All-Pairs Shortest Paths and Fine-Grained Complexity

Timothy Chan
UIUC

(joint work with Virginia Vassilevska Williams (MIT) & Yinzhan Xu (MIT))

(CALDAM'22)

APSP

- Given weighted dense graph $G = (V, E)$ with n vertices, for every pair $u, v \in V$, compute $D[u, v] =$ shortest path distance from u to v

APSP: Background

$O(n^3)$	Floyd–Warshall'62
$O^*(n^3/\log^{1/3} n)$	Fredman'75
$O^*(n^3/\sqrt{\log n})$	Takaoka'92, Dobosiewicz'90
$O^*(n^3/\log^{5/7} n)$	Han'04
$O^*(n^3/\log n)$	Takaoka'04, Zwick'04, C.'05
$O^*(n^3/\log^{5/4} n)$	Han'06
$O^*(n^3/\log^2 n)$	C.'07
$O^*(n^3/\log^3 n)$	C.'17
$O(n^3/c^{\sqrt{\log n}})$ rand.	Williams'14
$O(n^3/c^{\sqrt{\log n}})$ det.	C.–Williams'15

- APSP Hypothesis: no $O(n^{3-\varepsilon})$ alg'm (even for integer weights)

Hardness from APSP

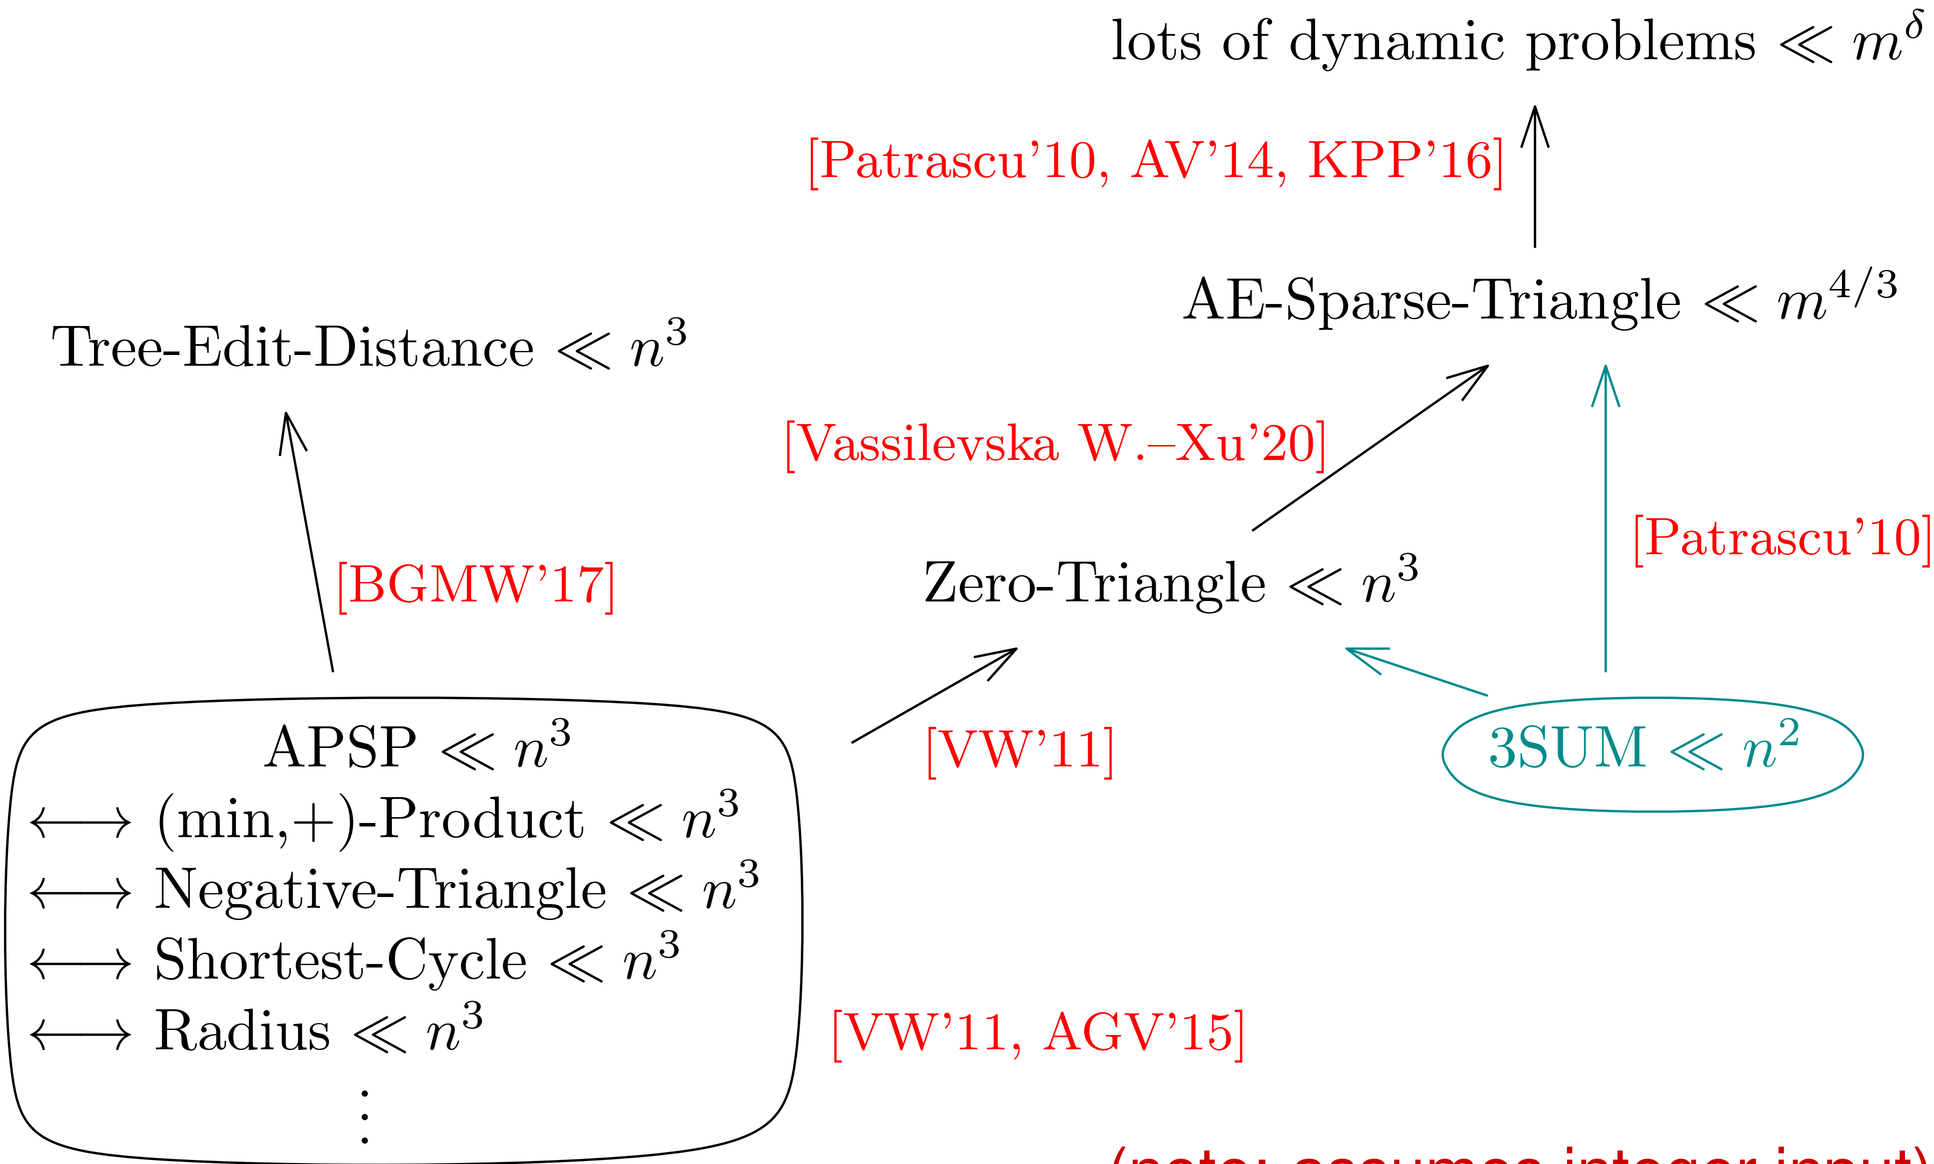

(note: assumes integer input)

This Talk

- hardness from APSP for unweighted graphs
  (or small $\widetilde{O}(1)$ integer weights)
  [C.–Vassilevska W.–Xu, ICALP'21]

- hardness from APSP for real-weighted graphs
  [C.–Vassilevska W.–Xu, STOC'22, to appear]

unwt-APSP: Background

unweighted undirected:

