DISTANCE-PRESERVING GRAPH CONTRACTIONS

A. Bernstein[1], K. Däubel[1], Y. Disser[2], M. Klimm[3], T. Mütze[1] and F. Smolny[1]
[1]TU Berlin, [2]TU Darmstadt, [3]HU Berlin

Setting

Given: A simple, undirected graph $G = (V, E)$ with edge lengths $\ell : E \to \mathbb{R}_{>0}$ and error parameters $\alpha \geq 1$ (multiplicative), $\beta \geq 0$ (additive).

Operation: Edge contractions, i.e given a set $C \subseteq E$ we obtain the graph G/C by:

1. contracting the edges in C and
2. deleting parallel edges and loops ($\Delta(C)$:= set of deleted edges).

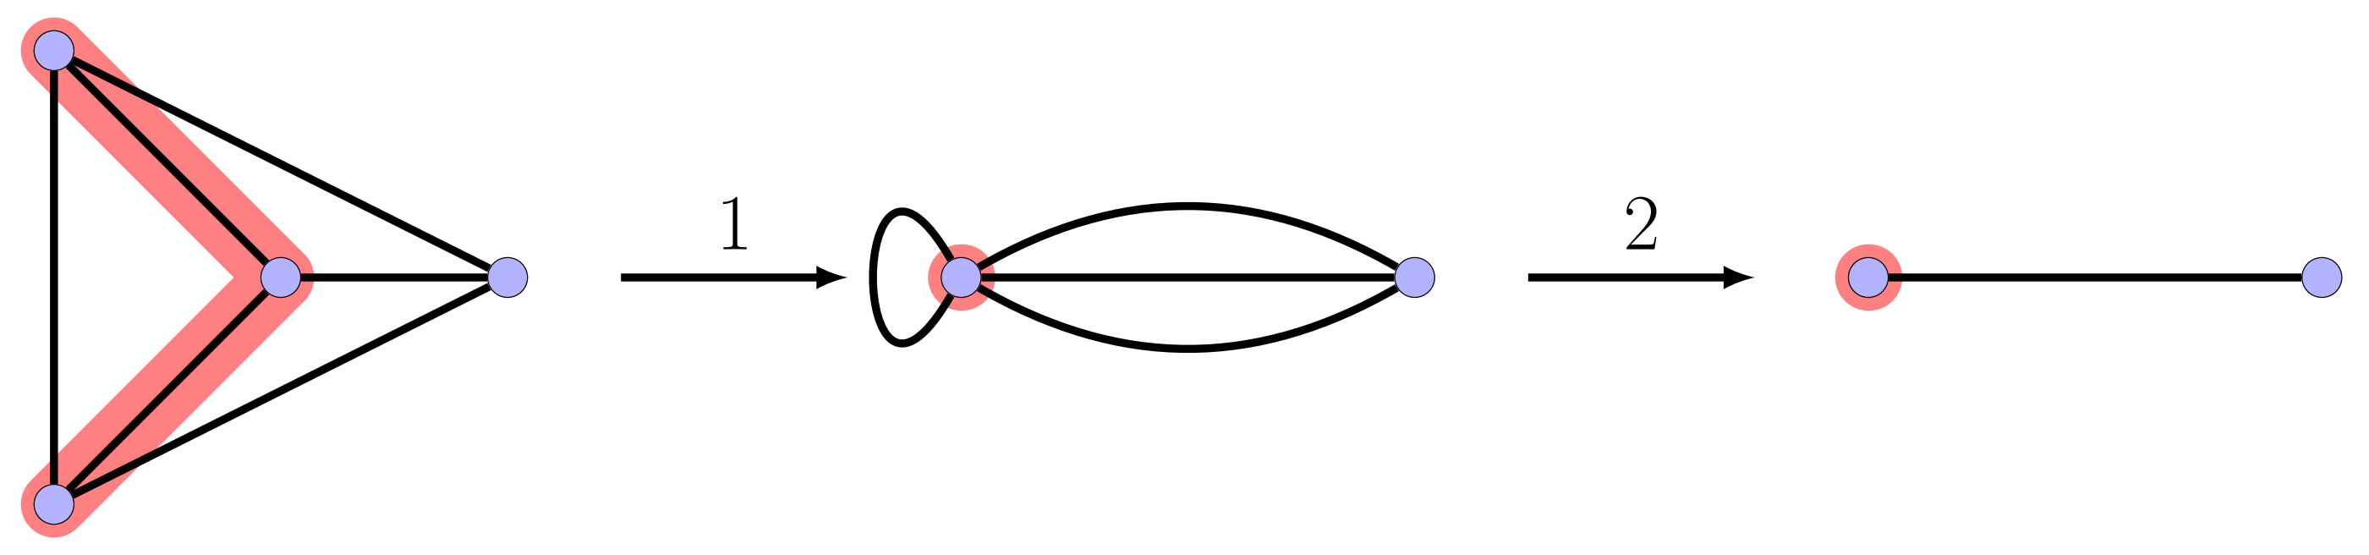

The number of vertices and edges in G/C coincides with

- the number of conn. comp. of (V, C) and
- $m(G) - |C| - |\Delta(C)|$, respectively.

