Max-Flow Min-Cut Theorem and Faster Algorithms in a Circular Disk Failure Model

Yusuke Kobayashi
University of Tokyo
Tokyo, 113-8656, Japan
kobayashi@mist.i.u-tokyo.ac.jp

Kensuke Otsuki
University of Tokyo
Tokyo, 113-8656, Japan
kensuke_otsuki@mist.i.u-tokyo.ac.jp

***Abstract*—Fault-tolerance is one of the most important factors in designing networks. Failures in networks are sometimes caused by an event occurring in specific geographical regions such as hurricanes, earthquakes, bomb attacks, and Electromagnetic Pulse (EMP) attacks. In INFOCOM 2012, Neumayer et al. introduced geographical variants of max-flow min-cut problems in a circular disk failure model, in which each failure is represented by a disk with a predetermined size. In this paper, we solve two open problems in this model: we give a first polynomial-time algorithm for the geographic max-flow problem, and prove a conjecture of Neumayer et al. on a relationship between the geographic max-flow and the geographic min-cut.**

I. Introduction

Fault-tolerance is one of the most important factors in designing networks. In most studies on fault-tolerance in networks, "connectivity" of the network is regarded as the measure of robustness (e.g. [1]–[3]). However, failures in networks are sometimes caused by an event occurring in specific geographical regions such as hurricanes, earthquakes, bomb attacks, and Electromagnetic Pulse (EMP) attacks. Recently, some models in which such localized failures are taken into consideration are proposed in [4]–[10].

In particular, Neumayer et al. [5] considered the model in which each failure is represented by a hole (disk) with a predetermined size (see Section II for details), and they gave a polynomial-time algorithm for computing the minimum number of failures that disconnect two specified nodes s and t, which they call the "geographic min-cut". They also formulated the problem of finding the maximum number of s-t paths such that no two paths can be disconnected by the same hole, which they call the "geographic max-flow". Similar problems are considered in [11], [12].

In this paper, we further develop theory and algorithms for geographic min-cut and geographic max-flow in this model. Our contributions are described as follows.

- We give a min-max theorem that characterizes a geographic max-flow (Theorem 1).
- We show that the geographic min-cut is at most the geographic max-flow plus one (Theorem 3), which was conjectured by Neumayer et al.
- We give a first polynomial-time algorithm for the geographic max-flow (Theorem 4).
- We give a polynomial-time algorithm for the geographic min-cut which is simpler (and probably faster) than known algorithms (Theorem 5).
- We implement our algorithms and confirm that they can solve large problems efficiently.

We emphasize here that the second and the third results solve two open problems raised by Neumayer et al. [5] in INFOCOM 2012.

We also note that we can extend our results to the case when each hole is of different shapes. See Section VI for details.

This paper is organized as follows. First, we describe formal problem settings in Section II. Next, in Section III, we discuss min-max relations between the geographic min-cut and the geographic max-flow, and prove the conjecture of Neumayer et al. In Section IV, we give algorithms for the geographic max-flow and the geographic min-cut based on the min-max relation. Then, experimental results are shown in Section V. Finally, we give concluding remarks in Section VI.

II. Problem Settings

Let $G = (V, E)$ be a graph drawn in the plane with a node set V, a link set E, and two distinct nodes $s, t \in V$, where each link is drawn as a line segment. In what follows, a link is sometimes called an edge. Let r_b be a hole radius and $r_p (> r_b)$ be a protection radius. A disk of radius r_p whose center is s or t is called a *protective disk*. Define $\mathcal{H}(r_b, r_p)$ as the set of all disks of radius r_b whose centers are not contained in protective disks of radius r_p. Each element of $\mathcal{H}(r_b, r_p)$ is called a *hole* in this paper.

We consider a geographic variant of the min-cut problem, which is defined as follows.

Geographical Min-Cut by Circular Disasters (GMCCD)

Input: a graph $G = (V, E)$ drawn in the plane, two distinct nodes s and t, a hole radius r_b, and a protection radius $r_p (> r_b)$.

Find: a minimum cardinality set of holes in $\mathcal{H}(r_b, r_p)$ that disconnect s from t.

Let MIN-CUT denote the optimal value of this problem. A set of holes in $\mathcal{H}(r_b, r_p)$ that disconnect s from t is called a *hole cut* in this paper. We can also consider a geographic variant of the max-flow problem.

978-1-4799-3360-0/14/$31.00 ©2014 IEEE

978-14799-3360-0/14/$31.00 ©2014 IEEE

Geographical Max-Flow by Circular Disasters (GMFCD)

Input: a graph $G = (V, E)$ drawn in the plane, two distinct nodes s and t, a hole radius r_b, and a protection radius $r_p (> r_b)$.

Find: a maximum cardinality set of s-t paths such that no hole in $\mathcal{H}(r_b, r_p)$ intersects a pair of these paths.

Let MAX-FLOW denote the optimal value of this problem. So far, no polynomial-time algorithm for GMFCD was known, whereas a polynomial-time algorithm for GMCCD was given in [5].

Example 1. Consider the graph as in Fig. 1 (which is an example given in [5]). Small circles represent holes of radius r_b and two large shaded (or yellow in the color version) circles are protective disks. In the graph, we can easily see that MAX-FLOW = 1 and MIN-CUT = 2.

III. Geographic Max-Flow Min-Cut Theorem

In this section, we investigate relations between MAX-FLOW and MIN-CUT. First we give a characterization of maximum flows in Sections III-A and III-B. Then, in Section III-C, we prove MIN-CUT $\leq$ MAX-FLOW + 1, which is the conjecture of Neumayer et al. Note that this bound is tight by Example 1, and it significantly improves previously known bound: MIN-CUT $\leq 2 \cdot$ MAX-FLOW + 2 given in [11].

A. Statement of the theorem

Let C be a closed curve in the plane that does not go through s or t. We define the *winding number* $w(C)$ of C as the number of times that C separates s and t. More precisely, let L be the line segment connecting s and t, and fix orientations of L and C. Let $w_1(C)$ be the number of times that C crosses L from left to right and $w_2(C)$ be the number of times that C crosses L from right to left. Then, define $w(C) := |w_1(C) - w_2(C)|$.

