Disjoint Paths in Densely Embedded Graphs

Jon Kleinberg* Éva Tardos†

October, 1995

Abstract

We consider the following *maximum disjoint paths problem* (MDPP). We are given a large network, and pairs of nodes that wish to communicate over paths through the network — the goal is to simultaneously connect as many of these pairs as possible in such a way that no two communication paths share an edge in the network. This classical problem has been brought into focus recently in papers discussing applications to routing in high-speed networks, where the current lack of understanding of the MDPP is an obstacle to the design of practical heuristics.

We consider the class of *densely embedded, nearly-Eulerian graphs*, which includes the two-dimensional mesh and many other planar and locally planar interconnection networks. We obtain a constant-factor approximation algorithm for the maximum disjoint paths problem for this class of graphs; this improves on an $O(\log n)$-approximation for the special case of the two-dimensional mesh due to Aumann–Rabani and the authors. For networks that are not explicitly required to be "high-capacity," this is the first constant-factor approximation for the MDPP in any class of graphs other than trees.

We also consider the MDPP in the on-line setting, relevant to applications in which connection requests arrive over time and must be processed immediately. Here we obtain an asymptotically optimal $O(\log n)$-competitive on-line algorithm for the same class of graphs; this improves on an $O(\log n \log \log n)$-competitive algorithm for the special case of the mesh due to Awerbuch, Gawlick, Leighton, and Rabani.

1 Introduction

We consider the following *maximum disjoint paths problem* (MDPP). We are given a large network, and pairs of nodes that wish to communicate over paths through the network — the goal is to simultaneously connect as many of these pairs as possible in such a way that no two communication paths share an edge in the network. This problem is well-known to be computationally difficult. Deciding whether all pairs can be so connected is one of Karp's

*E-mail: kleinber@theory.lcs.mit.edu. Laboratory for Computer Science, MIT, Cambridge MA 02139 USA. Supported by an ONR Graduate Fellowship.

†E-mail: eva@cs.cornell.edu. School of ORIE, Cornell University, Ithaca NY 14853. Supported in part by a Packard Fellowship and an NSF PYI award.

original NP-complete problems [12]; it remains NP-complete even when the underlying graph is the two-dimensional mesh [15].

Our interest in this problem comes from two main sources. First, establishing disjoint paths is fundamental to routing in high-speed networks (see for example the applications mentioned in [6, 8, 21], as well as applications to optical routing in [1, 3, 24]). Although the types of routing problems that arise in such settings tend to have additional side constraints (e.g. connections have limited duration and can bring varying amounts of "profit"), the formulation described in the first paragraph contains the essence of virtually all such real-life routing problems in which each connection consumes a large fraction of the bandwidth on a link. As such, the current lack of understanding of the disjoint paths problem is a major obstacle to the design of practical heuristics. Indeed, [6] notes that in practice, the greedy algorithm tends to be used for routing, despite its bad performance on a number of very common interconnection patterns. Moreover, robust ways are known for converting algorithms for the MDPP into algorithms that can handle connections of limited duration or variable value [5]; thus, the difficulties contained in these more elaborate routing problems seem to stem mainly from the intractability of the MDPP.

This problem is also of basic interest in algorithmic graph theory. A lot of work has been done on identifying special cases of the disjoint paths problem that can be solved in polynomial time, or for which simple min-max conditions can be stated; see the survey by Frank [10]. Much less work has been done, however, on approximation algorithms for the MDPP; we are interested in extending the classes of graphs for which good approximations can be obtained.

1.1 Our Results

To be precise, let $G = (V, E)$ be a graph on n vertices and $\mathcal{T} = \{(s_1, t_1), \ldots, (s_k, t_k)\}$ a collection of *terminal pairs* — pairs of vertices of G. We say that $\mathcal{T}$ is *realizable* in G if there exist mutually edge-disjoint s_i-t_i paths, for $i = 1, \ldots, k$. The problem is then to find a realizable subset of $\mathcal{T}$ of maximum cardinality.

Our first main result is a constant-factor approximation for the maximum disjoint paths problem in the class of *densely embedded, nearly-Eulerian graphs* (defined below), which includes many common planar and locally planar interconnection networks. This improves on an $O(\log n)$-approximation for the case of the two-dimensional mesh due to Aumann and Rabani [3] and an $O(\log n)$-approximation for a class of planar graphs including the mesh due to the authors [14]. Our present algorithm makes use of variants of a number of the techniques developed in our earlier paper [14].

The assumption that we know all the terminal pairs in advance is not reasonable in situations in which connection requests between pairs of nodes arrive over time and must be processed immediately. In such a setting, it makes sense to consider *on-line* routing algorithms. Such an algorithm is given the graph G, terminal pairs arrive in an arbitrary order, and for each such pair it must irrevocably reject it, or assign it a path in G. As is standard, we refer to the approximation ratio achieved by an on-line algorithm as its *competitive ratio*; such an algorithm is said to be c-competitive if its competitive ratio is at most c.

Our second main result, then, is an $O(\log n)$-competitive randomized on-line algorithm for the MDPP in densely embedded, nearly-Eulerian graphs. This improves on an $O(\log n \log\log n)$-competitive algorithm for the special case of the two-dimensional mesh due to Awerbuch, Gawlick, Leighton, and Rabani [6]; moreover, [6] proves that no randomized on-line algorithm for the two-dimensional mesh can be better than $\Omega(\log n)$-competitive, implying that our algorithm is asymptotically optimal.

We feel that an important feature of our algorithms, in addition to the improved bounds, is that they are not specific to the mesh; the advantage of developing algorithms that work on the somewhat larger class of "densely embedded graphs" is that are not sensitive to small variations in the structure of the underlying graph. This could be of value in the context of network routing, where the underlying network may have a "mesh-like" topology, but lack the completely regular structure of the mesh. In contrast, previous algorithms such as [6, 3] could not be applied to any network other than the mesh itself, since they required its fixed row/column structure.

The size of the constants in our algorithms as presented here, while not astronomical, pushes them outside the range of immediate practical utility. However, the previous best bounds — both off-line and on-line — for the two-dimensional mesh [3, 6, 14] involve similarly large constants inside the $O(\cdot)$ notation. Moreover, despite the large constants, some of the ideas used by the algorithms here may be of use in suggesting practical heuristics.

The rest of the paper is organized as follows. First, we present some preliminary algorithmic tools that we need in Section 2. Our routing algorithms are simpler to explain in the special case of the mesh, and we consequently present this case first. This will be done in Section 3. In Section 4, we then introduce the class of densely embedded graphs, and we present the algorithms for this class in Section 5. Finally, we show how to extend our algorithms to a slightly more general class of graphs in Section 6.

1.2 Previous Work

Much of the previous work on this problem has dealt with the case in which each path consumes only a small fraction of the available bandwidth on an edge; this can be modeled by requiring $\Omega(\log n)$ parallel copies of each edge. In this case, the randomized rounding technique of Raghavan and Thompson [23, 22] can be used to obtain an off-line constant-factor approximation. Awerbuch, Azar, and Plotkin give an on-line $O(\log n)$-competitive algorithm for this case [4], which they show is asymptotically tight for deterministic on-line algorithms.

As noted in [6] however, there are many applications in which each communication path consumes a large fraction of the available bandwidth on a link; thus it makes sense to consider approximation algorithms for graphs without a large number of parallel edges. The results here are much more restricted. For trees with parallel edges, Garg, Vazirani, and Yannakakis [11] obtain an off-line 2-approximation (the maximization problem is NP-complete, though deciding realizability is easy); Awerbuch et. al. [6] give an $O(\log d)$-competitive randomized on-line algorithm for trees of diameter d, extending an earlier result of Awerbuch et. al. [5]. Essentially the only approximation results known for graphs other than trees are those mentioned earlier for the mesh and related planar graphs [3, 6, 14]. Thus our result here is the first constant-factor approximation for any class of graphs other than trees, when one does not require $\Omega(\log n)$ parallel copies of each edge.

A different approach is taken in papers of Peleg and Upfal [21] and Broder, Frieze, and Upfal [8] (see also Broder et. al. [7]). Here the underlying graph G is assumed to have strong expansion properties; in this case one can prove that any set of terminal pairs of at most a given size must be realizable in G, and that corresponding paths can be found in (randomized) polynomial time. The results in [8] are strong enough that they implicitly provide a polylogarithmic approximation for the MDPP in sufficiently strong expanders of bounded degree.

In this context, it is worth mentioning the following closely related routing problem: one must route all terminal pairs so as to minimize the maximum *congestion* on any edge; that is, the maximum number of paths that contain an edge. Deciding whether $\mathcal{T}$ can be routed with congestion equal to 1 is the same as deciding whether $\mathcal{T}$ is realizable; but as an optimization problem, minimizing congestion is much better understood. Aspnes et al. [2] give an on-line $O(\log n)$-competitive algorithm for the problem; and a randomized rounding algorithm of Raghavan and Thompson [23] gives a routing off-line with congestion at most $OPT + o(OPT) + O(\log n)$, where OPT denotes the optimum congestion. This leaves open the question of whether a constant-factor approximation is achievable also for small values of OPT. In a companion paper with Satish Rao [13], we give a constant-factor approximation for densely embedded graphs; it was here that the idea of using a "simulated network" (see Section 3) was initially developed.

Cases in which the MDPP can be solved in polynomial time are surveyed in [10]; here we only discuss two specific results that we will use in handling densely embedded graphs. First, suppose G is planar, the terminals $\mathcal{T}$ lie on a single face of G, and the pair $(G, \mathcal{T})$ satisfies the following *parity condition*: the *augmented graph* formed by adding to G the edges corresponding to $\mathcal{T}$ must be Eulerian. In this case, a theorem of Okamura and Seymour [20] says that the realizability of $\mathcal{T}$ in G can be decided in polynomial time; and in fact the following *cut condition* is sufficient for realizability: one cannot remove j edges from G and separate more than j terminal pairs. A linear-time algorithm for this problem has recently been obtained by Wagner and Weihe [30]. We will use an extension of the Okamura–Seymour, due to Frank [9], which concerns the case in which the parity condition does need not to hold on the face containing the terminals.

We also use a theorem of Schrijver [28] that provides an algorithm for finding vertex-disjoint paths in a graph embedded on a compact surface Σ, such that the paths satisfy given homotopy constraints.

2 Preliminary Tools

2.1 The AAP Algorithm

We make use of a variant of an on-line MDPP algorithm of Awerbuch, Azar, and Plotkin [4]. If H is a graph with n nodes in which each edge has capacity at least $\log 2n$, the algorithm of [4] achieves a competitive ratio of $2 \log 4n$. For our purposes, we need to develop a strengthening of this "AAP algorithm": we want to be competitive against the fractional optimum; and when we deal with the more general case of densely embedded graphs, we want only to require capacities to be $\varepsilon \log n$, for an arbitrary $\varepsilon > 0$. Here, we show how to obtain such an algorithm.

Proposition 2.1 *If all edge capacities are at least $(\varepsilon \log n + 1 + \varepsilon)$, there is a deterministic on-line* MDPP *algorithm that is $O(2^{1/\varepsilon} \log n)$-competitive against the fractional optimum.*

Proof. We follow the AAP algorithm and its analysis very closely. We vary a little from their notation, since we only deal here with routing a maximal number of requests, each of infinite duration. Thus, the i^{th} request is specified by a pair (s_i, t_i) of terminals. We define the "profit" of the connection to be n; thus the total profit obtained by the on-line algorithm is simply n times the number of terminal pairs routed.

Define $\mu = 2^{1+1/\varepsilon} n$, so we have

$$\varepsilon \log \mu = \varepsilon \log n + 1 + \varepsilon.$$

Let u_e denote the capacity of edge e; thus we can assume that for all e,

$$u_e \geq \varepsilon \log \mu.$$

With this value of μ, we now run the AAP algorithm — for the sake of completeness, we state this algorithm here.

For $j = 1, 2, \ldots, k$:
Define λ_e^j to be the fraction of u_e consumed by paths already routed.
Define $c_e^j = u_e(\mu^{\lambda_e^j} - 1)$.
For a request (s_i, t_i), route it on any path P satisfying $\sum_{e \in P} \frac{1}{u_e} c_e^j \leq n$.
If no such path is available, then reject the request.

First we argue why the relative load on an edge will never exceed 1. At the moment before this happened, on edge e say, we had

$$\lambda_e^j > 1 - \frac{1}{u_e} \geq 1 - \frac{1}{\varepsilon \log \mu}.$$

Thus

$$\begin{aligned}
\frac{c_e^j}{u_e} &= \mu^{\lambda_e^j} - 1 \\
&> \mu^{1 - \frac{1}{\varepsilon \log \mu}} - 1 \\
&= \frac{\mu}{2^{1/\varepsilon}} - 1 = 2n - 1 \\
&\geq n.
\end{aligned}$$

So we have

$$\frac{c_e^j}{u_e} > n$$

and thus the connection could not have used this edge.

Suppose there are a total of k requests. Let A denote the set of requests routed by the AAP algorithm. Then we show

$$2^{1+1/\varepsilon} \log \mu \sum_{j \in A} n \geq \sum_e c_e^{k+1}. \tag{1}$$

As in the proof in [4] we show this by induction on the number of admitted requests, via proving that if the algorithm admits the j^{th} request, we have

$$\sum_e c_e^{j+1} - c_e^j \leq 2^{1+1/\varepsilon} n \log \mu.$$

So consider edge e on the j^{th} path used by the AAP algorithm. We have

$$c_e^{j+1} - c_e^j = u_e \left(\mu^{\lambda_e^j} (2^{(\log \mu)/u_e} - 1) \right).$$

Now the exponent on the 2 is at most $1/\varepsilon$, and for $x \in [0, 1/\varepsilon]$ we clearly have $2^x - 1 \leq 2^{1/\varepsilon} \cdot x$. Thus

$$\begin{aligned} c_e^{j+1} - c_e^j &\leq u_e \cdot \mu^{\lambda_e^j} \cdot 2^{1/\varepsilon} \cdot (\log \mu)/u_e \\ &= \mu^{\lambda_e^j} \cdot 2^{1/\varepsilon} \cdot \log \mu \\ &= 2^{1/\varepsilon} \cdot \log \mu \cdot \left[\frac{c_e^j}{u_e} + 1 \right]. \end{aligned}$$

Summing over all edges gives the desired bound.

Finally, we show that the expression

$$\sum_e c_e^{k+1} \tag{2}$$

is an upper bound on the profit of the *fractional* optimum minus the profit of the on-line algorithm. ([4] shows this for the integer optimum, but the proof is essentially the same.)

Let $\mathcal{Q}$ denote the set of indices which were rejected by the on-line algorithm but for which a positive fraction of the demand was routed by the optimum. For $j \in \mathcal{Q}$, suppose that the fractional optimum uses paths $P_j^1, \ldots, P_j^z$, with weights $\gamma_j^1, \ldots, \gamma_j^z$. Then since j was rejected by the on-line algorithm, and the costs are monotonic in the indices, we must have

$$n \leq \sum_{e \in P_j^i} \frac{c_e^{k+1}}{u_e}$$

for any i, j. Then for any edge e we have

$$\sum_{i,j: e \in P_j^i} \frac{\gamma_j^i}{u_e} \leq 1,$$

and hence we have

$$\begin{aligned}\sum_j \sum_i \gamma_j^i n &\leq \sum_j \sum_i \sum_{e \in P_j^i} \frac{\gamma_j^i c_e^{k+1}}{u_e} \\ &\leq \sum_e c_e^{k+1} \cdot \sum_{i,j: e \in P_j^i} \frac{\gamma_j^i}{u_e} \\ &\leq \sum_e c_e^{k+1}.\end{aligned}$$

Combining the bounds in Equations (1) and (2), we obtain the claimed competitive ratio. ■

A lower bound of [4] implies that the factor of $2^{1/\varepsilon}$ is unavoidable for deterministic on-line algorithms.

