Mountain Climbing, Ladder Moving, and the Ring-Width of a Polygon

JACOB E. GOODMAN*, *City College, City University of New York*
JÁNOS PACH**, *Hungarian Academy of Sciences*
CHEE K. YAP***, *Courant Institute, New York University*

JACOB ELI GOODMAN received his Ph.D. at Columbia University in 1967, under the direction of Heisuke Hironaka. In addition to algebraic geometry, he has worked in topological graph theory and discrete geometry, and has coauthored a series of papers with Richard Pollack focusing on the interplay between geometric and topological properties of configurations of points and arrangements of hyperplanes and their generalizations.

He is coeditor in chief of *Discrete & Computational Geometry*, a journal of mathematics and computer science published by Springer-Verlag.

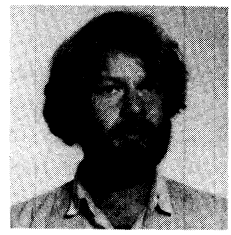

JÁNOS PACH received his Ph.D. at Eötvös University, Budapest, in 1980, under the direction of Miklós Simonovits. He is a senior research fellow at the Mathematical Institute of the Hungarian Academy of Sciences, working mainly in combinatorics, discrete geometry, and computational geometry. He is a frequent visitor at the Courant Institute, New York University.

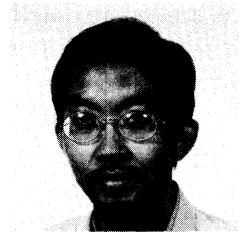

CHEE-KENG YAP was born in Singapore in 1952, grew up in Malaysia, and attended MIT (B.S., 1975) and Yale (Ph.D., 1980). Since 1981, he has been at the Computer Science Department of the Courant Institute of Mathematical Sciences, New York University. He is an editor of *S.I.A.M. Journal of Computing* and the *Journal of Computer and System Sciences*. His research interests are mainly (ultimately) algorithmic, whether in the guise of complexity theory, symbolic and algebraic computation, computational geometry, or robotic 'applications' as in this article.

Introduction. In this article we consider several geometric problems which at first glance seem only vaguely connected, but which on closer examination turn out to be intimately related. The first is well known in mathematical folklore (but the proof presented here seems to be new), the second is an amusing variant of the first, and the third is a motion-planning problem arising in the relatively new mathematical discipline of robotics; in fact it was this third problem, suggested in [**8**], which motivated our investigation. (See [7] for a general survey of the area of motion planning.)

*Supported in part by NSF grant DMS-8501492 and PSC-CUNY grant 666426
**Supported in part by Courant Institute (NYU) and Hungarian NSFR grant 1812
***Supported in part by NSF grants DCR-8401898 and DCR-8401633

(I) Two mountain climbers begin at sea level, at opposite ends of a (two-dimensional) chain of mountains. Can they find routes along which to travel, always maintaining equal altitudes, until eventually they meet?

(II) Two painters are carrying a 20-foot ladder, one at each end, along a garden path which begins and ends with long straight segments. Can they negotiate the path without violating the "keep off the grass" signs? Moreover, what happens if the path splits into two paths which subsequently reunite: can the painters separate and later rejoin each other?

(III) If an irregularly shaped body can be dragged completely (i.e., continuously translated and rotated) through two posts a distance d apart, can it be dragged completely through two posts separated by a distance greater than d?

Each of these problems seems to have an intuitively clear solution which, on closer examination, proves decidedly elusive; this is true especially of the second and third. However, by the time we are done, we will have seen not only that the problems are interconnected in that the solutions to the first two are needed for the solution to the third, but that the most important common feature of all of the problems is that once we translate them into the language of elementary plane topology, their solutions become natural and transparent.

First we present a simple topological lemma which lies at the heart of the proofs.

LEMMA 1. *Given a square S with consecutive vertices A, B, C, D, and a graph Γ (by which we mean here a finite system of [not necessarily simple] closed arcs and intersection points) lying in S and containing A and C and no other point of the boundary of S. If A and C are not in the same connected component of Γ then B and D can be joined by a path lying in S which avoids Γ.*

Proof. Replace Γ by a tubular neighborhood $\bar{\Gamma}$ of sufficiently small radius ε. Then, starting at A, run along the boundary of $\bar{\Gamma}$. (The path we follow is well defined, since—even though Γ may branch, as in Fig. 1—the boundary of $\bar{\Gamma}$ does not.) Eventually we must either return to A or arrive at C, since the total number of arcs is finite and we cannot repeat a previous position because no branching takes place. In the first case the path itself (together with part of the boundary of S) constitutes a path from B to D which avoids Γ, and in the second case the corresponding arcs of Γ constitute a path from A to C. (See Fig. 1 for illustrations of both possibilities.) □

(Notice that Lemma 1 can be thought of as a kind of weak converse to the Jordan Curve Theorem: From the Jordan theorem it follows immediately that if there is a path joining a pair of opposite corners of a square which lies entirely within the square, then any path joining the two remaining corners and lying in the square must cross the first; Lemma 1 asserts that if there is no path which can be pieced together from a given finite collection of arcs to connect a pair of opposite corners, then the remaining two corners can be connected by a path not touching any of the arcs.)

We shall begin the discussion of each of the three problems by giving it a precise formulation, and in some cases discuss the consequences of choosing our formulations with less care. The article concludes with the derivation of an algorithm which determines the minimum separation d for which the condition in Problem III holds. We thank Walter Daum for providing the example illustrated in Fig. 9, and we note,

