

Lecture 10: Decremental Strongly Connected Components and  
Single-Source Reachability in Near-Linear TimeLecturer: *Ilya Maier, Ijay Narang, Baris Onat*Scribe: *Hareton Song*

## Introduction and Problem Statement

- **Goal:** Maintain Strongly Connected Components (SCCs) of a directed graph  $G$  while supporting edge deletions.
  - The paper aims to maintain SCCs in a directed graph  $G$  while supporting edge deletions dynamically. SCCs are essential structures in graph theory, where each SCC represents a subset of vertices such that there is a directed path between any pair of vertices within the same subset. Efficiently maintaining SCCs is critical for dynamic graph applications.
  - In dynamic graphs, where edges are deleted over time (known as decremental graphs), SCCs need to be continuously updated. The challenge is to do this efficiently, avoiding the costly recomputation of SCCs after each deletion. The goal of this work is to develop an algorithm that supports these updates efficiently, especially in large-scale networks or systems.
- **Operations:**
  - **Delete( $u, v$ ):**
    - \* This operation removes a directed edge from vertex  $u$  to vertex  $v$  in the graph. After the deletion, SCCs must be updated to reflect the possible new disconnection or subdivision of existing SCCs.
    - \* The complexity of handling this deletion efficiently is one of the core focuses of the paper.
  - **Query( $u, v$ ):**
    - \* This operation checks if two vertices  $u$  and  $v$  are still part of the same SCC after one or more edge deletions. Efficient querying after a series of updates is crucial for maintaining the dynamic structure of the graph in real-time applications.

- \* The paper provides an efficient way to answer such queries in near-optimal time.
- **Main Result:** Achieve this in  $\tilde{O}(m)$  total time, where  $\tilde{O}$  hides polylogarithmic factors.
  - The paper presents an algorithm that achieves these operations in  $\tilde{O}(m)$  total time, where  $m$  is the number of edges in the graph and  $\tilde{O}$  hides polylogarithmic factors. This means the algorithm scales well with large graphs and achieves nearly linear time complexity when handling both deletions and queries.
  - This result represents a significant improvement over previous methods, which either had higher time complexity or could not handle general decremental graphs efficiently.

## Definitions

- **SCC:** A maximal subgraph where each vertex is reachable from every other vertex.
- **Condensation:** Contract each SCC into a single vertex, resulting in a DAG.
- **Diameter of Subgraph:** Maximum distance between any two vertices within the subgraph.

## Dynamic SCC Maintenance

- **Challenge:** Edge deletions may split SCCs, making efficient updates difficult without recalculating everything.
  - One of the central challenges in maintaining SCCs dynamically lies in handling edge deletions. When an edge is deleted from a directed graph, it can cause an SCC to split into smaller components. If this happens, recalculating the SCC structure from scratch would be computationally expensive, especially in large graphs. The goal is to efficiently update the SCCs without recomputing them entirely after each deletion.
  - **Why is this hard for directed graphs?** Unlike undirected graphs, where connectivity is symmetric (if a node is reachable

from another, the reverse is also true), directed graphs have a one-way notion of connectivity. This directional constraint means that the removal of a single edge can dramatically affect the reachability between vertices, which makes SCC maintenance more complex and costly in terms of computation.

- **Previous Work:** Efficient for undirected graphs; directed graphs have additional complexity due to directional constraints.
  - Algorithms for maintaining connectivity under edge deletions have been well-studied for undirected graphs, where methods can more easily handle updates. For undirected graphs, approaches like dynamic tree structures or union-find algorithms allow efficient connectivity maintenance.
  - Directed graphs, however, introduce additional complexities because the directed nature of edges introduces more nuanced changes in SCC structures upon deletions. As a result, maintaining SCCs in directed graphs has historically been more challenging, and previous algorithms have been either less efficient or less general in their approach to handling arbitrary deletions.
  - Earlier work on dynamic graph problems either focused on undirected graphs or failed to achieve optimal performance for large, decremental directed graphs. The need for an efficient method specifically tailored to decremental directed graphs—where only edge deletions are allowed, not insertions—has been a long-standing open problem.