2.2 Combining On-Line Algorithms

In routing connections on-line, we will adopt an approach in which the decision whether to accept a given connection is made by a combination of several algorithms — the connection is accepted if each of the individual algorithms accepts it. From the competitive ratios of these individual algorithms one can infer a competitive ratio for this combined algorithm; in this section we show how this can be done.

Let U denote a finite set, with $S_1, \ldots, S_n$ subsets of U such that $U = \cup_i S_i$. Let $\mathcal{F}_i$ denote a collection of subsets of S_i closed with respect to inclusion, and let

$$\mathcal{F} = \{C : \forall i\ (C \cap S_i) \in \mathcal{F}_i\}.$$

Given a set $U' \subset U$, define $\mu(U')$ to be the maximum size of a member of $\mathcal{F}$ contained in U'. We wish to design an algorithm for the following on-line maximization problem with respect to U and $\mathcal{F}$. Elements of some $U' \subseteq U$ arrive in an arbitrary order, and on each element our algorithm either accepts or rejects it; the goal is to accept a subset of these elements that is in $\mathcal{F}$ and as large as possible relative to $\mu(U')$. Our algorithm will be called c-competitive if it always accepts a set of size at least $\frac{1}{c}\mu(U')$.

We can define the corresponding on-line maximization problems with respect to S_i and $\mathcal{F}_i$, for each $i = 1, \ldots, n$, in exactly the same way. Say that for each i, we are given an algorithm A_i which is c_i-competitive for the problem associated with S_i and $\mathcal{F}_i$. Moreover, we assume that the state of A_i is completely determined by the set of elements it has accepted so far. We then define our "combined algorithm" $A = \wedge_{i=1}^n A_i$ for the on-line maximization problem with respect to U and $\mathcal{F}$ as follows. As each $u \in U'$ is presented to A, it accepts u iff for each i such that $u \in S_i$, A_i accepts u. The total set accepted so far, intersected with S_i, serves as the state for A_i. Let c^* denote the maximum competitive ratio of any of the algorithms A_i, and suppose each element of U appears in at most d of the S_i.

Proposition 2.2 *A is $c^* d$-competitive.*

Proof. Assume the algorithm A was presented with a set U' and it returned X. Let Y denote a member of $\mathcal{F}$ contained in U' of maximum size; we show that $|Y| \leq c^* d|X|$. Let R_i' denote the elements of $Y \setminus X$ that were rejected by algorithm A_i, $J_i = X \cap Y \cap S_i$ the elements of S_i accepted by *both* A and the optimal solution Y, and $R_i = J_i \cup R_i'$. Note that $Y = \cup_i R_i$, and $R_i \subset Y \cap S_i \in \mathcal{F}_i$.

We want to prove that $|R_i| \leq c^* |X \cap S_i|$ $(i = 1, \ldots, n)$. Set $U_i' = (X \cap S_i) \cup R_i$; these are the elements of $U' \cap S_i$ either accepted by A or rejected by A_i. Order U_i' as it appears in U', and present it as input to A_i. Then as in the running of the combined algorithm A, A_i will accept precisely the set $X \cap S_i$. Since A_i is c^*-competitive, and $R_i \in \mathcal{F}_i$, we have

$$|R_i| \leq c^* |X \cap S_i|.$$

We also have $|Y| \leq \sum_i |R_i|$, and by the definition of d we have $\sum_i |X \cap S_i| \leq d|X|$. Thus

$$|Y| \leq \sum_i |R_i| \leq c^* \sum_i |X \cap S_i| \leq c^* d|X|. \blacksquare$$

In the natural way, one can define a fractional version of this model; one then shows the following by a minor modification of the proof above.

Proposition 2.3 *If each A_i is c^*-competitive against the fractional optimum, then A is $c^* d$-competitive against the fractional optimum.*

3 The Two-Dimensional Mesh

In this section, let $G = (V, E)$ denote the $n \times n$ mesh. A very rough sketch of the algorithms is as follows. Since much stronger results are known for cases in which edges have capacity $\Omega(\log n)$, we want to model G by a high-capacity "simulated network" $\mathcal{N}$. To do this we choose, for a constant γ, a maximal set of $\gamma \log n \times \gamma \log n$ subsqures of G subject to the condition that the distance between subsquares is at least $2\gamma \log n$. These will serve as the nodes of $\mathcal{N}$. We now label some pairs of subsquares as "neighbors" and connect them with $\Omega(\log n)$ parallel edges; these are the high-capacity edges of $\mathcal{N}$.

On this network $\mathcal{N}$ we run the algorithm of Awerbuch, Azar, and Plotkin [4] in the on-line case, and the algorithm of Raghavan and Thompson [23] in the off-line case. The algorithms running in $\mathcal{N}$ will return routes consisting of a sequence of neighboring subsquares. To convert such a route into a path in G, we construct disjoint paths between neighboring subsquares. We link a sequence of neighboring pairs together by the natural crossbar structures surrounding each subsquare. This leaves us with the problem of routing from each original terminal to the boundary of its subsquare. In the on-line case we will route at most one terminal in each subsquare, and hence routing out of a subsquare will be easy. In the off-line case we use network flow techniques to route the appropriate subset of terminals to the boundary of the subsquare. Finally, to prove the approximation ratios, we argue that the number of pairs routed by the optimum in G is upper-bounded by the maximum number of pairs that can be routed in a copy of $\mathcal{N}$ in which all edges have capacity $O(\log n)$.

When G denotes the two-dimensional mesh, let $G[i,j]$ denote the vertex with row number i and column number j, and $G[i:i',j:j']$ the subsquare

$$\{G[p,q] : i \le p \le i',\ j \le q \le j'\}.$$

Let $d(u,v)$ denote the least number of edges in a u-v path. By the L_∞ distance between vertices $G[i,j]$ and $G[i',j']$, we mean $L_\infty(G[i,j],G[i',j']) = \max(|i-i'|,|j-j'|)$.

3.1 Building the Simulated Network

We choose a constant $\gamma > 1$ (any constant will do; it will have an influence on the approximation ratio we obtain). Our first goal is to choose a maximal set of "mutually distant" vertices around which to grow nodes of the simulated network. We divide the the mesh into $\gamma \log n$ by $\gamma \log n$ subsquares as follows.

Definition 3.1 *A subsquare of V is called a* γ-block *if for some natural numbers i and j, it is equal to the set*

$$G[(i-1)\gamma \log n : i\gamma \log n, (j-1)\gamma \log n : j\gamma \log n].$$

If X is a γ-block with associated natural numbers i and j, then the vertex

$$v = G[(i-\frac{1}{2})\gamma \log n, (j-\frac{1}{2})\gamma \log n]$$

will be called the center *of the γ-block, and we will denote X by C_v. A boundary vertex of X is one with maximal or minimal row or column number. We use X^* to denote the union of X with the (at most) eight other γ-blocks that share boundary vertices with X.*

By a *wall* of X, we mean a maximal set of boundary vertices having the same row or column number. A vertex of X is *internal* if it is not a boundary vertex.

Let V' denote the set of all centers of γ-blocks. We build a graph G' on V' by joining $u,v \in V'$ if the corresponding sets C_u^* and C_v^* intersect at an internal vertex. We now run a randomized version of Luby's maximal independent set algorithm [17] on this graph. That is, each vertex picks a random number between 1 and j, where j is large enough that the probability of ties is small. If v has a number higher than any of its neighbors', it enters the MIS and its neighbors drop out. We then iterate.

Let $M \subset V'$ denote the resulting MIS. For any $v \in V'$, C_v^* intersects C_u^* internally for at most 24 other vertices $u \in V'$; thus with probability $\frac{1}{25} - o(1)$, v will enter M on the first iteration. Moreover, if $u,v \in V'$ are at a distance of at least $11\gamma \log n$ from each other, then they have no common neighbors in G', and so these events are independent. Thus,

Lemma 3.2 *Let $u,v \in V'$ be such that $d(u,v) \ge 11\gamma \log n$. Then with constant probability, both u and v belong to the set M constructed above.*

If $v \in M$, we will call C_v a *cluster*. We now want to construct internally disjoint *enclosures* D_v around each C_v^*, for $v \in M$, such that every vertex of G belongs to some enclosure and

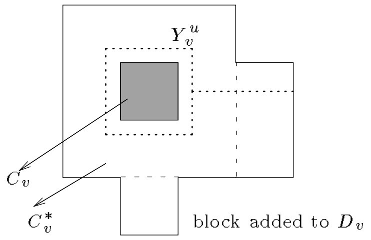

Figure 1: A cluster and its surroundings

such that each D_v is a union of γ-blocks. The sets C_v^* are disjoint and are unions of γ-blocks, but they do not cover all of G. However, by the maximality of M, any γ-block X that does not belong to C_v^* for some $v \in M$ must share a boundary vertex with such a C_v^*. For each such X, we pick such a C_v^* arbitrarily and add X to D_v. Thus the D_v now form a partition of G, and each D_v is a union of γ-blocks.

We now define $\mathcal{N}$ to be the graph on vertex set M, with u and v joined by an edge if some γ-block of D_u shares a wall with a γ-block of D_v. As argued above, any γ-block in D_v must belong to the 5×5 set of γ-blocks centered at C_v, and from this it is easy to argue that at most 20 γ-blocks not contained in D_v can share a wall with a γ-block of D_v. Thus we have

Lemma 3.3 *The degree of a vertex in $\mathcal{N}$ is at most* 20*.*

If we were to contract each enclosure D_v, the resulting graph would contain all the edges of $\mathcal{N}$ with multiplicity $\Theta(\log n)$ — enough to run a "high-bandwidth" algorithm. But given a routing in $\mathcal{N}$, we then run into the following problem: we need to convert the sequence of neighboring clusters connecting the clusters of s_i and t_i in the contracted graph back to a routing in G. For this we use the natural "crossbar" structures in the mesh in the on-line case, with additional flow techniques in the off-line case.

We start by developing the "crossbar" structures we use.

Definition 3.4 *A v-*ring *is a subgraph $G[X]$, where X is the set of vertices at L_∞ distance exactly r from v, for some r between $\frac{1}{2}\gamma \log n$ and $\gamma \log n$. Thus a v-ring is either a cycle or a path, depending on whether C_v has a wall on the boundary of G. If R and R' are v-rings, we say that R is* inside *(resp.* outside*) R' if the distance from R to v is less than (resp. greater than) the distance from R' to v.*

Note that the set of v-rings are the $\frac{1}{2}\gamma \log n$ disjoint cycles around v right outside the boundary of C_v, in the "inner half" of $C_v^* \setminus C_v$.

At this point we will need some additional notation: If $X \subset V$, let $\delta(X)$ denote the set of edges leaving S, and $\pi(X)$ the set of vertices of X incident to an edge of $\delta(X)$. For two sets $X, Y \subset V$, let $\delta(X, Y)$ denote the set of edges crossing from X to Y.

For each pair (v, w) that is an edge of $\mathcal{N}$, we choose, for a sufficiently small constant ρ, a set $\tau_{v,w}$ of $\rho \log n$ edges in $\delta(D_v, D_w)$. Let τ_v' denote the set of all vertices in D_v that are incident to an edge in some $\tau_{v,w}$. We also choose a set σ'_v of $\rho \log n$ vertices evenly spaced on the outer boundary of C_v. Now by Lemma 3.3, we can choose ρ small enough that $|\tau'_v \cup \sigma'_v| \leq \frac{1}{2}\gamma \log n$, and hence we can associate a different v-ring to each vertex in $\tau'_v \cup \sigma'_v$. Moreover, in a straightforward fashion we can construct edge-disjoint paths from each such vertex to its associated ring. We assign the outermost $\rho \log n$ rings to the vertices of σ'_v. For $u \in \tau'_v \cup \sigma'_v$, let Y_v^u denote the union of the ring associated with u with the path from u to this ring. Then we have

Lemma 3.5 *The (non-simple) paths Y_v^u are mutually edge-disjoint, and each pair meets at some vertex of $C_v^* \setminus C_v$.*

Proof. The paths are edge-disjoint by construction. Now suppose $u, w \in \tau'_v$, and that the ring associated with u is inside the ring associated with w. Then the path from u to its ring must intersect the ring associated with w, and so Y_v^u and Y_v^w meet at a vertex. The same argument applies if $u, w \in \sigma'_v$, and also if $u \in \sigma'_v$ and $w \in \tau'_v$ because in this case the ring associated with u lies outside the ring associated with w. ■

We are now ready to describe and analyze the routing algorithms themselves.

3.2 The On-Line Algorithm

Say that a terminal pair $(s_i, t_i) \in \mathcal{T}$ is *short* if $d(s_i, t_i) < 16\gamma \log n$, and *long* otherwise. The on-line algorithm makes an initial random decision whether to accept only short connections or only long connections; this costs at most a factor of two in the competitive ratio. Below we give $O(\log n)$-competitive algorithms for handling each type of connection.

3.2.1 Routing Long Connections

First, we only consider terminal pairs with both ends in sets of the form C_v — denote this set of terminal pairs by $\mathcal{T}_M$. If $\mathcal{T}^*$ denotes a realizable subset of maximum size, then by Lemma 3.2, the expected number of terminal pairs in $\mathcal{T}^*$ that belong to $\mathcal{T}_M$ is a constant fraction of $|\mathcal{T}^*|$. Thus we only lose a constant factor in the competitive ratio by restricting attention to $\mathcal{T}_M$.

Let $\mathcal{N}(c)$ denote the graph $\mathcal{N}$ in which each edge has been given a capacity of c. We now define an on-line routing problem in the simulated network $\mathcal{N}(\rho \log n)$. If $s_i \in C_v$, then we define its image in the "simulation" to be $\psi(s_i) = v$. The input will simply be the sequence of terminal pairs $(\psi(s_i), \psi(t_i))$, where (s_i, t_i) is the sequence of pairs presented to the algorithm running on G. Our algorithm for the problem in the simulated network is as follow: we route (v, w) if (i) the AAP algorithm on $\mathcal{N}(\rho \log n)$ accepts (v, w), and (ii) no connection with an end equal to either v or w has yet been accepted.

Lemma 3.6 *The above algorithm is $O(\log n)$-competitive against the fractional optimum in $\mathcal{N}(\rho \log n)$.*

Proof. Let (v, w) be a requested connection in $\mathcal{N}$. Following the approach developed in Section 2.2, we view (v, w) as being "judged" by three on-line algorithms: the AAP algorithm, an algorithm A_v which only permits one connection with an end equal to v, and an algorithm A_w which only permits one connection with an end equal to w. By Proposition 2.1, the AAP algorithm is $O(\log n)$-competitive; the algorithms A_v and A_w are also $O(\log n)$-competitive since the maximum fractional weight of connections that the optimum can accept originating at any one node of $\mathcal{N}$ is $O(\log n)$. Thus, applying Proposition 2.3 with $c^* = O(\log n)$ and $d = 3$, we see that the combined routing algorithm is $O(\log n)$-competitive against the fractional optimum. ■

Our on-line algorithm in G simply runs the above simulation; whenever $(\psi(s_i), \psi(t_i))$ is accepted, it routes the pair (s_i, t_i) in G using the paths constructed in Lemma 3.5. The following lemma says that it will not run out of "bandwidth" while doing this.

