

Dynamic Compression: Multicuts for Planar Graphs with Outer Terminals

Subhadip Mitra

subhadip_mitra@hotmail.com

Abstract

This paper presents a unified framework for dynamic graph compression with applications to the multicut problem in planar graphs where all terminals lie on the outer face. We develop an on-the-fly algorithm for computing and maintaining the maximum flow minimum cut gap in dynamically changing networks, achieving $O(\log n)$ update time for single-edge modifications. For planar graphs with k terminals on the outer boundary, we establish that the multicut problem admits an exact polynomial-time solution and provide an $O(n \log n)$ algorithm based on shortest-path separators. We extend these results to handle deep emulations of fault-free mesh architectures on meshes with random faults, proving an $O(\log n)$ slowdown bound for constant fault probability. Additionally, we develop compression techniques for adjacency list representations tailored to web graphs, mobile host networks, and social network structures, achieving 2.1 bits per edge for scale-free graphs. The algorithms are validated on real-world datasets including ClueWeb09, Twitter's follower graph, and mobile network traces from the CRAWDAD repository.

Keywords - *planar graphs; outer terminals; multicut; dynamic algorithms; graph compression; max-flow min-cut; mobile networks; social graphs*

I. INTRODUCTION

The interplay between graph structure and algorithmic efficiency has motivated extensive research into specialized algorithms for restricted graph classes. Planar graphs, in particular, admit substantially faster algorithms for many problems that are NP-hard or require polynomial but impractical running times on general graphs. The multicut problem, which asks for a minimum-weight set of edges whose removal disconnects specified terminal pairs, exemplifies this phenomenon: while NP-hard on general graphs, the problem becomes tractable when the input graph is planar and the terminals satisfy certain geometric constraints.

A particularly important special case arises when all terminals lie on the outer face of a planar embedding. This configuration occurs naturally in VLSI design, where pins on a chip boundary must be connected without crossings, and in geographic routing, where source and destination nodes often lie on the boundary of a region. The restriction to outer terminals transforms the multicut problem from NP-hard to polynomial-time solvable, yet the precise complexity and the design of efficient algorithms have remained incompletely understood.

Modern applications increasingly demand algorithms that operate on graphs undergoing continuous modification. Social networks experience edge insertions and deletions as friendships form and dissolve; web graphs evolve as pages are created, linked, and removed; mobile ad-hoc networks reconfigure as hosts move through physical space. These dynamics motivate the development of algorithms that efficiently update their outputs in response to graph modifications, rather than recomputing from scratch after each change.

This paper makes four principal contributions. First, we present a polynomial-time algorithm for the minimum multicut problem on planar graphs with all terminals on the outer face, establishing an $O(n \log n)$ running time through the use of shortest-path separators and a reduction to minimum-cost circulation. Second, we develop a dynamic

data structure that maintains the max-flow min-cut gap under edge insertions and deletions with $O(\log n)$ amortized update time. Third, we extend our techniques to the mesh emulation problem, showing how to embed a fault-free mesh into a faulty mesh with $O(\log n)$ slowdown when faults occur with constant probability. Fourth, we present compression algorithms for adjacency lists that exploit the structural properties of web graphs, mobile networks, and social graphs.

The remainder of this paper is organized as follows. Section II reviews related work on planar graph algorithms, dynamic graph data structures, and graph compression. Section III establishes formal definitions and notation. Section IV presents our algorithm for multicuts with outer terminals. Section V develops the dynamic max-flow min-cut data structure. Section VI addresses mesh emulation with faults. Section VII describes the compression techniques. Section VIII reports experimental results, and Section IX concludes.

II. RELATED WORK

A. Multicuts in Planar Graphs

The multicut problem on general graphs was shown to be NP-hard by Dahlhaus et al. [1], who also established APX-hardness, precluding a polynomial-time approximation scheme unless $P = NP$. The best known approximation ratio for general graphs is $O(\log k)$ due to Garg, Vazirani, and Yannakakis [2], achieved through region-growing techniques applied to the linear programming relaxation.

For planar graphs, the situation improves markedly. The work of Borradaile and Klein [3] on multiple-source shortest paths provides tools that enable efficient cut computation in planar graphs. When all terminals lie on a single face, the problem reduces to computing a minimum-cost set of non-crossing paths, as observed by Bentz [4]. The algorithms of Italiano et al. [5] for dynamic planar graphs provide a foundation for our dynamic data structures.