We say that a closed curve C can be represented as an *alternating curve of length* l if it is a concatenation of $J_1, L_1, J_2, L_2, \dots, J_l, L_l$ in this order such that

- $J_1, J_2, \dots, J_l$ are curves each contained in a face of G, and
- for each $i = 1, 2, \dots, l$, L_i is a line segment that can be covered by a hole in $\mathcal{H}(r_b, r_p)$.

Note that J_i and L_i might be a single point. For a closed curve C, let $l(C)$ be the minimum number l such that C can be represented as an alternating curve of length l.

By using these notations, we can give an exact min-max theorem for MAX-FLOW, whose proof is given in the next subsection.

Theorem 1. *Suppose we are given a graph $G = (V, E)$ drawn in the plane, two distinct nodes s and t, a hole radius r_b, and a protection radius $r_p (> r_b)$. Then,*

$$\text{MAX-FLOW} = \min \left\{ \left\lfloor \frac{l(C)}{w(C)} \right\rfloor \mid C \textit{ is a closed curve} \right\}.$$

Example 2. Consider again the graph as in Fig. 1. As shown in Fig. 2, there exists a closed curve C such that $l(C) = 3$, $w(C) = 2$, and hence $\left\lfloor \frac{l(C)}{w(C)} \right\rfloor = 1$. This value is equal to MAX-FLOW.

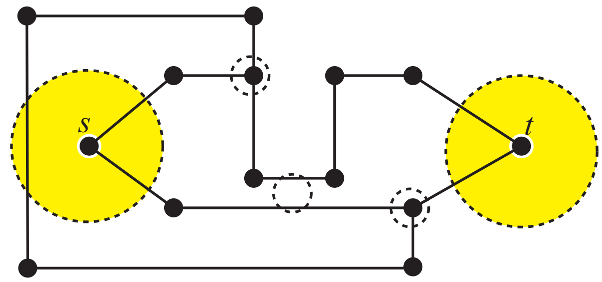

Fig. 1. Example of the problems

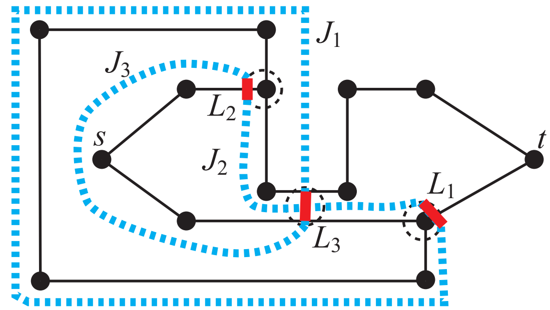

Fig. 2. Example of a closed curve C

B. Proof of Theorem 1

In this subsection, we give a proof of Theorem 1. Our proof is based on ideas in [13] (see also [14], [15]), which shows a min-max theorem for maximum induced disjoint s-t paths in plane graphs.

We say that two s-t paths are *separated*, if no hole in $\mathcal{H}(r_b, r_p)$ intersects both of these paths. For a pair of edges $e, e' \in E$, if there exists a hole in $\mathcal{H}(r_b, r_p)$ that intersects both edges, then we take two points $w_{e,e'}$ on e and $w_{e',e}$ on e' that are contained in the common hole arbitrarily. Define

$$W := \{w_{e,e'}, w_{e',e} \mid e, e' \in E \text{ contained in a common hole}\}.$$

Let $\mathcal{L}$ be the set of all line segments with both endpoints in W such that each line segment is contained in a hole in $\mathcal{H}(r_b, r_p)$. Note that a line segment might be a single point, that is, $(w, w) \in \mathcal{L}$ for $w \in W$. Then, two s-t paths P and P' are not separated if and only if there exists a line segment $(w, w') \in \mathcal{L}$ such that w is on P and w' is on P'. Theoretically, $|W|$ is bounded by $|E|^2$ which is a polynomial size. Practically, we can obtain W by adding a small number of points to V, because we have to add points to V only in some exceptional cases (see e.g. Fig. 3).

First, we show

$$\text{MAX-FLOW} \leq \min \left\{ \left\lfloor \frac{l(C)}{w(C)} \right\rfloor \mid C \text{ is a closed curve} \right\}.$$

Suppose we have s-t paths $P_1, \dots, P_k$ that are pairwise separated and let C be a closed curve that is a concatenation of $J_1, L_1, J_2, L_2, \dots, J_{l(C)}, L_{l(C)}$ in this order. Since C intersects

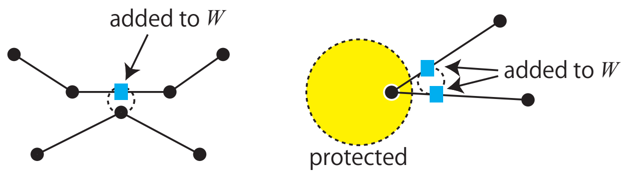

Fig. 3. Construction of W. In these cases, two edges can be covered by a common hole, but their endnodes cannot be covered by a common hole. In such cases, we need to add some points to W.

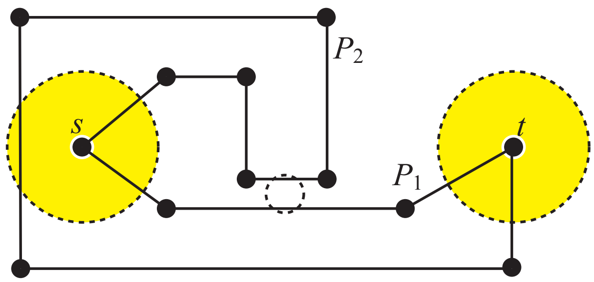

Fig. 4. A pair (P_1, P_2) is clockwise separated. Since there exists a hole dividing $R(P_2, P_1)$, a pair (P_2, P_1) is not clockwise separated.

each P_i at least $w(C)$ times, each P_i intersects at least $w(C)$ line segments of $L_1, L_2, \ldots, L_{l(C)}$. This means that $l(C) \geq k \cdot w(C)$. By the integrality of k, we have $k \leq \left\lfloor \frac{l(C)}{w(C)} \right\rfloor$.

