Motion Planning and the Design of Orienting Devices for Vibratory Part Feeders

Tomás Lozano-Pérez
Artificial Intelligence Laboratory
Massachusetts Institute of Technology
January, 1986

Abstract: We propose an approach to the design of orienting devices in mechanical part feeders. The basic idea is that the function of orienting devices can be described by the constraints they impose on the motion of parts being fed. These motion constraints are expressed as a diagram in the part's configuration space. Combinations of these diagrams serve both in describing a device's function and in designing devices with specified behavior. The design problem for orienting devices in part feeders is a kind of inverse of the motion planning problem in robotics. In both cases we know the shape of the moving part. In motion planning, we are given the obstacles and we must find a legal path. In our view of design, we are given the desired motion and are asked to find the shape of the obstacle(s), that is, the shape of the orienting device.

1. The general problem

The general goal of this paper is to explore the relationship between motion planning and design, in particular, to illustrate how devices may be specified by the way they constrain the motions of other objects [Lozano-Pérez 86]. We explore these concepts within the context of orienting devices for part feeders. In particular, we limit our attention to the *geometric* considerations in the design of orienting devices. The design of practical orienting devices depends heaviliy on exploiting the physics of feeders, but we ignore most of these considerations. Our specific goal, then, is developing a systematic way of generating candidate designs on purely geometric grounds.

The geometric design problem for orienting devices is a kind of inverse of the motion planning problem in robotics. In both cases we know the shape of the moving part. In motion planning, we are given the obstacles and we must find a legal path between the specified origin and destination. In our view of design, however, we are given the desired motion (actually a range of possible motions) and are asked to find a legal shape of the obstacle, that is, the orienting device.

One reason for our choice of orienting devices as a domain is that they provide a relatively simple, well-structured domain for investigating the relationship between shape and function. Another reason is the practical importance of these devices. In the absence of truly general robots, much of the work in automating assembly tasks goes toward the design of part feeders, assembly fixtures, and tools. A systematic approach to the design of these devices would be valuable [Asada and By 85]. Yet another motivating factor is that feeders and fixtures exploit intrinsic geometric and mechanical properties of object interactions; these same properties can be exploited during robot manipulation [Mason 82, 85, 86, Erdmann and Mason 86]. Last, but not least, a significant body of literature exists for feeders, notably [Boothroyd, Poli, and Murch 77], which includes hundreds of examples.

2. Vibratory Part Feeders

One of the most widespread problems in automation is that of feeding unoriented parts to automatic machinery that expect parts to be precisely oriented. Clever solutions to this problem abound, but one of the cleverest and most useful solutions is the vibratory bowl feeder, illustrated in Fig. 1. The basic idea is admirably simple. The unoriented parts are placed at the bottom of a vibrating bowl that has a track spiraling along its inside surface. The motion of the bowl causes the parts to climb the track. Near the top of the track there are a number of simple devices, called *orienting devices*, that either knock incorrectly oriented parts back into the bowl or re–orient them. In a sense, feeder bowls are programmable devices and the orienting devices are their programming.

A simple example of the use of passive orienting devices is shown in Fig. 1; the goal is to feed rectangular parts one at a time in orientation *a*. The bowl's vibration eliminates unstable orientations, so the orienting devices need only ensure that the other stable orientations of the part are discarded. This can be done by the combination of two simple, widely–used devices: one is the *wiper blade* and the other is the *narrowed track*. The wiper blade will push parts in orientations *c* and *d* off the track and back into the bowl. The height of the blade's bottom edge is chosen to allow parts in orientation *a* or *b* to pass, but to reject parts in orientations *c* and *d*. Note that the blade also rejects stacked parts in orientations *a* or *b*. The narrowed track then causes parts in orientation *b* to fall back in the bowl. The width of the track must be chosen so that the center of mass of parts in orientation *a* lies on the track but the center of mass for parts in orientation *b* lies off the track. I call both of these types of devices *filters*, because they allow some orientations to pass while rejecting others.

Orienting devices get more interesting as the parts become more complex, for example, see Figs. 2 and 3. The goal in both of these cases is to feed cylindrical parts with non–symmetrical ends in a preferred orientation. Assume that the parts have already been constrained to be in a grooved track, possibly by filters as above. There remains a two–way choice based on which end of the cylinder is first. We could construct a filter to select for one of these orientations but it is more efficient to actively re–orient the part to the desired orientation. Assuming that both part orientations are equally probable, re–orienting would

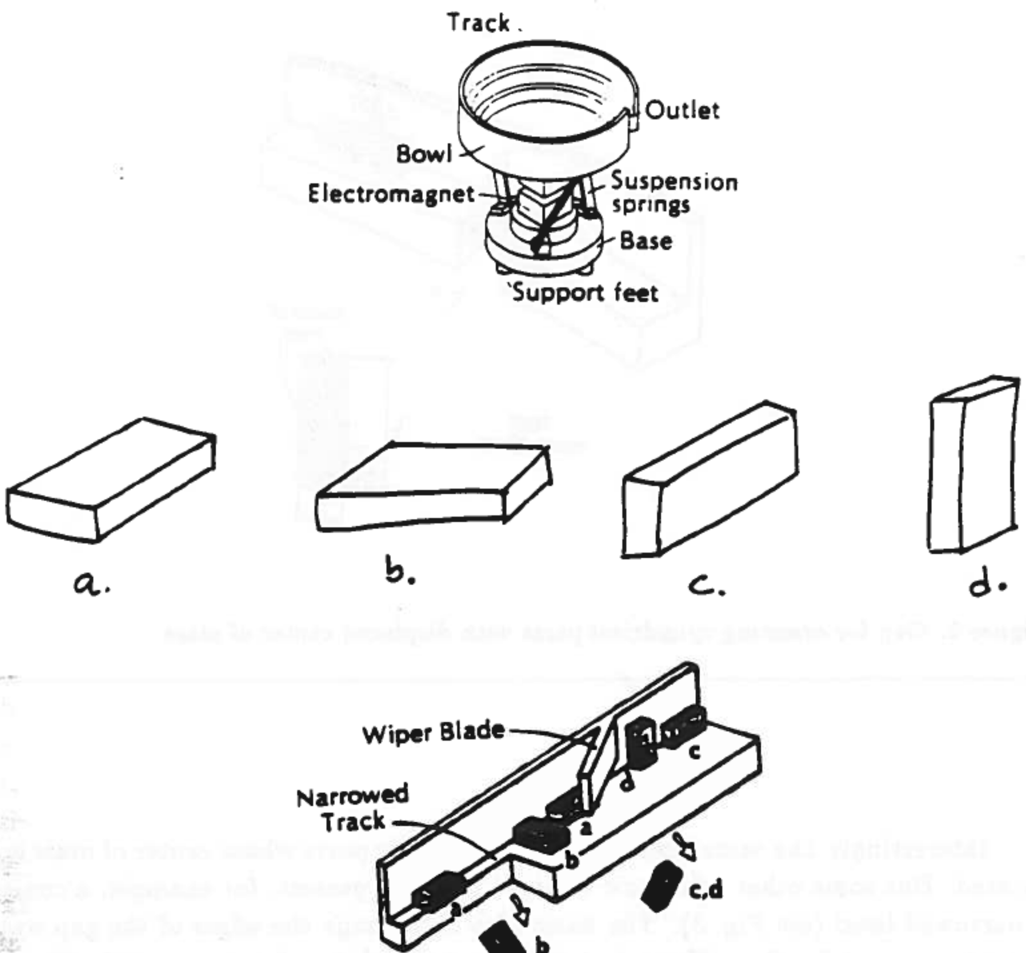

Figure 1. Vibratory bowl feeder with orienting devices to feed rectangular parts in the *a* orientation [Boothroyd, Poli, and Murch 77].

double the rate at which parts are fed, as compared to a filter. The devices shown in the figures re–orient the parts appropriately.

Assume that the cylinders have their center of mass displaced towards one end, possibly because of a hole at the other end. Let the center of mass be a distance l from one end and a distance L from the other, with $l < L$. Consider the effect of a gap in the track. If the heavy end is leading and the gap is wider than l, the part will fall into the gap, heavy end first. If the light end is leading and the gap width is less than L, then the part will advance beyond the gap. When the trailing end of the cylinder goes beyond the first edge of the gap, however, its center of mass will not be supported and the part will fall into the gap, again with its heavy end first. An output track, placed below the gap, collects the oriented parts.

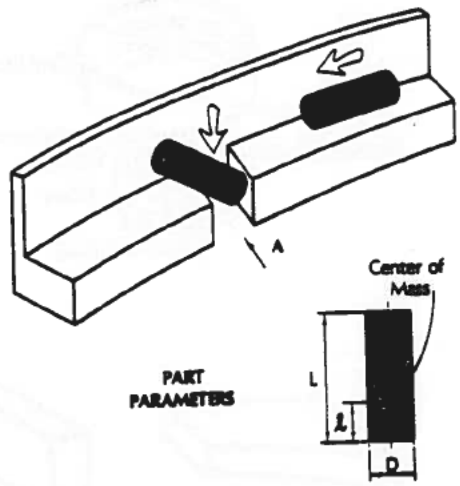

Figure 2. Gap for orienting cylindrical parts with displaced center of mass

Interestingly, the same effect can be achieved for parts whose center of mass is centrally located. But some other geometric features must be present, for example, a conical end or a narrowed head (see Fig. 3). The basic idea is to shape the edges of the gap so that they provide support for the wide end of the part but not for the narrow end. Thus, the part behaves (relative to this ***shaped gap***) as if the narrow end were not there and is reduced to the case of a displaced center of mass. Of course, the shape and dimensions of the gap must be chosen carefully for the specific part (see Section 4.3).

Clearly, the performance of bowl feeders depends critically on the orienting devices. These devices have to be designed specifically for each part since they exploit characteristics of the part shape and mass distribution. Designs are based on experience and on some repertoire of traditional techniques. There is currently no completely systematic approach to these designs. Some pioneering work by Boothroyd and his colleagues at U. of Mass at Amherst [Boothroyd, Poli, and Murch 77] and Redford at U. of Salford (England) [Redford and Boothroyd 67, Redford, Lo, and Killeen 83], however, has gone a long way towards understanding the operation of bowl feeders and their orienting devices. In particular, these researchers have developed group coding techniques that help systematize the classes of parts that can be fed and the types of orienting devices normally used to feed particular part groups. The examples of feeders in this paper are drawn from [Boothroyd, Poli, and Murch 73].