B. Dynamic Graph Algorithms

Dynamic graph algorithms maintain properties of a graph undergoing edge insertions and deletions. The seminal work of Sleator and Tarjan [6] on link-cut trees provides $O(\log n)$ amortized time for path queries in dynamic trees. Eppstein et al. [7] extended these ideas to maintain minimum spanning trees in general graphs. For planar graphs, the data structure of Eppstein [8] maintains connectivity in $O(\log n)$ time per operation.

Dynamic maintenance of flow-related quantities has proven more challenging. Goranci et al. [9] gave a dynamic algorithm for approximate max-flow in planar graphs with $O(n^{2/3})$ update time. Our contribution improves this bound to $O(\log n)$ for the specific case of maintaining the flow-cut gap rather than the flow value itself.

C. Graph Compression

The WebGraph framework of Boldi and Vigna [10] established that web graphs can be compressed to approximately 3 bits per edge by exploiting URL ordering and reference coding. Subsequent work by Chierichetti et al. [11] achieved further compression for social networks by exploiting community structure. Mobile network graphs present distinct challenges due to their temporal variability, addressed by Chaintreau et al. [12] through time-aggregated representations.

III. DEFINITIONS AND PRELIMINARIES

A. Graph-Theoretic Foundations

Definition 1 (Planar Graph). A graph $G = (V, E)$ is planar if it admits an embedding in the plane such that edges intersect only at their endpoints. A planar embedding Γ of G partitions the plane into faces, exactly one of which is unbounded (the outer face).

Definition 2 (Outer Terminal Configuration). Given a planar graph G with embedding Γ and terminal set $T = \{t_1, t_2, \dots, t_k\} \subseteq V$, we say T has an outer terminal configuration if all terminals lie on the boundary of the outer face. The terminals inherit a cyclic ordering $(t_{\pi(1)}, t_{\pi(2)}, \dots, t_{\pi(k)})$ from the embedding.

Definition 3 (Terminal Pairs and Multicut). A demand set $D = \{(s_1, t_1), \dots, (s_r, t_r)\}$ specifies pairs of terminals to be separated. A multicut for D is a set $M \subseteq E$ such that for each $(s_i, t_i) \in D$, removing M disconnects s_i from t_i . For edge weights $w: E \rightarrow \mathbb{R}^+$, the minimum multicut problem seeks M minimizing $\sum_{e \in M} w(e)$.

Definition 4 (Non-Crossing Partition). For terminals in outer configuration with cyclic order (t_1, \dots, t_k) , a partition P of $\{1, \dots, k\}$ is non-crossing if for any two blocks $B, B' \in P$, there do not exist $i < j < k < l$ with $i, k \in B$ and $j, l \in B'$. The set of non-crossing partitions is denoted $NC(k)$.

Definition 5 (Dual Graph). For a planar graph G with embedding Γ , the dual graph G^* has a vertex for each face of Γ and an edge e^* connecting faces f_1, f_2 whenever edge $e \in E(G)$ lies on the boundary of both f_1 and f_2 . Edge weights transfer via $w^*(e^*) = w(e)$.

B. Flow-Cut Duality

Definition 6 (Maximum Concurrent Flow). Given graph G with edge capacities $c: E \rightarrow \mathbb{R}^+$ and demands D , the maximum concurrent flow λ^* is the supremum of λ such

that λ units can be simultaneously routed for each demand pair without exceeding capacities.

Definition 7 (Flow-Cut Gap). Let C^* denote the minimum multicut weight and λ^* the maximum concurrent flow. The flow-cut gap is $\rho = C^*/\lambda^*$. By weak LP duality, $\rho \geq 1$.

Theorem 1 (Okamura-Seymour [13]). For a planar graph G with all terminals on the outer face and demands D such that each terminal appears in exactly one demand pair, the flow-cut gap $\rho = 1$. That is, max-flow equals min-cut.

C. Dynamic Graph Structures

Definition 8 (Dynamic Graph Problem). A dynamic graph problem maintains a property $P(G)$ of graph G under a sequence of operations: $\text{Insert}(e)$ adds edge e , $\text{Delete}(e)$ removes edge e , and $\text{Query}()$ returns $P(G)$. The update time is the worst-case or amortized time for Insert/Delete ; the query time is the time for Query .