$\widetilde{O}(n^{\omega})$ $\leq O(n^{2.373})$ Alon–Galil–Margalit'91, Seidel'92

unweighted directed:

$\widetilde{O}(n^{(\omega+3)/2})$ $\leq O(n^{2.687})$ Alon–Galil–Margalit'91

$\widetilde{O}(n^{2+\rho})$ $\leq \boxed{O(n^{2.529})}$ Zwick'98

(where ρ satisfies $\omega(1, \rho, 1) = 1 + 2\rho$)

- unwt-dir-APSP Hypothesis: no $O(n^{2.5-\varepsilon})$ alg'm

Prelims on Matrix Multiplication

- given $n_1 \times n_2$ matrix A and $n_2 \times n_3$ matrix B with integer entries in $[\ell] := \{0, \ldots, \ell\}$
- let $M(n_1, n_2, n_3)$ be time to compute standard product

$$(AB)[i,j] \; = \; \sum_k A[i,k]B[k,j]$$

- Current bds:

$M(n,n,n) = O(n^{2.373})$ Stothers'10, Vassilevska W.'12

$M(n,n^{0.313},n) = O(n^{2+o(1)})$ Le Gall–Urrutia'17

$M(n,n^{0.529},n) = O(n^{2.058})$ Le Gall–Urrutia'17

Prelims on Matrix Multiplication

- given $n_1 \times n_2$ matrix A and $n_2 \times n_3$ matrix B with integer entries in $[\ell] := \{0, \ldots, \ell\}$
- let $M(n_1, n_2, n_3)$ be time to compute standard product

$$(AB)[i,j] \;=\; \sum_k A[i,k]B[k,j]$$

- let $M^*(n_1, n_2, n_3 \mid \ell)$ be time to compute (min,+) product

$$(A * B)[i,j] \;=\; \min_k \left(A[i,k] + B[k,j]\right)$$

- Trivial: $M^*(n_1, n_2, n_3 \mid \ell) \;=\; O(n_1 n_2 n_3)$
- Fact: $M^*(n_1, n_2, n_3 \mid \ell) \;=\; \widetilde{O}(\ell \cdot M(n_1, n_2, n_3))$
- Proof: can compute (max,+) product from standard product of $A'[i,k] = U^{A[i,k]}$ & $B'[k,j] = U^{B[k,j]}$ (large $\widetilde{O}(\ell)$-bit integers)

Why (min,+) Product is Useful to APSP

- Suppose we have computed all shortest paths of length $\leq \ell/2$
- To compute all shortest paths of length $\leq \ell$:

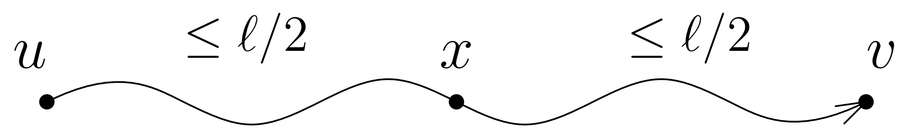

$$D_{\leq \ell}[u,v] \;=\; \min_{x \in V} \left(D_{\leq \ell/2}[u,x] + D_{\leq \ell/2}[x,v] \right)$$

$\Rightarrow$ time $O(M^*(n,n,n \mid \ell))$

- try all ℓ's that are powers of 2...
  but doesn't work well as ℓ gets large!

Hitting Set

- Observation: $\exists$ subset $H_\ell \subset V$ that hits all shortest paths of length $\geq \ell$, with $|H_\ell| = \widetilde{O}(n/\ell)$

- Proof: just take random $H_\ell \subset V$, with sampling prob. p
  $\Rightarrow$ for any fixed $u, v \in V$ with shortest path $\pi[u, v]$ of length $\geq \ell$,

$$\begin{aligned} \Pr[H_\ell \text{ not hit } \pi[u,v]] &\leq (1-p)^\ell \\ &\leq e^{-p\ell} \\ &\leq 1/n^c \qquad \text{by setting } p = (c/\ell)\ln n \end{aligned}$$

  (Alternate Proof: greedy hitting set)

First unwt-dir-APSP Alg'm

- short shortest paths of length $\le \ell_0$:

  $\Rightarrow$ total time $\widetilde{O}(M(n,n,n \mid \ell_0)) = \widetilde{O}(\ell_0 n^\omega)$

- long shortest paths of length $\ge \ell_0$:

  compute single-source shortest paths from each $x \in H_{\ell_0}$
  & single-sink shortest paths to each $x \in H_{\ell_0}$
  $\Rightarrow$ total time $O(|H_{\ell_0}| \cdot n^2) = O(n^3/\ell_0)$