Lemma 3.7 *The algorithm in G can route all the connections accepted by the simulation.*

Proof. When the simulation accepts $(\psi(s_i), \psi(t_i))$, it specifies a sequence of neighboring clusters $C_{v_1}, C_{v_2}, \ldots, C_{v_r}$, where $v_1 = \psi(s_i)$ and $v_r = \psi(t_i)$.

The algorithm in G routes (s_i, t_i) by simply moving from one cluster in this sequence to the next using paths of the form $Y^u_{v_i}$. More concretely, it first chooses any $w \in \sigma'_{v_1}$ and $w' \in \sigma'_{v_r}$ and sets $Z_0 = Y^w_{v_1}$ and $Z_r = Y^{w'}_{v_r}$. Then for each $j = 1, \ldots, r-1$, it chooses any edge $(u, u') \in \tau_{v_j, v_{j+1}}$ that has not yet been used for routing, and sets $Z_j = Y^u_{v_j} \cup Y^{u'}_{v_{j+1}}$. Since there are at least $\rho \log n$ such edges available, and the simulation accepts at most $\rho \log n$ pairs whose routes use the edge in $\mathcal{N}$ from v_j to v_{j+1}, this is always possible. Now, by Lemma 3.5, the union $Z_0 \cup \cdots \cup Z_r$ contains a path from the boundary of C_{v_1} to the boundary of C_{v_r}. ■

Finally, we have to show that optimum in G is not far from the optimum in the simulation.

Lemma 3.8 *For any realizable subset $\mathcal{T}'$ of $\mathcal{T}_M$, $\psi(\mathcal{T}')$ can be routed in $\mathcal{N}(5\gamma \log n)$.*

Proof. For each s_i-t_i path P in the optimal routing, construct the following path for $(\psi(s_i), \psi(t_i))$ in $\mathcal{N}$ — when P crosses from D_w into $D_{w'}$, add an edge from w to w'. Since $|\delta(D_w, D_{w'})| \leq 5\gamma \log n$, so at most this many paths in the constructed routing will use the edge (w, w') in $\mathcal{N}$. ■

Now, since the on-line algorithm is $O(\log n)$-competitive against the fractional optimum in $\mathcal{N}(\rho \log n)$, it is also $O(\log n)$-competitive against the fractional optimum in $\mathcal{N}(5\gamma \log n)$, which by Lemmas 3.2 and 3.8 is at least as large as the maximum realizable subset of $\mathcal{T}$. Thus we are $O(\log n)$-competitive in routing long connections.

3.2.2 Routing Short Connections

To handle short connections, we require the following two facts.

Proposition 3.9 *Let $H = (V, E)$ be an arbitrary graph of diameter d. Then there is a deterministic on-line* MDPP *algorithm that is $2 \cdot \max(d, \sqrt{|E|})$-competitive.*

Proof. Let $m = |E|$. The algorithm maintains a sequence of graphs $H_1, H_2, \ldots$ as follows. $H_1 = H$. The algorithm always routes the first request on a shortest path P_1, and sets $H_2 = H_1 \setminus P_1$. In general, when presented with request (s_i, t_i), the algorithm routes it on a shortest path P_i in H_i *if* $d(s_i, t_i) \leq \sqrt{m}$ in H_i. It then sets $H_{i+1} = H_i \setminus P_i$. Let p denote the total number of paths routed by the algorithm.

Let $d' = \max(d, \sqrt{m})$. Consider any routing for $\mathcal{T}$, consisting of paths $Q_1, \ldots, Q_q$. At most d' of the Q_j intersect each P_i, since the Q_i are all edge-disjoint and $|P_i| \leq d'$. Also, at most $\sqrt{m}$ of the Q_j fail to intersect any of the P_i, since the pair (s_j, t_j) associated with Q_j must have been rejected by the on-line algorithm, and hence $|Q_j| > \sqrt{m}$. Thus we have $q \leq d'p + \sqrt{m} \leq 2d'p$. ■

Lemma 3.10 *Let $X = G[(i-r):(i+r), (j-r):(j+r)]$ for some natural numbers i, j, and r; let $Y = G[(i-2r):(i+2r), (j-2r):(j+2r)]$; and let $\mathcal{T}'$ be a set of terminal pairs with both ends in X. Then the maximum size of a subset of $\mathcal{T}'$ that is realizable in $G[Y]$ is at least $\frac{1}{4}$ the maximum size of a subset of $\mathcal{T}'$ that is realizable in G.*

Proof. Let $\mathcal{T}'' \subset \mathcal{T}'$ denote a realizable set of maximum size, and consider the set of paths in a routing of $\mathcal{T}''$. Consider the set of paths leaving X. Since $|\delta(X)| \leq 8r$, there are at most $4r$ such paths. Delete the portion of each path between its first and last intersection with $\delta(X)$; using the r disjoint rings in $G[Y \setminus X]$ (as defined in the previous section), we can connect the resulting "breakpoints" of at least r of these paths using edge-disjoint paths in $G[Y \setminus X]$. Thus, we are still routing at least $\frac{1}{4}$ of $\mathcal{T}''$ in $G[Y]$. ■

The algorithm for short connections is now as follows. Set $r = 32\gamma \log n$, and run Luby's algorithm to find a subset M' of V, maximal subject to the property that $L_\infty(u, v) \geq 4r$ for $u, v \in M'$. For $u \in M'$, define X_u to be the set of all vertices whose L_∞ distance from u is at most r, and Y_u to be the set of all vertices whose L_∞ distance from u is at most $2r$.

Lemma 3.11 *If (s_i, t_i) is a short connection, then with constant probability there is a u such that $s_i, t_i \in X_u$.*

Proof. A sufficient condition for some X_u containing both s_i and t_i to enter the MIS in the first iteration is that among all vertices within L_∞ distance $4r$ of s_i, the one that picks the highest number has L_∞ distance at most $\frac{1}{2}r$ from s_i. This happens with constant probability. ■

Let $\mathcal{T}_u$ denote the set of short connections both of whose ends lie in X_u. We now run the algorithm of Proposition 3.9 on the (disjoint) subgraphs $G[Y_u]$ simultaneously, using the $\mathcal{T}_u$ as the sets of terminal pairs. By Proposition 3.9 and Lemma 3.10, we are $O(\log n)$-competitive in each subgraph. Thus, by Lemma 3.11, we are $O(\log n)$-competitive in routing short connections. Thus,

Theorem 3.12 *The on-line algorithm is $O(\log n)$-competitive in the two-dimensional mesh.*

3.3 The Off-Line Algorithm

For the constant-factor off-line approximation, we use a variant of the graph $\mathcal{N}$. In $\mathcal{N}(\gamma \log n)$, for any fixed constant γ, one can obtain a constant-factor approximation to the MDPP by the following *randomized rounding* algorithm of Raghavan and Thompson [23, 22]. First we solve the fractional relaxation of the MDPP instance (this can be done in polynomial time); from this, we obtain for each terminal pair (s_i, t_i) a collection of paths $P_i^1, \ldots, P_i^z$ and associated weights $y_i^1, \ldots, y_i^z \in [0,1]$ such that $x_i = \sum_j y_i^j \in [0,1]$. We now pick a *scaling factor* $\mu < 1$; independently for each terminal pair (s_i, t_i) we route it on path P_i^j with probability μy_i^j, and don't route it at all with probability $1 - \mu x_i$. If we do route it, we say that (s_i, t_i) has been *rounded up*. In [23, 22], it is shown that with constant probability, no capacity is violated by the selected paths, and the number of pairs that are rounded up is a constant fraction of the fractional optimum.

In particular, we require the following theorem from [22].

Theorem 3.13 (Raghavan) *Let $X_1, X_2, \ldots, X_r$ be independent Bernoulli trials with $EX_j = p_j$ and $\Psi = \sum_i X_i$; so $E\Psi = m = \sum_i p_i$. Then for $\delta > 0$ we have*

$$Pr[\Psi > (1+\delta)m] < \left[\frac{e^\delta}{(1+\delta)^{(1+\delta)}}\right]^m.$$

We specialize this to the form in which we will use it as follows.

Corollary 3.14 *Let $0 < \mu < 1$, and $p_1, \ldots, p_r \in [0,1]$. Let $X_1', X_2', \ldots, X_r'$ be independent Bernoulli trials with $EX_j' = \mu p_j$ and $\Psi' = \sum_i X_i'$. Let $m = \sum_i p_i$. Then*

$$Pr[\Psi' > m] < (e\mu)^m.$$

Proof. Apply the bound of Theorem 3.13 with m set to μm and δ set to $\mu^{-1} - 1$. ∎

Let us consider how to use this randomized rounding approach in routing long connections off-line. In the high-capacity network $\mathcal{N}$, this rounding approach is fine; but to get a constant-factor approximation we also have to be within a constant factor of the optimum in routing terminals out of the clusters (in the on-line algorithm it was enough to route only one). To this end, we build the following more complicated network $\mathcal{N}'$. Let z_v denote the node representing $v \in M$ in the network $\mathcal{N}$; we construct $\mathcal{N}'$ by attaching C_v to z_v via an edge from z_v to *each* node in $\pi(C_v)$. Let $\mathcal{N}'(\gamma)$ denote the network $\mathcal{N}'$ in which each edge *between nodes of the subgraph* $\mathcal{N}$ has capacity γ, and all other edges have unit capacity.

Let f^* denote the value of the fractional optimum solution to the MDPP in $\mathcal{N}'(\rho \log n)$. We now run the randomized rounding algorithm on $\mathcal{N}'(\rho \log n)$. With high probability the parts of all selected paths lying in the subgraph $\mathcal{N}(\rho \log n)$, taken together, do not violate any capacity constraint; and the number of pairs that are rounded up is within a constant factor of the fractional optimum f^*. We now must convert the selected paths in $\mathcal{N}'(\rho \log n)$ into s_i-t_i paths in G. We can use the technique of the previous section to produce, for each selected pair (s_i, t_i), a "global" path P_i that begins at $\tau'_{\psi(s_i)} \subset \pi(D_{\psi(s_i)})$ and ends at $\tau'_{\psi(t_i)} \subset \pi(D_{\psi(t_i)})$.

The real problem is how to find paths *within* the clusters such that each s_i (resp. t_i) that has been rounded up can reach the endpoint of this associated global path on $\tau'_{\psi(s_i)}$ (resp. $\tau'_{\psi(t_i)}$). For this, the paths returned by the randomized rounding are of no value, since the edges of $\mathcal{N}'$ within the clusters C_v have only unit capacity. Instead we argue as follows.

Let S_v denote the set of terminals in C_v that are rounded up. Each is trying to "get out" to its associated path that begins at τ'_v. Given the crossbar in $C^*_v \setminus C_v$, it is sufficient to route each terminal in S_v to any vertex in σ'_v.

So this leaves us with the following *escape problem.* We are given the set S_v of terminals that have been rounded up, and we want to route a large fraction of them to σ'_v. The following lemma, whose proof contains the central step of the algorithm, says that this can be done.

Lemma 3.15 *For a sufficiently small (constant) value of μ, there is a constant $c < 1$ so that we can find sets $S'_v \subset S_v$, such that*

(i) if one end of a pair (s_i, t_i) belongs to $\cup_v S'_v$ then so does the other,

(ii) each set S'_v can be linked to $\sigma'_v \subset \pi(C_v)$ via edge-disjoint paths.

(iii) the probability that $|\cup_v S'_v| > cf^*$ is at least a constant, where the probability is taken over the randomized rounding that produced the sets S_v.

Proof. We will first construct such a set with condition (iii) weakened to the requirement that each S'_v can be linked to $\pi(C_v)$ via edge-disjoint paths (rather than to the smaller set σ'_v). This is sufficient to imply the lemma, as follows. Suppose we obtain such sets S''_v. For each $s \in S''_v$, identify the vertex in σ'_v closest to s. We now build a graph on the set of terminal pairs in $\cup_v S''_v$, joining two if at either end they have the same closest vertex in some set σ'_v. We claim this graph has degree at most $8\gamma\rho^{-1}$; for the spacing between vertices of σ'_v on the boundary of C_v is $4\gamma\rho^{-1}$, and so at most twice this number of terminal pairs can compete with some (s_i, t_i) (at either end) for the same vertex of σ'_v. Thus this graph has an independent set of size at least $\frac{1}{8\gamma\rho^{-1}+1}|\cup_v S''_v|$; if we let the terminal pairs in this large independent set constitute the sets S'_v, we satisfy the requirements of the lemma.

This allows us to deal with a standard escape problem on a rectangular mesh: each terminal is allowed to "escape" to any vertex on the boundary. First observe the following fact: an escape problem on a rectangular mesh is feasible if and only if, for all p, q, any subrectangle of size $p \times q$ contains at most $2(p + q)$ terminals. To see this, note that we can reduce the escape problem to a single-source/single-sink maximum flow problem, and thus only have to verify the cut condition. On a rectangular mesh, the smallest rectangle enclosing any connected cut has no greater capacity, and contains at least as many terminals, as the original cut; thus the cut condition holds if and only if it holds for all subrectangles.

This suggests the following algorithm to construct the set S''_v. Call a rectangle *overfull* if it violates the cut condition; we go through each $s \in S_v$, deleting it if it is contained in any overfull rectangle — we also then delete its matching terminal in some other cluster. This results in a set S''_v, which by the argument of the previous paragraph can be completely routed to $\pi(C_v)$ on edge-disjoint paths.

It remains to lower-bound the expected size of $\cup_v S''_v$. Say that a terminal s *survives* if

(i) it is initially rounded up, and hence included in a set S_v, and

(ii) it is not deleted because it or its matching terminal is contained in an overfull rectangle.

So what is the probability that a terminal s, which is given weight $f_s \in [0,1]$ by the fractional optimum, survives? First we consider the probability that s is contained in a *fixed* overfull $p \times q$ rectangle R, *given* that s has been rounded up. In order for this rectangle R to become overfull, it must be that at least $2(p+q)$ terminals in R other than s were rounded up. But since the un-scaled fractional flow was feasible, the total fractional weight contained in R is at most $2\mu(p+q)$. Thus, setting $y = (e\mu)^2$, the probability that the rectangle R becomes overfull after rounding, given that s is rounded up, is at most y^{p+q}.

Now, since s is contained in at most pq rectangles of dimensions $p \times q$, the probability that s is contained in any overfull rectangle after rounding, given that it is rounded up, is clearly bounded by the infinite sum

$$\sum_p \sum_q pqy^{p+q} = \frac{y^2}{(1-y)^4}.$$

s can also be deleted if its matching terminal is contained in an overfull rectangle, so the probability of s being deleted, after having been rounded up, is at most $\frac{2y^2}{(1-y)^4}$. By taking μ small enough, we can make this last expression a constant less than $\frac{1}{2}$.

Finally, the probability that s survives is equal to the probability that it is rounded up, which is μf_s, times the probability that it is not deleted after being rounded up, which is at least $\frac{1}{2}$. Thus, the expected size of $\cup_v S''_v$ is at least $\frac{1}{2}\mu f^*$, for a small enough constant value of μ. So by Markov's inequality, the probability that $|\cup_v S''_v| \geq \frac{1}{4}\mu f^*$ is at least an absolute constant; and the lemma follows. ■

To turn the above lemma from a statement holding with constant probability to one with high probability, we repeat the randomized rounding $O(\log n)$ times and take the best routing obtained. A single run of randomized rounding fails if one of the following two bad events happens: (1) one of the high capacity edges of the simulated network is used by too many selected paths, or (2) the set $\cup_v S'_v$ selected by Lemma 3.15 is too small. The first event has inverse polynomially small probability, and the probability of the second is bounded by a constant, so the total probability of any bad event is bounded by a constant less than 1. This implies that out of the $O(\log n)$ tries, one must succeed with high probability.

