Optimizing and Approximating Geometric Covering Tours

Joseph S. B. Mitchell
Stony Brook University

Fields Conference on Discrete Geometry and Optimization, Sept 19-23, 2011

What are "geometric covering tours"?

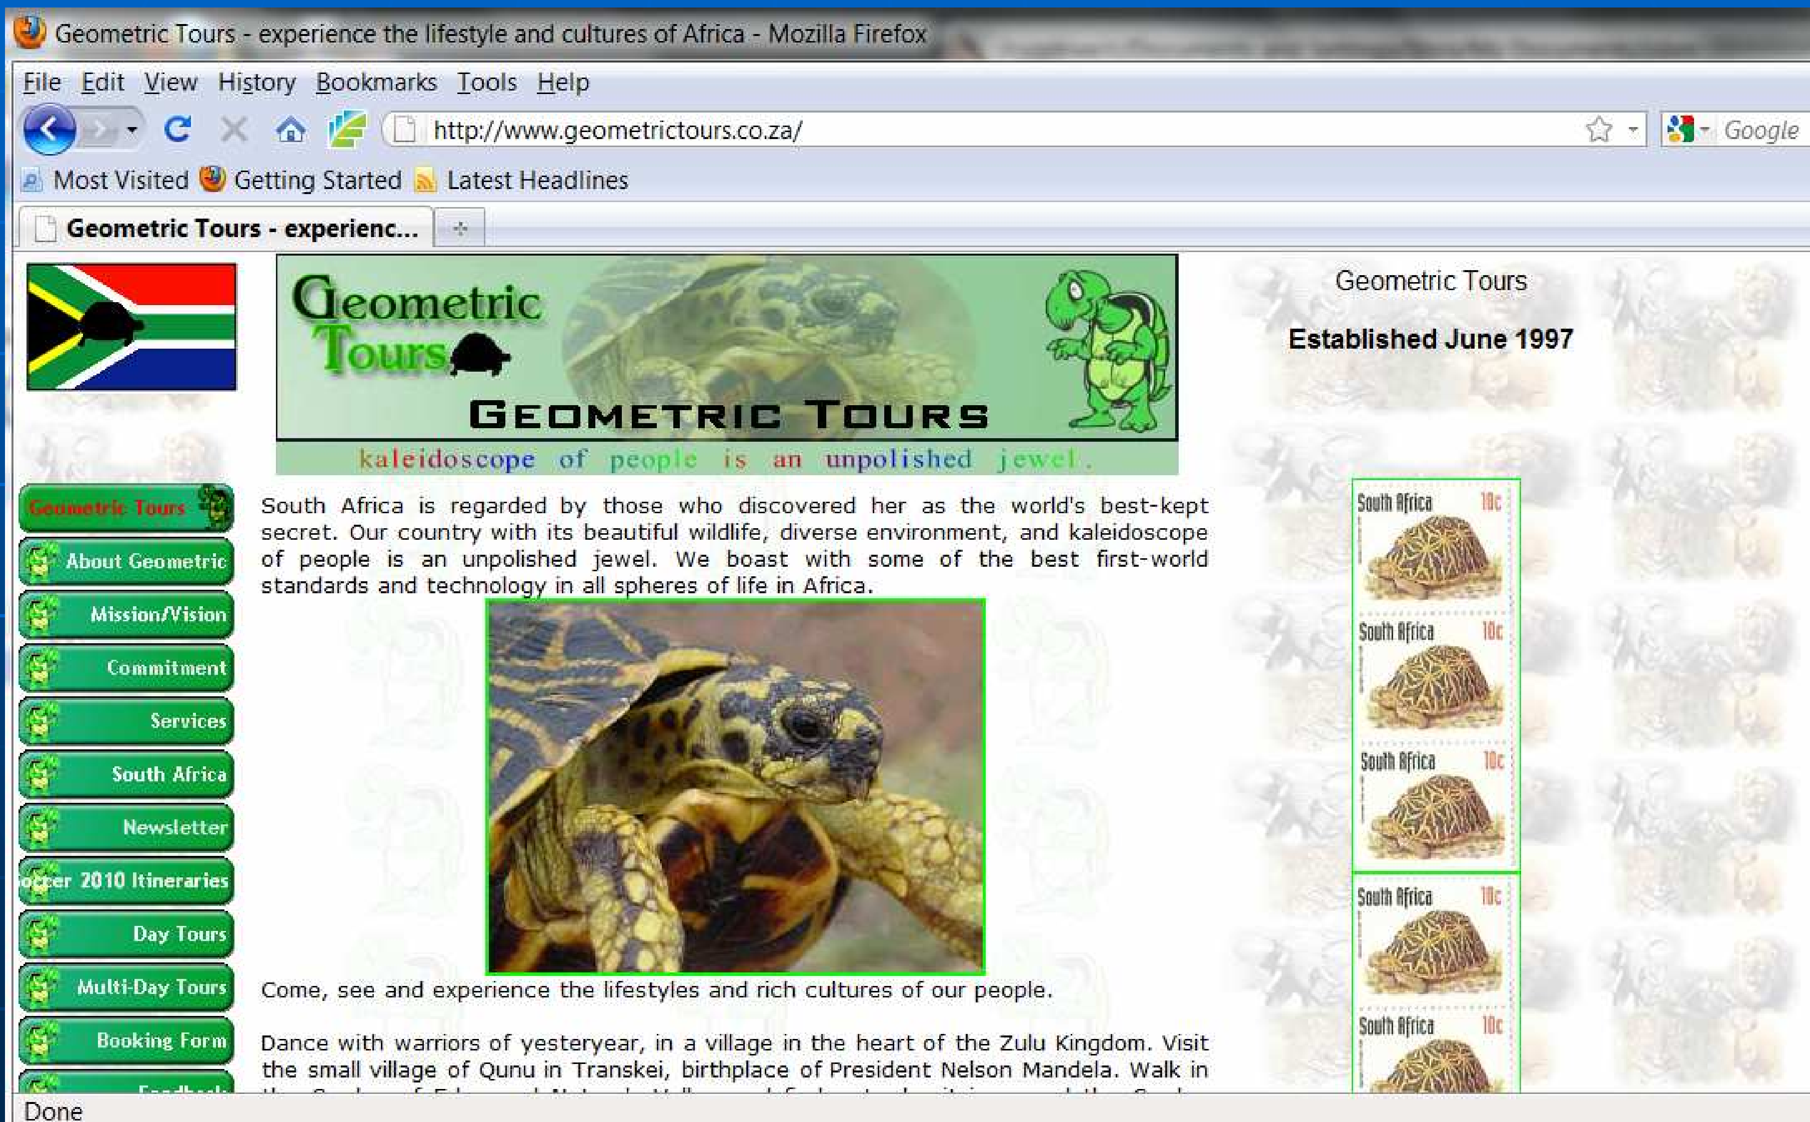

Covering Tours

- Cover a point set S

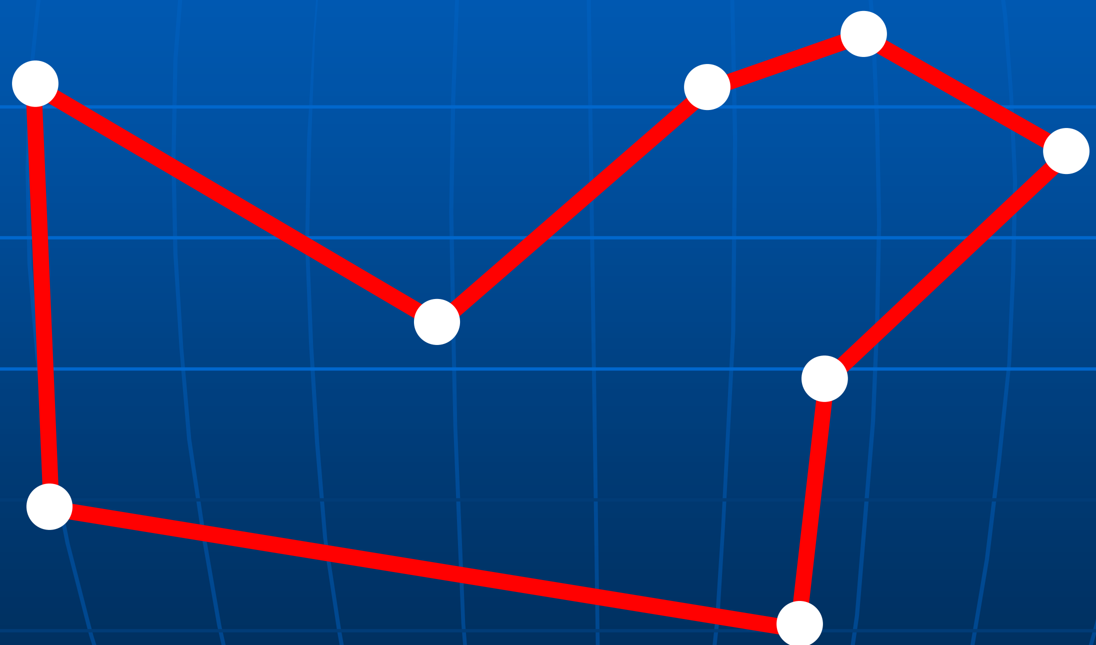

Just geometric TSP

Covering Tours

- Cover a set of disks

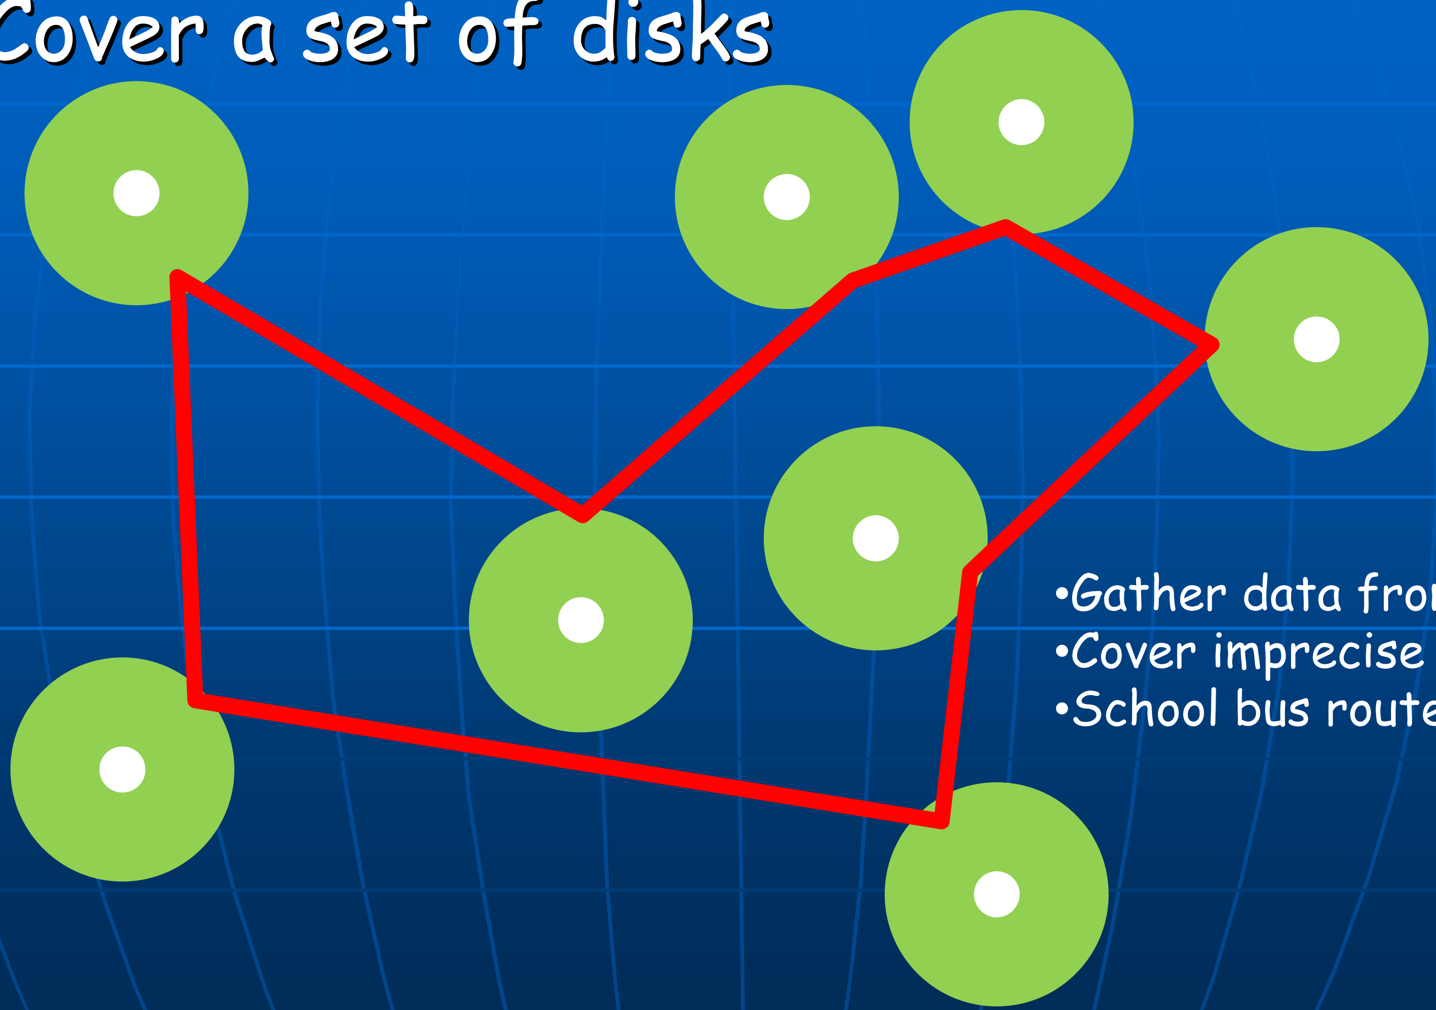

- Gather data from sensors
- Cover imprecise points
- School bus route

TSP with (circular) neighborhoods

Covering Tours

- Cover a set of polygons

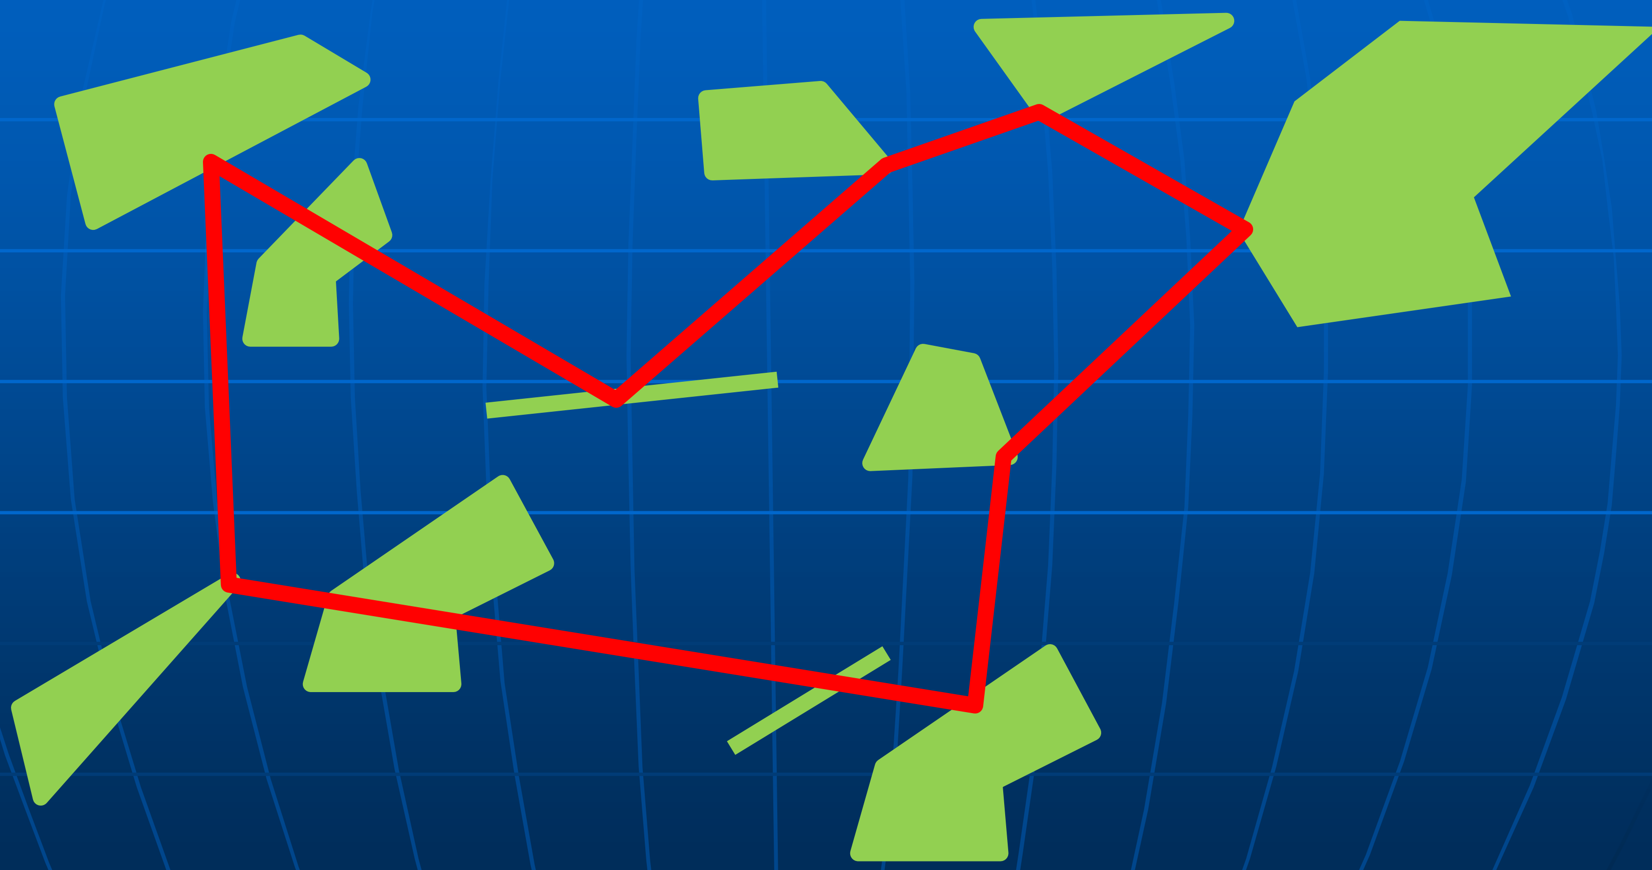

TSP with neighborhoods

Covering Tours

- Cover set of all visibility polygons

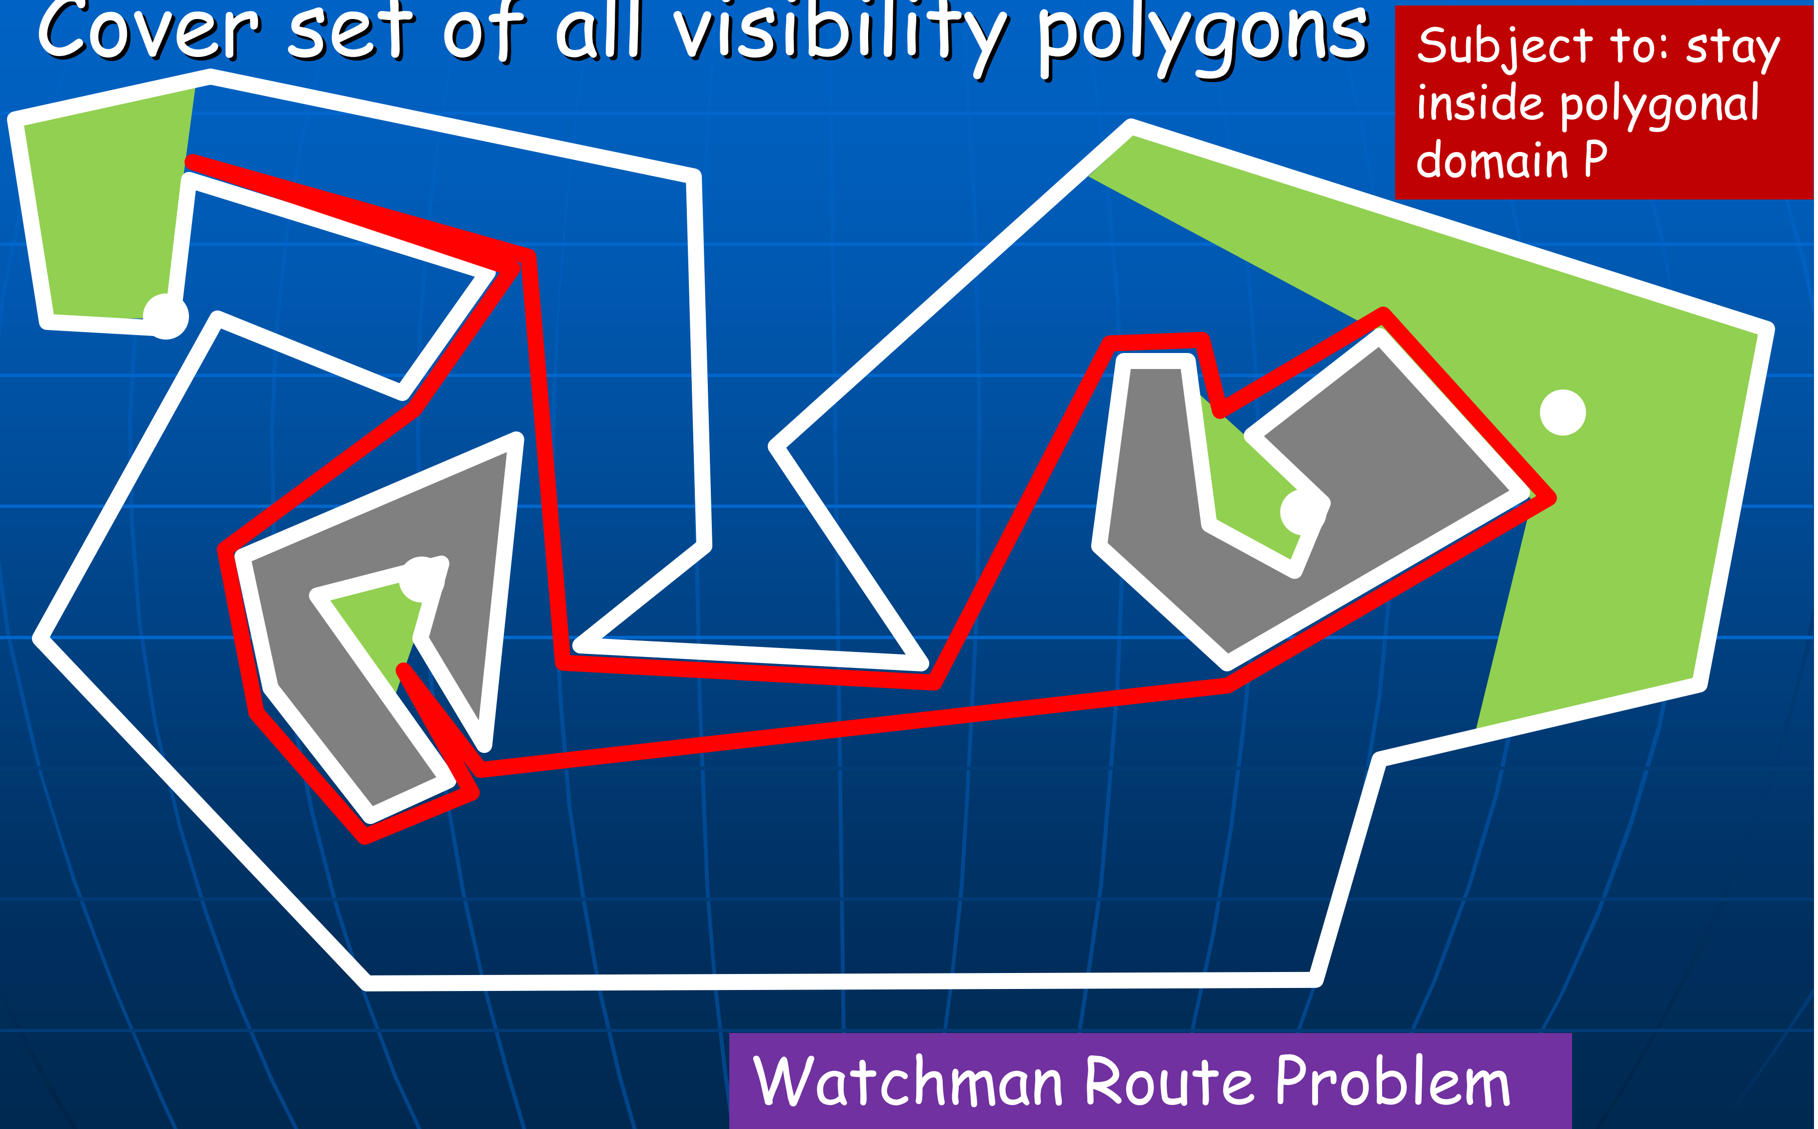

A Brief Taxonomy

- Type of network: path, tour, tree

Up to constants in approx, no difference

- What must be covered/visited: set S of points, regions, visibility polygons, distance-truncated visibility regions, sets of points/regions, etc.
- Objective function: min-total-length (Euclidean, L_1 , geodesic), bottleneck (min-max, max-min), number of links/Steiner points, total amount of turning, number of reflex vertices ("reflexivity"), other functions of edge lengths, etc