Distance preservation: We don't want distances in G/C to decrease too much, thus we call $C \subseteq E$ an (α, β)-contraction if

$$\boxed{\mathrm{dist}_{G/C}(u, v) \geq \mathrm{dist}_G(u, v)/\alpha - \beta.}$$

Goal: Find an (α, β)-contraction C with the maximum number $|C \cup \Delta(C)|$ of lost edges.

A different view and comparison to spanners

Contractions:	**Spanners:**
Contraction of edges in C:	Deletion of edges in $E \setminus E'$:
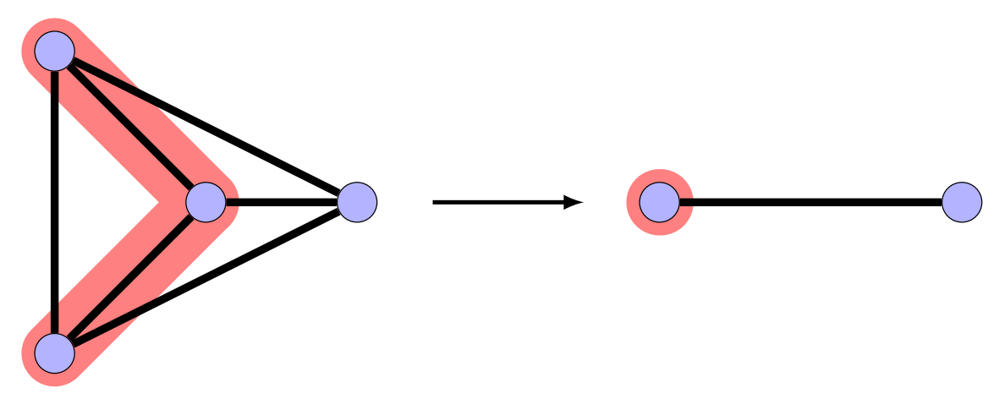	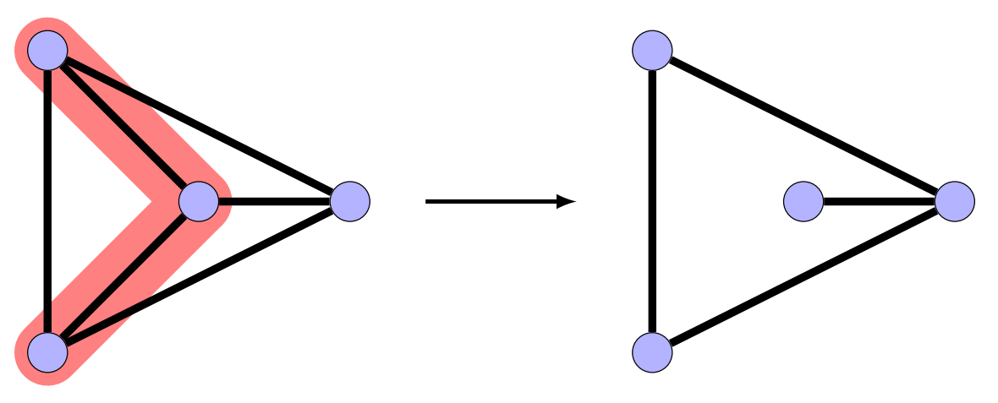
Equivalent to setting edge lengths in C to 0:	Equivalent to setting edge lengths in $E \setminus E'$ to ∞:
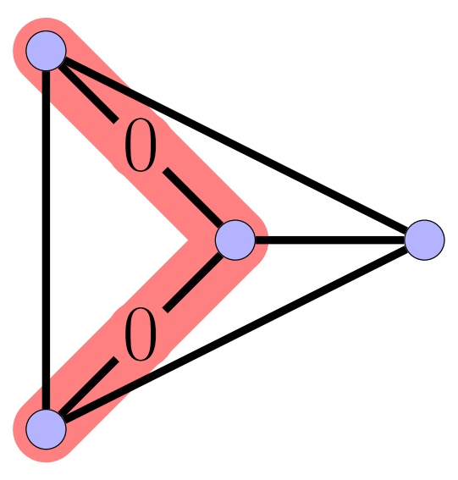 0 0	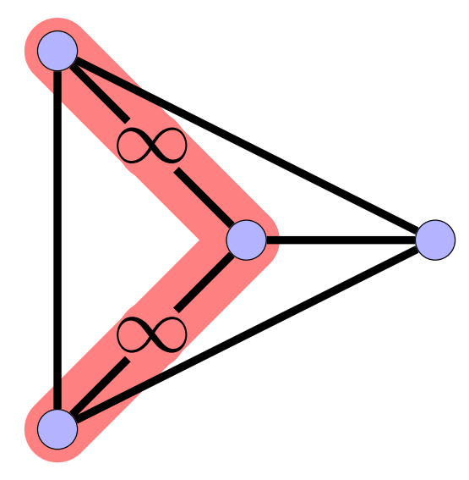 ∞ ∞
$\mathrm{dist}_0(u, v) \geq \mathrm{dist}_G(u, v)/\alpha - \beta$	$\mathrm{dist}_\infty(u, v) \leq \alpha\, \mathrm{dist}_G(u, v) + \beta$

Reminder: Spanner
A subgraph $H = (V, E'), E' \subseteq E$ is an (α, β)-spanner of G if $\boxed{\mathrm{dist}_H(u, v) \leq \alpha\, \mathrm{dist}_G(u, v) + \beta}$.

