Orienting Objects in a Minimum Number of Robot Sweeping Motions

Graham E. Deacon, Poh Lian Low, and Chris Malcolm

DAI Research Paper No. 619

This paper appears in the Proceedings of the Sixth International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems, pp.494–503, Edinburgh, Scotland, June 1993.

ORIENTING OBJECTS IN A MINUMUM NUMBER OF ROBOT SWEEPING MOTIONS

Graham E. Deacon, Poh Lian Low, and Chris Malcolm
Department of Artificial Intelligence, University of Edinburgh,
5 Forrest Hill, Edinburgh, EH1 2QL, Scotland.

Abstract

The technique described in this paper is an application of Brost's push-stability diagram, a graphical representation of the qualitative behaviour of an object when pushed in a horizontal plane by a flat fence, that is derived from some results of Mason. We describe an algorithm that uses this diagram to calculate the minimum number of robot sweeping motions needed to put a randomly oriented planar object into a known orientation, and what these motions are.

1 Introduction

This paper describes a system that uses the push-stability (P-S) diagram of Brost [2],[3], in order to derive a sequence of robot motions that will transform an object (lying on a horizontal planar surface) from an initial random orientation into a known orientation. The P-S diagram is a graphical representation of the qualitative behaviour of a sliding workpiece when pushed by a flat fence, and is derived from some basic results of Mason [11], [13]. We will briefly describe Mason's and Brost's work before describing the algorithm in detail.

2 Background

Under the assumption of quasi-static conditions, Mason [11], [13] derived a rule that predicts the direction of rotation of a pushed object. This rule uses the notion of a friction cone, a geometric expression of Coulomb's law of friction, which delimits the range within which (reactive) frictional forces must lie. At the point of contact between the pushing and pushed objects we construct three rays: two that represent the limits of the friction cone, and one in the direction of the pusher's motion. Mason's rule says that the direction that a part will rotate is decided by how these three rays 'vote' relative to the part's centre of friction (COF)[1]. If the majority pass to the left of the COF the object will rotate clockwise, to the right anticlockwise, and if there is one ray either side of the COF, and one directly through it, the object translates.

A corollary to this rule is that if we use a flat fence as the pusher certain object orientations (with respect to the fence) are 'stable' in the sense that if the orientation is changed slightly in either direction, subsequent pushing motion returns the part to the initial orientation: see figure 1 for an example. In addition, if such stable orientations exist (for a given value of the coefficient of friction between pushing and pushed objects) a sufficiently long sweep will bring the part to rest in one of them.

All the possible outcomes of Mason's rule for all possible orientations of a part are represented by Brost's push-stability diagram [2], [3]. Brost creates what he calls an operation space by plotting possible pushing directions, δ, against possible part orientations, ϕ. He then investigates the conditions under which each of the rays change their vote. This partitions the operation space according to 'decision boundaries'. Tallying the votes for each region decides which direction the part will rotate for any point (δ, ϕ) in that region. The partitions left in the space are the boundaries about which the sense of the part's rotational motion changes.

An example of a push-stability diagram and a diagram defining the notation used are given in figures 3 and 2 respectively. In constructing a P-S diagram Brost assumes straight line motion of the pusher, and that the fence is sufficiently large that the possibility of the object rolling off the end can be neglected. A consequence of this last assumption is that for the purposes of the P-S diagram an object's geometry is equivalent to its convex hull.

The thick arrows in figure 3 indicate the direction that the object will rotate if a (δ, ϕ) pair from that region applies. The thick lines that run horizontally across the diagram are known as convergent boundaries since the rotation directions either side of these lines lead to the object rotating towards this orientation. The thinner partitioning lines are known as divergent boundaries because the object tends to roll away from them. Convergent boundaries correspond to the

[1] When pushing an object the reactive force due to friction experienced by the pusher appears to be coming from just one point. This point Mason calls the centre of friction [11], [13]. Often this point can be approximated to by the centre of mass.

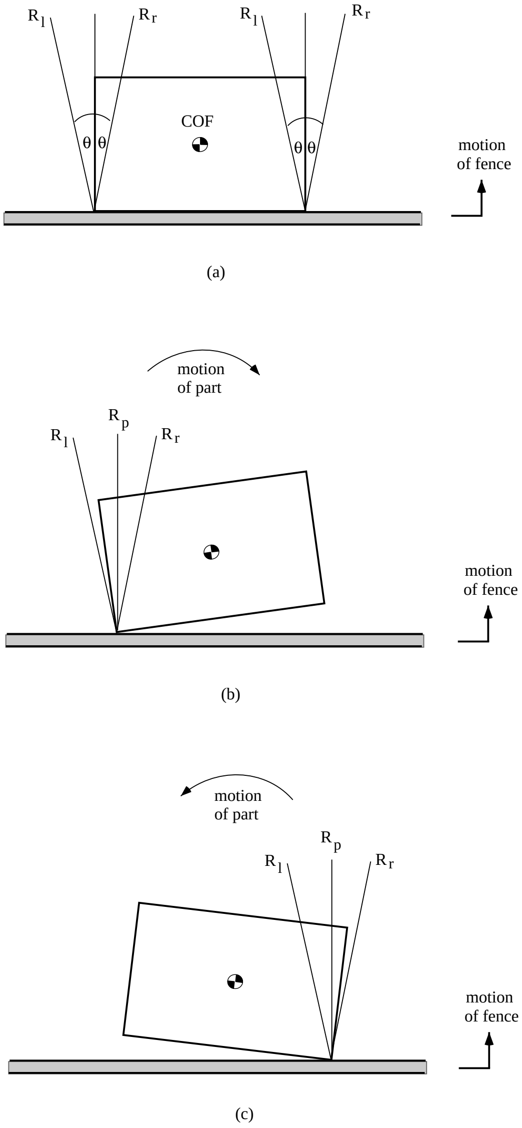

Figure 1: Pushing a block with a straight fence. R_l is the left ray of the friction cone, R_r the right ray, and R_p the the direction of pushing. θ is the friction cone half-angle between the fence and the block. (a) shows a stable orientation, and (b) and (c) slight 'perturbations' from it.

stable situations described above, and divergent boundaries to the situation where there is a ray either side of the COF and one straight through it.

