Tight Dynamic Problem Lower Bounds from Generalized BMM and OMv

Ce Jin and Yinzhan Xu (MIT)

Summary: We prove new tight fine-grained lower bounds for various **dynamic problems**, using combinatorial k-clique hypothesis and (generalization of) OMv hypothesis.

Fine-grained Hypotheses we use

Combinatorial k-clique hypothesis:
No *combinatorial* algorithm can detect k-clique in an n-node graph in $O(n^{k-\varepsilon})$ time, for any $\varepsilon > 0$.

Combinatorial algorithm: An informal notion. Refers to algorithms that do not use Fast Matrix Multiplication (e.g., Strassen's).

OuMv$_k$ hypothesis:
For constant integer $k \geq 2$, the following problem cannot be solved in $O(n^{1+k-\varepsilon})$ time, for any $\varepsilon > 0$:

- Pre-process a subset $M \subseteq \{1,2,\dots,n\}^k$
- Answer n online queries:
  Given k sets $U^{(1)}, U^{(2)}, \dots, U^{(k)} \subseteq \{1,2,\dots,n\}$,
  Is $(U^{(1)} \times U^{(2)} \times \dots \times U^{(k)}) \cap M$ non-empty?

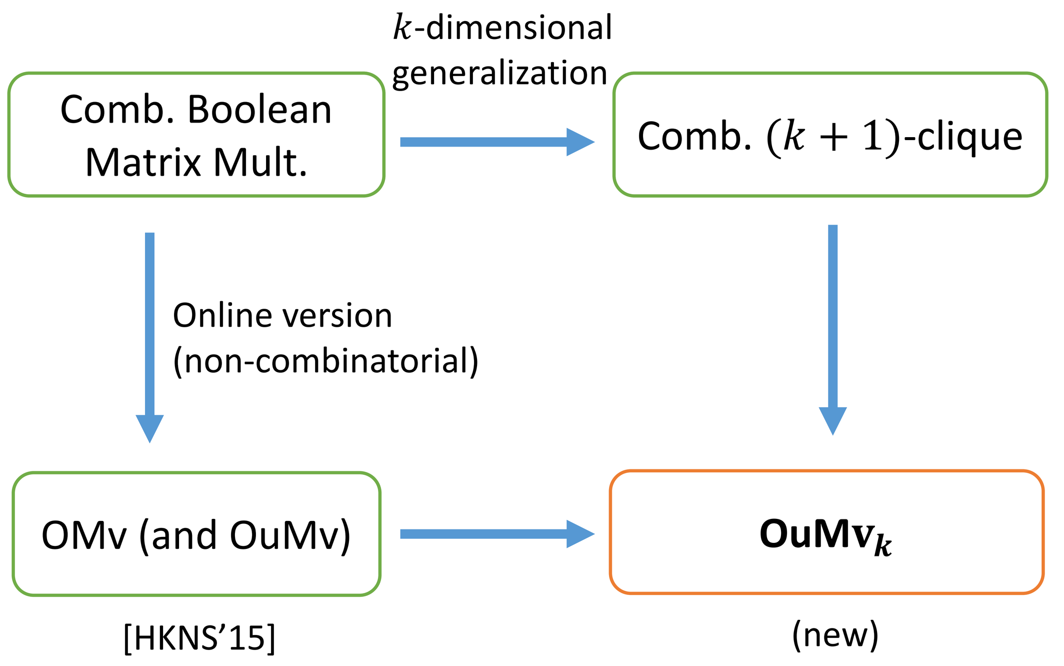

Problem Definitions

Dynamic Range-Mode:
- Maintain an integer array $A[1], A[2], \dots, A[n]$
- Support Insertions and Deletions
- Query l, r: what is the most frequent element in $A[l], A[l+1], \dots, A[r]$? (breaking ties arbitrarily)

Dynamic Subgraph Connectivity
- Preprocess a static undirected graph G with m edges
- Maintain a dynamic vertex subset S (vertices that are "on")
- Turn on vertex u: $S \leftarrow S \cup \{u\}$
- Turn off vertex u: $S \leftarrow S \setminus \{u\}$
- Query u, v: are u and v connected in the induced subgraph $G[S]$?