Definition 9 (Compressed Adjacency Representation). For graph $G = (V, E)$ with $|V| = n$, a compressed adjacency representation consists of: (1) a node ordering $\pi: V \rightarrow \{1, \dots, n\}$, (2) for each v , the compressed adjacency list $C(v)$ encoding $N(v)$ using reference nodes and gap sequences. The compression ratio is $|C|/(m \log n)$ where $|C|$ is total bits.

IV. MULTICUTS WITH OUTER TERMINALS

A. Structural Properties

Lemma 1. Let G be a planar graph with outer terminal configuration $T = \{t_1, \dots, t_k\}$. Any minimum multicut M for demands D induces a non-crossing partition of T , where terminals in the same block remain connected in $G - M$.

Proof. Suppose for contradiction that M induces a crossing partition: there exist $i < j < k < l$ with t_i, t_k connected in $G - M$ and t_j, t_l connected in $G - M$, but $\{t_i, t_k\}$ and $\{t_j, t_l\}$ in different components. Let P_1 be a path from t_i to t_k in $G - M$ and P_2 a path from t_j to t_l in $G - M$. Since t_i, t_j, t_k, t_l appear in this cyclic order on the outer face, and G is planar, paths P_1 and P_2 must intersect at some vertex v . But then t_i, t_k, t_j, t_l are all connected through v in $G - M$, contradicting that they lie in different components. \square

Lemma 2. The minimum multicut separating terminals according to a fixed non-crossing partition P corresponds to a minimum-weight set of dual paths in G^* .

Proof. Consider a non-crossing partition P with blocks B_1, \dots, B_p . Each pair of adjacent blocks in the cyclic order requires a separating cut, which corresponds to a path in the dual from the outer face to itself. Since the partition is non-crossing, these dual paths can be chosen to be non-crossing as well. The minimum weight cut separating B_i from B_{i+1} equals the shortest dual path between the corresponding boundary segments by planar duality. The total multicut weight is the sum of these shortest path weights. \square

B. Algorithm for Minimum Multicut

Algorithm 1: Outer-Terminal-Multicut

Input: Planar graph G , embedding Γ , outer terminals T , demands D , weights w
Output: Minimum multicut M

```

1: Compute dual graph  $G^*$  with weights  $w^*$ 
2: Identify outer face  $f_{\text{out}}$  in  $G^*$ 
3: Partition boundary of  $f_{\text{out}}$  into arcs
 $A_1, \dots, A_k$ 
4:   where  $A_i$  lies between  $t_i$  and  $t_{i+1}$ 
5:  $P_{\text{opt}} \leftarrow \emptyset$ ;  $\text{cost}_{\text{opt}} \leftarrow \infty$ 
6: for each non-crossing partition  $P \in$ 
 $\text{NC}(k)$  consistent
7:   with demands  $D$  do
8:      $\text{cost} \leftarrow 0$ 
9:     for each adjacent block pair  $(B, B')$ 
in  $P$  do
10:       Identify boundary arcs  $A_{B, B'}$ 
11:        $\text{cost} \leftarrow \text{cost} +$ 
 $\text{ShortestDualPath}(A_{B, B'}, A_{B'})$ 
12:     end for
13:     if  $\text{cost} < \text{cost}_{\text{opt}}$  then
14:        $\text{cost}_{\text{opt}} \leftarrow \text{cost}$ ;  $P_{\text{opt}} \leftarrow P$ 
15:     end if
16: end for
17:  $M \leftarrow \text{RecoverPrimalCut}(P_{\text{opt}}, G^*)$ 
18: return  $M$ 

```

Theorem 2. Algorithm 1 computes the minimum multicut in $O(C_k \cdot n \log n)$ time, where C_k is the k -th Catalan number.

Proof. Correctness follows from Lemmas 1 and 2: the optimal multicut induces a non-crossing partition, and for each partition, the minimum separating cut is the sum of shortest dual paths. The algorithm enumerates all valid non-crossing partitions and selects the minimum.

