., and Taylor, R. "Automatic synthesis of fine-motion strategies for robots", MIT AI Memo 759, 1983

[48] Lozano-Perez, T. "Spatial planning: A configuration space approach", *IEEE Transactions on Computers* **C-32**, February, 1983

[49] K.J. MacCallum "Intelligent CAD Research Issues" in Intelligent CAD, I, North-Holland, 1989. (From: Proceedings of the IFIP TC 5/WG 5.2 Workshop on Intelligent CAD, Boston, MA. October, 1987.)

[50] Mason, M. "Compliance and force control for computer controlled manipulators", MIT AI Lab. Technical Report TR-515, April, 1979

[51] McDermott, D. and Davis, E. "Planning routes through uncertain territory", *Artificial Intelligence*, **22**, pp. 107-156, 1984

[52] McDermott, J. and Larkin, J. H. , "Re-representing textbook physics problems", in Proceedings of the Second National Conference of the Canadian Society for Computational Studies of Intelligence, Toronto, 1978

[53] Mukerjee, A. and Joe, G. "A qualitative model for space", Proceedings of AAAI-90, August, 1990.

[54] Nielsen, P. "A qualitative approach to mechanical constraint", Proceedings of the National Conference on Artificial Intelligence, August 1988; also in *Qualitative Reasoning about Physical Systems*, Daniel S. Weld and Johan deKleer Eds., Morgan Kaufman Publishers, Inc., 1990

[55] Nielsen, P. "A qualitative approach to rigid body mechanics", Ph. D. Thesis, University of Illinois, 1988

[56] Nielsen, P. "The role of abstraction in place vocabularies", Proceedings of the Conference of the Cognitive Science Society, 1989

[57] Novak, G.S. "Computer Understanding of Physics Problems Stated in Natural Language" Technical Report NL-30 Computer Science Department, The University of Texas at Austin, 1976

[58] Pu, P. "Qualitative simulation of ordinary and intermittent mechanical mechanisms", Ph. D. thesis, Computer Science Department, University of Pennsylvania, 1989

[59] Randell, D. A., and Cohn, A.G. "Exploring naive toplogy: Modeling the force pump", *proceedings of QPW-89*, Palo Alto, California, August, 1989

[60] Saund, E. "The role of knowledge in visual shape recognition", Technical Report No. 1092, MIT Artificial Intelligence Laboratory, October, 1988.

[61] Schwartz, J. and Sharir, M. "On the piano movers problem. II. General techniques for computing topological properties of real algebraic manifolds", Advances in Applied Mathematics **4**, pp. 298-351, 1983

[62] Shoham, Y. "Naive Kinematics: One aspect of shape", IJCAI-85, Los Angeles, August, 1985

[63] Simmons, R. "Representing and reasoning about change in geologic interpretation", MIT Artificial Intelligence Lab TR-749, December, 1983

[64] Simon, H. and Larkin, J. "Why a diagram is (sometimes) worth ten thousand words", *Cognitive Science*, 11:65-99, 1987

[65] Stanfill, C. "The decomposition of a large domain: Reasoning about machines", AAAI-83, Washington, D.C., August, 1983

[66] Williams, B. "Qualitative analysis of MOS circuits", *Artificial Intelligence*, **24**, 1984

[67] Williams, B. "The use of contunity in a qualitative physics", Proceedings of the National Conference on Artificial Intelligence, 1984

[68] Ullman, S., "Visual routines", in Pinker, S. (Ed.) *Visual Cognition*, MIT Press, Cambridge, MA, 1985

[69] Waltz, D.L. and Boggess, L. *Visual Analog Representations for Natural Language Understanding*, Proceedings of the Sixth International Joint Conference on Artificial Intelligence, Tokyo, 1979

[70] Weld, D.S. "Comparative analysis", *Artificial Intelligence*, **36**, (1988), 333-374.