Universal Packet Routing

We are given a arbitrary topology network and we are given N packets to route. For each packet we have a source node and a destination node and the route. We assume that packet routes are *edge-simple*. In each time step only one packet can traverse an edge.

A *schedule* for a set of packets specifies the timing for the movement of the packets.

Goal is to produce a schedule for the packets that minimizes the total time and the maximum queue size needed to route all the packets to their destinations.

Congestion and Dilation

d - dilation: the maximum distance from any packet's source to its destination.

c - congestion: the maximum number of routes through any edge.

A lower bound on routing all the packets: $\Omega(c+d)$.

Can we find a schedule to route all packets in $O(c+d)$ time using only $O(1)$ size queues.

A Randomized online algorithm

Theorem 1. *There is an online algorithm for producing a schedule of length* $O(c + d\log(Nd))$ *using queues of size* $O(\log(Nd))$ *w.h.p.*

Proof.

The algorithm: For each packet, assign a delay chosen randomly, independently, and uniformly from the interval $[1, \frac{\alpha c}{\log(Nd)}]$, where α is a constant (will be fixed later).

First we analyze the *unconstrained* schedule: A packet that is assigned a delay of x waits in its initial queue for x time steps, and then moves on to its final destination without ever stopping.

The expected number of packets that go through an edge in a given time step in the above schedule is at most $log(Nd)/\alpha$.

Use the Chernoff bound to show that w.h.p no more than $O(\log(Nd))$ packets pass through any edge in a given time step.

Each step of the unconstrained schedule can be simulated by $O(\log(Nd))$ steps of a real schedule. This schedule length will be $O(c + d\log(Nd))$ steps using queues of size $O(\log(Nd))$. $\square$

A better schedule

Theorem 2. *For any set of packets whose paths are edge-simple and have congestion* c *and dilation* d*, there is a schedule having length* $(c+d)2^{O(\log^*(c+d))}$ *and maximum queue size* $\log(c+d)2^{O(\log^*(c+d))}$ *in which at most one packet traverses each edge at each step.*

Key Lemma

A T-frame is a sequence of T consecutive time steps. The *frame congestion*, C, in a T-frame is the largest number of packets that traverse any edge in the frame. The *relative congestion*, R in a T-frame is the ratio C/T of the congestion in the frame to the size of the frame.

Lemma 1. *There is a schedule of length* $O(c + d)$ *in which packets never wait in queues and in which the relative congestion in any frame of size* $\log d$ *or greater is at most 1.*

Lovasz Local Lemma

Let $A_1, \ldots, A_n$ be a set of bad events. We want to show that

$$\Pr(\cap_{i=1}^{n} A_i^c) > 0$$

1. If $\sum_{i=1}^{n} \Pr(A_i) < 1$ then $\Pr(\cap_{i=1}^{n} A_i^c) > 0$.

2. If all the A_i's are mutually independent and for all i, $\Pr(A_i) < 1$ then $\Pr(\cap_{i=1}^{n} A_i^c) > 0$.

3. If each A_i depends only on a few other events: The Lovasz Local Lemma.

Definition 1. *A event E is mutually independent of the events $E_1, \ldots, E_n$ if for any $T \subset [1, \ldots, n]$,*

$$\Pr(E | \cap_{j \in T} E_j) = \Pr(E)$$

Definition 2. *A dependency graph for a set of events $E_1, \ldots, E_n$ has n vertices $1, \ldots, n$. Events E_i is mutually independent of any set of events $\{E_j | j \in T\}$ iff there is no edge in the graph connecting i to any $j \in T$.*

Theorem 3. *Let $E_1, \ldots, E_n$ be a set of events. Assume that*

1. *For all i, $\Pr(E_i) \leq p$;*

2. *The degree of the dependency graph is bounded by d.*

3. *$4pd \leq 1$.*

Then

$$\Pr(\cap_{1=1}^{n} E_i^c) > 0$$

Proof

Let $S \subset \{1, \ldots, n\}$. We prove by induction on $s = 1, \ldots, n$ that if $|S| \leq s$, for all k

$$\Pr(E_k | \cap_{j \in S} E_j^c) \leq 2p$$

For $s = 0$, $S = \phi$, trivial.

W.l.o.g. renumber so that $S = \{1, \ldots, s\}$, and (k, j) is not an edge of the dependency graph for $j > d$.

$$\Pr(E_k | E_1^c, \ldots, E_s^c) = \frac{\Pr(E_k E_1^c \ldots E_s^c)}{\Pr(E_1^c \ldots E_s^c)}$$

$$= \frac{\Pr(E_k E_1^c \ldots E_d^c | E_{d+1}^c \ldots E_s^c)}{\Pr(E_1^c \ldots E_d^c | E_{d+1}^c \ldots E_s^c)}$$

$$\Pr(E_k E_1^c \ldots E_d^c | E_{d+1}^c \ldots E_s^c) \leq \Pr(E_k | E_{d+1}^c \ldots E_s^c) = \Pr(E_k) \leq p$$

Using the induction hypothesis:

$$\Pr(E_1^c \ldots E_d^c | E_{d+1}^c \ldots E_s^c) \geq 1 - \sum_{i=1}^{d} \Pr(E_i | E_{d+1}^c \ldots E_s^c) \geq 1 - \sum_{i=1}^{d} 2p \geq 1 - 2pd \geq 1/2$$

Thus,

$$\Pr(E_k | E_1^c, \ldots, E_s^c) \leq \frac{p}{1/2} = 2p$$

$$\Pr(E_1^c \dots E_n^c) = \Pi_{i=1}^n \Pr(E_i^c | E_1^c \dots E_{i-1}^c)$$

$$= \Pi_{i=1}^n (1 - \Pr(E_i | E_1^c \dots E_{i-1}^c)) \geq \Pi_{i=1}^n (1 - 2p) > 0$$