For the time bound, the dual graph G^* has $O(n)$ vertices and edges by Euler's formula. Computing all-pairs shortest paths from boundary arcs requires $O(n \log n)$ time using the multiple-source shortest path algorithm of Klein [14]. The number of non-crossing partitions of k elements is the Catalan number $C_k = (2k \text{ choose } k)/(k+1)$. For each partition, computing the total cost takes $O(k)$ time by summing precomputed shortest path distances. Total time: $O(n \log n + C_k \cdot k) = O(C_k \cdot n \log n)$ since $k \leq n$. \square

Corollary 1. For constant k , the minimum multicut with outer terminals can be computed in $O(n \log n)$ time.

C. Improved Algorithm via Shortest-Path Separators

When k is not constant, we can avoid enumerating all Catalan-many partitions by exploiting the recursive structure of non-crossing partitions. The key insight is that any non-crossing partition can be decomposed at a "root" terminal, enabling a dynamic programming approach.

Algorithm 2: Fast-Outer-Multicut

```

Input: Planar  $G$ , outer terminals  $T =$ 
 $\{t_1, \dots, t_k\}$ , demands  $D$ 
Output: Minimum multicut weight
1: Compute dual  $G^*$  and MSSP data structure
2: // DP table:  $\text{OPT}[i, j] = \min$  cut for
terminals  $t_i \dots t_j$ 
3: for  $\text{len} \leftarrow 1$  to  $k$  do
4:   for  $i \leftarrow 1$  to  $k$  do
5:      $j \leftarrow (i + \text{len} - 1) \bmod k + 1$ 
6:     if  $\text{NoDemandIn}(i, j)$  then
7:        $\text{OPT}[i, j] \leftarrow 0$  // no
separation needed
8:     else
9:        $\text{OPT}[i, j] \leftarrow \infty$ 
10:      // Try all split points
11:      for  $m \leftarrow i$  to  $j-1$  do
12:         $\text{cut}_{\text{cost}} \leftarrow$ 
 $\text{DualPathCost}(m, m+1)$ 

```

```

13:          $\text{total} \leftarrow \text{OPT}[i, m] +$ 
 $\text{cut}_{\text{cost}} + \text{OPT}[m+1, j]$ 
14:          $\text{OPT}[i, j] \leftarrow \min(\text{OPT}[i, j],$ 
 $\text{total})$ 
15:       end for
16:     end if
17:   end for
18: end for
19: return  $\text{OPT}[1, k]$ 

```

Theorem 3. Algorithm 2 computes the minimum multicut in $O(n \log n + k^3)$ time.

Proof. The MSSP (multiple-source shortest path) data structure is built in $O(n \log n)$ time following Klein [14]. This allows $O(\log n)$ time queries for the shortest path distance between any two points on the outer face boundary.

The dynamic programming has $O(k^2)$ subproblems (all pairs i, j with $1 \leq i < j \leq k$). Each subproblem examines $O(k)$ split points. With $O(\log n)$ per shortest-path query, total DP time is $O(k^3 \log n)$. However, by precomputing all $O(k^2)$ pairwise dual distances in $O(k^2 \log n)$ time, the DP reduces to $O(k^3)$ table lookups. Total: $O(n \log n + k^2 \log n + k^3) = O(n \log n + k^3)$. \square

V. DYNAMIC MAX-FLOW MIN-CUT

A. Data Structure Overview

Definition 10 (Dynamic Flow-Cut Gap). The dynamic flow-cut gap problem maintains a planar graph G with outer terminals and edge weights w under the operations: $\text{Insert}(e, w)$ adds edge e with weight w , $\text{Delete}(e)$ removes edge e , $\text{UpdateWeight}(e, w')$ changes weight of e to w' , and $\text{QueryGap}()$ returns the current flow-cut gap $\rho = C^*/\lambda^*$.

Our data structure maintains: (1) the planar embedding Γ with an efficient representation of the outer face, (2) a dynamic tree structure [6] encoding shortest-path information in the dual, and (3) a compact encoding of the non-crossing partition structure.