- Other constraints: obstacles, link lengths, order of visitation, convex tours, time windows, separate certain pairs of points, vehicle capacities, etc.
- Online vs. Offline

Outline

- Introduction: Geometric Covering Tours
- TSP with Neighborhoods
- Watchman Route
- Lawn mowing/milling
- Data Gathering in Sensor Networks
- Other Covering Tour Problems

TSPN: TSP with Neighborhoods

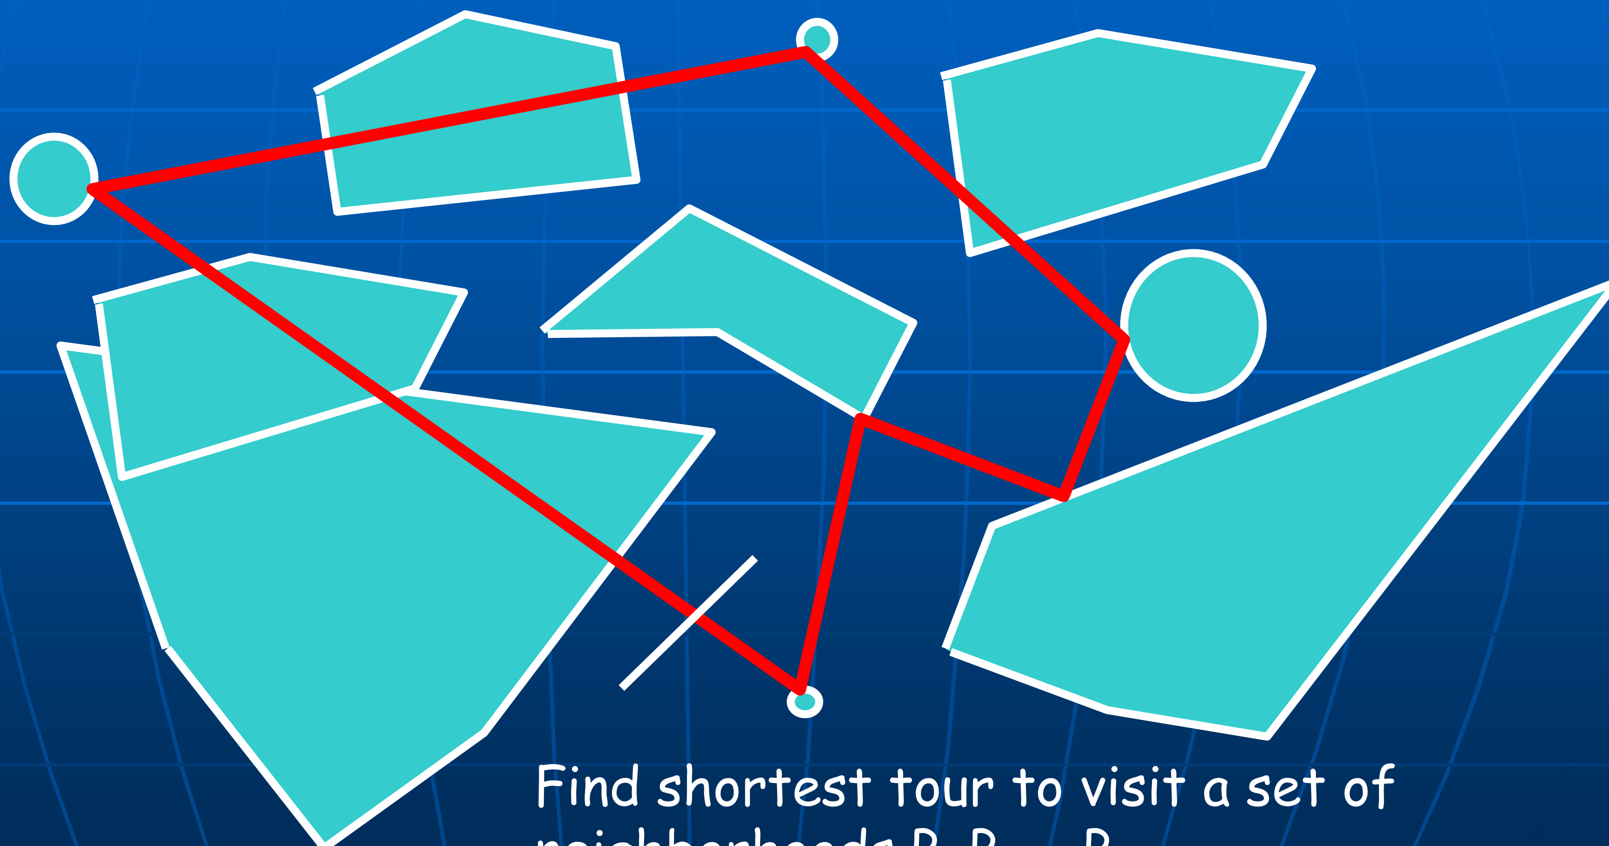

Find shortest tour to visit a set of neighborhoods $P_1, P_2, \ldots, P_n$

Background on TSPN

Generalizes 2D Euclidean TSP (thus, NP-hard)

Introduced by [Arkin & Hassin, 1994]

- "obvious" heuristics do not work:
  - TSP approx on centroids (as representative points)
  - Greedy algorithms (Prim- or Kruskal-like)
- O(1)-approx, time O(n + k log k), for "nice" regions:
  - Parallel unit segments
  - Unit disks
  - Translates of a polygon P
- Combination Lemma

TSPN: Approximation

- Hope: Since geometric TSP on points is "easy" to approximate (has O(n log n) PTAS), maybe TSPN does too!