Dynamic Skyline Points Counting
- Maintain a set P of n points in $\mathbf{R}^d$
- Insertion: $P \leftarrow P \cup \{x\}$
- Deletion: $P \leftarrow P \setminus \{x\}$
- Query: how many "skyline points" does P have?
  - $x \in P$ is a "skyline point" iff no other $y \in P$ dominates x (i.e. $y_i \geq x_i$ for all $i = 1,2,\dots,d$)

New LB for Dynamic Range Mode

Our result: No combinatorial algorithm for *Dynamic Range Mode* can achieve $\boldsymbol{n^{2/3-\varepsilon}}$ query time, $\boldsymbol{n^{2/3-\varepsilon}}$ update time and poly(n) preprocessing time, under combinatorial 4-clique hypothesis

A proof template for dynamic lower bounds:

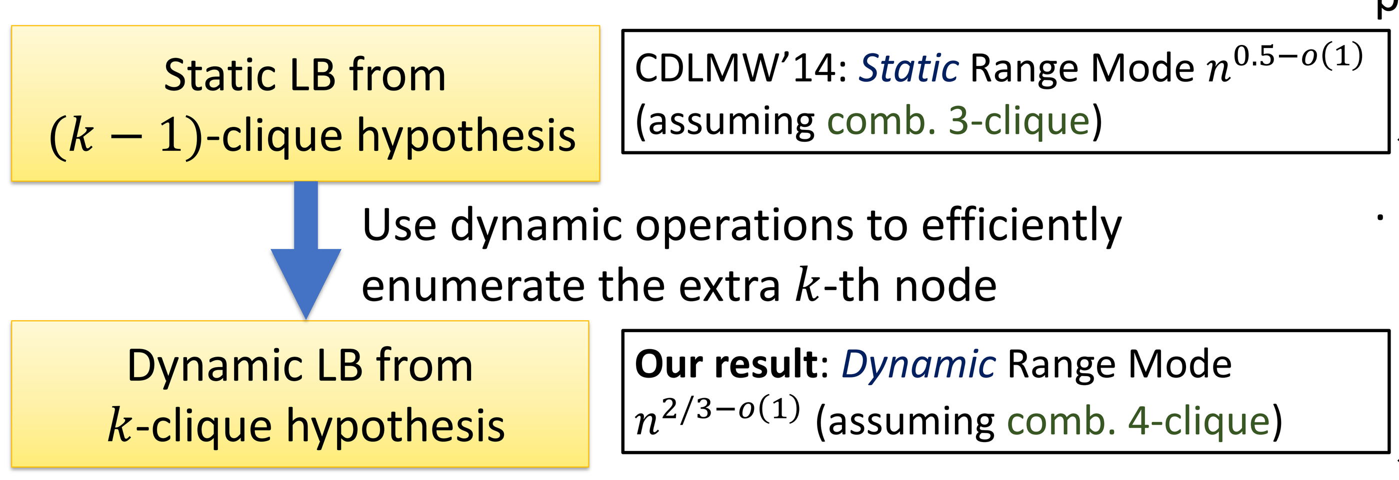

Other results: tight combinatorial LBs for more problems using this proof template:
- Dynamic 2D Orthogonal Range Color Counting
  $n^{2/3-o(1)}$ time (from comb. 4-clique)
- Dynamic d-Dimensional Orthogonal Range Mode
  $n^{1-\frac{1}{2d+1}-o(1)}$ time (from comb. $(2d+2)$-clique)
- Dynamic 2-Pattern Document Retrieval
  $n^{2/3-o(1)}$ time (from comb. 4-clique)
- ...

Proof:
- 4-clique on a (unbalanced) 4-partite graph with $V = A \uplus B \uplus C \uplus D$ where $|A| = |B| = n^{1/3}, |C| = n^{100}, |D| = n^{2/3}$ requires $|A||B||C||D| = n^{4/3} \cdot |C|$ time.
- $N_D(v)$ denotes the neighbors of v in D, and $\overline{N_D(v)}$ denotes the non-neighbors of v in D.
- Build the following array, where each blue block contains a permutation of D. There are $|A|$ blue blocks on the left and $|B|$ blue blocks on the right. (CDLMW'14)
  Then, the frequency of the mode in this range tells whether $N_D(a_i) \cap N_D(c) \cap N_D(b_j)$ is non-empty.
- For $c \in C$:
  Use $|D|$ updates to build the middle $N_D(c)$ block
  For all $(a_i, b_j) \in A \times B$ such that (a_i, b_j, c) is a K_3:
  Use one query to tell whether this K_3 extends to a K_4

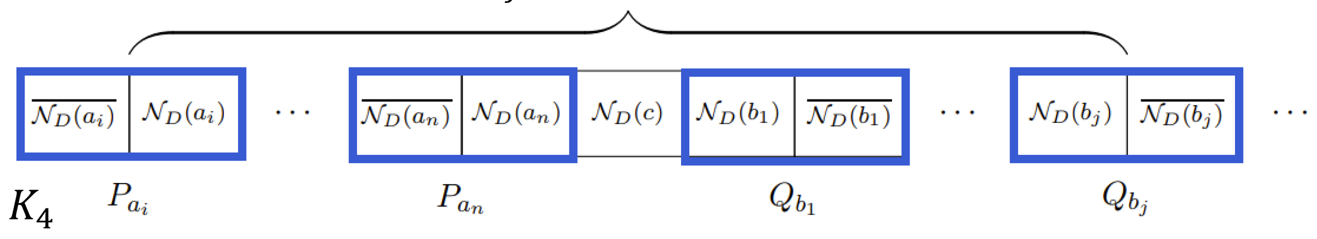

- Total #operations = $|C| \cdot (|D| + |A| \cdot |B|) = |C| \cdot n^{2/3}$
  If preprocessing takes $< n^{100}$ time, then each operation needs $n^{2/3}$ time.