This gives a constant-factor approximation for long connections: if (s_i, t_i) is rounded up, and $s_i, t_i \in \cup_v S'_v$, then we concatenate the paths from s_i to $\pi(C_{\psi(s_i)})$ (given by Lemma 3.15) to $\pi(D_{\psi(s_i)})$ (given by the crossbar in $C^*_v \setminus C_v$) to $\pi(D_{\psi(t_i)})$ (given by the edges of the path in $\mathcal{N}(\rho \log n)$, joined together with crossbars as in Lemma 3.7), and now symmetrically to $\pi(C_{\psi(t_i)})$ and to t_i.

We handle short connections as in Section 3.2.2. That is, we use a randomized algorithm to construct subsquares $X_u \subset Y_u$, with all Y_u disjoint, and only handle short connections both of whose ends lie in a single X_u. In this case, however, we now run the above algorithm recursively on each $G[Y_u]$.

Call a connection "medium" if it is now a long connection in this recursive call, and "small" otherwise. Medium connections are handled as described above. Small connections take place within clusters of size $O(\log\log n)$ and therefore can be simply solved to optimality by brute force. We can then take the largest realizable set we find among the long, medium, and small connections, obtaining

Theorem 3.16 *There is a randomized (off-line)* MDPP *algorithm in the two-dimensional mesh that produces a constant-factor approximation with high probability.*

4 Densely Embedded Graphs: Definition and Properties

We now want to extend the above algorithm to any graph that is sufficiently "mesh-like." We define the class of graphs here; in the following section we show how to extend the routing algorithms to this class.

We will need some additional notation and definitions: Recall that for $u, v \in V$, $d(u, v)$ denotes the least number of edges in a u-v path. Let $B_r(v) = \{u : d(u, v) \leq r\}$. Also recall that $\delta(X)$ and $\pi(X)$ denote the set of edges leaving X and the set of vertices of X incident to edges in $\delta(X)$. We write $X^o = S \setminus \pi(X)$. Observe that removing $\pi(X)$ from X disconnects it from the rest of the graph, and $\pi(B_r(v))$ consists of vertices at exactly distance r from v. If C is a connected subset of $G \setminus X$, we use $\Gamma(X, C)$ to denote the (unique) connected component of $G \setminus X$ containing C. The set of vertices in $\pi(X)$ which have a neighbor in $\Gamma(X, C)$ will be called the *segment of* $\pi(X)$ *bordering* C and denoted $\sigma(X, C)$. We say that a set $X \subset V$ is *simple* if $G \setminus X$ is connected.

Definition 4.1 *A graph H is an α-semi-expander if for every $X \subset V(H)$ for which $|X| \leq \frac{1}{2}|V(H)|$, we have $|\delta(X)| \geq \alpha\sqrt{|X|}$.*

Since our goal is to generalize the two-dimensional mesh, let us note the following properties of the mesh.

(i) It is a planar graph with bounded degree, and (aside from one "exceptional face") it is Eulerian and has bounded face size.

(ii) It is an α-semi-expander, for a constant $\alpha > 0$ based on the ratio of the two side lengths of the mesh.

(iii) Square sub-meshes of the mesh satisfy (i) and (ii).

In the arguments to follow, it is quite cumbersome — though not technically difficult — to deal with "exceptional faces" of the type in (i). Thus, in the following section we will work with the more restricted class of *uniformly densely embedded graphs*, where *all* faces have bounded size; and we will assume further assume that G is Eulerian. In Section 6.2, we show how to handle graphs with an exceptional face; in this way, our class of graphs will include the two-dimensional mesh.

First we need some preliminary topological definitions. Let Σ denote a compact orientable surface; it is well-known (see e.g. [18]) that Σ may be obtained from the 2-sphere by attaching a finite number of handles. By a *disc* we mean a set homeomorphic to the closed unit ball in $\mathbf{R}^2$, and by Σ-*disc*, we mean a subset of Σ homeomorphic to a disc. Our definition of graph embedding is standard; a *face* of an embedded graph G is a connected component of $\Sigma \setminus G$, and we say G is *strongly embedded* if the closure of each face is a Σ-disc, and each face is bounded by a simple cycle of G. Finally, we also use the the terms *curve* (continuous image of $[0,1]$) and *closed curve* (continuous image of S^1).

Our class of graphs is defined to satisfy analogues of properties (i), (ii), and (iii) locally.

Definition 4.2 *A graph $G = (V, E)$ is* uniformly densely embedded *with parameters α, λ, Δ, and ℓ if:*

(i) G is strongly embedded on a compact orientable surface Σ, it has maximum degree Δ, and each face is bounded by at most ℓ edges.

(ii) For each $r \leq \lambda \log n$ and each $v \in V$, the drawing of $G[B_r(v)]$ is contained in a Σ-disc.

(iii) For each $r \leq \lambda \log n$ and each $v \in V$, the graph $G[B_r(v)]$ is an α-semi-expander.

Thus, for the remainder of the paper aside from Section 6.2, we will assume that G is a simple Eulerian graph that is *uniformly densely embedded* on a surface Σ with parameters α, λ, Δ, and ℓ. In Section 6.2, we show how our algorithms can be adapted to handle graphs satisfying the following weaker definition; it is the same as the definition above, except that we allow an exceptional face.

Definition 4.3 *A graph $G = (V, E)$ is* densely embedded *and* nearly-Eulerian *with parameters α, λ, Δ, and ℓ if:*

(i) G is strongly embedded on a compact orientable surface Σ and has maximum degree Δ.

(i)′ G contains a face Φ^* such that all faces other than Φ^* are bounded by at most ℓ edges, and every vertex not incident to Φ^* has even degree.

(ii) For each $r \leq \lambda \log n$ and each $v \in V$, the drawing of $G[B_r(v)]$ is contained in a Σ-disc.

(iii) For each $r \leq \lambda \log n$ and each $v \in V$, the graph $G[B_r(v)]$ is an α-semi-expander.

The classes of graphs satisfying these definitions are incomparable to the class considered in our earlier paper [14]. The *semi-expansion* condition above will be shown to imply the *uniformly high-diameter* condition of [14] (see Lemma 4.4); however, in the current paper, we only require planarity and semi-expansion locally, and essentially no restrictions are placed here on the "global" structure of the graph. The examples of uniformly high-diameter graphs constructed in [14] are densely embedded as well; and in Section 4.2 we will discuss some related classes of graphs that are densely embedded.

4.1 Some Basic Properties

We now show that our definition implies G has some additional mesh-like properties. First of all, for any $v \in V$ and $r \leq \lambda \log n$, the fact that $G[B_r(v)]$ is a bounded-degree semi-expander implies that the set $B_r(v)$ has size at least quadratic in r; by also using the planarity of $G[B_r(v)]$, one can show an analogous upper bound. We summarize this as follows.

Lemma 4.4 *There are constants $\bar{\alpha}$ and β depending only on α and Δ such that the following holds. For each $r \leq \lambda \log n$ and each $v \in V$, we have $\bar{\alpha} r^2 \leq |B_r(v)| \leq \beta r^2$.*

Proof. Fix $r \leq \lambda \log n$ and $v \in V$, and let $S = B_r(v)$. To see the lower bound, note that for any $i \leq r$, if $x_i = |B_i(v)|$, then by the semi-expansion of H we have

$$x_i \geq x_{i-1} + \frac{\alpha}{\Delta - 1}\sqrt{x_{i-1}}.$$

For at least $\alpha\sqrt{x_{i-1}}$ edges leave $B_{i-1}(v)$, and at most $\Delta - 1$ are incident to any one vertex. Let $\nu = \frac{\alpha}{\Delta-1}$; now one verifies by induction that $x_i \geq \frac{1}{16}\nu^2 i^2$:

$$\begin{aligned} x_i &\geq \frac{1}{16}\nu^2(i-1)^2 + \frac{1}{4}\nu^2(i-1) \\ &= \frac{1}{16}\nu^2(i-1)(i+3) \\ &\geq \frac{1}{16}\nu^2 i^2. \end{aligned}$$

To see the upper bound, we observe that $G[S]$ is planar and has diameter at most $2r$. Let $n = |S|$. By the Lipton-Tarjan planar separator theorem [16], there is a set of at most $4r+1$ vertices whose removal breaks H into components each of size at most $\frac{2}{3}n$. Let X be a union of these components of size between $\frac{1}{3}n$ and $\frac{1}{2}n$. Then

$$\begin{aligned} \frac{\alpha\sqrt{n}}{\sqrt{3}} &\leq \alpha\sqrt{|X|} \\ &\leq |\delta(X)| \\ &\leq (\Delta - 1)(4r+1) \end{aligned}$$

from which the result follows. ∎

The following two facts are quite useful; the first essentially relates the size of the "perimeter" of a set $B_r(v)$ ($r \leq \lambda \log n$) to its radius.

Lemma 4.5 *Let $c > 1$ and r a positive integer be such that $cr < \lambda \log n$. Then for some r' between r and cr, we have $|\pi(B_{r'}(v))| \leq \beta \cdot \frac{c^2}{c-1} \cdot r$.*

Proof. Since $\pi(B_{r'}(v)) \subset \{u : d(v,u) = r'\}$, the sets $\pi(B_r(v)), \pi(B_{r+1}(v)), \ldots, \pi(B_{cr}(v))$ are all disjoint and contained in $B_{cr}(v)$. Since $|B_{cr}(v)| \leq \beta c^2 r^2$, one of these sets has size at most $\beta \cdot \frac{c^2}{c-1} \cdot r$. ∎

Using this, we show that we can extend any small enough set U to a simple set with at most a constant-factor increase in its radius. Recall from the beginning of Section 4 that a set X is simple if $G \setminus X$ is connected.

Lemma 4.6 *There is a constant ξ such that the following holds. Let $U \subset B_r(v)$, where $r \leq \frac{1}{\xi}\lambda \log n$. Then there is a component Γ of $G \setminus U$ and a planar simple set U' such that $U \subset U' \subset B_{\xi r}(v)$, $G \setminus U' = \Gamma$, and $\sigma(U, \Gamma) = \sigma(U', \Gamma)$.*

Proof. Choose r' between r and $2r$ so that $|\pi(B_{r'}(v))| \le 4\beta r$. Let $U_0 = B_{r'}(v)$, and $G \setminus U_0$ have components $\Gamma_1, \ldots, \Gamma_p$.

Now set $s = 8\bar{\alpha}^{-1/2}\alpha^{-1}\beta\Delta$ and $\xi = 2s + 2$. We claim that all but one of the components Γ_i are contained in $B_{\xi r}(v)$. For suppose not; then for $i \ne j$ there are $w \in \Gamma_i$ and $w' \in \Gamma_j$ such that w and w' are each at distance s from U_0, $B_s(w) \subset \Gamma_i$, and $B_s(w') \subset \Gamma_j$. Now consider the edge cut of size at most $4\beta\Delta r$ formed by $\delta(U_0)$; one of the two spheres $B_s(w)$ and $B_s(w')$, say the latter, is contained in a small component of this cut in $G[B_{\xi r}(v)]$. But then the semi-expansion of $G[B_{\xi r}(v)]$ requires that $4\beta\Delta r \ge \alpha\sqrt{|B_s(w')|}$, which is a contradiction since by Lemma 4.4 we have $|B_s(w')| \ge \bar{\alpha}s^2$.

So for some i, only Γ_i is not contained in $B_{\xi r}(v)$. Now let $\Gamma'_1, \ldots, \Gamma'_q$ denote the components of $G \setminus U$; so Γ_i is contained in one of these, say Γ'_1, and $\Gamma'_2, \ldots, \Gamma'_q$ are all contained in $B_{\xi r}(v)$. Thus we have

$$U' = U \cup \bigcup_{j>1} \Gamma'_j \subseteq B_{\xi r}(v).$$

In particular, U' is planar since $\xi r \le \lambda \log n$, and it is simple since $G \setminus U'$ has only the component Γ'_1. Thus U' satisfies the conditions of the lemma. ■

Finally, we show a general property of planar graphs H with small face size: if the distance between two nodes in H is large, than the value of any edge cut which contains both in the same segment of its boundary must also be relatively large.

Lemma 4.7 *Let H be a planar graph, with distinguished faces $\Phi_1, \ldots, \Phi_r$ bounded by cycles $Q_1, \ldots, Q_r$ respectively. Suppose that all faces other than $\Phi_1, \ldots, \Phi_r$ are bounded by at most ℓ edges, and for a constant d' and all $i \ne j$ we have $d(Q_i, Q_j) \ge d'$.*

Let $U \subset V(H)$ and $v, w \in \sigma(U, C)$ for some component C of $G \setminus U$. Then

$$|\delta(U)| \ge \min(\ell^{-1}d', \ell^{-1}d(v, w)).$$

Proof. Let $S = \sigma(U, C) \subset U$. In the graph $H[U]$, S lies on a single facial cycle Q. Traversing Q in a clockwise direction starting at v, we encounter faces $R_1, \ldots, R_p$ whose boundaries contain vertices both of U and of $H \setminus U$.

Suppose that among the $\{R_i\}$ there are two distinct large faces Φ_m and $\Phi_{m'}$. Choose such a pair for which $R_a = \Phi_m$, $R_b = \Phi_{m'}$, and $R_c \notin \{\Phi_i\}$ for $a < c < b$. Let P denote the corresponding maximal subpath of Q whose internal vertices are incident only to faces R_c, for $a < c < b$. Then among every ℓ consecutive vertices of P, there must be one incident to an edge in $\delta(U)$; since $|P| \ge d'$ by the hypotheses of the lemma, this implies the claimed bound.

Otherwise, there is a single large face Φ_m among the $\{R_i\}$; note that Φ_m may appear several times on the traversal of Q. Now there are two sub-paths of Q from v to w, which we denote P_0 and P_1. Since v, w border the same component of $H \setminus U$, the face Φ_m does not appear in a traversal of one of P_0 or P_1 — suppose it is P_0. So as above, among every ℓ consecutive vertices of P_0, there must be one incident to an edge in $\delta(U)$; and we have $|P_0| \ge d(v, w)$. ■

4.2 Related Classes of Graphs

In this section, we show a natural construction which produces uniformly densely embedded graphs; it is an extension of the definition of *geometrically well-formed graphs* in our earlier paper [14]. The material in this section is independent of the rest of the paper.

We wish to define a notion of a *surface* being *locally planar*, in the following sense. Let Σ be a compact orientable surface, embedded in $\mathbf{R}^3$. For $x \in \Sigma$, let $B'_d(x)$ denote the set of all points of Σ whose distance from x (as measured on Σ) is at most d.

Definition 4.8 *A set $X \subset \Sigma$ is (γ_0, γ_1)-*flat *for some positive constants γ_0, γ_1, if there is a Σ-disc D such that*

(i) $X \subseteq D$.

(ii) For all points $x \in X$ and $s \geq 0$ such that $B'_s(x) \subset X$, the surface area of $B'_s(x)$ is at least $\gamma_0 s^2$ and at most $\gamma_1 s^2$.

(iii) For all Σ-disc D' such that $D' \subseteq D$, if D''s boundary has length s, then the surface area of D' is at most $\gamma_1 s^2$.

*We say that Σ is (r, γ_0, γ_1)-*locally flat *if it is orientable, and for all $x \in \Sigma$ the set $B'_r(x)$ is (γ_0, γ_1)-flat.*

Of course all these properties hold if Σ, for example, is the unit sphere in $\mathbf{R}^3$.

Now we say that a graph is *locally well-formed* if it is drawn on a locally flat surface, and each face has geometrically about the same (small) size.