General Connected Regions

O(log k)-approx [Mata & M, SoCG'95]

Use guillotine rectangular subdivisions, DP

(*non* – disjoint: regions may overlap)

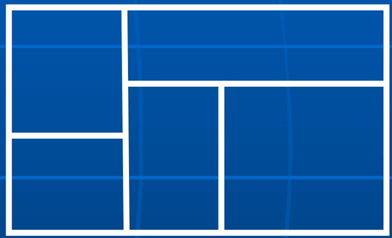

- $O(n^5)$ time

[Mata & M, SoCG'95]

- $O(n^2 \log n)$

[Gudmundsson & Levcopoulos, 1999]

k = # regions n = # vertices of all regions

O(1)-Approximations

- Unit disks, parallel unit segments, translates of *P*

  [Arkin & Hassin, 1994]

- Connected regions of comparable size

  [Dumitrescu & M, SODA'01]

- Disjoint **fat** regions of *any* size [de Berg, Gudmundsson, Katz, Levcopoulos, Overmars, van der Stappen, ESA'02]

- Discrete point sets within disjoint, fat, *non*-convex regions

  [Elbassioni, Fishkin, Mustafa, Sitters, ICALP'05]

- *Non* - disjoint, convex, fat, comparable size

  [Elbassioni, Fishkin, Sitters, ISAAC'06]

- *Arbitrary* (nearly) disjoint connected regions

  [M, SoCG'10]

$(1+\varepsilon)$-Approximations

- Disjoint (or nearly disjoint) **fat** regions of comparable size
  [Dumitrescu & M, SODA'01]
- Point clusters within disjoint **fat** regions of comparable size in R^d [Feremans, Grigoriev, EWCG'05]
- PTAS: Disjoint (or nearly disjoint) **fat** regions of *arbitrary* sizes. (Def: P is **fat** if area(P) = Ω(diam^2(P))) [M, SODA'07]
  Weaker notion than usual "fatness"
- QPTAS: Disjoint, α-fat, arbitrary sizes in R^d
  - With const probability, $(1+\varepsilon)$-approx in time $\mathrm{Exp}(O(1/\varepsilon)^{O(d)}\ O(\alpha)^{O(d^2)} \log^{O(d)} n)$ [Chan, Elbassioni, '08]
  - Also, similar for doubling dimension d

Related Work: APX-hardness

- General connected regions (overlapping):
  - No c-approx with $c<391/390$, unless P=NP

    [de Berg, Gudmundsson, Katz, Levcopoulos, Overmars, van der Stappen, ESA'02]

    (from MinVertexCover)
  - No c-approx with $c<2$

    [Safra, Schwartz, ESA'03]

    (from Hypergraph VertexCover)
- Line segments, comparable length

  [Elbassioni, Fishkin, Sitters, ISAAC'06]
- Pairs of points (disconnected) [Dror, Orlin, 2004]

Exact Poly-Time Solutions

TSPN for a set of infinite lines in 2D:

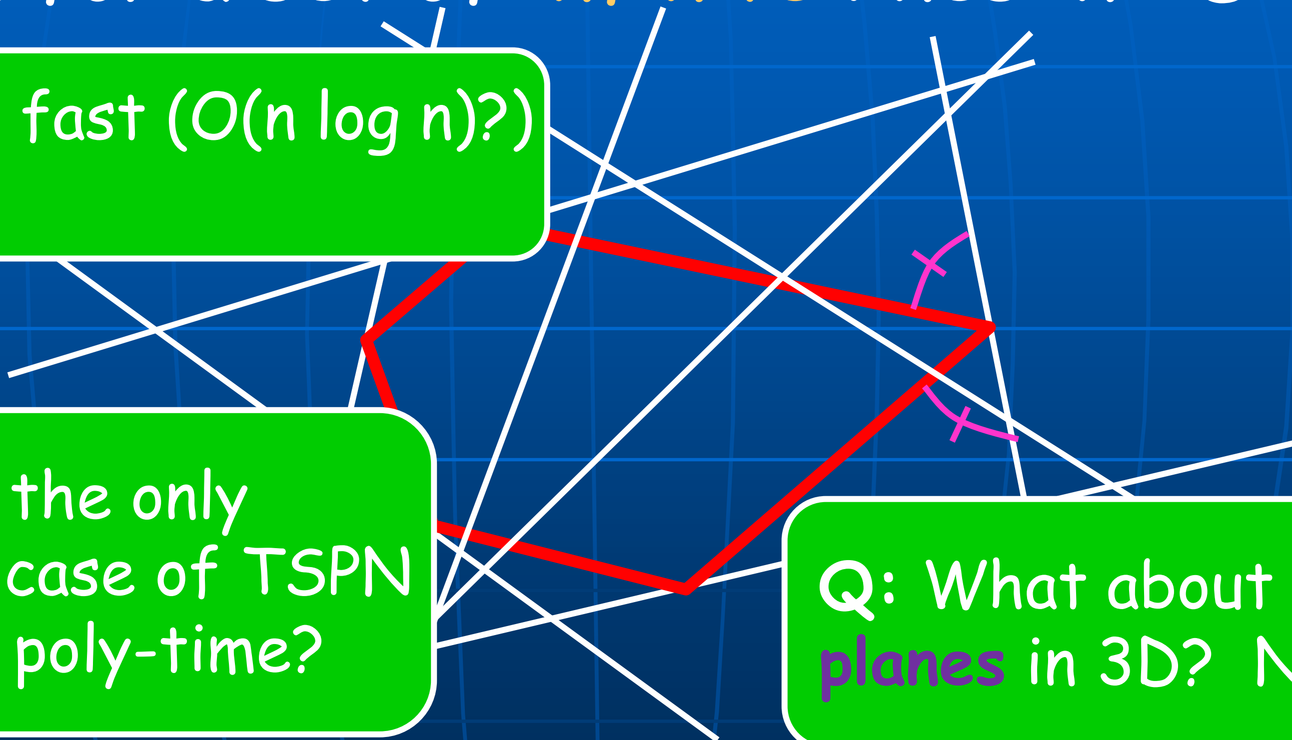

Q: Simple, fast (O(n log n)?) algorithm?

Q: Is this the only nontrivial case of TSPN solvable in poly-time?

Q: What about visiting **planes** in 3D? NP-hard?

Solved in $O(n^4 \log n)$ time using Watchman Route solution

[Dror, Efrat, Lubiw, M, STOC'03]

Difficulty in Applying TSP Methods to TSPN / MSTN

Consider a subproblem (rectangle):

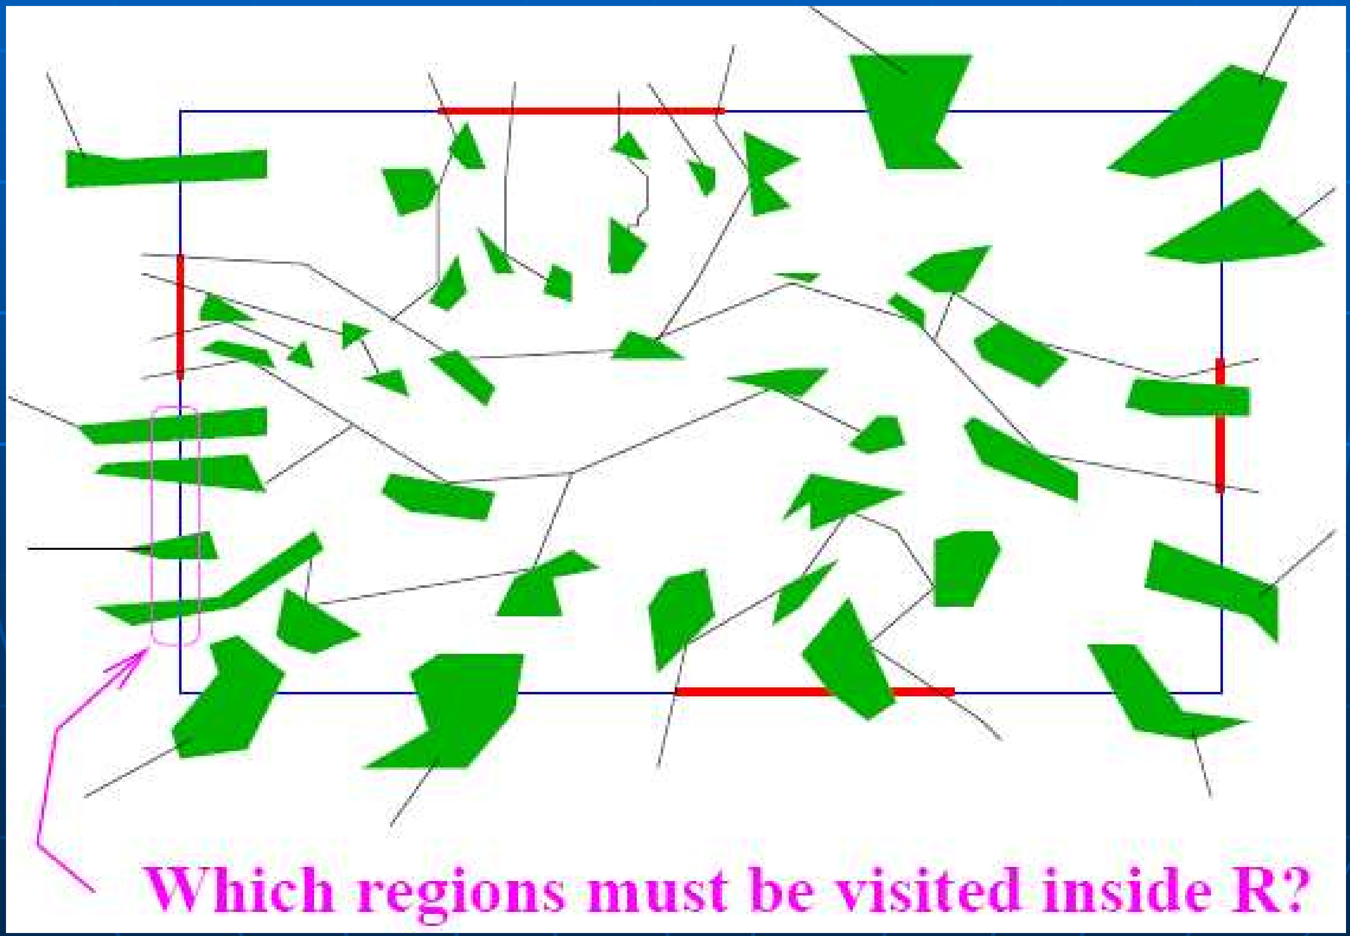

Approximation of 2D TSPN: Connected Regions

	Fat Regions	non-Fat Regions
Comparable sizes	Disjoint: PTAS Non-Disjoint: Newest PTAS, ~~O(1)~~	Disjoint: O(1) Non-Disjoint: O(1), APX-hard
Arbitrary size	Disjoint: Recent PTAS, ~~O(1)~~ Non-Disjoint: Newest O(1), PTAS?, ~~O(log n)~~	Disjoint: Newest O(1), ~~O(log n)~~ Non-Disjoint: O(log n), APX-hard

Conjecture: PTAS for all

Conjecture: O(1) for all

Recent Results [SoCG'10]

- An O(1)-approximation for TSPN for disjoint (or sufficiently disjoint) connected regions in the plane.

Previous: O(log n)-approximation

Also applies to **non**-disjoint convex regions

- A c-approx for TSPN for fat connected regions **implies** an O(c)-approx for TSPN for arbitrary connected regions.

Thus, enough to get O(1)-approx for fat regions to get O(1)-approx for arbitrary connected regions.

Outline of Method

- Replace each (possibly skinny) input region with its four **fat** "directional hulls" (which can overlap)

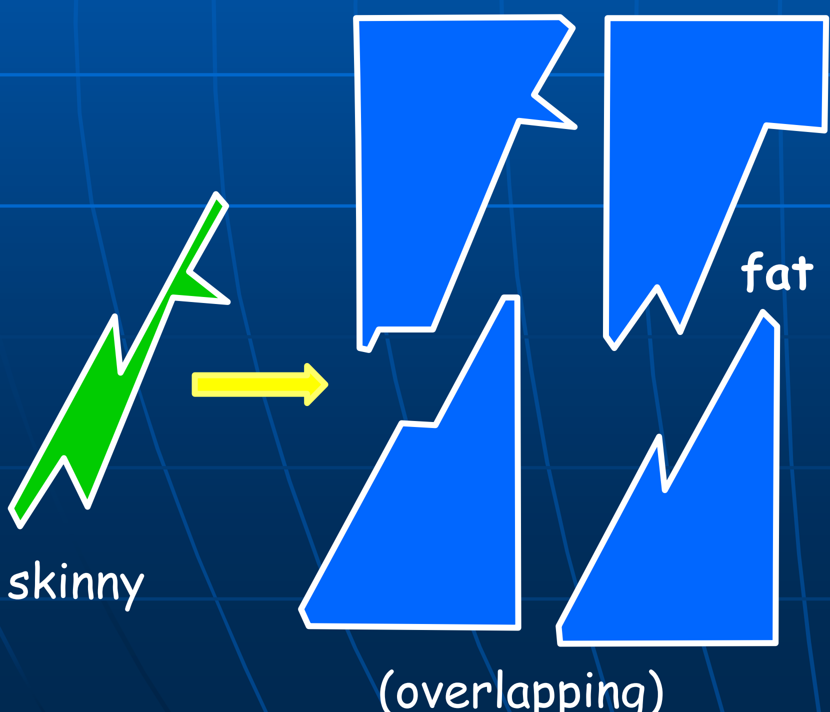

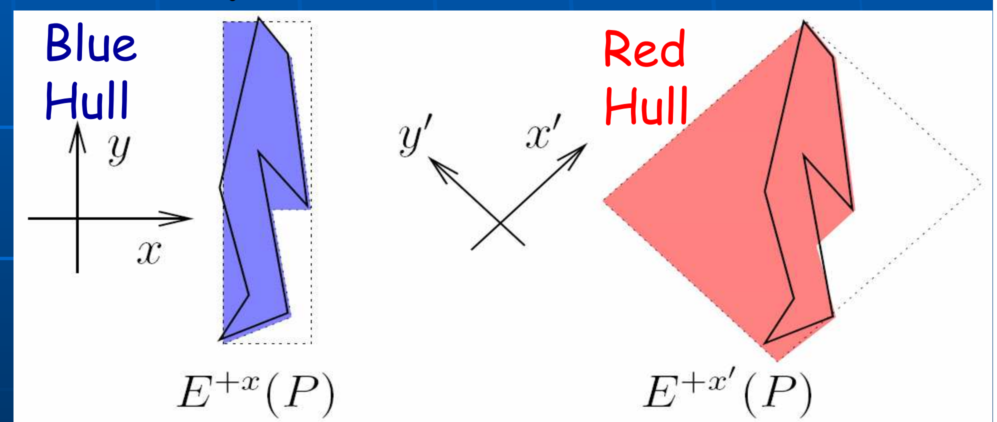

Lemma: Either all 4 blue hulls are **fat** or all 4 **red** hulls are fat

Def: X is ***fat*** if area(X) = Ω(diam^2(X))

Outline of Method

- Select a disjoint subset, $\mathcal{E}_0$, of these hulls: Greedily select in order of increasing size.