New LB for Dynamic Subgraph Connectivity

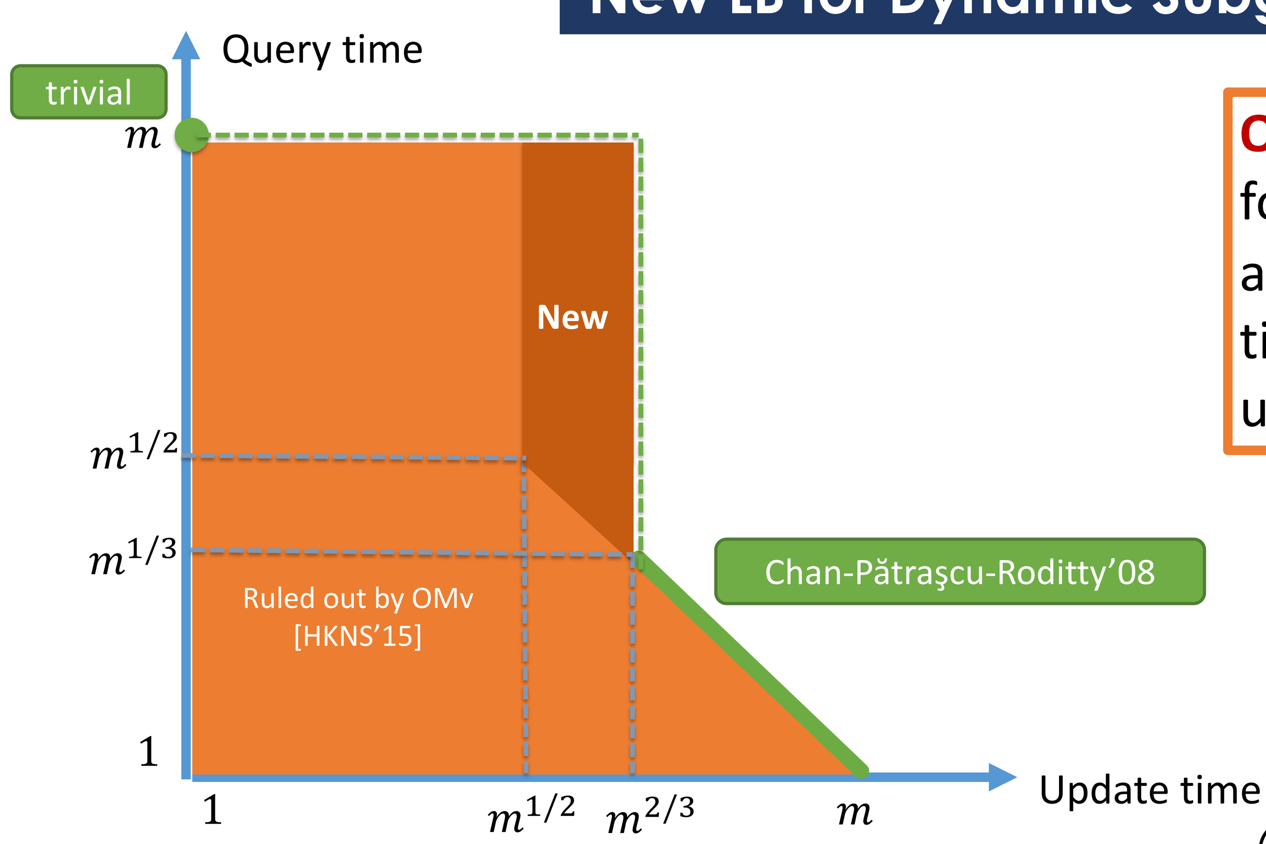

Our result: No combinatorial algorithm for *Dynamic Subgraph Connectivity* can achieve $\boldsymbol{m^{2/3-\varepsilon}}$ update time, $\boldsymbol{m^{1-\varepsilon}}$ query time, and poly(m) preprocessing time, under comb. 4-Clique hypothesis

Proof:
- 4-clique on a (unbalanced) 4-partite graph with $V = A \uplus B \uplus C \uplus D$ where $|B| = |C| = m^{1/3}, |A| = m^{2/3}, |D| = m^{100}$ requires $|A||B||C||D| = m^{4/3} \cdot |D|$ time.
- Construct the following (static) graph with $O(m)$ edges.
- For each $d \in D$:
  Use $|A|$ updates so that $a \in A_2$ is on iff $(a, d) \in E$
  For $b \in B$ such that $(b, d) \in E$:
    Turn on b and turn off all $B \setminus \{b\}$ (only $O(1)$ updates)
    For $c \in C$:
      Let c be on iff $(c, d), (c, b) \in E$
    If s, t are connected then return True
  Return False