Dynamic Programming on Trees

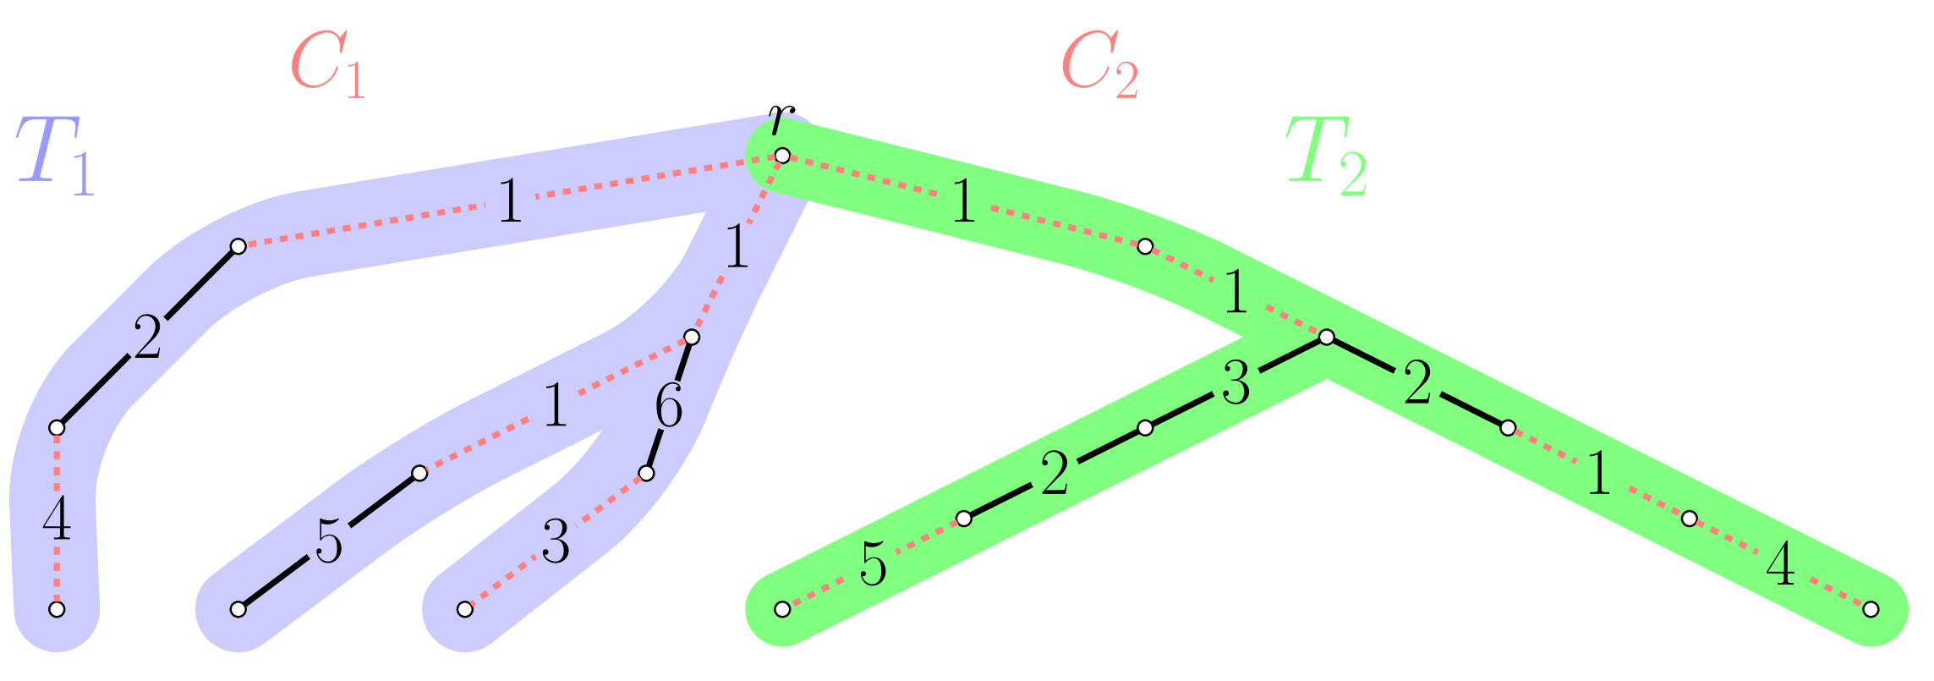

Split T at the root r.
If C_1, C_2 are feasible, when is $C := C_1 \cup C_2$?