Definition 4.9 *A graph H drawn on Σ is* locally well-formed *with constant parameters $\Delta, \ell, \gamma_0, \gamma_1, \rho_0, \rho_1$ if it has maximum degree Δ and there is an $r > 0$ such that*

(i) Σ is $(r \log n, \gamma_0, \gamma_1)$-locally flat,

(ii) The maximum number of edges on a face in the drawing of H is ℓ, and

(iii) for each face Φ of G there is an $x \in \Sigma$ so that $B'_{\rho_0 r}(x) \subset \Phi \subset B'_{\rho_1 r}(x)$.

We now want to show that every locally well-formed graph is uniformly densely embedded. To show this, the following routine lemma is useful: in Definition 4.1 it is enough to require semi-expansion for cuts that produce only two components.

Lemma 4.10 *H is an α-semi-expander if and only if the condition of Definition 4.1 holds for all sets X for which $H[X]$ and $H \setminus X$ are both connected.*

Proof. We proceed by induction on the number of connected components of $H[X]$ and $H \setminus X$. Assume $H[X]$ is not connected, and let $\Gamma_1, \ldots, \Gamma_p$ be components. Then each of the sets $\Gamma_1, \ldots, \Gamma_p$ must satisfy Definition 4.1 by the induction hypothesis. From this we get

$$|\delta(X)| = \sum_i |\delta(\Gamma_i)| \geq \alpha \sum_i \sqrt{|\Gamma_i|} \geq \alpha \sqrt{|\cup_i \Gamma_i|} = \alpha \sqrt{|X|}.$$

If H connected but $H \setminus X$ is not, and each connected component has size at most $\frac{1}{2}|V(H)|$, then the above argument applies to the components of $H \setminus X$. Otherwise, apply Definition 4.1 to the cut defined by the single large component, both of whose sides are connected. ∎

Proposition 4.11 *If G is locally well-formed with parameters $\Delta, \ell, \gamma_0, \gamma_1, \rho_0, \rho_1$, then there are positive constants α and λ such that G is uniformly densely embedded with parameters α, λ, Δ, and ℓ.*

Proof. Let G be locally well-formed with the given parameters. Then for any $v \in V$, if $s \leq \rho_1^{-1} \log n$, the set $B_s(v)$ is contained in $B'_{r \log n}(v)$ and hence in a Σ-disc. Now let $X \subset B_s(v)$; we wish to show that it satisfies the semi-expansion requirement in $G[B_s(v)]$. By Lemma 4.10, we may assume that both $G[X]$ and $G[B_s(v) \setminus X]$ are connected. Thus $\delta(X)$ lies on a single face of $G[X]$. Let $q = |\delta(X)|$; then there is a closed curve L on Σ of length at most $\rho_1 r q$ that bounds a Σ-disc containing X. Thus, X is contained in a disc of area at most $\gamma_1 \rho_1^2 r^2 q^2$. But each face in $G[X]$ has area at least $\gamma_0 \rho^2 r^2$, so X has at most $\gamma_1 \rho_1^2 \gamma_0^{-1} \rho_0^{-2} q^2$ faces, and hence at most ℓ times this many vertices. ■

In a series of papers proving, among other things, that the disjoint paths problem for a fixed number of terminal pairs is solvable in polynomial time [27], Robertson and Seymour make use of another notion of "denseness" of surface embeddings — namely *representativity.* It turns out that our definition of uniformly densely embedded graphs could also have been expressed in these terms. We say that a closed curve on Σ is *null-homotopic* if it is homotopic to a point; it is well-known (see e.g. [25]) that a closed curve is null-homotopic if and only if it is contained in a Σ-disc. Now we say that a drawing of G on Σ is *c-representative* [25, 26] if any non-null-homotopic closed curve on Σ meets the drawing of G at least c times.

In this terminology, we could have replaced the condition that each $G[B_r(v)]$ ($r \leq \lambda \log n$) be contained in a Σ-disc by the condition that the drawing of G be $\Omega(\log n)$-representative. More precisely,

Proposition 4.12 *If G satisfies parts (i) and (iii) of Definition 4.2, and the drawing of G is $(\lambda \log n)$-representative, then there is a constant λ' such that G is uniformly densely embedded with parameters α, λ', Δ, and ℓ. Conversely, if G is uniformly densely embedded with parameters α, λ, Δ, and ℓ, then there is a constant λ' such that the drawing of G is $(\lambda' \log n)$-representative.*

Proof. The converse statement is easier. If G is uniformly densely embedded with parameters α, λ, Δ, and ℓ, then any closed curve $\mathcal{R}$ on Σ meeting G at fewer than $\ell^{-1}\lambda \log n$ vertices meets it only at vertices contained in $B_{\lambda \log n}(v)$ for some $v \in V$. Thus $\mathcal{R}$ is contained in a Σ-disc and is null-homotopic.

Now suppose G satisfies parts (i) and (iii) of Definition 4.2, and the drawing of G is $(\lambda \log n)$-representative. We must show that for some λ', every $G[B_{\lambda' \log n}(v)]$ is drawn in a Σ-disc. Choose $\lambda' < \frac{1}{4}\lambda$ and let $U = B_{\lambda' \log n}(v)$ for some $v \in V$. We claim that every simple cycle of $G[U]$ is null-homotopic in Σ. For suppose not, and choose the shortest non-null-homotopic cycle Q contained in $G[U]$. Say for simplicity that Q contains an even number of vertices, $v_0, \ldots, v_k, \ldots, v_{2k} = v_0$, and let Q_0 and Q_1 denote the two sub-paths of Q with ends equal to v_0 and v_k. Now suppose there were some path P in $G[U]$ with ends equal to v_0 and v_k of length less than k; then one of $Q_0 \cup P$ or $Q_1 \cup P$ would contain a non-null-homotopic simple cycle of $G[U]$ of length less than $2k$, contradicting our choice of Q. Thus $k \leq 2\lambda' \log n$, and so $|Q| \leq 4\lambda' \log n < \lambda \log n$. Now since G is strongly embedded, there is a simple closed curve $\mathcal{R}$ on Σ, meeting G at precisely the vertices of Q, that is homotopic to Q in Σ; but since $\mathcal{R}$ meets G fewer than $\lambda \log n$ times, it is null-homotopic in Σ. This contradicts our assumption that Q is non-null-homotopic.

Thus $G[U]$ contains only null-homotopic simple cycles. By Theorems (11.2) and (11.10) of [25], this implies that $G[U]$ is contained in a Σ-disc. ∎

5 Densely Embedded Graphs: The Routing Algorithms

One encounters a number of difficulties in extending the algorithms of Section 3 to densely embedded graphs in general. Some of these are easily taken care of — for example, we cannot define "subsquares" of G anymore; but we can use balls of the form $B_r(v)$ instead, and we have seen above that these behave in much the same way. We similarly may choose a maximal set of mutually distant balls and grow enclosures around them. The major problems are the following. (1) We used the natural crossbars inside a mesh for routing; do these enclosures have similar crossbars inside them? (2) Where is the high-capacity simulated network $\mathcal{N}$?

To build crossbars inside the enclosures we use the Okamura-Seymour theorem [20], analogously to a construction in our earlier paper [14]. To define the high-capacity simulated network $\mathcal{N}$ we want to grow the enclosures out until they touch. However, at this point their boundaries might not be "smooth" enough to allow us to build crossbars inside them; additionally, there is no reason for enclosures that do touch to have $\Omega(\log n)$ edges in their common boundary.

Nevertheless, it is still possible to build a simulated network $\mathcal{N}$, as follows. We grow enclosures that have smooth boundaries, and are large enough that they contain large crossbars, but we keep them mutually distant from one another. Then we define the notion of a *Voronoi partition* of G to allow us to determine which clusters are "neighbors"; we define the simulated network $\mathcal{N}$ by putting $\Omega(\log n)$ parallel edges between neighboring clusters (whether or not they have that many edges in the common boundary of their Voronoi regions).

We show that the collection of these parallel edges "represents" the graph G well enough that it can be used as the network $\mathcal{N}$. In particular, we need to show that all paths accepted by the simulated network can be routed in G. For this we make use of a theorem of Schrijver [28]; we show that there exist $\Omega(\log n)$ paths in G between the neighboring enclosures, such that all paths between all pairs are mutually disjoint.

We make no attempt to optimize constants here. Set $\lambda_0 = \lambda$, and choose positive constants $\lambda_1, \lambda_2, \ldots$ so that $\lambda_{j+1} \ll \lambda_j$ (the exact relationship between these constants is easy to determine from the analysis below). A connection (s_i, t_i) is *short* if $d(s_i, t_i) \leq \lambda_2 \log n$ and *long* otherwise. As before, we handle long and short connections separately; for now we concentrate on pre-processing the graph as described above for handling long connections.

5.1 Building the Simulated Network

5.1.1 Clusters and Enclosures

We wish to choose a maximal set of mutually distant vertices around which to grow clusters. Let G^r denote the graph obtained from G by joining u and v if $d(u, v) \leq r$. We first run

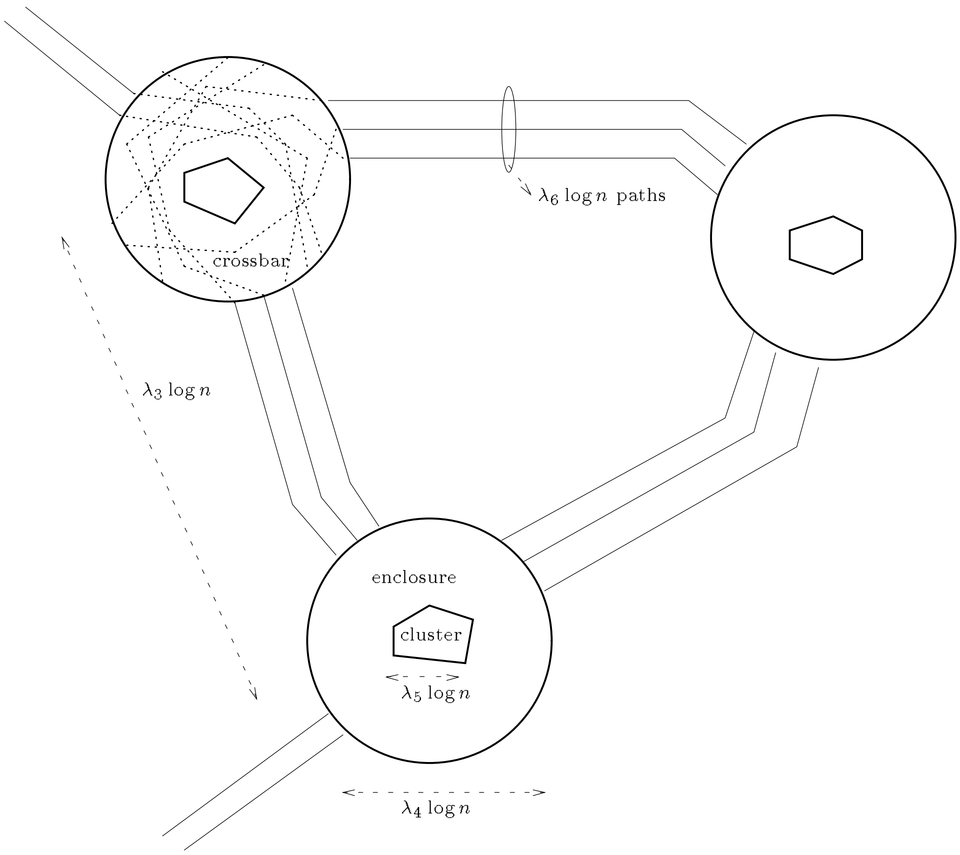

Figure 2: Building the simulated network

Luby's randomized maximal independent set algorithm [17] in $G^{\lambda_3 \log n}$.

Let M denote the resulting MIS. For any $x \in V$, some vertex within $\lambda_5 \log n$ of it will enter M on the first iteration if the largest number chosen in $B_{2\lambda_3 \log n}(x)$ is chosen by a vertex in $B_{\lambda_5 \log n}(x)$. This happens with constant probability, by Lemma 4.4. Moreover, if $d(x, y) \geq \lambda_2 \log n$, then these events are independent for x and y. Thus,

Lemma 5.1 *Let $x, y \in V$ be such that $d(x, y) \geq \lambda_2 \log n$. Then with constant probability there are $u, v \in M$ such that $d(x, u) \leq \lambda_5 \log n$ and $d(y, v) \leq \lambda_5 \log n$.*

Around each $v \in M$ we now grow a *cluster* of radius roughly $\lambda_5 \log n$, and an *enclosure* around each cluster, with "smooth" boundaries. We need the following facts. Let H denote an arbitrary graph, and Q a simple cycle of H. For $u, v \in Q$, let $d_Q(u, v)$ denote the shortest-path distance from u to v on Q — that is, the length of the shorter of the two u-v paths on Q.

Definition 5.2 *We say that Q is* ε-smooth *if for all $u, v \in Q$ we have $\varepsilon d_Q(u, v) \leq d(u, v)$.*

Definition 5.3 *If U and W are two subsets of $V(H)$, we say that U is ε'-close to W if for each $u \in U$ there is a $w \in W$ such that $d(u, w) \leq \varepsilon' |W|$.*

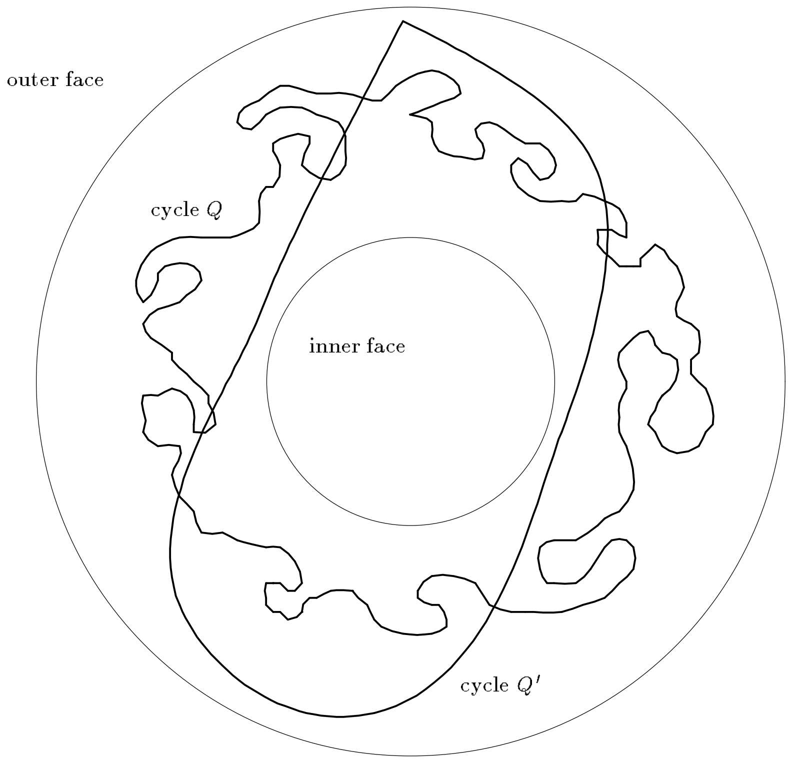

Figure 3: Smoothing a cycle

The following fact is quite similar to, but more general than, Theorem 4.4 of our earlier paper [14]; the proof is very similar as well. In effect, it says that given a cycle Q in a planar graph H that encloses (in the sense of homotopy) the "hole" formed by some internal face, then for a small $\varepsilon > 0$ we can find a cycle Q' of no greater length that is ε-close to Q, $\Omega(\varepsilon)$-smooth, and also encloses this hole. See Figure 3.

