Integrating Facility Location and Network Design

Amitabh Sinha

(Joint work with R. Ravi)

GSIA, Carnegie Mellon University

Motivation: UFL with cable capacities

- Facility Location
- Input: Set of clients & facilities (with opening costs) in a metric space.

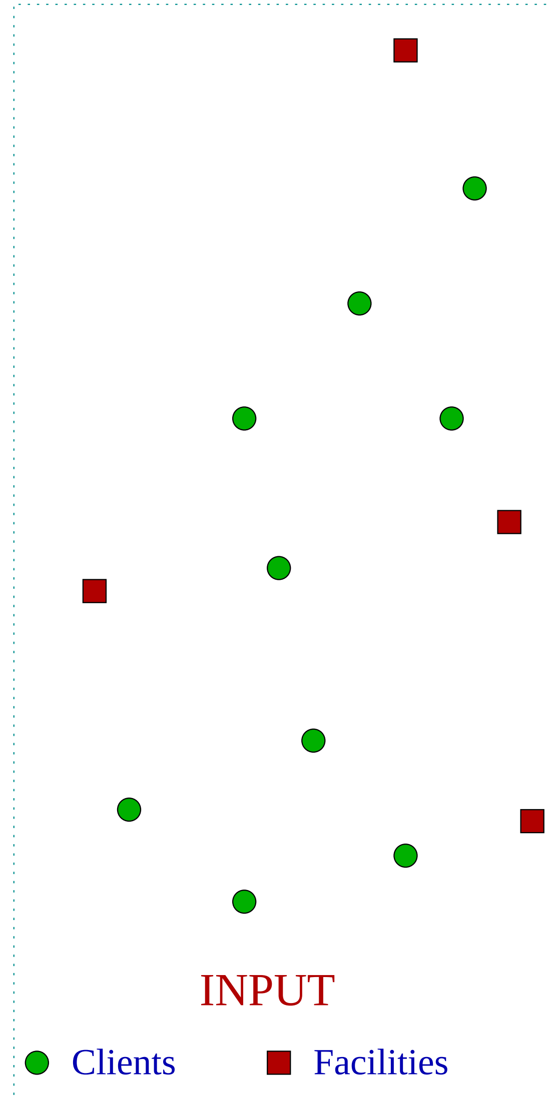

Motivation: UFL with cable capacities

- Facility Location

- Input: Set of clients & facilities (with opening costs) in a metric space.
- Objective: Open some facilities to serve clients. Client service cost: distance to nearest open facility. Minimize total cost.

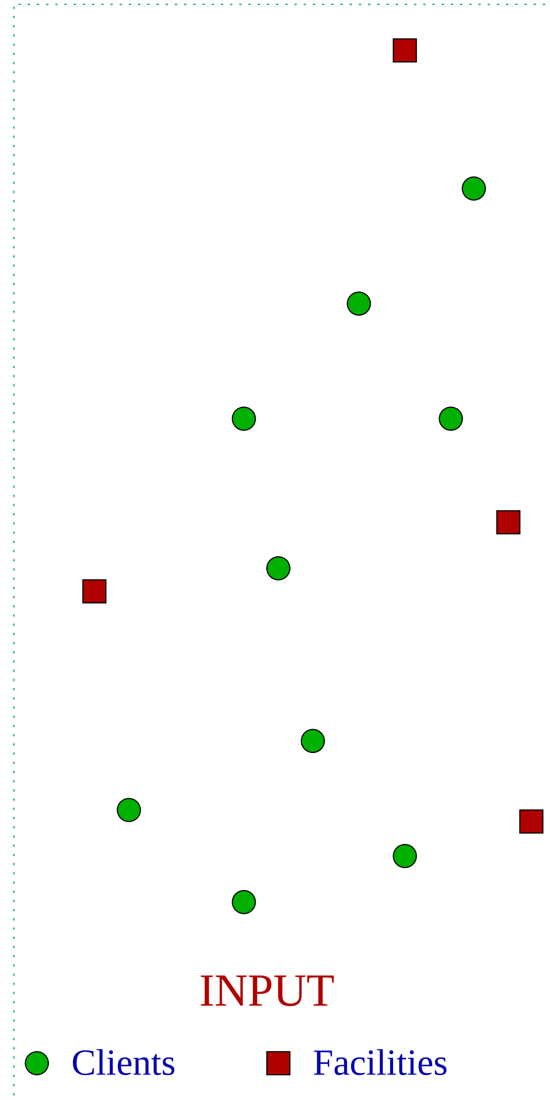

Motivation: UFL with cable capacities

- Facility Location
- Input: Set of clients & facilities (with opening costs) in a metric space.
- Objective: Open some facilities to serve clients. Client service cost: distance to nearest open facility. Minimize total cost.

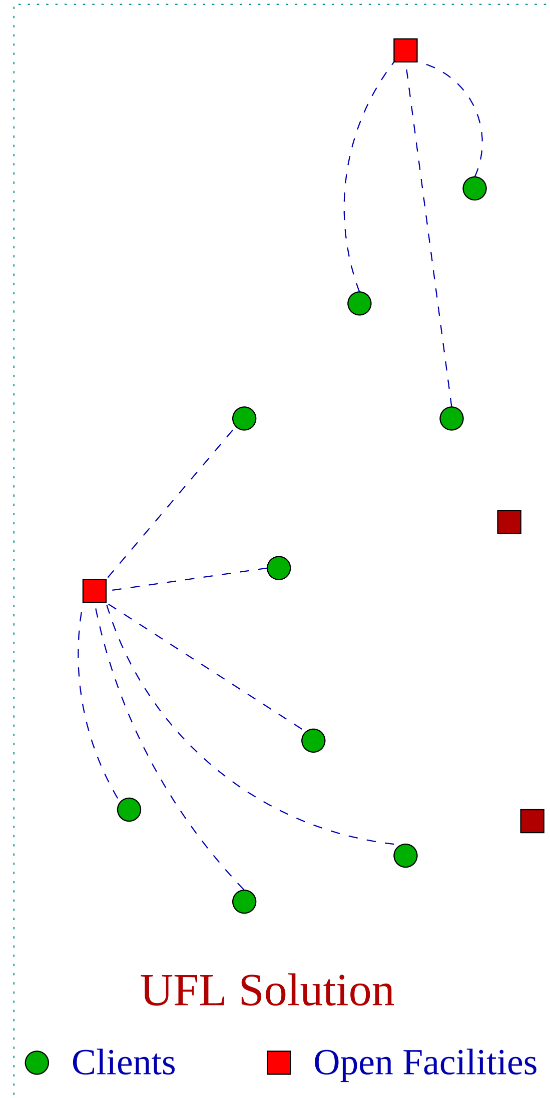

Motivation: UFL with cable capacities

- Facility Location with cable capacities
- Input: Set of clients & facilities (with opening costs) in a metric space.
- Objective: Open some facilities to serve clients. Client service cost: distance to nearest open facility. Minimize total cost.
- New twist: Clients connect to facilities via (capacitated) cables. Service cost becomes more complicated.

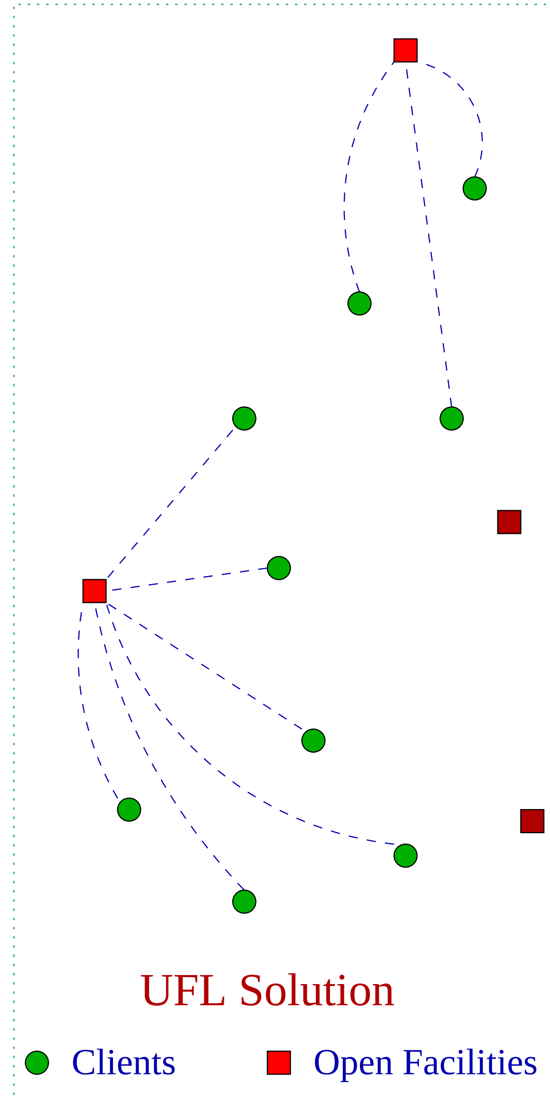

Motivation: UFL with cable capacities

- Facility Location with cable capacities
- Input: Set of clients & facilities (with opening costs) in a metric space.
- Objective: Open some facilities to serve clients. Client service cost: distance to nearest open facility. Minimize total cost.
- New twist: Clients connect to facilities via (capacitated) cables. Service cost becomes more complicated.