Next we show

$$\text{MAX-FLOW} \geq \min\left\{ \left\lfloor \frac{l(C)}{w(C)} \right\rfloor \mid C \text{ is a closed curve} \right\}$$

by giving an algorithm for finding either k pairwise separated s-t paths or a closed curve C with $\frac{l(C)}{w(C)} < k$ for any k. For any s-t path P, we suppose that it is oriented from s to t and P^{-1} is oriented from t to s. For two s-t paths P' and P'' without crossings, let $R(P', P'')$ denote the closed region encircled by $P' \cdot (P'')^{-1}$ in clockwise orientation, where $P' \cdot (P'')^{-1}$ is the closed curve obtained by concatenating P' and $(P'')^{-1}$. For two s-t paths P' and P'' without crossings, a pair (P', P'') is *clockwise separated* if for any hole H in $\mathcal{H}(r_b, r_p)$, $R(P', P'') - H$ is connected. Obviously, P' and P'' are separated if and only if a pair (P', P'') and a pair (P'', P') are both clockwise separated (see Fig. 4 for an example). In what follows, we show the inequality by the induction on k.

1) Induction step: First, we consider the case $k \geq 3$ under the assumption that we have $k-1$ pairwise separated s-t paths $P_1, \ldots, P_{k-1}$. We may assume that these paths do not cross each other, and the first edges of $P_1, \ldots, P_{k-1}$ occur in this order clockwise at s. Let P_k be an s-t path in $R(P_{k-1}, P_1)$ such that (P_{k-1}, P_k) is clockwise separated[1]. In our algorithm, we start with the k paths $P_1, \ldots, P_{k-1}, P_k$, and we replace one path with a new path, repeatedly. Our algorithm is described in Algorithm 1 (see Fig. 5 for an example).

If (P_l, P_{l-k+1}) is clockwise separated in line 2 of Algorithm 1, then $P_{l-k+1}, \ldots, P_{l-1}, P_l$ are pairwise sepa-

[1]This condition is equivalent to "separated" when $k \geq 3$. Since we will use the same argument for the case of $k = 2$ later, we use the condition "clockwise separated" here.

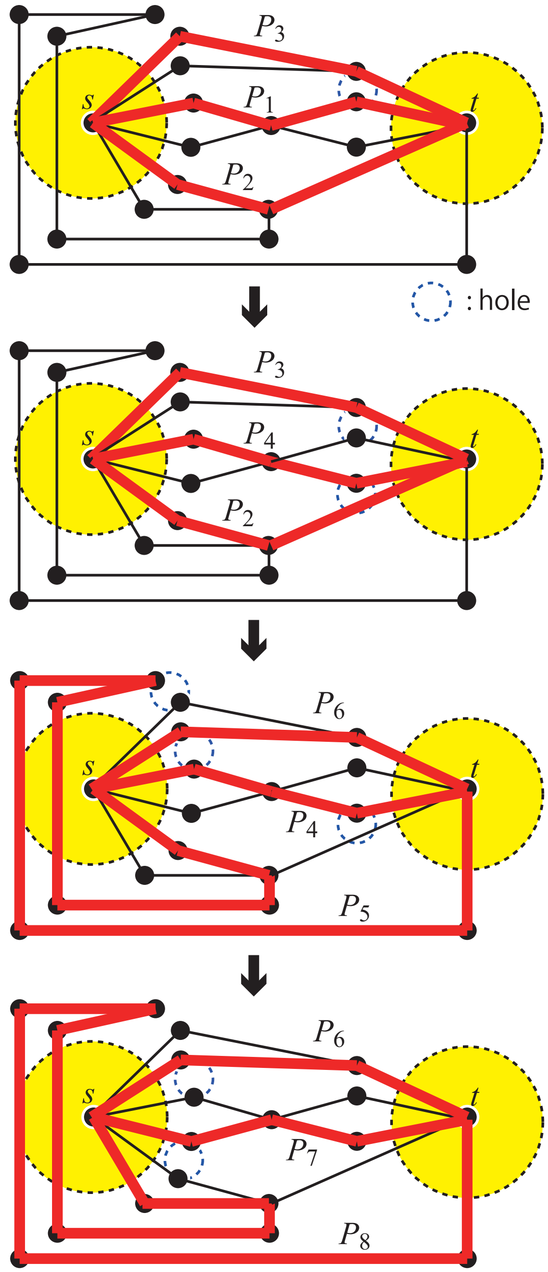

Fig. 5. Example of iterations in Algorithm 1 ($k = 3$). We begin with three s-t paths P_1, P_2, and P_3 such that (P_1, P_2) and (P_2, P_3) are clockwise separated, but (P_3, P_1) is not clockwise separated (the first figure). In the first iteration, we take an s-t path P_4 such that (P_3, P_4) is clockwise separated and $R(P_4, P_2)$ is maximized under this condition (the second figure). Since (P_4, P_2) is not clockwise separated, we take P_5 such that (P_4, P_5) is clockwise separated. Then, since (P_5, P_3) is not clockwise separated, we take P_6 such that (P_5, P_6) is clockwise separated (the third figure). Similarly, we take P_7 and P_8. Since (P_8, P_6) is clockwise separated, P_6, P_7, and P_8 are desired paths.

Algorithm 1 Induction step

Input: pairwise separated $k-1$ s-t paths $P_1, \dots, P_{k-1}$ and an s-t path P_k in $R(P_{k-1}, P_1)$ such that (P_{k-1}, P_k) is clockwise separated

Output: pairwise separated k s-t paths or a closed curve C with $\frac{l(C)}{w(C)} < k$

```
1: for l = k, k+1, ..., k+|W|+1 do
2:   if (P_l, P_{l-k+1}) is clockwise separated[1] then
3:     return P_{l-k+1}, ..., P_{l-1}, P_l that are separated paths
4:   else
5:     let P_{l+1} be the s-t path in R(P_{l-k+1}, P_{l-k+2})
       such that (P_l, P_{l+1}) is clockwise separated and
       R(P_{l+1}, P_{l-k+2}) is maximized under this condition
6:   end if
7: end for
8: return a closed curve C with l(C)/w(C) < k
```

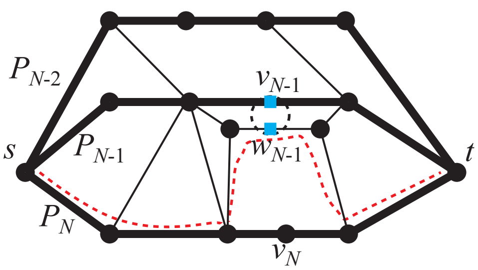

Fig. 6. Definition of w_{N-1} and v_{N-1}. If we reroute P_N along the border of a face containing v_N, then the obtained path contains a node w_{N-1} which is close to a node v_{N-1} on P_{N-1}. Therefore, $(w_{N-1}, v_{N-1}) \in \mathcal{L}$, and there exists a curve from v_N to w_{N-1} contained in a face of G.

rated, because (P_i, P_{i+1}) is clockwise separated for $i = l-k+1, \dots, l-1$. Therefore, it suffices to return $P_{l-k+1}, \dots, P_{l-1}, P_l$. In what follows, we give a procedure for finding a closed curve C with $\frac{l(C)}{w(C)} < k$ (line 8) when such paths do not appear while $l = k, k+1, \dots, k+|W|+1$.