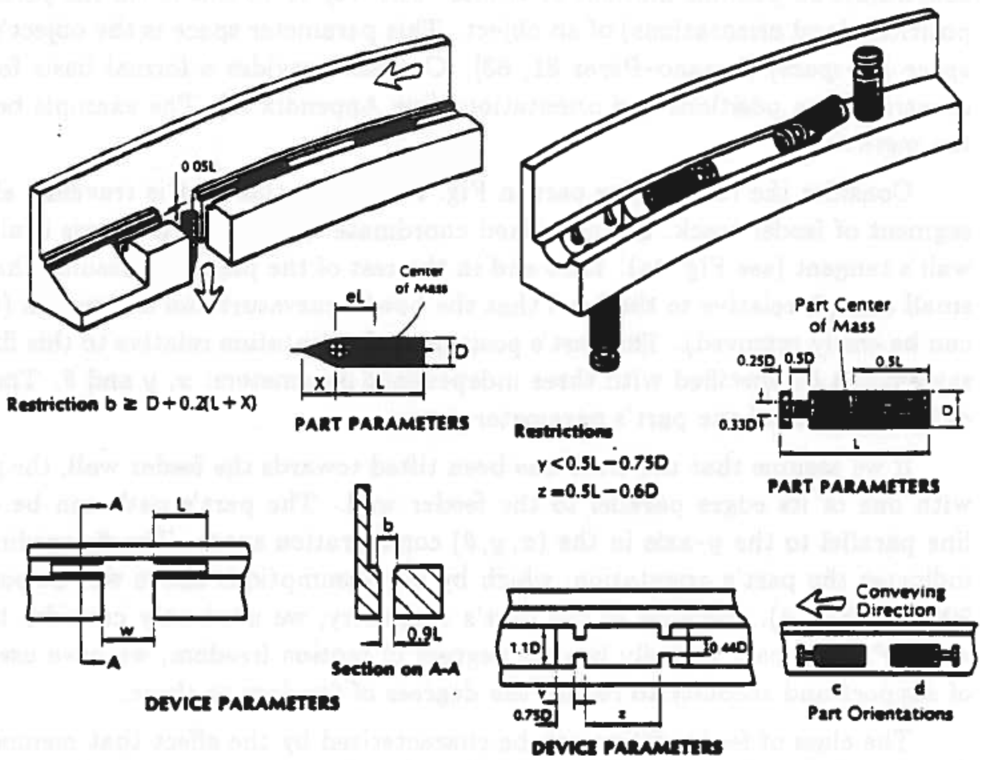

Figure 3. Shaped gaps [Boothroyd, Poli, and Murch 77]

3. Functional Representations of Orienting Devices

Before considering the *design* of orienting devices, we develop a methodology for *describing* the function of orienting devices. We will use this type of descriptions later as specifications of the desired function of filters.

3.1 Characterizing the function of feeder filters

What is the function of feeder filters? One simple definition is the following: *A feeder filter removes from the track those parts traveling in undesired orientations.* This characterization is in terms of the orientations and motions of parts, not their shape. It indicates the desired motion of parts in a specified range of positions and orientations.

This definition suggests a formalization of filter function in terms of positions and orientations rather than shape. Traditional pictorial representations combine the shape and position of bodies inextricably. For our purposes, we need to represent directly the constraints on possible motions of bodies. One way to do this is via the parameter space of positions (and orientations) of an object. This parameter space is the object's *configuration space (C–space)* [Lozano–Pérez 81, 83]. C–space provides a formal basis for dealing with constraints on positions and orientations (see Appendix A). The example below illustrates the method.

Consider the rectangular part in Fig. 1. Assume this part is traveling along a straight segment of feeder track. Define a fixed coordinate system whose y–axis is aligned with the wall's tangent (see Fig. 4a). Here and in the rest of the paper we assume that the object is small enough relative to the bowl that the bowl's curvature can be ignored (this restriction can be easily removed). The part's position and orientation relative to this fixed coordinate system can be specified with three independent parameters: x, y and θ. These parameters define the axes of the part's parameter space.

If we assume that the track has been tilted towards the feeder wall, the part will travel with one of its edges parallel to the feeder wall. The part's path can be described as a line parallel to the y–axis in the (x, y, θ) configuration space. The θ coordinate of the line indicates the part's orientation, which by our assumptions above will be some multiple of 90^o (see Fig. 4). Because of the part's symmetry, we need only consider the lines for 0^o and 90^o. The part actually has six degrees of motion freedom, we have used assumptions of support and stability to reduce the degrees of freedom to three.

The class of feeder filters can be characterized by the effect that members of the class have on the configuration of parts being fed. Let us temporarily limit our attention to motion in the plane of the track. The narrowed track filter we saw in Fig. 1 allows orientation 1 ($\theta = 0^o$) to pass, but not orientation 2 ($\theta = 90^o$). This filter produces the C–space shown in Fig. 4. For convenience, position of the center of mass is chosen to characterize the x, y, z position of the part. The cross–hatched area in the figure represents the positions of the rectangular part in which the track provides support, that is, those positions for which the center of mass of the part is on the track. Therefore, any position of the part outside of the cross–hatched area will cause it to fall back into the bowl.

As the part falls, its z position and its orientation will change, therefore, we cannot represent it's motion in the diagram we have been using. We will, however, indicate the direction of rotation in the diagram; this will prove useful later. Should one wish to, another two-dimensional diagram can be used to show the change in height and orientation as the part falls. Conveniently, we do not have to represent the falling motion in detail; the filter's function is simply to remove the part from the track.

The C–space diagrams for the feeder filters shown in Figs. 4 and 5, in the C–space of the rectangular part, show that these devices satisfy the general characterization of filters. In both cases, parts in orientation 1 ($\theta = 0^o$) are allowed to pass and parts in orientation 2 ($\theta = 90^o$) lose support and fall off the track. The C–space diagrams also reveal the differences between the filters: the possible motions of the part are different. Note that in Fig. 5 the side wall provides some of the part's support; it is possible for the correctly oriented part to fall through the hole if its center of mass loses support. Therefore this

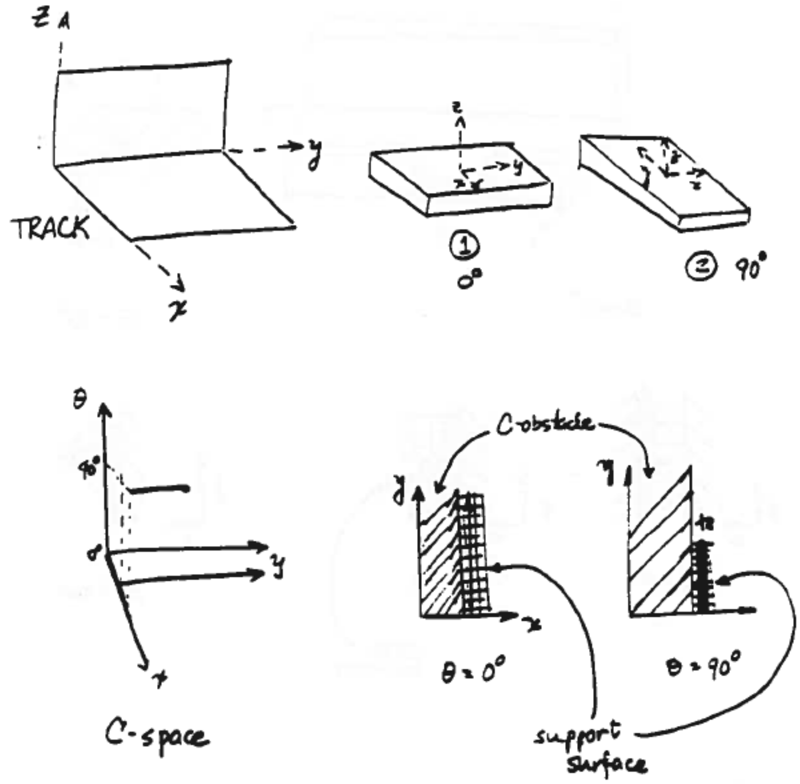

Figure 4. The C–space for a rectangular part on a narrowed feeder track.

particular feeder is only useful for parts whose length is greater than twice their width $l > 2w$.

One crucial characteristic of these C–space diagrams is that many filters produce similar diagrams, that is, they differ only in the positions of edges in the diagrams. The filter shown in Fig. 6 illustrates this point. The C–space diagram for this filter is essentially identical to that of the filter in Fig. 5 even though the shape of the track is different. The essential difference between the two filters is that in Fig. 6 correctly oriented parts fall into a groove cut in the track. Therefore, the side wall does not provide support for both orientations of the part, as it did in Fig. 5. Because of this difference, the filter can be used for a wider range of parts. The filter is applicable when $l > w$, rather than $l > 2w$ for the previous filter. In practice, however, some significant margin must exist between the part's length and width to prevent incorrectly oriented parts from falling into the groove. In the two–dimensional C–space diagrams this difference shows up as a difference in the pairing between surfaces in the diagrams corresponding to different orientations. The right–hand C–space surfaces (C–surfaces) in Fig. 5 both arise because of the same feeder surface. In

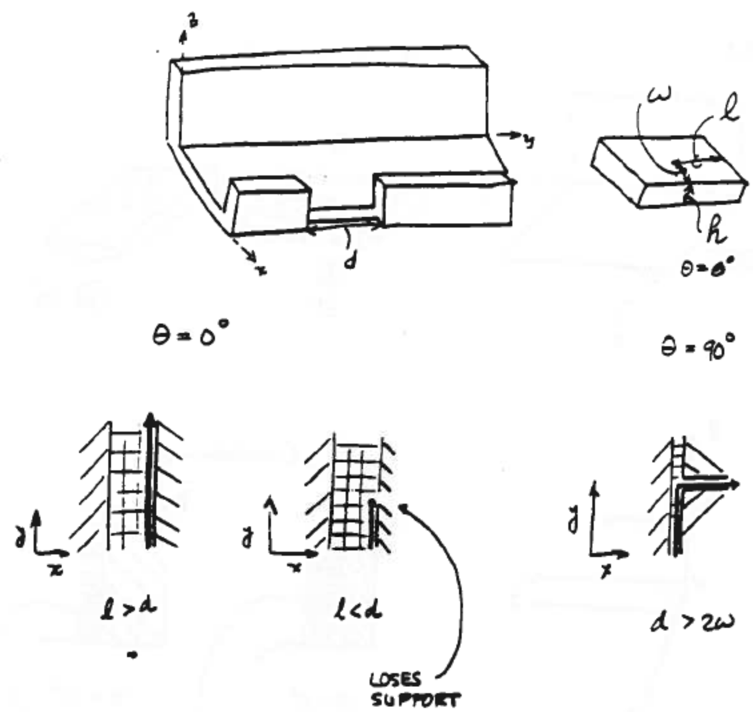

Figure 5. Feeder filter for a rectangular part: (a) Sloped track with a narrow hole in the inner track wall. (b) C–space diagrams for orientations 1 and 2 of the rectangular part.