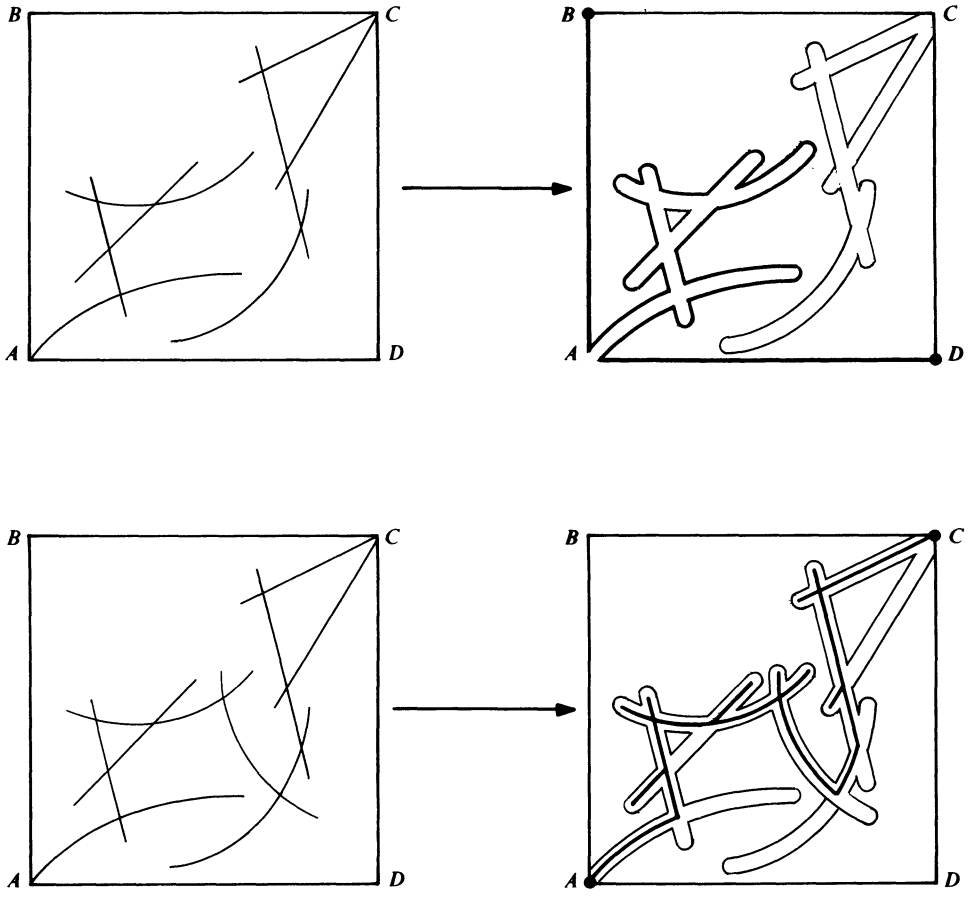

FIG. 1

in addition, that Theorem 2 and some weaker versions of Theorem 4 have also been proved independently by Fan Chung and Paul Seymour [2], David Eppstein [3], and John Kutcher and Joseph O'Rourke [4], respectively.

1. The mountain climbers' problem. Our assumption will be that the silhouette of the mountain chain is piecewise monotone. As a matter of fact, we can then also suppose it to be piecewise linear, by replacing each maximal monotone portion by the line segment joining its endpoints. If we now select a point of maximum altitude and reparametrize, it becomes sufficient to prove the following theorem:

THEOREM 1. *Let $s(\tau)$ and $s'(\tau)$ be continuous, piecewise linear functions from* $[0,1]$ *to* $[0,1]$, *with* $s(0) = s'(0) = 0$ *and* $s(1) = s'(1) = 1$. [*Note: the "prime" does not represent differentiation here*!] *Then there exist two other continuous, piecewise linear functions*

$$\tau(x), \tau'(x) : [0,1] \to [0,1]$$

such that

$$\begin{aligned} \tau(0) &= \tau'(0) = 0, \\ \tau(1) &= \tau'(1) = 1, \\ s(\tau(x)) &= s'(\tau'(x)) \quad \textit{for every } x \in [0,1]. \end{aligned} \tag{1}$$

Moreover, τ and τ' have a total of $O(n^2)$ critical points, where n denotes the total number of critical points of s and s'.

(In other words, if A and B are the mountain climbers, we have rephrased the problem so that they each start at time 0 and height 0 and want to wind up at time 1 and height 1, maintaining the same height as they move. A and B can backtrack, of course, and their motion is represented by the functions $\tau(x)$ and $\tau'(x)$ which we seek.)

Proof. The essence of the argument is to consider the set Γ of pairs of positions which are at the same height, and to show that we can get from the bottom (both at height 0) to the top (both at height 1) without ever leaving Γ. Hence let Γ be the set of pairs (τ, τ') for which $s(\tau) = s'(\tau')$. Then both $(0,0)$ and $(1,1)$ are contained in Γ, and we must show that they belong to the same arcwise connected component. There are two ways to proceed. The first, which is simpler, uses Lemma 1 after first observing that nothing changes if we assume that the only points on the silhouettes at height 0 or 1 are the beginning and the end. The simple proof based on these hints is left to the reader. However, in order to gain greater insight into the situation, we will instead do a little extra work, and show that Γ is the union of a finite number of isolated points and 1-manifolds with boundary and is made up of at most n^2 line segments, and that in fact the only boundary points arising are $(0,0)$ and $(1,1)$ themselves.

To this end, we examine a pair of points $(\tau, s(\tau))$ and $(\tau', s'(\tau'))$ with $s(\tau) = s'(\tau')$, i.e., a pair of positions of A and B which are at the same height. There are several cases, each illustrated in Fig. 2, where arrows with the same number of arrowheads represent corresponding motions of A and B and of the point of Γ which represents their joint position.

(a) If A and B are each at interior points of maximal linear segments of s and s', respectively, then they can each move in two directions (either both forward or both backward if the segments have the same monotonicity, or one forward and the other backward if the segments have opposite monotonicity). Hence in the neighborhood of the corresponding point of Γ, Γ is a 1-manifold (in fact a line segment!) (Fig. 2a).

(b) If one of A and B, say A, is at a local maximum (or minimum) of s, and the other is at an interior point of a segment of s', then A can move in two directions and B in one, so again Γ is a 1-manifold in a neighborhood of the corresponding point (Fig. 2b).

(c) If A and B are each at local maxima or each at local minima of s (resp. s'), then there are precisely four ways for them to proceed maintaining the same height: A and B can each move either forward or backward. Thus the corresponding point of Γ looks like a simple crossing (Fig. 2c).

(d) If A is at a local maximum of s and B at a local minimum, or vice versa, there is no way for them to proceed (or, for that matter, to have gotten there in the first place!), so the corresponding point of Γ is isolated (Fig. 2d).