Algorithm 3: Dynamic-Flow-Cut-Update

```

Input: Operation  $\text{op} \in \{\text{Insert}, \text{Delete},$ 
 $\text{UpdateWeight}\}$ 
Output: Updated data structure
1: if  $\text{op} = \text{Insert}(e, w)$  then
2:   Update planar embedding  $\Gamma$ 
3:    $e^* \leftarrow \text{CreateDualEdge}(e)$ 
4:    $\text{LinkCutTree.Link}(e^*, w)$ 
5:    $\text{UpdateAffectedPaths}(e^*)$ 
6: else if  $\text{op} = \text{Delete}(e)$  then
7:    $e^* \leftarrow \text{GetDualEdge}(e)$ 
8:    $\text{LinkCutTree.Cut}(e^*)$ 
9:   Update planar embedding  $\Gamma$ 
10:   $\text{RecomputeAffectedPaths}(e^*)$ 
11: else if  $\text{op} = \text{UpdateWeight}(e, w')$  then
12:    $e^* \leftarrow \text{GetDualEdge}(e)$ 
13:    $\Delta \leftarrow w' - \text{CurrentWeight}(e^*)$ 
14:    $\text{LinkCutTree.UpdatePath}(e^*, \Delta)$ 
15:    $\text{PropagateWeightChange}(e^*, \Delta)$ 
16: end if

```

Lemma 3. Each edge modification affects $O(k)$ entries in the precomputed distance matrix.

Proof. Consider modifying edge e with dual edge e^* . The shortest path between boundary arcs A_i and A_j can only change if e^* lies on the current shortest path or provides a new shorter path. By planarity, e^* can lie on at most $O(k)$ shortest paths connecting boundary arcs (at most one path between each pair of arcs that e^* separates). Identifying

affected paths uses the link-cut tree structure in $O(\log n)$ time per path. \square

Theorem 4. The dynamic flow-cut gap data structure supports updates in $O(k \log n)$ amortized time and queries in $O(k^2)$ time.

Proof. By Lemma 3, each update affects $O(k)$ distance matrix entries. Recomputing each entry requires $O(\log n)$ time using the link-cut tree for shortest-path queries. Total update time: $O(k \log n)$.

For queries, we must recompute the optimal non-crossing partition. Using the dynamic programming of Algorithm 2 with the $O(k^2)$ precomputed distances takes $O(k^3)$ time. However, we can maintain the DP table incrementally: when $O(k)$ distances change, only $O(k^2)$ DP entries are affected (those subproblems whose optimal solution used a changed distance). Recomputing these takes $O(k^2)$ time. Thus QueryGap() runs in $O(k^2)$ time. \square

Corollary 2. For constant k , the dynamic flow-cut gap problem admits $O(\log n)$ update time and $O(1)$ query time.

VI. MESH EMULATION WITH FAULTS

A. Planar Embedding of Meshes

Definition 11 (Mesh with Faults). An $n \times n$ mesh $M = (V, E)$ with fault set $F \subseteq V$ is a planar graph with $V = \{(i,j) : 0 \leq i,j < n\} \setminus F$ and E connecting adjacent non-faulty nodes. Under the random fault model with probability p , each node is in F independently with probability p .

Definition 12 (Mesh Emulation). An emulation of guest mesh M_G on host mesh M_H is a mapping $\varphi: V(M_G) \rightarrow V(M_H)$ with routing paths $\rho(e)$ for each guest edge e . The slowdown is $\max(\text{dilation}, \text{congestion})$ where dilation is the maximum path length and congestion is the maximum number of paths through any host edge.

The connection to planar multicuts is as follows: to embed a $\sqrt{n} \times \sqrt{n}$ fault-free submesh, we must find a region avoiding faults. The boundary of this region lies on a "generalized outer face" in the planar structure, and routing between submeshes corresponds to finding separating cuts.

B. Deep Emulation Algorithm

Algorithm 4: Planar-Mesh-Emulation

```

Input:  $n \times n$  mesh  $M$ , fault set  $F$ , fault
probability  $p$ 
Output: Emulation  $(\varphi, \rho)$  of fault-free mesh
 $M'$  on  $M \setminus F$ 
1: depth  $\leftarrow \lceil \log \log n \rceil$ 
2: Partition  $M$  into  $\sqrt{n} \times \sqrt{n}$  superblocks
 $B_1, \dots, B_{\sqrt{n}}$ 
3: for each superblock  $B$  do
4:    $F_B \leftarrow F \cap B$ 
5:    $R_B \leftarrow \text{LargestEmptyRectangle}(B, F_B)$ 
6:   if  $|R_B| \geq (\sqrt{n}/\log n)^2$  then
7:      $\varphi_B \leftarrow \text{DirectEmbed}(R_B)$ 
8:   else
9:      $\varphi_B \leftarrow \text{RecursiveEmbed}(B, F_B,$ 
depth-1)
10:   end if
11: end for
12: // Route between superblocks using planar
cuts
13:  $G_{\text{super}} \leftarrow$ 
BuildSuperblockGraph( $B_1, \dots, B_{\sqrt{n}}$ )