Fig. 6, the corresponding C–surfaces arise from different feeder surfaces.

We can think of each orientation of the part as traveling in its own channel, a volume of space swept out by the part as it moves along the track. Then, the difference between the feeders in Fig. 5 and 6 is in the relative positions of these channels. In Fig. 5, the channels shared their support surfaces and right–hand walls; in Fig. 6, they are displaced both horizontally and vertically. The choice of placement of these channels relative to each other is quite important and accounts for much of the variation among feeders.

The C–space diagrams also reveal that the same feeder — the feeder in Fig. 5 — will work for some shapes and not others (see Fig. 7). This brings up an interesting question: What is the class of parts for which a particular filter works? The same type of possible motion analysis can be used to answer this question, but we will not attempt an answer here.

In general, a rigid object has six degrees of motion freedom; the C–space for such an object must be six dimensional. In the examples above, we used the support constraints and stability assumption to make the C–space three dimensional, with one dimension having only discrete values. These assumptions make the analysis very simple, but they introduce some limitations. We have already seen one of the limitations: we cannot model the behavior of

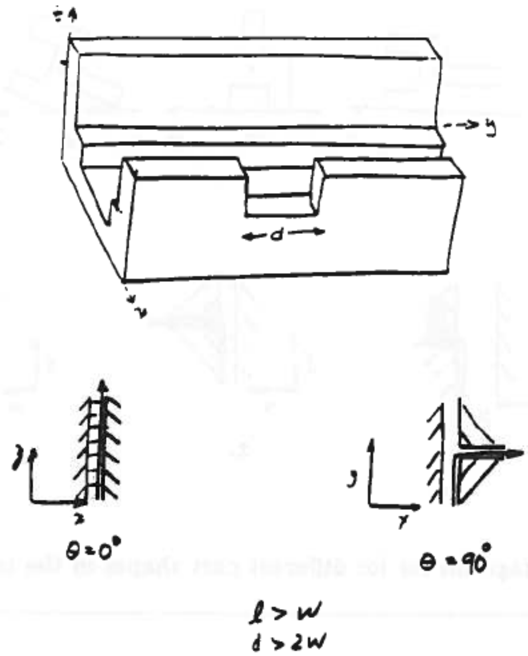

Figure 6. (a) Sloped track with groove and gap in the supporting wall. (b) C-space diagram for the rectangular part.

the part when it loses support. A more serious limitation is our inability to deal with parts that may reach the orienting devices in a whole range of orientations instead of a few stable orientations. In particular, two parts may be stacked on each other. There is nothing inherent in the C-space representation that prevents us from dealing with a continuous range of orientations explicitly [Brooks and Lozano-Pérez 82, Canny 84, Donald 84], but the analysis becomes more complex. For reasons of clarity, we will continue to use these simplifying assumptions throughout the paper.

Another issue that may at first appear to be a problem is representing the behavior of solid objects whose stable orientations span the whole range of rotational degrees of freedom. We can deal with this problem by using multiple (x, y, θ) cross sections of the C-space, each corresponding to different values of the other two rotational degrees of freedom and z (see Fig. 8). In some cases, such as when designing a feeder, the treatment across these different C-spaces must be coordinated. But, this can be done using our knowledge of the parameters that define each C-space cross section. This is essentially equivalent to the way we treated C-space diagrams for different values of θ. In fact, this treatment of different orientation generalizes directly to treating different parts; the techniques discussed in this paper for separating orientations also apply to separating parts.

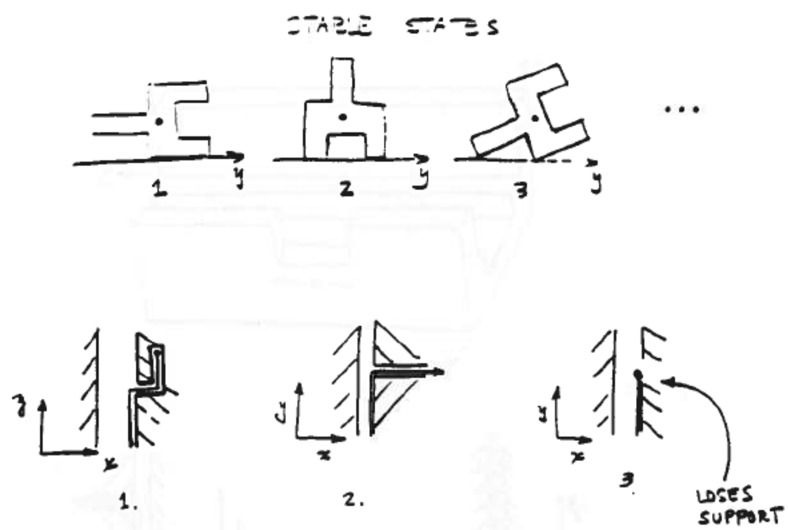

Figure 7. C–space diagrams for for different part shapes in the track from Fig. 5.

It is important to note that the C–space diagrams we have been using can be obtained automatically given geometric descriptions of parts and feeders [Brooks and Lozano–Pérez 82, Lozano–Pérez 81, Canny 84, Donald 84]. So the C–space representation could be used in an automatic system that verified whether a feeder filter was effective for a given part.

3.2 Classifying filters

C–space descriptions can be used to classify feeder filters into sub–classes. Any filter can be characterized by the manner in which it accepts and rejects parts. The patterns of possible motions are determined by the interactions of part surfaces with feeder surfaces. These surface interactions give rise to C–space surfaces that constrain the possible motions of a part in the feeder.

The advantage of C–space representations is that they represent possible motions directly; therefore all the motion constraints on a part are purely local. In the object space, non–local interactions may determine a part's possible motions. The example in Fig. 9 illustrates this general point. The example shows two complex parts that are to be assembled, using only translations. During the assembly, different sets of surfaces constrain the motion of the objects; at different times the constraining surfaces may be widely separated. In the C–space, this assembly can be modeled as moving a point, representing one object's position relative to the other, into a "hole". The local environment of the configuration point contains all the relevant geometric constraints on possible motions.

It is not a coincidence that this assembly operation generates a hole–like pattern in the C–space. It is almost definitional that an assembly operation generates this pattern. An

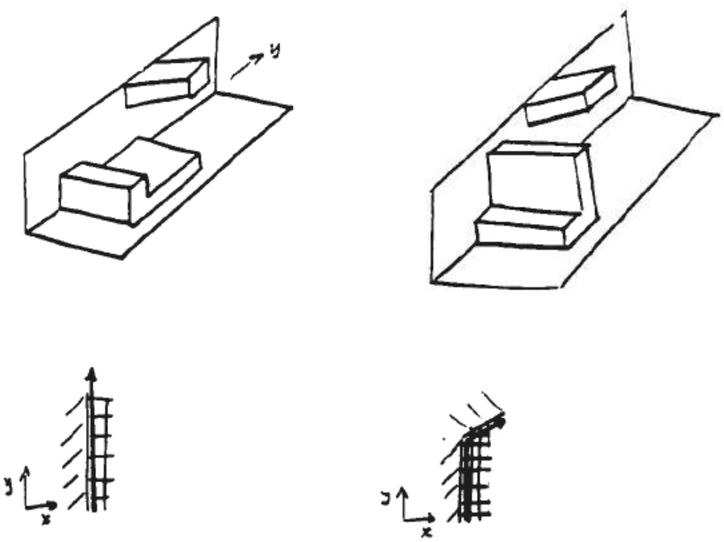

Figure 8. (a) Three-dimensional part rotated out of the plane of the track interacting with a wiper blade. (b) Different three-dimensional cross-sections of the part's six-dimensional C-space show the characteristic filter patterns.

assembly operation typically takes an object from an unconstrained position to a position where its motion (including rotations) is constrained in all directions except the one needed to reach it. This characterization describes a hole in the C-space. One exception to this rule is part alignment for subsequent fastening operations. But, unconstrained alignments are undesirable, since they are not inherently stable. The use of alignment pins and tabs again converts alignment operations into a hole-like pattern in the C-space. In fact, many design-for-assembly rules can be formalized in this fashion.

Filters and other orienting devices can be classified by the local patterns of C-surfaces that they introduce into a part's C-space or, dually, the pattern of motions that are unhampered by C-surfaces. The important point is that only a small number of local patterns are possible. Therefore, we can construct a complete and systematic classification system based on these patterns. The C-space diagrams of Figs. 4, 5, 6, and 8 are instances of a fairly general class of patterns shown in Fig. 10. This is not a complete class of patterns; these only deal with translation in planes parallel to the track and some simple rotations parallel or perpendicular to the feeding direction. More general patterns require considering more general rotations (possibly by means of higher-dimensional C-surfaces). It is interesting that many of the filters listed in [Boothroyd, Poli, and Murch 77] can be characterized by combinations of these simple patterns.