Let $N := k + |W| + 1$. By the assumption, (P_N, P_{N-k+1}) is not clockwise separated, and hence there exists a node $v_N \in P_N \setminus P_{N-k}$. Note that a path is regarded as a subset of the plane. Since P_N maximizes $R(P_N, P_{N-k+1})$, if we reroute P_N so that the obtained path does not go through v_N, then the path contains a node close to P_{N-1}. More precisely, as in Fig. 6, we can find a pair of nodes $w_{N-1} \in W$ and $v_{N-1} \in P_{N-1} \cap W$ such that $(w_{N-1}, v_{N-1}) \in \mathcal{L}$ (i.e., w_{N-1} and v_{N-1} are covered by a common hole in $\mathcal{H}(r_b, r_p)$) and v_N and w_{N-1} can be connected by a curve J_N contained in a face of G. Furthermore, we can see that $v_{N-1} \notin P_{N-k-1}$, because w_{N-1} is strictly to the right of P_{N-k} (when we walk from s to t along P_{N-k}). By repeating the same argument, we can find v_i, w_i, and J_i for $i = N-1, N-2, \dots, k+1$ such that

- $w_i \in W$ and $v_i \in (P_i \setminus P_{i-k}) \cap W$ with $(w_i, v_i) \in \mathcal{L}$, and
- J_i is a curve from v_i to w_{i-1} contained in a face of G.

By pigeonhole principle, $v_i = v_j$ for some $k+1 \le i < j \le N$. Let C be a closed curve obtained by concatenating

$$(v_i, w_i), J_{i+1}, (v_{i+1}, w_{i+1}), J_{i+2}, \dots, (v_{j-1}, w_{j-1}), J_j$$

in this order, where (x, y) is the line segment connecting x and y. We will show that this curve C satisfies $\frac{l(C)}{w(C)} < k$, which is equivalent to $u := \left\lfloor \frac{j-i}{k} \right\rfloor < w(C)$, because $l(C) = j - i$. If $u = 0$, then the inequality is trivial. Otherwise, v_j is strictly to the right of P_{j-k}. When we consider a curve from v_i on P_i to v_{j-k} on P_{j-k} along C, it separates s and t at least $u-1$ times, because $j - k \ge i + (u-1)k$. Therefore, C separates s and t more than u times, that is, $w(C) > u$.

By the above procedure, we can find a closed curve C with $\frac{l(C)}{w(C)} < k$ in line 8 of Algorithm 1.

2) Base cases: Next we deal with the base cases ($k = 1, 2$) of the induction. Since the case of $k = 1$ is trivial, we consider the case when $k = 2$. We say that a closed curve C *separates* s and t if every curve connecting s and t intersects with C. We use the following lemma, whose proof is given in the appendix.

Lemma 2. *Suppose that s and t are connected in G. We can find in polynomial time either*

- *a hole in $\mathcal{H}(r_b, r_p)$ that disconnects s from t, or*
- *an s-t path P in G such that for any line segment $L \in \mathcal{L}$, the union of P and L (which is regarded as a subset of the plane) contains no closed curve separating s and t.*

If we have a hole in $\mathcal{H}(r_b, r_p)$ that disconnects s from t, then we have a closed curve C with $l(C) = w(C) = 1$. Otherwise, we have an s-t path P in G as in the second case of Lemma 2. In this case, define $P_1 = P_2 = P$ and assume that P_2 is to the left of P_1. Then, since (P_1, P_2) is clockwise-separated, we can apply Algorithm 1 to obtain pairwise separated two s-t paths or a closed curve C with $\frac{l(C)}{w(C)} < 2$, which completes the proof for the case of $k = 2$.

C. Proof of the conjecture

By using Theorem 1, in this subsection we give a proof of the conjecture of Neumayer et al.

Theorem 3. *Suppose we are given a graph $G = (V, E)$ drawn in the plane, two distinct nodes s and t, a hole radius r_b, and a protection radius $r_p (> r_b)$. Then,*

$$\text{MAX-FLOW} \le \text{MIN-CUT} \le \text{MAX-FLOW} + 1.$$

Proof: Since MAX-FLOW $\le$ MIN-CUT is obvious, we prove MIN-CUT $\le$ MAX-FLOW $+ 1$. By Theorem 1, we can take a closed curve C such that $\left\lfloor \frac{l(C)}{w(C)} \right\rfloor =$ MAX-FLOW. Hence, it suffices to find a hole cut of size $\left\lfloor \frac{l(C)}{w(C)} \right\rfloor + 1$ (i.e., a set of $\left\lfloor \frac{l(C)}{w(C)} \right\rfloor + 1$ holes in $\mathcal{H}(r_b, r_p)$ that disconnect s from t).

If $w(C) \ge 2$, then C must contain a self-crossing and we can decompose C into two closed curves C_1 and C_2 by uncrossing procedures (see Figs 7, 8, and 9). Obviously, $w(C_1) + w(C_2) = w(C)$. To evaluate the total length $l(C_1) + l(C_2)$, we consider the following three cases.

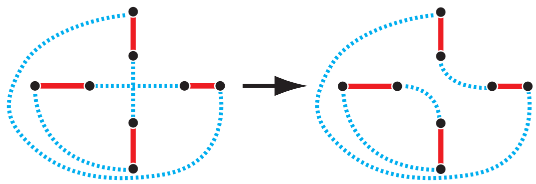

Fig. 7. Uncrossing procedure (case 1). Two curves J_i and J_j are replaced with two new curves.

