

is backboneed.

Consider a problem Π , and let the optimal solution to the given instance on G with edge costs $c(\cdot)$ be a subgraph $H \subseteq G$. Then, from the low-stretch property of the random embedding, the expected cost of H under the cost function \hat{c}_T is $\mathbb{E}_T[\sum_{e \in H} \hat{c}_T(e)] \leq \tilde{O}(\log n) \cdot \sum_{e \in H} c(e)$. Therefore the expected cost of any optimal solution under edge costs \hat{c}_T is at most $\tilde{O}(\log n) \cdot \sum_{e \in H} c(e)$. Consequently, any β -approximation algorithm for the problem Π on backboneed graphs would return a subgraph $H' \subseteq G$, with expected cost (with respect to \hat{c}_T) at most $(\beta \times \tilde{O}(\log n)) \cdot \sum_{e \in H} c(e)$. Since $\hat{c}(e) \geq c(e)$ for any edge $e \in E(G)$, the expected cost of the subgraph H' with respect to edge costs $c(\cdot)$ is also at most $(\beta \times \tilde{O}(\log n)) \cdot \sum_{e \in H} c(e) \leq (\beta \times \tilde{O}(\log n))c(H)$. Therefore, the solution H' is a randomized $\beta \times \tilde{O}(\log n)$ -approximate solution on the original edge costs $c(\cdot)$. ■

B Proofs from Section 4

B.1 2-CFL on Non-Metric Instances We now consider instances where the connection cost for the clients is given by some distance function $d(\cdot, \cdot)$ which may itself not satisfy triangle inequality, and the edge costs for building the 2-connected core is $c(\cdot)$.

We show how we can get poly-logarithmic approximations for the above “non-metric” 2-CFL problem using essentially the same techniques we used for 2-ECGS. We first guess one facility which the optimal solution opens and call it r . The LP is almost identical to the one given for 2-CFL on general graphs, except for the client-facility connection cost being some arbitrary function $d(\cdot, \cdot)$ instead of the tree distances $c(\cdot)$. Here is a brief overview of the rounding algorithm for 2-CFL. We skip the details of the proofs as they are very similar to the ones given in the earlier sections.

- (i) Solve the LP relaxation optimally. Then filter the client connection costs: If we let $D_u^* = \sum_{v \in V} d(u, v)z_{uv}^*$, it must be that $\sum_{v \in V \mid d(u, v) \leq 2D_i^*} z_{uv}^* \geq \frac{1}{2}$. Set $z_{uv}^* \leftarrow 0$ if $d(u, v) > 2D_i^*$ and scale the solution by factor 2.
- (ii) For each $u \in \mathcal{D}$, create a group $g_u = \{v \in V \mid z_{uv}^* \neq 0\}$ of facilities associated with this client. It is easy to check that the solution (x^*, y^*) is a feasible solution for the 2-ECGS LP with these groups.
- (iii) Perform Stage I and Stage II of the 2-ECGS algorithm once; if a group g_u is 2-connected to the root, open a facility at the representative vertex v_{g_i} . Because the 2-ECGS algorithm ensures that $\Omega(\frac{1}{\log n})$ groups are 2-connected to the root, and we open facilities for these groups, we know that $\Omega(\frac{1}{\log n})$ clients have a facility opened near them. A similar

analysis as the one for the 2-ECGS problem can be used to see that the total cost spent in this step is at most $O(\log n)\text{LPOpt}$.

- (iv) We can then repeat this process $O(\log^2 n)$ times and output the union of all previous partial solutions to guarantee with high probability a feasible solution to the 2-CFL problem.

THEOREM B.1. *Non-metric 2-CFL admits an $O(\log^3 n)$ approximation algorithm on backboneed graphs, and an $O(\log^4 n)$ approximation algorithm for general graphs.*