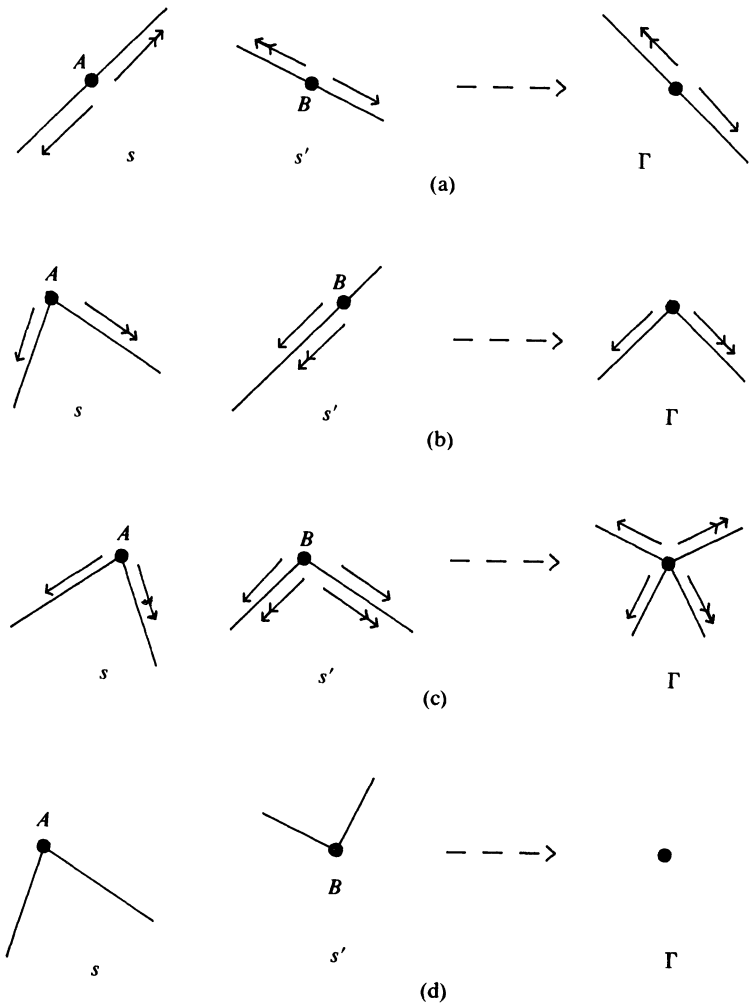

FIG. 2

Now suppose A and B are moving along s and s', respectively. Then they continue their motion until one of them hits a critical point. Since the other is at the same height, and since each of s and s' has $O(n)$ critical points (hence $O(n)$ segments), this can happen in at most $O(n^2)$ ways. The only situations in which there is only one way for A and B to move are when both are at $(0,0)$ or both at $(1,1)$; hence the corresponding points of Γ are the only boundary points. Thus Γ is the union of $O(n^2)$ line segments and isolated points, and (throwing away the isolated points), each vertex of Γ has even degree except for $(0,0)$ and $(1,1)$, which means that these must lie in the same connected component of Γ. □

Notice that if we allowed the mountain silhouette to dip below sea level, the conclusion would no longer be valid, as is shown by Fig. 3.

We mention also, for completeness, that there is yet another proof, which proceeds by induction on the number of monotone segments in the silhouette. We

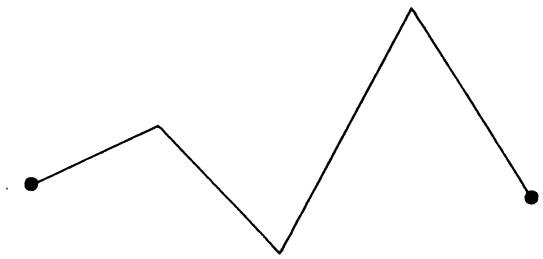

FIG. 3

have chosen to disregard this proof, however, since we wish to focus on a topological approach that applies to all three problems.

2. The ladder movers' problem. To answer the first part of Problem II of the Introduction, we prove:

THEOREM 2. *Suppose* $x = (x_1(s), x_2(s))$, $0 \leqslant s \leqslant L$, $L > 1$ *is a polygonal curve, parametrized by its arc length, whose initial and final segments each have length at least* 1, *and which has the property that the only points within unit distance of* $x(0)$ *are the points* $x(s)$ *with* $0 \leqslant s \leqslant 1$, *and the only points within unit distance of* $x(L)$ *are the points* $x(s)$ *with* $L - 1 \leqslant s \leqslant L$. *Then a line segment of unit length can be moved continuously from* $\overline{x(0), x(1)}$ *to* $\overline{x(L-1), x(L)}$ *so that its endpoints A and B never leave the curve.*

Proof. This time we let Γ be the set of pairs (s, t) such that $|x(s) - x(t)| = 1$. A line segment $\overline{AB}$ of unit length makes four angles with the curve x: two at A (call them A_1 and A_2), and two of B (B_1 and B_2). Assume for the moment that no two segments of the polygonal path are parallel and exactly one unit apart. Again examining the several cases that may arise, we see that either A and B can move together in two directions (if exactly one of the A_i and exactly one of the B_i is acute, as in Fig. 4a, the general case), or one can move in two directions and the other in one (if exactly one or three of A_1, A_2, B_1, B_2 is acute, as in Fig. 4b), or each can move in either of two directions, giving a total of four possibilities (if both A_i are acute, and neither B_i is, or vice versa, as in Fig. 4c), or neither can move (if all of the A_i and B_i are acute, or none of them is, so that the situation is actually impossible to reach, as in Fig. 4d). There is one additional possibility: it may happen that two segments of the polygonal path x are parallel, and precisely one unit apart. In this case, if A (say) is at a vertex, there will be either one or three ways the motion can proceed: Fig. 4e illustrates the latter possibility. Thus the situation is very much as in the case of Theorem 1, although the curves which comprise Γ are now either arcs of (or complete) ellipses or line segments, but with the major difference that the degrees of some vertices of Γ may now be odd. But—except for the "starting" and "ending" vertices of Γ, i.e., the points $(0, 1)$ and $(L - 1, L)$—this happens only when two segments of x are parallel and distance 1 apart, so that the corresponding arc of Γ beginning at such a vertex of odd degree must be a line segment which also ends at a vertex of odd degree. Thus we can remove all such offending arcs to obtain a new graph Γ', which shares all the topological properties of the graph Γ in the proof of Theorem 1, and the conclusion follows. □

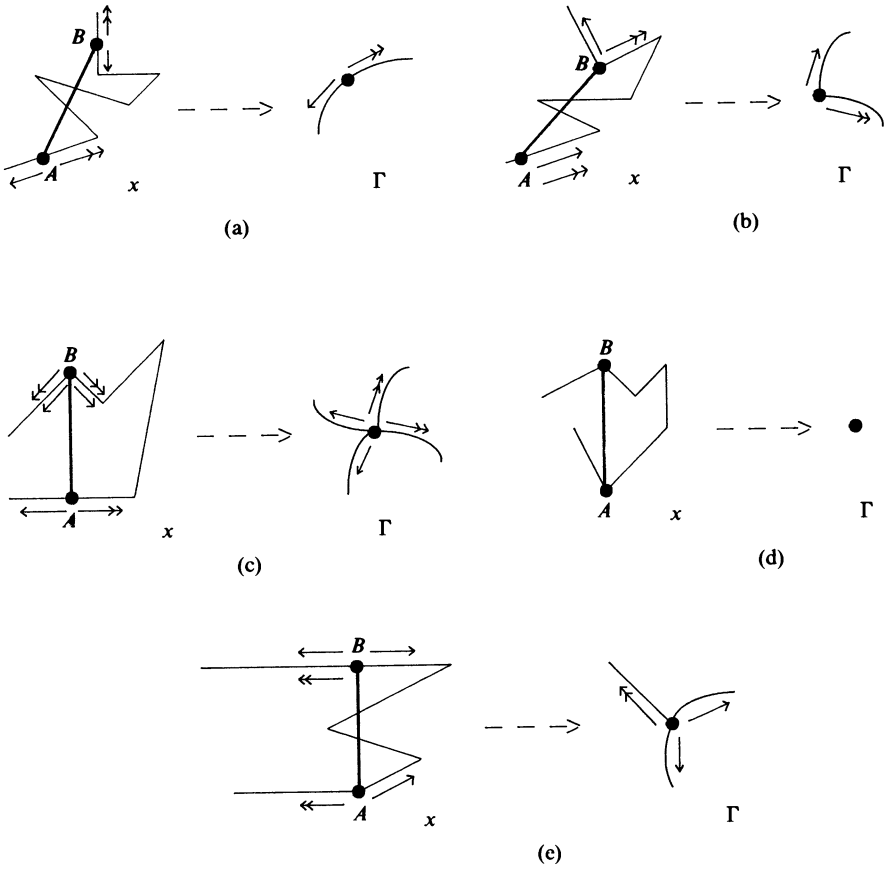

FIG. 4

The second part of Problem II will be addressed by Theorem 3, which may be paraphrased as follows: Suppose the two paths, which we again assume polygonal, diverge at point P and rejoin at point Q. If two people, one walking from P to Q along one branch, the other from Q to P along the other branch, cannot move in such a way that they stay more than 20 feet apart all the way, then the painters can indeed carry the 20-foot ladder along the two paths from P to Q.