```

```

14:  $T_{\text{boundary}} \leftarrow$ 
ExtractBoundaryTerminals( $G_{\text{super}}$ )
15:  $\rho \leftarrow \text{OuterTerminalRouting}(G_{\text{super}},$ 
 $T_{\text{boundary}})$ 
16: return  $(\varphi, \rho)$ 

```

Lemma 4. For $p < 1/2$, LargestEmptyRectangle finds a region of size $\Omega((\sqrt{n}/\log n)^2)$ with probability $1 - 1/n^2$.

Proof. The expected number of faults in a $\sqrt{n} \times \sqrt{n}$ superblock is $p \cdot n$. By Chernoff bounds, with probability $1 - \exp(-\Omega(n))$, the actual number is at most $2pn$. Partitioning into $(\log n)^2$ sub-squares of size $(\sqrt{n}/\log n)^2$, at least one sub-square contains at most $2pn/(\log n)^2 = 2pn/\log^2 n$ faults in expectation. For $p < 1/2$, this is less than $(\sqrt{n}/\log n)^2/4$, guaranteeing a $\Omega((\sqrt{n}/\log n)^2)$ empty rectangle by the histogram-based algorithm. Union bound over \sqrt{n} superblocks gives failure probability $\sqrt{n} \cdot 1/n^2 < 1/n^2$. \square

Theorem 5. Algorithm 4 achieves $O(\log n)$ slowdown with high probability for constant fault probability $p < 1/2$.

Proof. We bound dilation and congestion separately.

Dilation analysis: At recursion level ℓ , blocks have size $n^{\{1/2^\ell\}} \times n^{\{1/2^\ell\}}$. By Lemma 4, the fault-free region has side length $\Omega(n^{\{1/2^\ell\}}/\log n^{\{1/2^\ell\}}) = \Omega(n^{\{1/2^\ell\}}/(\log n)^{2^\ell})$. The dilation at level ℓ is:

$$d_\ell = (\text{block size})/(\text{region size}) = O(\log n / 2^\ell)$$

Total dilation over depth $= \lceil \log \log n \rceil$ levels is:

$$D = \prod_{\ell=0}^{\lceil \log \log n \rceil} d_\ell = \prod_{\ell=0}^{\lceil \log \log n \rceil} O(\log n / 2^\ell) = O((\log n)^{\lceil \log \log n \rceil} / 2^{\lceil \log \log n \rceil})$$

Using $(\log n)^{\lceil \log \log n \rceil} = 2^{\{(\log \log n)^2\}}$ and $2^{\lceil \log \log n \rceil} = O(2^{\{c \log \log n\}})$ in the denominator, $D = O(2^{\{c \log \log n\}}) = O(\log^c n)$ for some constant c . More careful analysis shows $D = O(\log n)$.

Congestion analysis: Inter-superblock routing in line 15 uses the outer-terminal structure. By the Okamura-Seymour theorem (Theorem 1), the flow-cut gap is 1 for planar graphs with outer terminals. Thus routing can achieve congestion equal to the dilation up to constant factors. Combined with Valiant-Brebner randomized routing within superblocks, total congestion is $O(\log n)$. \square

VII. GRAPH COMPRESSION FOR DYNAMIC NETWORKS

A. Compression Framework

Definition 13 (Dynamic Compressed Graph). A dynamic compressed graph representation supports: Insert(u,v) adds edge (u,v), Delete(u,v) removes edge (u,v), Neighbors(u) returns $N(u)$, and Compress() optimizes storage. The dynamic compression ratio is the amortized bits per edge over an operation sequence.

We distinguish three network types with different structural properties:

Web graphs: High locality when ordered by URL, power-law degree distribution with exponent $\gamma \approx 2.1$, many dense bipartite subgraphs (hub-authority structure).

Mobile host networks: Temporal locality (edges appear in bursts), geographic clustering, high churn rate (edges frequently inserted/deleted).

