Qualitative Kinematics in Mechanisms*

Boi Faltings

Laboratoire d'Intelligence Artificielle, Département d'Informatique, Ecole Polytechnique Fédérale de Lausanne (EPFL), Ecublens, 1015 Lausanne, Switzerland

ABSTRACT

Reasoning about shapes and their interaction is an important unsolved problem in artificial intelligence. In this paper, we present a theory of qualitative reasoning about kinematic interactions in fixed-axis mechanisms. A key idea of the approach is that a mechanism is represented not by the set of its parts, but by the pairwise interactions between them. Using such symbolic representations of kinematic pairs *as the basis for reasoning bypasses the difficult problem of finding a symbolic shape representation which is powerful enough for inferring kinematic properties.*

We introduce a qualitative representation of the kinematic function of pairs of parts, called a place vocabulary, *and show how it can serve as a spatial substrate for envisionments and causal explanations. By composition of place vocabularies, complex mechanisms such as mechanical clocks can be analyzed.*

1. Introduction

The function of a mechanical device depends on the way that motion and forces are transmitted through contacts between its parts, commonly called its *kinematics*. This paper presents a first-principles theory of *qualitative* kinematics. We introduce a representation, called the *place vocabulary*, which provides the complete set of kinematic states and inference rules to reason qualitatively about the kinematics of a fixed-axis mechanism. The theory is generative and applies to any two-dimensional mechanism whose parts can be described using straight lines and circular arcs.

The place vocabulary is a spatial substrate which allows envisioning the mechanism's behavior in response to external influences. In contrast to numerical simulation or observation of the device, using the place vocabulary makes it possible to give *causal* explanations of behavior, the basis for characterizing

*The research described in this paper was carried out while the author was with the Qualitative Reasoning Group at the University of Illinois.

Artificial Intelligence **44** (1990) 89–119
0004-3702/90/$03.50 © 1990 — Elsevier Science Publishers B.V. (North-Holland)

function. Furthermore, as each element of the place vocabulary is defined in terms of numerically precise mechanism configurations, there always exists a principled way of relating the symbolic representation to actual observations, a prerequisite for diagnosis and machine learning.

Similar to other work in qualitative physics [6, 7, 14, 15, 18, 20], the place vocabulary is qualitative in that it only distinguishes directions of motion and forces, not their magnitudes. Furthermore, contact points are distinguished only if they are either between different object parts, or result in different inference rules for the transmission of motion or forces. Because it represents kinematics at this qualitative level, the place vocabulary is *complete*: all possible qualitative behaviors are described in one symbolic structure. Such a finite functional representation is essential for problems such as verification and explanation of mechanisms or automatic computer-aided design.

1.1. The place vocabulary approach

Why has reasoning about kinematics in general proven to be such a hard problem? The main obstacle is finding an appropriate problem representation. We do not know of a general symbolic representation scheme for physical shapes that reasoning could be based on. In a purely symbolic shape representation, symbols represent entire classes of shapes, and numeric dimensions are abstracted away. But in the interaction of shapes, very fine distinctions are often crucial. Another problem is that shapes are multidimensional, and their metric dimensions influence the kinematic behavior in nonlinear and highly interdependent ways. The influence of the different metric dimensions cannot be decomposed as in many other domains of qualitative physics. Therefore, reasoning with intervals of quantity values is not successful for kinematics.

The key idea of the place vocabulary theory is that these problems can be avoided by using symbolic representations of object *pairs*, rather than individual objects. For a particular pair of objects, it is possible to distinguish the details of the shapes that are important in their interaction from those that are not. This makes it possible to construct a meaningful symbolic representation of this interaction. Complete mechanisms consisting of many parts can be analyzed by composition of the representations of the pairwise interactions within them.

In this paper, we introduce a qualitative representation of the kinematics of pairs of objects, the place vocabulary. By showing how the place vocabulary can be defined and computed in a principled way based on the configuration space of a mechanism, we develop a qualitative theory of mechanism kinematics.

The theory can be applied in two ways. First, for novel devices, place vocabularies can be computed based on the part geometries. Second, for common devices, such as gearwheels, place vocabularies can serve as a representation of knowledge about their kinematics. A particular interaction can then be analyzed by instantiating its place vocabulary from a kinematic knowledge base. As the place vocabulary is based on first principles, such a knowledge base could be constructed automatically as a collection of analyzed examples.

1.2. Reasoning about mechanism kinematics

The purpose of this section is to define the framework of mechanism kinematics and illustrate more precisely the concept of place vocabularies. Consider the mechanical clock shown in Fig. 1. Assume that we do not know how the device works and want to understand its function. The first step is to determine the components and their freedom of motion. In the case of the clock, each part is attached to the frame by a rotational joint, and therefore its only possible motion is rotation around a fixed axis. This property makes the clock an example of a *fixed-axis* mechanism. Parts may also be hinged to one

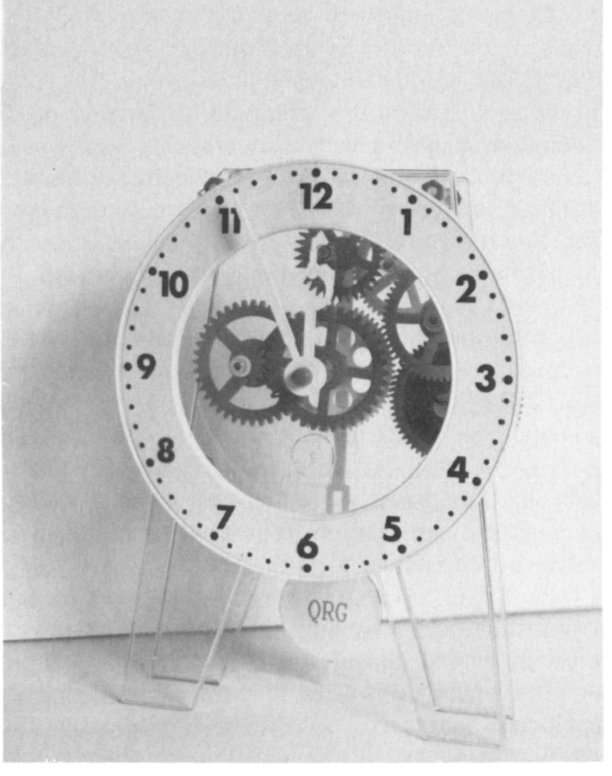

Fig. 1. The QRG clock.

another and thus give more complicated freedom of motion. The identification of what freedom of the parts results from particular types of joints and hinges is one of the subjects of research by Leo Joskowicz [19].

In a typical mechanism, each part can only interact with a small number of others. Once the components are identified, the next step is to determine the pairs of parts which interact with each other. Following the classic work of Franz Reuleaux [27, 28], these are called the *kinematic pairs* of the mechanism. Reuleaux distinguishes between two types of kinematic pairs: *lower* and *higher* pairs. In lower pairs, the objects touch each other at a pair of surfaces, and this contact is maintained throughout the possible motions of the objects. There are only six possible types of lower pairs, which correspond to revolute, prismatic, helical, cylindric, spheric and planar joints between objects [27, 28]. Their symbolic representation is straightforward, and there exists an exact mathematical analysis [5] that can be used as a basis for reasoning about them. In higher pairs, the contact can vary during the motion of the objects, and there are an infinite number of them. Examples of higher pairs are cams, ratchets, escapements, and pairs of gearwheels. In this paper, we restrict ourselves to the analysis of higher pairs only.

A mechanism's kinematic pairs are arranged in sequences, called *kinematic chains*, which transmit the motion of an input part to that of an output part. For example, the clock contains a kinematic chain of gearwheels which transmits the motion of the escape wheel to the hands of the clock. The function of a kinematic chain is given as a composition of the functions of its constituent kinematic pairs. In the case of gearwheels, the function of each pair is to transmit rotation, reversing the direction and changing the speed by a certain ratio. The function of a chain of gears is given as the composition of these ratios into a single ratio of transmission from the input to the output gear.

Gearwheels are common enough that we could assume that the function of gearwheel pairs could be a primitive in a theory of kinematics. However, a clock also contains a kinematic pair whose function is very specific and cannot reasonably be assumed to be a kinematic primitive: the escapement shown in Fig. 2.