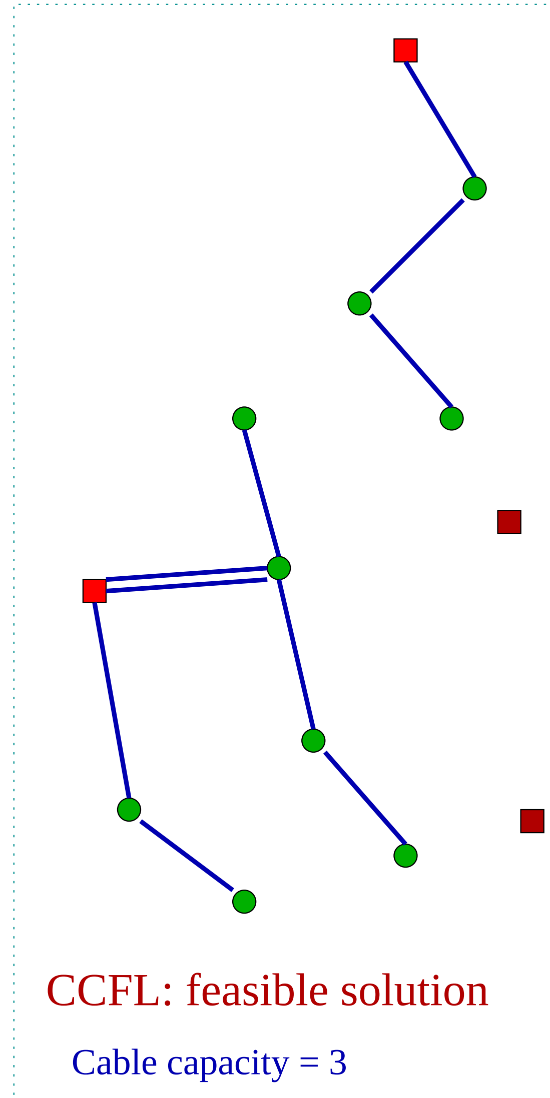

Outline

- Define CCFL: Capacitated cable facility location.
- Lower bounds for CCFL.
- Approximation algorithm for CCFL.
- Define KCFL: k-cable facility location.
- Thoughts on approximating KCFL.

Problem definition

- Capacitated Cable Facility Location (CCFL):
- Graph (metric), Edge weights c_e, Clients $D \subseteq V$, Facilities $\mathcal{F}$ with costs ϕ_j, and Cable capacity u.

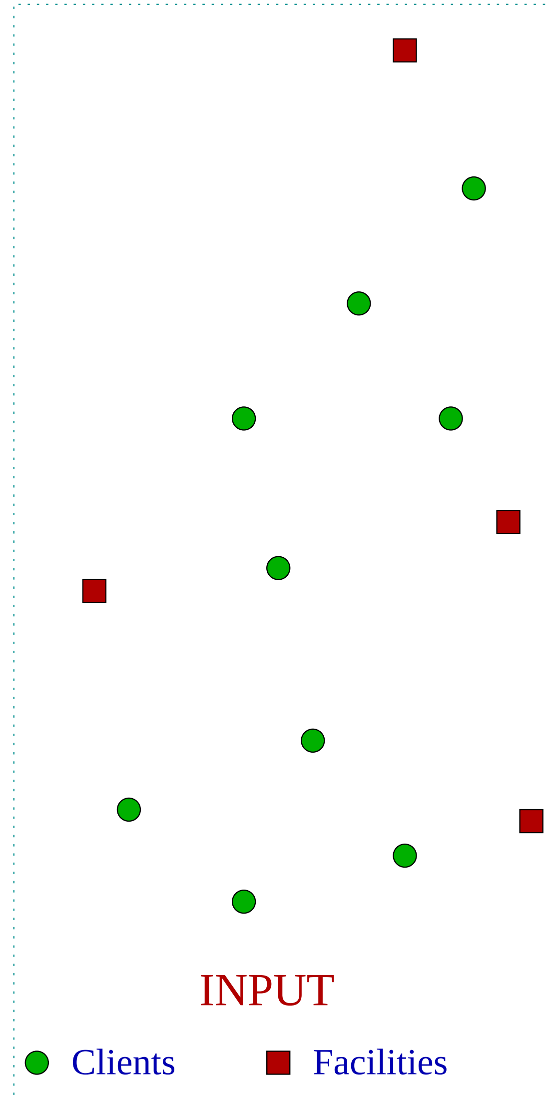

Problem definition

- Capacitated Cable Facility Location (CCFL):
- Graph (metric), Edge weights c_e, Clients $D \subseteq V$, Facilities $\mathcal{F}$ with costs ϕ_j, and Cable capacity u.
- Goal: Open some facilities, and install cables on edges, to support 1 unit of flow from each client to some open facility.

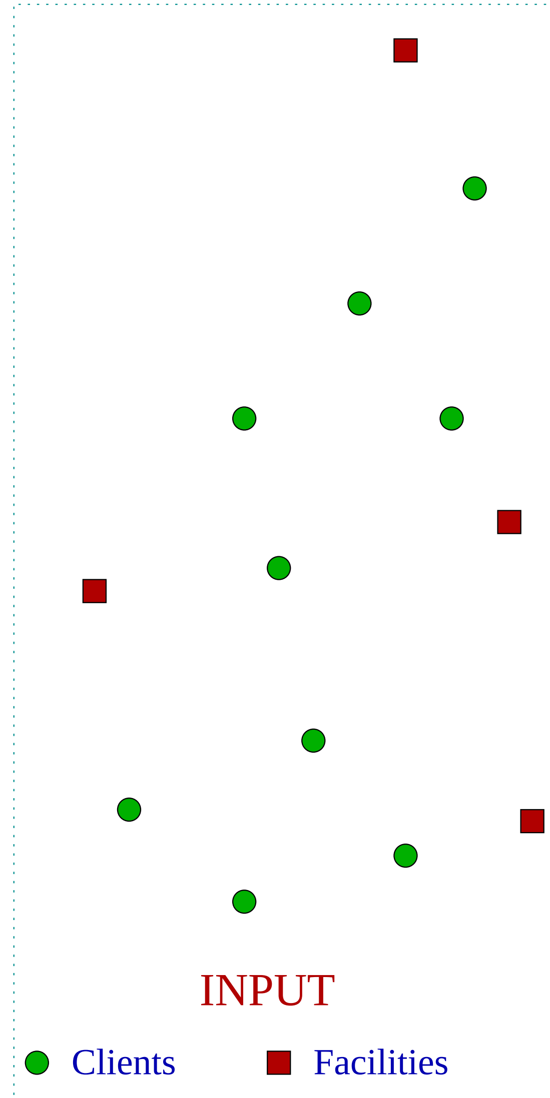

Problem definition

- Capacitated Cable Facility Location (CCFL):
- Graph (metric), Edge weights c_e, Clients $D \subseteq V$, Facilities $\mathcal{F}$ with costs ϕ_j, and Cable capacity u.
- Goal: Open some facilities, and install cables on edges, to support 1 unit of flow from each client to some open facility.
- Objective: Minimize total cost (facilities + cables).

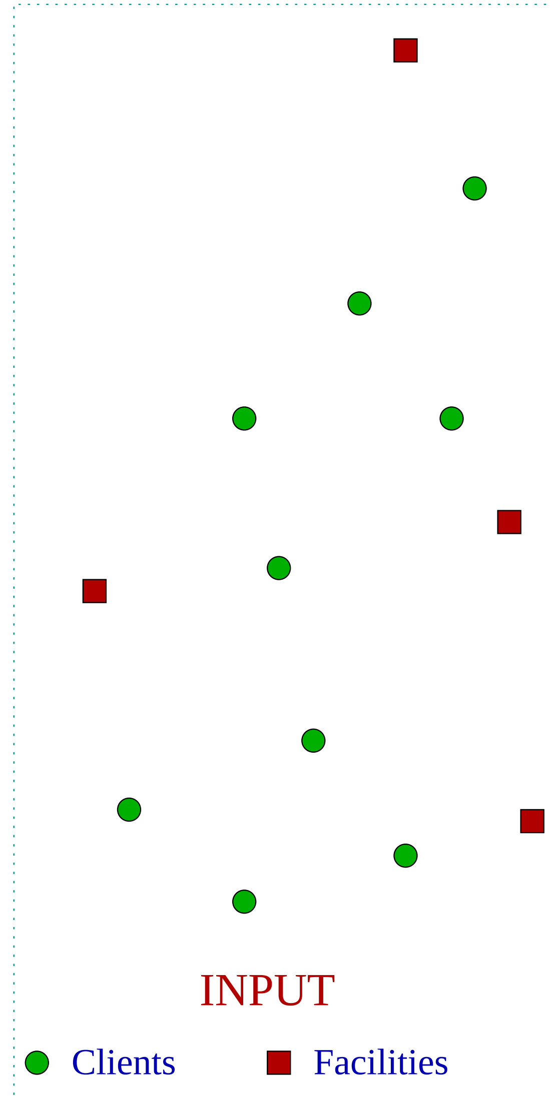

Problem definition

- Capacitated Cable Facility Location (CCFL):
- Graph (metric), Edge weights c_e, Clients $D \subseteq V$, Facilities $\mathcal{F}$ with costs ϕ_j, and Cable capacity u.
- Goal: Open some facilities, and install cables on edges, to support 1 unit of flow from each client to some open facility.
- Objective: Minimize total cost (facilities + cables).