Assuming that parts enter the filter in one of a small number of distinct orientations,

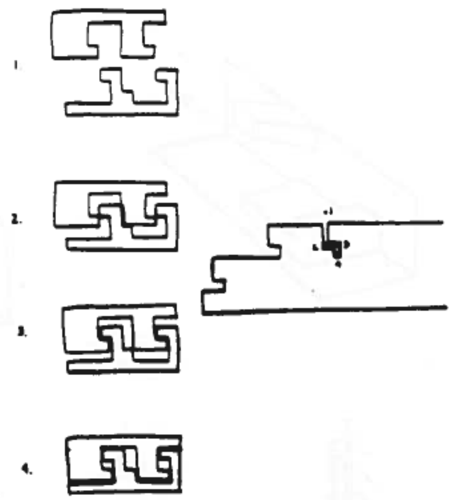

Figure 9. Example that illustrates that all possible relative motions of a part can be determined by examining the local environment in C-space.

the filter's operation can be characterized by associating one of the accept or reject C-space patterns in Fig. 10 to each orientation. For example, the simple filters we considered in the last section deal with two stable orientations; they can be classified as follows:

- Fig. 4: 0° Accept 1, 90° Reject 5A.
- Fig. 5: 0° Accept 2, 90° Reject 4.
- Fig. 6: 0° Accept 2, 90° Reject 4.
- Fig. 8: Accept 1, Reject 1.

There are 14 combinations of one accept and one reject pattern; we have seen filters corresponding to some of these combinations and we will see a few others in the rest of the paper.

In summary, the C-space diagrams represent the range of possible motions of the part in the feeder. The function of orienting devices is to constrain the motions of the parts being fed. The C-space diagrams make these motion constraints explicit and thereby capture the function of the devices. Any specific C-space diagram will contain some details specific to the dimensions of the parts and the orienting devices, but only a few general patterns of possible motions are useful. It is these patterns that determine the classes of orienting devices.

Many types of feeders re-orient parts, that is, they take the configuration of the part from one region of the C-space to another; one simple example is the shaped gap of Section 2. But, shaped gaps only deal with two distinct orientations; they can be characterized in a manner analogous to feeder filters (see Section 4.3). Re-orientation is the subject of Section 4.4.

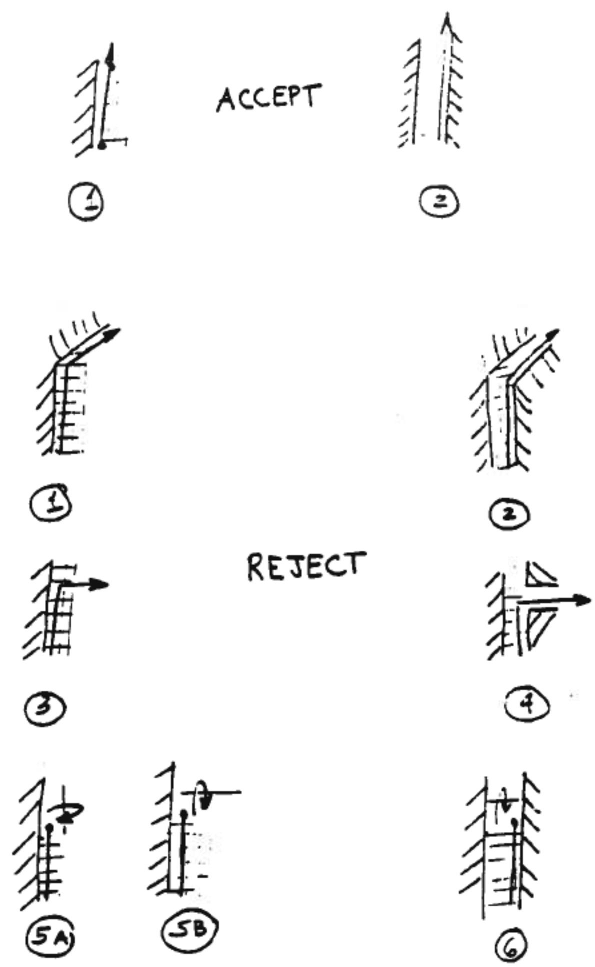

Figure 10. Some basic two-dimensional C-space filter patterns: Accept patterns and reject patterns

4. Design of Orienting Devices

The abstract representation of orienting devices in terms of possible motion constraints in C-space suggests a general approach to the design of orienting devices. But, a more constrained method can be developed using additional knowledge about classes of orienting devices.

4.1 A general approach

The basic idea for a general approach to feeder design is that the C–space constraints in the specification of orienting devices must be enforced by the surfaces of C–space obstacles. For example, imagine that the feeder specification calls for one orientation of a part to be pushed off the track. This means that the orienting device must create an obstacle, in the C–space of the part, with a C–surface that prevents motion in this undesirable orientation but not in any of the desirable orientations. The interaction of these C–space requirements constrains the shape of the orienting device.

In this view, the feeder design problem is a kind of inverse of the motion planning problem. In both cases we know the shape of the moving part. In motion planning, we are given the obstacles and we must find a legal path between the specified origin and destination. In our view of feeder design, however, we are given the desired motion (actually a range of possible motions) and are asked to find a legal shape of the obstacle, that is, the orienting device.

The major difficulty in inferring the obstacle from the path is that the general filter specification does not provide a unique path for the part. The definition of a filter simply specifies that certain orientations be pushed off the track. This type of specification leaves too many options; simple algorithms cannot choose the best one without some additional information.

One way of addressing the problem is by exploiting the classification system we developed in the last section. This classification is a knowledge base of possible filter solutions. The key point is that the *form* of these solutions are expressed independently of the shape of parts; it is a knowledge base about feeder filters. Thus, the size of the knowledge base does not need to grow to deal with more parts; it would only need to grow to represent different classes of possible motions, such as rotations.

4.2 Using possible motions to design feeder filters

To design filters, we need more detailed specifications of the motions of the part when accepted or rejected by the filter. Assuming that only feeder surfaces are used to constrain the motion of parts, detailed part paths give strong constraints on the shape of the filter. In fact, the C–space patterns we used to classify filters provide precisely the information we need to design them; these patterns specify the possible motions of accepted and rejected parts.

Let's take the C–space diagrams of the filter shown in Fig. 5 as a model. Because we are dealing with translating polyhedra, the shape of the C–surfaces in these specifications is relatively insensitive to part shape. The relative position of the C–surfaces for different orientations does depend strongly on part shape. Fig. 11 shows a parameterization of the accept and reject patterns that makes this dependence explicit.

How can we go from C–space patterns and a new part description to designs for orienting devices that have the desired behavior? The answer follows from the definition of C–space obstacles.

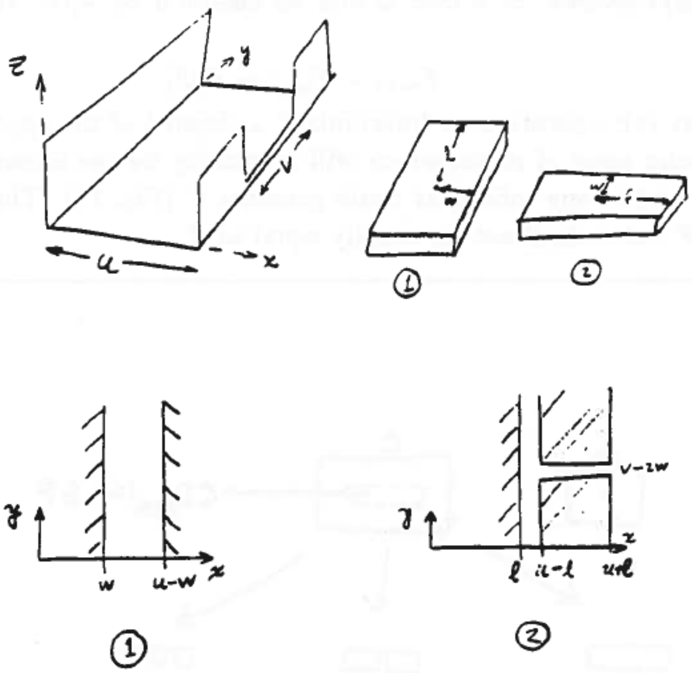

Figure 11. A more general parameterization of the filter in Fig. 5.

Every point inside of each C-space obstacle in one of the patterns represents a position of the part A that overlaps *some* part of the feeder F. This definition cannot be inverted uniquely, that is, given C-space obstacles, C_i and a moving part, A, there isn't a unique solid, F_i, that gives rise to that C-space obstacle (Fig. 12). When considering only translations, the maximal solid, F_{max}, that generates a given C-space obstacle, C, for the given part A is given by

$$F_{max} = \{x \mid (\ominus A)_x \subseteq C\}$$

where $\ominus A = \{-a \mid a \in A\}$ and $(\ominus A)_x$ refers to $\ominus A$ at position x. This is precisely the definition of $CI_{\ominus A}(C)$ [Lozano-Pérez 83]. Essentially, this is the inverse of the process used to compute configuration space obstacles [Lozano-Pérez 83].

Given polygonal objects A and C, F_{max} can be computed by using the fact that each point outside of *all* the C-space obstacles represents a position of the part A that does not contact *any* part of the feeder (except the track surface). Let O be the set of all the points outside of all the C-space obstacles; each of these points represents a position of A. Let F_{total} be an initial solid that is guaranteed to totally contain the feeder. We can construct F_{max} by subtracting the volume occupied by A in each position in O from F_{total}. We can *sweep* the shape of the part over the positions in O to obtain a volume of space where no piece of the actual feeder may intrude. *Sweeping* the part requires that we take the union of the space occupied by the part in all the specified configurations. The resulting volume is called the ***swept volume**** of A over O and we denote it by $A[O]$. Hence

$$F_{max} = F_{total} - A[O]$$

We can visualize this operation as "machining" a channel of the appropriate shape for the part out of a solid piece of metal which will eventually be the feeder track. F_{max} is only guaranteed to *contain* any solid that could generate C (Fig. 12). Thus, F_{max} contains the desired feeder F but it need not be exactly equal to F.