Key definition: $\underline{\mathrm{load}(P) := \ell(P)/\alpha - \ell_{G/C}(P)}$ for any path P

Properties:

- Feasibility $\Leftrightarrow$ $\mathrm{load}(P) \leq \beta$ for all P
- Additivity: $\mathrm{load}(P_1) + \mathrm{load}(P_2) = \mathrm{load}(P_1 + P_2)$

Define $\mathrm{load}(T, v) := \max\{\mathrm{load}(P) : P \text{ ends in } v\}$
Lemma: C feasible $\Leftrightarrow$ C_1, C_2 feasible and $\mathrm{load}(T_1, r) + \mathrm{load}(T_2, r) \leq \beta$

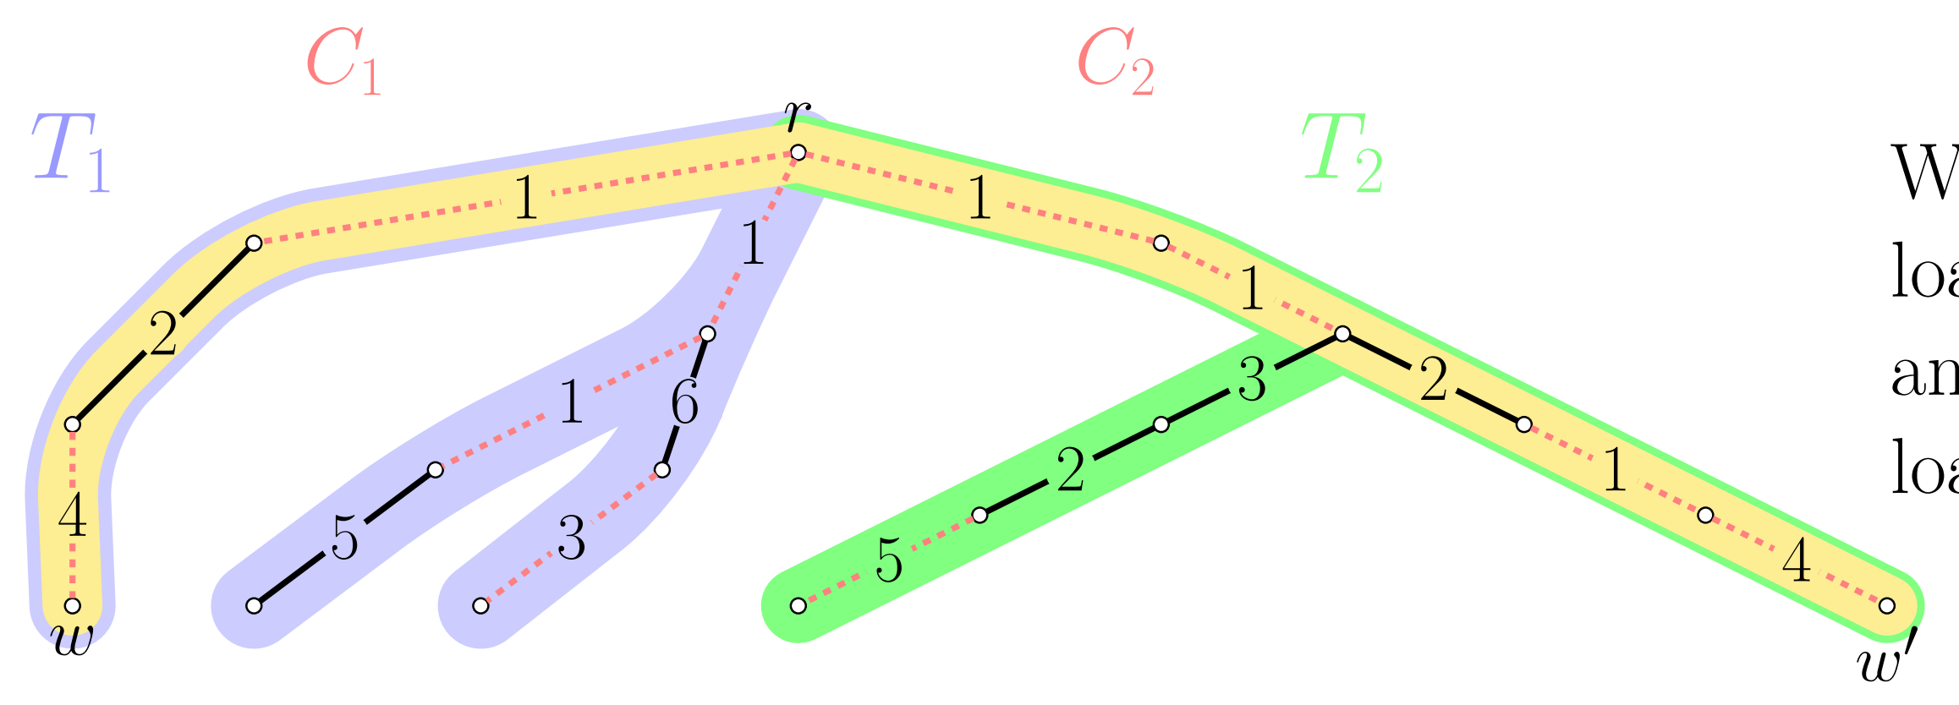

With $\alpha = 2$,
$\mathrm{load}(T_1, r) = \mathrm{load}(w, r) = 1.5$
and
$\mathrm{load}(T_2, r) = \mathrm{load}(w', r) = 2.5.$