THEOREM 3. *Let ρ_1, ρ_2 be two polygonal paths between points c_1 and c_2, i.e., piecewise linear arcs ρ_i: $[0,1] \to \mathbf{R}^2$ with $\rho_1(0) = \rho_2(0) = c_1$, $\rho_1(1) = \rho_2(1) = c_2$. Assume further that*

$$\begin{aligned}
&\rho_1\left(\tfrac{1}{3}\right) = \rho_2\left(\tfrac{1}{3}\right) = c_1^*, \rho_1\left(\tfrac{2}{3}\right) = \rho_2\left(\tfrac{2}{3}\right) = c_2^* && \textit{for some } c_1^*, c_2^*,\\
&|c_1 - c_1^*| = |c_2 - c_2^*| = d && \textit{for some } d > 0,\\
&\rho_1(t) = \rho_2(t) && \textit{for every } t \in \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right],\\
&|c_1 - \rho_i(t)| > d \textit{ and } |c_2 - \rho_i(1-t)| > d && \textit{for every } t \in \left(\tfrac{1}{3}, 1\right], i = 1, 2.
\end{aligned} \tag{2}$$

Then at least one of the following statements is true:

(i) *Point A can move from* c_1 *to* c_2^* *along* ρ_1, *and point B can move from* c_1^* *to* c_2 *along* ρ_2, *so that their distance is always equal to* d.

(ii) *A can move from* c_1 *to* c_2 *along* ρ_1, *and B can move from* c_2 *to* c_1 *along* ρ_2, *so that their distance is always larger than* d.

Proof. (See Fig. 5.) Suppose (ii) does not hold. Let Γ represent the pairs of points a unit distance apart; this time

$$\Gamma = \{(s,t) \mid 0 \leqslant s, t \leqslant 1|, |\rho_1(s) - \rho_2(t)| = d\}.$$

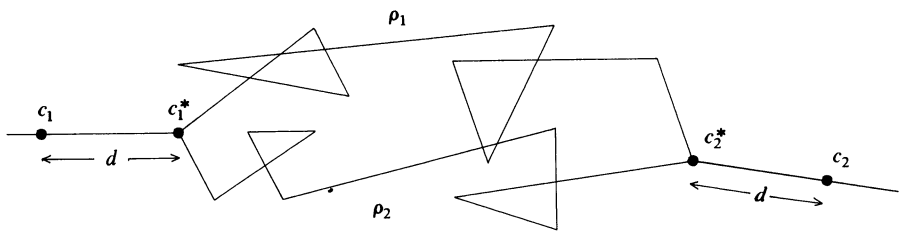

FIG. 5

Again, as in the proof of Theorem 2, Γ consists of only finitely many arcs of ellipses, line segments, and isolated points. This time, however, there are four points at which Γ meets the boundary of the square $0 \leqslant s \leqslant 1, 0 \leqslant t \leqslant 1$, as Fig. 6 shows. We must show that $(0, 1/3)$ can be joined to $(2/3, 1)$ along Γ.

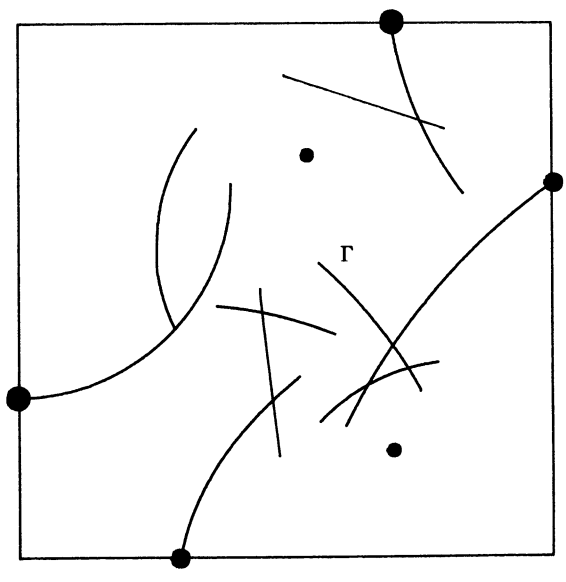

FIG. 6

Considering the possible critical points of the motion of A and B, we see that exactly the same cases arise as in the proof of Theorem 2 (it doesn't matter, locally, whether portions of the curves on which the points A and B move are on "parallel" tracks, or not). Hence, as in the proof of Theorem 2, after removing the parts of Γ arising from possible motion along parallel line segments (as well as any isolated points), we see that Γ contains no vertex of odd degree except for the four points on the boundary of the square. Hence the only cases arising are those shown (schematically) in Fig. 7. In case (a) it follows from Lemma 1 that there is a path from $(0,1)$ to $(1,0)$ which avoids Γ, and this means that alternative (ii) holds, contrary to our assumption. Thus we must be in either case (b) or case (c), but in each of these cases there is a path in Γ with the desired property. □

The condition that the curve in Theorem 2 is polygonal is not superfluous, as is shown by

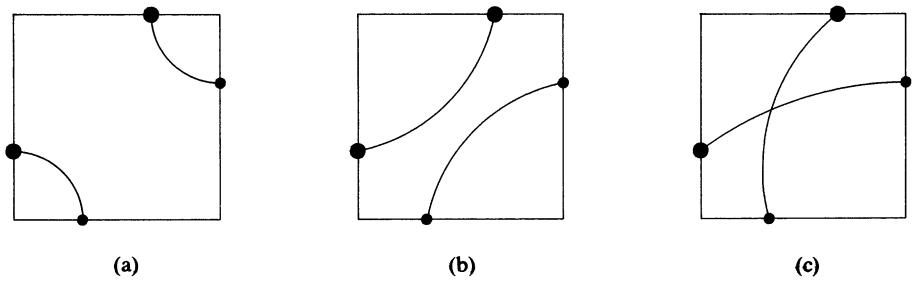

FIG. 7

PROPOSITION 1. *There is a simple (i.e., non-self-intersecting) continuous arc x^*, which starts and ends with nonparallel half lines h_1 and h_2, and which admits no continuous motion of a line segment of unit length from h_1 to h_2 along x^*.*

Proof. Let c_p, c_q, and c_r denote the unit circles around points p, q, and r, respectively, with $p, q \in c_r$. Let h_1 and h_2 be two half lines with endpoints p and r, respectively, and let x^* be defined as the union of h_1, h_2, the portion of c_r between p and q, and an arc connecting q and r, which meets c_p and c_q infinitely many times, as in Fig. 8.