The escapement is a kinematic pair which consists of a *pallet* and an *escape wheel*. Both are hinged so as to allow rotational freedom only. The pallet is rigidly connected to a pendulum and swings back and forth, while the escape wheel is driven by a spring to turn clockwise. The escapement is the heart of the clock, with the function of allowing the wheel to turn by one tooth each time the pallet oscillates. As the period of oscillation of the pendulum and pallet is constant, the wheel turns at a constant rate and can be used to drive the hands of the clock. A natural explanation of the function of the escapement is that of a sequence of states, as shown in Fig. 3. The states A through G illustrate the duty cycle of the escapement. In state A, the wheel is held back

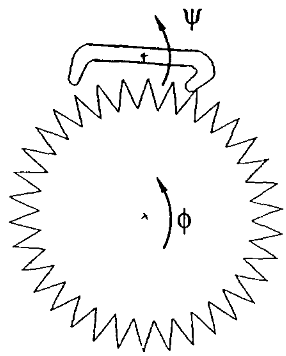

Fig. 2. The escapement of the QRG clock.

by the contact with the right end of the pallet. The pallet is turning counter-clockwise, and lifts off from the wheel. Under the influence of the spring, the wheel then turns, as shown in state B, until it is stopped by contact with the left end of the pallet, as shown in state C. The pallet reaches the extremum of its oscillation and starts to turn clockwise, as shown in state D. It again lifts off the wheel, allowing it to turn, until it is stopped by contact with the right end of the pallet, reaching state F. The pallet reverses its motion in state G, and the cycle recommences. The wheel has turned by one tooth during the cycle.

The explanation we have just given is similar to what can be found in mechanism books. It explains the typical qualitative behavior of the kinematic pair in terms of kinematic states and directions of motion and force. In order to construct such an *envisionment* of the behavior, we need a model of the possible kinematic states, the inference rules they define, and the possible

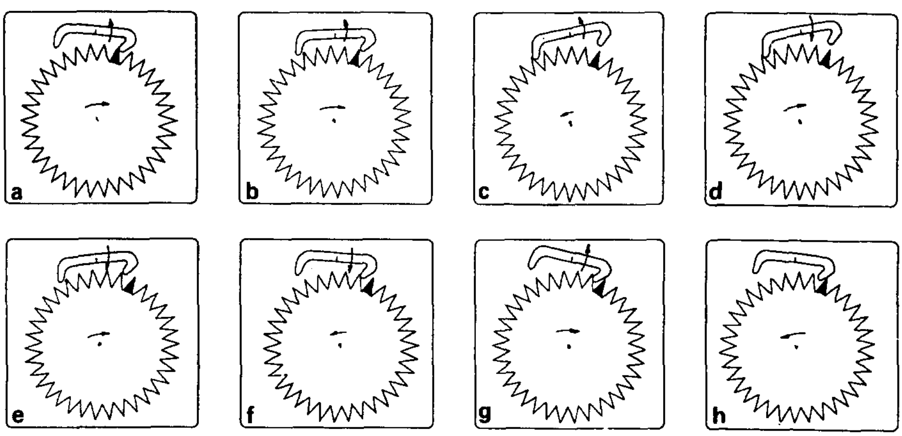

Fig. 3. Different states in the duty cycle of the escapement. One of the teeth of the wheel is darkened to indicate the rotation.

transitions between them. In this paper, we propose the *place vocabulary* as a representation that can serve as such a model. Each set of configurations which share the same kinematic state and the same qualitative inference rules is represented as a *place*. Possible transitions are represented by arranging the places in a graph in which each edge is labeled with the direction of motion that corresponds to the transition. The place vocabulary is a spatial substrate sufficiently powerful to allow theories of qualitative physics to generate an explanation such as that given above,[1] including the causal interpretations of the predicted behavior. In recent work, Paul Nielsen has computed a complete envisionment of the behavior of the clock shown in Fig. 1 [25].

The place vocabulary represents only the possible contact relationships, not the actual motion of the parts. As a result, different states in the envisionment may share the same place, but be distinguished by the direction of motion of the parts. This is the case for states A and F in the example. Furthermore, the place vocabulary also includes states that cannot be reached given the way the parts move in normal operation, but would be reachable given other motions of the parts. For example, the configuration shown in state H can be reached by counterclockwise motion of the wheel from state G, but this does not happen in normal operation.

We have developed and implemented algorithms for computing place vocabularies from shape descriptions. The implementation covers the kinematic interactions of objects of arbitrarily complicated two-dimensional shapes with boundary curves composed of straight lines and arcs (segments of circles). Furthermore, it is restricted to the analysis of kinematic chains of higher kinematic pairs in fixed-axis mechanisms. The restriction of the types of boundary curves and the number of dimensions was imposed in order to simplify the implementation, and is not essential to the theory or the algorithms. The restriction to higher kinematic pairs, on the other hand, is essential in order to efficiently compute place vocabularies, although it too is not a fundamental requirement for the theory itself.

1.3. Related work on reasoning about shapes

Because of the difficulties outlined earlier, previous work on reasoning about shapes has used one of the following two approaches:

- reasoning with symbolic representations of a fixed and restricted set of objects,
- reasoning using numerical approximations.

As an example of the first approach, a CAD system might have a finite set of symbols to describe the standardized types and sizes of nuts and bolts along

[1] Provided a qualitative model of external influences, in this case, the oscillation of the pendulum attached to the pallet.

with rules that state how these symbols can fit together. The problem is that this approach lacks generality. Only a predefined and limited set of shapes can be represented, and the possible motions may also have to be restricted. Craig Stanfill [30] has used this approach in a system that reasons about interactions of shapes restricted to cylinders. More recently, Andrew Gelsey [17] has developed a system for analyzing mechanisms consisting of kinematic pairs of a fixed vocabulary of types which includes joints, gears and cams. In Gelsey's system, the models of the known kinematic pairs are used to obtain the algebraic relations between the motions of the mechanism's parts. These relations are then reasoned about to obtain an analysis of the mechanism's behavior. While his system is quite capable of reasoning about the complicated shapes that occur in mechanisms, the set of shapes is limited to those for which an analysis has been defined in the program.

The second approach, using numeric processing, has been better developed; two examples are the ACRONYM system of Rodney Brooks [3] and the MERCATOR system of Ernest Davis[2] [4]. In both systems, the multidimensional problem is reduced to reasoning with scalar numeric quantities. This reduction works only for problems in which the influence of the parameters can be decomposed, and neither system is capable of reasoning about complex interactions such as those which occur in mechanism kinematics. Another variant is the use of numerical simulations of analogical models, such as in [16]. While this approach is very general, it lacks the completeness of the qualitative representation, and cannot provide any causal explanations for behavior.

2. Place Vocabularies

In this section, we define place vocabularies in two different ways. First, we give a *functional* specification, concentrating on the use of the place vocabulary for reasoning tasks. We then show how this functional specification can be satisfied by a definition of the place vocabulary as a decomposition of *configuration space*. Finally, we illustrate the place vocabulary for the escapement example presented earlier.

2.1. A functional description

The representation we proposed in the introduction, a graph of kinematic states, can be formally defined as the *envisionment* [8, 15] of the behavior of a kinematic pair. Each state in an envisionment is characterized by a particular combination of signs of the dynamic parameters, and a particular set of relations between them. In the case of mechanisms, the dynamic parameters

[2] Although MERCATOR reasons about spatial relations, not kinematics, it addresses a similar geometric problem.

are the positions of the two objects in the kinematic pair, and the forces they exert on each other. The relations between dynamic parameters are given by the type of *contact* between the objects.

A full explanation of a device's function usually depends on its dynamic behavior and therefore has to include assumptions about the forces and movements imposed externally on the parts. The device's *place vocabulary*, on the other hand, describes only the influence of the object geometries; it makes no assumptions about external influences. It is an *undirected* graph of *places*, where each place corresponds to a different type of contact between the objects. This representation of the changing contact relationships in a mechanism defines the qualitative relations between the dynamic parameters in the different configurations. It thus provides the *spatial substrate* upon which the full envisionment can be built.

The distinction between place vocabulary and envisionment is made for reasons of plausibility. The assumptions about external forces necessary for computing the envisionment depend on the context of use of the device, and involve high-level reasoning. The kinematic analysis of shape interactions, on the other hand, is closely tied to their visual perception and should be considered a low-level process. The place vocabulary is an explicit model of the information provided to the high-level reasoner by low-level processes. The place vocabulary idea itself was first proposed by Ken Forbus in [13, 14]. This paper describes the extension of the original concept, which was limited to representing motion of point masses, to kinematic interactions of geometric objects with finite dimensions.