Lemma $\Rightarrow$ C feasible if $\beta \geq 4$
Bottom-up dynamic programming finds solutions of all possible sizes minimizing their respective loads, in $\mathcal{O}(n^3)$ time.

Examples of Iterative Contractions

Two iterations with $\alpha = 5/4$ and $\beta = 3$ on a tree:

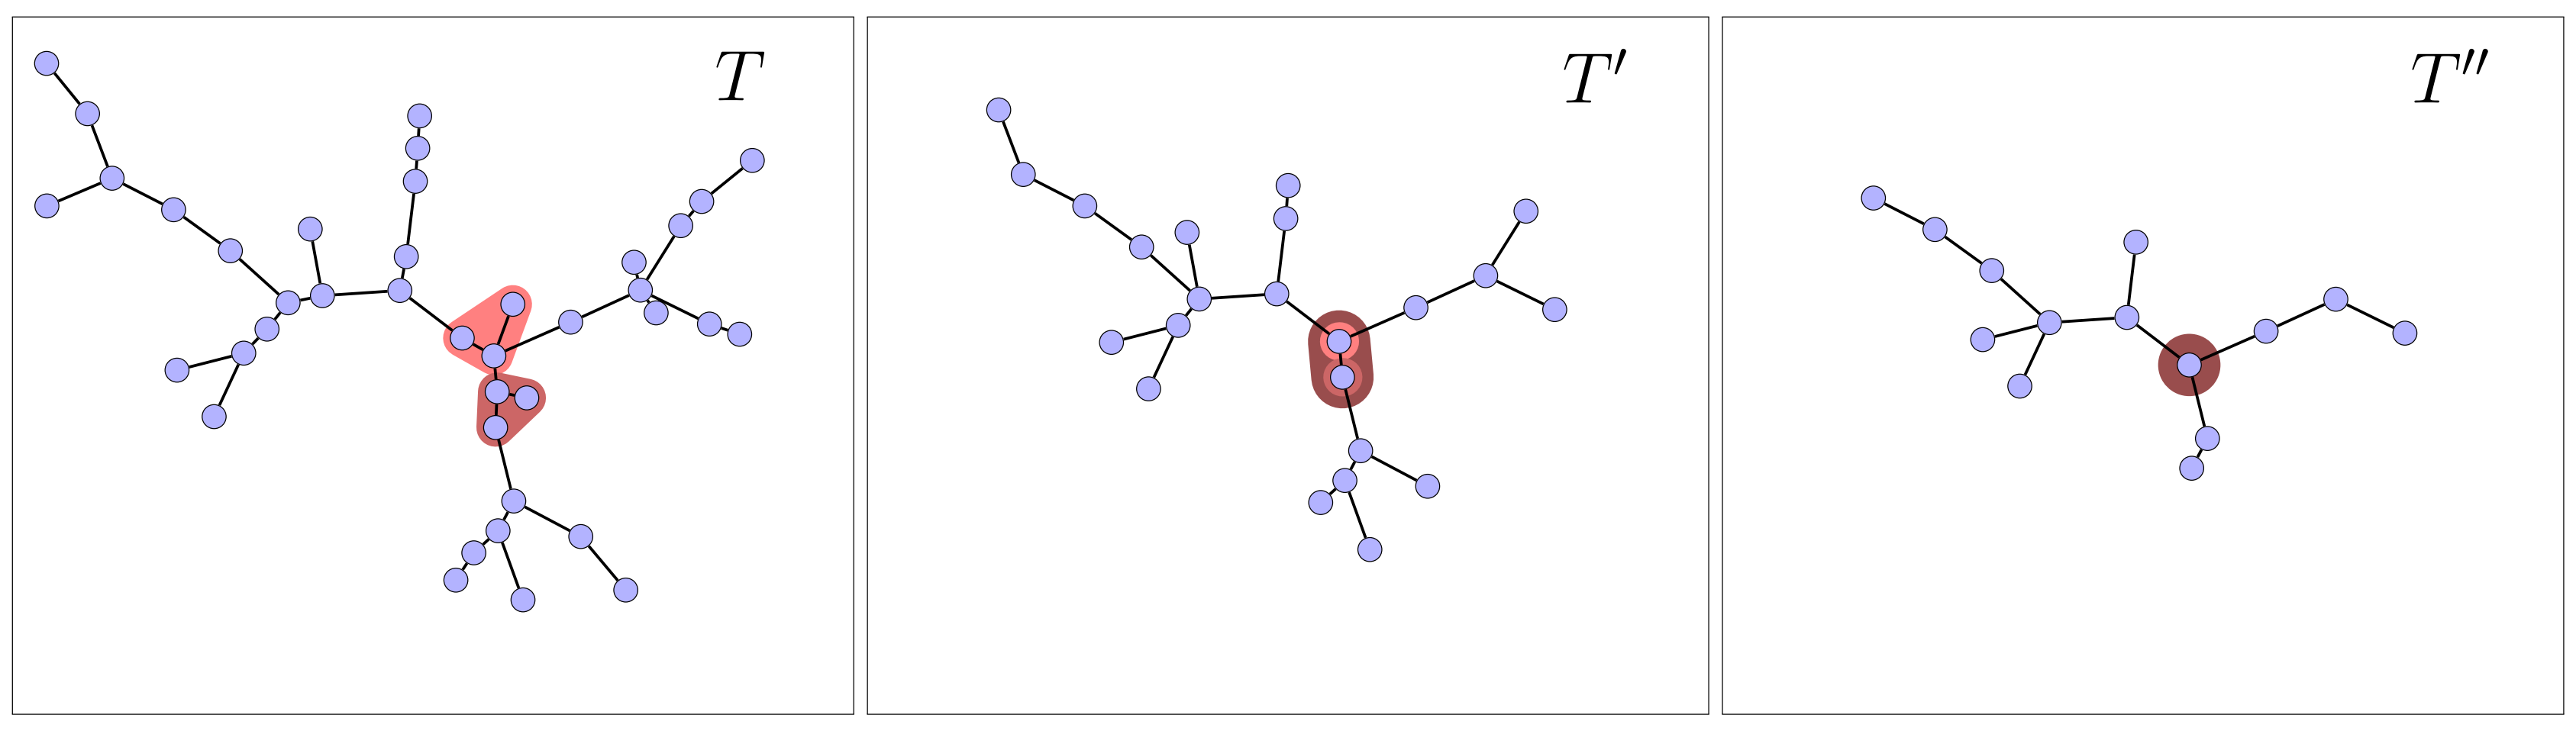

Two iterations with $\alpha = 4/3$ and $\beta = 3$ on a planar graph:

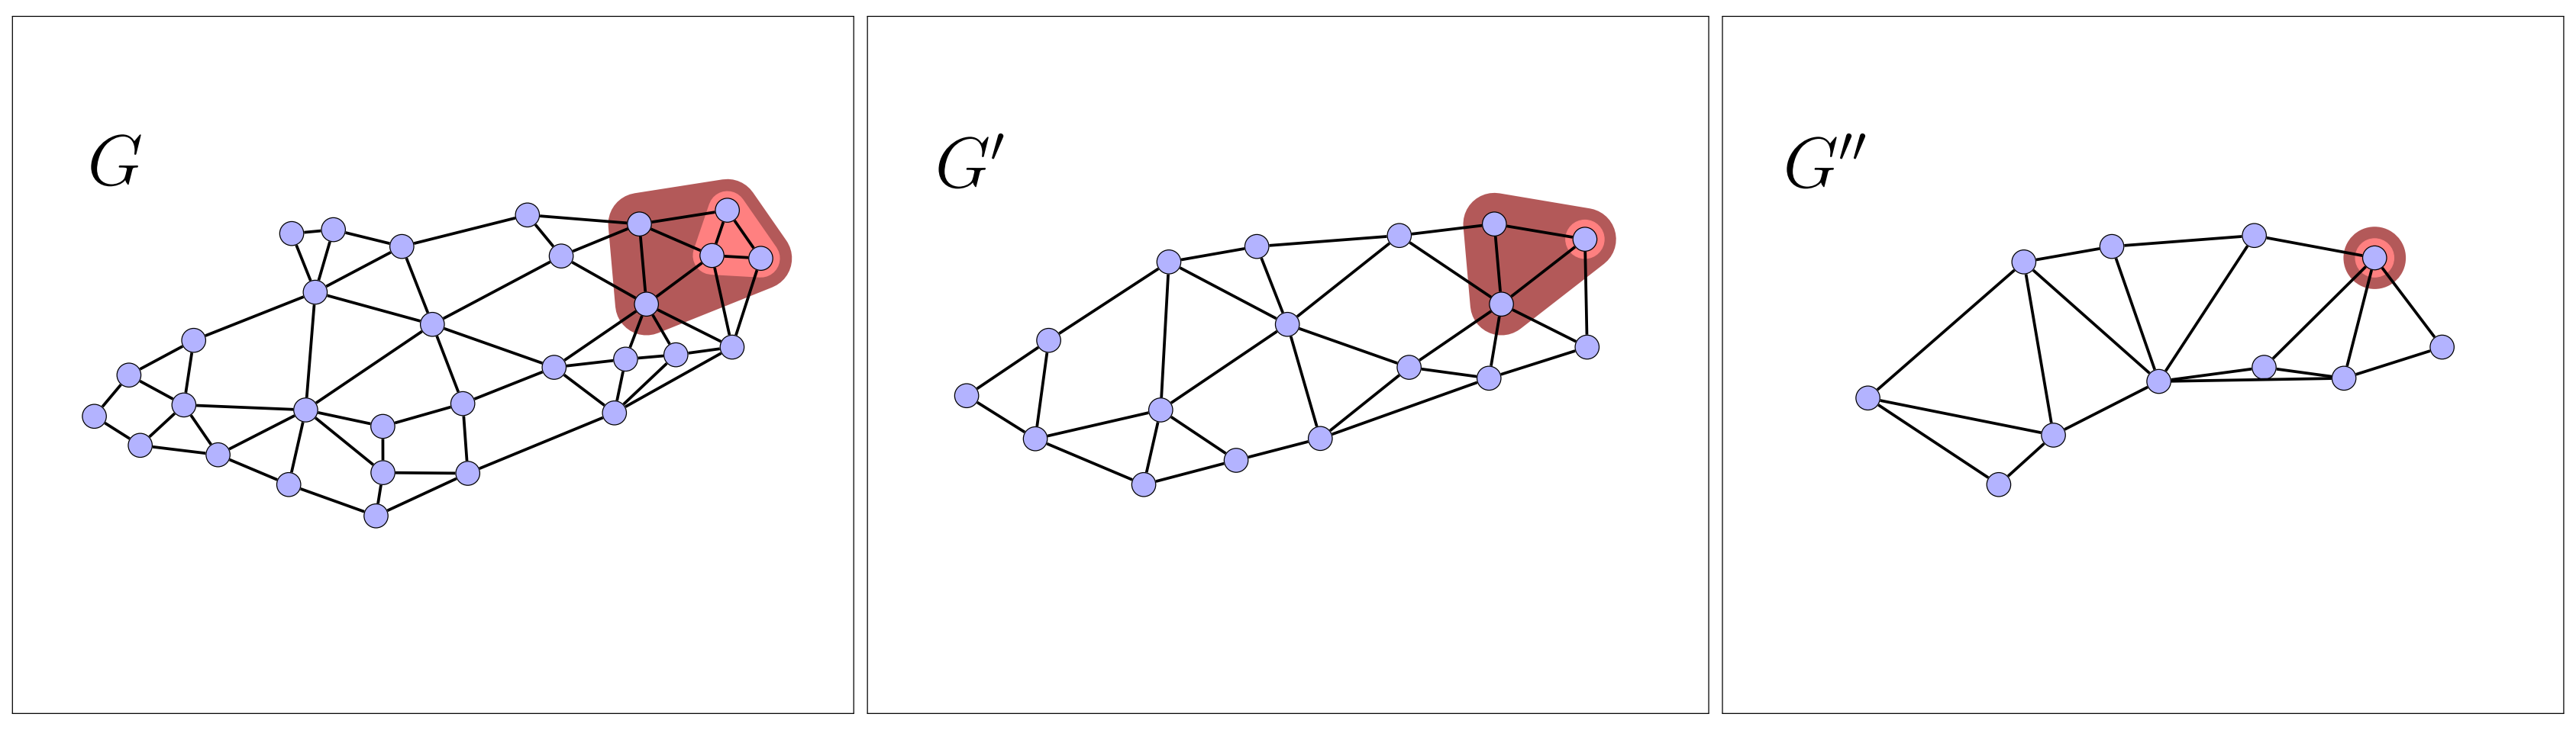

Distances are geometric and some contracted sets of vertices are highlighted.

Results

Algorithms and hardness

	Graph classes			
	Path	Tree	Cycle	General
addit. (α=1), unit lg.	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$\mathcal{O}(n)$	$m^{\frac{1}{2}-\varepsilon}$-inapx.
affine (α, β), unit lg.				
addit. (α=1)	$\mathcal{O}(n^3)$		NP-hard	$n^{1-\varepsilon}$-inapx
affine (α, β)				

Greedy Algorithms

Even for bipartite graphs and $\beta = 1$

Dynamic Programming (see below)

Asymptotic bounds with unit lengths

	# of edges in G/C	Time
$(\alpha, \beta) = (2k - 1, 1)$	$n^{1+1/k}$	$\mathcal{O}(m)$
$(\alpha, \beta) = (2 \log_2 n - 1, 1)$	$2n$	$\mathcal{O}(m)$
$(\alpha, \beta) = (k - 1, 1)$	$\Omega(n^{1+1/k})$	—
$(\alpha, \beta) = (1, k)$	$m - km/(2n)$	$\mathcal{O}(m)$
$(\alpha, \beta) = (1, k)$	$\mathcal{O}(n^2/k)$	$\mathcal{O}(m)$
$(\alpha, \beta) = (1, \mathcal{O}(1))$	$\Omega(n^{4/3-o(1)})$ [Abboud, Bodwin '16]	
min. degree D	# of vertices in G/C	Time
$(\alpha, \beta) = (5, 1)$	n/D	$\mathcal{O}(m)$
$(\alpha, \beta) = (k, 1)$	$\Omega(n/(kD))$	—