Since the value of δ is restricted to being constant throughout a pushing operation the motion of the part when pushed can be represented as a straight line trajectory on the P-S diagram. Figure 3 includes three starting configurations labelled (a), (b), and (c). In situations (a) and (c) the initial orientations of the part are the same, but the final stable orientations are different because of different pushing directions. In situations (b) and (c) the pushing directions are the same, but the difference in initial orientations results in different stable orientations.

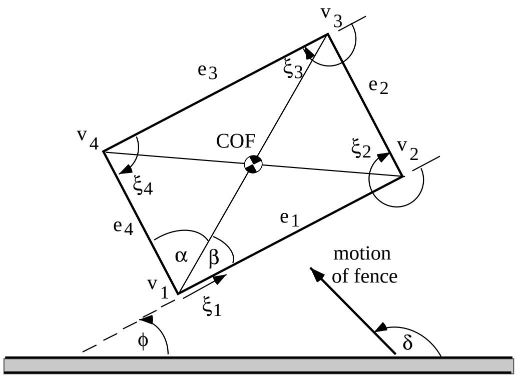

Figure 2: Notation used in the P-S diagram. The edges are labelled $e_1, e_2, \ldots, e_n$ in an anticlockwise direction starting from the longest edge. The vertices are labelled in a similar manner with vertex v_i on the anticlockwise edge of e_i. In addition, each edge has an edge angle ξ_i, which is the angle measured clockwise from e_1 to e_i. ϕ is the angle of orientation and δ the angle of pushing. The interior angles of the block, α and β, are also shown since they can be used to parameterise the P-S diagram (see figure 3).

3 Using the P-S Diagram to Orient a Part Uniquely

3.1 The Basics of the Approach

The P-S diagram is in fact an encoding of a type of situation that Mason has called a manipulation funnel [12]:

> "Just as a common funnel eliminates uncertainty in the location of a grain of sand, a manipulation funnel eliminates uncertainty in the location of objects. Using a funnel, the goal, to position an object, can be accomplished despite variation in the initial locations and shapes of the objects."

We can use the P-S diagram to tell us what ranges of part orientation and pushing directions will result in the same final orientation: all (δ, ϕ) that lie in the same region of the P-S diagram will, after moving the fence a sufficient distance, result in the object resting stably against the fence in a particular orientation. This does not completely constrain the position of the part since it may lie anywhere along the fence due to its slipping or rolling relative to the fence. However, the questions of whether a part will slip, and if it does how far, and

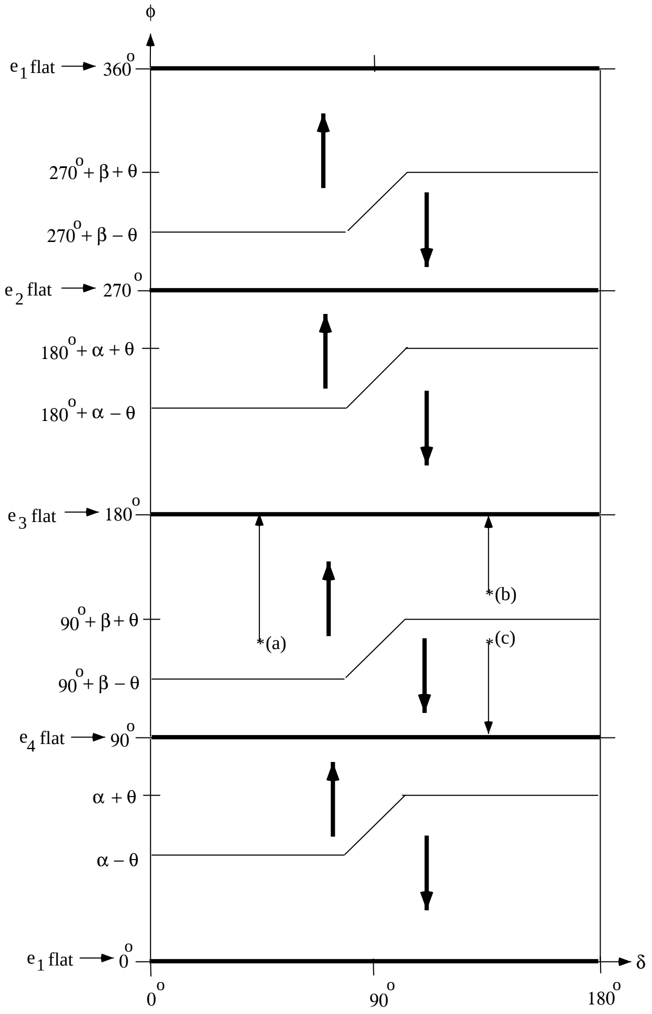

Figure 3: The push-stability diagram for a block. The label "e_i flat" denotes the situation where edge e_i lies flat against the pushing fence. Because of the way the polygon is labelled, the values of ϕ at which an edge of the object is flat against the fence are just the ξ_i angles (see figure 2).

just how far is a 'sufficient distance' to move the fence so that alignment is guaranteed, have been addressed by Peshkin and Sanderson [15]. For the moment what is important is that the P-S diagram tells us what orientations are stable, and what the conditions are that lead to them.

We will illustrate the basic technique that we are proposing with an example. This example will use a P-S diagram derived by Brost [2], [3] for a triangular object whose COF is not coincident in the sliding plane with its centre of area, see figure 4. For the moment we will confine ourselves to considering the vertical band of this P-S diagram defined by $\delta \in]90^o + \theta, 180^o]$. Initially the part is in an unknown orientation, but we know from inspection of the P-S diagram that a push where $\delta \in]90^o + \theta, 180^o]$ that is sufficiently long, will result in the part being in either stable orientation e_1, or stable orientation e_3. We can't know which one since we don't know the initial orientation. Let's call the pushing orientation used δ_p. In this case the part will have 'come to rest' in the operation space at either (δ_p, e_1), or (δ_p, e_3). These points are marked with the label "(1)" next to them in figure 4(a).

The basic observation that is the key to this technique is that if we reorient the fence after a sweep, by an amount $-\gamma$, it is like shifting the possible resting orientations by an amount γ on the P-S diagram.

