# Optimizing **Dynamic Graph Processing with** the Locality-First Strategy



#### Helen Xu Lawrence Berkeley National Laboratory **ACDA Workshop**

#### **Real-World Graphs Are Sparse**

**Sparse graphs**, which have many fewer edges than the total possible number of edges, underlie most real-world applications.



Social networks

Sparsity **disrupts locality** due to the presence of many zeroes in the data.

Computational biology



**Road networks** 

<sup>...</sup>and others!





#### **Real-World Graphs Are Also Dynamic**

Furthermore, many real-world sparse graphs are **dynamic**: they **change** over time.





Systems for processing dynamic graphs support updates (e.g. edge insertions and deletions) and queries (algorithms run on the graph).

colocating and updating data (e.g., in CSR).



Updates

- Dynamic graphs disrupt locality because of the inherent tradeoff between



### Multicore Optimization Enables Fast Graph Queries and Updates

Despite these challenges to locality, high-performance dynamic-graphprocessing systems such as Aspen [DhulipalaShBI19] have taken huge steps towards **efficient queries and updates**.

On 48 cores, Aspen runs the following queries on Twitter (2.4B edges):





Breadth-first search 0.32s

Times are human-measurable even with parallelism, demonstrating the importance of efficient processing



PageRank 24.03s Betweenness centrality 4.72s



#### **Query Speed in Dynamic Graph Systems** Terrace [PandeyWhXuBu21], a dynamic graph processing system, optimizes further with a "locality-first design" that takes advantage of graph structure.



[PandeyWhXuBu21] Pandey, Wheatman, Xu, Buluç. "Terrace: A Hierarchical Graph Container for Skewed Dynamic Graphs." SIGMOD '21.

Both systems support parallelization.

Both systems run the same algorithms by implementing the Ligra [ShunBI13] abstraction.

Surprisingly, in some cases, **Terrace achieves speedup** on queries over Ligra [ShunBI13], a static graph system.





### Updatability of Dynamic Graph Systems

Terrace achieves the best of both worlds in query and update performance by taking advantage of locality.



Edges were generated using an rMAT distribution [ChakrabatiZhFa04] and added in batches using the provided API.

Terrace achieves up to 48M inserts per second and up to 9M deletes per second. Future work includes optimizing batch deletions.





#### **Dynamic Graph Processing and** the Locality-First Strategy



#### Add **parallelism** into data







#### **Understanding Locality in Graph Queries**

Systems for processing dynamic graphs must support fast graph queries.

Vertex scans, or the **processing of a vertex's incident edges**, are a crucial step in graph queries [ShunBI13].





#### Most Graph Systems Separate Neighbor Lists for Parallelization

Existing dynamic graph systems optimize for parallelism first with separate per-vertex data structures e.g. trees [DhulipalaBISh19], adjacency lists

[EdigerMcRiBa12], and others [KyrolaBIGu12, BusatoGrBoBa18, GreenBa16].

**Vertex IDs Pointers to edges** 

Edges



Weakness: Separating the data structures **disrupts spatial locality**.



#### **Dynamic Graph Processing and** the Locality-First Strategy



**Problem:** Dynamic graph processing



**Understand** locality in dynamic graph processing

\$







#### Add **parallelism** into data structures







### Enhancing Spatial Locality by **Colocating Neighbor Lists**

Idea: Colocate neighbor lists in the same data structure, which avoids cache **misses** when traversing edges in order.



Question: Do these misses actually affect performance, or are they a low-order term?

[WheatmanXu21] Wheatman and Xu. "A Parallel Packed Memory Array to Store Dynamic Graphs." ALENEX '21.



11

### **Dynamic Graphs Are Often Skewed**

Real-world dynamic graphs, e.g. social network graphs, often follow a **skewed** (e.g. power-law) distribution with a **few high-degree vertices and many low-degree vertices** [BarabasiAl99].

Example power law:



10 neighbors	% < 1000 neighbors		
64.56	99.51		
Twitter follo		These high de for e maxim the Tw abo [Bea	graphs exhibit gree variance: xample, the um degree in vitter graph is ut 3 million amerAsPa15]

12

#### Insight: Locality-First Skew-Aware Design

# Next step: refine the solution with a hierarchical design that takes





**Problem: High-degree** vertices slow down updates for all vertices in the shared data structure

advantage of skewness while maintaining locality as much as possible.





#### Implementing the Hierarchical Skew-Aware Design

Terrace implements the locality-first hierarchical design with **cache-friendly data structures.** 



14

#### **Selecting Data Structures for Dynamic Graphs**

Given a cache block size B and input size N, B-trees and PMAs take  $\Theta(N/B)$ block transfers to scan.

B-tree inserts take  $O(\log_R(N))$  transfers, while PMA inserts take  $O(\log^2(N))$ .



In theory, B-trees [BayerMc72] asymptotically dominate Packed Memory Arrays (PMA) [ItaiKoRo81, BenderDeFa00] in the classical external-memory model [AggarwalVi88].

#### **Exploiting Skewness for Cache-Friendliness**

# The locality-first design in Terrace reduces cache misses during graph







## Terrace: Applying the Locality-First Strategy to Dynamic Graph Processing

In practice, Terrace is about **2x faster on graph query algorithms** than Aspen while **maintaining similar updatability**.

Terrace's **cache-friendly design** demonstrates the impact of **the localityfirst strategy** in graph processing.

\$



**Problem:** Dynamic graph processing

Understand locality: opportunities for spatial locality due to skewness



Exploit spatial locality with a cache-friendly skewaware data structure



Implementation of Terrace, a **parallel** dynamic-graphprocessing system based on the **skew-aware design** [PandeyWhXuBu21]

https://github.com/PASSIONLab/terrace