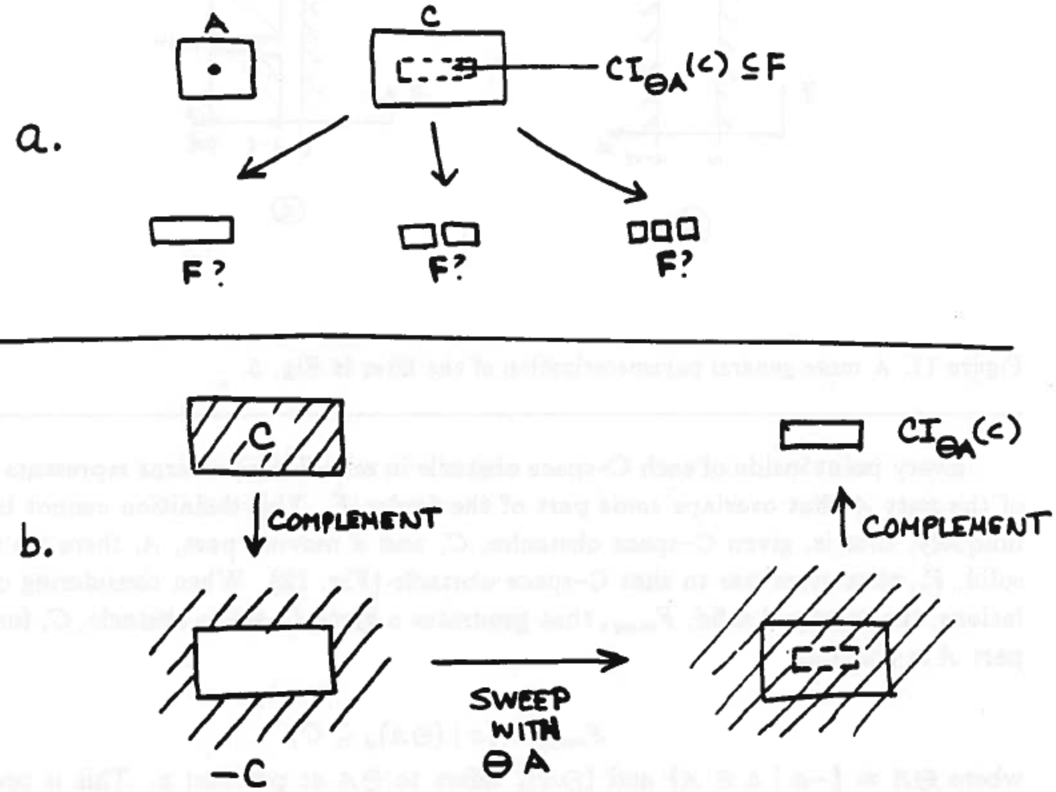

Figure 12. (a) Inverting a C-space obstacles does not produce a unique object. (b) The CI operation can be computed by subtracting the empty volume from a solid.

*Swept volume algorithms are of interest in several areas of robotics and CAD, see for example [dePennington, Bloor, Balila 83].

Each accept and reject pattern, P_i, generates two important volumes: $A[O_i]$ and $F_{i,max}$. These volumes constrain the shape of the feeder in the following way: the C–space obstacles for A due to F satisfy the accept and reject patterns and that $F \subseteq F_{i,max}$ for all i. An equivalent formulation is: $F \cap A[O_i] = \emptyset$ for all i, that is, that is, the feeder solid does not intrude into any of the swept volumes for any of the patterns. This can be computed by

$$F_{total} - \bigcup_i A[O_i]$$

We can carry out the procedure outlined above, say for the pattern in Fig. 5 and the rectangular part. The resulting solids (see Fig 13) are $A[O_a]$, $F_{a,max}$, $A[O_r]$, $F_{r,max}$, where the a subscript refers to accept and the r to reject. Before these solids can be composed to form the final feeder, we must choose their positions relative to each other. We saw an example earlier where the difference between two filters with the same accept and reject patterns, Figs. 5 and 6, was due solely to the relative position of the swept volumes. The space of feeders with identical accept and reject patterns is defined by the correspondence of features such as feeder walls. In our example, we choose to align the right–hand walls.

In our simple example, we find that the swept volume $A[O_r]$ obtained from the reject pattern intersects the maximum feeder $F_{a,max}$ obtained from the accept pattern (Fig. 13). In fact, this will be the case for all legal relative positions of these volumes. On the surface, this would mean that it is not possible to build a feeder that achieves both the desired accept and reject behavior for the rectangular part. Furthermore, it is easy to see that this will be the case for all part shapes and this particular combination of patterns. We have already seen, however, a filter that implements these patterns (see Fig. 5). What has gone wrong?

The answer is simple. Recall that we have only computed $F_{a,max}$, that is, the maximal volume for the feeder. If we simply subtract the swept volume $A[O_r]$ from $F_{a,max}$, the resulting solid still gives rise to the requisite C–obstacles in the accept pattern. Other examples illustrating the same phenomenon can be seen in Fig. 13.

It is not possible to solve all cases of intersecting volumes in this simple way; Fig. 14 is one such case. The overlap here is too large, and yet we know that a filter can implement precisely this specification by using a groove in the track (Fig. 14). The problem in this case is due to the assumptions we made in defining two–dimensional accept and reject patterns. We had to fix the other four degrees of freedom of the part: three rotational degrees of freedom and z. In reconstructing the shape of the feeder from the lower–dimensional C–space specifications, we cannot ignore the values of these parameters, especially z. The example in Fig. 14 is problematic because it involves motion on two support planes, that is, different z values. The examples in Fig. 13, on the other hand, only require motion in a single plane.

The problem illustrated in Fig. 14 can be corrected simply by requiring that the $A[O_a]$ volume be lowered by the height of the part. This defines the expected groove in the track. Many combinations of patterns require this type of operation to define a legal feeder. In some cases, as in Figs. 5 and 6, the different choices of z values for the volumes produce legal filters that are only slightly different.

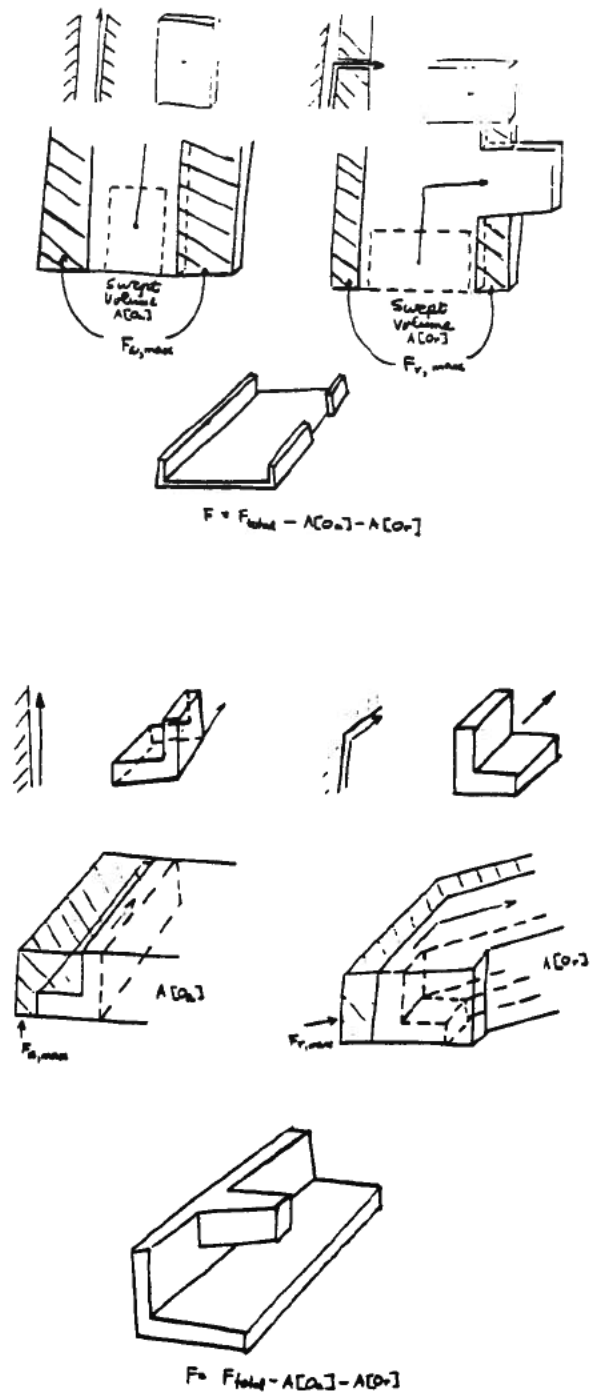

Figure 13. Two examples where the volumes from the accept and reject patterns intersect.

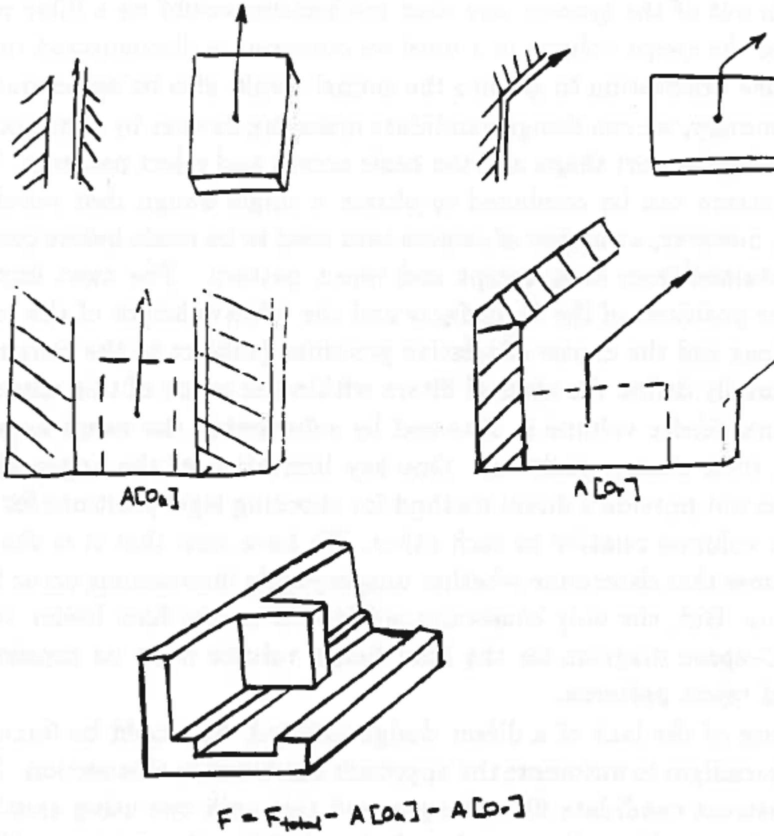

Figure 14. An example that requires different support planes

There are limits, however, to the displacement between the volumes. The essence of feeder filters is that all the part orientations travel along common segments of track before reaching the orienting device, where they are separated by the filter. Thus, the swept volumes for different part orientations must overlap sufficiently so that the parts can enter them from the common track. The filter, however, must also guarantee that parts do not enter the swept volume reserved for other part orientations.