As an example of a place vocabulary, let us return to the escapement shown earlier in Fig. 2. Each of the qualitatively different contact relations, for states A (= F), B (= E), C (= D), G and H are represented by a different place. Based on the contact configuration, there are three types of places:

- places where no contact between the two objects exists,
- places where the objects touch in a single point,
- places where the objects touch in two different points.

In the first case, no relation between the dynamic parameters can be asserted. In the second case, the objects can mutually push each other and thus move in coordinated motion. Except in singular cases, the third case corresponds to a single particular configuration of the objects, reached in the transition between two places which each exhibit one of the two contact points.

In the general case, a place where a single contact between the objects exists is characterized by two inference rules each for the net force and motion. As an example, consider the place which underlies the kinematic states C and D in Fig. 3. Letting $[\delta x]$ indicate the sign of the derivative of quantity x, and ϕ and ψ be the rotation angles of the wheel and the pallet as shown in Fig. 2, we can assert that

$$[\delta\psi] = + \quad \Rightarrow \quad [\delta\phi] = + \,,$$
$$[\delta\phi] = - \quad \Rightarrow \quad [\delta\psi] = - \,.$$

Similarly, for the forces, a positive net moment on the pallet will cause a positive net moment on the wheel, and a negative net moment on the wheel will cause a negative net moment on the pallet. In mechanism analysis, friction between parts in kinematic pairs is minimized and can be ignored. In this case, the relations between the forces or moments are always identical to the relations between the derivatives of the object positions. Note the unidirectional character of the inferences: it reflects the intuitive fact that one can push on a contact point, but not pull. The inferences in mechanism analysis thus have a clearly defined *causal* direction.

The places are arranged in a graph, the place graph, which reflects the possible transitions between the places. Because no assumptions about the dynamic behavior are made, a transition is possible if and only if there exists some coordinated motion of the two objects which results in a change from the contact configuration of the first place to that of the second. A possible transition between places is characterized by two elements:

- the limit points of one or both object positions where the transition can occur,
- the direction of motion which can cause the transition.

The positions of the limit points can be analyzed to reduce ambiguities in the envisionment; this is discussed later in the paper. The second element, the direction of motion, governs the dynamic behavior. Each transition in the place vocabulary is associated with a list of qualitative directions of motion which may allow the transition. In the dynamic analysis, a transition from one place to another is possible whenever the actual qualitative direction of motion is among those that allow the transition.

The place vocabulary thus determines the qualitative relations between dynamic parameters that are required to reason about the mechanism's behavior. In this section, we have concentrated on this functional aspect of the place vocabulary. In the next section, we will show how such a place vocabulary can be formally defined as a decomposition of the mechanism's *configuration space*. This formal definition shows how the characteristics of the places can be computed, and defines the precise semantics of the representation.

2.2. Place vocabularies in configuration space

In its original formulation in the FROB system [13, 14], the place vocabulary was used to represent motion of point masses in a two-dimensional space bounded by straight lines. As the position of point masses is completely represented by their coordinates in physical space, the sets of qualitatively equivalent positions were just regions of the physical space. The place vocabulary used in FROB consisted of a graph representing the decomposition of the physical space by horizontal and vertical straight lines.

In this paper, we show how a similar idea can be applied to analyze the kinematics of objects of finite dimensions. The kinematic constraints given by the condition that physical objects may not overlap can be expressed as constraints on a point moving in *configuration space* [2, 10, 23]. The configuration space can be decomposed into regions of equivalent configurations of the objects. Interestingly, the place vocabulary specification of the previous section corresponds exactly to the most natural decomposition of the configuration space into a *cell complex*. This provides the necessary framework for the computation of the place vocabulary, and defines the semantics of the representation.

2.2.1. *Configuration space*

The position of a physical object can be described by a small set of parameters. In the case of unrestricted motion in three dimensions, an object's position is completely specified by three Euclidean position parameters and three orientation parameters. The motion of mechanism parts is restricted by joints, and with very few exceptions a single parameter suffices to specify the position of a part. A configuration of a set of objects is defined as a particular combination of positions for each object and can be specified as a vector of all the position parameters. The space of all such vectors is called the *configuration space* [2, 10, 23], or C-space. Each dimension of C-space corresponds to a different position parameter of one of the objects.

Kinematic interactions arise from the fact that physical objects may not overlap. The configuration space can be divided into two sets of points:

- legal configurations where no overlap between objects exists,
- illegal configurations where an overlap exists.

We call the union of all legal regions the *free space* and that of the illegal regions the *blocked space*. The condition for a configuration to fall on the boundary between free and blocked space is that there is both a configuration in blocked space and a configuration in free space that can be reached by an arbitrarily small motion. This is the case if and only if the two objects are in contact but do not overlap.

There are two possible ways that plane objects O_1 and O_2 can touch:

- A vertex of O_1 touches a boundary segment of O_2, or vice versa. Note that this subsumes the case where a pair of vertices touch each other.
- A boundary segment of O_1 touches a boundary segment of O_2.

Each possible instance of these cases defines a *configuration space constraint* in the configuration space for O_1 and O_2. We call the first *vertex constraints*, and the second *boundary constraints*. A configuration space constraint is an algebraic curve containing all the configurations where the defining condition of touch is satisfied. As each point of touch rules out motion in the direction of the surface normal at the point of touch, it reduces the freedom of the objects by 1 degree. The dimensionality of the constraints is thus exactly 1 less than that of the configuration space itself. The constraints are defined by the condition that a certain point on one object satisfies the equation of the algebraic curve of the boundary on the other object. As the boundary in general covers only a part of the complete curve, the constraint curve contains a segment where the constraint is applicable and a part where it is not. We use *applicability constraints* [10] to delimit the region of configuration space where the constraint curve is applicable. The second condition for a configuration to fall on the boundary between free and blocked space is that there may be no other overlap between the objects. The actual boundary thus consists of the *envelope* of the applicable segments of the constraints. The inside of the envelope consists of blocked space; constraint segments are said to be *subsumed* whenever they fall within this region.

The idea of configuration space was originally invented by Heinrich Hertz and has been used in mechanical engineering and physics as an aid for formalizing complex motions and kinematics. As a computational device, it was developed by Tomas Lozano-Pérez [2, 22, 23] for the problem of robot motion planning, with further extensions by Bruce Donald [10]. These earlier computational formulations of configuration space restricted themselves to polyhedral approximations, for which only vertex constraints were needed. However, when we consider qualitative kinematics, an approximation of a curve by line segments results in spurious discontinuities of behavior. A wheel that is approximated by a polygon will not run smoothly, contrary to its actual behavior. To avoid such bogus qualitative descriptions, we have extended the theory to allow curved boundaries as well.

As a mechanism consists of a large number of parts, its configuration space has a large but finite number of dimensions. However, as any given part only interacts with a small number of others, the problem can be decomposed. In particular, we consider a mechanism to be a kinematic chain of pairwise interactions. Each of these *kinematic pairs* consists of two objects with one degree of freedom each, which define a two-dimensional C-space. Analyzing this configuration space is a problem of plane geometry. Subsequently, the analyses for kinematic pairs can be composed into a complete analysis of the kinematic chain, as described later in the paper. In the following definition of place vocabularies, we refer to the analysis of kinematic pairs in two-dimensional configuration spaces.

The topology of the configuration space is not always Euclidean. Rotational freedom results in a parameter which wraps around, homotopic to a circle.

This results in the three different cases of the configuration space topology being:

- a Euclidean plane, when both objects have translational freedom;
- a cylinder surface, when one object translates and one rotates;
- a torus surface, when both objects rotate.

In the general case, when the objects have more than a single degree of freedom, more complicated topologies are possible, but are not treated by the current implementation of the theory.

2.3. Place vocabularies: A qualitative representation of C-space

Each configuration space constraint is defined by a point of touch between the objects, and defines a one-dimensional curve in the two-dimensional configuration space. Recall that the place vocabulary contains three different types of places, corresponding to zero, one or two points of contact between the objects. In configuration space, each point of contact is specified by a constraint, and so the three types of places correspond to:

- two-dimensional regions of free space, where no contact exists,
- segments of constraint curves, where one contact exists,
- intersections between constraint curves, where two contact points exist.

Because of their form in configuration space, we often refer to the different places as full-dimensional faces (FULLFACES), constraint segments (CSEGs) and intersection points (IPs).

The configuration of the mechanism can leave a place either by establishing a new point of contact or by breaking one. Therefore, all adjacencies in the place vocabulary are between places whose number of contacts differs by one. In configuration space, the two-dimensional regions are bounded by one-dimensional constraints, which are in turn bounded by their zero-dimensional intersections. In mathematical topology, such a decomposition is called a *cell complex*, the most general way of decomposing a space into contiguous regions.