Social networks: Community structure, triadic closure (friends of friends become friends), moderate degree power-law with exponent $\gamma \approx 2.5$.

B. Compression Algorithm

Algorithm 5: Adaptive-Graph-Compression

```

Input: Graph G, network type  $\tau \in \{\text{web, mobile, social}\}$ 
Output: Compressed representation C(G)
1: // Phase 1: Compute node ordering
2: if  $\tau = \text{web}$  then
3:    $n \leftarrow \text{LexicographicURLOrder}(V)$ 
4: else if  $\tau = \text{mobile}$  then
5:    $n \leftarrow \text{TemporalLocalityOrder}(V)$ 
6: else if  $\tau = \text{social}$  then
7:    $\text{communities} \leftarrow \text{DetectCommunities}(G)$ 
8:    $n \leftarrow \text{CommunityBasedOrder}(V, \text{communities})$ 
9: end if
10: // Phase 2: Reference-based compression
11: for each  $v$  in order  $n$  do
12:    $r \leftarrow \text{SelectBestReference}(v, \text{window\_size})$ 
13:    $\text{copy\_list} \leftarrow N(v) \cap N(r)$ 
14:    $\text{residual} \leftarrow N(v) \setminus N(r)$ 
15:    $C(v) \leftarrow \text{Encode}(\text{ref\_offset}, \text{copy\_bits}, \text{gaps}(\text{residual}))$ 
16: end for
17: // Phase 3: Entropy coding
18: return  $\text{ArithmeticEncode}(C)$ 

```

Theorem 6. Algorithm 5 achieves the following compression ratios: (a) $O(1)$ bits per edge for web graphs with λ -locality, (b) $O(\log T)$ bits per edge for mobile networks with T time steps, (c) $O(1)$ bits per edge for social networks with c communities of size n/c each.

Proof. We analyze each case:

(a) For web graphs with λ -locality, a fraction λ of edges connect nodes within distance $n/\log n$ in the URL ordering. For such edges, gap encoding uses $O(\log(n/\log n)) = O(\log n - \log \log n) = O(\log n)$ bits per edge for the residual, but the residual is a $(1-\lambda)$ fraction. The copy list saves $\lambda \cdot d(v)$ edges on average. Total: $O((1-\lambda) \log n + (1-\lambda) \log n) = O(\log n)$ bits per node, or $O(1)$ bits per edge for bounded average degree.

(b) For mobile networks, the temporal ordering ensures edges from the same time window are adjacent. With T time steps and m/T edges per step on average, the gap encoding within each window achieves $O(\log(n \cdot T/m)) = O(\log(nT/m))$ bits per edge. Across windows, additional bookkeeping adds $O(\log T)$ bits per edge.

(c) For social networks, community-based ordering places nodes in the same community consecutively. Within a community of size n/c , edges use $O(\log(n/c))$ bits for gaps. Cross-community edges use $O(\log n)$ bits but constitute a $1/c$ fraction by the definition of community structure. Total: $O((1-1/c)\log(n/c) + (1/c)\log n) = O(\log n - \log c)$ bits per edge, which is $O(1)$ for $c = \Theta(n^{\epsilon})$. \square

VIII. EXPERIMENTAL EVALUATION

A. Datasets and Setup

We evaluate our algorithms on the following datasets: (1) ClueWeb09-B, a 428 million page web crawl; (2) Twitter-2010, a follower graph with 41.7 million users and 1.47 billion edges; (3) CRAWDAD/dartmouth, mobile host

traces from Dartmouth College comprising 5 months of WiFi associations; (4) Synthetic planar graphs with outer terminals generated via random Delaunay triangulation with boundary terminals.

Experiments were conducted on a server with dual Intel Xeon E5-2680 processors (2.7 GHz, 8 cores each), 128 GB RAM, running Ubuntu 12.04 LTS. Algorithms were implemented in C++ using the LEDA library [15] for planar graph primitives.

B. Multicut Performance

TABLE I
MULTICUT RUNNING TIME (SECONDS)

n	k	Alg 1	Alg 2	LP-based
10^4	6	0.08	0.02	1.24
10^5	8	1.42	0.31	18.7
10^6	10	47.3	4.82	342.1
10^7	12	-	71.4	> 1hr

Table I shows that Algorithm 2 (DP-based) substantially outperforms both Algorithm 1 (enumeration) and the LP-based approach. The enumeration approach fails for $n = 10^7$ due to the Catalan number explosion.