We can illustrate both of these restrictions using a variation of the filter in Fig. 13. Interchange the assignments of accept and reject patterns for the two orientations of the rectangular part in Fig. 13, that is, accept the $\theta = 90^o$ orientation and reject the $\theta = 0^o$ orientation. Applying our approach produces a feeder track analogous to that in Fig. 6 but the groove produced by the swept area of the part in $\theta = 90^o$ orientation is too wide. If the swept areas were connected, we would need a mechanism for keeping parts in the $\theta = 0$ orientation out of the groove; any such mechanism would be a filter in itself. Similarly, if we displace the swept volume in z until we construct a disconnected tunnel, any method of selecting one orientation to go into the tunnel would also be an adequate filter on its own.

In summary, we can design candidate orienting devices by constructing maximal feeder volumes using the part shape and the basic accept and reject patterns. The volumes defined by each pattern can be combined to obtain a single design that satisfies all the patterns. There are, however, a number of choices that need to be made before combining the maximal feeders obtained from each accept and reject pattern. The most important of these are: the relative positions of the C–surfaces and the relative height of the volumes. The pattern combinations and the choice of relative positions (subject to the constraints defined above) combinatorially define the class of filters within the scope of this approach.

We final feeder volume is obtained by subtracting the swept areas from the maximal feeders at their chosen positions. One key limitation of the approach described above is that it does not provide a direct method for choosing legal positions for the maximal feeders and swept volumes relative to each other. We have seen that it is the relative positions of these volumes that determine whether unacceptable interactions occur between the different orientations. But, the only constraint we have is on the final feeder volume itself, namely, that the C–space diagram for the final feeder volume must be consistent with the chosen accept and reject patterns.

Because of the lack of a direct design method, we would be forced to use a *generate–and–test* paradigm to automate the approach described in this section. That is, an algorithm would construct candidate filter designs and test each one using existing C–space analysis tools (see Section 3) until one is found that satisfies the input specification. Fortunately, the main strength of the possible–motion representation is that the class of filter designs is relatively small. Therefore, a generate–and–test approach to automation is feasible.

In the rest of the section we present a more detailed discussion of some specific kinds of orienting devices.

4.3 Gaps and Narrowed Tracks

All feeder filters ultimately work by removing the support of some parts and not others. We have been using a very conservative model of support in the previous sections; essentially requiring that a surface be available under the complete object. In this subsection we briefly examine a more general model of support and how it impacts the design of gaps and narrowed tracks.

Roughly speaking, a body is stably supported if its center of mass, projected along the direction of gravity, lies on a support surface. Support surface is used here in a special sense: the convex hull of the points of contact between the body and the track. Thus, any three non–collinear points of contact that span the projection of the part's center of mass constitute a support surface. It follows that if there exists some half space containing all the support points for a body but that does not include the projection of the center of mass, then the body is not stably supported. More generally, the part is stably supported when all its possible incremental displacements raise the part's center of gravity (see Appendix B).

One of the most common feeder filters is the narrowed track. The basic idea for the narrowed track is to remove the part's support so as to allow the part to rotate around an axis parallel to the direction of feeding (see Fig. 10, reject pattern 5A). Another method is to introduce a gap where the part rotates around an axis perpendicular to the direction of feeding (see Fig. 10, reject patterns 5B and 6).

When the desired axis of rotation is known, we can construct a swept volume which guarantees that there exists* a position where the part will rotate around the axis. Simply pass the rotation axis through the center of mass, then rotate the part 90^o around the axis. The swept volumes produced in this way are used to define maximal feeders in the same manner as the swept volumes in the previous section. Applying this method to a rectangular part produces the traditional narrowed track illustrated in Fig. 1.

To be effective, a filter must not only guarantee that parts in the reject orientation fall off the track but also that other orientations of the part do ***not*** fall off. Therefore, the maximal feeder obtained from the rotational swept-volume must guarantee support for those other orientations. Note that whereas it is sufficient to ensure the existence of one position at which the reject orientation falls off the track, the feeder must provide support for all positions of the other orientations. Clearly, filters are not possible for completely symmetric shapes such as squares and circles unless the presence of an internal weight or cavity shifts the center of mass from the symmetry axis. Even asymmetric parts, however, can fail the conditions given above. Appendix B outlines how a test for support over a range of orientations can be implemented using C-space computations.

For a complex part shape, the swept-volume method may produce excessively complex track edges. The swept-volume, however, can be simplified by enclosing it in a simpler volume. Clearly, any volume that completely includes the rotational swept-volume cannot support the reject-orientation. So, for example, any swept volume can be conservatively approximated by an enclosing rectangular solid, thus producing a narrowed track. But, such radically simplified swept volumes may produce tracks that fail to support the accept-orientations.

Holes or concavities in the part's rotational swept-volume can generate protrusions, called ***scallops***, in the edge of the feeder track. These scallops can be effective in situations where a simple narrowed track is not (see Fig. 15).

The swept-volume method can also be used to design shaped gaps of the type illustrated in Fig. 3. Although the shaped gap is not strictly a filter, it can be designed using the same methods. The basic idea of a shaped gap is to allow the part to rotate (around an axis perpendicular to direction of feeding) only towards one of its ends. In designing shaped gaps, one assumes that the part has already been constrained to move along a line, usually by means of a groove in the track, as shown in Fig. 3.

*This is a sufficient but not necessary condition for rotation around the specified axis.

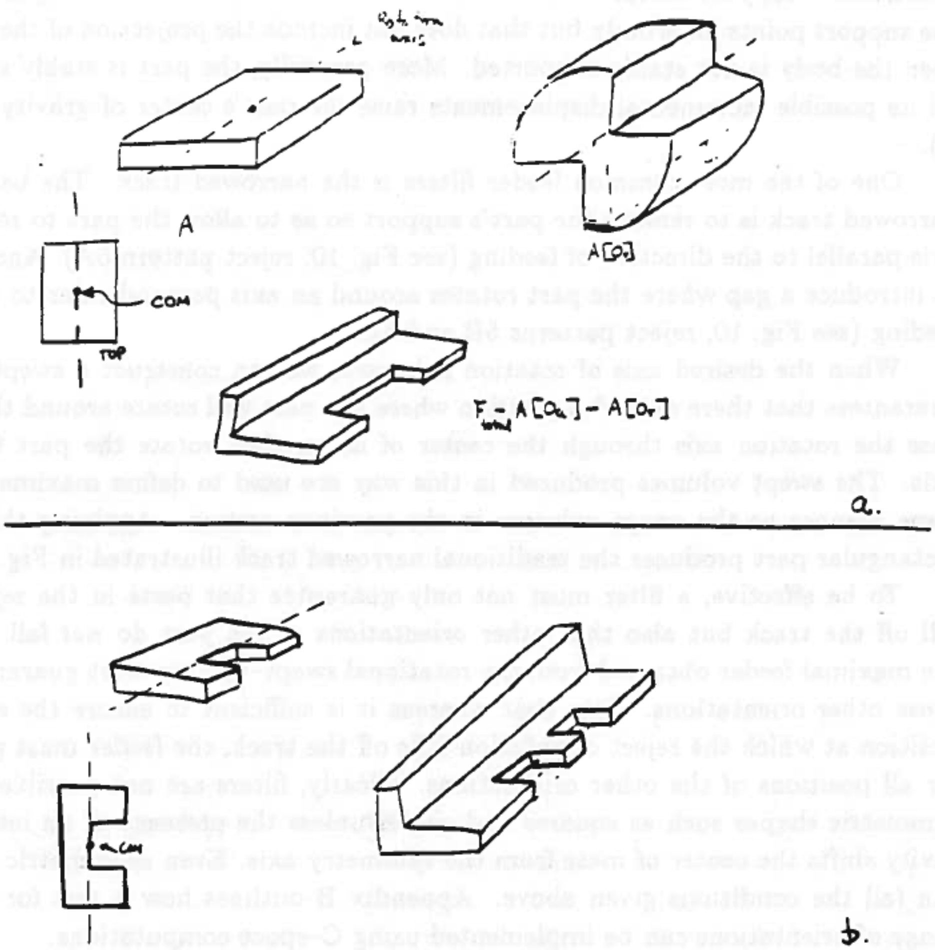

Figure 15. Examples of two simple feeder filters defined by a rotational swept volumes

Each of the legal rotations of the part generates a swept volume and, hence, a gap shape. The gaps could be separately placed in the track, with separate output channels below them. But, they can also be superimposed into one gap as in Fig. 3. The superposition is governed by the constraint that the combined gap must allow only the desired rotations (Fig. 16).