In our example, after the first push the part either has orientation e_1 or e_3 with respect to the fence. Rotating the fence by $-\gamma$ would mean that at the start of the next push, the part either has orientation $e_1 + \gamma$, or $e_3 + \gamma$. Let the next sweep be in the direction δ_q. Consequently, for the next sweep we must have either $(\delta_q, e_1 + \gamma)$, or $(\delta_q, e_3 + \gamma)$ as the initial conditions. Now, if we arrange, by an appropriate choice of γ, that both of these points lie in regions where pushing leads to the same convergent boundary, then after the second sweep the part must be in a known orientation (with respect to the fence). For our example, $\delta_q \in]90^o + \theta, 180^o]$, and we choose γ as shown in figure 4(a), where the points with the label "(2)" next to them are the initial conditions for the second sweeping motion. This results in a final resting point in the operation space of (δ_q, e_3), which is the point labelled with a "(3)" in figure 4(a).

Basically the technique we are proposing is to repeatedly exploit the manipulation funnel(s) that the P-S diagram encodes: the first sweeping action puts the part into one of a finite number of stable orientations, and subsequent sweeping actions reduce the number of possible orientations until such time as there is only one possibility, or no further reduction in the number of possible stable orientations can be achieved. Each change in state, where a state is defined by the set of possible orientations before or after a pushing operation, that use of the P-S diagram can 'implement', can be represented as a re-write rule. From these rules we can construct the tree that represents all reachable states. This can tell us if it is possible to uniquely orient a part by pushing alone, and if so, the 'routes', in terms of sequences of robot sweeping actions.

3.2 Generating the Search Tree

In order to get down to specifics we will find it useful to develop some notation.

First, however, we note that, as can be shown by considering the construction of the P-S diagram [4], any diagonal segments that occur in a P-S diagram only occur on divergent boundaries and must lie in the range $\delta \in [90^o - \theta, 90^o + \theta]$, *i.e.* in situations where δ points

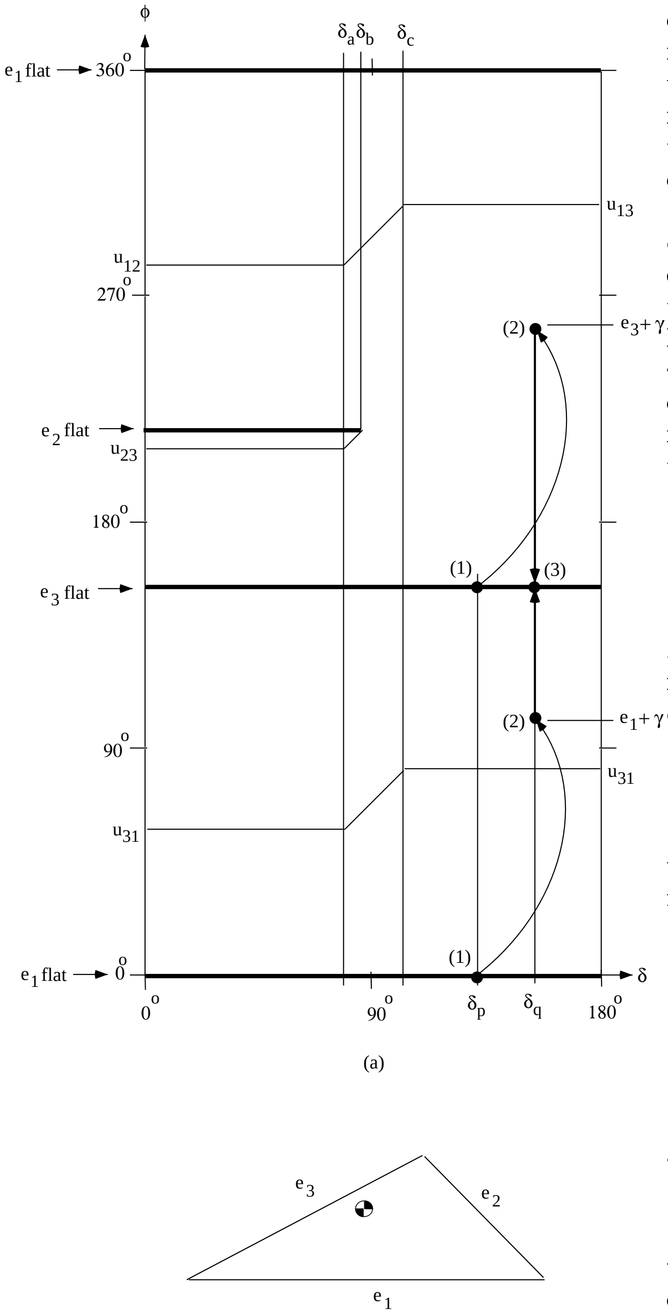

(a)

(b)

Figure 4: (a) Brost's push-stability diagram for the object shown in (b).

into the friction cone. Furthermore, the slope of such lines is always unity [2], [3].

Let $E = \{e_1, e_2, \ldots, e_n\}$ be the set of all possible stable orientations of the part, where we use e_i to denote edge i resting stably against the fence. (n is the number of e_i appearing on the P-S diagram, *i.e.* the number of edges in the part's convex hull.)

We will find it useful to be able to refer to the range of δ over which a variable is defined, so let $\mathcal{R}(e_i)$ denote the range of δ over which e_i exists. In addition, variables may have different values dependent on what range of δ we are working with. We will use superscripts to indicate the range of δ over which a particular value of a variable is valid.