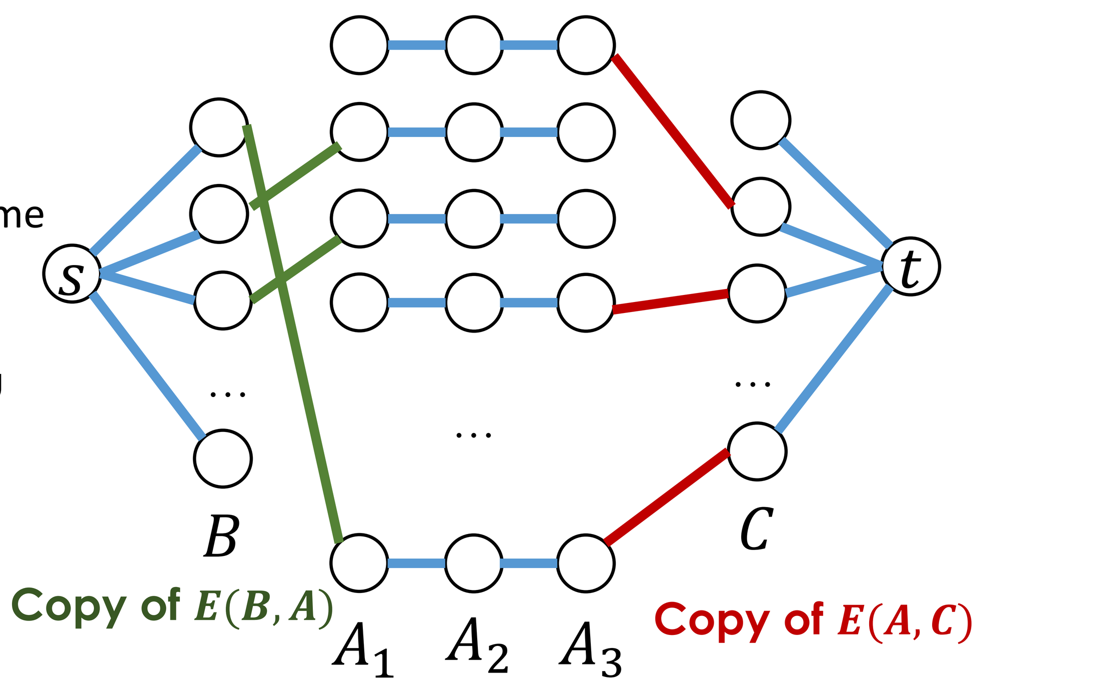

- #updates = $|D| \cdot (|A| + |B| \cdot (1 + |C|)) = |D| \cdot m^{2/3}$
- #queries = $|D| \cdot |B| = |D| \cdot m^{1/3}$
  If preprocessing takes $< m^{100}$ time, then either update time $\geq m^{2/3}$ or query time $\geq m$.

LBs for Geometry problems from OuMv_k Hypothesis

We consider **OuMv$_k$ hypothesis** which directly generalizes OMv and OuMv [Henzinger-Krinninger-Nanongkai-Saranurak'15] to higher dimensions. This leads to tight combinatorial LBs for various dynamic geometric problems. Consider the following example:

Chan'03 (adapted) gave a *semi-online* algorithm for *Dynamic Skyline Points Counting* in $\mathbf{R}^{2k-1}$ with $\tilde{O}(n^{1-1/k})$ update time.
(Semi-online model is between online and offline: When x is inserted, we are told when x will be deleted in the future)
Our result: this running time is tight under OuMv$_k$ hypothesis

(The $k = 2$ case (in $\mathbf{R}^3$) based on OMv was recently independently proved by Dallant & Iacono (2021))

Other tight LBs from OuMv$_k$ hypothesis:
- Chan's Halfspace problem in $\mathbf{R}^k$ (Chan'03)
- k-dimensional Erickson's problem
- $(k+1)$-dimensional Langerman's problem
- ...

Some Open Questions

- Can Dynamic Subgraph Connectivity have update time better than $m^{2/3}$ using fast matrix multiplication?
- What is the optimal update time for Dynamic Skyline Points Counting in $\mathbf{R}^{2k}$? (offline algorithms allowed)
- What is the optimal update time for Dynamic Skyline Points Counting in $\mathbf{R}^3$ without assuming semi-online model?
  (Chan'20: amortized $\tilde{O}(n^{2/3})$ algorithm. Our LB: $n^{1/2-o(1)}$)