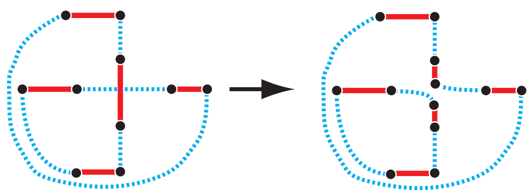

Fig. 8. Uncrossing procedure (case 2). A curve J_i and a line segment L_j are replaced with two line segments and two curves.

1) If two curves J_i and J_j are crossing, then we can easily uncross C without increasing the length, that is, $l(C_1) + l(C_2) = l(C)$ (see Fig. 7).
2) If a curve J_i and a line segment L_j are crossing, then we can uncross C by using two line segments instead of L_j, that is, $l(C_1)+l(C_2) \le l(C)+1$ (see Fig. 8).
3) Suppose that two line segments L_i and L_j are crossing. Then, the hole containing L_i also contains an endnode of L_j or the hole containing L_j also contains an endnode of L_i. Therefore, we can uncross C by using at most three line segments instead of L_i and L_j, that is, $l(C_1) + l(C_2) \le l(C) + 1$ (see Fig. 9).

In each case, we have $l(C_1) + l(C_2) \le l(C) + 1$. By repeating uncrossing procedures, we have closed curves $C_1, C_2, \dots, C_{w(C)}$ such that $w(C_i) = 1$ for each i and $\sum_i l(C_i) \le l(C) + w(C)$. Since we have

$$\min_i\{l(C_i)\} \le \left\lfloor \frac{1}{w(C)} \sum_i l(C_i) \right\rfloor \le \left\lfloor \frac{l(C)}{w(C)} \right\rfloor + 1,$$

there exists a closed curve C_i such that $w(C_i) = 1$ and $l(C_i) \le \left\lfloor \frac{l(C)}{w(C)} \right\rfloor + 1$. This shows the existence of a hole cut of size at most $\left\lfloor \frac{l(C)}{w(C)} \right\rfloor + 1$. ■

IV. Algorithms

In this section, we discuss algorithmic results on the GMFCD and the GMCCD. First, by the constructive proof of Theorem 1 in Section III-B, we obtain a polynomial-time algorithm for computing MAX-FLOW.

Theorem 4. *An optimal solution of the GMFCD and a closed curve C minimizing $\lfloor \frac{l(C)}{w(C)} \rfloor$ can be computed in polynomial time.*

Note that this is the first polynomial-time algorithm for the GMFCD. The most time consuming part is Algorithm 1 that runs in $O(|W|^2)$ time. Since we execute Algorithm 1 at most $k :=$ MAX-FLOW times, the total running time is $O(k|W|^2)$. In most practical cases, since k is small and $|W| = O(|V|)$, the running time is $O(|V|^2)$.

Next, we propose a new algorithm for the GMCCD, which is simpler (and probably faster) than known algorithms.

Theorem 5. *An optimal solution of the GMCCD can be computed in polynomial time.*

Proof: By Theorem 4, we can compute s-t paths $P_1, \dots, P_k$ that are pairwise separated, where $k :=$ MAX-FLOW. Furthermore, by Theorem 3, we can also obtain a hole cut of size $k+1$ (i.e., a set of $k+1$ holes in $\mathcal{H}(r_b, r_p)$ that disconnect s from t). Since MAX-FLOW $\le$ MIN-CUT, our remaining task is to find a hole cut of size k if one exists.

Since the case of $k = 1$ is easy, in what follows we suppose $k \ge 2$. We may assume that $P_1, \dots, P_k$ do not cross each other, and the first edges of $P_1, \dots, P_k$ occur in this order clockwise at s. Recall that $R(P_{i-1}, P_i)$ is the closed region encircled by the closed curve $P_{i-1} \cdot (P_i)^{-1}$ in clockwise orientation. For $i = 1, \dots, k$, let $\mathcal{F}_i$ be the set of all faces of G contained in $R(P_{i-1}, P_i)$, where $P_0 := P_k$.

We observe that a hole cut of size k exists if and only if there exists a closed curve C with $w(C) = 1$ and $l(C) = k$ that is represented as a concatenation of $J_1, L_1, J_2, L_2, \dots, J_k, L_k$ in this order, where J_i is a curve contained in a face of $\mathcal{F}_i$, and $L_i \in \mathcal{L}$ is a line segment connecting $R(P_{i-1}, P_i)$ and $R(P_i, P_{i+1})$. Note that $P_{k+1} := P_1$ and $\mathcal{F}_{k+1} := \mathcal{F}_1$. To check the existence of such a curve, we construct a directed graph $D = (\mathcal{F}, A)$, where

$$\begin{aligned}\mathcal{F} &:= \bigcup_i \mathcal{F}_i \\ A &:= \{(F_i, F_{i+1}) \mid i \in \{1, \dots, k\},\ F_i \in \mathcal{F}_i,\ F_{i+1} \in \mathcal{F}_{i+1}, \\ &\qquad \exists H \in \mathcal{H}(r_b, r_p) \text{ intersecting } F_i,\ P_i, \text{ and } F_{i+1}\}.\end{aligned} \tag{*}$$

Then, finding a hole cut of size k is equivalent to finding a directed cycle of length k in D.

With this observation, we can compute MIN-CUT by Algorithm 2, and it is obvious that it runs in polynomial time. ■

V. Experimental Results

In this section, we describe experimental results. We implemented our algorithms for the GMFCD and the GMCCD, and evaluated their performance by computational experiments. Our experiments were conducted on a computer with Intel Core i7, 2.8 GHz and 8 GB of memory. All programs are written in Java.

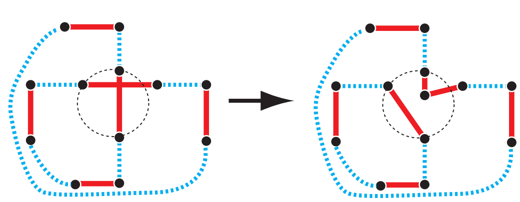

Fig. 9. Uncrossing procedure (case 3). Two line segments L_i and L_j are replaced with three line segments.