It is also convenient to divide the P-S diagram up into distinct vertical bands, *i.e.* specific ranges of δ, where each band contains a different set of stable orientations to the others. To achieve this let $\mathcal{M}(e_i) = m_i$ be the length of e_i, *i.e.* if $\mathcal{R}(e_i) = r_i = [\delta_{lower}, \delta_{upper}]$ then $m_i = \delta_{upper} - \delta_{lower}$. Now let M_b be the ordered set of the lengths of those e_i that include $\delta = 0^o$ in their range, and let M_t be the ordered set of the lengths of those e_i that include $\delta = 180^o$ in their range, *i.e.*

$$\begin{aligned} M_b &= \{m_c, m_d, \ldots, m_t\}, \text{ where } m_c \leq m_d \leq \ldots \leq m_t \\ M_t &= \{m_f, m_g, \ldots, m_w\}, \text{ where } m_f \leq m_g \leq \ldots \leq m_w \end{aligned}$$

Once we have these orderings we form the sets R_b and R_t by replacing the elements of M_b and M_t by the corresponding ranges, and adding superscripts to act as indices representing where in the sequences they fall, *i.e.*

$$\begin{aligned} R_b &= \{r_c^1, r_d^2, \ldots, r_t^p\} \\ R_t &= \{r_f^1, r_g^2, \ldots, r_w^q\} \end{aligned}$$

where $p = |R_b|$ and $q = |R_t|$. From these we derive the following

$$\begin{aligned} G_b &= \{g_b | g_b = r_i^x \backslash r_j^y \text{ for } 1 < x \leq p,\ x - y = 1, \\ &\qquad \text{where } r_i^0 \stackrel{\text{def}}{=} \emptyset, \text{ and } r_i^x, r_j^y \in R_b\} \\ G_t &= \{g_t | g_t = r_i^x \backslash r_j^y \text{ for } 1 < x \leq q,\ x - y = 1, \\ &\qquad \text{where } r_i^0 \stackrel{\text{def}}{=} \emptyset, \text{ and } r_i^x, r_j^y \in R_t\} \end{aligned}$$

and we can use these to form

$$G = \{g_i | g_i = g_t \cap g_b \ \ \forall g_t \in G_t,\ \forall g_b \in G_b\}$$

which is the set of distinct ranges of δ. We next form the ordered sets E^{g_i}, where $g_i \in G$, which record the order that stable edge orientations occur on the P-S diagram for a particular distinct range of δ. In fact, because points on the ϕ axis are points on a circle, we require that the ordering on E^{g_i} be cyclic, *i.e.* the ordering 'wraps around'.

There is also a natural division in the P-S diagram arising from the values where δ is pointing into the friction cone. We will define

$$\begin{aligned} \Delta_b &= [0^o, 90^o - \theta[\\ \Delta_f &= [90^0 - \theta, 90^o + \theta] \\ \Delta_t &=]90^o + \theta, 180^o] \end{aligned}$$

in order to be able to distinguish these ranges readily.

Let $e_{ij} = e_i - e_j$ such that e_j lies clockwise of e_i, and $e_{ij} \in [0^o, 360^o[$. This gives the distance between two convergent boundaries, which we stipulate to be clockwise since the points of ϕ lie on a circle and there are two different distances between them (unless they are diametrically opposite). In order to avoid negative distances, if $e_i - e_j < 0$ we add 360^o to the result to give e_{ij}. Note that e_i and e_j do not need to be nearest neighbours on the P-S diagram.

Let u_{ij} be the divergent (or unstable) boundary between e_i and e_j, where e_j is the next stable orientation clockwise of e_i.

Let $c_{ijk} = u_{ij} - u_{jk}$ be the 'capture range' for e_j. In other words, this is a measure of the 'width' of the manipulation funnel that leads to stable orientation e_j.

Note that e_{ij}, u_{ij}, and c_{ijk} may only be defined over finite ranges of δ. Starting with $\mathcal{R}(e_i) = r_i$, which can be found from the calculations for constructing the P-S diagram, the ranges of existence for the other variables can be calculated according to:

$$\begin{aligned}\mathcal{R}(e_{ij}) &= \mathcal{R}(e_i) \cap \mathcal{R}(e_j)\\ \mathcal{R}(u_{ij}) &= \mathcal{R}(e_i) \cap \mathcal{R}(e_j)\\ \mathcal{R}(c_{ijk}) &= \mathcal{R}(u_{ij}) \cap \mathcal{R}(u_{jk})\end{aligned}$$

Let a state description, S_d^L, be the set of possible stable orientations that the part can be in after d sweeping motions, *i.e.* $S_d^L \subseteq E$, where L is a list of ordered pairs that records the ranges of values that the γ's and δ's that led to the current state must have been elements of. Consequently, for any state, the route to get there can be found from inspection of L.

In order to find the set of re-write rules that represent the action of the manipulation funnels we calculate whether

$$e_{cx} < c_{pqr}$$

for all e_{cx} and all c_{pqr} of the P-S diagram. For those variables where this inequality applies we can derive a 'convergence' rule of the form

$$(\delta, e_c + \gamma), (\delta, e_d + \gamma), \ldots, (\delta, e_x + \gamma) \longmapsto (\delta, e_q)$$

or alternatively,

$$\{e_c, e_d, \ldots, e_x\} \overset{(\gamma,\delta)}{\longmapsto} \{e_q\}$$

where

$$\delta \in \mathcal{R}(c_{pqr})$$

and if $\delta \in \Delta$, where $\Delta \in \{\Delta_b, \Delta_t\}$, then

$$u_{qr}^{\Delta} - e_x < \gamma < u_{pq}^{\Delta} - e_c$$

whereas, if $\delta \in \Delta_f$ then either

$$u_{qr}^{\Delta_b} - e_x < \gamma - \delta + (90^o - \theta) < u_{pq}^{\Delta_b} - e_c$$

or

$$u_{qr}^{\Delta_t} - e_x < \gamma - \delta + (90^o + \theta) < u_{pq}^{\Delta_t} - e_c$$

depending on the ranges over which the divergent boundaries are defined. Although we know what range δ must lie in we do not know which e_i lie in the head of the rule unless we know which g_i range δ actually lies in. The g_i that δ could possibly lie in are those that satisfy

$$\bigcup_i g_i = \mathcal{R}(c_{pqr})$$

i.e. those g_i that constitute a covering of $\mathcal{R}(c_{pqr})$. If we choose a δ from each of these g_i then we can generate all the distinct changes in state reachable from the current state by varying δ, but using the same capture region. It is these g_i ranges that need to be stored in L. The $e_k \in E$ that appear in the head of the convergence rules for each of these choices can be found by inspecting the relevant E^{g_i} to see whether there are any possibly stable orientations lying between e_c and e_x on the P-S diagram for $\delta \in g_i$. If there is, then these would lie in the same capture region and be 'converged' onto e_q along with e_c and e_x.