Two configurations can be considered equivalent only if the qualitative relations between the dynamic parameters are equal at both configurations. This provides the second criterion for the decomposition of configuration space. Based on the precise definition of the inference rules given in the section below, it turns out that this additional criterion amounts to a *monotonicity* condition on the constraint segment, which sometimes requires additional breakups of the constraints.

The adjacencies of the places are already given by the arrangement in the cell complex. The directions of motion which may allow a transition between the places are given by the set of all directions which can transverse the boundary between them in configuration space.

2.4. Qualitative motions: Pushing and support relationships

We specify motion of the mechanism by its *qualitative direction* in configuration space. A qualitative direction is given by a pair of signs (S_1, S_2), where each sign can be either $+$, 0 or $-$, indicating the direction of change of the motion parameter of the corresponding object. There are eight different qualitative directions, as indicated in Fig. 4. We distinguish between *precise* directions, where one of the components of the vector is 0, and *fuzzy* directions, where both of the components are either $+$ or $-$.

The qualitative relations between motion parameters apply only to places with one point of contact, which specify a coordinated motion of the two objects that maintains the contact. The direction of this coordinated motion is given by the qualitative direction of the constraint curve in configuration space.

We define the qualitative direction of a constraint at a configuration as the qualitative direction of the vector which is tangent to the constraint and directed so that blocked space lies on its right. If we denote the two parameters of the configuration space as x and y, this convention implies that a *fuzzy* qualitative direction (a, b) defines the following inference rules:

$$\begin{aligned}[\delta x] = b \quad &\Rightarrow \quad [\delta y] = a\,,\\ [\delta y] = -a \quad &\Rightarrow \quad [\delta x] = -b\,.\end{aligned}$$

For a *precise* qualitative direction (a, b), the possible inference is given by the constraint:

$$\begin{aligned}[\delta x] &\neq b\,, \quad \text{if } a = 0\,,\\ [\delta y] &\neq -a, \quad \text{if } b = 0\,.\end{aligned}$$

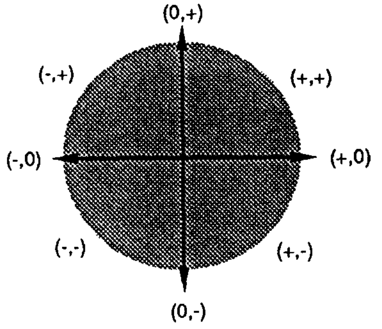

Fig. 4. The eight possible qualitative directions. Note that the four vectors containing 0 indicate a precise direction, whereas the other four describe all directions within a sector, a fuzzy direction.

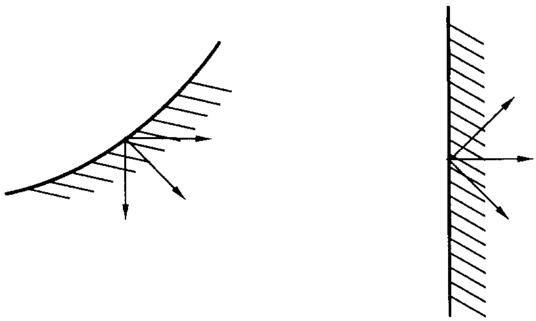

Fig. 5. A constraint rules out three different qualitative directions of motion. On the left, a constraint with a fuzzy direction, on the right, a constraint with a precise direction.

In configuration space, the inference rules have the effect of ruling out precisely the three qualitative directions of motion that a constraint makes illegal, as shown in Fig. 5. The formulation as inference rules makes explicit the causal interpretation of the motion. In the case of a fuzzy direction, we can *push* one object with the other, as expressed by the two inference rules. A constraint with a precise direction expresses the concept of *supportedness*.

It has been shown [11, 24] that forces and moments in physical space are equivalent to forces on the corresponding point in configuration space. In the case of a constraint with a fuzzy direction, the two inference rules for the qualitative directions of motion apply equally to the transmission of forces or moments through the contact point. In the case of a constraint with a precise direction, where one of the components is zero, no force can be transmitted. The constraint on the motion defined by the constraint can be interpreted as a rule that an arbitrary force in this direction is *absorbed* by the contact. These two relations between forces express the intuitive concepts of pushing and supportedness in a precise way, based on first principles.

The relations between motion parameters and forces must be constant within each place. Because both are given by the qualitative direction of the constraint, this direction must be constant within each place. The constraint segments corresponding to places therefore must be *monotone* subsegments.

Consider a constraint formed by a vertex on object O_1 touching a boundary segment on object O_2, and let the position/orientation parameters of the objects be p_1 and p_2. As the constraints are well-behaved algebraic curves with a continuous derivative, the monotonicity requirement will be satisfied whenever a constraint has no zero crossings of the derivatives $\mathrm{d}p_1/\mathrm{d}p_2$ and $\mathrm{d}p_2/\mathrm{d}p_1$ except possibly at the endpoints. The presence of direction changes on a constraint can be detected directly by making tests on the object dimensions. Additional subdivisions can then be introduced at each such direction change. The tests and their derivation are described in detail in [12].

2.5. Adjacencies and their classification

The adjacencies of the places are defined by the adjacencies of the corresponding regions in configuration space. To provide the necessary information for the dynamic analysis, we need to label each adjacency with the directions of motion that will allow a transition along it. We now describe how these directions are found for the three different cases of transition from a constraint segment, an intersection point, and a two-dimensional region.

Transition from a constraint segment into the adjacent free space region is possible in three qualitative directions of motion: along the normal pointing into free space, and along the two directions adjacent to it. This is illustrated in Fig. 6.

The normal vector can be found from the qualitative direction of the constraint segment. Motion in the direction opposite to these three vectors is ruled out by the constraint. When the motion is in one of the remaining two directions, parallel to the constraint, there exists an ambiguity: the motion may either remain on the constraint, eventually leading to a transition to one of the intersection points at the ends of the constraint, or move off the constraint into free space. This ambiguity cannot be avoided except by variation of the coordinate system to make one of the axes parallel to the constraint. It does not exist if the constraint is parallel to one of the coordinate axes.

A zero-dimensional place is a point in configuration space. The point is the intersection of two constraint segments, and there are thus three adjacent places: the two constraint segments and the region of free space. In the most ambiguous case, all of them fall in the same qualitative sector and thus can be reached by the same qualitative motion. We then have three possible transitions for that particular direction of motion, as shown in Fig. 7.

In the case of two-dimensional places, there can exist considerably more ambiguity. Consider the situation in Fig. 8. In the case of P_0, motion in the $(0, -)$ direction (straight down in the diagram) can lead either to P_1, P_2 or P_3. Because there can be an arbitrary number of constraint segments at the lower boundary of P_0, this ambiguity can become arbitrarily and unacceptably large. It can be reduced by breaking up P_0 using additional divisions parallel to the coordinate axes as indicated by the dotted lines. The divisions can be under-

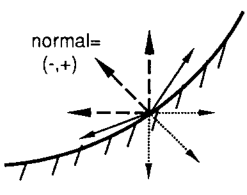

Fig. 6. The possible transitions from a one-dimensional place.

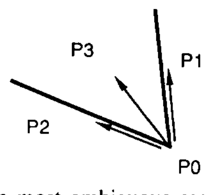

Fig. 7. The most ambiguous case of transitions from a zero-dimensional place.

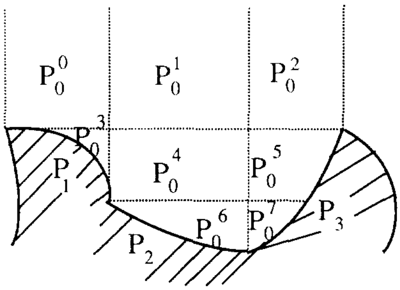

Fig. 8. Example of places in confuguration space. Qualitative directions are specified as pairs of (x, y) in the coordinate system shown.

stood as landmark values in each of the two parameters, defined by the coordinates of the zero-dimensional places and points of constraint non-monotonicity. Each region corresponds to a legal combination of intervals between two landmark values. From each of these regions, motion in any of the four exact directions $(0, +)$, $(0, -)$, $(+, 0)$, $(-, 0)$ will lead to a unique place transition, and motion in one of the remaining four directions can result in at most two different place transitions. This latter ambiguity arises because it is not known which landmark value will be crossed first. As an example, from place P_0^6, we find the transition table shown in Table 1.