Algorithm 2 Find-MIN-CUT

Input: pairwise separated k s-t paths $P_1, \ldots, P_k$ and a hole cut of size $k+1$

Output: a hole cut of minimum size

construct the directed graph $D = (\mathcal{F}, A)$ defined as (*)
for all $F \in \mathcal{F}_1$ **do**
  by a breadth-first-search from F, try to find a directed cycle of length k in D containing F
  if such a directed cycle exists **then**
    return a hole cut of size k corresponding to the directed cycle
  end if
end for
return a hole cut of size $k+1$

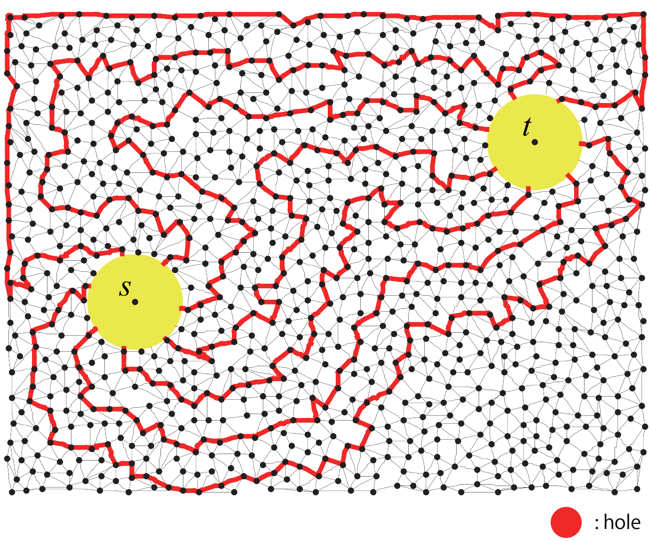

Fig. 10. Experimental result (MAX-FLOW)

As we have seen before, our algorithm for computing MAX-FLOW consists of the induction step (Algorithm 1) and the base cases (Lemma 2). Practically, since most short s-t paths satisfy the second condition of Lemma 2, we do not need an implementation of the algorithm in Lemma 2. Therefore, we can compute MAX-FLOW by just applying Algorithm 1, repeatedly. We generated input plane graphs with 1000 nodes randomly in a 300×400 rectangular and applied our algorithm to them. Then we can solve the GMFCD in a few seconds. As an example, a computational result with 1000 nodes and 7 paths is shown in Fig. 10, where we set $r_b = 10$ and $r_p = 30$. Note that we regarded the nodes on the boundary of the protective disks as the terminals to make the figure easier to see.

We also implemented Algorithm 2, and applied it to randomly generated graphs. For graphs with 1000 nodes, Algorithm 2 computes MIN-CUT in a few seconds. Fig. 11 is a computational result with 1000 nodes, where we set $r_b = 7$ and $r_p = 30$. In this case, we can see MIN-CUT = MAX-FLOW = 8.

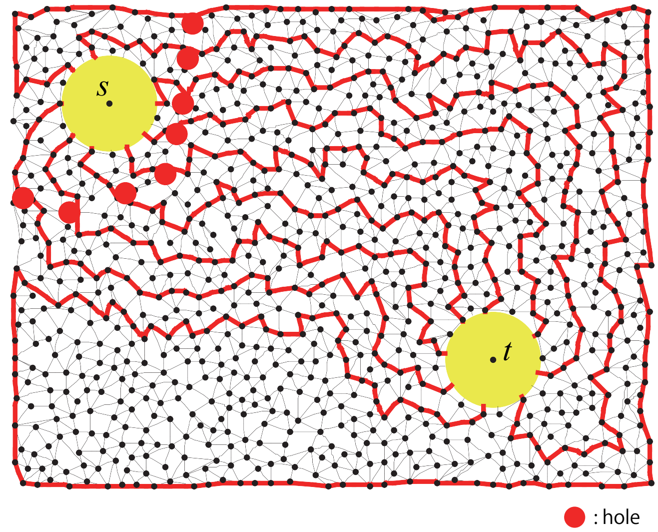

Fig. 11. Experimental result (MIN-CUT)

VI. Concluding Remarks

In this paper, we discussed the geographical min-cut and the geographical max-flow in the model, in which every hole is a disk of the same radius r_b. We proved a min-max theorem and gave polynomial-time algorithms for the GMFCD and the GMCCD that can be applied to large graphs.

Our results can be extended to the case with holes of different shapes. Suppose that $\mathcal{H}$ is a set of convex shapes (holes) satisfying the following property.

> **Property:** Suppose that two line segments L_1 in $H_1 \in \mathcal{H}$ and L_2 in $H_2 \in \mathcal{H}$ are crossing. Then, H_1 also contains an endpoint of L_2 or H_2 also contains an endpoint of L_1. (See the case analysis of the proof of Theorem 3.)

In Theorems 1 and 3, we can replace $\mathcal{H}(r_b, r_p)$ with any set $\mathcal{H}$ satisfying the above property. For example, $\mathcal{H}$ can be a set of disks of different sizes or a set of axis parallel squares. In particular, by setting $\mathcal{H}$ as the set of all edges (not incident to s and t), we obtain a min-max theorem for the maximum induced disjoint s-t paths [13] as a special case of Theorem 1.

Note that we cannot extend our proofs to the case when $\mathcal{H}$ is a set of general connected shapes violating the property above. Actually, computing MAX-FLOW becomes NP-hard when $\mathcal{H}$ does not necessarily satisfy the property [11].

Acknowledgements

The authors thank anonymous reviewers for helpful comments. This work is supported by JST, ERATO, Kawarabayashi Large Graph Project, Japan. The first author is supported by Grant-in-Aid for Scientific Research, MEXT, Japan.

References

[1] C. Bettstetter and C. Hartmann, "Connectivity of wireless multihop networks in a shadow fading environment," *Wireless Networks*, vol. 11, pp. 571–579, 2005.

[2] Y.-C. Cheng and T. G. Robertazzi, "Critical connectivity phenomena in multihop radio models," *IEEE Transactions on Communications*, vol. 37, pp. 770–777, 1989.

[3] P. Gupta and P. Kumar, "Critical power for asymptotic connectivity in wireless networks," *Stochastic Analysis, Control, Optimization and Applications*, pp. 547–566, 1998.

[4] M. F. Habib, M. Tornatore, F. Dikbiyik, and B. Mukherjee, "Disaster survivability in optical communication networks," *Comput. Commun.*, vol. 36, pp. 630–644, 2013.

[5] S. Neumayer, A. Efrat, and E. Modiano, "Geographic max-flow and min-cut under a circular disk failure model," in *Proceedings of the 31st IEEE International Conference on Computer Communications (IEEE INFOCOM 2012)*, 2012, pp. 2736–2740.