It is probably worth pointing out that if a state includes both e_c and e_x in its definition, then any other e_k in its definition that lie between e_c and e_x in the range $\delta \in g_i$ *can* be converged, but they don't *need* to be a part of the convergence. That is, if these other e_k are not in the state definition this does not mean that the rule in question is not applicable. This is because a rule of the form

$$\{e_c, e_d, e_e\} \overset{(\gamma,\delta)}{\longmapsto} \{e_q\} \tag{1}$$

can be considered to be constituted of the following set of separate rules

$$\begin{aligned}\{e_c\} &\overset{(\gamma,\delta)}{\longmapsto} \{e_q\}\\ \{e_d\} &\overset{(\gamma,\delta)}{\longmapsto} \{e_q\}\\ \{e_e\} &\overset{(\gamma,\delta)}{\longmapsto} \{e_q\}\end{aligned}$$

where the γ and δ apply over the same range of values as they do in (1). Each of these rules are valid individually, but the convergence rule captures the fact that they are valid simultaneously. As a consequence the rule

$$\{e_c, e_d\} \overset{(\gamma,\delta)}{\longmapsto} \{e_q\}$$

is redundant, providing that the ranges of γ and δ are the same as those in (1), since it is subsumed by (1). In fact, if H is the set of stable orientations in the head of a rule R, then any state description S, that is a member of the power set of H, can be converged by the rule R, and any rules other than R, whose heads are members of the power set of H, are redundant (providing that the ranges of γ and δ for which they are applicable are the same as for R).

The bounds on the angle of reorientation, γ, are derived from the following considerations: we know that after reorientation by γ we want the 'uppermost' possibly stable orientation, e_c, to be less than the uppermost bound of the capture region, *i.e.*

$$e_c + \gamma < u_{pq}$$

Also, after reorientation by γ, the 'lowermost' stable edge, e_x, must be greater than the lowermost bound of the capture region, *i.e.*

$$e_x + \gamma > u_{qr}$$

Collecting these conditions we have

$$u_{qr} - e_x < \gamma < u_{pq} - e_c$$

This applies providing the divergent boundaries are horizontal. In the range where the divergent boundaries run diagonally across the P-S diagram their values can be related to their horizontal segment values by

$$u_{qr}^{\Delta_f} = u_{qr}^{\Delta_b} + \delta - (90^o - \theta)$$

or, since $u_{qr}^{\Delta_t} = u_{qr}^{\Delta_b} + 2\theta$,

$$u_{qr}^{\Delta_f} = u_{qr}^{\Delta_t} + \delta - (90^o + \theta)$$

This leads to the second set of constraints for γ given above, where we add what is effectively an offset that is a function of δ to γ, in order to remain in the same capture region.

Let $U \in [0^o, 360^o[$ be the range of possible orientations of the part with respect to the fence.

Given this notation the possible first sweeping motions can be represented by

$$U \overset{(0,\delta\in g_i)}{\longmapsto} S_1^{\langle(0,g_i)\rangle} \quad \forall g_i$$

If there are m different elements in G then we have m possible nodes at a depth of $d = 1$. Note that the first sweeping motion is a special case since there is no constraint on the part's possible orientation at the start of the sweeping motion. As a consequence of our definitions we know that after the first sweeping motion

$$S_1^{\langle(0,g_i)\rangle} = E^{g_i}$$

except that the ordering is not important. We know that for subsequent sweeping motions the part must start in one of a set of possible stable orientations. From here we apply the convergence rules to the members of a particular state S_d^L to reduce the 'uncertainty' in the orientation of a part by reducing the number of possible orientations (*cf.* [5]).

It is possible that the initial orientation of the part is such that it lies on a divergent boundary of the P-S diagram, and, in theory, it could stay in this orientation no matter how long the initial sweep applied is. It is possible to develop a simple strategy to guarantee avoiding this problem [4]. However, in practice, this was not found necessary [9]: imperfections in the horizontal surface across which the part is being pushed 'perturb' the part so that it rotates in one direction or the other.

Let F be any set of elements chosen from the current state, *i.e.* $F \subseteq S_d^L$. Let H_i be the set of stable edges that are in the head of rule i. In order for rule i to be applicable to take us from the current state to the next we must have $F \subseteq H_i$. However, for each application of a convergence rule to a state S_d^L, if $F \subset S_d^L$, then there must be a $\bar{F} = S_d^L \backslash F$. We will require that $\bar{F} \cap H_i = \emptyset$, so that $\bar{F}$ is the set of all the possible stable orientations that do not appear in the head of rule i, but which are elements of S_d^L. In which case subsequent states are found from

$$F \overset{(\gamma,\delta)}{\longmapsto} C$$
$$\bar{F} \overset{(\gamma,\delta)}{\longmapsto} D$$
$$S_{d+1}^L = C \cup D$$

where the value of C follows from the definition of the convergence rule, and D can be found from the P-S diagram. In order to generate the whole tree we apply all rules that satisfy these conditions. There is no point, however, applying a rule to a state if it results in transforming only one of the stable orientations, since this is equivalent to a pure rotation of the sweeper. As a consequence the cardinality of subsequent states consistently decreases.

In the case where all the elements in a state description are in the head of a rule then we have $H_i = F = S_d^L$ implying that $\bar{F} = \emptyset$, which results in $D = \emptyset$, which in turn implies $S_{d+1}^L = C$, giving

$$S_d^L \overset{(\gamma,\delta)}{\longmapsto} S_{d+1}^{L'}$$

as we would expect, with L being updated to L' to include the ordered pair of the permissible ranges of γ and δ to implement this transition.

There is a further complication however. The elements of $\bar{F}$ may map onto different stable edges depending on what value of γ in the allowed range is chosen. These different outcomes are different states in the tree of possible moves, and if we are to guarantee generating the complete tree we need to take this possibility into account.