For simplicity, assume that both orientations of the part are symmetric around an axis through the center of mass which is parallel to the feeding direction. Cut the part into two along a line through the center of mass which is perpendicular to the feeding direction. This is illustrated in Fig. 16. The constraints on the combined gap are as follows:

- **There must be a part position for which the leading half of orientation 1 is completely inside the gap (Fig. 16a). That is, there must be a position where orientation 1 would fall into the gap leading edge first.**

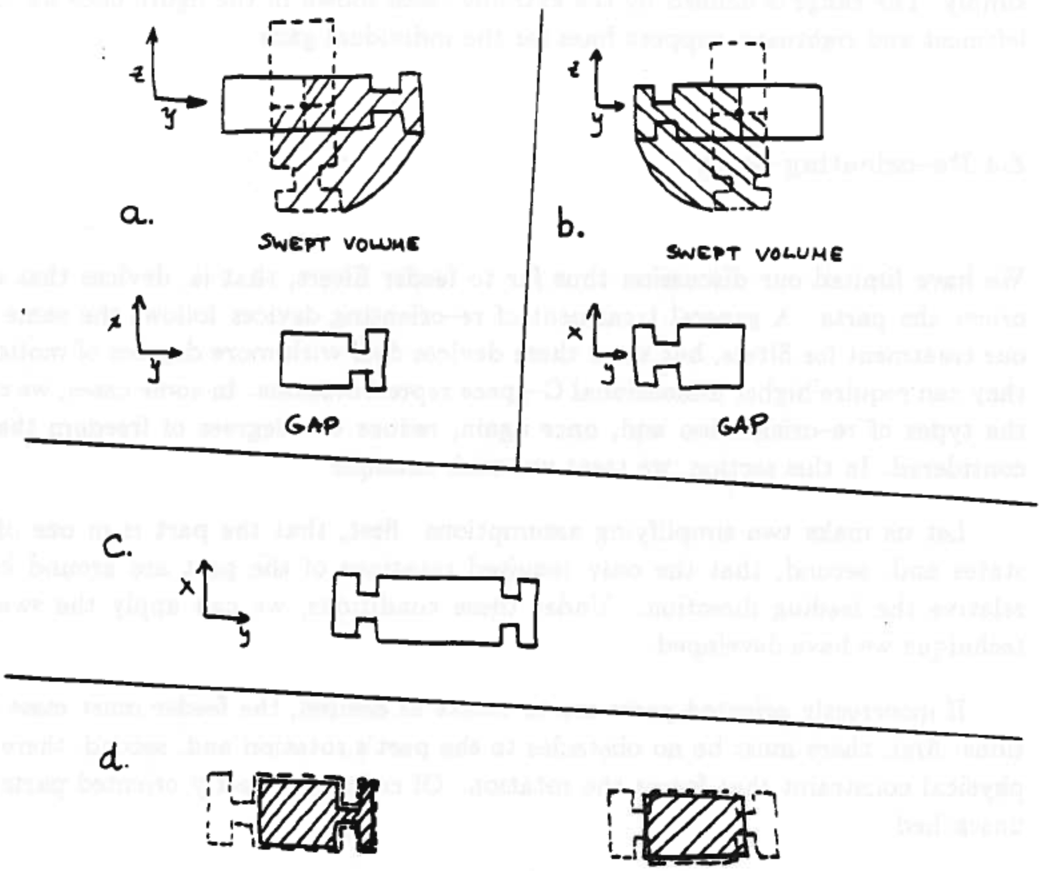

Figure 16. The design of shaped gaps: (a) the gap obtained by sweeping the leading half of orientation 1, (b) the gap obtained by sweeping the leading half of orientation 2, (c) the combined gap, (d) extreme cases for gap sizing.

- **There must *not* be a part position for which the leading half of orientation 2 is completely inside the gap (Fig. 16b). That is, there must not be a position where orientation 2 would fall into the gap leading edge first.**
- **There must be a part position for which the trailing half of orientation 2 is completely inside the gap (symmetric case of Fig. 16a). That is, there must be a position where**

orientation 2 would fall into the gap trailing edge first.

The legal range of relative positions of the two gaps shown in Fig. 16 can be computed simply. The range is defined by the extreme cases shown in the figure once we identify the leftmost and rightmost support lines for the individual gaps.

4.4 Re–orienting parts

We have limited our discussion thus far to feeder filters, that is, devices that do not re-orient the parts. A general treatment of re–orienting devices follows the same outline as our treatment for filters, but since these devices deal with more degrees of motion freedom they can require higher dimensional C–space representations. In some cases, we can restrict the types of re–orientation and, once again, reduce the degrees of freedom that must be considered. In this section, we treat one such example.

Let us make two simplifying assumptions: first, that the part is in one of its stable states and, second, that the only required rotations of the part are around known axes relative the feeding direction. Under these conditions, we can apply the swept–volume technique we have developed.

If incorrectly oriented parts are to rotate as desired, the feeder must meet two conditions: first, there must be no obstacles to the part's rotation and, second, there must be a physical constraint that forces the rotation. Of course, correctly oriented parts must pass unscathed.

The first of these conditions can be met by the now familiar technique of defining a swept volume for the part as it simultaneously rotates and translates. This swept volume can then be used to define a maximal feeder. Fig. 17 illustrates swept volumes generated from the rectangular part in this fashion. Subtracting this swept volume from the potential feeder solid creates a channel for the part along its path. The second condition requires that the feeder's motion be able to push the part through this channel without jamming the part. It is the likelihood of frictional jamming that makes re–orienting devices difficult to design and somewhat unreliable. It is possible, however, to exploit frictional constraints to orient parts; see [Mason 82, 85, 86, Erdmann and Mason 86, Brost 85, Mani and Wilson 85] for detailed developments of this idea.

Fig. 17 also shows the type of devices used to re–orient rectangular parts in practical feeders. The actual device can be derived from the method suggested here by exploiting gravity. In particular, the feeder does not need a top surface to hold the part, gravity will serve; this will also reduce the likelihood of jamming. Note that the swept–volume method would also generate the channel next to the feeder wall, which allows correctly oriented parts to pass through.

The dimensions of this volume are determined from bounds on the magnitude of jumps and the "wobble" of the part. These dimensions in turn determine the minimum dimension of any channels in the C–space of the part relative to the filter.

The effect of enlarging the C–space channels is equivalent to further sweeping the original swept volume over the range of positions and orientations in the uncertainty volume. The resulting passages and gaps will allow the part to pass in any position and orientation within the uncertainty range. This process also identifies marginal designs that depend on very small distinctions in shape and location. Note that this additional sweeping can be used to take into account the effect of manufacturing tolerances in the part and in the feeder.

6. Prospects for Automatic Design of Part Feeders

We have attempted to show that the sometimes complex geometry of orienting devices arises from the interaction of some simple part motions. This led us to the use of swept volumes over ranges of possible motions as the basic design mechanism for orienting devices. We have been unable as of yet to incorporate all the constraints on the relative positions of these swept volumes into a direct design method. It does seem, however, that a generate–and–test approach to automatic feeder design may be feasible.

This investigation of feeder design, like the general investigation of shape from function of which it is a part, is in its early stages. A large number of practical considerations have been neglected. Nevertheless, it is noteworthy that the majority of the feeder designs in the U. Mass Handbook of Feeding and Orienting Techniques for Small Parts [Boothroyd, Poli, and Murch 77] are within the scope of the simple methods outlined in this paper.

I believe that the specific method for feeder design described here holds promise both as a design methodology for humans and as a basis for automated design.

Acknowledgments

My thanks to Matt Mason for his help in formulating the problem, to Bruce Donald and Mike Erdmann for their helpful comments on an earlier draft, and to Neil Singer for discussions on parts feeders.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's Artificial Intelligence research is provided in part by a grant from the Systems Development Foundation, and in part by the Advanced Research Projects Agency under Office of Naval Research contracts N00014-80-C-0505 and N00014-82-K-0334. The author's research is also supported by an NSF Presidential Young Investigator grant.

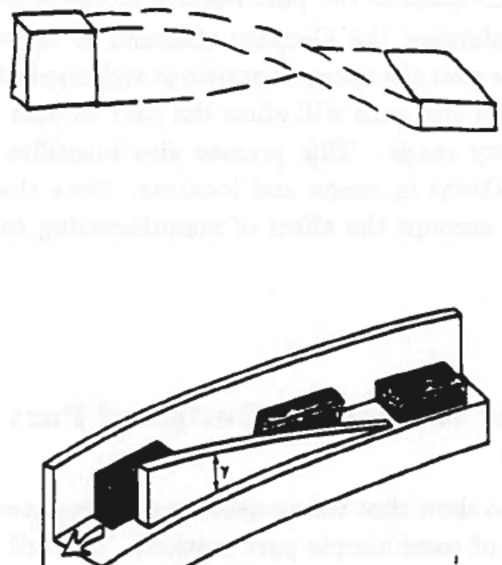

Figure 17. Constructing devices to re–orient parts [Boothroyd, Poli, and Murch 82]

5. Motion Uncertainty

In this paper we have dealt only with geometric considerations in the analysis and design of orienting devices. In practice, the physics of feeders is crucial in understanding which designs are feasible and in determining their relative efficiency [Boothroyd, Poli, and Murch 82]. A thorough treatment of the impact of the physics of feeders on this approach to design of orienting devices awaits further work. There is, however, one important physical consideration that cannot be neglected.

The motion of parts in vibratory bowl feeders is composed of sliding motions interspersed with jumps. Although it is difficult to model the detailed motion of parts, it is sufficient for our purposes to understand that parts are not in continuous contact with the feeder surfaces. So, for example, a gap that relied on removing a part's support at a single specific orientation could not work reliably; the part may jump over that position.

This effect can be readily modeled in the C–space of the part. We can treat the part in its C–space not as a point (a single position and orientation) but as a volume (a range of positions and orientations). This C–space volume represents our uncertainty in the position and orientation of the part[Lozano–Pérez, Mason and Taylor 84, Erdmann 84].

References

H. Asada and A. B. By, "Kinematics of Workpart Fixturing", IEEE Int. Conf. on Robotics and Automation, St. Louis, MO, March 1985.

G. Boothroyd, C. R. Poli, and L. E. Murch, *Handbook of Feeding and Orienting Techniques for Small Parts*, Dept. of Mech. Eng., U. of Mass., Amherst, MA., 1977 (updated July 1981).

G. Boothroyd, C. R. Poli, and L. E. Murch, *Automatic Assembly*, Marcel Dekker, Inc., New York, 1982.

R. A. Brooks and T. Lozano-Pérez, "A Subdivision Algorithm in Configuration Space for Findpath with Rotation", to appear in *IEEE Transactions on Systems, Man and Cybernetics.*

R. C. Brost, "Planning Robot Grasping Motions in the Presence of Uncertainty," Computer Science Department and the Robotics Institute, Carnegie-Mellon University, CMU-RI-TR-85-12, 1985.

J. F. Canny, "Collision Detection for Moving Polyhedra", MIT AI Memo 806, October 1984.

A. dePennington, M. S. Bloor, and M. Balila, "Geometric Modeling: A Contribution Towards Intelligent Robots", Seventh Int. Symp. on Ind. Robots, Chicago IL., April 1983, pp. 7-35 – 7-54.

B. R. Donald, "Motion Planning with Six Degrees of Freedom", MIT AI Technical Report 791, May 1984.

M. Erdmann, "On Motion Planning with Uncertainty", MIT AI Technical Report 810, August 1984.

M. Erdmann and M. T. Mason, "An Exploration of Sensorless Manipulation", *Proceedings of the 1986 IEEE Int. Conf. on Robotics and Automation*, San Francisco, April 1986.