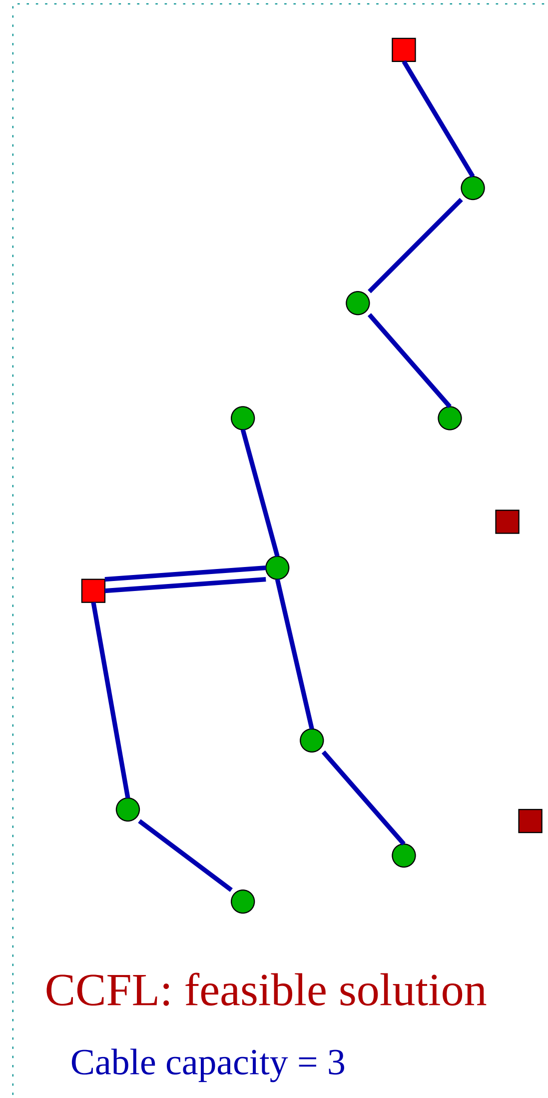

Special cases and past work

- $u = 1$: UFL; $\rho_{UFL} = 1.52$ [MYZ 02]. Others: [STA 97, JV 99, AGKMMP 01, JMS 02].

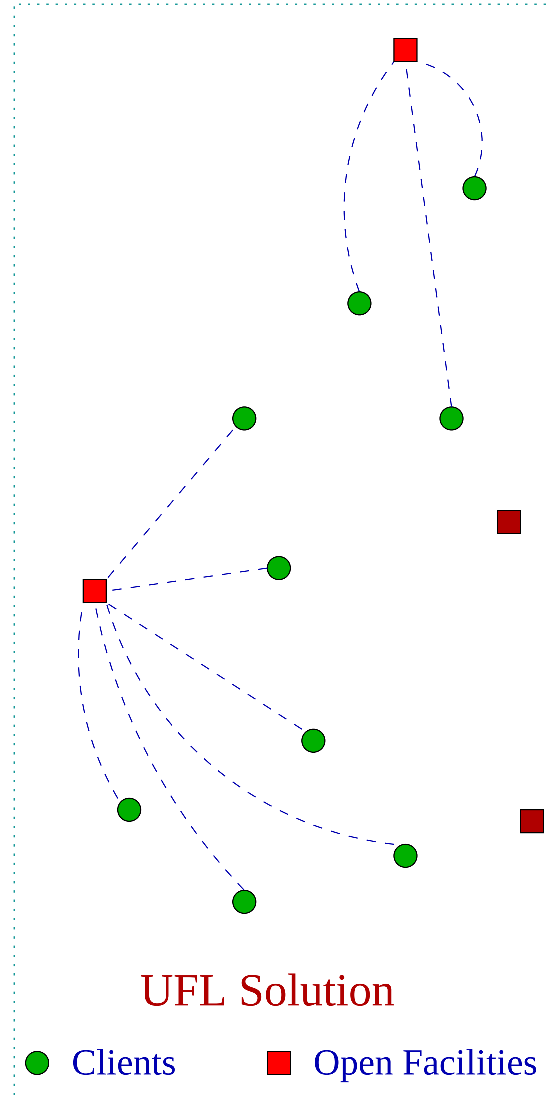

Special cases and past work

- $u = 1$: UFL; $\rho_{UFL} = 1.52$ [MYZ 02]. Others: [STA 97, JV 99, AGKMMP 01, JMS 02].
- $u = \infty$: Steiner tree; $\rho_{ST} = 1.55$ [RZ 99]. Others: [TM 80, AKR 95, Zel 95, HP 99].

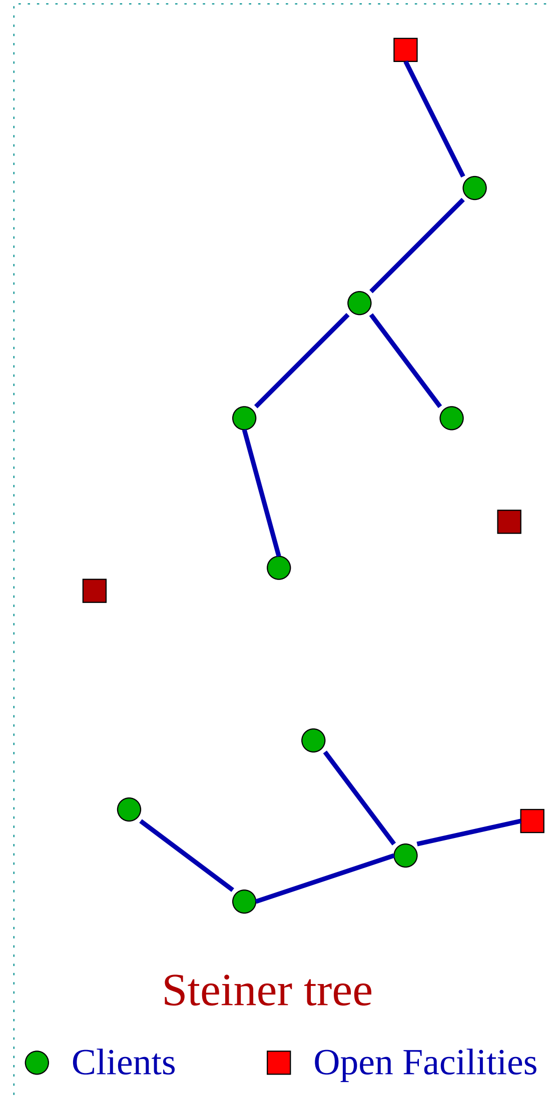

Special cases and past work

- $u = 1$: UFL; $\rho_{UFL} = 1.52$ [MYZ 02]. Others: [STA 97, JV 99, AGKMMP 01, JMS 02].
- $u = \infty$: Steiner tree; $\rho_{ST} = 1.55$ [RZ 99]. Others: [TM 80, AKR 95, Zel 95, HP 99].
- $|\mathcal{F}| = 1$: Single sink single cable edge installation; $\rho_{SS} = 3$ [HRS 00]. Others: [AA 97, AZ 98, GKKRSS 01, GMM 01, Tal 02].