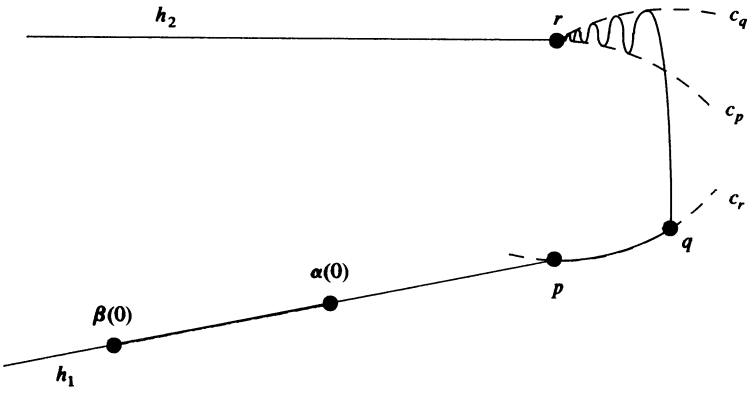

FIG. 8

Assume now, in order to obtain a contradiction, that there are continuous functions $\alpha(t), \beta(t) : [0,1] \to x^*$ such that $|\alpha(t) - \beta(t)| = 1$ for every $t \in [0,1]$ and such that $\alpha(0), \beta(0) \in h_1$ and $\alpha(1), \beta(1) \in h_2$, with $\alpha(0)$ between $\beta(0)$ and p. Let

$$t^* = \inf\{t \in [0,1] \mid \alpha(t) = r\}.$$

Then, given any $\varepsilon > 0$, $\alpha(t) \in c_p$ for infinitely many values $t \in (t^*, t^* - \varepsilon)$, and similarly $\alpha(t) \in c_q$ for infinitely many values $t \in (t^*, t^* - \varepsilon)$. But $\beta(t) = p$ whenever $\alpha(t) \in c_p$, and similarly $\beta(t) = q$ whenever $\alpha(t) \in c_q$; hence $\lim_{t \to t^*} \beta(t)$ cannot exist, contradicting our assumption. □

By a slight modification of this construction we can obtain a smooth (i.e., differentiable) simple curve which does not permit a continuous, rectifiable motion of the ladder.

Moreover, the assumptions about $x(0)$ and $x(L)$ in Theorem 2 are essential, as shown by the example illustrated in Fig. 9.

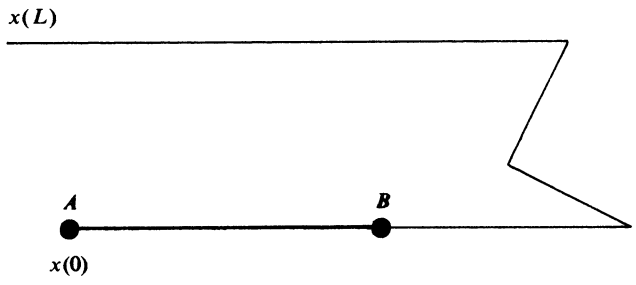

FIG. 9

3. The ring-width problem. Let P be a simple closed polygon in the Euclidean plane $\mathbf{R}^2$ with two handles (half lines), h_1 and h_2, sticking out at two points of its boundary. Assume that these handles do not point in the same direction, and that they do not meet each other or P at any other points (see Fig. 10).

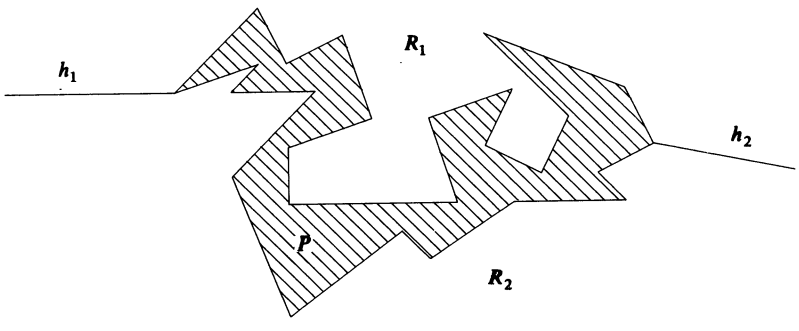

FIG. 10

With some distortion of language (which may be justified by the 3-dimensional analogue of the problem), the *ring-width* $w_R(P)$ of P is defined as the smallest

number w such that a "ring of diameter w" can be pulled from one handle of P to the other. To be precise, let R_1 and R_2 denote the closures of the two unbounded regions of $\mathbf{R}^2 \setminus (h_1 \cup P \cup h_2)$. Then $w_R(P)$ is the least w having the property that there exist two continuous arcs π_1: $[0,1] \to R_1$, π_2: $[0,1] \to R_2$ satisfying

$$|\pi_1(t) - \pi_2(t)| = w \quad \text{for every } t \in [0,1], \tag{3}$$

and such that

$$\overline{\pi_1(0), \pi_2(0)} \cap (h_1 \cup P \cup h_2) = \overline{\pi_1(0), \pi_2(0)} \cap h_1 \neq \varnothing$$

and

$$\overline{\pi_1(1), \pi_2(1)} \cap (h_1 \cup P \cup h_2) = \overline{\pi_1(1), \pi_2(1)} \cap h_2 \neq \varnothing .$$

The *elastic ring-width* of P, $w_{ER}(P)$, is defined similarly, the only difference being that the ring is now allowed to contract and expand up to size w; i.e., instead of (3) only

$$|\pi_1(t) - \pi_2(t)| \leqslant w \quad \text{for every } t \in [0,1] \tag{3'}$$

is required. Obviously $w_{ER}(P) \leqslant w_R(P)$. We will prove the following assertions, the first of which addresses Problem III of the Introduction:

THEOREM 4. *Any ring of diameter $d \geqslant w_R(P)$ can be pulled from one handle of P to the other.*

THEOREM 5. $w_R(P) = w_{ER}(P)$ *for any P.*

(Warning: It takes some time to realize that neither of the two statements above is trivial! The notions of ring-width and elastic ring-width were introduced in [**8**] and both of these results were conjectured, but for polygons without handles; that case still remains open.)

Let B_i denote the boundary of R_i $(i = 1,2)$. A pair of arcs (π_1, π_2) with the properties above, including (3) rather than (3′), is called *a contact motion of a ring* of diameter w along P if, in addition, π_i: $[0,1] \to B_i$ $(i = 1,2)$. *A contact motion of an elastic ring* is defined similarly, except that now condition (3) should be replaced by (3′).