T. Lozano-Pérez, "Automatic Planning of Manipulator Transfer Movements", *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-11, No. 10, 681 – 698, October 1981.

T. Lozano-Pérez, "Spatial Planning: A Configuration Space Approach", *IEEE Transactions on Computers*, Vol C-32, No. 2, 108 – 120, February 1983.

T. Lozano-Pérez, "Shape from Function", MIT AI Memo, January 1986.

T. Lozano-Pérez, M. T. Mason, and R. H. Taylor, "Automatic Synthesis of Fine-Motion Strategies for Robots", *International Journal of Robotics Research*, Vol. 3, No. 1, 3 – 24, 1984.

M. Mani, and W. Wilson, "A Programmable Orienting System for Flat Parts," *Proc., NAMRI XIII*, Berkeley, Calif., May 19–22, 1985.

D. Marr, *Vision*, W. H. Freeman and Co., San Francisco, CA., 1982.

M. T. Mason, "Compliance and Force Control for Computer-Controlled Manipulators". *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-11, No. 6, June 1981.

M. T. Mason, "Manipulator Grasping and Pushing Operations", MIT AI Technical Report 690, June 1982.

M. T. Mason, "The Mechanics of Manipulation," *Proceedings of the 1985 IEEE Int. Conf. on Robotics and Automation*, St. Louis, Missouri, March 25–28, 1985, pp. 544–548.

M. T. Mason, "Mechanics and Planning of Manipulator Pushing Operations," *International Journal of Robotics Research* 5, 1 (1986).

A. H. Redford and G. Boothroyd, "Vibratory Feeding", *Proc. I. Mech. Eng.*, Vol 182, Part 1, No. 6, 1967–1968.

A. H. Redford, E. K. Lo, and P. J. Killeen, "Parts Feeder for a Multi-Arm Assembly Robot", 15th CIRP Seminar on Manufacturing Systems, Amherst, MA., June 1983.

P. H. Winston, *Artificial Intelligence*, Second Edition, Addison–Wesley Publishing Co., Reading, MA., 1984

Appendix A. Configuration Space

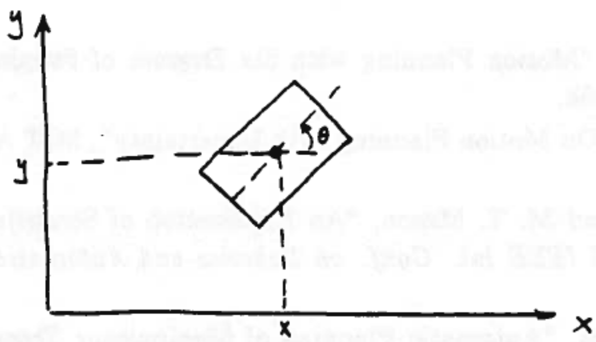

Figure 18. The configuration parameters for a rectangular part in the plane

Consider a parameter space of positions and orientations of the part. A point in this space represents a unique position and orientation of the part; each point is known as a *configuration*. Such a parameter space of positions and orientations is known as the *configuration space (C–space)* of the part [Lozano–Pérez 81, 83]. The position and orientation of a rectangular part that is free to move in a plane, for example, can be specified by three parameters, x, y, and θ. In this description, the vector (x, y) indicates the position of some point on the part relative to a coordinate system fixed in the plane, and θ indicates the angle between some edge of the part and the coordinate system's x–axis (see Fig. 18). There are many possible parameter spaces that can serve as configuration spaces, for example, using polar coordinates, instead of cartesian coordinates, encode the position of a point on the part.

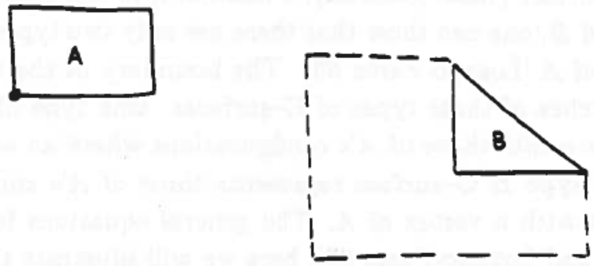

Figure 19. The configuration space obstacle for a moving rectangle due to the presence of a triangular obstacle.

More generally, specifying the position and orientation of a rigid body in space requires six parameters, for example, three linear displacements and three angular displacements.

The presence of other bodies in the environment of a moving body creates obstacles in the moving body's C–space. These *configuration space obstacles* are the collection of all configurations of the moving body in which it would overlap some other body [Lozano–Pérez 81,83, Brooks and Lozano–Pérez 83, Donald 84]. Consider a rectangular polygon that is free to translate, but not rotate, in the plane (see Fig. 19). The rectangle's configuration can be specified by the (x, y) coordinates of the specified point relative to a fixed coordinate system. The presence of a triangular obstacle generates a forbidden region of the C–space; this is the C–space obstacle.

The surfaces of these C–space obstacles can be used to characterize geometrical and physical interactions among bodies [Mason 81, Lozano–Pérez, Mason, and Taylor 84, Erdmann 84]. These surfaces behave much as actual surfaces, including the generation of reaction and frictional forces. In fact, an adequate intuitive model is to treat the moving object in C–space as a point mass interacting with rigid obstacles. This view is, under fairly general conditions, equivalent to the original situation [Erdmann 84].

Appendix B. Characterizing Stability in C–space

Consider an object A whose motion is constrained by other objects B_i so that: (a) any pure translation of A in the direction of gravity is blocked by the B_i, and (b) any pure rotation causes A's center of mass to move up relative to the direction of gravity. It is easy to see that A is stably supported by the B_i. These conditions are geometric constraints on the configuration of A relative to other objects. These constraints can be represented in the configuration space of A.

We saw that the configurations of an object A for which it is in contact with another object B form a surface (more generally, a manifold) in A's configuration space. For polygonal objects A and B, one can show that there are only two types of these C–surfaces in the (x, y, θ) C–space of A [Lozano-Pérez 83]. The boundary of the C–space obstacle due to B is made up of patches of these types of C–surfaces. One type of C–surface, conventionally called type A, represents those of A's configurations where an edge of A is in contact with a vertex of B. A type B C–surface represents those of A's configurations where an edge of B is in contact with a vertex of A. The general equations for these C–surfaces can be found in [Brooks and Lozano-Pérez 82]; here we will illustrate their use with a simpler set of equations.

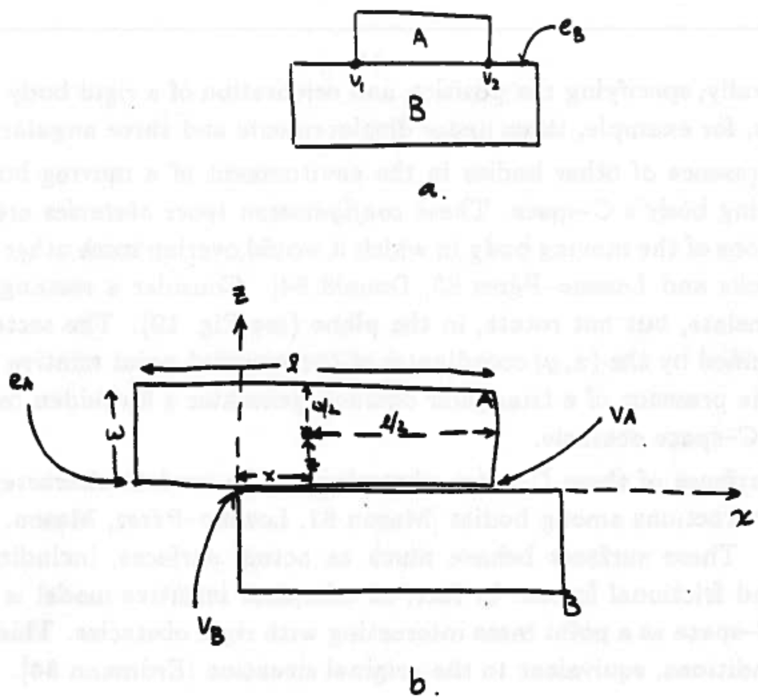

Figure 20. (a) Intersection of two type B C–surfaces. (b) Intersection of type A and type B C–surfaces

Fig. 20 shows two cases of a movable object A and a stationary obstacle B, with A in edge–edge contact with B. In Fig 20a, the configuration of A is at the intersection of two type B C–surfaces. This follows from the fact that two vertices of A, v_1 and v_2, are in contact with an edge of B, e_B. In Fig 20b, the configuration of A is at the intersection of a type A C–surface and a type B C–surface. This follows from the fact that two vertices of A, that an edge of A, e_A, is in contact with a vertex of B, v_B, and an edge of B, e_B, is in contact with a vertex of A, v_A. As A rotates counterclockwise from its initial position in Fig. 20b, its configuration leaves the type B C–surface but remains on the type A C–surface. As it rotates clockwise, its configuration leaves the type A C–surface remains on the type B C–surface.

To show the stability of A, the interesting constraint is showing that any rotation of A raises the position of its center of mass. Using the location of A's center of mass to characterize the position of A, we want to show that $dz/d\theta$ is positive for legal motions of A. In this case, A's motions are constrained to be on C–surfaces as above.

We can derive an expression that relates the configuration parameters x, z, θ on a type A C–surface using the nomenclature in Fig. 20b:

$$-x \sin\theta + z \cos\theta = \frac{w}{2}$$

Hence

$$z = \frac{x \sin\theta + \frac{w}{2}}{\cos\theta}$$

This expression is valid only for $\theta \geq 0$; we can see that $dz/d\theta$ is positive for $\theta \geq 0$ as long as $x > 0$. This captures the fact that A will become unstable when the center of mass is no longer over the support edge.

We can derive another expression that relates the configuration parameters on a type B C–surface:

$$z = \frac{w}{2}\cos\theta - \frac{l}{2}\sin\theta$$

This expression is valid only for $\theta \leq 0$; in this case, $dz/d\theta$ is positive for $\theta \leq 0$.

These conditions guarantee that A is stably supported in Fig. 20b. A similar analysis can be used in three dimensions.

The more general methods reported in [Brooks and Lozano-Pérez 81, Donald 84, Canny 84] can be adapted to determine numerically the stability of polyhedral parts along along arbitrary paths.