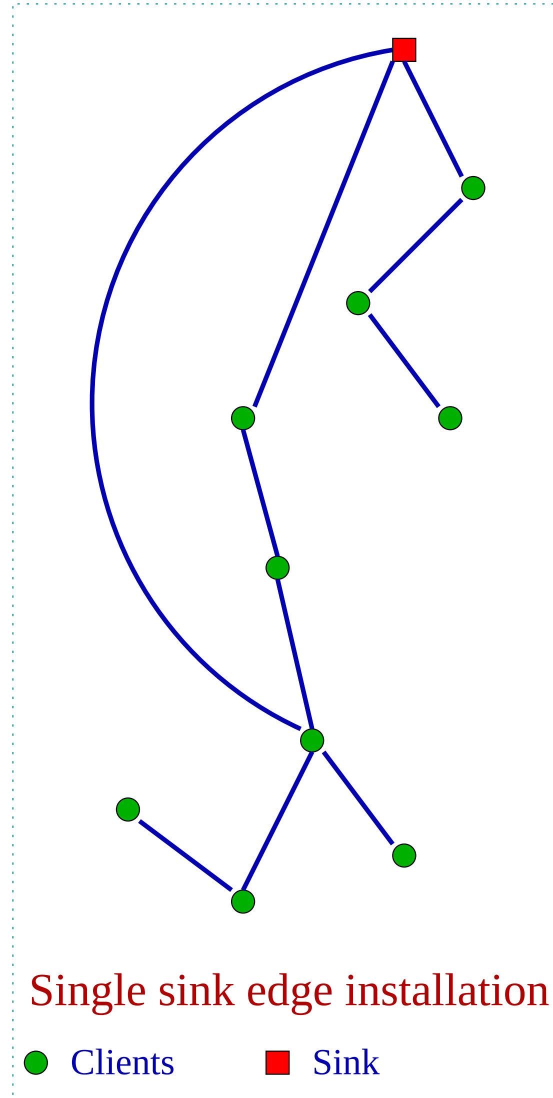

Special cases and past work

- $u = 1$: UFL; $\rho_{UFL} = 1.52$ [MYZ 02]. Others: [STA 97, JV 99, AGKMMP 01, JMS 02].
- $u = \infty$: Steiner tree; $\rho_{ST} = 1.55$ [RZ 99]. Others: [TM 80, AKR 95, Zel 95, HP 99].
- $|\mathcal{F}| = 1$: Single sink single cable edge installation; $\rho_{SS} = 3$ [HRS 00]. Others: [AA 97, AZ 98, GKKRSS 01, GMM 01, Tal 02].
- CCFL: $O(\log n)$ [MMP 00], also for KCFL. This paper: 3.07

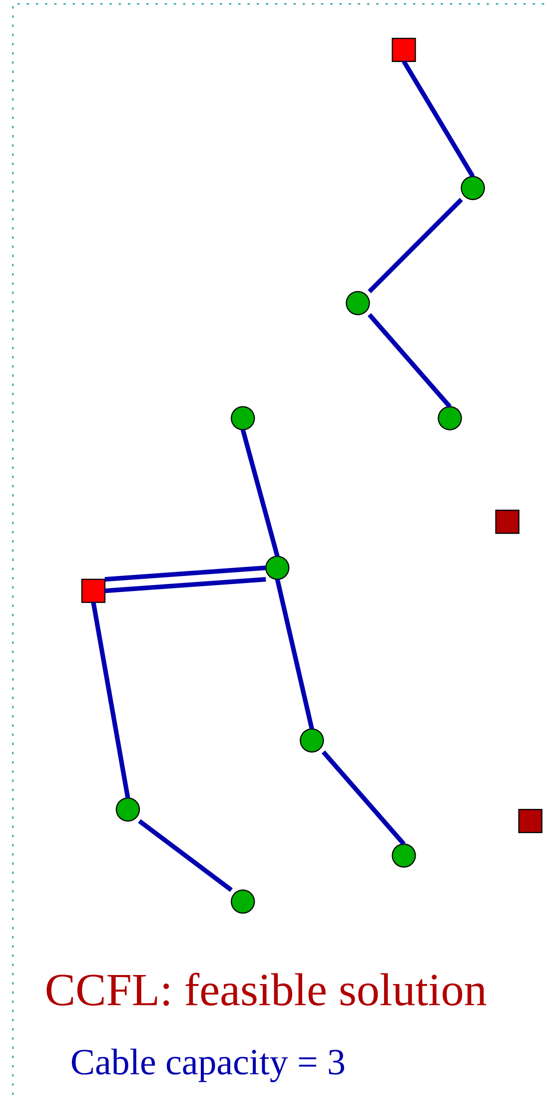

Lower bound: Routing

- New UFL instance: Scale edge costs to $c'_e = c_e/u$.
- $OPT(UFL) \leq OPT(CCFL)$.
- Reason: In CCFL, each client incurs service cost at least $1/u$ of the cost of its path to its facility.

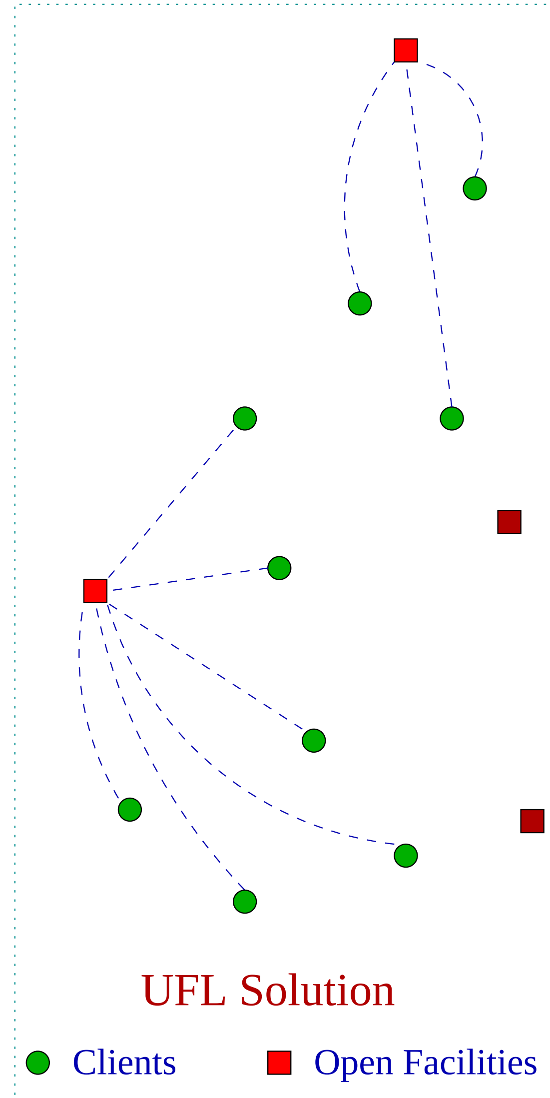

Lower bound: Connectivity

- New Steiner tree instance: Add *root* r, connect to each facility with edge cost ϕ_j. Terminals: $D \cup \{r\}$.

- $OPT(ST) \leq OPT(CCFL)$.

- Reason: In CCFL, each client must have a connection to some facility.

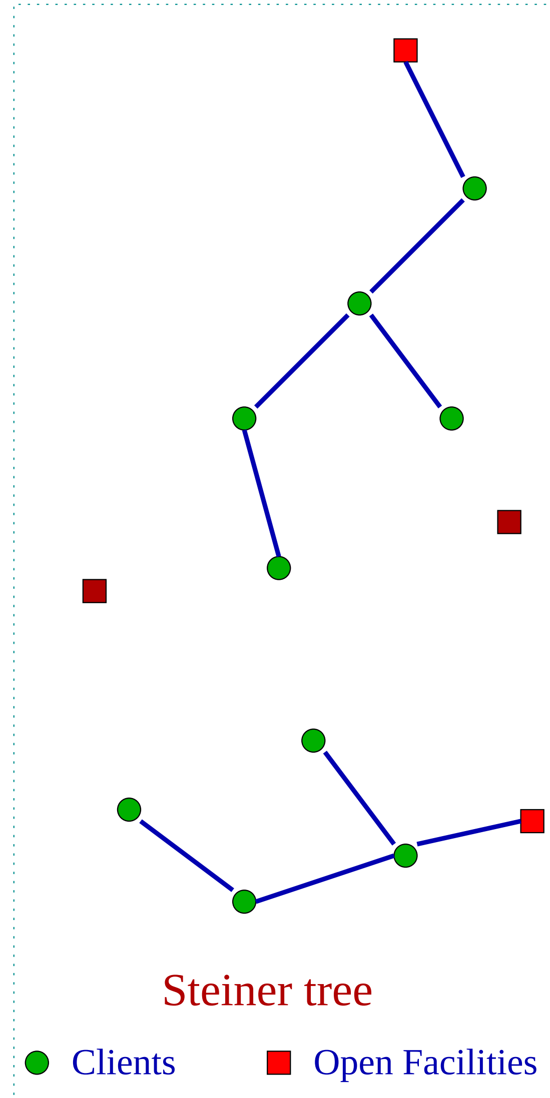

Algorithm motivation

- Routing LB: Good for high demand, bad for low demand.
- Connectivity LB: Bad for high demand, good for low demand.
- How to combine them?