Except in a special case, the region structure can be assembled by purely combinatorial methods given the orderings of the points in each of the two coordinates and the adjacencies in the place vocabulary. An ambiguity occurs when there are divisions in both coordinates which intersect the same constraint segment. However, the dynamic predictions that can be made in each of the two possible cases differ only in very esoteric cases.

2.6. Example of a place vocabulary

Now that place vocabularies have been defined, it is time to look at a complete example. Consider the escapement mechanism explained in the introduction,

Table 1
The transitions out of P_0^6 for motion in the different qualitative directions

Direction	Possible transitions
$(-, -)$	$\{P_2\}$
$(-, 0)$	$\{P_2\}$
$(-, +)$	$\{P_2, P_0^4\}$
$(0, +)$	$\{P_0^4\}$
$(+, +)$	$\{P_0^4, P_0^5, P_0^7\}$
$(+, 0)$	$\{P_0^7\}$
$(+, -)$	$\{P_2, P_0^7\}$
$(0, -)$	$\{P_2\}$

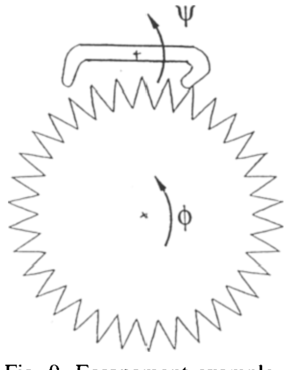

Fig. 9. Escapement example.

again shown in Fig. 9. In this section, we show how the place vocabulary we referred to in the introduction is derived based on the configuration space of the mechanism, shown in Fig. 10. The two parameters spanning the configuration space are the orientation ϕ of the wheel in the horizontal and the orientation ψ of the pallet in the vertical direction. The space shown is a torus surface; its left end connects to the right end and its bottom connects to its top. The shaded region in the configuration space is the blocked space within which all configurations are illegal. The boundary between free and blocked space is made up by segments of constraints corresponding to different configurations of touch between the objects. Note that this configuration space contains two

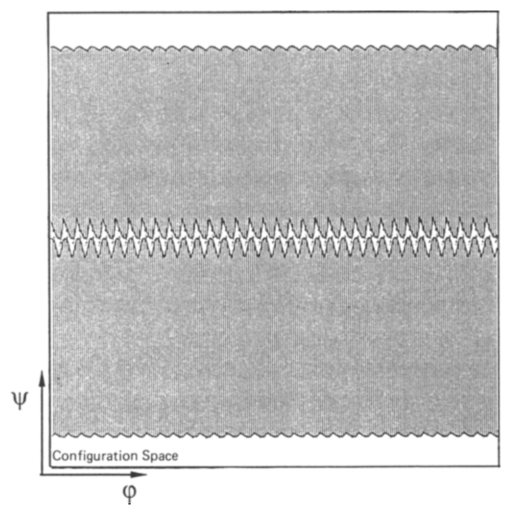

Fig. 10. The configuration space for the escapement example. The horizontal parameter is the orientation ϕ of the wheel, the vertical parameter the orientation ψ of the pallet. The indicated square region is described in detail later in the text.

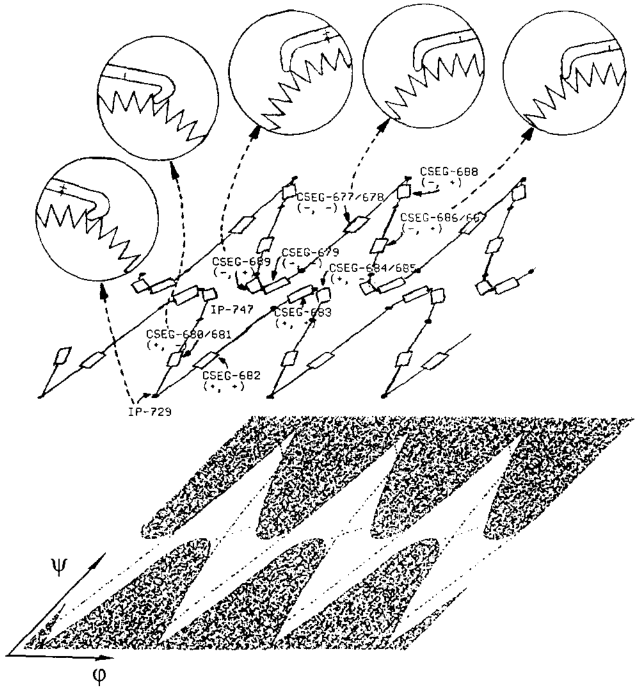

Fig. 11. Derivation of the place vocabulary for a region of configuration space.

disjoint regions of free space. The region in the center corresponds to the normal operation of the escapement. In the other region's configurations, the pallet is turned over so that it can touch the wheel with its back side. In the following discussion, we focus on the region of normal operation.

Since the escape wheel is periodic, so is its configuration space, and ultimately its place vocabulary. Figure 11 illustrates how the places are defined based on the configuration space. The configuration space region shown on the bottom is part of the one indicated in Fig. 10. The part of the place vocabulary which corresponds to this region is indicated above it. The black dots indicate zero-dimensional places, defined in configuration space as constraint intersection points or points where a constraint's qualitative direction changes. The boxes indicate one-dimensional places, defined as constraint segments. Some of the places are labeled for reference: CSEGs refer to constraint segments and IPs stand for intersection points. The qualitative direction of the constraint segments is also indicated as a pair of signs $([\delta\phi], [\delta\psi])$, indicating the directions of motion of the wheel and the pallet.

For reasons of simplicity, Fig. 11 does not show the places corresponding to the two-dimensional free-space regions. All constraint segments and intersection points are adjacent to the region of free space between them. In the representation computed by the program, this free space region is broken up into *quasiconvex* regions, which are described later in the paper. These divisions are indicated by the dashed lines in the configuration space. They sometimes cause a constraint segment to be broken into several pieces, resulting in the ambiguous numbers for some of the labeled places.

Figure 11 shows sample configurations corresponding to some of the places. Starting from the left, the first configuration shown corresponds to the intersection point IP-729. The configuration is characterized by two points of contact between the pallet and the wheel. Note that the simultaneous contact holds only in this particular configuration and cannot be maintained when the objects move. If we turn the wheel counterclockwise (in the negative direction), we reach CSEG-680/681, the next configuration shown, which exhibits only one of the contact points of IP-729. The qualitative direction of CSEG-680 is $(+, -)$, a fuzzy direction which gives the following two inference rules:

$$[\delta\phi] = - \quad \Rightarrow \quad [\delta\psi] = + ,$$
$$[\delta\psi] = - \quad \Rightarrow \quad [\delta\phi] = + .$$

The first rule expresses the fact that pushing the wheel clockwise pushes the pallet counterclockwise (making it rise). The second rule states the inverse pushing relation: turning the pallet clockwise will push the wheel in the counterclockwise direction. These are the two causal inferences that the contact allows.

The next configuration shows the contact exhibited in CSEG-689, which is the same as that in CSEG-679. The two segments differ in their qualitative direction: in CSEG-689, it is $(-, +)$, while in CSEG-679, it is $(-, -)$. Thus, in CSEG-689 we have the rule

$$[\delta\phi] = + \quad \Rightarrow \quad [\delta\psi] = - ,$$

while in CSEG-679, this rule becomes

$$[\delta\phi] = - \quad \Rightarrow \quad [\delta\psi] = - .$$

In order to push the pallet upwards in a clockwise direction, the wheel has to be turned either counterclockwise in CSEG-689 or clockwise in CSEG-679. Note that in both cases, the pallet can only be pushed upward, never downward. The configuration shown is very close to the boundary between the two segments: it is difficult to tell from the figure in which direction the wheel has to be turned to push up the pallet.

Continuing from CSEG-679, we reach CSEG-677/678, where the wheel now touches the inside of the pallet, and finally CSEG-686/687. Movement from one state to the other is possible in two ways: either the objects stay in contact

Table 2
The possible transitions from IP-729 for motion in the different qualitative directions

Direction	Possible transitions
$(-, +)$	{CSEG-680/681, free space}
$(0, +)$	{free space}
$(+, +)$	{CSEG-682, free space}

and pass through a configuration of double contact, or the pallet lifts off the wheel and touches it again in the other configuration. The possibility of representing the latter type of motion, called an intermittent motion, is a strong advantage of the place vocabulary theory over heuristic approaches, which often have to rely on a continuous contact between the objects.