We will use this theorem to smooth out the boundaries of the clusters and the enclosures around them.

Theorem 5.4 *For each $\varepsilon > 0$ the following holds. Let Σ_1 be a compact surface (possibly with boundary), H a graph embedded on Σ_1, and Q a simple cycle of H that is non-null-homotopic on Σ_1. Then in polynomial time one can find an $\frac{\varepsilon}{1+\varepsilon}$-smooth simple cycle Q' such that*
(i) $|Q'| \le |Q|$,
(ii) Q' is ε-close to Q, and
(iii) Q' is also non-null-homotopic on Σ_1.

Proof. For $u, v \in Q$, let $[u, v]_Q$ denote the shorter of the two u-v paths contained in Q (ties broken arbitrarily), and let $\bar{\varepsilon} = \frac{\varepsilon}{1+\varepsilon}$. If Q is not $\bar{\varepsilon}$-smooth, then there are $u, v \in Q$ such that

$$\bar{\varepsilon} d_Q(u, v) > d(u, v). \tag{3}$$

Moreover, we can efficiently find such a u and v so that there is a shortest u-v path P_{uv} in H that is vertex-disjoint from Q (for example, the pair u, v satisfying (3) for which $|P_{uv}|$ is minimum).

Now one of the two simple cycles $[u,v]_Q \cup P_{uv}$ and $(Q \setminus [u,v]_Q) \cup P_{uv}$ is not null-homotopic on Σ_1; and each is shorter than Q. We thus update Q, replacing it with the cycle from among these two that is not null-homotopic.

We now iterate this process of "slicing off" parts of Q using short paths through H. Since the length of the cycle decreases with each iteration, this process must terminate in a cycle Q' that is $\bar{\varepsilon}$-smooth. Moreover, each iteration maintains the invariant that the current cycle is non-null-homotopic on Σ_1. Thus, we only have to verify that the final cycle is ε-close to Q.

This is clearly true after the first iteration: since $|P_{uv}| < \bar{\varepsilon} d_Q(u,v) < \varepsilon|Q|$, every vertex on the updated cycle can reach a vertex of Q by a path of length at most $\varepsilon|Q|$. Now, let Q_i denote the cycle obtained after i iterations of slicing off. As long as portions of Q remain on Q_i, we say that we are in the "first phase"; other phases will be defined below. In the first phase, Q_i consists of alternating intervals $Q_{i1}, P_{i1}, Q_{i2}, \ldots, Q_{ir}, P_{ir}$, where $Q_{ij} \subset Q$, and the interval Q'_{ij} of Q lying between Q_{ij} and $Q_{i,j+1}$ has been sliced off by P_{ij}. We can show by induction on the number of iterations that $|P_{ij}| \le \bar{\varepsilon}\left|Q'_{ij}\right|$ — as was true after the first iteration.

This is done by the following case analysis. In the $(i+1)^{\text{st}}$ iteration, we find a new path; there are three cases to consider.

1. One end of P lies on Q_{ij} and the other on Q_{ik}, where possibly $j = k$. Then the property clearly continues to hold, since $|P|$ is at most $\bar{\varepsilon}$ times the number of current cycle vertices cut off, which is in turn at most the number of original vertices of Q between the endpoints of P.

2. One end of P lies on P_{ij} and the other on Q_{ik} (so P_{ij} is lengthened). Suppose that the amount of original cycle cut off *in addition to* Q'_{ij} is equal to x, and the amount of P_{ij} that is cut off by P is y. Then if $P_{i+1,j}$ denotes P_{ij} after this iteration, we have

$$\begin{aligned}
|P_{ij}| &\le \bar{\varepsilon}\left|Q'_{ij}\right| \\
|P| &\le \bar{\varepsilon}(x+y) \\
|P_{i+1,j}| &= |P_{ij}| + |P| - y \\
&\le \bar{\varepsilon}(\left|Q'_{ij}\right| + x + y) - y \\
&\le \bar{\varepsilon}(\left|Q'_{ij}\right| + x)
\end{aligned}$$

3. One end of P lies on P_{ij} and the other lies on P_{ik} (so P glues some of the new paths together). There are two subcases.

   (i) $j = k$. Then $|P_{ij}|$ goes down while $|Q_{ij}|$ is not affected, so the property still holds.

   (ii) $j \neq k$. Again suppose that the amount of original cycle cut off in addition to Q'_{ij} and Q'_{ik} is equal to x, the amount of P_{ij} cut off by P is y, the amount of P_{ik} cut off by P is z, and the new interval is denoted $P_{i+1,j}$. Then

$$\begin{aligned}
|P_{ij}| &\leq \bar{\varepsilon}\left|Q'_{ij}\right| \\
|P_{ik}| &\leq \bar{\varepsilon}|Q'_{ik}| \\
|P| &\leq \bar{\varepsilon}(x+y+z) \\
|P_{i+1,j}| &= |P_{ij}|+|P|+|P_{ik}|-y-z \\
&\leq \bar{\varepsilon}(\left|Q'_{ij}\right|+x+y+z+|Q'_{ik}|)-y-z \\
&\leq \bar{\varepsilon}(\left|Q'_{ij}\right|+x+|Q'_{ik}|)
\end{aligned}$$

If the iterations come to an end before the end of the first phase, then indeed Q' is ε-close to Q — any vertex on P_{ij} can reach Q by a path of length at most $\bar{\varepsilon}\left|Q'_{ij}\right| \leq \bar{\varepsilon}|Q|$. Otherwise, consider the iteration in which the first phase comes to an end. By analogous arguments, we obtain a cycle Q^1 such that $|Q^1| \leq \bar{\varepsilon}|Q|$ and every vertex on Q^1 can reach Q by a path of length at most $\bar{\varepsilon}|Q|$.

Each phase now proceeds exactly like the previous one, except that it begins with a cycle whose length has been reduced by at least a factor of $\bar{\varepsilon}$. Thus when the process terminates, all vertices on Q' will be able to reach Q by a path of length at most

$$|Q| \cdot \sum_{i=1}^{\infty} \bar{\varepsilon}^i = \varepsilon|Q|.$$

Thus Q' is ε-close to Q. ∎

We use the following procedure to build the clusters and the enclosures around each node v in M. Let $K_v = B_{\lambda_5 \log n}(v)$.

(i) Choose a radius r between $2\lambda_5 \log n$ and $3\lambda_5 \log n$ so that $|\pi(B_r(v))| \leq 9\beta\lambda_5 \log n$. Set $C_v = B_r(v)$.

(ii) Now extend C_v to a simple set as in Lemma 4.6; since $\lambda_3 > 2c\lambda_5$, no C_v is grown enough that it overlaps any other by this process.

(iii) We now apply the ε-smoothing algorithm of Theorem 5.4 to the facial cycle Q_v of $G[C_v]$ containing $\pi(C_v)$. Here $H_v = G[B_{\lambda_3 \log n}(v)] \setminus K_v^o$ plays the role of H, and the cylinder formed by removing the portions of Σ on which $G[K_v^o]$ and $G \setminus B_{\lambda_3 \log n}(v)$ are drawn plays the role of Σ_1. Now for a constant ε the resulting cycle Q'_v is ε-smooth in this subgraph H_v of G, and it is also $\frac{\varepsilon}{1-\varepsilon}$-close to Q_v. If we choose $\varepsilon < \frac{1}{18\beta+1}$, then since Q_v is initially at least $\frac{1}{9\beta}|Q_v|$ away from K_v, we know that any path with both ends on Q'_v that passes through K_v must have length at least $\frac{1}{9\beta}|Q_v| \geq \frac{1}{9\beta}|Q'_v|$. Thus, there are no "short cuts" between vertices of Q'_v that make use of K_v; hence Q'_v is in fact ε-smooth in G.

The smooth cycle Q'_v encloses a set S of vertices containing K_v. Update C_v to be this set S.

We now grow an *enclosure* $D_v \supset C_v$ by the same three-step process, except that we now use the constant λ_4 in place of λ_5, and the set C_v^o in place of K_v^o Thus, we have clusters of radius $\approx \lambda_5 \log n$, enclosures of radius $\approx \lambda_4 \log n$, and they are separated by a mutual distance of $\approx \lambda_3 \log n$.

Following the algorithm of Section 3, we now must build crossbar structures in the enclosures to replace the natural crossbars of the mesh. We build crossbars using an extension of the Okamura–Seymour theorem [20] due to Frank [9], along the same lines as was done in [14]. To be precise,

Definition 5.5 *If $X \subset V$, we say a* crossbar *anchored in X is a set of edge-disjoint paths, each with at least one end in X, such that every pair of paths meets in at least one vertex.*

Let $\sigma_v = \pi(C_v)$, and recall that by Q'_v we mean the facial cycle of $G[C_v]$ containing σ_v. Analogously, let $\tau_v = \pi(D_v)$, and φ_v the facial cycle of $G[D_v]$ containing τ_v. We wish to build a crossbar in $G[D_v \setminus C_v^o]$, anchored in $\sigma_v \cup \tau_v$, of size at least a constant fraction of $|\sigma_v \cup \tau_v|$. For a large enough constant κ depending on ε, we choose a set σ'_v of $|\sigma_v|/\kappa$ vertices on σ_v spaced about κ apart, and a set τ'_v of $|\tau_v|/\kappa$ on τ_v spaced about κ apart.

Lemma 5.6 *There is a crossbar anchored in $\sigma'_v \cup \tau'_v$, such that each vertex of $\sigma'_v \cup \tau'_v$ is the endpoint of a distinct path of the crossbar.*

Proof. Consider the planar graph $G[D_v \setminus C_v^o]$. This graph has only two large faces — the outer face bounded by φ_v, and the inner face left by the deletion of C_v^o. We find a shortest path P^* in the planar dual graph whose endpoints are equal to these two large faces, and we delete the edges used by this path. Denote the resulting graph by G_v; note that it has only a single large face, bounded by a cycle φ'_v which contains $\sigma'_v \cup \tau'_v$.

We claim that the cycle φ'_v is ε'-smooth in the graph G_v, where

$$\varepsilon' = \min(\frac{1}{2}\varepsilon, (27\beta\ell)^{-1}).$$

To see this, suppose that P is a path in G_v with endpoints $u, w \in \varphi'_v$. If u and w both belong to Q'_v, or both belong to φ_v, then this follows from the ε-smoothness of these two cycles in $G[D_v \setminus C_v^o]$. If one belongs to each, then $|P| \geq (\lambda_4 - 4\lambda_5)\log n$, while $|\varphi'_v| \leq 9\beta\ell(2\lambda_4 + \lambda_5)$, and again the bound follows. Finally, if both u and w lie on the short dual path P^*, then in fact $|P| \geq d_{\varphi'_v}(u, w)$, while if exactly one lies on P^* then it is easily verified that $|P| \geq \frac{1}{2}\varepsilon d_{\varphi'_v}(u, w)$.

Write $X = \sigma'_v \cup \tau'_v$. Let us assume for simplicity that $|X|$ is odd. Now for $u \in X$, define u^+ to be the vertex on φ'_v that is $\frac{1}{2}\kappa$ steps clockwise from u, and write $X^+ = \{u^+ : u \in X\}$. Now $|X \cup X^+|$ is even, so we can pair each $u \in X \cup X^+$ with its unique "antipodal" point $\tilde{u} \in X \cup X^+$ under the clockwise ordering of φ'_v. Note that vertices in X are paired with vertices in X^+, and vice versa. We now define a disjoint paths problem in G_v, with the set of terminal pairs $\mathcal{T}_v$ equal to $\{(u, \tilde{u}) : u \in X\}$. Note that all terminals are at least $\frac{1}{2}\kappa$ from one another on the cycle φ'_v.

We now want to show that $\mathcal{T}_v$ is realizable in G_v. First, say that a cut is *non-trivial* if it separates at least one pair of terminals. We are dealing with a disjoint paths problem in a planar graph with all terminals on the outer face. Moreover, every node of G_v not on the outer face has even degree. In such a case, Frank's extension of the Okamura-Seymour theorem [9, 20] says that following *strict cut condition* is sufficient for the realizability of $\mathcal{T}_v$ — every non-trivial cut has more capacity than the number of terminal pairs it separates.

It is enough to consider non-trivial cuts of the form $\delta(U)$ with both $G_v[U]$ and $G_v \setminus U$ connected. For such a set U, there must be two vertices $v, w \in \pi(U)$ such that $v, w \in \varphi'_v$. Suppose that the distance from v to w in G_v is d; then by Lemma 4.7, we have $|\delta(U)| \geq \ell^{-1}d$. Since the facial cycle φ_v is ε'-smooth, the number of terminal pairs disconnected by $\delta(U)$ is at most $2\varepsilon'^{-1}d/\kappa$. Thus taking $\kappa > 2\ell\varepsilon'^{-1}$ ensures that the strict cut condition will be satisfied.

Finally, observe that the edge-disjoint paths in a realization of $\mathcal{T}_v$ provide the crossbar required by the lemma, since each pair of paths must meet at some vertex of G_v. ∎

In the crossbar just constructed, let Y_v^u denote the path with an endpoint equal to u, where $u \in \sigma'_v \cup \tau'_v$.

5.1.2 The simulated network $\mathcal{N}$

We now construct the *simulated network*; the nodes of this network are the clusters, which we represent by the vertices in M. We define a neighbor relation on the clusters using the notion of a *Voronoi partition*; two clusters will be joined by an edge in $\mathcal{N}$ if they are neighbors in this sense.

Let H be a graph and $S \subset V(H)$. We fix a lexicographic ordering $\preceq$ on the elements of S. For $s \in S$, let

$$\mathcal{U}_s = \{v \in G : \forall s' \in S : d(v,s) \leq d(v,s') \text{ and } \forall s' \preceq s : d(v,s) < d(v,s')\}.$$

That is, $\mathcal{U}_s$ is the set of vertices that are at least as close to s as to any other s', with ties broken based on $\preceq$.

Definition 5.7 *The* Voronoi partition $\mathcal{V}(H,S)$ *of H with respect to S is the partition $\{\mathcal{U}_s : s \in S\}$.*

The following fact is immediate.

Lemma 5.8 *For each $s \in S$, $H[\mathcal{U}_s]$ is connected.*

Proof. Suppose $v \in \mathcal{U}_s$; we claim that any shortest s-v path P is contained in $\mathcal{U}_s$. For suppose not, and let $v' \in P$ be the closest vertex to s that lies in $\mathcal{U}_{s'}$ for some $s' \neq s$. Then $d(s',v) \leq d(s,v)$, and in fact $d(s',v) < d(s,v)$ if $s \preceq s'$. It follows that $v \in \mathcal{U}_{s'}$, a contradiction. ∎

We can now build a graph $\mathcal{N}(H,S)$ on the vertices in S, joining two if their Voronoi cells share an edge.

Definition 5.9 *The* neighborhood graph *of S in H, denoted $\mathcal{N}(H,S)$, is the graph with vertex set S, and an edge (s,s') iff there is an edge of H with endpoints in $\mathcal{U}_s$ and $\mathcal{U}_{s'}$.*

The simulated graph we use will be the neighborhood graph $\mathcal{N}(G,M)$ with every edge given capacity $\approx \lambda_6 \log n$. Let $\mathcal{V}$ and $\mathcal{N}$ denote $\mathcal{V}(G,M)$ and $\mathcal{N}(G,M)$ respectively, and $\mathcal{N}(\gamma)$ the graph $\mathcal{N}$ in which each edge is given capacity γ.