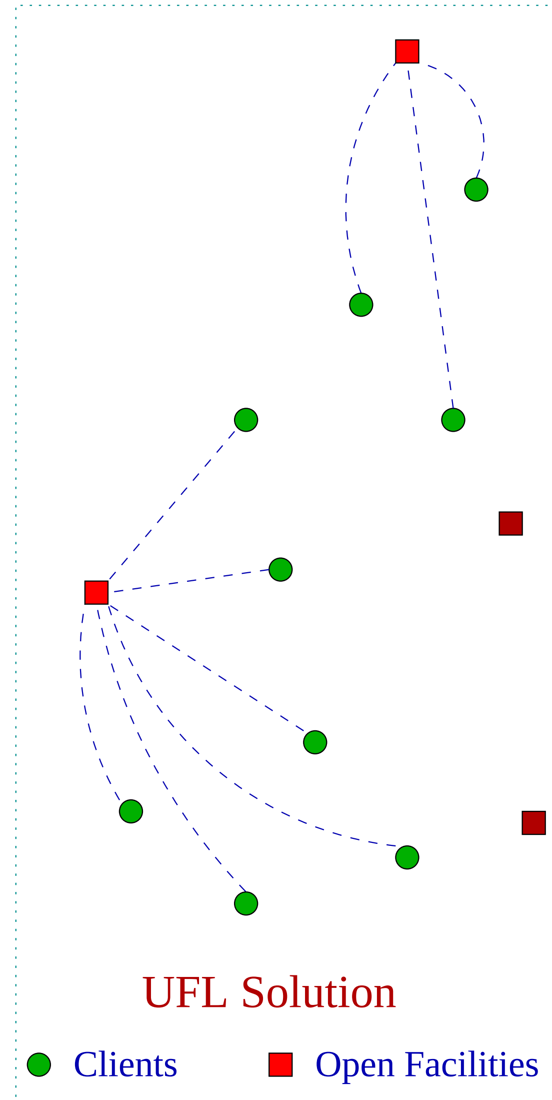

Algorithm motivation

- Routing LB: Good for high demand, bad for low demand.
- Connectivity LB: Bad for high demand, good for low demand.
- How to combine them?

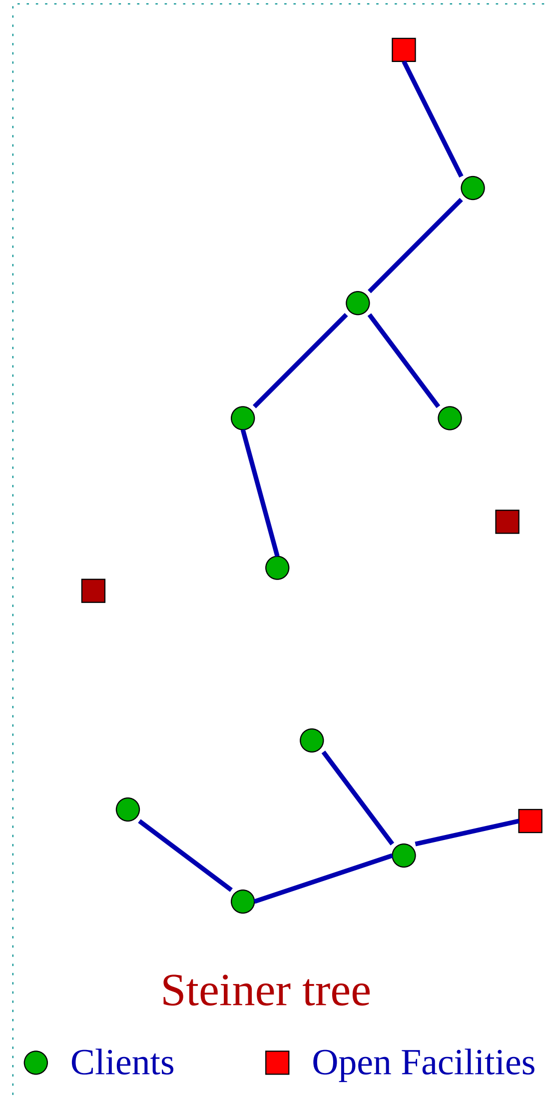

Algorithm motivation

- Routing LB: Good for high demand, bad for low demand.
- Connectivity LB: Bad for high demand, good for low demand.
- How to combine them?
- Use ideas from single sink edge installation algorithm!

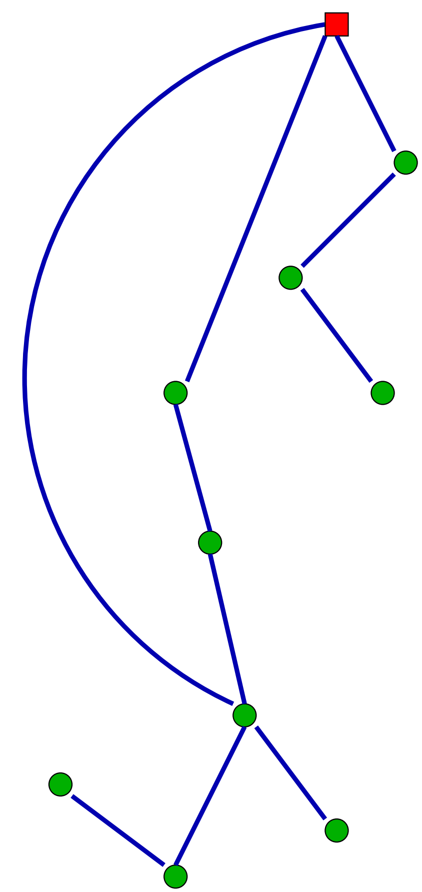

Single sink edge installation

Algorithm description ... 1

1. Solve scaled UFL ($c'_e = c_e/u$).

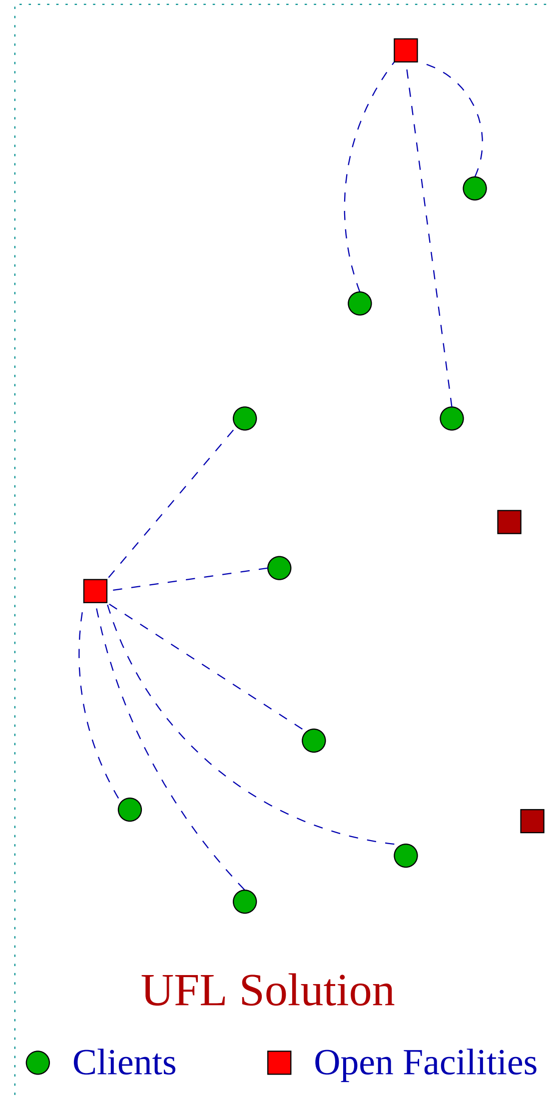

Algorithm description ... 1

1. Solve scaled UFL ($c'_e = c_e/u$).
2. Solve Steiner tree instance.

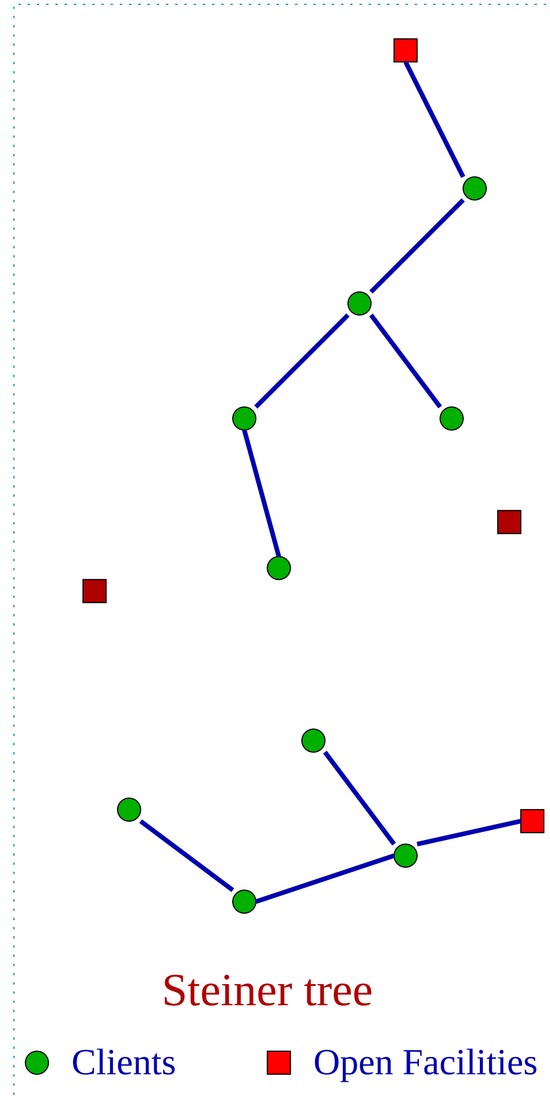

Algorithm description ... 1

1. Solve scaled UFL ($c'_e = c_e/u$).
2. Solve Steiner tree instance.
3. Open facilities of both stages. Install cables of Steiner tree stage.