For a given place, each possible direction of motion can be associated with sets of adjacent places to which a transition may occur. For example, in the case of IP-729, we have the table of possible transitions shown in Table 2. The directions not shown in Table 2 have no adjacent places; motion in any of these directions is not possible at all. For CSEG-680/681, the possible transitions are shown in Table 3.

Note that if a motion in direction (a, b) can lead from place X to place Y, then motion in the $(-a, -b)$ direction can lead from Y to X. The possible transitions from the two-dimensional free-space region are determined from all the adjacent intersection points and constraint segments by inverting the directions of transition. Because there is a large number of possible transitions into the two-dimensional region from other elements of the place vocabulary, the number of possible transitions *from* this region is very large, and its representation as a single place causes enormous ambiguities. Furthermore, it allows many incorrect sequences of transitions. For example, by moving from CSEG-682 in the $(-, +)$ direction (turning the wheel clockwise, and the pallet counterclockwise), we could end up in CSEG-686 by just moving in the configuration space region shown in Fig. 11. But CSEG-686 has a larger ϕ coordinate than CSEG-682, and lies in the $(+, +)$ direction!

The solution to this problem is to define landmark values in both parameters

Table 3
The possible transitions from CSEG-680/681; directions not shown are impossible motions

Direction	Possible transitions
$(-, +)$	{IP-747, free space}
$(0, +)$	{free space}
$(+, +)$	{free space}
$(+, 0)$	{free space}
$(+, -)$	{IP-729, free space}

and keep track of the relative position with respect to them, as described earlier. Several landmark values in the ϕ direction exist between CSEG-682 and CSEG-686, and the incorrect sequence of transitions can be ruled out because it could not cross them in the right order.

As a last point, the representation of the place vocabulary takes into account the periodicity of the structure by representing only a prototype of each different place. The periodic copies are represented by tokens which refer to the explicit representation of the original. Adjacencies to places which are represented as periodic copies are indicated by *periodic links*, which refer to the prototypical copy and an offset. The envisionment then does not have to differentiate the different teeth of the wheel, and only represents the behavior for the prototypical tooth.

3. The Two Different Aspects of Place Vocabularies

For the purpose of high-level reasoning, a place in the place vocabulary is just a set of rules for inferring directions of motion and place transitions. For many applications, it is also important to relate this symbolic description to actual observations and actions performed on the device. For example, in trouble-shooting or machine learning, it is necessary to compare the observed behavior with the symbolic model, and to propose concrete tests or experiments to gather additional information. In some cases, it may even be necessary to verify the model itself [26]. Such tasks require the capability to classify observations in terms of the place vocabulary and to generate sample configurations for places, which is not provided by the symbolic representation alone.

The required perceptual aspect of the place vocabulary can be provided by attaching to each place a precise algebraic representation of the configuration space region that it corresponds to. As it turns out, the region decomposition of free space which is most practical for reasoning is very different from that which is required for a useful algebraic description. In this section, we discuss the requirements for the two different aspects of the place vocabulary and how they are satisfied in the actual representation.

3.1. The place vocabulary as a substrate for reasoning

As a substrate for reasoning about mechanism kinematics, the place vocabulary is a purely symbolic representation. Each place is a structure which contains the following information:

- a mapping from qualitative motions to possible place transitions;
- for one-dimensional places, the qualitative direction (as described earlier, this defines the inference rules for the transmission of force and motion);
- the locations of the extremal points in each of the configuration space parameters, defined by quantity spaces.

The locations of the extremal points are required for place composition, described later in the paper. They are also needed for the straight-line subdivision of free space regions that is used to reduce transition ambiguities. For both purposes, it is sufficient to represent the precise values as points in a quantity space [15], which must be circular in the case of a rotational parameter.

The places are arranged in the place graph, which is a purely symbolic representation of the kinematics and contains no numerical or algebraic elements. Furthermore, reasoning with it requires no reference to the underlying configuration space, which is only used as an intermediate formalization.

3.2. The place vocabulary as a perceptual representation

In order to link the place vocabulary with precise, nonqualitative configurations of the mechanism, an algebraic description of the configuration space region corresponding to each place is attached to its symbolic representation. In order to be useful for classifying an observed configuration, it is necessary that these regions satisfy some form of convexity criterion. As we are dealing with algebraic curves, not straight lines, we use an analogous *quasiconvexity* criterion. A quasiconvex region is defined by a set of functions $\{\mathscr{F}_i\}$ and a set of signs $\{\mathscr{S}_i\}$, where each $\mathscr{S}_i$ is either $+$, 0, or $-$. The set of points $\{x\}$ belonging to the region is defined by the condition that for each of the $\mathscr{F}_i$, $\mathrm{sign}[\mathscr{F}_i(x)] = \mathscr{S}_i$. For typical mechanisms, the free space regions in general do not satisfy this criterion, and further subdivisions have to be imposed, as illustrated in Fig. 12. Because of the algebraic characteristics of configuration space, it turns out to be very difficult to impose suitable subdivisions in the configuration space domain. Instead, we delimit the region of applicability of each constraint by *applicability constraints* [10], curves in configuration space which can be computed in a principled manner. The applicability constraints form a well-defined subdivision which guarantees the quasiconvexity property.

It would be desirable to re-use the applicability constraints as a division of free space which reduces transition ambiguities. Unfortunately, the fact that they are not parallel to one of the coordinate axes makes this impossible.

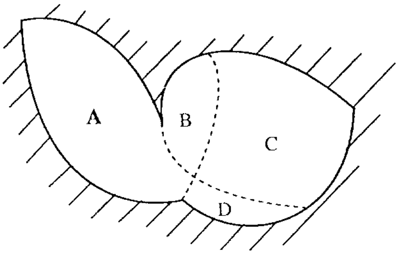

Fig. 12. A cell decomposed into regions A–D by applicability constraints.

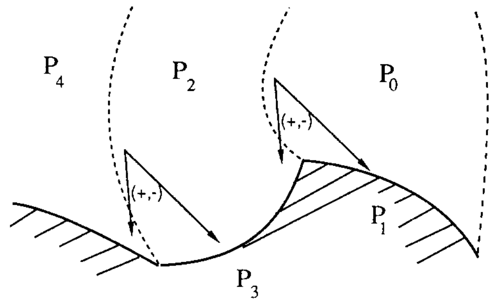

Fig. 13. Curved divisions of free space do not reduce spurious transitions.

Consider the situation shown in Fig. 13. A transition from P_0 in the $(+, -)$ direction may lead either to P_1 or P_2. From P_2, the same motion may lead to P_3 or P_4. It is thus possible to move from P_0 to P_4 by continuous motion directed to the right in Fig. 13. But P_4 lies on to the left of P_0, and should not be reachable by such a motion! Clearly, the curved divisions to not help very much with reducing transition ambiguities, and could be ignored in high-level reasoning. Other decompositions, such as the one defined by landmark values discussed earlier, have to be substituted or superimposed depending on the application of the place vocabulary.

The exact representation of the configuration space regions explicitly defines the semantics of the symbolic description in terms of actual, numerically specified configurations of the objects. This component distinguishes the place vocabulary from other forms of knowledge representation, where the link between symbolic and perceptual levels is never formally represented.

4. Analysis of Complete Mechanisms

In the introduction, we defined a mechanism as an implementation of a *kinematic chain*, following the classic work of Reuleaux [27, 28]. The kinematic chain is formed by rigid connections between the objects involved in successive higher pairs. The first step in analyzing a kinematic chain is to identify the kinematic pairs that it consists of, and compute place vocabularies for them. In this section, we describe how to analyze the complete kinematic chain based on the place vocabularies for the kinematic pairs.

4.1. Analysis of kinematic chains

Successive kinematic pairs are related by a shared object or a pair of objects which are rigidly connected. Their configuration spaces share the common position parameter of this object or pair of objects. These common parameters form the basis for *composing* the place vocabularies of the successive kinematic pairs, and analyze the complete chain. There are two different ways to approach this composition. We can construct a combined place vocabulary for the complete kinematic chain and base the envisionment on it. Alternatively, we can compute envisionments for each kinematic pair and specify the dependencies separately.

The complexity analysis of the number of places in the place vocabulary for a kinematic chain [12] reveals that Reuleaux's intuitions in dividing the analysis into kinematic pairs was a good one. It shows that, in general, the number of places grows exponentially with the number of composed interactions. This makes the computation of a single place vocabulary for an entire kinematic chain impractical. In general, we can assume that the size of the envisionment grows in proportion to the size of the place vocabulary and thus conclude that a single envisionment of a complete kinematic chain is impractical.