Outline of Method

- Compute an approximately optimal tour, T, of the disjoint fat regions $\mathcal{E}_0$

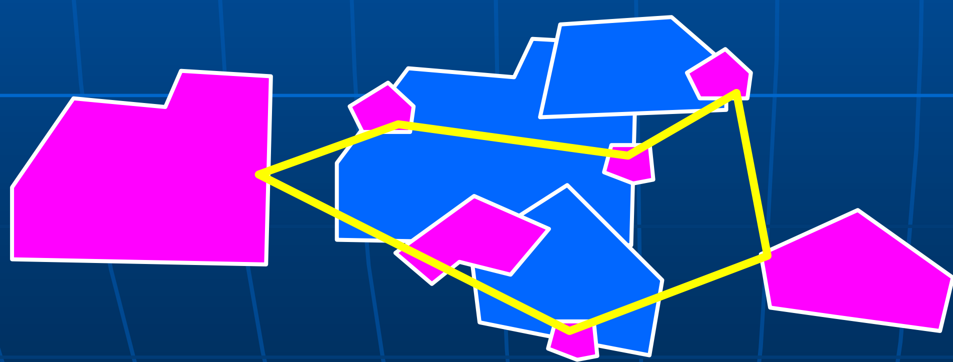

Outline of Method

- Convert T to a polygonal subdivision, G, having histogram faces

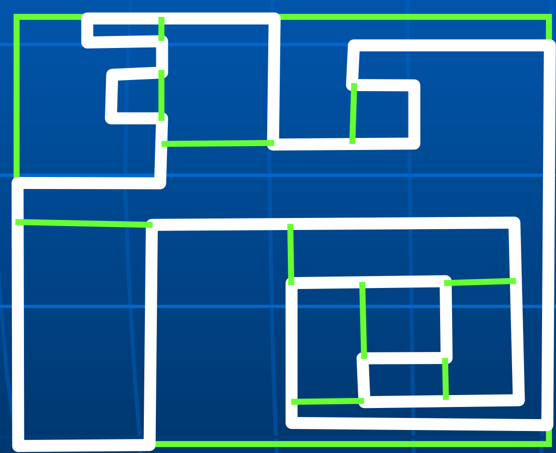

O(1)-factor

Outline of Method

- Any region P_i not visited by G must be close (within distance O(diam(P_i)) of the boundary of the face, H, containing P_i

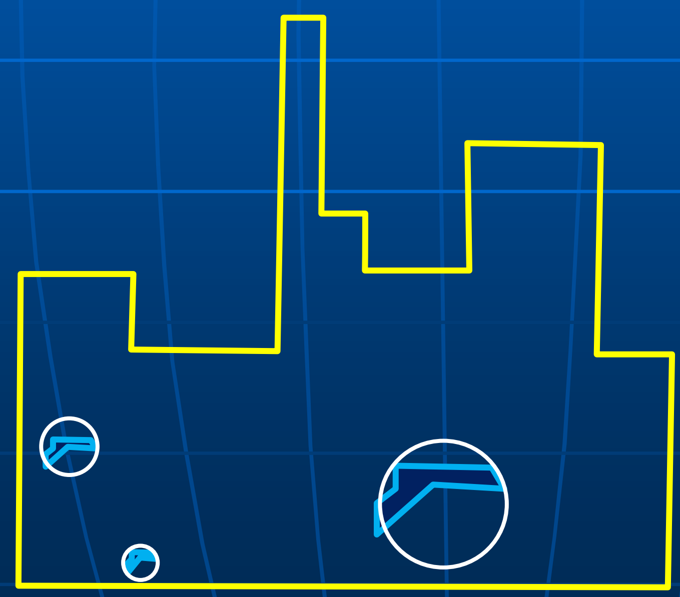

Outline of Method

- New problem for H: Find a min-length forest, F^*, that spans all regions R_H , so that $F^* \cup \partial H$ is connected

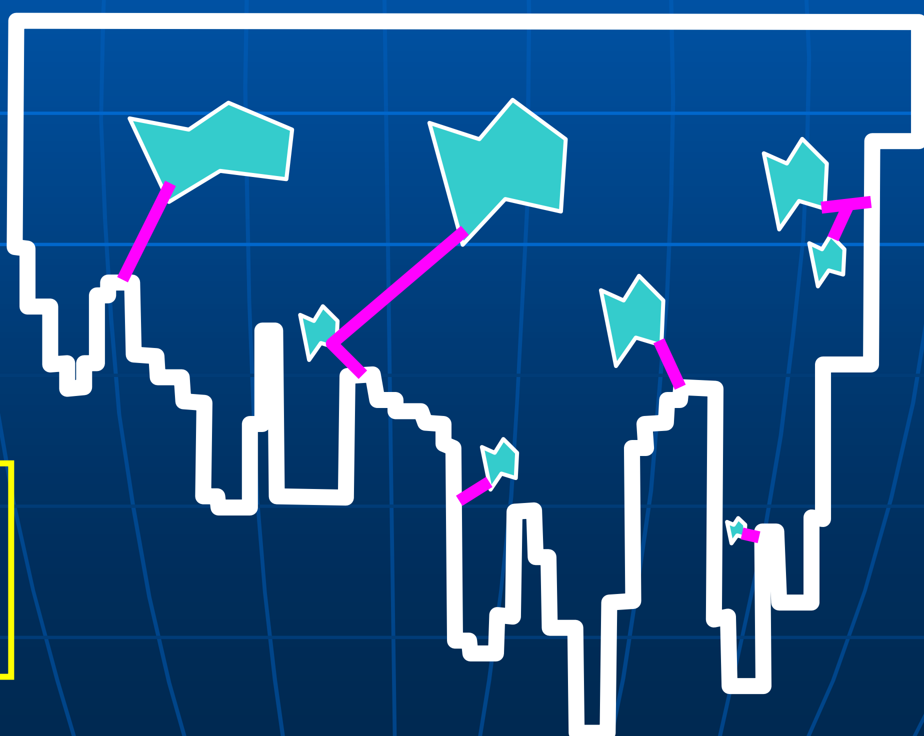

Find min-length forest F linking all regions in H to ∂H.

Outline of Method

- Define stratified grid for *H*

Convert forest problem to special form of set cover using stratified grid

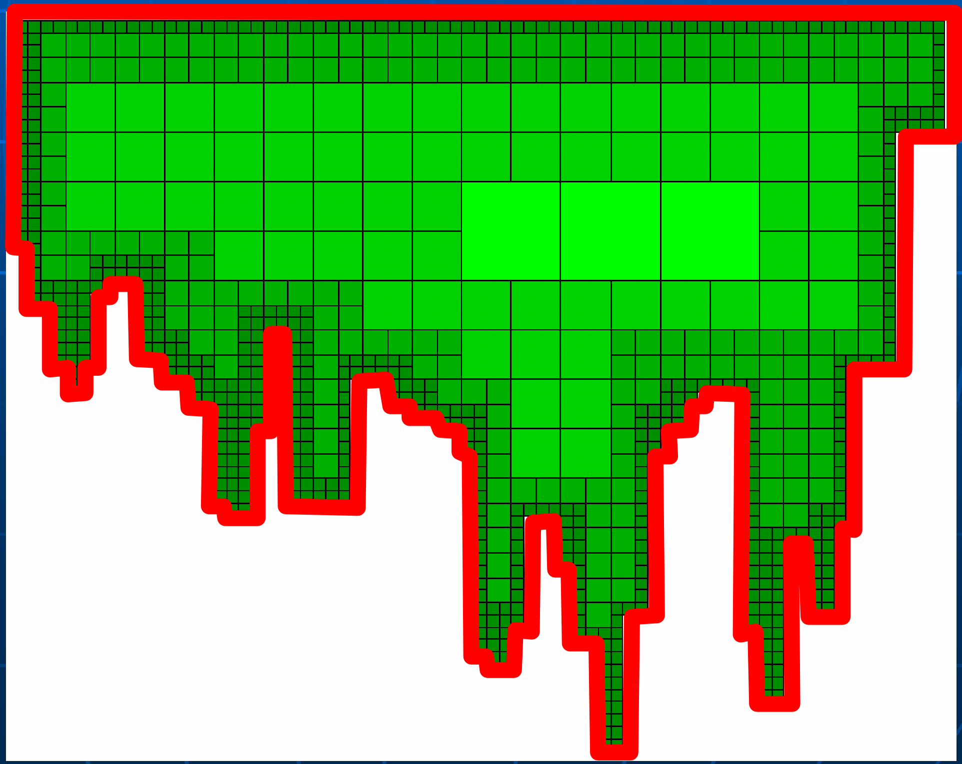

Outline of Method

- Approximation algorithm:
  - Lower bound: OPT $\geq \Omega$(total sizes of grid cells intersecting F^*)
  - Algorithm: DP to find min-weight covering set of grid cells that intersect all regions within H

Related Covering Problem:
Cover segs with fewest rays
DP solves [Katz,M,Nir]

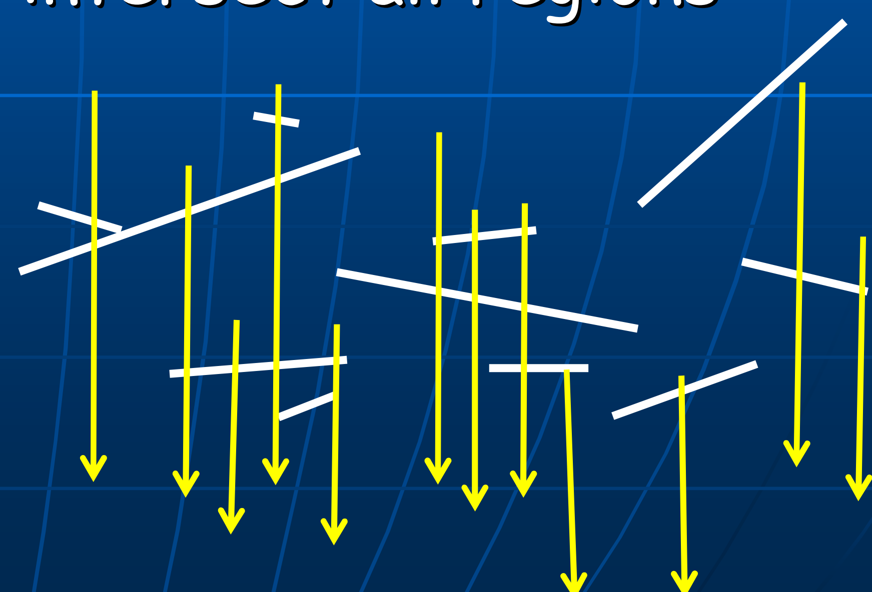

Open Problems

- O(1)-approx for arbitrary connected regions in the plane?

  (known for (nearly) disjoint or convex regions)
- PTAS for arbitrary disjoint connected regions in the plane? (now known for **fat** regions)

  (APX-hardness relies on overlap)
- O(1)-approx for disconnected regions in the plane? (group (class) Steiner, 1-of-a-set TSP)

  (APX-hard for pairs of points)
- What about obstacles? (geodesic metric)

(recent: $\Omega(\log k)$-hard; $O(\log k)$-approx in some cases)

- Higher dimensions: Lines or planes in 3D?

Convex Covering Tours

- Input: Set S of geometric objects in 2D

Goal: Determine if there exists a convex transversal (stabber)

Arik Tamir (3/13/87); parallel segs by DP, [Goodrich-Snoeyink]

Related: Allow objects *interior* to tour [Rappaport, et al]

Convex Covering Tours

- Settle the open problem in 2D: [ADKMPSS'11]
  - Deciding existence of a convex transversal is NP-complete, in general
  - If objects S are disjoint, or form set of pseudodisks, then poly-time algorithm to decide, and to max # objects stabbed

Assumes candidate set *P* of corners of stabber is given.

- 3D: NP-complete, even for disjoint disks

Hard even for **terrain** stabbers!

- Used to compute "convexity" measure

Ordered Covering Tours/Paths

- Order given [DELM]

  Convex: poly-time

  Non-convex, overlapping: NP-hard

- Related to 3D shortest paths

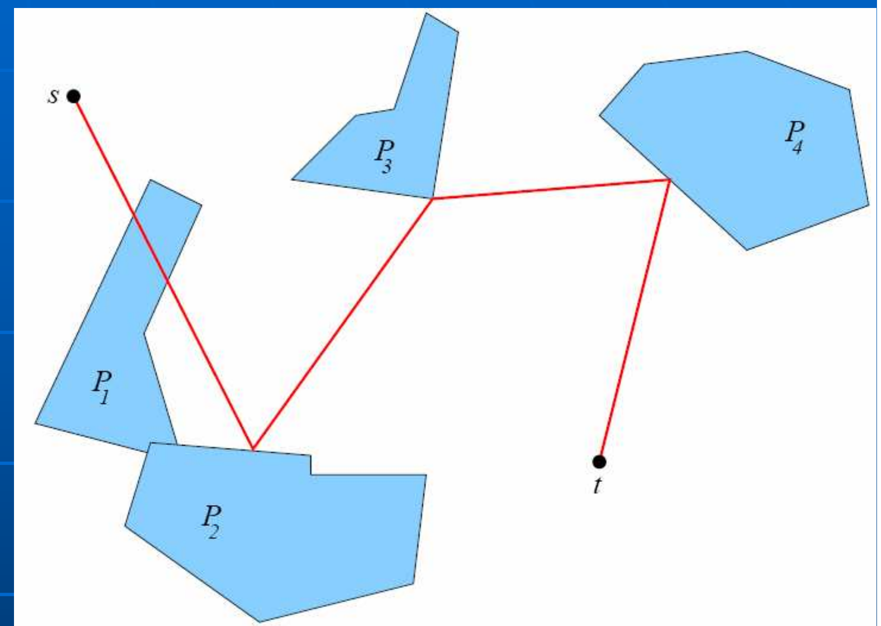

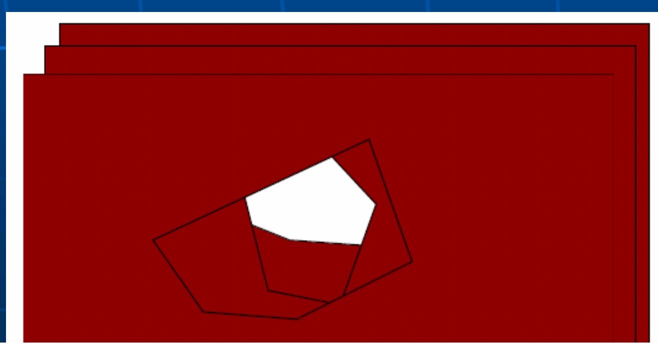

Q: Disjoint non-convex?

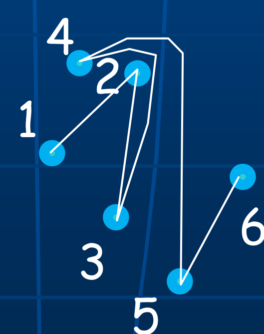

Q: Shortest **simple** tour, even for points?

Touring Regions: SOCP Solution:

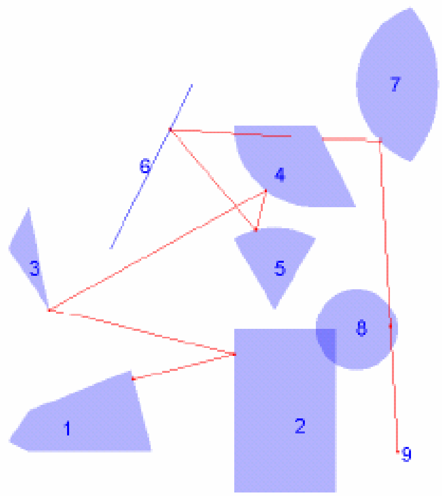

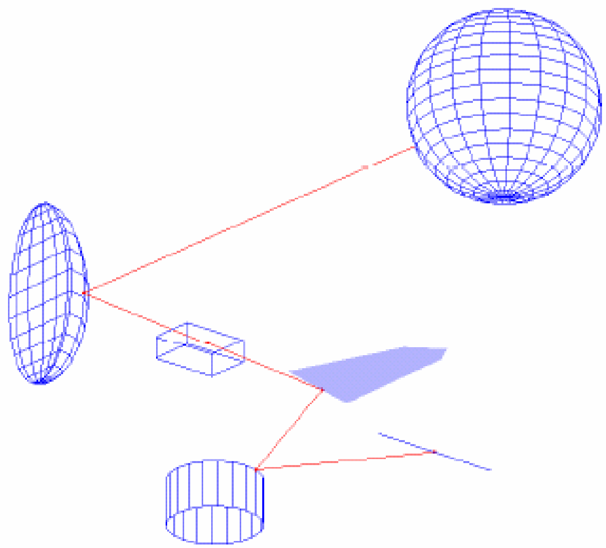

SOCP formulation: Solve in MATLAB (SeDuMi package)

$O(d^2 n^{1.5} K^2 \log 1/\epsilon)$ for $(1+\epsilon)$-approx in $\Re^d$

Can incorporate distance constraints on links, etc.