The following two facts about $\mathcal{N}$ are easy to establish. First, by the maximality of M, we have

Lemma 5.10 *For all $v \in M$, $\mathcal{U}_v \subset B_{\lambda_3 \log n}(v)$.*

Proof. Suppose $u \in \mathcal{U}_v$ but $d(v,u) > \lambda_3 \log n$. Then $d(v',u) > \lambda_3 \log n$ for *all* $v' \in M$; this contradicts the fact that M is a maximal independent set in $G^{\lambda_3 \log n}$. ∎

This, along with Lemma 4.4, implies

Lemma 5.11 *The degree of a vertex in $\mathcal{N}$ is at most $\Delta' \leq 16\bar{\alpha}^{-1}\beta$.*

Proof. Let U denote the neighbors of v in $\mathcal{N}$, including v itself. Then by Lemma 5.10,

$$\bigcup_{u \in U} \mathcal{U}_u \subset B_{2\lambda_3 \log n}(v),$$

and hence

$$\sum_{u \in U} |\mathcal{U}_u| \leq 4\beta\lambda_3^2 \log^2 n.$$

But around each $u \in U$ there is a disjoint ball of radius $\frac{1}{2}\lambda_3 \log n$, which contains at least $\bar{\alpha}\lambda_3^2 \log^2 n/4$ vertices. Thus $|U| \leq 16\bar{\alpha}^{-1}\beta$. ∎

5.1.3 Inter-Cluster Paths

The goal of this part is, for a constant λ_6, to construct $\lambda_6 \log n$ disjoint paths between each pair of enclosures D_v, D_w where (v,w) is an edge in $\mathcal{N}$. This will allow us to convert a routing in the simulated network $\mathcal{N}(\lambda_6 \log n)$ into actual disjoint paths in G. Recall that the outer facial cycle of $G[D_v]$ is denoted φ_v, and it contains a set τ'_v of vertices evenly spaced at distance κ.

Theorem 5.12 *There exist vertex-disjoint paths in G, each with ends in sets τ'_v and otherwise disjoint from all D_v, such that for $(v,w) \in \mathcal{N}$, there are at least $\lambda_6 \log n$ such paths with one end in τ'_v and the other in τ'_w.*

Proof. The proof is based on the following theorem of Schrijver [28]. Let Σ_1 be a surface (possibly with boundary), H a graph embedded on Σ_1, and $\{A^i : i = 1, \ldots, k\}$ a set of disjoint curves on Σ_1 each of which is either closed or anchored at vertices of H on the boundary of Σ_1. The problem is to find vertex-disjoint paths and cycles $\{P^i\}$ in G so that P^i is homotopic to A^i for each i.

Call a collection of curves $\overline{\mathcal{R}}$ on Σ_1 *essential* if it consists either of a single closed curve that is not null-homotopic, or it is a finite union of curves each with endpoints on the boundary of Σ_1. Schrijver [28] proves that such vertex-disjoint paths and cycles exist if for each essential collection of curves $\overline{\mathcal{R}}$, there are curves $\{B^i\}$, where B^i is homotopic to A^i, such that $\overline{\mathcal{R}}$ intersects the drawing of H more than it intersects the curves $\{B^i\}$. (The main result of [28] is in fact a necessary and sufficient condition; this weaker statement suffices for our purposes. Additionally, [28] is stated for the special case of surfaces without boundary, but the extension we use here follows immediately from [28].)

Say that a curve is G-normal if it meets the drawing of G only at vertices, and define its G-length to be the number of times it meets the drawing. For each v, v' that are neighbors

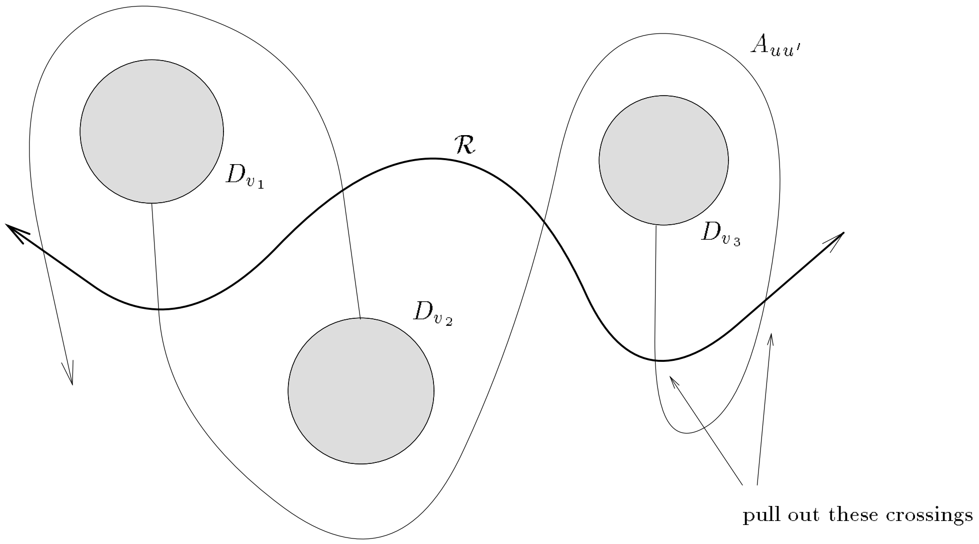

Figure 4: Pulling out a crossing

in $\mathcal{N}$, we draw a G-normal arc $\mathcal{A}_{vv'}$ on Σ with endpoints v and v'. We can ensure that all these arcs are disjoint, since each $\mathcal{U}_v$ is connected, and for each $(v, v') \in E(\mathcal{N})$, there is at least one edge of G with endpoints in $\mathcal{U}_v$ and $\mathcal{U}_{v'}$. Choose a small constant $\lambda_6 \leq |\tau'_v|/\Delta'$ (say, less than $\frac{1}{8}\alpha^2\bar{\alpha}\Delta'^{-2}\Delta^{-2}\lambda_3^{-1}\lambda_4^2$; the reason for this will become clear below). Now suppose we have in fact $\lambda_6 \log n$ copies of each $\mathcal{A}_{vv'}$, all running "parallel" to one another. By pushing them apart appropriately, we can assume that each arc runs through a different vertex in τ'_v and $\tau'_{v'}$. Let $A_{uu'}$ denote the G-normal arc that runs through $u \in \tau'_v$ and $u' \in \tau'_{v'}$.

Define G' to be the graph obtained by deleting D_v^o for each $v \in M$. We now cut Σ along the facial cycles φ_v to obtain Σ', a surface with boundary. Note that G' is properly embedded on Σ'. The theorem is now a consequence of the following claim. ■

Claim 5.13 *There exist vertex-disjoint paths $P_{uu'}$ in G' such that $P_{uu'}$ is homotopic to $A_{uu'}$.*

Proof. If curves $\mathcal{R}$ and $\mathcal{R}'$ are homotopic, we write $\mathcal{R} \sim \mathcal{R}'$. We extend this notation to finite collections of curves $\overline{\mathcal{R}}$, $\overline{\mathcal{R}}'$ in the obvious way. Following notation of Schrijver [28], the number of crossings of $\mathcal{R}_0$ and $\mathcal{R}_1$ is denoted $\mathrm{cr}(\mathcal{R}_0, \mathcal{R}_1)$, and we define

$$\mathrm{mincr}(\mathcal{R}_0, \mathcal{R}_1) = \min\{\mathrm{cr}(\mathcal{R}'_0, \mathcal{R}'_1) : \mathcal{R}_0 \sim \mathcal{R}'_0, \mathcal{R}_1 \sim \mathcal{R}'_1\}.$$

By the theorem of Schrijver [28] given above, the desired paths exist if for each essential collection of curves $\overline{\mathcal{R}}$ on Σ', one has

$$\mathrm{cr}(\overline{\mathcal{R}}, G') > \sum_{(u,u')} \mathrm{mincr}(\overline{\mathcal{R}}, A_{uu'}). \tag{4}$$

Note that in verifying this inequality, we may assume $\overline{\mathcal{R}}$ is G'-normal, and that $\overline{\mathcal{R}}$ has no self-crossings.

For a collection of curves $\overline{\mathcal{R}}$, define its *index* to be

$$\operatorname{cr}(\overline{\mathcal{R}}, G') - \sum_{(u,u')} \operatorname{mincr}(\overline{\mathcal{R}}, A_{uu'}).$$

So it is enough to consider collections of curves $\overline{\mathcal{R}}$ whose index is minimum in their homotopy class, and to show that such $\overline{\mathcal{R}}$ in fact have positive index.

Set $r = \frac{1}{2}\Delta^{-1}\alpha\bar{\alpha}^{1/2}\lambda_4 \log n$. We claim that no G-normal closed curve of G-length less than $2r$ can enclose a set of the form D_v. For if it did, then D_v could be disconnected by an edge cut of size less than $\alpha\bar{\alpha}^{1/2}\lambda_4 \log n$, which is not possible since $|D_v| \geq \bar{\alpha}\lambda_4^2 \log^2 n$. From this it follows that any G'-normal closed curve of G'-length less than $2r$ must be null-homotopic in Σ'.

Note that we can view the expression $\operatorname{cr}(\overline{\mathcal{R}}, G') - \sum_{(u,u')} \operatorname{cr}(\overline{\mathcal{R}}, A_{uu'})$ as a sum over the finitely many arc-components of $\overline{\mathcal{R}}$; if the value of this expression is not positive, we show how modify the curves $A_{uu'}$ so that it increases. We do this by considering each arc-component of $\overline{\mathcal{R}}$ in turn. Let $\mathcal{R}$ denote a single arc-component of $\overline{\mathcal{R}}$; we consider two cases, based on the G'-length of $\mathcal{R}$.

Case 1. $\operatorname{cr}(\mathcal{R}, G') \leq r$. Then $\mathcal{R}$ must have both endpoints on the same facial cycle (it is too short to touch two such cycles, and if it were a closed curve it would have to be null-homotopic, by the above argument.) But then it is easy to produce arcs $\{A'_{uu}\}$ for which $\operatorname{cr}(\mathcal{R}, G') > \sum_{(u,u')} \operatorname{cr}(\mathcal{R}, A'_{uu'})$ since φ_v is ε-smooth.

Case 2. $\operatorname{cr}(\mathcal{R}, G') > r$. Again, we just have to exhibit arcs $A'_{uu'} \sim A_{uu'}$ lying on Σ' so that

$$\operatorname{cr}(\mathcal{R}, G') > \sum_{(u,u')} \operatorname{cr}(\mathcal{R}, A'_{uu'}), \tag{5}$$

without increasing the number of crossings of these arcs with the other components of $\overline{\mathcal{R}}$. If the set $\{A_{uu'}\}$ satisfies (5), we are done; otherwise, we show how to modify this set of arcs so that it does. See Figure 4.

If the set $\{A_{uu'}\}$ does not satisfy Inequality (5), then there is some interval $\mathcal{R}'$ of $\mathcal{R}$ of G'-length r for which (5) is violated. Let us consider such an $\mathcal{R}'$.

Observe that each arc $A_{uu'}$ has G'-length at most $2\lambda_3 \log n$, and hence at most $\Delta'^2\lambda_6 \log n$ of these arcs can meet $\mathcal{R}'$, since at most Δ'^2 pairs of clusters have at least one end close enough to $\mathcal{R}'$. Now suppose the total number of crossings of these arcs with $\mathcal{R}'$ exceeds

$$\frac{(2\lambda_3 \log n)(\Delta'^2\lambda_6 \log n)}{r} < \operatorname{cr}(\mathcal{R}', G).$$

Then some arc $A_{uu'}$ meets $\mathcal{R}'$ more than $2\lambda_3 \log n/r$ times, and hence the interval of $A_{uu'}$ between some pair of consecutive crossings with $\mathcal{R}'$ has G'-length less than r.

Suppose that this pair of consecutive crossings occurs at vertices w and w'. Let $\mathcal{R}''$ denote the G'-normal curve formed from this interval of $A_{uu'}$ and the portion of $\mathcal{R}'$ between w and w'. $\mathcal{R}''$ has G'-length less than $2r$, and so it must be null-homotopic by the argument given above.

Now, since $\mathcal{R}$ has minimum index over all curves in its homotopy class, the portion of $A_{uu'}$ between w and w' meets G' at least as many times as the portion of $\mathcal{R}'$ between w and w'. We can therefore modify $A_{uu'}$ so that it runs along $\mathcal{R}'$ for this interval. This does not increase the G'-length of $A_{uu'}$; and it decreases the number of crossings of $\mathcal{R}$ — as well as $\overline{\mathcal{R}}$ (since it has no self-crossings) — with $A_{uu'}$.

Thus this process terminates; when it does, we have a set of arcs $\{A'_{uu'}\}$ for which Inequality 5 holds. ∎

Let us denote one such path with ends u and u' by $Z_{uu'}$. Moreover, we have $u \in \tau'_v$ and $u' \in \tau'_{v'}$, and they are ends of paths Y_v^u and $Y_{v'}^{u'}$ respectively. Denote by $\bar{Z}_{uu'}$ the concatenation of the three paths Y_v^u, $Z_{uu'}$, and $Y_{v'}^{u'}$.

5.2 The On-Line Algorithm

With the simulated network $\mathcal{N}$ in place, the routing algorithms themselves are essentially the same as those for the mesh; here we just describe what must be modified.

First of all, the analogue of Lemma 3.10 is the following.

Lemma 5.14 *Let $r \leq \lambda_1 \log n$, $U \subset B_r(v)$ for some $v \in V$, and $\mathcal{T}'$ a set of terminal pairs in U. Then the maximum size of a subset of $\mathcal{T}'$ that is realizable in $B_{8\xi^2 r}(v)$ is within a constant factor of the maximum size of a subset of $\mathcal{T}'$ that is realizable in G.*

Proof. First choose a radius r' between $2r$ and $3r$ for which $|\pi(B_r(v)| \leq 9\beta r'$. Then construct a simple set extension of $B_{r'}(v)$ as in Lemma 4.6; and ε-smooth its outer facial cycle to obtain a set $U' \supset U$ contained in $B_{4\xi r}(v)$. Let $U'' \subset B_{8\xi^2 r}$ denote a simple set extension of $B_{8\xi r}(v)$, as in Lemma 4.6. For a constant κ', we can pick a set S of vertices on the outer facial cycle of U' spaced κ' apart, and use Frank's theorem [9] as in Lemma 5.6 to construct a set of edge-disjoint paths connecting "antipodal" pairs in S, such that all paths stay within $U'' \setminus U'$. Note that we must take care to ensure that the parity condition is met, since the outer facial cycle of U'' can contain odd-degree vertices. To do this we remove sub-paths of this cycle between consecutive pairs of the (necessarily even number of) odd-degree vertices; U'' is large enough that the strict cut condition will remain satisfied.

Consider the set of paths in a realization of a maximum-size subset of $\mathcal{T}'$ in G. Of the paths that meet $\pi(U')$, we can select a constant fraction of paths such that the the pairs of first and last intersections with $\pi(U')$ can be connected along the outer facial cycle of U' to different vertices in S. We can then use the crossbar of the previous paragraph to connect all of these pairs together; the resulting paths are within a constant fraction of the maximum number achievable in G, and they do not leave U''. ∎

The algorithm for short connections is now as follows. We run a randomized version of Luby's algorithm, this time in $G^{\lambda_1 \log n}$. Let M' denote the resulting MIS. With constant probability, both ends of a short connection are within $\frac{1}{16\xi^2}\lambda_1 \log n$ of the same $v \in M'$, as in Lemmas 3.11 and 5.1. We now let U_v denote $B_{\lambda_1 \log n / 16\xi^2}(v)$ and only route connections both of whose ends lie in the same U_v. To route such connections, we run the algorithm of Proposition 3.9 in each $B_{\frac{1}{2}\lambda_1 \log n}(v)$; by Lemma 5.14, this is within $O(\log n)$ of optimal in each U_v.