- set $\ell_0 = n^{(3-\omega)/2} \Rightarrow \boxed{\widetilde{O}(n^{(\omega+3)/2})}$

Zwick's unwt-dir-APSP Alg'm ('98)

- Suppose we have computed all shortest paths of length $\leq 2\ell/3$
- To compute all shortest paths of length in $(2\ell/3, \ell]$:

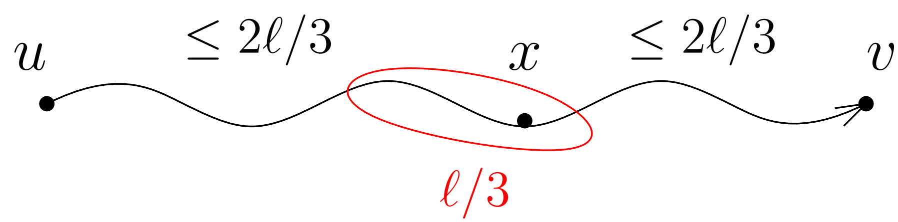

$$D_{\leq \ell}[u,v] = \min \begin{cases} D_{\leq 2\ell/3}[u,v] \\ \min_{x \in H_{\ell/3}} (D_{\leq 2\ell/3}[u,x] + D_{\leq 2\ell/3}[x,v]\} \end{cases}$$

$\Rightarrow$ time $O(M^*(n, |H_{\ell/3}|, n \mid \ell)) = \widetilde{O}(M^*(n, n/\ell, n \mid \ell))$

- try all ℓ's that are powers of $3/2$

$\Rightarrow$ total time $\boxed{\widetilde{O}\left(\max_{\ell} M^*(n, n/\ell, n \mid \ell)\right)}$

Zwick's unwt-dir-APSP Alg'm ('98)

$$\widetilde{O}\left(\max_{\ell} M^*(n, n/\ell, n \mid \ell)\right)$$
$$\leq \widetilde{O}\left(\max_{\ell} \min\{\ell \cdot M(n, n/\ell, n),\ \ n \cdot (n/\ell) \cdot n\}\right)$$
$$\leq \widetilde{O}\left(\ell_0 \cdot M(n, n/\ell_0, n) + n^3/\ell_0\right)$$

- set $\ell_0 = n^{0.471} \Rightarrow \boxed{O(n^{2.529})}$

can Zwick's alg'm be improved?

is there a more graph-theoretical approach?

New: Zwick's Alg'm is "Optimal"!

- Zwick's Alg'm: $\textsf{unwt-dir-APSP}(n) \le \widetilde{O}\left(\max_{\ell} M^*(n, n/\ell, n \mid \ell)\right)$

- Claim: $\max_{\ell} M^*(n, n/\ell, n \mid \ell) \le O(\textsf{unwt-dir-APSP}(n))$

New: Zwick's Alg'm is "Optimal"!

- Claim: $\max_{\ell} M^*(n, n/\ell, n \mid \ell) \;\leq\; O(\text{unwt-dir-APSP}(n))$
- Proof: (min,+) product of $n \times n/\ell$ matrix A and $n/\ell \times n$ matrix B with entries in $[\ell]$ reduces to APSP on this unweighted graph:

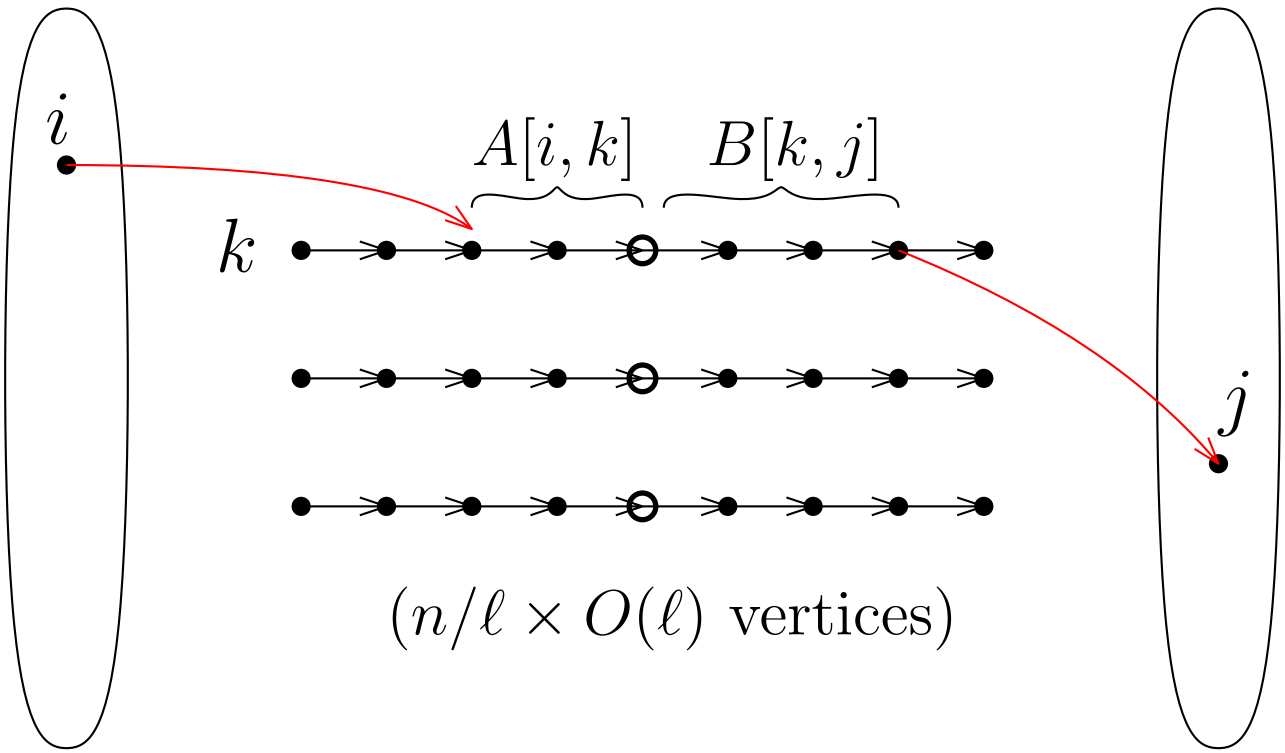

Consequences: Equivalences

1. dir-APSP with small $\widetilde{O}(1)$ integer wts is equally hard as unwt.
2. dir-APSP with small $\widetilde{O}(1)$ integer wts is equally hard with or without negative wts
3. unwt-dir-APSP is equally hard for DAGs as for general dir. graphs
4. unwt-APSP and unwt-APLP are equally hard for DAGs
5. for unwt-dir-APSP, computing distances is equally hard as paths
6. unwt-(≤ 2)-red-APSP is equally hard for undir. as for dir. graphs (related to all-pairs lightest shortest paths [Zwick'99])
7. approximate unwt-dir-APSP with $\widetilde{O}(1)$ additive error is equally hard as exact (related to [Roditty–Shapira'08])
8. unique unwt-dir-APSP is equally hard as unwt-dir-APSP (related to APSP counting)

Hardness from unwt-dir-APSP: Example

- Variants of matrix product for $n \times n$ matrices having alg'ms with "intermediate" complexity:
  - min-witness product ($\min\{k : A[i,k] \wedge B[k,j]\}$) in $O(n^{2.529})$
  - dominance product ($\bigwedge_k [A[i,k] < B[k,j]]$) in $O(n^{2.684})$ [Matoušek'91]
  - equality product ($\bigvee_k [A[i,k] = B[k,j]]$) in $O(n^{2.684})$
  - min-witness equality product ($\min\{k : A[i,k] = B[k,j]\}$) in $O(n^{2.688})$
  - (min,=) product ($\min\{A[i,k] : A[i,k] = B[k,j]\}$) in $O(n^{2.688})$
  - (min,max) product ($\min_k \max\{A[i,k], B[k,j]\}$) in $O(n^{2.688})$ [Duan–Pettie'09]

- Claim: unwt-dir-APSP reduces to min-witness equality product

Hardness from unwt-dir-APSP: Example

- Claim: unwt-dir-APSP reduces to min-witness equality product
- Proof: (min,+) product of $n \times n/\ell$ matrix A and $n/\ell \times n$ matrix B with entries in $[\ell]$ reduces to min-witness equality product of:

$$A'[i,(k,z)] = A[i,k] \quad \text{and} \quad B'[(k,z),j] = z - B[k,j]$$

where we order $(k,z) \in [n/\ell] \times [O(\ell)]$ by z

Hardness from unwt-dir-APSP

(min,max) product $\ll n^{2.5}$

$\longleftrightarrow$ AP-Bottleneck-Paths $\ll n^{2.5}$

min-witness equality product $\ll n^{2.5}$

new

[LPV'20]

unwt-dir-APSP $\ll n^{2.5}$

$\longleftrightarrow$ unwt-DAG-APSP $\ll n^{2.5}$

$\longleftrightarrow$ unwt-DAG-APLP $\ll n^{2.5}$

$\longleftrightarrow$ (≤ 2)-red-unwt-undir-APSP $\ll n^{2.5}$

$\vdots$

new

- Open: min-witness product? equality product?