Safari Problem

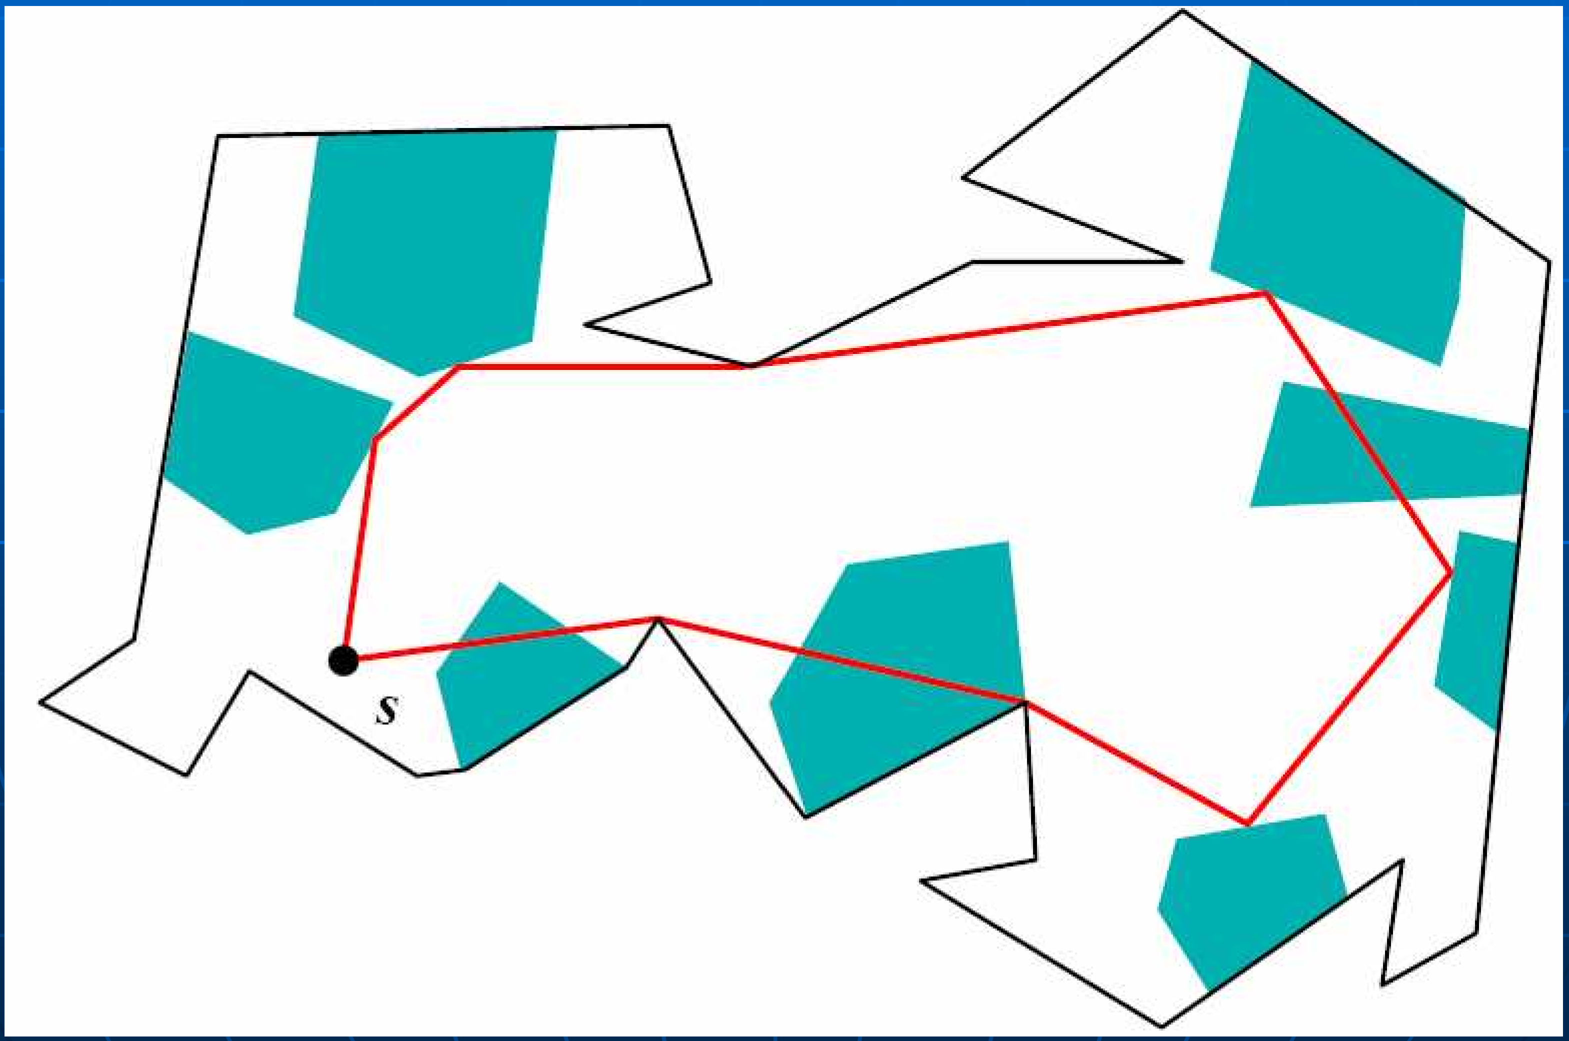

Zookeeper Problem

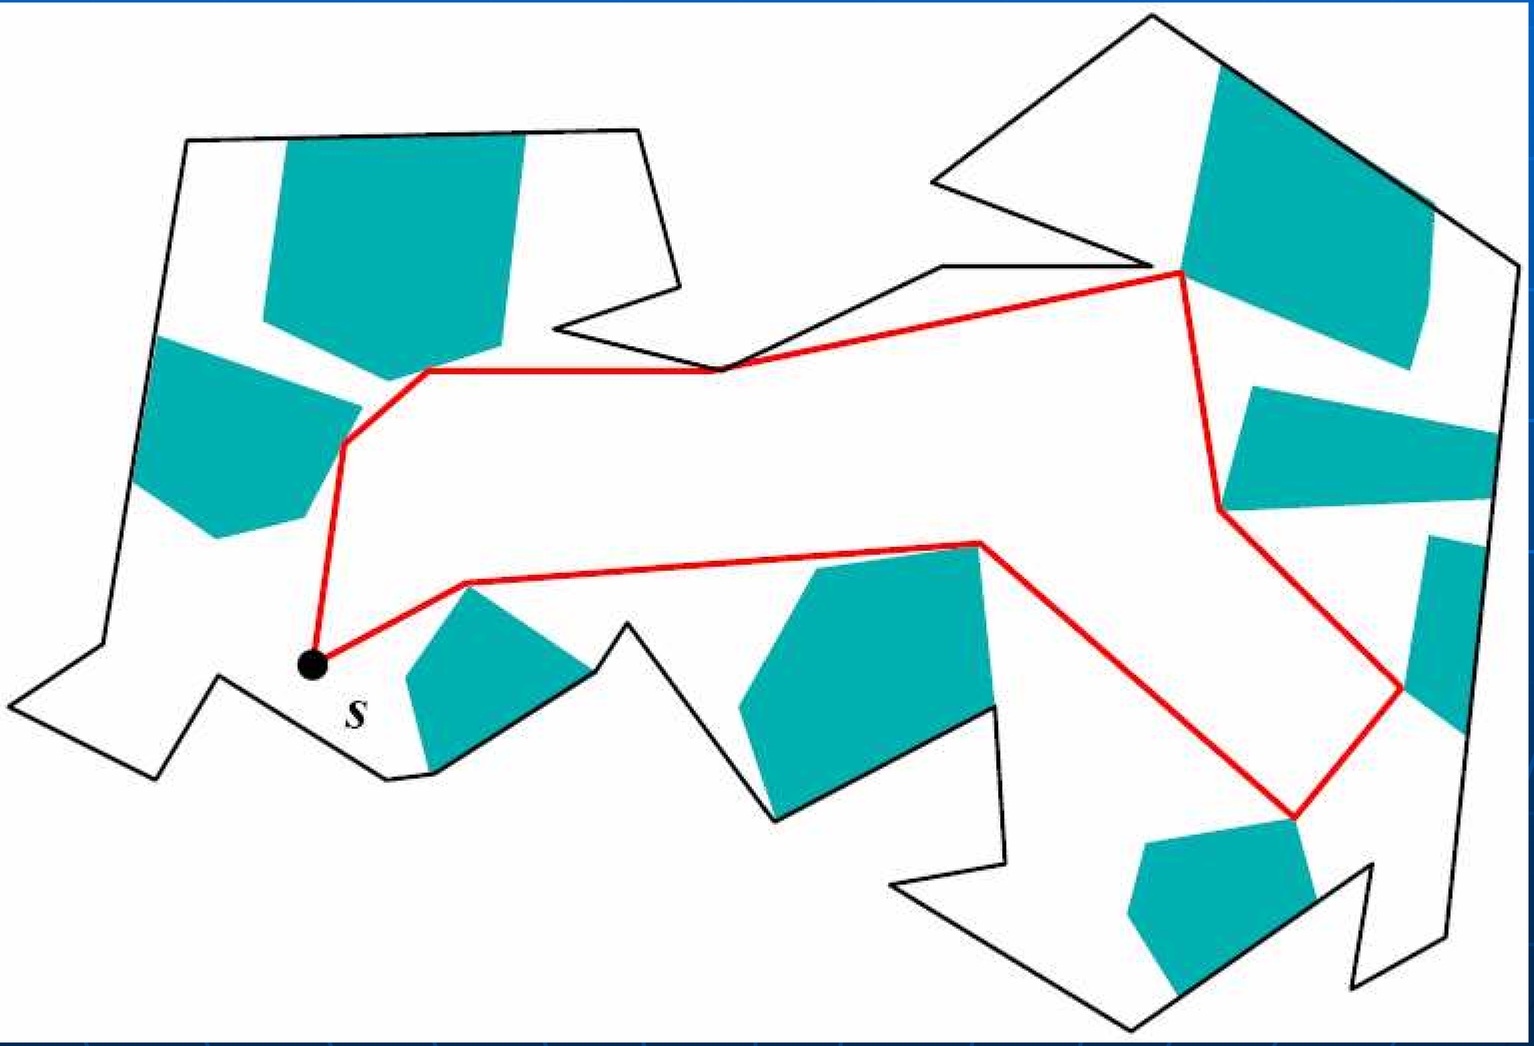

Watchman Route Problem

- Find a shortest tour for a guard to be able to see all of the domain

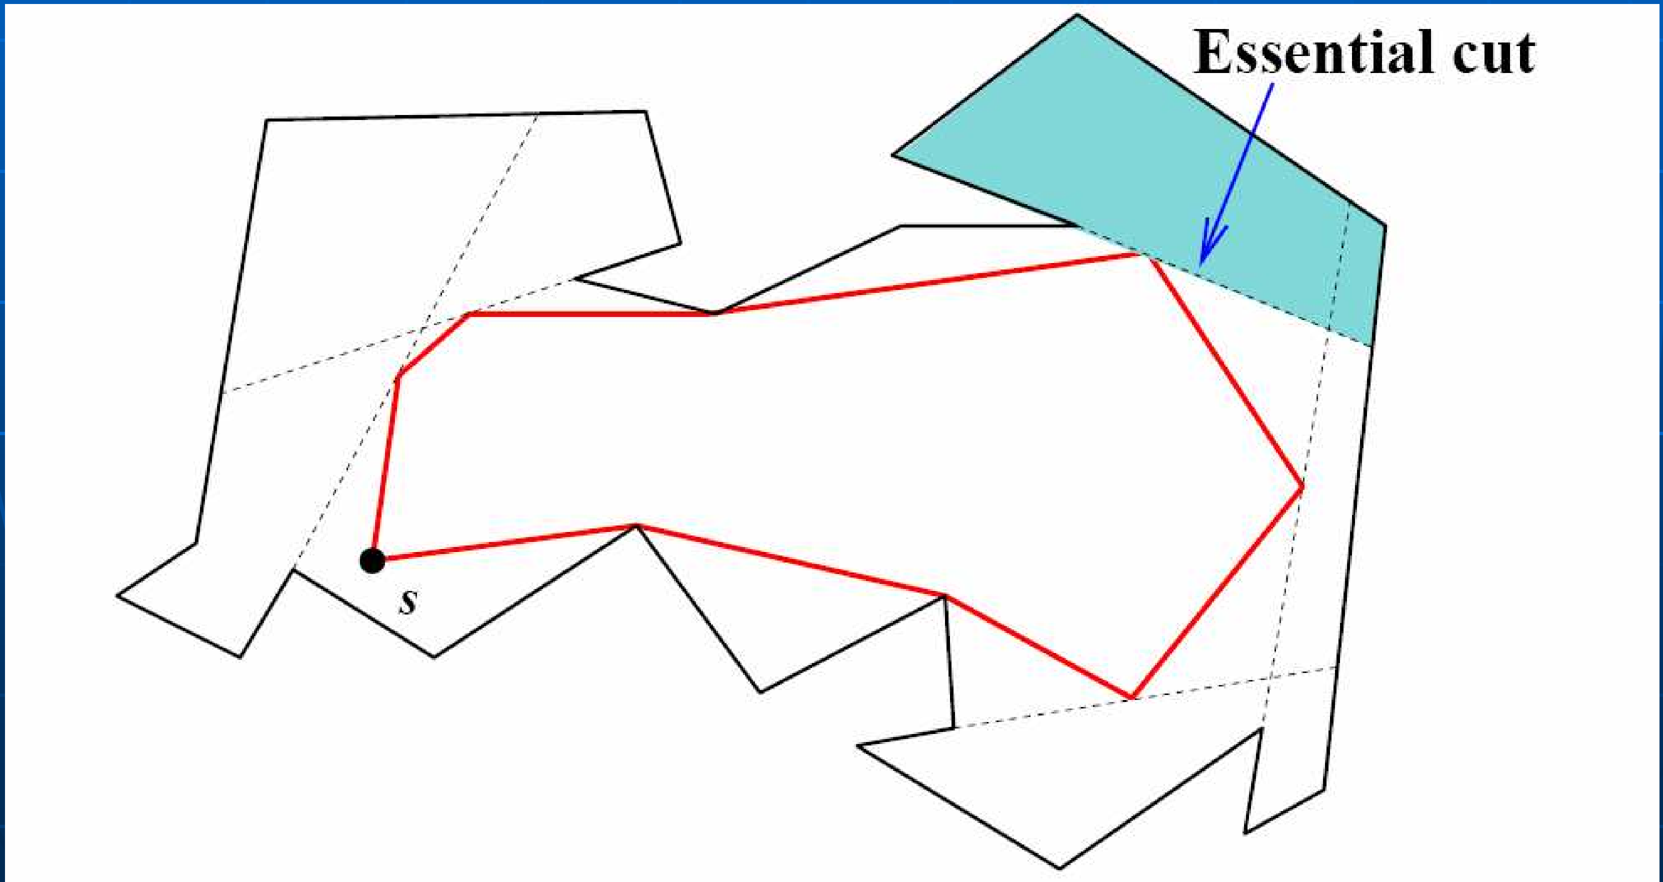

Fact: The optimal path visits the essential cuts in the order they appear along ∂P.