THEOREM 6. *There is a contact motion of a ring of diameter d along P if and only if there is a contact motion of an elastic ring of diameter d along P.*

Theorems 4–6 are easy consequences of:

THEOREM 7. *Every polygon P with handles permits a contact motion of a ring of diameter d, for any $d \geqslant w_{ER}(P)$.*

Before we can prove Theorem 7, we need one more lemma, which can be paraphrased as follows: If A and B are walking from city c_1 to city c_2 along roads γ_1 and γ_2, respectively, and they hold a rope stretched between them, then at some moment they will necessarily catch C who wants to get from c_2 to c_1 at the same time within the region enclosed by γ_1 and γ_2.

LEMMA 2. *Let $\gamma : [0,1] \to \mathbf{R}^2$ be a simple arc connecting two points c_2 and c_1, with $\gamma(0) = c_2$ and $\gamma(1) = c_1$, and let h_1 and h_2 be two disjoint half lines whose endpoints are c_1 and c_2, respectively. Assume further that no interior point of γ is on h_i, and let*

R_1 and R_2 denote the closures of the connected components of $\mathbf{R}^2 \setminus (h_1 \cup h_2 \cup \gamma)$. Let $\gamma_i : [0,1] \to R_i$ $(i = 1,2)$ be (possibly self-intersecting) arcs with $\gamma_i(0) = c_1$, $\gamma_i(1) = c_2$.

Then there exists $x \in [0,1]$ such that $\gamma(x)$ lies on the line segment $\overline{\gamma_1(x), \gamma_2(x)}$.

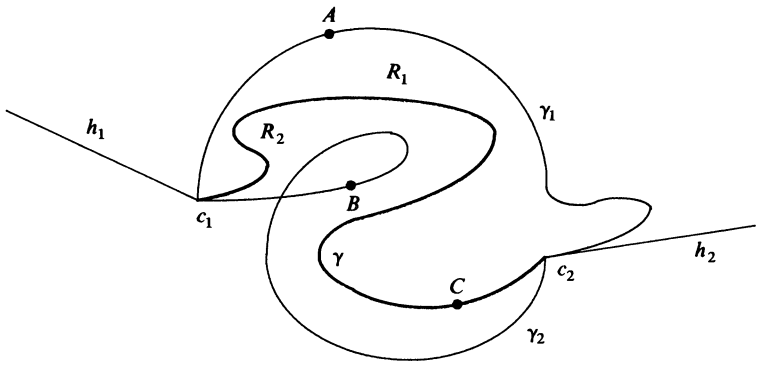

FIG. 11

Proof. (See Fig. 11.) For any $t \in [0,1]$, let us denote $\gamma_1(t)$, $\gamma_2(t)$, $\gamma(t)$ by A, B, C, respectively. If, during the motion, C ever coincides with either A or B, we are done. So suppose not. As soon as A, B, and C are under way, i.e., at time t with $0 < t < 1$, consider the directed arc consisting of h_1 (reversed), γ_1 from c_1 to A, the directed line segment $\overrightarrow{AB}$, γ_2 from B to c_2, and h_2. For t close to 0, point C is in the region "above" this directed path, while for t close to 1, C is "below" it (see Fig. 12). Since everything varies continuously with t, C must cross the path somewhere. But it never does so outside of the segment $\overline{AB}$, hence we are done.

(Notice that what we are really talking about is the winding number about C of the closed curve Γ consisting of the directed path described above plus an arc of a large circle going counterclockwise from a point far out on h_2 to a point far out on

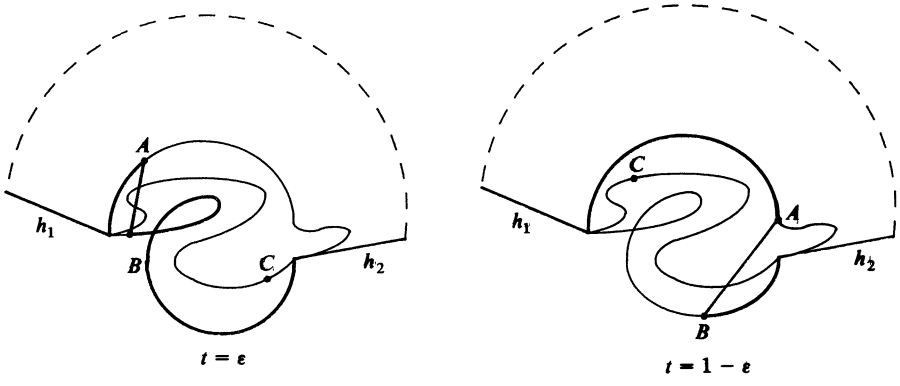

FIG. 12

h_1. Since this winding number,

$$\frac{1}{2\pi i}\int_\Gamma \frac{dz}{z-\gamma(t)},$$

as long as it is defined, varies continuously with t as t goes from ε to $1-\varepsilon$ and changes from 1 to 0, remaining locally constant (since it is an integer!), it must fail to be defined for some value of t between 0 and 1, which implies the result.)

Proof of Theorem 7. Let us fix $d \geqslant w_{ER}(P)$, and consider the motion of an elastic ring of diameter $w_{ER}(P) = w$ along P, i.e., a pair of continuous maps π_1: $[0,1] \to R_1$, π_2: $[0,1] \to R_2$, such that (3′) holds. (See Fig. 10.) Assume without loss of generality that the graph of π_i is a polygonal path $\bar{\rho}_i$ satisfying conditions (2) for some $c_i, c_i^* \in h_i$ $(i = 1,2)$. That is,

$$\pi_1(\tau) = \bar{\rho}_1(s(\tau)), \qquad \pi_2(\tau) = \bar{\rho}_2(t(\tau)) \quad \text{for all } \tau \in [0,1], \tag{4}$$

for some continuous, piecewise monotone functions $s(\tau)$, $t(\tau)$: $[0,1] \to [0,1]$ with

$$s(0) = t(0) = 0, s(1) = t(1) = 1. \tag{5}$$

Further, Let β_i: $[0,1] \to B_i$ be a piecewise linear "parametrization" of the portion of B_i between c_1 and c_2 such that

$$\beta_i(\tau) = \bar{\rho}_1(\tau) = \bar{\rho}_2(\tau) \quad \text{for every } \tau \in \left[0, \tfrac{1}{3}\right] \cup \left[\tfrac{2}{3}, 1\right], \qquad i = 1,2.$$

First we are going to show that a person A can walk from c_1 to c_2^* along $\bar{\rho}_1$, and B can walk from c_1^* to c_2 along β_2, so that their distance is always equal to d. Roughly speaking, this means that there is a motion of a ring of diameter d around P, which is a "contact motion from below."