[6] S. Neumayer and E. Modiano, "Network reliability with geographically correlated failures," in *Proceedings of the 29th IEEE International Conference on Computer Communications (IEEE INFOCOM 2010)*, 2010.

[7] ——, "Network reliability under random circular cuts," in *Proceedings of the IEEE Global Telecommunications Conference (IEEE GLOBECOM 2011)*, 2011.

[8] A. Sen, B. H. Shen, L. Zhou, and B. Hao, "Fault-tolerance in sensor networks: A new evaluation metric," in *Proceedings of the 25th IEEE International Conference on Computer Communications (IEEE INFOCOM 2006)*, 2006.

[9] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity - a new paradigm for design of fault-tolerant networks," in *Proceedings of the 15th International Conference on High Performance Switching and Routing (HPSR 2009)*, 2009, pp. 94–100.

[10] S. Trajanovski, F. A. Kuipers, and P. V. Mieghem, "Finding critical regions in a network," in *Proceedings of the 5th International Workshop on Network Science for Communication Networks (NetSciCom 2013)*, 2013, pp. 223–228.

[11] D. Bienstock, "Some generalized max-flow min-cut problems in the plane," *Mathematics of Operations Research*, vol. 16, pp. 310–333, 1991.

[12] N. Sonnerat and A. Vetta, "Galaxy cutsets in graphs," *Journal of Combinatorial Optimization*, vol. 19, pp. 415–427, 2010.

[13] C. McDiarmid, B. Reed, A. Schrijver, and B. Shepherd, "Induced circuits in planar graphs," *Journal of Combinatorial Theory, Series B*, vol. 60, pp. 169–176, 1994.

[14] ——, "Non-interfering dipaths in planar digraphs," Centrum voor Wiskunde en Informatica, Tech. Rep., 1991.

[15] ——, "Non-interfering network flows," in *Proceedings of the Third Scandinavian Workshop on Algorithm Theory (SWAT 1992)*, 1992, pp. 245–257.

APPENDIX

In the appendix, we give a proof of Lemma 2, which we restate here.

Lemma. *Suppose that s and t are connected in G. We can find in polynomial time either*

- *a hole in $\mathcal{H}(r_b, r_p)$ that disconnects s from t, or*
- *an s-t path P in G such that for any line segment $L \in \mathcal{L}$, the union of P and L (which is regarded as a subset of the plane) contains no closed curve separating s and t.*

Proof: To simplify the argument, we discuss the case when $\mathcal{L}$ is replaced with $\mathcal{L}'$, where $\mathcal{L}'$ is the set of all line segments contained in holes in $\mathcal{H}(r_b, r_p)$. Note that there is no essential difference between $\mathcal{L}$ and $\mathcal{L}'$.

Assume that G contains no s-t path satisfying the condition in the second case. That is, for any s-t path P, there exists a line segment $L \in \mathcal{L}'$ such that $P \cup L$ contains a closed curve separating s and t.

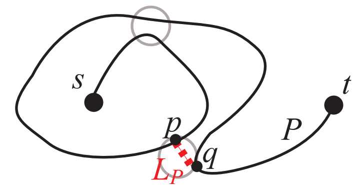

Fig. 12. Definition of L_P. There exist some line segments satisfying the condition (1) and L_P is the one satisfying the condition (2).

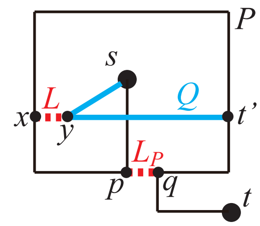

Fig. 13. A curve Q violating the second condition. There exists a line segment L such that $P[x, t'] \cup Q[y, t'] \cup L$ forms a closed curve separating s and t.

For a path P and for two points $x, y \in P$, the part of P between x and y is denoted by $P[x, y]$. For an s-t path P and for two points $x, y \in P$, we denote $x \prec_P y$ if x appears earlier than y when we move from s to t along P.

For each s-t path P, we take a line segment $L_P \in \mathcal{L}'$ connecting $p, q \in P$ with $p \prec_P q$ satisfying the following conditions (see Fig. 12).

(1) $P[p, q] \cup L_P$ forms a closed curve separating s and t.
(2) Subject to (1), q is maximal with respect to $\prec_P$.
(3) Subject to (1) and (2), p is minimum with respect to $\prec_P$.

Among all s-t paths, we choose an s-t path P such that $P[q, t]$ is maximal (i.e. it is as long as possible).

We show the following claim, which will be used to reroute the path P.

Claim 6. *Let P be a path chosen as above. If there is no hole in $\mathcal{H}(r_b, r_p)$ that disconnects s from t, then there exists a path Q from s to a node t' on $P[p, q] \setminus \{p, q\}$ with the following properties:*

- *Q does not intersect with $P[p, q] \cup L_P$ except its endnodes (i.e., Q is contained in the inside of cycle $P[p, q] \cup L_P$), and*
- *there exists no line segment $L \in \mathcal{L}'$ connecting $x \in P[p, q]$ and $y \in Q$ such that $x \prec_P t'$ and $P[x, t'] \cup Q[y, t'] \cup L$ forms a closed curve separating s and t (see Fig. 13).*

Proof: Suppose that there is no hole in $\mathcal{H}(r_b, r_p)$ that disconnects s from t. This implies that there exists a path Q from s to $P[p, q] \setminus \{p, q\}$ such that Q does not intersect with $P[p, q] \cup L_P$ except its endnodes.

In order to derive a contradiction, we assume that every path Q from s to $P[p,q]\setminus\{p,q\}$ violates the second condition, that is, for each Q, there exists a line segment $L'_Q \in \mathcal{L}$ connecting $x \in P[p,q]$ and $y \in Q$ such that $x \prec_P t'$ and $P[x,t'] \cup Q[y,t'] \cup L'_Q$ forms a closed curve separating s and t, where t' is the endnode of Q on $P[p,q]\setminus\{p,q\}$. Among such line segments, we choose L'_Q as the shortest one. (If there exist more than one shortest line segments, we choose L'_Q so that y is as small as possible with respect to $\prec_P$.) Among all paths from s to $P[p,q]\setminus\{p,q\}$, we choose a path Q such that L'_Q is as long as possible.

