In the following we design oblivious algorithms that obtain close to optimum congestion (no bounds on dilation).

We always assume that we route a flow (instead of packet routing).

We can also assume this is a randomized path-selection scheme that guarantees that the expected load on an edge is close to the optimum congestion.

Hierarchical Decompositions

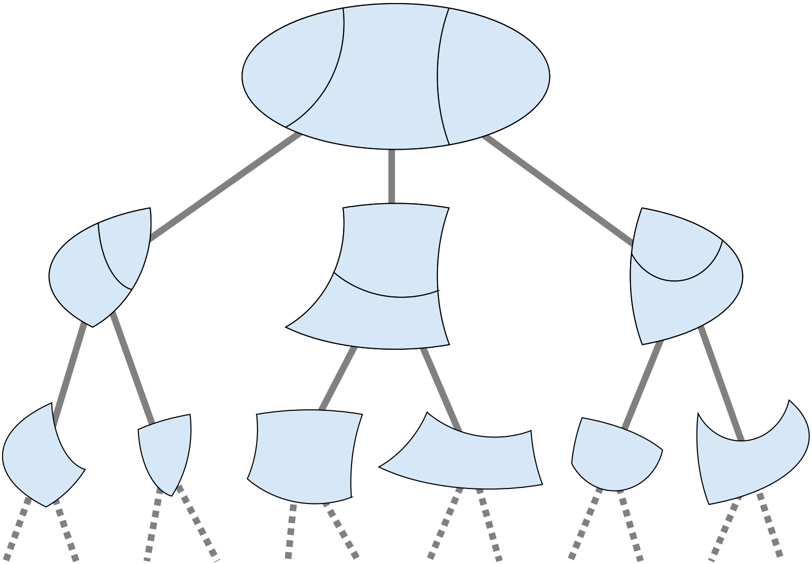

Hierarchical Decompositions & Oblivious Routing

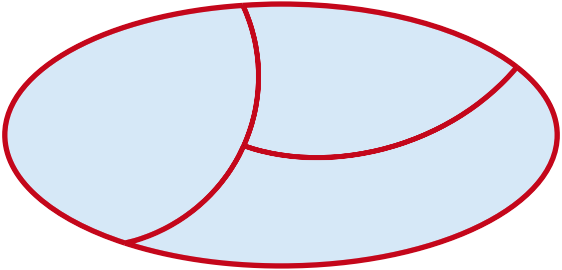

define multicommodity flow problem for every cluster:

- every border edge of a sub-cluster injects one unit and distributes it evenly to all others

Formally

- cluster S partitioned into clusters $S_1, \dots, S_\ell$
- weight $w_S(v)$ of node v is total capacity of edges connecting v to nodes in other sub-clusters or outside of S
- demand for pair $(x,y) \in S \times S$

$$\frac{w_S(x)w_S(y)}{w_S(S)}$$

- gives flow problem for every cluster
- if every flow problem can be solved with congestion C then there is an oblivious routing scheme that always obtains congestion

$$\mathcal{O}(\mathrm{height}(T) \cdot C \cdot C_{\mathrm{opt}}(\mathcal{P}))$$

Oblivious Routing Scheme

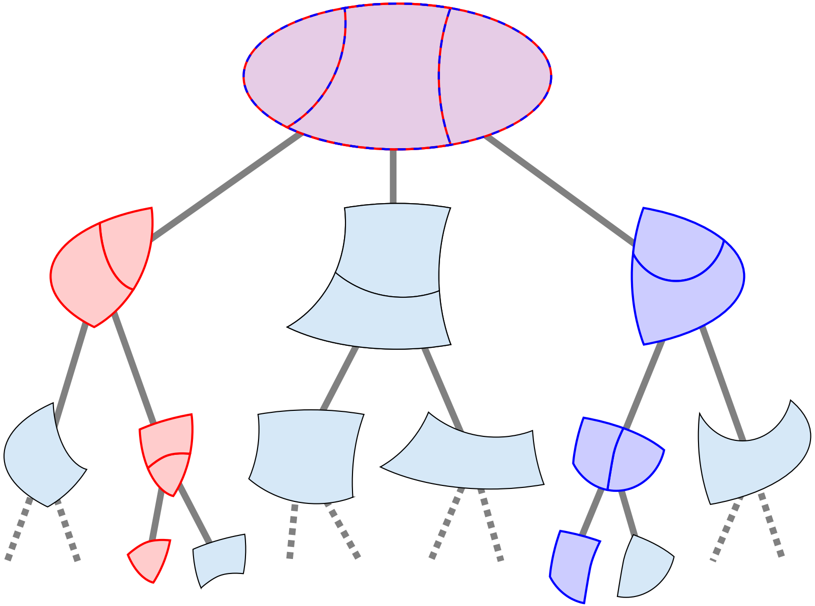

Oblivious Routing Scheme — A Single Cluster S

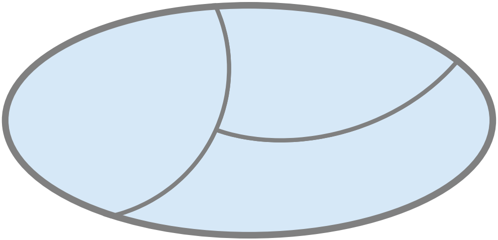

Input:
Messages from sub-clusters have been routed to random border-edges of corresponding sub-cluster.