   This is infeasible!

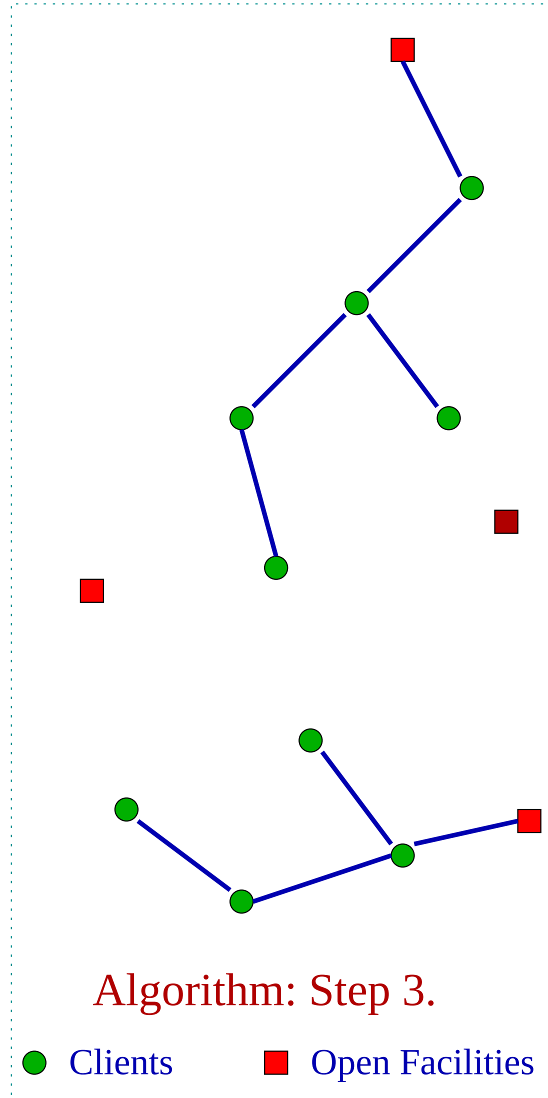

Algorithm: Step 3.

Algorithm description ... 1

1. Solve scaled UFL ($c'_e = c_e/u$).
2. Solve Steiner tree instance.
3. Open facilities of both stages. Install cables of Steiner tree stage.
4. Convert to feasible solution by aggregating demand and installing new cables.

   (Details coming up.)

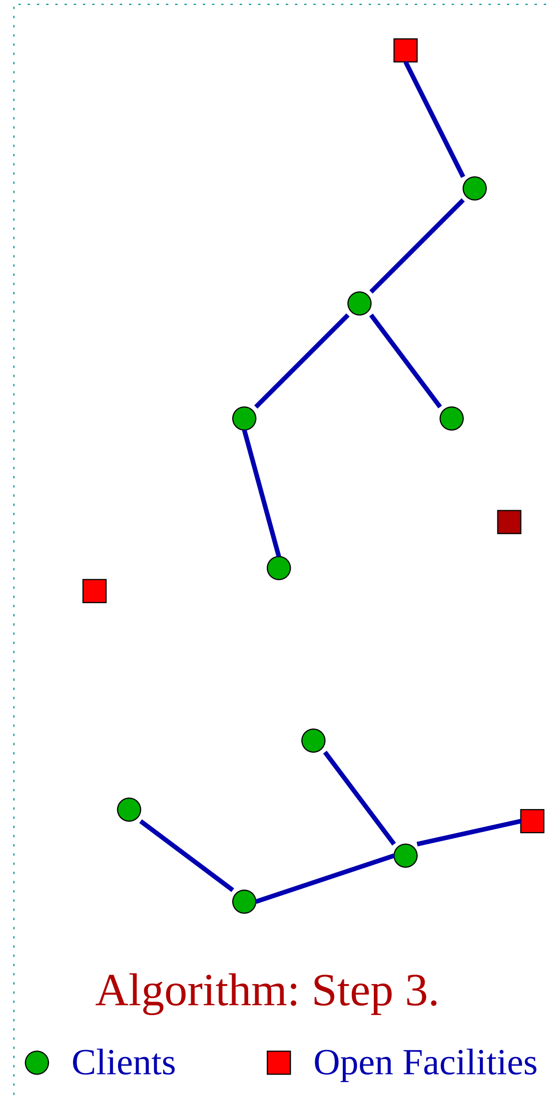

Algorithm description ... 2

4. Installing new cables to make solution feasible:

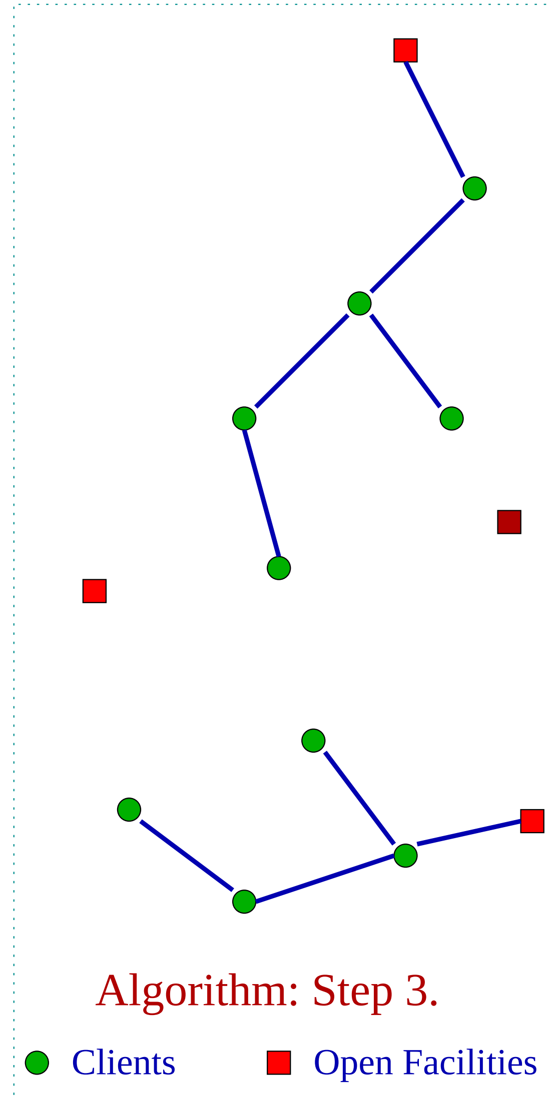

Algorithm description ... 2

4. Installing new cables to make solution feasible:
   For each tree in forest:
   (a) Identify "lowest" node with demand $\geq u$.
       Form "clump" of u nodes in such a subtree.

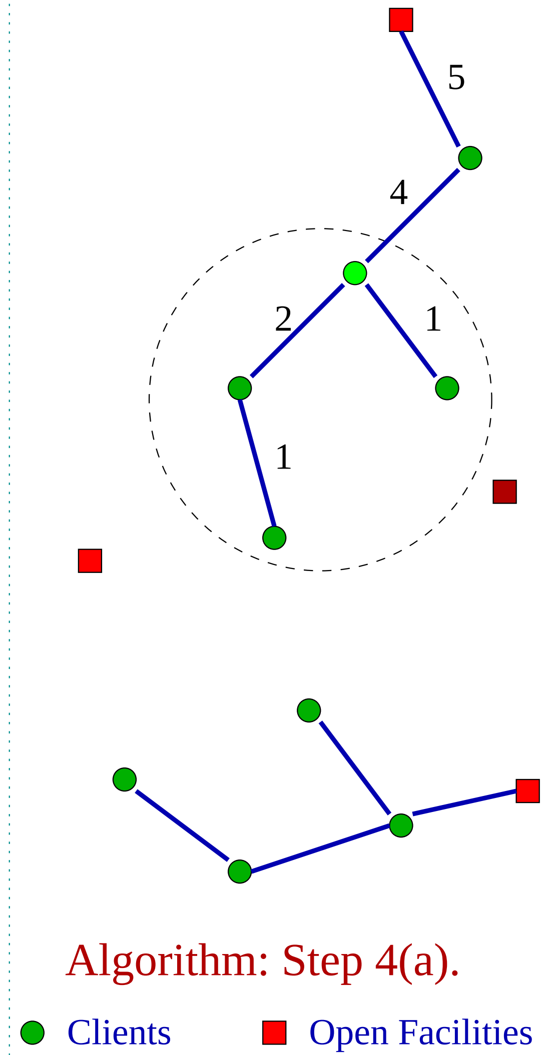

Algorithm: Step 4(a).

Algorithm description ... 2

4. Installing new cables to make solution feasible:
   For each tree in forest:
   (a) Identify "lowest" node with demand $\geq u$.
   Form "clump" of u nodes in such a subtree.
   (b) In this "clump", install a new cable to connect nearest client-facility pair.

