
Communications Network Design

lecture 22

Matthew Roughan

<matthew.roughan@adelaide.edu.au>

Discipline of Applied Mathematics
School of Mathematical Sciences
University of Adelaide

March 2, 2009

Communications Network Design: lecture 22 – p.1/19

This lecture considers the problems of routing and network design where the critical input data (e.g. the traffic matrix) is missing or incomplete.

Communications Network Design: lecture 22 – p.1/19

Network Design without complete information

What can we do where the critical input data (e.g. the traffic matrix) is missing or incomplete? The answer is to optimize with respect to all possible traffic matrices. There are several possible algorithms: oblivious routing and Valiant network design.

Communications Network Design: lecture 22 – p.2/19

Communications Network Design: lecture 22 – p.2/19

Oblivious

- ▶ I have mentioned
 - ▷ it can be hard to measure traffic demands
 - ▷ they will have errors
 - ▷ even when we measure precisely, there are errors in forecasts
- ▶ what can we do?
 - ▷ oblivious routing
 - ▷ Valiant network design
- ▶ design principles that are **oblivious** to the traffic
 - ▷ they will be suboptimal for any particular traffic
 - ▷ on average they will do better than any particular optimization approach (that requires knowledge of the traffic)

Oblivious routing [1, 2]

- ▶ assume traffic matrix $T = [t_{pq}]$ is unknown
- ▶ network $G(N, E)$, with link capacities r_e
- ▶ optimization objective:
 - ▷ **minimize maximum utilization**
 - ▷ link utilization $u_e = f_e/r_e$
- ▶ if we knew T we could write out the standard routing optimization (with a new cost function)

Min-max utilization routing

Formulation:

$$\begin{aligned} \text{minimize } C(\mathbf{f}) &= \max_{e \in E} \frac{f_e}{r_e} \\ \text{such that} \\ f_e &= \sum_{\mu \in P: e \in \mu} x_\mu, & \forall e \in E \\ x_\mu &\geq 0, & \forall \mu \in P \\ \sum_{\mu \in P_{pq}} x_\mu &= t_{pq}, & \forall [p, q] \in K \\ f_e &\leq r_e, & \forall e \in E \end{aligned}$$

Min-max utilization routing

- ▶ if we knew T , we can solve this routing problem
- ▶ call the solution $\mathbf{f}^{\text{opt}}(T)$
- ▶ denote the cost of this solution by

$$C_{\text{opt}}(T) = C(\mathbf{f}^{\text{opt}}(T)) = \max_{e \in E} \frac{f_e^{\text{opt}}}{r_e}$$

- ▶ given any other routing \mathbf{f} we can compare costs by computing

$$\text{PERF}(\mathbf{f}, \{T\}) = \frac{C(\mathbf{f})}{C_{\text{opt}}(T)}$$

- ▶ this is the relative cost of routing \mathbf{f} with respect to the optimal possible routing, given T

Oblivious routing

- ▶ but we don't know T
 - ▷ assume T can take any possible value in a set \mathcal{T} of possible traffic matrices
 - ▷ in the worst case, the cost is

$$\text{PERF}(\mathbf{f}, \mathcal{T}) = \max_{T \in \mathcal{T}} \text{PERF}(\mathbf{f}, \{T\}) = \max_{T \in \mathcal{T}} \frac{C(\mathbf{f})}{C_{\text{opt}}(T)}$$

- ▶ so write a new optimization problem

$$\text{minimize } \text{PERF}(\mathbf{f}, \mathcal{T}) = \min_{\mathbf{f}} \max_{T \in \mathcal{T}} \frac{C(\mathbf{f})}{C_{\text{opt}}(T)}$$

Oblivious routing

- ▶ there are now several ways to solve these [1, 3, 2]
 - ▷ polynomial time algorithms exist
- ▶ also several theoretical bounds [1, 3, 2]
- ▶ note $\text{PERF}(\mathbf{f}, \mathcal{T}) \geq 1$
 - ▷ the oblivious routing is always suboptimal compared to the optimal routing for a known traffic matrix
 - ▷ larger values mean that oblivious routing is worse than optimal routing
 - ▷ bounds grow polylogarithmically with $|N|$
 - ▷ so in theory the method gets worse for larger networks


Oblivious routing

- ▶ Applegate and Cohen [2] show that typical values for real(ish) networks are about 1.4-1.9
 - ▷ $\text{PERF}(\mathbf{f}, \mathcal{T}) = 1.4$ corresponds to 40% extra capacity needed in a network (40% extra costs)
 - ▷ $\text{PERF}(\mathbf{f}, \mathcal{T}) = 1.9$ corresponds to 90% extra capacity needed in a network (90% extra costs)
- ▶ this is a significant extra cost
 - ▷ best you can do if you don't know the traffic matrix
 - ▷ better than many other approaches (given lack of knowledge about traffic matrix)
 - ▷ better than theoretical bounds