1. forward messages to random intra sub-cluster edge
2. delete messages for which source and target are in S
3. forward remaining messages to random border edge

all performed by applying flow problem for cluster several times

Sparsest Cut

Definition 1

Given a multicommodity flow problem $\mathcal{P}$ with demands D_i between source-target pairs s_i, t_i. A sparsest cut for $\mathcal{P}$ is a set S that minimizes

$$\Phi(S) = \frac{\text{capacity}(S, V \setminus S)}{\text{demand}(S, V \setminus S)} \ .$$

$\text{demand}(S, V \setminus S)$ is the demand that crosses cut S.

$\text{capacity}(S, V \setminus S)$ is the capacity across the cut.

Sparsest Cut

Clearly,

$$1/\Phi_{\text{min}} \le C_{\text{opt}}(\mathcal{P})$$

For single-commodity flows we have $1/\Phi_{\text{min}} = C_{\text{opt}}(\mathcal{P})$.

In general we have

$$\frac{1}{\Phi_{\text{min}}} \le C_{\text{opt}}(\mathcal{P}) \le \mathcal{O}(\log n) \cdot \frac{1}{\Phi_{\text{min}}} \ .$$

This is known as an approximate maxflow mincut theorem.

LP Formulation

Maximum Concurrent Flow:

$$
\begin{array}{llrcl}
\max & & \lambda & & \\
\text{s.t.} & \forall i & \sum_{p \in P_{s_i,t_i}} f_p & \geq & D_i \\
 & \forall e \in E & \sum_{p: e \in p} f_p & \leq & c(e) \\
 & & f_p, \lambda & \geq & 0
\end{array}
$$

$\mathcal{P}_{s,t}$ is the set of path that connect s and t.

The Dual:

$$
\begin{array}{llrcl}
\min & & \sum_e c(e)\ell(e) & & \\
\text{s.t.} & \forall p \in \mathcal{P} & \sum_{e \in P} \ell(e) & \geq & \text{dist}_i \\
 & & \sum_i D_i \text{dist}_i & \geq & 1 \\
 & & \text{dist}_i, \ell(e) & \geq & 0
\end{array}
$$

Duality

Primal:

$$\begin{array}{llll} \max & c^t x & & \\ \text{s.t.} & Ax & \leq & b \\ & x & \geq & 0 \end{array}$$

Dual:

$$\begin{array}{llll} \min & b^t y & & \\ \text{s.t.} & A^t y & \geq & c \\ & y & \geq & 0 \end{array}$$

Metric Embeddings

Definition 2

A metric (V, d) is an ℓ_1-embeddable metric if there exists a function $f : V \to \mathbb{R}^m$ for some m such that

$$d(u,v) = \|f(u) - f(v)\|_1$$

Definition 3

A metric (V, d) embeds into ℓ_1 with distortion α if there exists a function $f : V \to \mathbb{R}^m$ for some m such that

$$\frac{1}{\alpha}\|f(u) - f(v)\|_1 \leq d(u,v) \leq \|f(u) - f(v)\|$$

Theorem 4

Any metric (V, d) *on* $|V| = n$ *points is embeddable into* ℓ_1 *with distortion* $\mathcal{O}(\log n)$.

Theorem 5

For any flow problem $\mathcal{P}$ one can obtain at least a throughput of $\Phi_{min}/\log n$, where $\Phi_{\min}$ denotes the sparsity of the sparsest cut. In other words

$$C_{opt}(\mathcal{P}) \leq \mathcal{O}(\log n)\frac{1}{\Phi_{min}}$$

LP Formulation

The optimum throughput is given by

$$\begin{array}{lrl} \min & \sum_e c(e)d(e) & \\ \text{s.t.} & d \text{ metric} & \\ & \sum_i D_i d(s_i,t_i) & \geq \; 1 \end{array}$$

or

$$\begin{aligned} C_{\text{opt}}(\mathcal{P}) &= \frac{\sum_i D_i d(s_i,t_i)}{\sum_{e=(u,v)} c(e)d(u,v)} \\ &\leq \alpha \frac{\sum_i D_i \cdot \|f(s_i)-f(t_i)\|}{\sum_{e=(u,v)} c(e)\cdot\|f(u)-f(v)\|} \\ &= \alpha \frac{\sum_i D_i \cdot \sum_S \gamma_S \chi_S(s_i,t_i)}{\sum_{e=(u,v)} c(e)\cdot \sum_S \gamma_S\chi_S(u,v)} \\ &= \alpha \frac{\sum_S \gamma_S \sum_i D_i \chi_S(s_i,t_i)}{\sum_S \gamma_S \sum_{e=(u,v)} c(e)\chi_S(u,v)} \\ &\leq \alpha \max_S \frac{\sum_i D_i \chi_S(s_i,t_i)}{\sum_{e=(u,v)} c(e)\chi_S(u,v)} = \alpha \cdot \frac{1}{\Phi_{\min}} \end{aligned}$$

Fréchet Embedding

Given a set A of points we define a mapping

$$f(x) := d(x, A)$$

The mapping f is contracting this means

$$\|f(x) - f(y)\| \leq d(x, y)$$

Suppose we have a probability distribution p over sets $A_1, \dots, A_k$:

Then define $f : V \to \mathbb{R}^k$ by

$$f(x)_i : V = p(A_i) \cdot d(x, A_i)$$

f is still contracting.

We use a probability distribution over sets such that the expected distance between x and y is at least

$$d(x,y)/\mathcal{O}(\log n)$$