Since we assume that there is no hole separating s and t, there exists a path R from s to $(P[x,t'] \cup Q[y,t'])\setminus\{x,y\}$ that does not intersect with L'_Q. Let r_2 be the endpoint of R that is different from s, and let r_1 be the point in $Q \cap R$ such that no inner nodes of $R[r_1,r_2]$ are contained in Q. Define the path Q^* as $Q^* := Q[s,r_1] \cup R[r_1,r_2]$ if $r_2 \in P[x,t']\setminus\{x\}$ and $Q^* := Q[s,r_1] \cup R[r_1,r_2] \cup Q[r_2,t']$ if $r_2 \in Q[y,t']\setminus\{y\}$.

By the assumption, there exists a line segment $L'_{Q^*} \in \mathcal{L}$ connecting $x^* \in P[p,q]$ and $y^* \in Q^*$ such that $x^* \prec_P t^*$ and $P[x^*,t^*] \cup Q^*[y^*,t^*] \cup L'_{Q^*}$ forms a closed curve separating s and t, where t^* is the endnode of Q^* on $P[p,q]\setminus\{p,q\}$. Since L'_{Q^*} must intersect with Q, L'_{Q^*} is longer than L'_Q, which contradicts the choice of Q. This completes the proof. ■

If there is a hole in $\mathcal{H}(r_b,r_p)$ that disconnects s and t, then we are done. Therefore, we may assume that we have a curve Q from s to t' as in Claim 6. Note that since the proof of Claim 6 is constructive, we can find such Q in polynomial time. Consider the s-t path P' obtained by concatenating $Q[s,t']$ and $P[t',t]$. Since P' is an s-t path, there exists a line segment $L' \in \mathcal{L}$ connecting $p',q' \in P'$, where $p' \prec_{P'} q'$, such that $L' \cup P'[p',q']$ forms a closed cycle separating s and t. Among such line segments, we choose L' such that

(a) $q \prec_P q'$ or $q = q'$, and
(b) subject to (a), p' is maximum with respect to $\prec_{P'}$.

Note that if there exists no line segment satisfying the condition (a), that is, for any line segment L', its endnode q' satisfies $q' \prec_P q$, then $P[q,t]$ is strictly contained in $P'[q',t]$, which contradicts the choice of P.

Since q is defined as a maximum node with respect to $\prec_P$, cycle $P[p,q] \cup L_P$ separates p' and t. Furthermore, since $P[q,t]$ is as long as possible by the choice of P, either $P[p,q] \cup L_P$ separates s and q' or $q' = q$. Therefore, L' intersects with either $P[p,q]$ or L_P.

Case 1. Assume that $q' \neq q$ and L' intersects with $P[p,q]$. Let x be a node which is an intersection of L' and $P[p,q]$ (see Fig. 14). By Claim 6, $P'[p',t'] \cup L'[p',x] \cup P[x,t']$ does not separate s and t. By combining this with the fact that $P'[p',q'] \cup L'$ separates s and t, we can see that $P[x,q'] \cup L'[q',x]$ separates s and t, which contradicts the maximality of q.

Case 2. Assume that L' intersects with L_P. Let x be a node which is an intersection of L' and L_P. In this case, $C' = L'[p',x] \cup L_P[x,q] \cup P'[p',q]$ forms a closed curve. We divide this case into the following two cases.

Case 2-1. We consider the case when C' does not separate s and t (see Fig. 15). By Claim 6, p and p' are not contained in a common hole in $\mathcal{H}(r_b,r_p)$. Then, by the property of $\mathcal{H}(r_b,r_p)$ described in Section VI, q and q' are contained in a common hole in $\mathcal{H}(r_b,r_p)$. This shows that by concatenating $P[q,q']$ and the line segment connecting q and q', we obtain a closed cycle separating s and t, which contradicts the maximality of q.

Case 2-2. We consider the case when C' separates s and t (see Figs 16 and 17). Since C' also separates s and p, $P[s,p]$ intersects $P'[p',t']$ or $L'[p',x]$. Let y be the first intersecting point of $P[s,p]$ and $P'[p',t'] \cup L'[p',x]$ when we move from p to s along $P[s,p]$. We consider the following two cases separately.

If y is on $L'[p',x]$, then $P[y,q'] \cup L'[y,q']$ forms a closed cycle separating s and t (see Fig. 16). This contradicts the maximality of q when $q \prec_P q'$, and it contradicts the

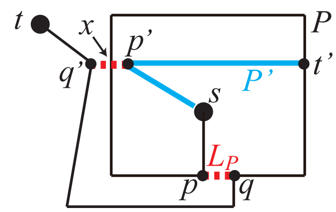

Fig. 14. In the case 1, $P[x,q'] \cup L'[q',x]$ separates s and t.

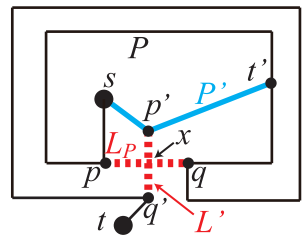

Fig. 15. In the case 2-1, by concatenating $P[q,q']$ and the line segment connecting q and q', we obtain a closed cycle separating s and t.

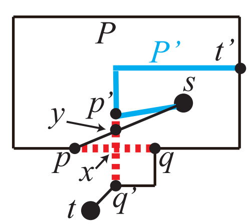

Fig. 16. In the case 2-2, y is on $L'[p',x]$.

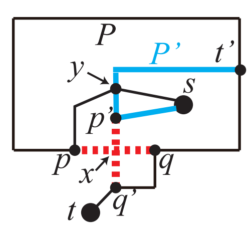

Fig. 17. In the case 2-2, y is on $P'[p',t']$.

minimality of p when $q = q'$.

If y is on $P'[p', t']$, then $P'' := P[s, y] \cup P'[y, t'] \cup P[t', t]$ forms an s-t path (see Fig. 17). By the minimality of p and by the maximality of p' (see condition (b)), if $P''[p'', q''] \cup L''$ forms a closed curve separating s and t for some line segment $L'' \in \mathcal{L}'$ connecting p'' and q'', where $p'' \prec_{P''} q''$, then $P[q, t]$ is strictly contained in $P''[q'', t]$. This contradicts the choice of P.

By repeating this procedure, we can find a hole or an s-t path P satisfying the condition. ■