Watchman Route Problems

- Closely related to TSPN: visit VP(p), for all p in P

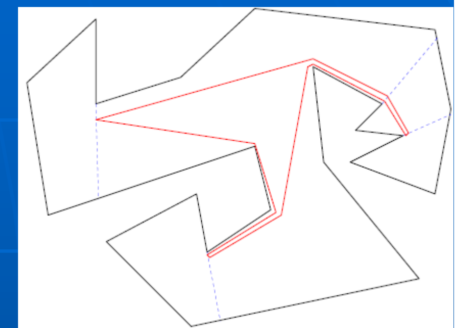

- Poly-time in simple polygons [CN,DELM]
  - Best time bound: $O(n^3 \log n)$ [DELM]
- NP-hard in polygons with holes

  - No approx algorithm known in general
  - Rectilinear visibility: O(log n)-approx [MM'95]
  - Progress:
    - PTAS for some fat obstacle cas
    - $\Omega(\log n)$-lower bound, in general
    - O(log n)-approx with a "bounded per assumption" New: general case
- 3D: Depends on 3D TSPN [ADDFM]

Q: Approx for planar domain, standard visibility?

Q: Approx for guard on a terrain surface?

Special Case

- Watchman on an arrangement of lines
  - Exact polytime algorithm (DP to search for CH)
  - Simpler 2-approx [Dumitrescu-Zylinski'11]

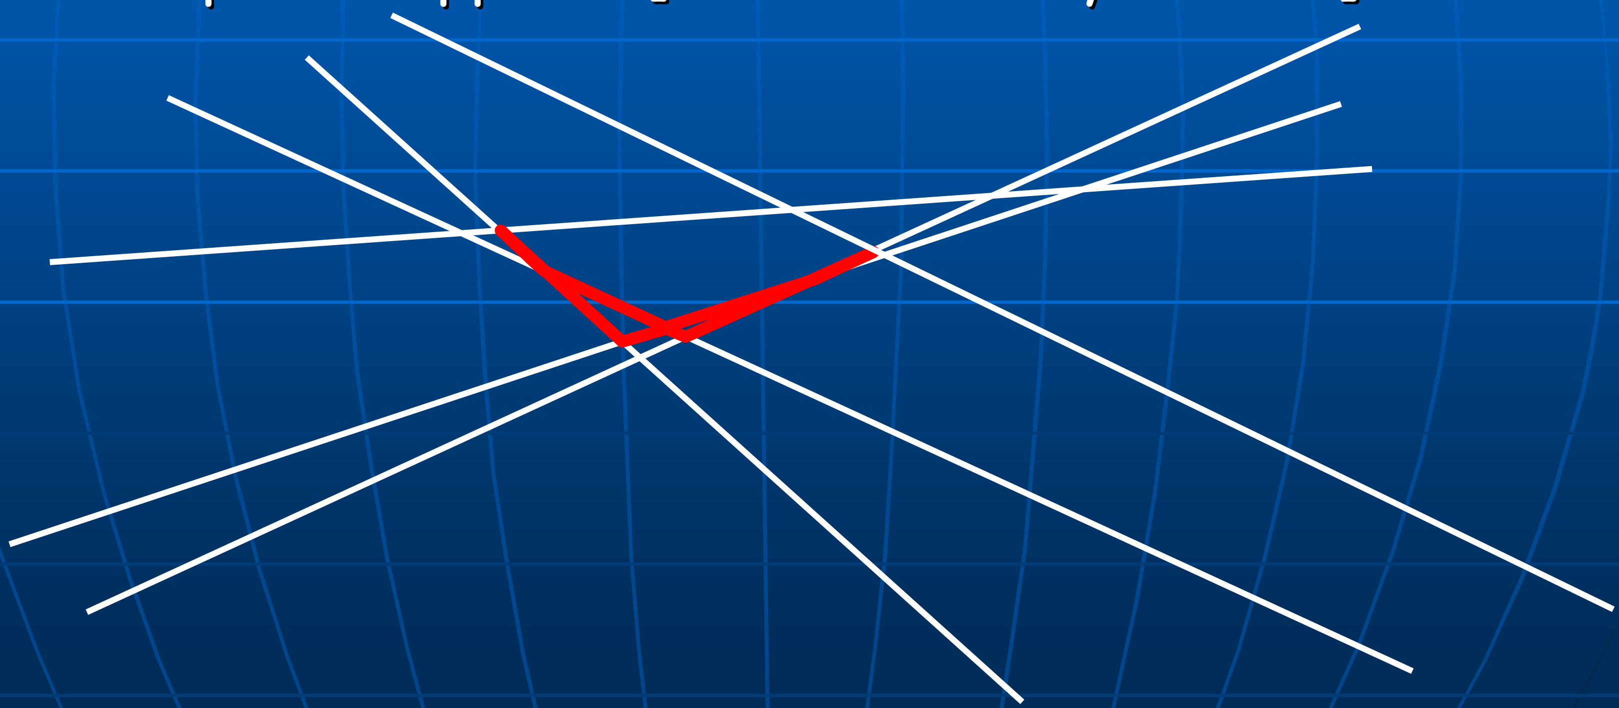

Special Case

- Thin polygons: PSLG's

"Frank's Problem"

NP-hard

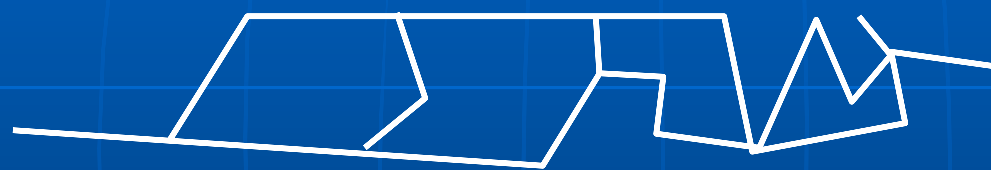

- polylog-approx using one-of-a-set TSP on sets of collinear vertices along straight paths

$O(\log^2 n \log\log n \log k)$-approx [CCGG, GKR]
1.5c-approx if straight corridors have < c vertices [Slavik]

- O(1)-approx if no straight corridors (collinear adjacent edges)

Connected vertex cover, [AHH]

Hardness of Approximation: Watchman Route Problem

- $\Omega(\log n)$: From Set-Cover: Sets S_1 , S_2 , …, S_M , and elements $U=\{x_1, x_2, …, x_N\}$

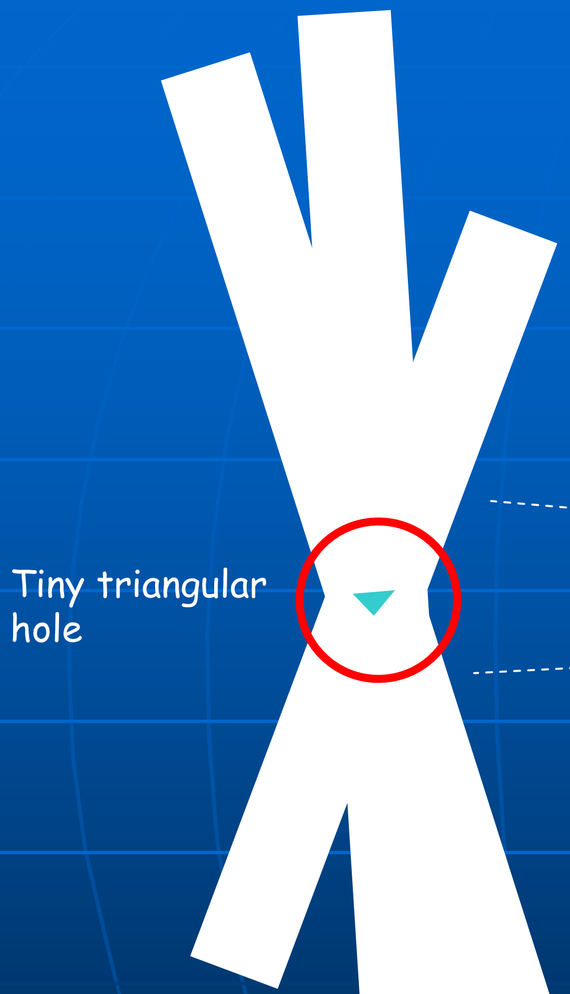

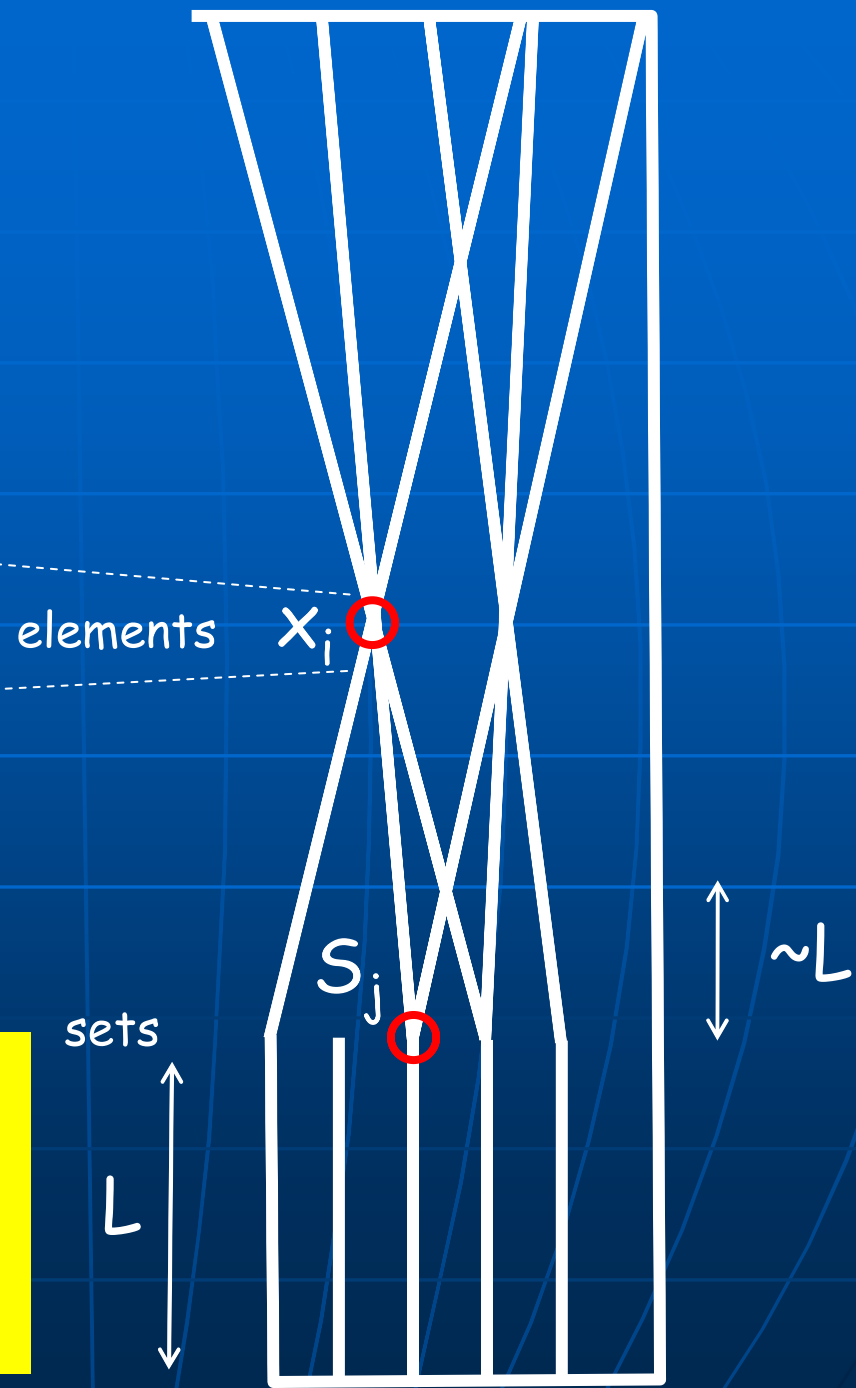

Can cover with k sets iff there is a watchman tour of length at most (2k+C)L, where L is very large

O(log n)-Approx Algorithm

- Input: Multiply connected polygonal domain P, having n vertices, satisfying the bounded perimeter assumption (BPA): perim(VP(p)) = O(diam(VP(p)), for every p in P

e.g., bounded degree corridor domains

- Compute the visibility polygons, Q_i = VP(v_i), of each vertex v_i

O(log n)-Approx Algorithm

- Consider the (simple, star-shaped) polygons Q_i in order, from smallest diameter to largest, and build a greedy, "well separated sequence", S, of such polygons
- Compute an approximately optimal tour, T_S, of the (disjoint) polygons S within P: TSPN with obstacles (P).

Visiting the Rest of the Regions, VP(p)

- So far, we have an approx tree, T_S, that visits the subset, S, of well-separated disjoint regions
- Remaining regions have special structure: Any point p that is not seen by T_S has VP(p) within geodesic distance diam(VP(p)) of T_S
- Generalize the "stratified grid" to geodesic metric within P

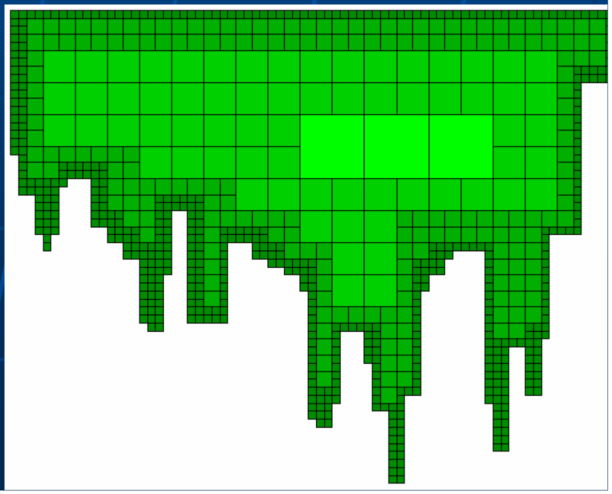

Special Case: Single Convex Obstacle

Sailor-in-the-Fog Problem

A sailor rows a mile out to sea, throws an anchor and a fishing line, and promptly falls asleep. By the time he wakes up, a dense fog has surrounded him. Knowing the distance to the shore but not knowing the direction, he wants to devise a path that is guaranteed to reach shore and that minimizes the distance travelled in the worst case. In other words, he would like to find the shortest curve that starts at the origin and intersects all lines at distance 1 from the origin. [Chan, Golynski, Lopez-Ortiz, Quimper, 2003]

Solution:
Length = 6.3972

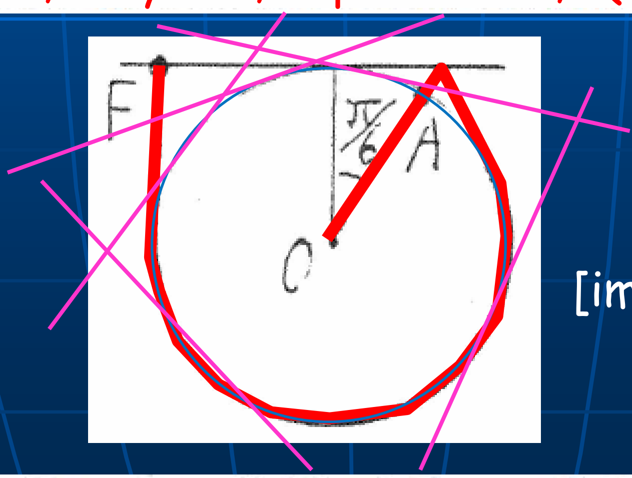

[image: Zalgaller, 2005]

J. R. Isbell, "An optimal search pattern," *Naval Research Logistics Quart.*, 4, 357–359 (1957)

Watchman Outside a Disk

- 2D paths: Depending if (a) neither, (b) both, or (c) exactly one endpoint of path must lie on the surface of the disk:

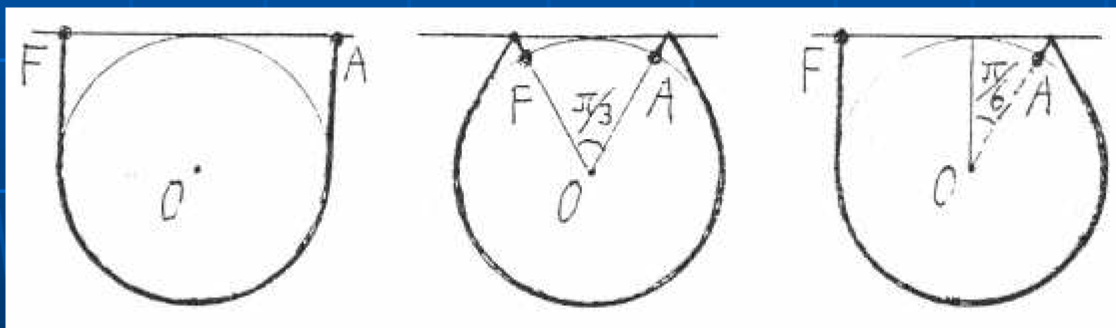

(a) (b) (c)

SHORTEST INSPECTION CURVES FOR THE SPHERE

V. A. Zalgaller* *Journal of Mathematical Sciences, Vol. 131, No. 1, 2005*

External Watchman Path for a Sphere

- Short Path
  Length 11.08

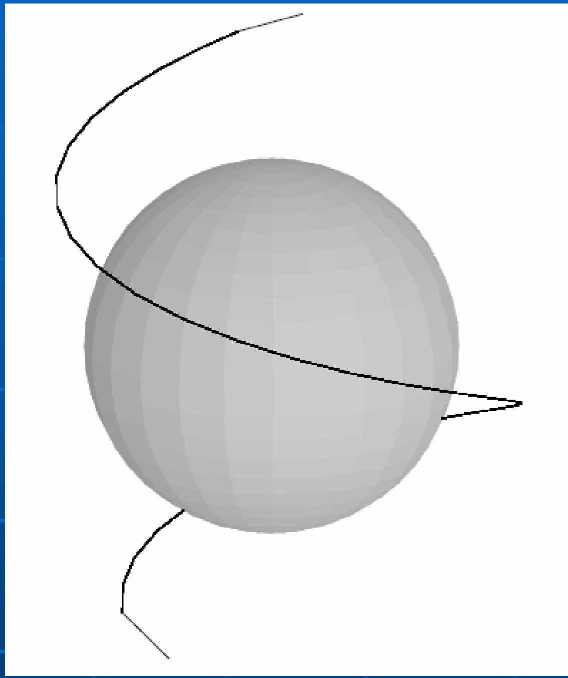

Two segments and a spiral:

$$\{((1-at^2)\sin(b\pi t),\ (1-at^2)\cos(b\pi t),\ ct) \mid -1 \le t \le 1\}$$

Fatten spiral near middle

$$a = 0.4,\ \ b = 1.18,\ c = 1.12,\ x_0 = -0.37,\ y_0 = -0.199,\ z_0 = 1.24$$

By computer search

The Asteroid Surveying Problem and Other Puzzles

Timothy M. Chan   Alexander Golynski   Alejandro López-Ortiz   Claude-Guy Quimper

[SoCG'03 video]

External Watchman Path for a Sphere

- Short Path
  Length 10.726

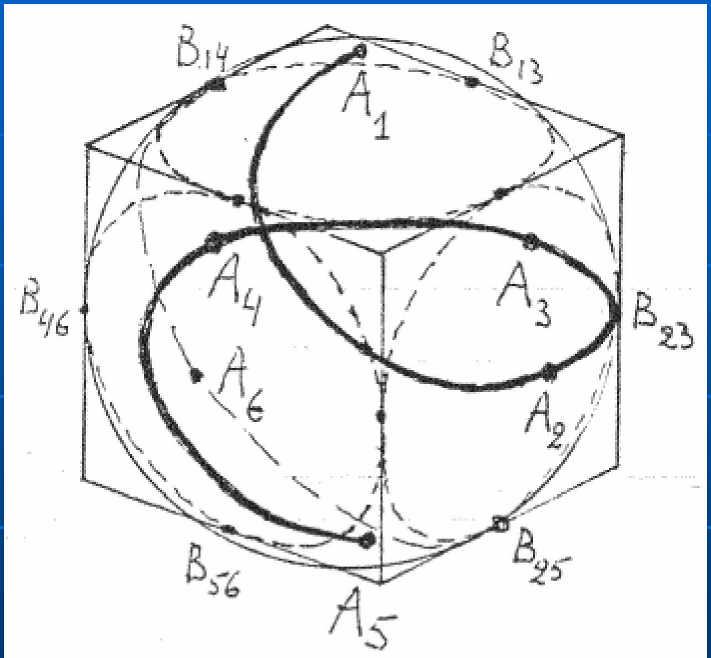

a "rather short" inspection curve that lies at the constant altitude of $\sqrt{2}-1$

$$L = \pi(2+\sqrt{2}) \approx 10.726$$

SHORTEST INSPECTION CURVES FOR THE SPHERE

V. A. Zalgaller* *Journal of Mathematical Sciences, Vol. 131, No. 1, 2005*

External Watchman Cycle for a Sphere

Shortest **Cycle** ?

"Shortest Inspection Curves for the Sphere"
V. A. Zalgaller

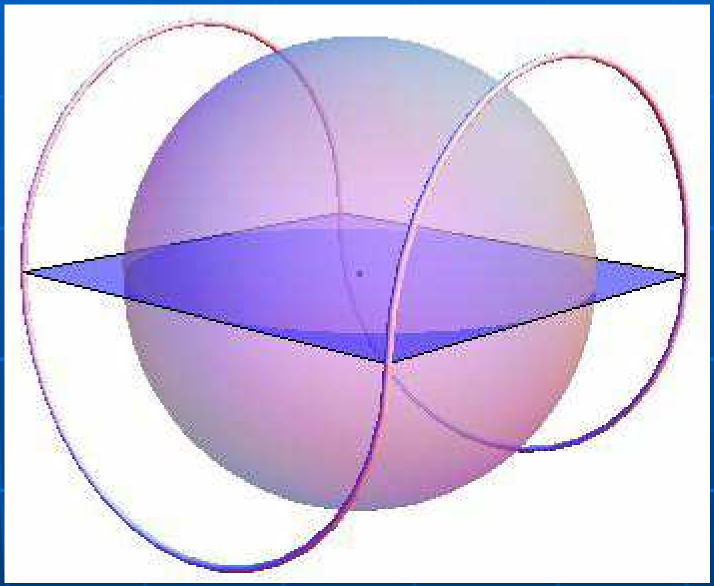

"baseball stitch curve"

[discussions: Jin-ichi Itoh, Joe O'Rourke, Anton Petrunin, Y. Tanoue, Costin Vilcu]

108 double stitches

Special Case: Single Convex Obstacle

- Watchman on the surface of a convex polytope in 3D:
  - Require to stay on the surface ("bug")
  - Allow to fly over the surface (but not through the obstacle) – "space ship"
  - Require to stay within a given altitude
- Exact solutions for special cases (e.g., Platonic solids) [Itoh, M, O'Rourke]
- PTAS in general: TSPN for a set of planes in 3D [Arkin, Demaine, Demaine, M]

Many Obstacles Special Case: Watchman PTAS

- Watchman with
  - Fat obstacles
  - Limited view distance, R
  - Robot of radius r, with R/r constant
- Method: m-guillotine

"Realistic domains"

Watchman PTAS

or "How to see the forest for the trees"

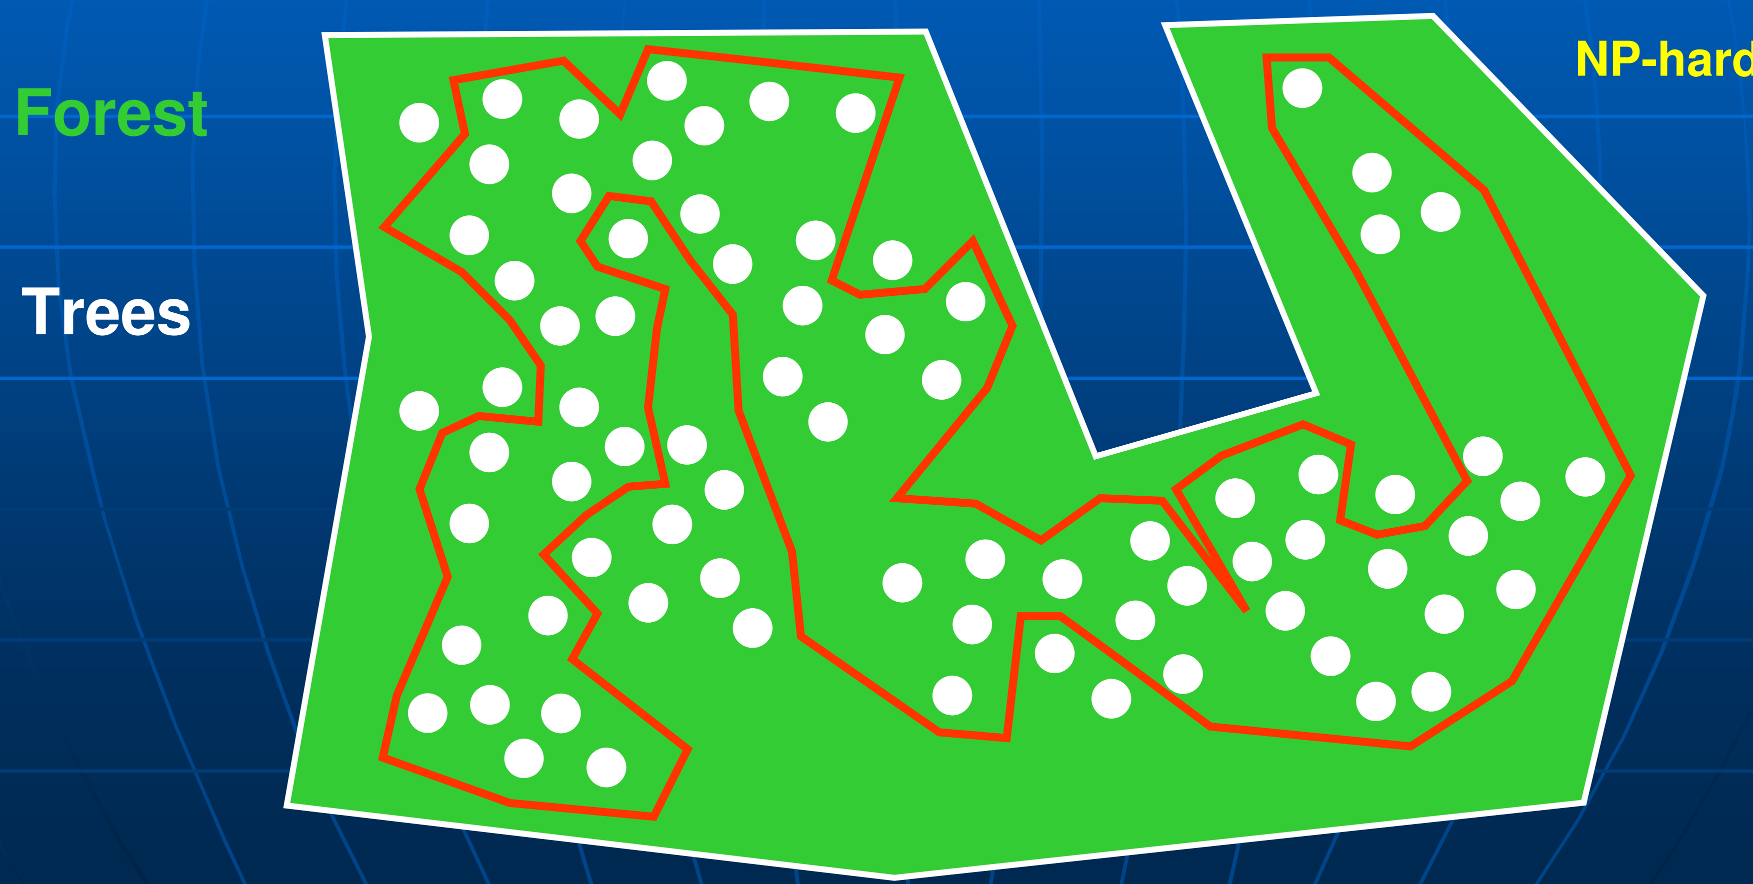

Forest Assumptions

Either: (1) limited view distance

Require robot to get within distance ***R*** of a point ***p*** in order to see it

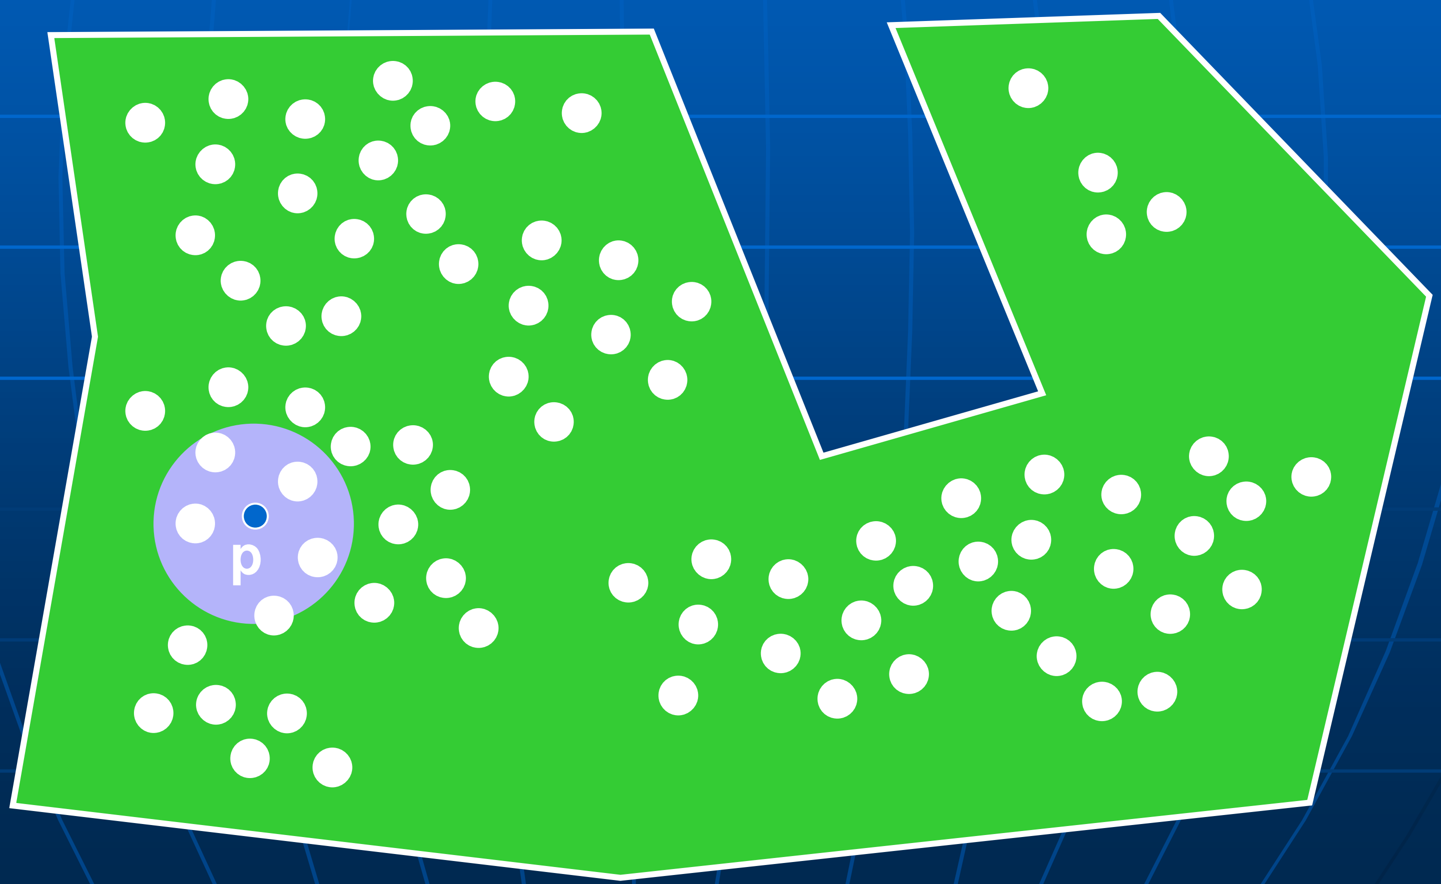

Forest Assumptions

Highest ever polynomial time bound?? (for a 2D problem)