As the function of a kinematic chain is to transmit motion, the result of the analysis should specify how this motion is translated. The analysis thus consists of the mapping of an input history to an output history. The input history is a sequence of state transitions in the envisionment for the input pair, and the output history is the same for the output pair. For example, in the QRG clock, the motion of the pendulum results in a certain history in the envisionment for the escapement. This history is determined by analysis of the dynamic parameters based on the place vocabulary. The motion is propagated through the kinematic chain to the gearwheel driving the hour hand of the clock. Since the hand is connected rigidly to the final gearwheel, its motion is given by a history in the envisionment of the last kinematic pair.

We now outline how a history can be propagated through a kinematic chain. Consider two successive kinematic pairs P_1 and P_2 in the kinematic chain, and two successive states A and B in the history for P_1. By place composition (Section 42), we find the set of states in the envisionment for P_2 that P_2 could be in given that P_1 is in state A. We call this set $\mathcal{A}$ and let $\mathcal{B}$ be the similar set corresponding to B. Among the members of $\mathcal{A}$ and $\mathcal{B}$, there will be only few pairs such that a transition between them is allowed in the place vocabulary (or envisionment). Only these pairs are possible in the history for P_2. We mark these pairs and the transitions between them as possible and carry out this process for all transitions in the input place vocabulary (or envisionment) for P_1. The possible histories in P_2 are given as the set of possible complete chains of marked pairs. To reduce the number of ambiguities to an acceptable level, the procedure should be combined with the dynamic analysis according to Newtonian mechanics. This is the subject of current research by Paul Nielsen [25].

To carry out this propagation, the following information is required:

- a set of place vocabularies (or envisionments) for the kinematic pairs, and
- for each kinematic pair, a mapping function from histories in its place vocabulary to simultaneously possible histories in the place vocabularies of adjacent pairs.

While histories can be defined either in envisionments or in place vocabularies, the mapping required for the propagation is based on the place vocabulary only. The next section describes how to find this mapping.

4.2. Place composition

We consider the problem of computing the place vocabulary for a kinematic chain made up of two pairs. This place vocabulary is defined in the three-dimensional configuration space spanned by the motion parameters of the three parts. The place vocabularies for the kinematic pairs are given in two-dimensional subspaces. Their representation in the full space is formed by their projection through the missing parameter, as illustrated in Fig. 14. The set of places in the composition are all the intersection volumes formed by these projections. For zero-dimensional places, the backprojected places intersect whenever they agree exactly in the common configuration space parameter. Figure 14 illustrates the intersection of the projections of one-dimensional places. An intersection between the projections exists whenever there is an overlap in the intervals of the common configuration space parameter covered by the curves. Determining the set of intersections between one-dimensional places is therefore a problem of interval intersection, and can be handled by the quantity space technique of qualitative process theory [15].

The intersection volume of two backprojected two-dimensional places is defined by the intersections of the backprojections of the bounding one-dimensional places.[3] An intersection between the places exists whenever at least one pair of one-dimensional places or free-space divisions in their boundaries intersect each other. Rather than carry out the interval intersection test on each pair of curves separately, we can make use of transitivity and directly test the intervals covered by the places themselves.

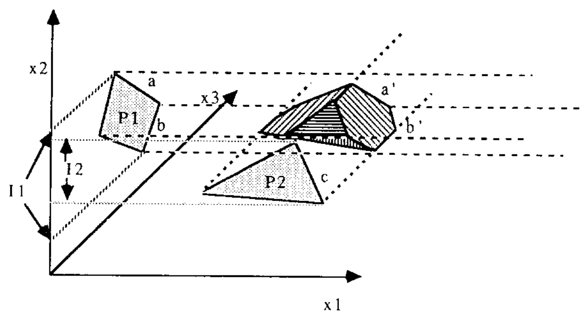

Fig. 14. Places P_1 and P_2, with common configuration space parameter x_2, compose into the intersection volume shown. The one-dimensional places a and b intersect place c to form places a' and b' in the intersection. Note that an intersection between P_1 and P_2 exists if and only if the intervals I_1 and I_2 covered in the common parameter intersect.

[3] Note that the backprojected places cannot contain one another.

The mapping between the successive place vocabularies can be found by carrying out the interval intersection tests for each combination of places in the two place vocabularies. For each place in the place vocabulary in the first pair, the set of places that the second pair can be in simultaneously is the set of places where the interval test succeeds. The composition of place vocabularies for the analysis of kinematic chains thus requires only interval intersection tests between the extremal points of the places, and can be carried out based on purely symbolic representations.

4.3. Complexity considerations

Many qualitative physics systems suffer from excessive numbers of states and intractably high branching factors in the envisionment. The place vocabularies for kinematic pairs do not suffer from these problems. The number of places corresponds to the number of different contact configurations, a manageably small number. The envisionment is built on the place vocabulary by adding the dynamic parameters; in the case of kinematic pairs, these are the directions of motion of the two objects. Given particular directions of motion, the number of possible transitions to different places is limited by the adjacencies. For zero- and one-dimensional places, it is limited to at most five. For two-dimensional places, it can be limited to at most three by proper introduction of subdivisions. The possible transitions due to changes in the dynamic parameters may cause additional ambiguities. However, as in past examples of qualitative analyses of Newtonian mechanics problems [6, 9, 13, 14], practical experience with place vocabularies has not revealed a problem with excessive ambiguities. In summary, for kinematic pairs the ambiguities are in general limited to a manageable amount.

An envisionment of complete kinematic chains, however, leads to a combinatorial explosion of states. This is because the common envisionment of the complete mechanism has to include all combinations of states in each of the kinematic pairs. However, such a complete envisionment ignores the natural decomposition of the device. A better idea for representing the behavior of the complete mechanism is as a series of related envisionments for each kinematic pair. Another possibility is to use summarized descriptions, as explored by Paul Nielsen in the analysis of the QRG clock [25].

5. Implementation and Further Examples

The theory described in this paper has been tested using an implementation which computes place vocabularies from geometric descriptions of kinematic pairs. The computation proceeds in the following steps:

(1) Process the geometric descriptions to instantiate the constraints and their endpoints.
(2) Find the intersections between constraints to determine their envelopes.

(3) Find the connected regions of free space.
(4) Compute the quasiconvex decompositions and construct the output representation.

For details of the algorithms used, the reader is referred to [12].

In the rest of this section, we present two additional examples which we have analyzed using the implementation: gears as the most common kinematic pair in mechanisms, and a cylinder escapement as an example of a device difficult to analyze using heuristic methods. The program has also analyzed a large number of other examples, including all the kinematic pairs in the QRG clock shown in Fig. 1, ratchets, cams, and different types of escapements. Some of these additional examples are discussed in [12]. On all the examples, the place vocabulary has been a sufficiently precise representation for a correct qualitative analysis of the device's kinematics.

5.1. Gears

Consider the pair of gears shown in Fig. 15. This gear pair occurs in the mechanism driving the hands of the QRG clock mentioned earlier. The configuration space for this example is again two-dimensional, the two dimensions being made up of the orientation angle ϕ for the lower gear and the angle ψ for the upper gear. The lower gear has 12 teeth, while the upper one has 36 teeth. As the greatest common divisor of these two numbers is 12, the configuration space, shown in Fig. 16, has 12 identical doubly connected "channels," corresponding to different initial choices for pairs of interacting teeth. Each of these channels results in a disjoint component of the place graph.

As the interactions between the teeth are all periodic, so is the place vocabulary. In this case, we have $12 \cdot 36 = 432$ possible choices for a pair of teeth, so there exist 432 periodic copies of each place. The places corresponding to the same channel are arranged in cycles in the place graph. Traversing a cycle corresponds to the upper gear performing one rotation, and the lower gear performing three rotations. The number of rotations can be tracked by observing the transitions of landmark values. This allows inferring the fun-

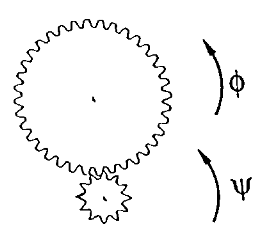

Fig. 15. A pair of gears.