k-Partition (see below)

conditional on girth conjecture

Asymptotic Bounds

Small (α, β)-contractions in unweighted graphs can be constructed similary to spanners:

k-Partition

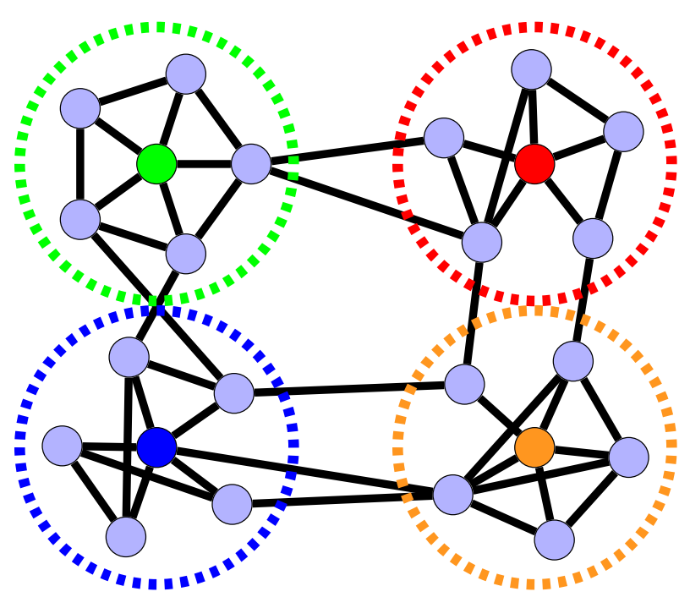

Example of $k = 2$-partition P with:

- $P = \{(●, P_1), (●, P_2), (●, P_3), (●, P_4)\}$
- vertices have distance at most $1 = k - 1$ to their respective centers
- 'few' intra cluster edges [Awerbuch '85]

Contractions:		**Spanners:**	
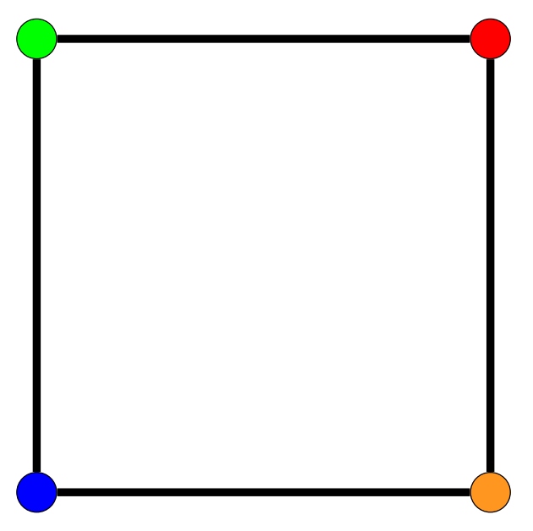	$(2k - 1, 1)$-contraction with at most $n^{1+1/k}$ edges	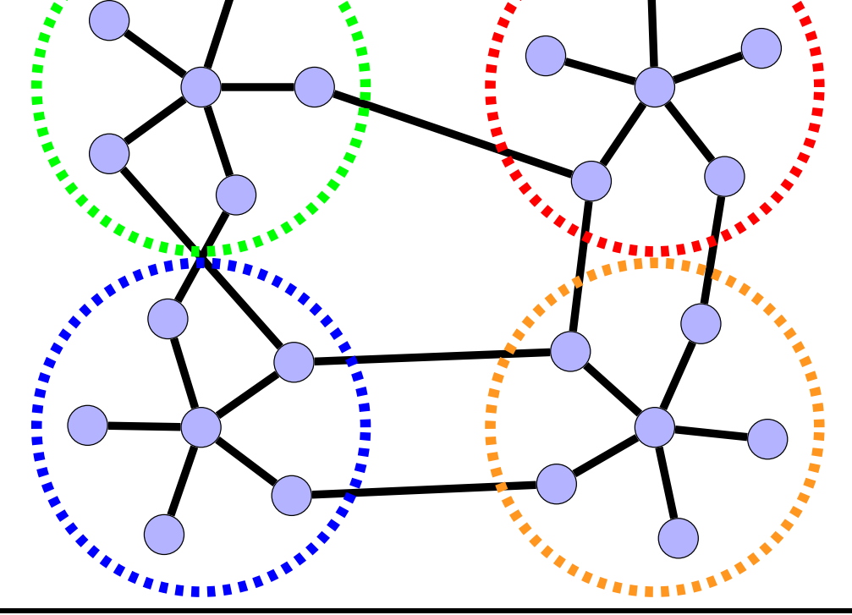	$(2k - 1, 0)$-spanner with at most $\mathcal{O}(n^{1+1/k})$ edges [Baswana et al. '12]

LaTeX TikZposter