Time: $O(n^{O(R)})$

Dark Forest Conjecture: For R < const, there exists a dark point p

Recently proved!: R < const [Dumitrescu and Jiang]

$R < 2*10^{108}$

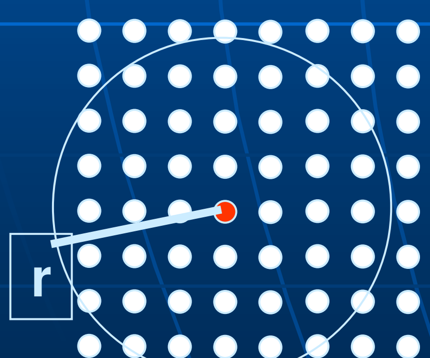

Dark if tree radius > 1/r

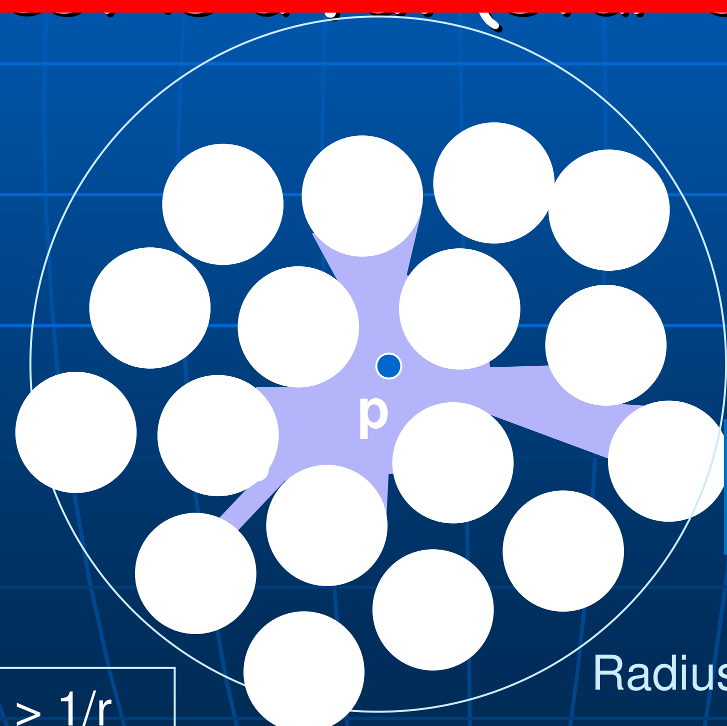

Related to Polya's Orchard Problem

Olber's paradox [1826]

Lawnmower/Milling Problem

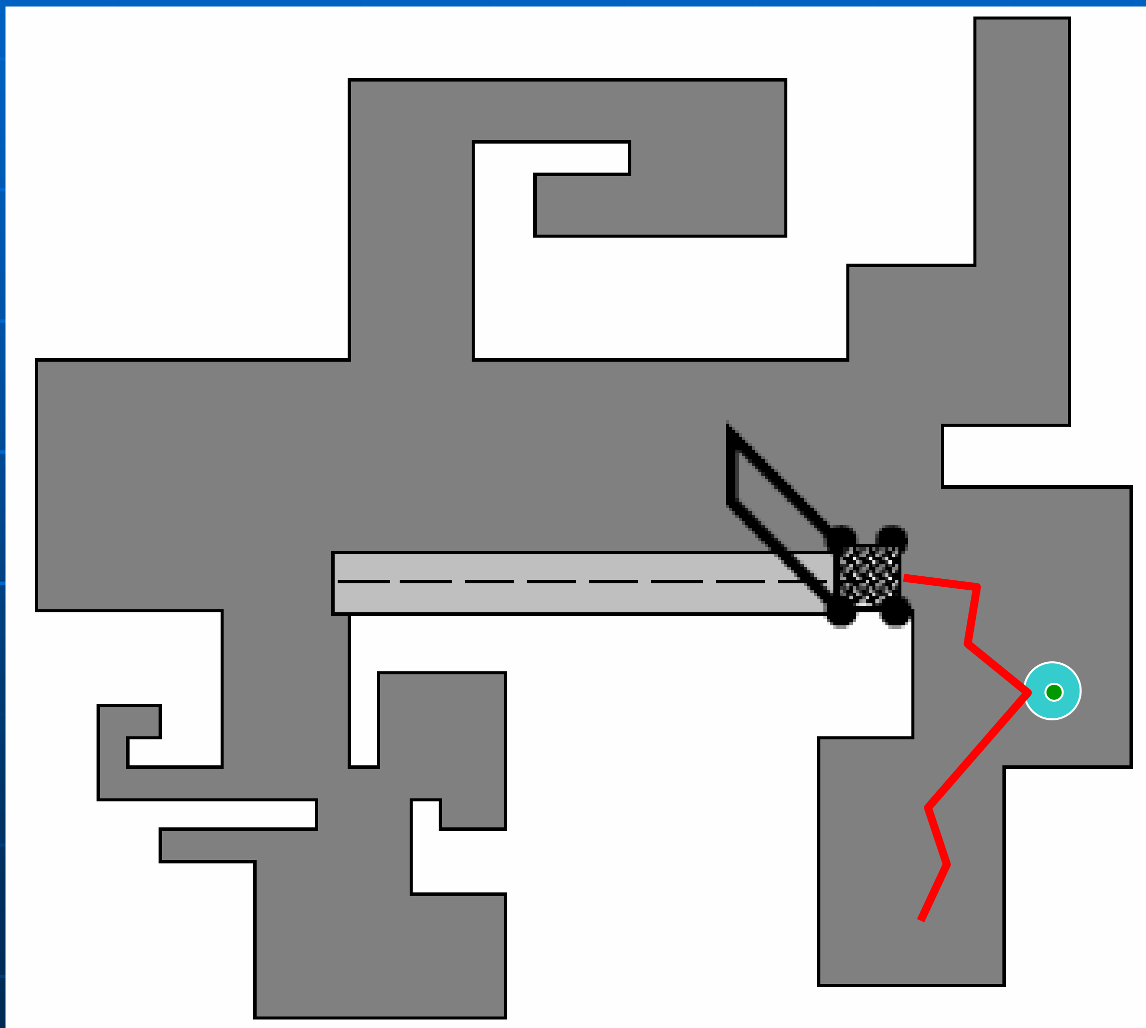

Best method of mowing the lawn?

TSPN: Visit the disk centered at each blade of grass

Lawnmowing/Milling Results

- NP-hard, in general
- 2.5-approx for milling (stay inside P) [AFM]
- (3+ε)-approx for lawnmowing [AFM]
  - Recent: PTAS
- 6/5-approx for d-sweepe
- 3.75-approx for min-turn milling in integral rectilinear polygon [ABDMS]
- PTAS to min (tour length)+C*(# turns)
- PTAS to min (tour length)+C*(# scans) "discrete vision cost" model [FMS]

Q: PTAS for milling a polygonal domain?

Q: Mill a **simple** polygon in poly-time?

Sensor Network Application: Cover Tour Problem

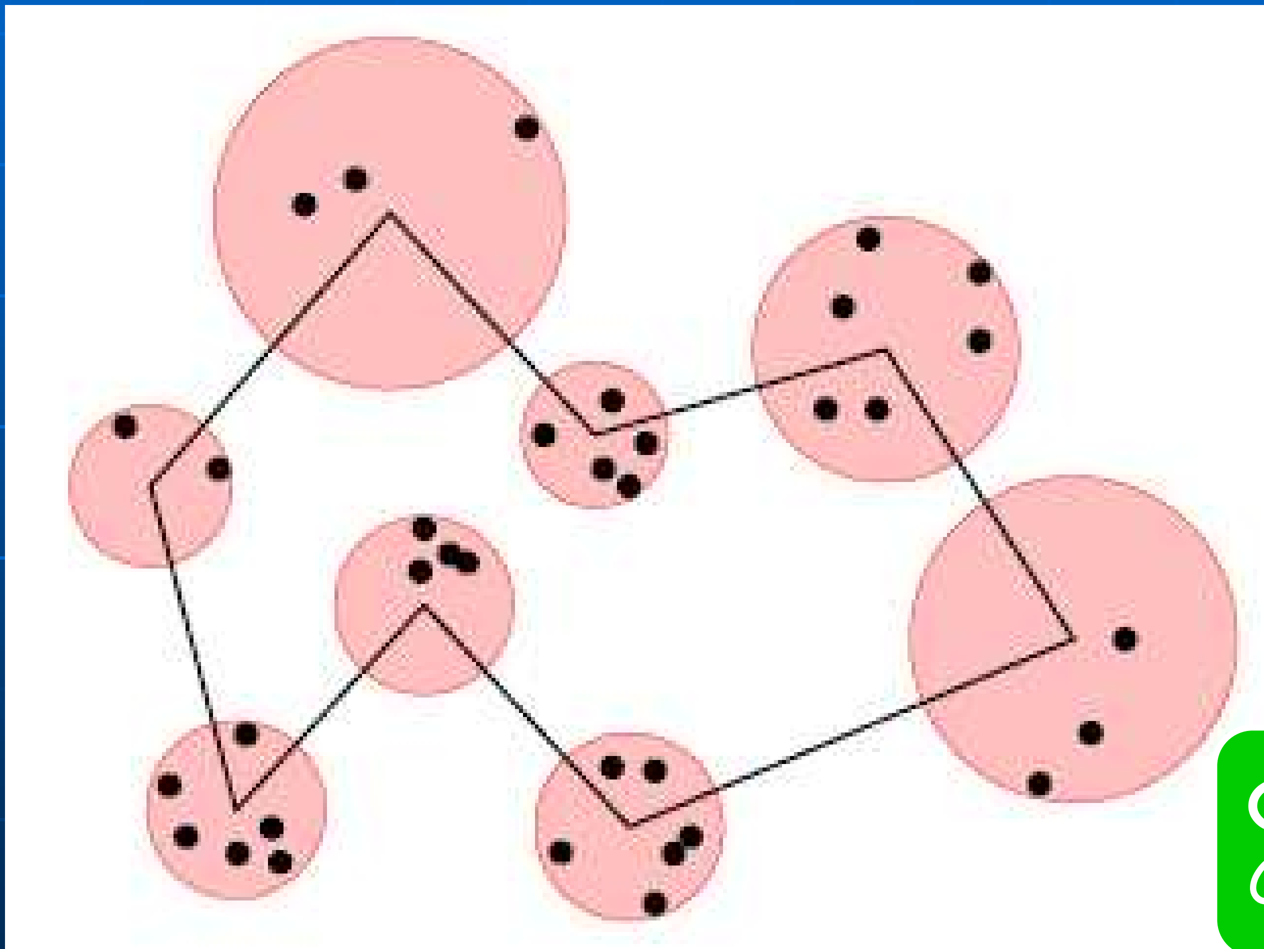

Min: Tour length + C * (sum of radii)

Result: PTAS

C > 4 ; else OPT is a single disk

Q: Min Tour length + C * (sum of radii^2)?

Alt, Arkin, Bronnimann, Erickson, Fekete, Knauer, Lenchner, M,Whittlesey, SoCG'06

Point Pair Separation

- Colored: n red, blue points
  - Find a shortest separator
  - PTAS
  - Covering tour:

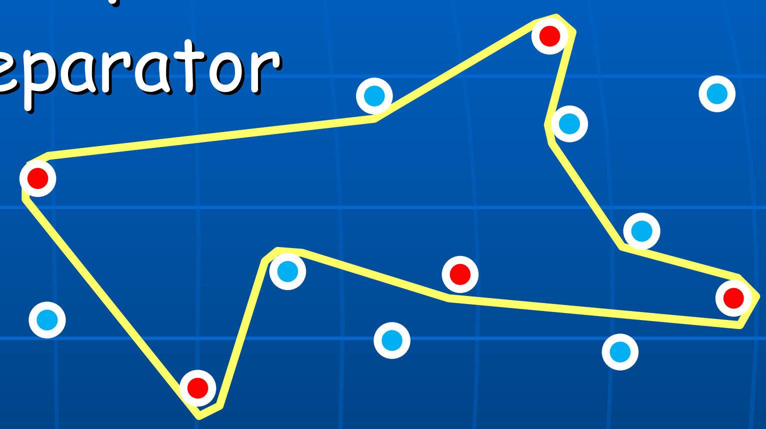

  - Visit/**cross** every edge of complete bipartite graph
- Uncolored: n point pairs (uncolored) [Jie Gao]
  - Find shortest cycle separating each pair

Q: Good approx?

Orienteering

- Given a length bound on tour, visit as many sites as possible
- O(1)-approx [AMN, SoCG'98]
- PTAS, for rooted case, based on improved analysis of m-guillotine method for k-TSP [CH, SoCG'06]

Q: O(n log n)?

Variations on the Classic TSP:

- Max TSP: max tour length
  - 5/7-approx in metric spaces
  - PTAS in $\Re^d$ for L_p metrics [Ba96]
  - $O(n^{f-2} \log n)$ for fixed d [Ba*98]

    (f = # facets for "disk")

    e.g., L_1 or L_∞ in $\Re^2$: $O(n^2 \log n)$
  - $O(n)$ for L_1 or L_∞ in $\Re^2$ [Fe98]
  - NP-hard for L_2 in $\Re^d$, $d \geq 3$ [Fe98]

OPEN: Complexity of Max TSP in Euclidean *plane*?

OPEN: Complexity of Max noncrossing TSP in Euclidean plane?

Variations on the Classic TSP (cont):

- bottleneck TSP: min the max edge length

  2-approx in metric spaces (best possible)

  (no $O(1)$-approx without $\Delta \neq$)

  NP-hard in Euclidean plane (Ham. cycle in grid graphs)

  OPEN: Better than 2-approx in E^2?

Variations on the Classic TSP (cont):

- max scatter TSP: max the min edge length

  NP-complete in metric spaces

  2-approx in metric spaces (best possible)

  (no $O(1)$-approx without $\Delta \neq$)

OPEN: Complexity of max scatter TSP in the plane?

OPEN: Better than 2-approx using geometry?

Variations on the Classic TSP (cont):

- minimum latency: "traveling repairman problem"

  Given starting point

  Goal: Min the sum of the arrival times at all other points

  NP-hard in E^2; 3.59-approx (10.78-approx, metric spaces)

  quasipoly-time approx scheme in $\Re^2$ $O(n^{O(\log n \log \log n/\epsilon^2)})$, [ArKa99]

  OPEN: Complexity of min latency in Euclidean plane?

  (metric version is MAX-SNP-hard)

- counting polygonalizations

  OPEN: Complexity?

Variations on the Classic TSP (cont):

- Kinetic TSP: points S are moving (known trajectories)

  Studied by Hammar and Nilsson [99]

- All velocites the same: PTAS

- Various velocites:

  No c-approx with $c < 2$, even if only two points move

  No approx factor better than $2^{\Omega(\sqrt{n})}$, even if the max velocity is bounded

Area Optimization:

Min-Area TSP (resp., Max-Area TSP):

find a *simple* tour on $S \subset \Re^2$ of min (max) area

Both problems are NP-complete [FP]

- Max-Area TSP

  (1/2)-approx in $O(n \log n)$ time (surround $\geq \frac{1}{2}$ area of $CH(S)$)

  (NP-complete to determine if $> \frac{2}{3} + \epsilon$ of area can be obtained)

- Min-Area TSP

  OPEN: Is there a poly-time approx algorithm for Min-Area TSP?

  (none for min-area disjoint triangle matching on $3n$ points)