This Talk

- hardness from APSP for unweighted graphs
  (or small $\widetilde{O}(1)$ integer weights)
  [C.–Vassilevska W.–Xu, ICALP'21]

- hardness from APSP for real-weighted graphs
  [C.–Vassilevska W.–Xu, STOC'22, to appear]

Hardness from ~~int-~~real-APSP

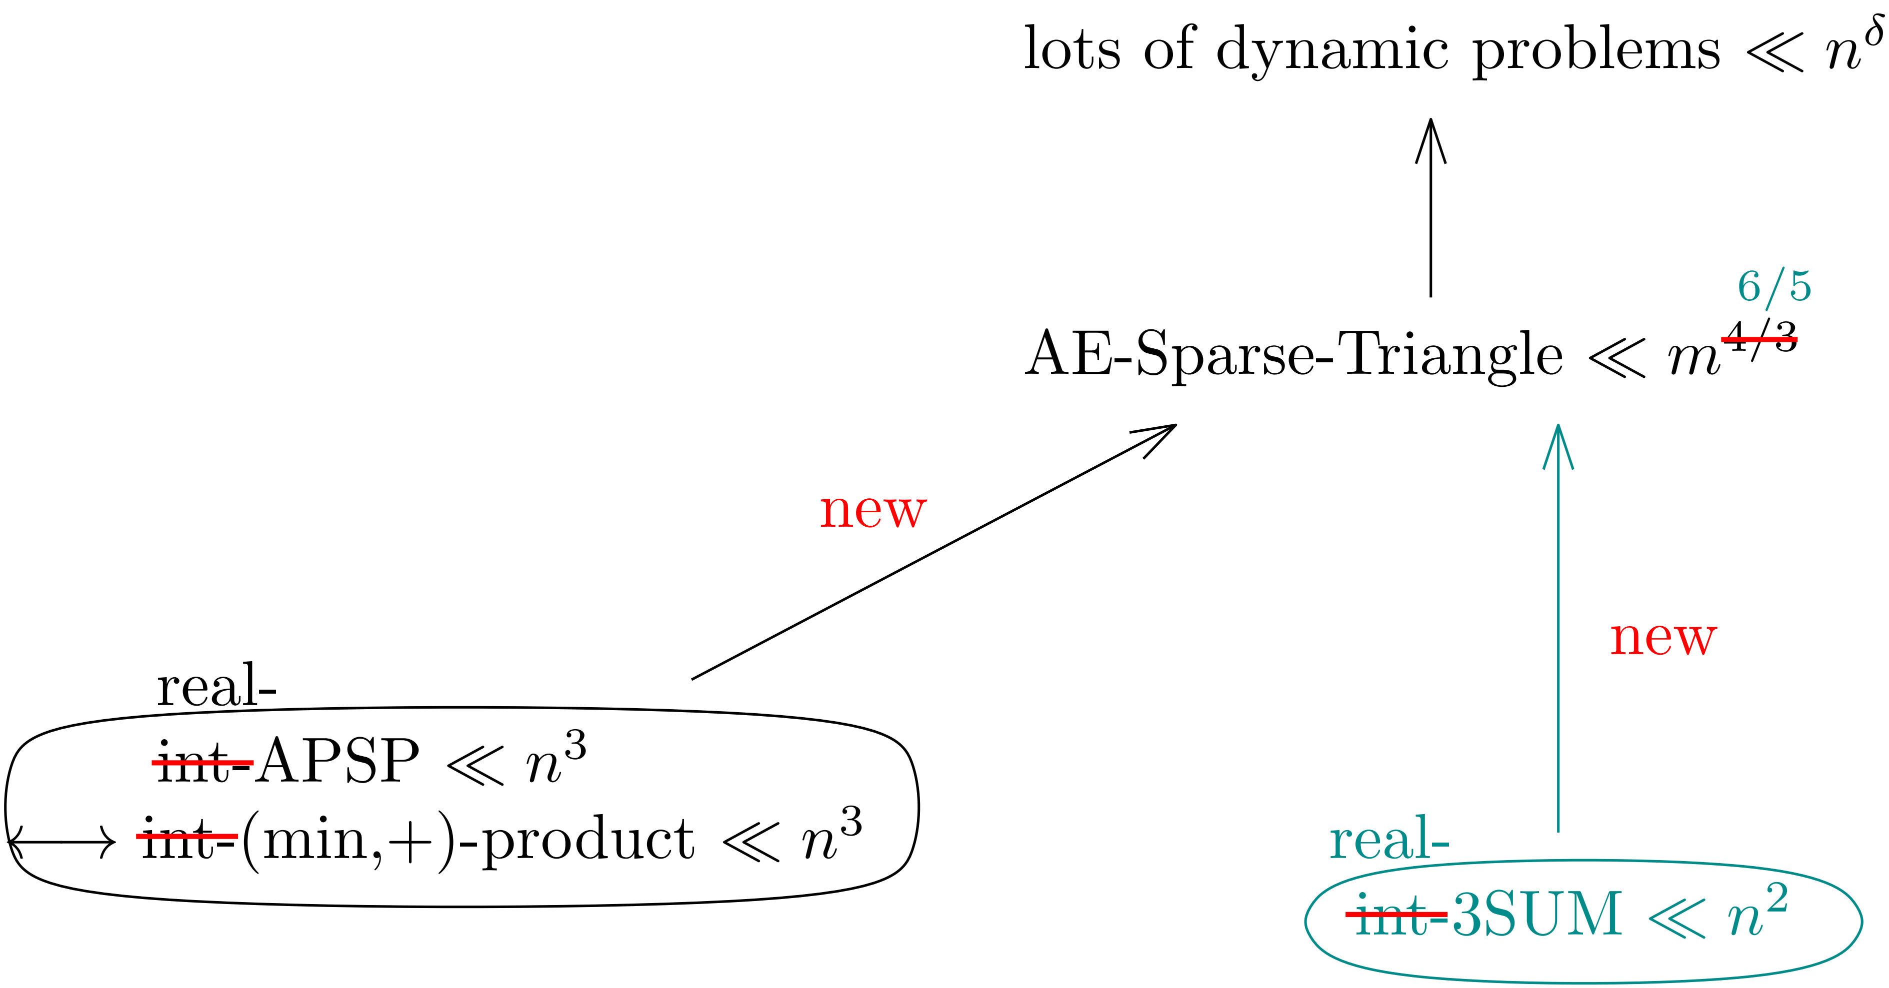

AE-Sparse-Triangle

- Given graph $G = (V, E)$ with m edges,
  for every edge $uv \in E$, decide $\exists$ triangle through uv

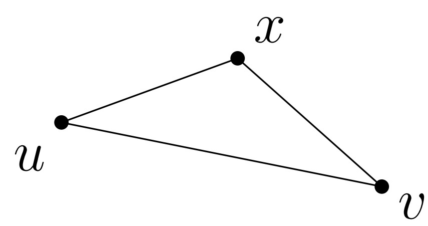

- Note: related to set disjointness queries ($N(u) \cap N(v) \neq \emptyset$?)
- Alon–Yuster–Zwick'94: $O(m^{2\omega/(\omega+1)}) \leq O(m^{1.408})$

Fredman's real-APSP "Alg'm" ('75)

- by known reduction,
  real-APSP$(n) \;\le\; \widetilde{O}(M^*(n,n,n)) \;\le\; \widetilde{O}(n/d \cdot M^*(n,d,n))$
- suffice to compute (min,+) product of $n \times d$ matrix A and $d \times n$ matrix B
- Fredman's trick: $A[i,k'] + B[k',j] < A[i,k] + B[k,j]$
  $$\iff A[i,k'] - A[i,k] < B[k,j] - B[k',j]$$
- sort all $A[i,k'] - A[i,k]$ and all $B[k,j] - B[k',j]$ in $\widetilde{O}(d^2 n)$ time
- afterwards, can compute $A * B$ without any more comparisons!
- total # comps $\widetilde{O}((n/d) \cdot (d^2 n + n^2))$
- set $d = \sqrt{n} \;\Rightarrow\;$ $\boxed{\widetilde{O}(n^{5/2})}$ comps (but runtime still n^3!)

New: Fredman as a Reduction!