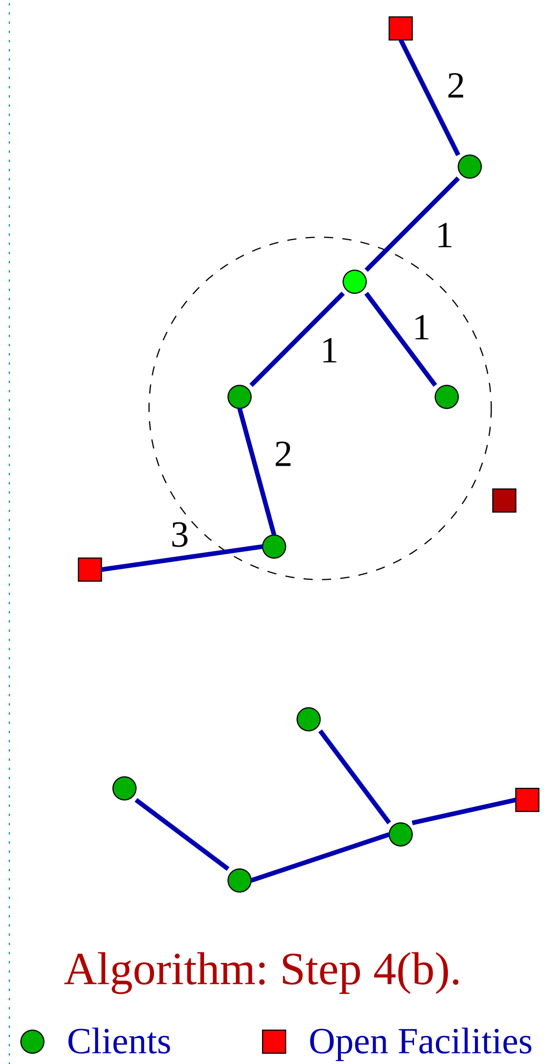

Algorithm: Step 4(b).

Algorithm description ... 2

4. Installing new cables to make solution feasible:
   For each tree in forest:
   (a) Identify "lowest" node with demand $\geq u$.
       Form "clump" of u nodes in such a subtree.
   (b) In this "clump", install a new cable to connect nearest client-facility pair.
   (c) Reroute flow appropriately.

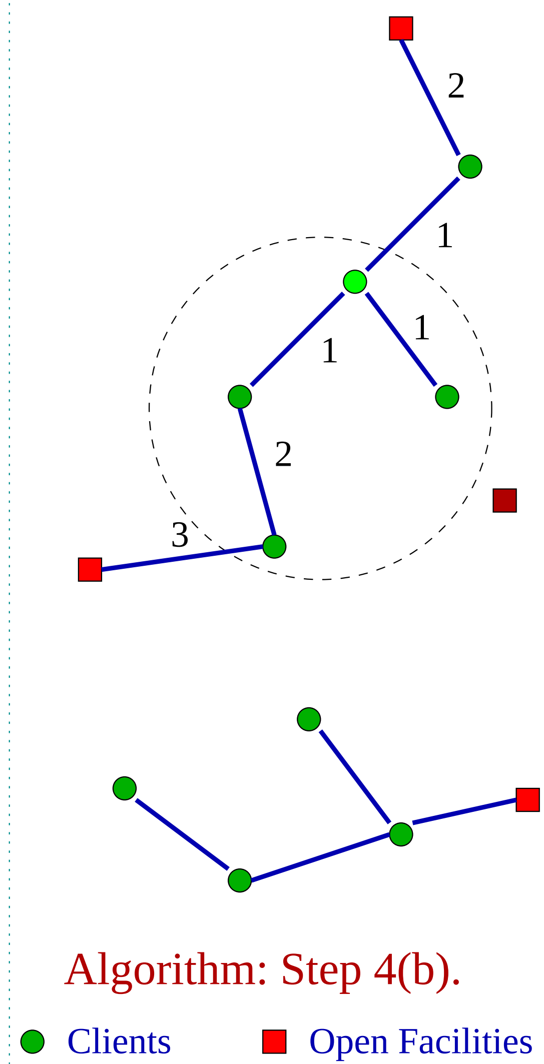

Algorithm: Step 4(b).

Algorithm description ... 2

4. Installing new cables to make solution feasible:
   For each tree in forest:
   (a) Identify "lowest" node with demand $\geq u$.
       Form "clump" of u nodes in such a subtree.
   (b) In this "clump", install a new cable to connect nearest client-facility pair.
   (c) Reroute flow appropriately.

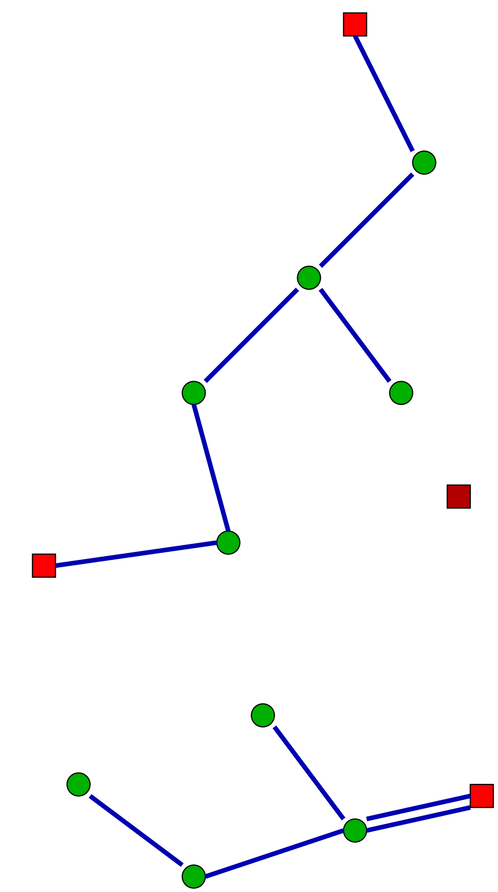

Algorithm: Final solution.

Performance analysis

- Theorem [HRS 00]: Aggregation-and-rerouting produces feasible solution.

Performance analysis

- Theorem [HRS 00]: Aggregation-and-rerouting produces feasible solution.
- Facility cost: Paid by the two lower bounds.

Performance analysis

- Theorem [HRS 00]: Aggregation-and-rerouting produces feasible solution.
- Facility cost: Paid by the two lower bounds.
- Cables on Steiner tree: Paid by Steiner tree lower bound.

Performance analysis

- Theorem [HRS 00]: Aggregation-and-rerouting produces feasible solution.
- Facility cost: Paid by the two lower bounds.
- Cables on Steiner tree: Paid by Steiner tree lower bound.
- New cables from "clumps": Paid by routing (service) cost component of UFL solution, since each client in UFL solution incurs c/u service cost and each clump has u clients.

Performance analysis

- Theorem [HRS 00]: Aggregation-and-rerouting produces feasible solution.
- Facility cost: Paid by the two lower bounds.
- Cables on Steiner tree: Paid by Steiner tree lower bound.
- New cables from “clumps”: Paid by routing (service) cost component of UFL solution, since each client in UFL solution incurs c/u service cost and each clump has u clients.
- Theorem: The algorithm is a $\rho_{ST} + \rho_{UFL}$ (≈ 3.07) approximation for CCFL.

Thoughts

- Natural IP formulation has good relaxation. Our algorithm yields a $gap_{ST} + gap_{UFL}$ (≈ 5) LP-rounding approximation algorithm for CCFL.

Thoughts

- Natural IP formulation has good relaxation. Our algorithm yields a $gap_{ST} + gap_{UFL}$ (≈ 5) LP-rounding approximation algorithm for CCFL.
- Generalizes to non-uniform demands at clients. If demand is *splittable*, performance ratio remains same (≈ 3.07).
  For unsplittable demand, the aggregate-and-reroute step needs a little more work. Performance ratio is now $\rho_{ST} + 2\rho_{UFL}$ (≈ 4.59).

Thoughts

- Natural IP formulation has good relaxation. Our algorithm yields a $gap_{ST} + gap_{UFL}$ (≈ 5) LP-rounding approximation algorithm for CCFL.
- Generalizes to non-uniform demands at clients. If demand is *splittable*, performance ratio remains same (≈ 3.07).
  For unsplittable demand, the aggregate-and-reroute step needs a little more work. Performance ratio is now $\rho_{ST} + 2\rho_{UFL}$ (≈ 4.59).
- No tight example known.
  Lower bound on approximation ratio is 1.46, coming from UFL.

KCFL: k-cable facility location

- Variant of CCFL: k cable types to choose from.

  Depending on flow, one particular type of cable may be most economical.