## Approach Overview

- Use **Hierarchical Decomposition** of the graph to maintain SCCs efficiently:
  - The graph is decomposed into multiple levels  $G_0, G_1, \dots, G_{\log n+1}$ , with each level containing a subset of the graph's edges. The SCCs at each level are incrementally condensed into smaller nodes, creating a more simplified version of the graph at higher levels. This structure allows the algorithm to handle updates at different levels of granularity.
  - **Why Hierarchical Decomposition Works:**

- \* By partitioning the graph and condensing SCCs into smaller components, the algorithm avoids global recomputation. At each level, updates only affect a localized subset of the graph, making the algorithm more efficient.
  - \* The idea of condensing SCCs into single nodes at higher levels helps reduce the problem size. Since fewer edges and nodes remain at higher levels, the algorithm can process updates more quickly. This is crucial for large graphs where recalculating reachability globally would be prohibitively expensive. When an edge is deleted, the updates only need to affect the specific SCCs at lower levels. The hierarchy helps localize these updates to the relevant subgraph, minimizing the need for global changes.
- Maintain low-diameter SCCs at each level to ensure efficient updates.
    - Generalized ES Trees (GES Trees):**
      - The GES Trees are an extension of the traditional Even-Shiloach Trees (ES Trees). They help maintain reachability information for condensed nodes at different levels in the hierarchy. GES Trees maintain shortest-path trees for both in-reachability and out-reachability. When an edge is deleted, GES Trees help update the reachability information for all vertices in the affected SCC. The trees dynamically adjust to reflect the new structure of the graph without recalculating reachability from scratch.
      - **S-Distance:** GES Trees utilize a modified metric called S-Distance, which is used to calculate the shortest paths between condensed nodes. The S-Distance considers both forward and backward paths, ensuring that reachability is maintained efficiently across the condensed SCCs.
        - \* **How S-Distance Works:** S-Distance computes distances between condensed nodes, ensuring that the trees can efficiently update reachability information across all levels of the hierarchy. It modifies the classical notion of distance to incorporate both in-reachability and out-reachability within a condensed SCC, which is critical for efficient SCC maintenance.

## Key Components:

1. **ES Trees (Even-Shiloach Trees):**
  - Maintain *reachability information* within SCCs dynamically.
  - Use two ES trees per SCC: *in-tree* and *out-tree* for managing incoming and outgoing edges, respectively.
2. **Generalized ES Trees (GES Trees):**
  - Extend ES Trees to handle *condensed nodes* at various levels.
  - Introduce S-Distance: A modified metric to manage distances between condensed nodes.
3. **Separators:**
  - Separators are used to split large SCCs when their diameter exceeds a threshold.
  - Ensures SCCs remain small and manageable, facilitating efficient updates.

## Handling Edge Deletions:

- **Deletion Workflow:**
  1. **Remove the Edge:** When an edge  $(u, v)$  is deleted, the algorithm first removes the edge from the graph.
  2. **Update GES Trees:** After the edge removal, the GES Trees are updated to reflect the new reachability information. The algorithm checks whether the deletion affects the SCC by recalculating the in-tree and out-tree for the affected vertices.
  3. **SCC Splitting:** If the deletion causes the SCC to split into smaller SCCs, the separator mechanism is triggered. This ensures that any large SCC that has grown beyond a manageable size is split into smaller, low-diameter SCCs.
  4. **Recompute the GES Trees:** Once the SCCs are split, new GES Trees are constructed for each of the newly formed SCCs. This allows the algorithm to efficiently maintain the updated reachability information.
- **Separator Lemma:** Guarantees that SCCs can be split efficiently while maintaining low diameter.
  - \* **Separators in SCC Maintenance:**

- A separator is a small set of vertices that, when removed, breaks an SCC into smaller components. The Separator Lemma guarantees that any SCC whose diameter exceeds a set threshold can be efficiently split into smaller SCCs. The algorithm uses separators to keep SCCs small and manageable, ensuring that updates remain efficient even after large changes to the graph.
- \* **Separator Lemma Guarantees:**
  - The Separator Lemma states that after removing a small set of vertices (the separator), the remaining SCCs will have low diameter, making it easier to maintain reachability and efficiently handle further deletions.