For long connections, we run the same simulation as in the case of the mesh, this time in the graph $\mathcal{N}(\lambda_6 \log n)$. The analogue of Lemma 3.6 holds exactly as before, as does Lemma 3.7 — making use of the paths $\bar{Z}_{uu'}$, since all such paths incident to the same enclosure must cross. Finally, the analogue of Lemma 3.8 is the following.

Lemma 5.15 *There is a constant γ such that for any realizable subset $\mathcal{T}'$ of $\mathcal{T}$, $\psi(\mathcal{T}')$ can be routed in $\mathcal{N}(\gamma \log n)$.*

Proof. Set $\gamma = 9\beta(\Delta'^3\lambda_5 + \lambda_3)$. For each s_i-t_i path P in the optimal routing, construct the following path for (u_i, v_i) in $\mathcal{N}$ — when P crosses from $\mathcal{U}_w$ into $\mathcal{U}_{w'}$, add an edge from w to w'. Now consider how many paths in our constructed routing use the edge (w, w'). Each such corresponds to a path in the original routing that used an edge in $\delta(\mathcal{U}_w, \mathcal{U}_{w'})$. We can't bound the size of this set directly, since it could be quite "meandering." But consider the following argument. $\delta(\mathcal{U}_w, \mathcal{U}_{w'}) \subset B_{2\lambda_3 \log n}(x)$ for some vertex $x \in G$; thus there is some r between $2\lambda_3 \log n$ and $3\lambda_3 \log n$ so that $|\pi(B_r(x))| \leq 9\beta\lambda_3 \log n$. So at most this many paths with both ends more than r away from x can use edge in $\delta(\mathcal{U}_w, \mathcal{U}_{w'})$. But closer than this, there are at most Δ'^3 clusters, each of which is the origin of at most $9\beta\lambda_5 \log n$ paths. ■

Thus the optimum in G is bounded by the fractional optimum in $\mathcal{N}(\gamma \log n)$, which is at most a constant factor more than the fractional optimum in $\mathcal{N}(\lambda_6 \log n)$, which is at most an $O(\log n)$ factor more than the number of pairs routed by the on-line algorithm in G. Thus we have

Theorem 5.16 *The on-line algorithm is $O(\log n)$-competitive in any uniformly densely embedded Eulerian graph G.*

5.3 The Off-Line Algorithm

The off-line algorithm too is essentially the same, now working with the larger simulated network $\mathcal{N}'(\lambda_6 \log n)$. The only change required is in the proof of Lemma 3.15. Here, we are no longer able to talk about "overfull rectangles"; however, we can define a *round cut* to be a set of the form $B_r(w) \cap C_v$. Since a given $u \in C_v$ is only contained in $O(r^2)$ round cuts of radius r, the following lemma establishes that round cuts can be used instead of rectangles in the analogue of Lemma 3.15 for densely embedded graphs.

Lemma 5.17 *Let G'_v denote $G[C_v]$ with an additional vertex z_v joined by an edge to each vertex in σ'_v. Then there is a constant ξ_1 such that for every $U \subset G'_v$ not containing z_v, there is a round cut $R \supseteq U$ satisfying $|\delta(R)| \leq \xi_1 |\delta(U)|$.*

Proof. Set $\xi'_1 = \kappa_1 + \bar{\alpha}^{-1/2}\alpha^{-1}$ and $\xi_1 = 4\Delta\xi'_1(\beta + \varepsilon^{-1})$. Let $U \subset G'_v$ be a set not containing z_v, and write $p = |\delta(U)|$; we construct a round cut R containing U for which $|\delta(R)| \leq \xi_1 p$.

Now if we contract $G'_v \setminus U$ to a single vertex, we obtain a planar graph with maximum face size κ_1 (as opposed to ℓ; this is due to the large spacing of the vertices of σ'_v). So by Lemma 4.7, the maximum distance between two points on $\pi(U)$ is at most $\kappa_1 p$.

Next we claim that every vertex in U must be within distance $\bar{\alpha}^{-1/2}\alpha^{-1}p$ of $\pi(U)$. The reason for this is analogous to the proof of Lemma 4.6 — if not, then U would contain a ball of more than this radius, which would contain more than $\alpha^{-1}p^2$ vertices; a contradiction since $|\delta(U)| = p$.

Thus for any $u \in \pi(U)$, we have $U \subset B_r(u)$, where

$$r = (\kappa_1 + \bar{\alpha}^{-1/2}\alpha^{-1})p = \xi_1' p.$$

Now by Lemma 4.5, there is an r' between r and $2r$ such that

$$|\delta(B_{r'}(u))| \leq 4\beta\Delta\xi_1' p.$$

Now let $R \supset U$ denote the round cut $B_{r'}(u) \cap C_v$. Every edge of $\delta(R)$ is an edge of either $\delta(B_{r'}(u))$ or of $\delta(C_v \cap R)$. The former quantity was just shown to be at most $4\beta\Delta\xi_1' p$. To bound the latter quantity, note that any two vertices in $\pi(C_v \cap R)$ are at most $2r' \leq 4\xi_1' p$ apart; since the facial cycle containing $\pi(C_v)$ is ε-smooth, this means that $\pi(C_v \cap R)$ contains at most $4\varepsilon^{-1}\xi_1' p$ vertices, and hence

$$|\delta(C_v \cap R)| \leq 4\Delta\varepsilon^{-1}\xi_1' p.$$

The claim now follows since

$$|\delta(R)| \leq 4\Delta\xi_1'(\beta + \varepsilon^{-1})p = \xi_1 p.$$

■

Thus we have

Theorem 5.18 *There is a randomized (off-line)* MDPP *algorithm in uniformly densely embedded Eulerian graphs that produces a constant-factor approximation with high probability.*

6 Extensions

6.1 Durations and Profits

In the on-line algorithm we assume that (i) all connections have infinite duration, and (ii) all connections have the same "value" (i.e. our objective function could have been a weighted sum of the set of pairs we accept, rather than an unweighted sum). However, there are general transformation techniques due to Awerbuch et. al. [5] that allow us to convert our results to on-line algorithms that can handle connections of limited duration and variable value, at the cost of additional logarithmic terms in the competitive ratio. Specifically, we pay $O(\log T)$ and $O(\log P)$, where T and P are the ratios between the largest and smallest durations and values respectively.

6.2 Graphs with an Exceptional Face

In this section, we sketch the extension of our algorithms to densely embedded, nearly-Eulerian graphs. Recall from Definition 4.3 that such a graph satisfies the properties of a uniformly densely embedded Eulerian graph, except that it is allowed to contain an "exceptional" face Φ^*, with facial cycle Q^* that may have length greater than ℓ and may contain vertices not of even degree.

For the remainder of this section, let G denote a densely embedded, nearly-Eulerian graph with parameters α, λ, Δ, and ℓ. For simplicity, we assume that the facial cycle Q^* is sufficiently large that it is not contained in any set $B_{\lambda \log n}(v)$; it is not difficult to remove this assumption.

The changes required in the algorithm come from the fact that there can now be a G-normal curve joining two distant vertices in G that intersects G relatively few times — this is because it can pass through the large face Φ^*. This has consequences in the proofs of Lemma 5.6 (and its relatives) and Claim 5.13. However, by requiring the outer cycles of clusters and enclosures to satisfy a more restrictive notion of ε-smoothness, these facts will follow as before.

We define our more restrictive type of ε-smoothness as follows. Let G/Q^* denote the graph G with a single additional node q^* joined by length-0 edges to each vertex of the long facial cycle Q^*. Then a small cut passing through two distant vertices, as described in the previous paragraph, *does* correspond to a short path in G/Q^* — it simply makes use of the additional node q^*. Now it is straightforward to show that we need only require the outer cycles of the clusters and enclosures to be ε-smooth in the graph G/Q^*; and this can be accomplished by running the ε-smoothing algorithm in this graph instead of in G.

This introduces a further difficulty, however. Say that we have just smoothed some cycle Q, obtaining a cycle Q'. While Q' will be ε-close to the original cycle Q in G/Q^*, there is no reason why this means it will be ε-close in G.

To handle this, we strengthen the statement of Theorem 5.4, as follows. For a vertex u, define the *u-restricted distance* $d^u(v, w)$ between two vertices v and w to be the minimum length of a v-w path avoiding u. From the proof of Theorem 5.4, one sees that we can always find a short path from the smooth cycle Q' back to the original cycle Q that avoids any prescribed vertex; i.e. Q' is ε-close to Q with respect to any u-restricted distance function. In particular, Q' is ε-close to Q in G/Q^* with respect to the q^*-restricted distance function; that is, Q' is ε-close to Q in the original graph G.

Thus, we can obtain clusters and enclosures with boundaries which are ε-smooth in G/Q^*, and which are not far from the original boundaries in G. The sets τ'_v will now consist of evenly spaced vertices only on the part of an enclosure's outer facial cycle that does not lie on Φ^*. The proof of Lemma 5.6 now follows exactly as before. When we use Schrijver's theorem to construct inter-cluster paths, we also cut the surface Σ along the long facial cycle Q^* so as to remove Φ^* from the surface. Now an essential curve can be anchored on the boundary of Φ^* as well as on the boundary of an enclosure; but this poses no problem since the enclosure boundaries are ε-smooth in G/Q^*.

Once the simulated network $\mathcal{N}$ has been set up, the on-line and off-line algorithms work exactly as before. (In particular, Lemma 5.17 follows without modification, since the large face Φ^* is incorporated into the hypotheses of Lemma 4.7.) Thus we have,

Theorem 6.1 *There is an $O(\log n)$-competitive on-line* MDPP *approximation in any densely embedded nearly-Eulerian graph.*

Theorem 6.2 *There is a randomized (off-line)* MDPP *algorithm in any densely embedded nearly-Eulerian graph that produces a constant-factor approximation with high probability.*

Acknowledgements

We thank Yuval Rabani for his contribution to this work. In particular, he suggested the general approach of performing randomized rounding on graphs with low-capacity edges and throwing out parts of the routing on which the capacity constraints become violated. Thanks also to Baruch Awerbuch, Michel Goemans, David Karger, Tom Leighton, and Lex Schrijver for their comments on these problems.

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, M. Sudan, "Efficient routing and scheduling algorithms for optical networks," *Proc. 5th ACM-SIAM SODA*, 1994, pp. 412–423.

[2] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin and O. Waarts, "On-line Load Balancing with Applications to Machine Scheduling and Virtual Circuit Routing" *Proc 25th ACM STOC*, 1993.

[3] Y. Aumann, Y. Rabani, "Improved bounds for all-optical routing," *Proc. 6th ACM-SIAM SODA*, 1995, pp. 567–576.

[4] B. Awerbuch, Y. Azar, S. Plotkin, "Throughput-Competitive Online Routing," *Proc. 34th IEEE FOCS*, 1993, pp. 32–40.

[5] B. Awerbuch, Y. Bartal, A. Fiat, A. Rosén, "Competitive non-preemptive call control," *Proc. 5th ACM-SIAM SODA.*

[6] B. Awerbuch, R. Gawlick, F.T. Leighton, Y. Rabani, "On-line admission control and circuit routing for high performance computing and communication," *Proc. 35th IEEE FOCS*, 1994, pp. 412–423.

[7] A. Broder, A. Frieze, S. Suen, E. Upfal, "Optimal construction of edge-disjoint paths in random graphs," *Proc. 5th ACM-SIAM SODA*, 1994, pp. 603–612.

[8] A. Broder, A. Frieze, E. Upfal, "Existence and construction of edge-disjoint paths on expander graphs," *SIAM J. Computing*, 23(1994), pp. 976–989.

[9] A. Frank, "Edge-disjoint paths in planar graphs," *J. Comb. Theory Ser. B*, 39(1985), pp. 164–178.

[10] A. Frank, "Packing paths, cuts, and circuits — a survey," in *Paths, Flows, and VLSI-Layout*, B. Korte, L. Lovász, H.J. Prömel, A. Schrijver, Eds., Berlin: Springer-Verlag, 1990. pp. 49–100.

[11] N. Garg, V. Vazirani, M. Yannakakis, "Primal-dual approximation algorithms for integral flow and multicut in trees, with applications to matching and set cover," *Proc. ICALP*, 1993, pp. 64–75.

[12] R.M. Karp, "Reducibility among combinatorial problems," *Complexity of Computer Computations*, R.E. Miller, J.W. Thatcher, Eds., New York: Plenum Press, 1972. pp. 85–103.

[13] J. Kleinberg, S. Rao, É. Tardos, "Routing with low congestion in densely embedded graphs," in preparation.

[14] J. Kleinberg, É. Tardos, "Approximations for the disjoint paths problem in high-diameter planar networks," *Proc. 27th ACM STOC*, 1995.

[15] M.E. Kramer, J. van Leeuwen, "The complexity of wire routing and finding the minimum area layouts for arbitrary VLSI circuits," *Advances in Computing Research 2: VLSI Theory*, F.P. Preparata, Ed., London: JAI Press, 1984. pp. 129–146.

[16] R. Lipton, R. Tarjan, "A separator theorem for planar graphs," *SIAM J. Applied Math.* 36(1979), pp. 177–189.

[17] M. Luby, "A simple parallel algorithm for the maximal independent set problem," *SIAM J. Computing*, 15(1986), pp. 1036–1053.

[18] W. Massey, *Algebraic Topology: An Introduction*, New York: Springer-Verlag, 1967.

[19] M. Middendorf, F. Pfeiffer, "On the complexity of the disjoint paths problem," *Combinatorica*, 13(1993). pp. 97–107.

[20] H. Okamura, P. Seymour, "Multicommodity flows in planar graphs," *J. Comb. Theory Ser. B*, 31(1981), pp. 75–81.

[21] D. Peleg, E. Upfal, "Disjoint paths on expander graphs," *Combinatorica*, 9(1989), pp. 289–313.

[22] P. Raghavan, "Probabilistic construction of deterministic algorithms: approximating packing integer programs," *J. Computer and System Sciences*, 37(1988), pp. 130–143.

[23] P. Raghavan, C.D. Thompson, "Randomized rounding: a technique for provably good algorithms and algorithmic proofs," *Combinatorica*, 7(1987), pp. 365–374.

[24] P. Raghavan, E. Upfal, "Efficient all-optical routing," *Proc. 26th ACM STOC*, 1994, pp. 134–143.

[25] N. Robertson, P.D. Seymour, "Graph Minors VII. Disjoint paths on a surface," *J. Combinatorial Theory Ser. B*, 45(1988), pp. 212–254.

[26] N. Robertson, P.D. Seymour, "Graph Minors XI. Circuits on a surface," *J. Combinatorial Theory Ser. B*, 60(1994), pp. 72–106.

[27] N. Robertson, P.D. Seymour, "Graph Minors XIII. The disjoint paths problem," *J. Combinatorial Theory Ser. B*, 63(1995), pp. 65–110.

[28] A. Schrijver, "Disjoint circuits of prescribed homotopies in a graph on a compact surface," *J. Combinatorial Theory Ser. B*, 51(1991), pp. 127–159.

[29] D. Sleator, R. Tarjan, "Amortized efficiency of list update and paging rules," *Comm. ACM*, 23(1985), pp. 202–208.

[30] D. Wagner, K. Weihe, "A linear time algorithm for multicommodity flow in planar graphs," *Proc. First European Symposium on Algorithms*, 1993, pp. 384–395.