Consider the case when $|\bar{F}| = 1$. Assume that γ needs to be chosen from $[\gamma_{min}, \gamma_{max}]$, and let $\bar{F} = \{e_p\}$. Say that reorienting by γ_{max} maps e_p onto e_x, and reorienting by γ_{min} maps e_p onto e_c. We need to check E^g,

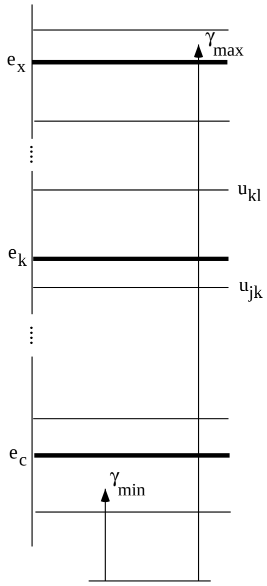

Figure 5: Possible stable reorientations for an element of $\bar{F}$

where g is the range from which δ is to be chosen, in order to find out if there are any stable orientations between e_c and e_x that e_p would map onto given an appropriate choice of γ in the range $[\gamma_{min}, \gamma_{max}]$ (see figure 5).

Let the set of possible stable orientations that e_p could be mapped onto be $P^g \subseteq E^g$, where the ordering on E^g is retained by P^g. For each element of P^g, for e_p to be mapped onto that element, *e.g.* e_k, we must have

$$u^g_{jk} < e_p + \gamma < u^g_{kl}$$

where u_{jk} and u_{kl} are the lower and upper divergent boundaries for e_k respectively. From this we can find the ranges within the range $[\gamma_{min}, \gamma_{max}]$ that lead to all the different mappings of e_k in P^g.

When $|\bar{F}| > 1$ we need to perform the above calculations for each member of $\bar{F}$. Each member of $\bar{F}$ will have its own ranges of γ that lead to different outcomes when the object is pushed with a fence orientation $\delta \in g$. We now have the problem of how these ranges 'interact'. That is to say, that if we use each range of γ generated for each element of $\bar{F}$ in a re-write rule we may get state mappings repeated. All we actually want are distinct state mappings.

Let e_p and e_q be two elements of $\bar{F}$, and let $\gamma_{p,1}, \gamma_{p,2}, \ldots, \gamma_{p,e}$ be the ranges of γ which lie within the range $[\gamma_{min}, \gamma_{max}]$ that lead to different outcomes for e_p, whilst $\gamma_{q,1}, \gamma_{q,2}, \ldots, \gamma_{q,f}$ are the ranges of γ lying in the range $[\gamma_{min}, \gamma_{max}]$ which give different outcomes for e_q. If we form

$$\gamma_{p,i} \cap \gamma_{q,j} \quad \forall i,\ \forall j$$

then the resulting ranges give all distinct combinations of e_p and e_q mappings. It is these ranges that need to be stored in L. This extends to an arbitrary number of elements in $\bar{F}$.

For any range of γ giving a distinct state transformation we can choose to use the middle of that range so as to be most tolerant to uncertainty. Similarly for δ.

4 Implementation

We tested the above analysis on the parts of the SOMA-4 world. See figure 6. SOMA cube parts are made by joining cubes together according to the following rule: each individual cube, called a cubie, can be joined by one or more face to other cubies (of the same size), such that the overall resulting shape is irregular, *i.e.* it includes a concavity. Parts in the SOMA-4 world are constructed using four or less cubies.

An offline planner has been written in SICStus PROLOG and runs on an 8 Mbyte Sun3/160. The planner uses three parameters: the part's name, the cubie size, and the friction cone half-angle. The part's name is used to extract geometric information about the part from a database. From this information the part's centre of mass and convex hull are calculated, and the P-S diagram for the given coefficient of friction is generated and may be viewed using GNUPLOT. The convergence rules are then derived and used to generate the graph of states in a breadth first manner until such time as a state description includes only one possible stable state or no further reduction is possible. The planner takes on average $3\frac{1}{2}$ minutes to arrive at a plan for a SOMA-4 part. The fact that the plans are generated off-line means that the time taken to find a plan isn't really an issue. The amount of memory required was, however, so when a node was generated it was checked to see whether it had been generated before by some other route in order to avoid duplicate branches. In fact we experienced a "Memory Allocation Failure" for part *fork2*, because there are many paths leading to a solution for this part. For this object a depth first search was used, though from inspection of its P-S diagram it seems likely that an admissible heuristic could be derived and used to guide the search of the iterative-deepening-A* algorithm [8], resulting in a significant reduction in the time and space required, whilst still finding an optimal solution.

Initially feasible plans are checked to see that the motions involved do not require the robot moving beyond its limits. The final output of the planner is a file containing parameter initialisations and calls to a sequence of VAL-II procedures. This information is transferred to the controller of a five degree of freedom Adept 1 robot via an RS-232 line. The initialisation process instantiates parameters such as the size of the parts and

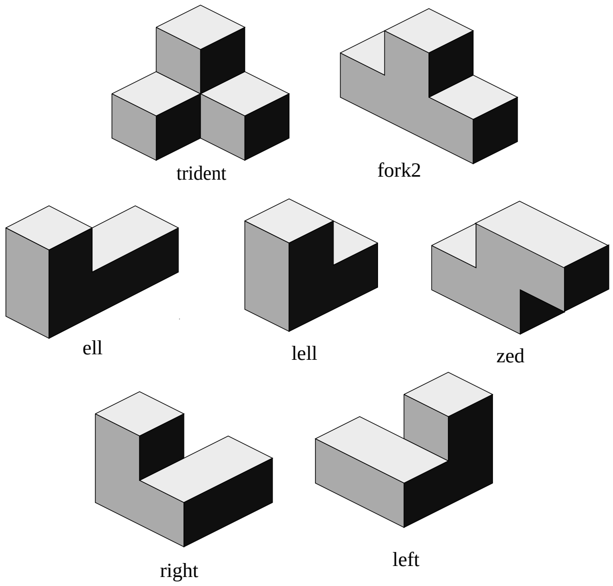

Figure 6: Objects in the SOMA-4 set.

the length of sweeping motion to be used. The VAL-II procedures are parameterised and implement such functions as simple arm motions and offset calculations.

Planned sequences of pushing motions for all the SOMA-4 parts have been found and successfully implemented. This was achieved without adopting Brost's strategy of shrinking the regions on the P-S diagram [2], [3]. The system's tolerance to spatial uncertainty during the execution of a plan is dependent on the size of the regions on the P-S diagram and where in these regions the trajectory describing the part's motion falls. By perturbing parts in various situations it was found that the tolerable errors in both part orientation and the line of pushing were at most $\pm 2^o$ away from the values predicted by the P-S diagram.