## Algorithmic Highlights:

### – Hierarchy Construction:

- \* Start with the original graph.
- \* Incrementally add edges and form SCCs, then *condense* them into nodes at each level.

This ensures that the graph structure is maintained efficiently as the algorithm progresses through different levels of granularity.

### – GES Tree Operations:

- \* *Node Splitting*: When SCCs split, corresponding nodes in the GES Tree are split, and paths are recalculated. The algorithm reassigns the vertices to new nodes, ensuring that the reachability information remains consistent.
- \* *Distance Maintenance*: Use S-Distance to ensure consistency in shortest paths among SCCs. This allows the system to handle the changes efficiently without recalculating the entire graph structure.

## Complexity:

Overall Summary of Complexity:

- Total update time is  $\tilde{O}(m)$ , which is near-linear and efficient for large-scale applications.

- The hierarchical approach avoids the need for *global recalculation* after each edge deletion, which would be impractical for large graphs.

Details of this section:

The total update time for the algorithm is  $\tilde{O}(m)$ , which is nearly linear. This makes the approach highly efficient for large-scale graphs, especially those where edges are frequently deleted. The hierarchical decomposition ensures that the algorithm only updates the necessary parts of the graph after each deletion, avoiding the need for global recomputation.

### **Key Complexity Insights:**

- The GES Tree operations are efficient because they focus on localized updates within each SCC.
- Separators are used to keep the SCCs small, reducing the overhead of maintaining large, complex SCCs.
- The hierarchical structure minimizes the number of nodes and edges that need to be recalculated, contributing to the overall efficiency.

### **Applications:**

- **Network Resilience:** Maintaining connectivity in communication networks when links fail.
- **Dependency Management:** Tracking interdependencies in software modules as components are updated or removed.
- **Social Networks:** Analyzing community structures as relationships change.

### **Conclusion:**

- The hierarchical approach, combined with specialized data structures like Generalized ES Trees and separators, enables efficient maintenance of SCCs under edge deletions.
- Future work may include handling fully dynamic graphs (both insertions and deletions) and improving deterministic solutions for SCC maintenance.

## Some Q&As in the Lecture

1. **Q: How does reachability update when an edge is deleted?**
  - **A:** ES Trees are updated for both *in-tree* and *out-tree* to determine which nodes are still connected. If an SCC splits, each new component is restructured with its own ES Tree. This ensures that reachability within each new SCC is maintained efficiently.
2. **Q: When do you decide to split an SCC using separators?**
  - **A:** An SCC is split if its *diameter* exceeds a set threshold. In such cases, a *separator* is identified to divide the SCC into smaller parts, which are more manageable. This helps keep the system efficient, ensuring no SCC becomes too complex.
3. **Q: Can this approach be applied to software dependency analysis?**
  - **A:** Yes, SCCs in a dependency graph represent tightly interconnected modules. If a dependency (edge) is removed, the approach helps determine which modules are still functional or if they are isolated from the rest. It's an effective way to manage module dependencies during system changes.
4. **Q: What is the practical implication of maintaining low diameter in SCCs?**
  - **A:** Keeping SCCs with *low diameter* ensures that reachability queries and updates are efficient. In practice, it means that the subgraphs remain small and easy to manage, which is critical for performance in large-scale systems.
5. **Q: How scalable is this approach for large graphs?**
  - **A:** The near-linear runtime  $\tilde{O}(m)$  makes it highly scalable, especially for large, sparse graphs. The hierarchical decomposition ensures that only parts of the graph are updated at each deletion, preventing the need for full recomputation. Practical concerns include managing memory for ES Trees and ensuring separators are efficiently handled.
6. **Q: What are the limitations of this approach for highly dynamic graphs?**
  - **A:** The current decremental focus means it's best suited for scenarios where *edges are only removed*. Extending it to handle frequent insertions would require different strategies, as



maintaining SCCs in a fully dynamic setting is significantly more complex. Practical application in environments with *frequent insertions* remains an open area for research.