- first "guess" answers $k_{ij} = \arg\min_k (A[i,k] + B[k,j])$
- for each $k \in [d]$, solve AE-Sparse-Triangle on this graph G_k:
  1. for each $i, j \in [n]$ with $k_{ij} = k$, create edge $u[i]v[j]$
  2. for each $i \in [n], k' \in [d]$ and dyadic interval I, create edge $u[i]x[k', I]$ if rank$(A[i,k'] - A[i,k])$ is in left half of I
  3. for each $j \in [n], k' \in [d]$ and dyadic interval I, create edge $x[k', I]v[j]$ if rank$(B[k,j] - B[k',j])$ is in right half of I

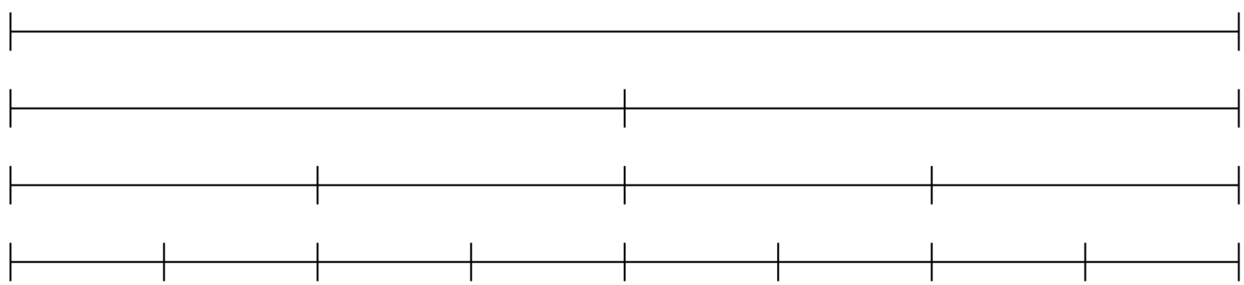

dyadic intervals

New: Fredman as a Reduction!

- first "guess" answers $k_{ij} = \arg\min_k (A[i,k] + B[k,j])$
- for each $k \in [d]$, solve AE-Sparse-Triangle on this graph G_k:
  1. for each $i, j \in [n]$ with $k_{ij} = k$, create edge $u[i]v[j]$
  2. for each $i \in [n], k' \in [d]$ and dyadic interval I, create edge $u[i]x[k', I]$ if $\text{rank}(A[i,k'] - A[i,k])$ is in left half of I
  3. for each $j \in [n], k' \in [d]$ and dyadic interval I, create edge $x[k', I]v[j]$ if $\text{rank}(B[k,j] - B[k',j])$ is in right half of I
- Observation: $\exists$ triangle through $u[i]v[j]$

$$\iff \exists k',\ A[i,k'] - A[i,k] < B[k,j] - B[k',j]$$

$$\iff \exists k',\ A[i,k'] + B[k',j] < A[i,k] + B[k,j]$$

$\iff k$ isn't correct answer for k_{ij}!

New: Fredman as a Reduction!

- first "guess" answers $k_{ij} = \arg\min_k (A[i,k] + B[k,j])$
- for each $k \in [d]$, solve AE-Sparse-Triangle on this graph G_k:
  1. for each $i, j \in [n]$ with $k_{ij} = k$, create edge $u[i]v[j]$
  2. for each $i \in [n], k' \in [d]$ and dyadic interval I, create edge $u[i]x[k', I]$ if $\text{rank}(A[i,k'] - A[i,k])$ is in left half of I
  3. for each $j \in [n], k' \in [d]$ and dyadic interval I, create edge $x[k', I]v[j]$ if $\text{rank}(B[k,j] - B[k',j])$ is in right half of I

- Analysis: G_k has $O(n^2/d + dn \log n)$ edges on average

  $\Rightarrow$ total time $\widetilde{O}((n/d) \cdot d \cdot \text{AE-Sparse-Triangle}(n^2/d + dn))$
- set $d = \sqrt{n}$ $\Rightarrow$ $\boxed{\widetilde{O}(n \cdot \text{AE-Sparse-Triangle}(n^{3/2}))}$

New: Fredman as a Reduction!

- Theorem: real-APSP$(n) \ \leq\ \widetilde{O}(n \cdot \text{AE-Sparse-Triangle}(n^{3/2}))$
- Corollary: if AE-Sparse-Triangle$(m) \leq O(m^{4/3-\varepsilon})$, then real-APSP$(n) \leq O(n \cdot (n^{3/2})^{4/3-\varepsilon}) \leq O(n^{3-\varepsilon'})$

Conclusions

- Moral: reinterpret known alg'ms as reductions!
- Many open questions:
  - relationship between int-APSP, unwt-dir-APSP, & real-APSP hypotheses?
  - better understanding of problems with intermediate complexity between n^2 and n^3...
  - counting variant of APSP in $\widetilde{O}(n^3)$ time for weighted graphs?
    (note: counts may be large $\widetilde{O}(n)$-bit numbers!)

THE END