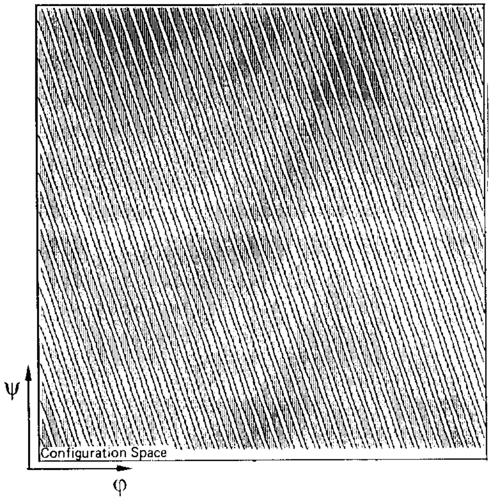

Fig. 16. The configuration space for the gear pair.

damental property of the gears, their ratio of rotation speed, from the place vocabulary. At the same time, the individual places (of which only the original versions are represented) allow a causal explanation of the interaction between the teeth which achieves the gearwheel function.

5.2. Cylinder escapement

An example which particularly illustrates the power of the approach is the analysis of a cylinder escapement, shown in Fig. 17. It consists of two parts: a cut-out cylinder, on the left, which rotates around its center, and a spoked wheel. Note that of the first part, only the hollow cylinder is actually in the same plane as the wheel. In the plane of our analysis, its center of rotation thus lies outside the object, although in some other plane it must be fixed by a joint. The wheel is attached to a spring that drives it in the counterclockwise direction. The cylinder is attached to a balance spring and rotates back and forth, and lets the wheel slip forward by one tooth in each oscillation. This

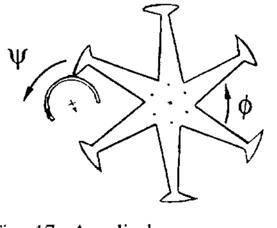

Fig. 17. A cylinder escapement.

example is very difficult to analyze using heuristic techniques, because its functioning depends crucially on the intermittent motion between the contacts. However, using the place vocabulary theory its behavior can be represented without difficulty.

The examples shown so far all had both of the objects rotating. The implementation also covers cases where one object or both objects are translating along a fixed axis, such as cams. However, from a qualitative point of view, a cam does not exhibit a very interesting behavior: its place vocabulary is simply a loop of places where both objects move alternatingly in the same and opposite directions.

6. Conclusions and Future Work

In this paper, we have described a theory of qualitative kinematics based on the concept of a place vocabulary, a compact representation of the kinematic interactions between the mechanism's parts. The place vocabulary provides the spatial substrate upon which envisionments and causal explanations of the device's behavior can be based. We have seen how place vocabularies can be computed from an exact description of the shapes of the objects. This computation has been implemented, and we have shown examples of the results. In all the examples we have analyzed, the place vocabulary has provided a sufficiently powerful representation for a correct qualitative analysis.

In contrast to earlier research in reasoning about shapes, the place vocabulary approach has the advantages that it is generative and complete. We consider the current theory the high ground from which we can proceed towards representations that may be easier to compute by weakening the assumptions of exact algebraic computations and accurate input data. As a model of human reasoning, the current theory models human competence, while weaker theories may provide better models of actual performance.

The place vocabulary representation can be applied to general kinematic problems involving three-dimensional motion with arbitrary degrees of freedom. However, we have not yet investigated how such place vocabularies might be computed. Algebraic cell decomposition techniques based on decision methods [1, 21, 29, 31] could be used to provide algorithms for these cases, but they are rather opaque and inefficient, and implausible as a model of human abilities. A promising research strategy for generalizing the current algorithms to more general cases is to decompose the problem into simpler subproblems, such as two-dimensional projections.

ACKNOWLEDGEMENT

I would like to thank Ken Forbus for much helpful advice and comment, Tom Galloway and Paul Nielsen for commenting on drafts of this paper, IBM for partial support through an IBM graduate fellowship, and ONR for partial support through contract no, N00014-85-K-0225.

REFERENCES

1. Ben-Or, D. Kozen and J. Reif, The complexity of elementary algebra and geometry, *J. Comput. Syst. Sci.* **32** (1986) 251–264.
2. R.A. Brooks and T. Lozano-Pérez, A subdivision algorithm in configuration space for findpath with rotation, in: *Proceedings IJCAI-83*, Karlsruhe, FRG (1983) 799–806.
3. R.A. Brooks, Symbolic reasoning among 3-D models and 2-D images, *Artificial Intelligence* **17** (1981) 285–348.
4. E. Davis, The MERCATOR representation of spatial knowledge, in: *Proceedings AAAI-83*, Washington, DC (1983).
5. J. Denavit and R.S. Hartenberg, A kinematic notation for lower-pair mechanisms based on matrices, *J. Appl. Mech.* **22** (1955).
6. J. de Kleer, Qualitative and quantitative knowledge in classical mechanics, MIT AI Lab, Tech. Rept. 352, Cambridge, MA (1975).
7. J. de Kleer, Causal and teleological reasoning in circuit recognition, MIT AI Lab, Tech. Rept. 529, Cambridge, MA (1979).
8. J. de Kleer and J.S. Brown, A qualitative physics based on confluences, *Artificial Intelligence* **24** (1984) 7–83.
9. J. de Kleer and D.G. Bobrow, Qualitative reasoning with higher-order derivatives, in: *Proceedings AAAI-84*, Austin, TX (1984).
10. B.R. Donald, Motion planning with six degrees of freedom, MIT AI Lab, Tech. Rept. 791, Cambridge, MA (1984).
11. M.A. Erdmann, On motion planning with uncertainty, MIT AI Lab, Tech. Rept. 810, Cambridge, MA (1984).
12. B. Faltings, The place vocabulary theory of qualitative kinematics in mechanisms, Tech. Rept. UIUCDCS-R-87-1360, University of Illinois, Urbana, IL (1987).
13. K.D. Forbus, Spatial and qualitative aspects of reasoning about motion, in: *Proceedings AAAI-80*, Stanford, CA (1980).
14. K.D. Forbus, A study of qualitative and geometric knowledge in reasoning about motion, MIT AI TR 615, Cambridge, MA (1981).
15. K.D. Forbus, Qualitative process theory, *Artificial Intelligence* **24** (1984) 85–168.
16. F. Gardin and B. Meltzer, Analogical representations of naive physics, *Artificial Intelligence* **38** (1989) 139–159.
17. A. Gelsey, Automated reasoning about machine geometry and kinematics, in: *Proceedings Third IEEE Conference on AI Applications*, Orlando, FL (1987).
18. P. Hayes, The naive physics manifesto, in: D. Michie, ed., *Expert Systems in the Micro-Electronic Age* (Edinburgh University Press, Edinburgh, 1979).
19. L. Joskowicz, A framework for the kinematic analysis of mechanical devices, Tech. Rept. 313, New York University, New York (1987).
20. B. Kuipers, Commonsense reasoning about causality: Deriving behavior from structure, *Artifical Intelligence* **24** (1984) 169–203.
21. D. Kozen and C.K. Yap, Algebraic cell decomposition in NC, Extended abstract, *Proceedings 26th FOCS* (1985).
22. T. Lozano-Pérez, M.T. Mason and R.H. Taylor, Automatic synthesis of fine-motion strategies for robots, *Int. J. Rob. Res.* **3** (1984).
23. T. Lozano-Pérez and M.A. Wesley, An algorithm for planning collision-free paths among polyhedral obstacles, *Commun. ACM* **22** (1979) 560–570.
24. M.T. Mason, Compliance and force control for computer controlled manipulators, *IEEE Trans. Syst. Man. Cybern.* **11** (1981) 418–432.
25. P. Nielsen, A qualitative approach to rigid body mechanics, Ph.D. Dissertation, University of Illinois, Urbana, IL (1988).

26. P. Dague, O. Raiman and P. Deves, Troubleshooting: When modelling is the trouble, in: *Proceedings AAAI-87*, Seattle, WA (1987).
27. F. Reuleaux, *Theoretische Kinematik* (Vieweg, Braunschweig, 1875).
28. F. Reuleaux, *The Kinematics of Machinery* (Macmillan, London, 1876).
29. J.T. Schwartz and M. Sharir, On the piano movers' problem, II: General techniques for computing topological properties of real algebraic manifolds, *Adv. Appl. Math.* **4** (1983) 298–351.
30. C. Stanfill, The decomposition of a large domain: Reasoning about machines, in: *Proceedings AAAI-83*, Washington, DC (1983).
31. A. Tarski, *A Decision Method for Elementary Algebra and Geometry* (University of California Press, Berkeley, CA, 1948).

Received January 1988; revised version received June 1989