Assume that this is not true. Then we can apply Theorem 3 with $\rho_1 = \bar{\rho}_1$, $\rho_2 = \beta_2$ to conclude that A can walk from c_1 to c_2 along $\bar{\rho}_1$, and B can walk from c_2 to c_1 along β_2, so that their distance is always larger than d. That is,

$$|\bar{\rho}_1(s'(\tau)) - \beta_2(t'(\tau))| > d \qquad \text{for all} \tau \in [0,1] \tag{6}$$

for some continuous, piecewise monotone functions $s'(\tau)$, $t'(\tau)$: $[0,1] \to [0,1]$ with

$$s'(0) = t'(1) = 0, s'(1) = t'(0) = 1. \tag{7}$$

By (5) and (7), we can apply Theorem 1 to obtain the fact that there exist two continuous, piecewise monotone functions $\tau(x)$, $\tau'(x)$: $[0,1] \to [0,1]$ satisfying (1). Then (4) and (3′) imply that

$$|\bar{\rho}_1(s(\tau(x))) - \bar{\rho}_2(t(\tau(x)))| = |\pi_1(\tau(x)) - \pi_2(\tau(x))| \leqslant w \leqslant d.$$

On the other hand, by (6) and (1),

$$|\bar{\rho}_1(s(\tau(x))) - \beta_2(t'(\tau'(x)))| = |\bar{\rho}_1(s'(\tau'(x))) - \beta_2(t'(\tau'(x)))| > d.$$

Applying Lemma 2 to the three curves $\gamma_1(x) = \bar{\rho}_1(s(\tau(x)))$, $\gamma_2(x) = \bar{\rho}_2(t(\tau(x)))$, and $\gamma(x) = \beta_2(t'(\tau'(x)))$, $x \in [0,1]$, we obtain the fact that $\gamma(x)$ lies on the segment $\overline{\gamma_1(x), \gamma_2(x)}$ for some $x \in [0,1]$. This contradicts the last two inequalities.

Thus, we can conclude that there is a motion of a ring of diameter d along P, which is a "contact motion from below." Performing the same trick again (with β_1, instead of β_2), we obtain a motion of a ring of diameter d along P, which is a contact motion (both from below and from above). This completes the proof. □

4. Computing the ring-width. Since our original motivation arose from the practical subject of robotics, let us now address the question of computation. We shall exploit the existence of contact motions (Theorem 7) to derive an algorithm for determining the ring-width of P. If P has n vertices then the algorithm runs in time $O(n^2 \log n)$. (Recall that while $O(f(n))$ means $\leqslant c_1 f(n)$, $\Omega(f(n))$ means $\geqslant c_2 f(n)$ for some constants $c_1, c_2 \neq 0$ and n sufficiently large; we shall use $\Omega(f(n))$ more loosely to mean $\geqslant c_2 f(n)$ for infinitely many values of n.) But first it is worthwhile pointing out that time $\Omega(n^2)$ is unavoidable for any algorithm that actually computes a contact motion attaining the minimum width (our algorithm has this property).

Consider the polygonal path shown in Fig. 13.

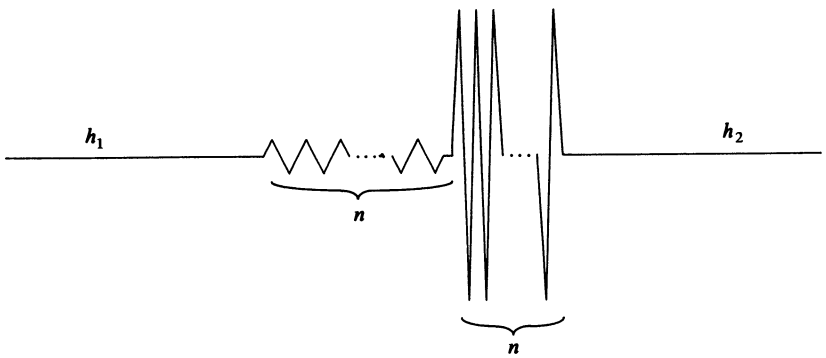

FIG. 13

The path has n "small" sawteeth followed by n "large" sawteeth. The small (resp. large) sawteeth have amplitude $1/n$ (resp. 1) and wavelength $1/n$ (resp. $1/n^2$). Clearly any contact motion of a unit-width ring from one handle h_1 to the other handle h_2 must make $\Omega(n^2)$ "basic moves" where we may define a basic move to be a motion that is monotone in both x and y. One can easily turn this example into a polygon with handles with the property that passing a minimum width ring from one handle to the other requires $\Omega(n^2)$ contact moves. Let us note that this lower bound is invalid for the example if we allow noncontact motion or if we allow an elastic ring; it is then unknown if $\Omega(n^2)$ basic moves can be forced.

Now we address the problem of computing $w_R(P)$. Regard P as the union of two polygonal paths γ_1, γ_2 (the "upper" and "lower" paths). The initial and final portions of the two paths coincide to form the handles. Let the vertex set and edge set of γ_i $(i = 1, 2)$ be V_i and E_i, respectively. A *position* $p = (x_1, x_2)$ is a pair with $x_1 \in \gamma_1$, $x_2 \in \gamma_2$. The (*combinatorial*) *type* of p is a pair

$$T(p) = (u_1, u_2) \in (V_1 \cup E_1) \times (V_2 \cup E_2),$$

where $x_1 \in u_1$, $x_2 \in u_2$. We regard an edge $e_i \in E_i$ as an open segment, so the combinatorial type of a position is unambiguous. Furthermore, if h_1, h_2 are the two handles of P, we assume they are in both E_1 and E_2.

Let $\pi = (\pi_1, \pi_2)$ $(\pi_i\colon [0,1] \to \gamma_i)$ be a contact motion of an elastic ring. The width of π is $\sup_t |\pi_1(t) - \pi_2(t)|$. From π, we get a collection of combinatorial types

$\{T(\pi(t)) \mid t \in [0,1]\}$. We may assume that π is well-behaved so that there are a finite number of changes of combinatorial type as t varies. (Otherwise, by standard arguments, we can replace π by a well-behaved π' with width $(\pi) \geqslant$ width (π') and this π' suffices for our purposes.) Thus we derive from π a finite sequence of combinatorial types

$$T_1, T_2, \ldots, T_k$$

such that the unit interval $[0,1]$ is divided into k time intervals (open, closed or half open)

$$I_1, I_2, \ldots, I_k$$

such that for each $t \in I_i$, $\pi(t)$ is of type T_i.

For elements $u, u' \in V_i \cup E_i$, we say u, u' are *adjacent* if either $u = u'$, or u and u' are incident to each other (so that one is a vertex and the other an edge). For combinatorial types $(u, v), (u', v') \in (V_1 \cup E_1) \times (V_2 \cup E_2)$, we say (u, v) and (u', v') are *adjacent* if both u, u' are adjacent and v, v' are adjacent.

For each combinatorial type (u, v), choose a *canonical position* $C(u, v)$ to be any position (x, y) where $x \in \bar{u}$, $y \in \bar{v}$ such that $|x - y|$ is minimized. Here, $\bar{u}$ is the topological closure of u, so an edge u becomes a closed segment $\bar{u}$. It is not hard to see that $C(u, v)$ is unique unless u, v are parallel edges, in which case we choose (x, y) so that at least one of x or y is a vertex. Also, we allow $x = y$, a possibility which arises if u, v are adjacent edges.