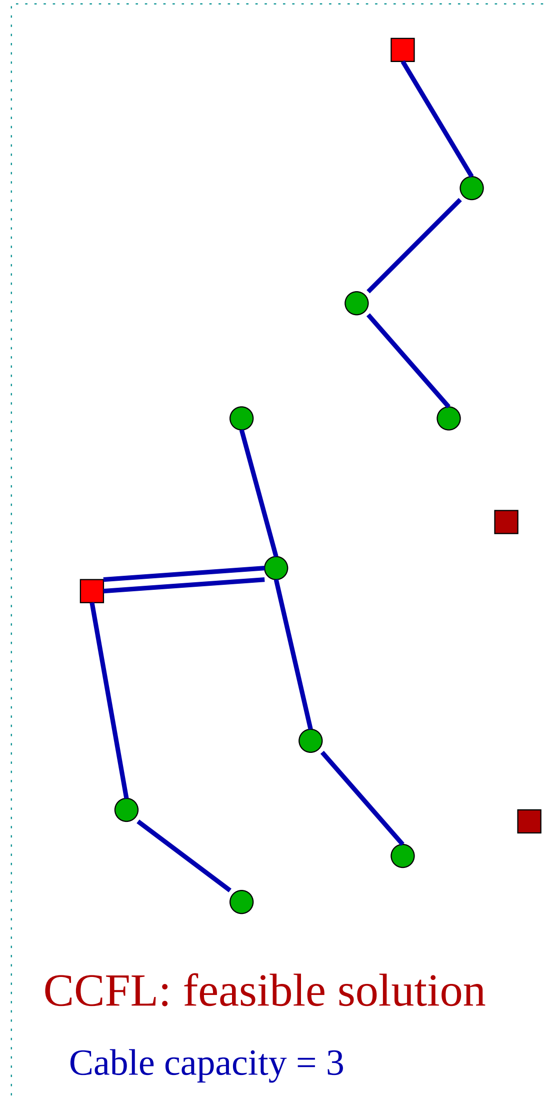

KCFL: k-cable facility location

- Variant of CCFL: k cable types to choose from.

  Depending on flow, one particular type of cable may be most economical.

- Current status:

  $O(\log n)$ due to [MMP 00],

  $O(k)$ due to [RS 02].

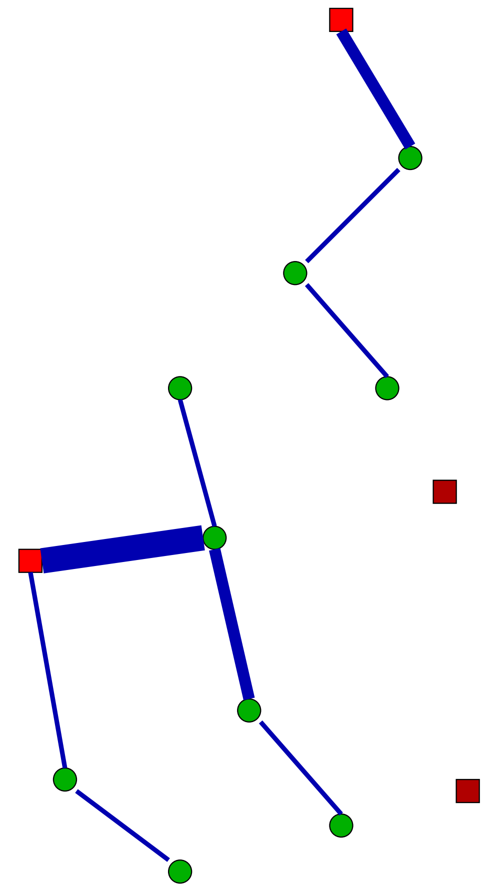

KCFL: feasible solution

k-cable single sink edge installation

- Single sink: $|\mathcal{F}| = 1$.
- $O(1)$ approximation due to [GMM 01].
  Combinatorial, randomized algorithm, using same structural lower bounds (routing and connectivity).

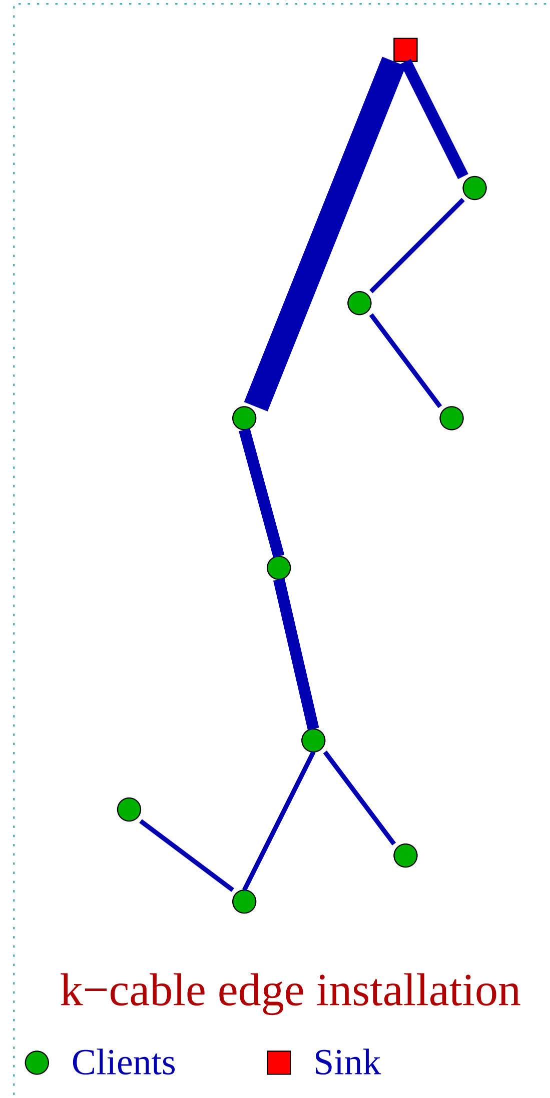

k-cable single sink edge installation

- Single sink: $|\mathcal{F}| = 1$.
- $O(1)$ approximation due to [GMM 01].
  Combinatorial, randomized algorithm, using same structural lower bounds (routing and connectivity).
- Improved $O(1)$ approximation due to [Tal 02].
  LP rounding, improves on $O(k)$ of [GKKRSS 01].

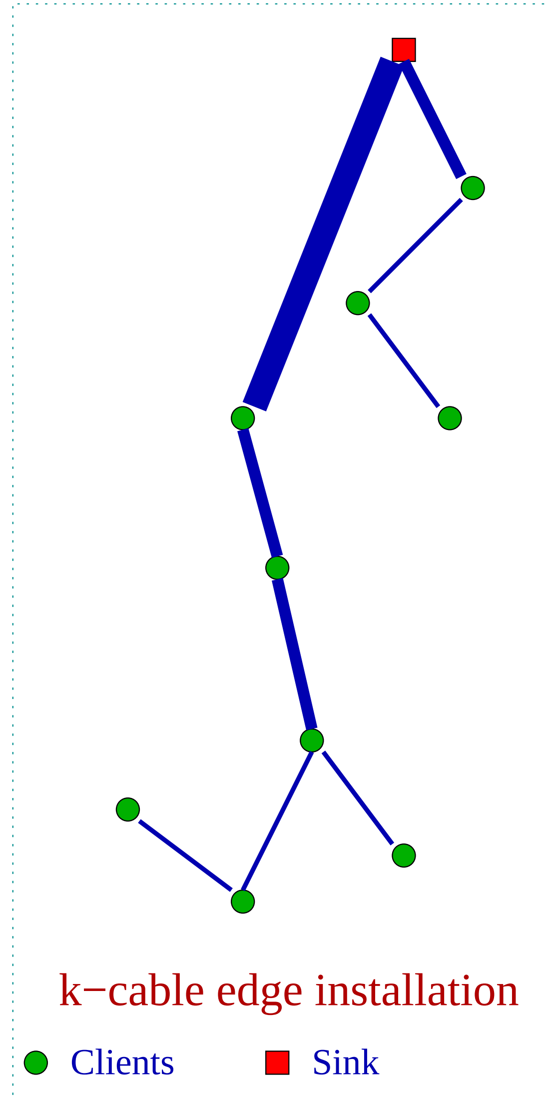

Thoughts on approximating KCFL

- Open: $O(1)$ for KCFL?

Thoughts on approximating KCFL

- Open: $O(1)$ for KCFL?
- Extend [GMM 01]: yields $O(k)$ [RS 02], with $O(1)$ on cable cost. $O(k)$ is only due to facility costs.
  Better algorithm / analysis may yield $O(1)$.

Thoughts on approximating KCFL

- Open: $O(1)$ for KCFL?
- Extend [GMM 01]: yields $O(k)$ [RS 02], with $O(1)$ on cable cost. $O(k)$ is only due to facility costs.
  Better algorithm / analysis may yield $O(1)$.
- Slight modification of LP of [Tal 02] yields formulation of KCFL.
  Open: Rounding or gap for LP.

Thoughts on approximating KCFL

- Open: $O(1)$ for KCFL?
- Extend [GMM 01]: yields $O(k)$ [RS 02], with $O(1)$ on cable cost. $O(k)$ is only due to facility costs.
  Better algorithm / analysis may yield $O(1)$.
- Slight modification of LP of [Tal 02] yields formulation of KCFL.
  Open: Rounding or gap for LP.
- Open: CCFL / KCFL with capacitated facilities.