Because the larger the parts used the longer the length of push required to guarantee that the part will come to rest on a convergent boundary of the P-S diagram, and due to the constraint of the workspace envelope of the Adept 1 robot, SOMA-4 parts made from 1cm. plastic cubies were used (*cf.* [1]). For parts of this size we found that a pushing distance of 16cm. was reasonable. Although this length of sweep didn't always leave a part resting stably against the fence, usually the next sweeping motion was able to tolerate the error involved. Overall this sweep length meant that plans were successful 97% of the time.

The maximum convex hull perimeter of the SOMA-4 parts that we used is 10cm. We used a fence of length 40cm, which, given that the part was initially placed close to the fence and near its middle, proved to be acceptable.

5 Comparison to Related Work

Mani and Wilson [10] also tried to generate a sequence of sweeping motions to put a part into a known orientation using what they called an edge stability map (ESM). The main difference between their ESM and the P-S diagram is that the ESM doesn't explicitly include the stable orientations on the diagram, that is to say, they didn't include the horizontal e_i lines. Instead they just drew the divergent (or unstable) boundaries and labelled the resulting regions according to the stable orientations that they led to.

Mani and Wilson base their method on the following observation: assume that the part is in one of two orientations, e_1 and e_2, then by shifting the ESM by e_1 and e_2 and overlaying the results, we arrive at a diagram that gives the outcomes of any subsequent sweeping operations in terms of the labels of the overlaid areas. That is to say, that the new diagram is divided up according to the superposition of the areas of the two shifted diagrams, and the possible outcomes of these areas are the labels of the areas from the e_1 and e_2 shifted diagrams that gave rise to them. If both labels for an area on this new diagram are the same, then a sweeping motion with parameters from this region will result in both possible initial orientations converging on the one final orientation given by the labels. In principle this technique can be applied to any number of possible initial orientations, but the greater the number of possible initial orientations, the less the likelihood of finding one resultant region that will transform them all into one final orientation. In this case a search is needed to try and find a suitable sequence of sweeping motions. To guide this search Mani and Wilson present a set of heuristics. Because we explicitly encode the manipulation funnels in the P-S diagram we can explicitly explore the possibilities for manipulation that they allow.

The P-S diagram is also used in some recent work by Akella and Mason [1] in order to plan a sequence of sweeping motions that take a part from a known initial position and orientation into a desired position and orientation. The translational and rotational components are separated, but since the object starts off in a known orientation no search tree is required to put the part into another orientation.

In [5] Erdmann and Mason investigate the task of orienting objects on a walled tray by tilting the tray in various directions. This allows the sliding object three degrees of freedom, so the friction cone is extended to cover the object's rotational degree of freedom. This work also constructs a tree of possible states where a state is defined in terms of a set of possible orientations of the part being manipulated. However, for our case $|S_{d+1}^{L'}| < |S_d^L|$, whereas for Erdmann and Mason the cardinality of their states may increase.

Peshkin and Sanderson [14] have investigated orienting a part using a conveyer belt with fences protruding onto it at various angles. This is analogous to pushing an (initially) stationary object with a moving fence, except that they allow their object to roll off the end of the fences. This introduces an additional complication that Brost's assumption of an infinitely long fence avoids: they can only guarantee the part coming to rest within some range of orientations.

Peshkin and Sanderson construct what they call configuration maps, which are graphical representations of the final orientation of an object (to within some bounded range), given the initial orientation relative to the fence. With these maps it is possible to construct a state transition tree describing the effects of successive operation of the fences on the part. However, the growth of their tree is not guided, as ours is, by encoding the operation of manipulation funnels. Rather there is a set of possible operations that can be applied, one for each configuration map, and they are applied in turn to the current state to find the ensuing next state (although some pruning of the tree is possible).

The work of Peshkin and Sanderson is based on the calculation of the locus of all possible centres of rotation of a pushed object [15]. From this they calculate worst case bounds on the length of push that needs to be used to guarantee an object coming into a stable orientation. An alternative is to find out empirically what length of sweep is needed. This was the approach that we took and it turned out to be a reasonable alternative since Akella and Mason [1] report that the "Peshkin distance" needed to be increased by 15% to be sure of the desired effect.

There has also been some work related to ours by Goldberg and co-workers who have been interested in orienting a part using a parallel jaw gripper. If it is assumed that both jaws touch the object at the same time there occurs a pure squeezing operation [2], [3]. By removing the friction between one of the gripper jaws and the object it is possible to deduce the reorienting behaviour of the object from a diameter function [6], [7], [16]. This can be derived by plotting, for all orientations of the object, the distance that the parallel jaws of a gripper would be apart if they were just touching the object. The resulting plot is piecewise sinusoidal where the maxima are analogous to divergent boundaries of the P-S diagram, and the minima are analogous to convergent boundaries. In fact the P-S diagram can be shown to be related to a radius function, which is also piecewise sinusoidal, and describes the locus of the COF of an object as it is rolled along a flat surface [6], [11].

The diameter function can be used to derive a step function, called the squeeze function, that represents the fact that the continuum of initial object orientations between local maxima on the diameter function lead to the same final orientation when the object is squeezed (assuming a frictionless gripper). This step function can be used to generate a sequence of squeezes that will orient an object (up to symmetry)[2]. The continuum of initial orientations that will lead to some particular stable orientation are considered to be the pre-image of that orientation. The stable orientation with the largest pre-image is chosen as the starting point. This pre-image then becomes the goal that the squeeze function has to map other stable orientations onto. The mapping that possesses the largest pre-image is chosen and the process repeated. Each mapping is physically achieved by re-orienting the gripper by the appropriate amount before squeezing, and is exploiting a manipulation funnel.

Since both the diameter function of an object and its radius function can be mapped on to a function that is a monotonically non-decreasing step function on the space of planar rotations[3], Goldberg has also used his algorithm to derive plans that initially push the object before grasping it [6]. In the situations where the object is pushed until it comes to rest in a stable configuration, ours and Goldberg's algorithm are almost equivalent. One difference is the fact that Goldberg restricts the motion of his gripper to be perpendicular to the jaws. This is equivalent to working with the slice of the P-S diagram that lies along $\delta = 90^o$. Consequently his algorithm can only exploit the manipulation funnels that lie along this line. By being able to use the whole of the P-S diagram we can capitalise on any larger convergence regions that lie on either side of this line. Another difference is that because we perform an exhaustive search our plans need not end on any particular stable orientation.