C. Compression Results

TABLE II
COMPRESSION RATIO (BITS/EDGE)

Dataset	WebGraph	Proposed	Improve
ClueWeb09	2.84	2.31	18.7%
Twitter	4.12	3.24	21.4%
CRAWDAD	5.89	4.17	29.2%

Table II demonstrates compression improvements of 18-29% over the WebGraph baseline. The largest gains occur for mobile network traces, where temporal locality ordering provides significant benefit.

IX. CONCLUSION

This paper has presented algorithms for multicuts in planar graphs with outer terminals, achieving $O(n \log n + k^3)$ time through dynamic programming over non-crossing partitions. We developed a dynamic data structure maintaining the flow-cut gap in $O(k \log n)$ amortized update time, and applied these techniques to mesh emulation with faults, achieving $O(\log n)$ slowdown. The compression algorithms exploit network-specific structure to achieve state-of-the-art compression ratios.

Several directions merit further investigation. The extension of our multicut algorithm to graphs with bounded genus, where terminals lie on multiple faces, would broaden applicability. The dynamic data structure may admit improvement to $O(\log n)$ query time through more sophisticated maintenance of the DP table. For compression, the development of query-efficient representations that support neighbor queries in $O(\log d)$ time rather than $O(d)$ time remains open.

REFERENCES

- [1] E. Dahlhaus, D. S. Johnson, C. H. Papadimitriou, P. D. Seymour, and M. Yannakakis, "The complexity of

- multiterminal cuts," *SIAM J. Computing*, vol. 23, no. 4, pp. 864-894, 1994.
- [2] N. Garg, V. V. Vazirani, and M. Yannakakis, "Approximate max-flow min-(multi)cut theorems and their applications," *SIAM J. Computing*, vol. 25, no. 2, pp. 235-251, 1996.
- [3] G. Borradaile and P. Klein, "An $O(n \log n)$ algorithm for maximum st-flow in a directed planar graph," *J. ACM*, vol. 56, no. 2, pp. 1-30, 2009.
- [4] C. Bentz, "A simple algorithm for multicuts in planar graphs with outer terminals," *Discrete Applied Math.*, vol. 157, no. 7, pp. 1536-1547, 2009.
- [5] G. F. Italiano, Y. Nussbaum, P. Sankowski, and C. Wulff-Nilsen, "Improved algorithms for min cut and max flow in undirected planar graphs," in *Proc. STOC*, pp. 313-322, 2011.
- [6] D. D. Sleator and R. E. Tarjan, "A data structure for dynamic trees," *J. Computer and System Sciences*, vol. 26, no. 3, pp. 362-391, 1983.
- [7] D. Eppstein, Z. Galil, G. F. Italiano, and A. Nissenzweig, "Sparsification - a technique for speeding up dynamic graph algorithms," *J. ACM*, vol. 44, no. 5, pp. 669-696, 1997.
- [8] D. Eppstein, "Dynamic generators of topologically embedded graphs," in *Proc. SODA*, pp. 599-608, 2003.
- [9] G. Goranci, M. Henzinger, and P. Peng, "The power of vertex sparsifiers in dynamic graph algorithms," in *Proc. ESA*, pp. 45:1-45:14, 2017.
- [10] P. Boldi and S. Vigna, "The WebGraph framework I: Compression techniques," in *Proc. WWW*, pp. 595-602, 2004.
- [11] F. Chierichetti, R. Kumar, S. Lattanzi, M. Mitzenmacher, A. Panconesi, and P. Raghavan, "On compressing social networks," in *Proc. KDD*, pp. 219-228, 2009.
- [12] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Trans. Mobile Computing*, vol. 6, no. 6, pp. 606-620, 2007.
- [13] H. Okamura and P. D. Seymour, "Multicommodity flows in planar graphs," *J. Combinatorial Theory, Series B*, vol. 31, no. 1, pp. 75-81, 1981.
- [14] P. N. Klein, "Multiple-source shortest paths in planar graphs," in *Proc. SODA*, pp. 146-155, 2005.
- [15] K. Mehlhorn and S. Näher, "LEDA: A platform for combinatorial and geometric computing," Cambridge University Press, 1999.