Valiant network design [4, 5, 6]

- ▶ previous example was for routing
- ▶ can we do network design where we don't know the traffic matrix?
- ▶ Yes! Use a trick from design of router backplanes

Valiant network design




Communications Network Design: lecture 22 – p.11/19

Communications Network Design: lecture 22 – p.11/19

Valiant network design

Abstract the access network to have capacity C



Communications Network Design: lecture 22 – p.12/19

Communications Network Design: lecture 22 – p.12/19


Valiant network design

Simple case (which can be generalized)

- ▶ assume access capacity C to each backbone node
- ▶ maximum value of offered traffic between two nodes is $\max t_{pq} = C$
 - ▷ in fact $\sum_q t_{pq} \leq C$ and $\sum_p t_{pq} \leq C$
 - ▷ but we don't know t_{pq}
- ▶ route traffic demand t_{pq} as follows
 - ▷ divide it into $|N|$ even groups
 - ▷ route group i as follows $p - i - q$
 - ▷ load balance across all of the possible 2 hop routes
 - ▷ do the same for all $p, q \in N$

Communications Network Design: lecture 22 – p.13/19

Valiant network design



Communications Network Design: lecture 22 – p.14/19

Communications Network Design: lecture 22 – p.13/19

Communications Network Design: lecture 22 – p.14/19

Valiant network design

- ▶ calculate backbone link capacity requirements
 - ▷ assume maximal traffic, e.g. $\sum_q t_{pq} = C$
 - ★ total traffic entering the network at node p fills the capacity of the access network coming into p
- ▶ for all p the first hop divides this traffic evenly amongst all $|N|$ links from $p-i$, e.g. creates loads

$$f_{p,i} = \frac{C}{|N|}$$

note we include a dummy link $p-p$ here

Valiant network design

- ▶ the second hop divides this traffic evenly amongst all $|N|$ links from $i-q$
 - ▷ its the dual of the previous step
- ▶ so traffic from node p creates loads $\frac{C}{|N|}$ on all links $p-i$ and $i-q$.
- ▶ we sum over all $|N|$ source nodes p and we create loads

$$\frac{2C}{|N|}$$

on all links in the network.

Valiant network design

- ▶ total capacity in the backbone is $2C|N|$
 - ▷ compare to each link with capacity C , total $C|N|^2$
- ▶ assume cost is proportional to bandwidth
 - ▷ not distance (linear cost model has $\alpha_e = 1, \beta_e = 0$)
- ▶ optimal network has capacity $\sum_{p,q} t_{pq} = C|N|$
- ▶ we have introduced factor of 2 extra bandwidth
- ▶ this design is provably the best oblivious network design [6] (given above costs).
- ▶ it also has great advantages for survivability
 - ▷ can survive any combination of node failures
 - ▷ highly robust to link failures as well
 - ▷ only need marginal increases in link capacities

Communications Network Design: lecture 22 – p.17/19

Better yet

- ▶ both of the above approaches assume we know don't know the traffic matrix
 - ▷ they are oblivious
- ▶ but in reality we know something
 - ▷ e.g. SNMP measurements of traffic on links
- ▶ can we design a network using the information we have, but taking into account the information we are missing?
 - ▷ **this is a current research challenge!**

Communications Network Design: lecture 22 – p.18/19

References

- [1] H. Räcke, "Minimizing congestion in general networks," in FOCS, no. 43, 2002.
- [2] D. Applegate and E. Cohen, "Making intra-domain routing robust to changing and uncertain traffic demands: Understanding fundamental tradeoffs," in ACM SIGCOMM, (Karlsruhe, Germany), pp. 313-324, August 2003.
- [3] Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke, "Optimal oblivious routing in polynomial time," in STOC, (San Diego, CA, USA), 2003.
- [4] L. G. Valiant, "A scheme for fast parallel communication," SIAM Journal on Computing, vol. 11, no. 2, pp. 350-361, 1982.
- [5] R. Zhang-Shen and N. McKeown, "Designing a predictable Internet backbone," in HotNets III, (San Diego, CA), November 2004.
<http://tiny-tera.stanford.edu/~nickm/papers/index.html>.
- [6] R. Zhang-Shen and N. McKeown, "Designing a predictable Internet backbone with Valiant load-balancing," in Thirteenth International Workshop on Quality of Service (IWQoS), (Passau, Germany), June 2005.
<http://tiny-tera.stanford.edu/~nickm/papers/index.html>.

Communications Network Design: lecture 22 – p.17/19

Communications Network Design: lecture 22 – p.19/19