6 Conclusion

In this paper we have presented a way of exploiting what Mason has called a manipulation funnel. The technique is based on Brost's push-stability diagram, which encodes the qualitative behaviour of a pushed object sliding on a horizontal planar surface. From this diagram we have suggested constructing the tree of state transitions, where a state is a list of possibly stable orientations. If we construct the whole tree it can be used to tell us:–

- whether it is possible to uniquely orient the object by pushing alone (for a particular value of the coefficient of friction between fence and object),
- the minimum number, and parameters, of the sweeping motions needed to put a part into a

[2]Using a parallel jaw gripper actually introduces an additional symmetry into the problem.

[3]Recent work has derived plans for objects that are made up of piecewise linear and circular edges by relaxing the need for the squeeze function to be a step function [16].

known orientation, and

- the minimum number, and parameters, of the sweeping motions needed to put a part into some particular desired orientation.

If we are just trying to put a part into an arbitrary known orientation then a breadth-first search suffices.

There are some limitations to our approach. Firstly, parts with symmetry will never lead to states with only one possible stable orientation. We can turn this to our advantage by reducing the P-S diagram so that it doesn't repeat itself, thus reducing the search space. For example, we need only work with either the top or bottom half of figure 3. (Mani and Wilson make an equivalent observation [10].)

Secondly, we have only managed to orient the part. Its position is unknown other than the fact that it must be somewhere along the line of the fence after the last sweeping motion. If the part possesses a stable edge that lies orthogonal to this line we could apply another sweeping motion that contacts this edge in order to leave the part in a known position and orientation. Alternatively we could arrange that the object slide down the fence to a projecting stop [11].

Thirdly, for some objects you need to know which way round they are about a horizontal axis. Erdmann and Mason [5] note this as well, whilst trying to orient an Allen key: whether the object starts off in a right or left-handed configuration is important because switching between the two is not possible using only planar motion.

Finally we note that this approach finds the minimum number of sweeping motions for a given P-S diagram. It is possible to change the number of motions required by changing the friction between the pushing and pushed objects, thus changing the P-S diagram, and therefore the search space for the task. It was found that with a high enough value of the coefficient of friction all of the members of the SOMA-4 set of parts could be uniquely oriented in one sweeping motion [9].

Acknowledgements

We would like to thank Andrew Fitzgibbon for making it possible to submit this paper, an anonymous reviewer for referring us to the work by Goldberg et al., and to Ken Goldberg for providing us with a copy of [6]. GED gratefully acknowledges the support of British Aerospace Plc.

References

[1] S. Akella and M.T. Mason, "Posing Polygonal Objects in the Plane by Pushing", *Proc. IEEE Int. Conf. Robotics and Automation*, Nice, France, May 1992, pp.2255–2262.

[2] R.C. Brost, *Planning Robot Grasping Motions in the Presence of Uncertainty*, Technical Report CMU-RI-TR-85-12, The Robotics Institute, Carnegie-Mellon University, Pittsburgh, PA, July 1985.

[3] R.C. Brost, "Automatic Grasp Planning in the Prescence of Uncertainty", *Int. J. Robotics Research*, Vol.7, No.1, 1988, pp.3–17.

[4] G.E. Deacon, *An Algorithm for Orienting Objects in a Minimum Number of Robot Sweeping Motions*, (Unpublished) Working Paper, Dept. of Artificial Intelligence, University of Edinburgh, June 1992.

[5] M.A. Erdmann and M.T. Mason, "An Exploration of Sensorless Manipulation", *IEEE J. Robotics and Automation*, Vol.4, No.4, 1988, pp.369–379.

[6] K.Y. Goldberg, "Orienting Polygonal Parts Without Sensors", *Algorithmica*, to appear 1993.

[7] K.Y. Goldberg, M.T. Mason, and M.A. Erdmann, "Generating Stochastic Plans for a Programmable Parts Feeder", *Proc. IEEE Int. Conf. Robotics and Automation*, Sacramento, CA, April 1991, pp.352–359.

[8] R.E. Korf, "Search: A Survey of Recent Results", in *Exploring Artificial Intelligence*, edited by H.E. Shrobe & American Association for Artificial Intelligence, Morgan Kaufmann Publishers Inc., San Mateo, CA, 1988, pp.197–237.

[9] P.L. Low, *Sensorless* SOMA *Part Orientation*, M.Sc. Dissertation, Dept. of Artificial Intelligence, University of Edinburgh, Sept. 1992.

[10] M. Mani and W.R.D. Wilson, "A Programmable Orienting System for Flat Parts", *Proc. North American Manufacturing Research Institute Conf. XIII*, Berkeley, CA, May, 1985, pp.427–432.

[11] M.T. Mason, *Manipulator Grasping and Pushing Operations*, Ph.D. Thesis, Dept. of Electrical Engineering and Computer Science, MIT, June 1982. Published in: M.T. Mason and J.K. Salisbury, *Robot Hands and the Mechanics of Manipulation*, MIT Press, Cambridge, MA, 1985.

[12] M.T. Mason, "The Mechanics of Manipulation", *Proc. IEEE Int. Conf. Robotics and Automation*, St. Louis, MO, March 1985, pp.544–548.

[13] M.T. Mason, "Mechanics and Planning of Manipulator Pushing Operations", *Int. J. Robotics Research*, Vol.5, No.3, 1986, pp.53–71.

[14] M.A. Peshkin and A.C. Sanderson, "Planning Robotic Manipulation Strategies for Workpieces that Slide", *IEEE J. Robotics and Automation*, Vol.4, No.5, 1988, pp.524–531.

[15] M.A. Peshkin and A.C. Sanderson, "The Motion of a Pushed, Sliding Workpiece", *IEEE J. Robotics and Automation*, Vol. 4, No. 6, 1988, pp.569–598.

[16] A.S. Rao and K.Y. Goldberg, "Orienting Generalized Polygonal Parts", *Proc. IEEE Int. Conf. Robotics and Automation*, Nice, France, May 1992, pp.2263–2268.