It is then easy to see:

LEMMA 3. *From any position $(x, y) \in (\gamma_1, \gamma_2)$, there is a motion $\pi_{x,y}$ from (x, y) to $C(T(x, y))$ such that the type of $\pi_{x,y}(t)$ for $t \in (0,1)$ is $T(x, y)$ and the width of $\pi_{x,y}$ is attained at $\pi_{x,y}(0) = (x, y)$.*

Using this property, we may now modify π to π' so that for each time interval I_i, there is a moment $s_i \in I_i$ when $\pi'(s_i) = C(T_i)$. Note that $T(\pi'(s_i))$ need not be equal to T_i, but for all $t \in I_i - \{s_i\}$, $\pi'(t)$ has type T_i as before. Furthermore width(π') = width(π).

Therefore we see that width(π') is attained in the portion of π' between two canonical positions $\pi'(s_i)$ and $\pi'(s_{i+1})$. By a further modification, we can convert π' to π'' so that width(π'') is attained precisely at some canonical position. The basis of this modification is the next lemma.

LEMMA 4. *For any pair of adjacent canonical positions (x, y) to (x', y'), there is a "canonical motion" $\pi_{x,y,x',y'}$ from (x, y) to (x', y') whose width is attained as $|x - y|$ or $|x' - y'|$.*

Proof. It is easy to see that there is a pair of edges e and f such that $x, x' \in \bar{e}$, $y, y' \in \bar{f}$. Then the motion

$$\pi_{x,y,x',y'}(t) = (tx' + (1-t)x, ty' + (1-t)y)$$

has the desired property. □

Repeated application of this lemma, by replacing the portion of π' between consecutive canonical positions $\pi'(s_i)$ and $\pi'(s_{i+1})$, gives us our final motion π''.

Furthermore, we assume that $\pi''(0) = (x_0, y_0)$ with $x_0 = y_0$ in h_1, and likewise $\pi''(1) = (x_1, y_1)$ with $x_1 = y_1$ in h_2. We say that such a π'' is in *canonical form*.

COROLLARY. *For any polygon P with handles, $w_R(P)$ is equal to the distance between some pair u, v with $u \in V_1 \cup E_1$ and $v \in V_2 \cup E_2$, where at least one of u or v is a vertex.*

We need one more concept before we present the algorithm. Let us define the *initial* and *final* positions. The polygon P has handles h_1 and h_2. Let v_i be the vertex where h_i attaches to P. Then the canonical positions (v_1, v_1) and (v_2, v_2) are called the initial and final positions respectively.

To derive an algorithm, we proceed as follows: we sort the set of all values $|x - y|$ where (x, y) range over all canonical positions. Let the sorted values be

$$0 = r_0 \leqslant r_1 \leqslant \cdots \leqslant r_k,$$

where there are k canonical positions and we break ties indiscriminately. Let (x_i, y_i) be the canonical position associated with $r_i = |x_i - y_i|$. We partition these canonical positions into a collection of disjoint classes. These classes represent an equivalence relation (called *mutual accessibility*) among the canonical positions. Initially, each canonical position is put into its own singleton set, so that they are all inaccessible from each other.

Now we process these canonical positions in stages, where in the ith stage $(i = 0, 1, \ldots)$, we process (x_i, y_i): when processing (x_i, y_i), we check each canonical position (x_j, y_j) adjacent to (x_i, y_i) (there are at most 6 such). If $j < i$ (so that (x_j, y_j) has already been processed) then we make (x_i, y_i) and (x_j, y_j) mutually accessible. In other words, if they are not yet mutually accessible, then we merge the two classes containing (x_i, y_i) and (x_j, y_j), respectively. We say stage i is *terminal* when the initial position and the final position are mutually accessible. This termination condition is checked just before we start each new stage. The terminal value of $r_i = |x_i - y_i|$ gives us the minimum width.

The justification of this algorithm comes from the following observation:

After we process a canonical position (x_i, y_i), two canonical positions with indices $j_1, j_2 \leqslant i$ are mutually accessible using motions of width at most r_i if and only if they belong to the same class.

Let us briefly note the time complexity of this algorithm: The sorting of the r_i's takes time $O(n^2 \log n)$ since there are $O(n^2)$ canonical positions. To process each (x_i, y_i), we need to maintain a collection of disjoint sets, quickly form the union of two disjoint sets, and determine for any two elements whether they belong to the same set. All of this can be accomplished very efficiently using the well-known *union-find* data structure [1]. In particular, any sequence of t such operations takes time $O(t\alpha(m))$ where m is the total number of elements in the sets and $\alpha(m)$ is a very slowly growing function (the so-called inverse Ackermann function). Since $m = O(n^2)$ and $t = O(n^2)$ in our case, the overall time complexity is $O(n^2 \log n)$.

REFERENCES

1. A. V. Aho, J. E. Hopcroft, and J. D. Ullman, The Design and Analysis of Computer Algorithms, Addison-Wesley, Reading, MA, 1974.
2. Fan Chung and Paul Seymour, personal communication.

3. David Eppstein, personal communication.
4. John Kutcher and Joseph O'Rourke, personal communication.
5. J. T. Schwartz and M. Sharir, Efficient motion planning in environments of bounded local complexity, Tech. Rep. No. 164, NYU-Courant Inst.
6. G. Strang, The width of a chair, *Amer. Math. Monthly*, 89 (1982) 529–535.
7. Chee K. Yap, Algorithmic motion planning, in Advances in robotics: Volume 1 (J. T. Schwartz and Chee K. Yap, eds.), Lawrence Erlbaum Assoc's, Hillsdale, NJ, 1987.
8. _____, How to move a chair through a door, *IEEE Trans. Automat. Control*, to appear.
9. _____, How to move a chair through a door, *IEEE J. Robotics and Automation*, RA-3 (1987) 172–181.

LETTERS TO THE EDITOR

Editor,

D. H. Lehmer's article "A New Approach to the Bernoulli Polynomials" [*Monthly*, Dec. 88] makes very pleasant reading but there is a small gap in the presentation. To define the Bernoulli polynomials via Lemma 1 requires that the b_r as given by Equation (11) should be *independent of m*. This has not been shown. In fact, Lemmas 1 and 2 show that, *for each fixed m*, there exists a unique monic polynomial of degree n, say $B_n^{(m)}(x)$, satisfying (6). Since the Fourier series $\phi_n(x)$ of Theorem 4 is *independent of m* and satisfies (6) *for all m*, the superscript in $B_n^{(m)}(x)$ can be dropped.

Sincerely,

DAVID CALLAN
University of Bridgeport
Bridgeport, CT 06601

Editor,

While I do wish to thank J. M. Patin for sharing his short proof of Stirling's formula [**3**] with readers of the MONTHLY, it should be noted that his proof is essentially the same as the one which appears as "Problem #95" in Donald J. Newman's *A Problem Seminar* [**1**; p. 20 (problem), p. 38 (hint), p. 104 (solution)], but in greater detail. In his review of Newman's book, Ivan Niven [**2**] states that this "challenging book for problem buffs" has only one "really misplaced problem," namely this one, requesting "a proof of one form of Stirling's formula, which is surely a textbook matter."