CSE 190, Great ideas in algorithms: Randomized routing

1 Routing

Let $G = (V, E)$ be an undirected graph, where nodes represent processors and edges represent communication channels. Each node wants to send a message to another node: $v \to \pi(v)$, where π is some permutation on the vertices. However, messages can only traverse on edges, and each edge can only carry one message at a given time unit. A routing scheme is a method of deciding on pathes for the messages obeying these restrictions, which tries to minimize the time it takes for all messages to reach their destination. Here, we will focus on the hypercube graph H, which is a common graph used in distributed computation. It has $V = \{0,1\}^n$, and edges allow to flip one bit, so $E = \{(x, x \oplus e_i) : x \in \{0,1\}^n, i \in [n]\}$, where e_i is the i-th unit vector, and $\oplus$ is bitwise xor.

An oblivious routing scheme is a scheme where the path of sending $v \to \pi(v)$ depends just on the endpoints $v, \pi(v)$, and not on the targets of all other messages. Such schemes are easy to implement. A very simple one is the "bit-fixing scheme". To route $v = (v_1, \ldots, v_n)$ to $u = (u_1, \ldots, u_n)$, we flip the bits in order if necessary. So for example, the path from $v = 10110$ to $u = 00101$ is

$$10110 \to 00110 \to 00100 \to 00101.$$

If more than one packet wishes to traverse an edge, then only one packet does so, and the rest are queued for later time steps. The order of sending the packets does not matter much. For example, you can assume a FIFO (First In First Out) queue on every edge.

2 Deterministic routing is bad

Although the maximal distance between pairs of vertices in H is n, routing based on the bit-fixing scheme can incur a very large overhead, due to the fact that edges can only carry one message at a time.

Lemma 2.1. *There are permutations $\pi : \{0,1\}^n \to \{0,1\}^n$ for which the bit-fixing scheme requires at least $2^{n/2}/n$ time steps to transfer all messages.*

Proof. Assume n is even, and write $x \in \{0,1\}^n$ as $x = (x', x'')$ with $x', x'' \in \{0,1\}^n$. Consider any permutation $\pi : \{0,1\}^n \to \{0,1\}^n$ which maps $(x', 0)$ to $(0, x')$ for all $x' \in \{0,1\}^{n/2}$. These $2^{n/2}$ pathes all pass through a single vertex $(0,0)$. As it has only n outgoing edges, we need at least $2^{n/2}/n$ time steps to send all these packages. □

In fact, a more general theorem is true, which shows that any deterministic oblivious routing scheme is equally bad. Here, a deterministic oblivious routing scheme is any scheme in which if $\pi(v) = u$ then the path from v to u depends only on v, u, and more is decided in some deterministic fixed way.

Theorem 2.2. *For any deterministic routing scheme, there exists a permutation* $\pi : \{0,1\}^n \to \{0,1\}^n$ *which requires at least* $2^{n/2}/\sqrt{n}$ *time steps.*

We will not prove this theorem. Instead, we will see how randomization can greatly enhance performance.

3 Randomized routing is good

We will consider the following scheme, which we call the two-step bit-fixing scheme. It uses randomness on top of the bit-fixing routing scheme.

(i) For each $v \in \{0,1\}^n$, choose uniformly an intermediate target $t(v) \in \{0,1\}^n$. Note that t is not necessarily a permutation.

(ii) Route v to $t(v)$ using the bit-fixing routing scheme.

(iii) Route $t(v)$ to $\pi(v)$ using the bit-fixing routing scheme.

The main theorem we prove is the following.

Theorem 3.1. *With high probability, all packets will be routed in at most* $14n$ *time steps.*

In preparation for proving it, we first prove a few results about the deterministic bit-fixing routing scheme.

Claim 3.2. *Let* $v, v', u, u' \in \{0,1\}^n$. *Let* P *be the directed path between* v *and* u *obtained by the bit-fixing routing scheme, and* P' *the directed path from* v' *to* u'. *If the pathes separate at some point, they never intersect again.*

Proof. Let the vertices in P be $v = w^{(1)}, w^{(2)}, \ldots, w^{(m)} = u$ and the vertices in P' be $v' = w'^{(1)}, w'^{(2)}, \ldots, w'^{(m')} = u'$. Assume that $w = w^{(i)} = w'^{(j)}$ but $w^{(i+1)} \neq w'^{(j+1)}$. Then $w^{(i+1)} = w^{(i)} \oplus e_a$ and $w'^{(j+1)} = w'^{(j)} \oplus e_b$ with $a \neq b$. Assume without loss of generality that $a < b$. Then $w_a^{(\ell)} = w_a^{(i+1)} = w_a \oplus 1$ for all $\ell \geq i+1$, while $w_a'^{(\ell)} = w_a'^{(j+1)} = w_a$ for all $\ell \geq j+1$. Thus, the pathes never intersect again. □

Lemma 3.3. *Fix any deterministic oblivious routing scheme. Assume that the path from v to $\pi(v)$ is $e_1, \ldots, e_k$. Let S be the set of all other $v' \in V$ such that the path from v' to $\pi(v')$ traverses one of the edges $e_1, \ldots, e_k$. Then the packet sent from v to $\pi(v)$ will reach its destination after at most $k + |S|$ steps.*

Proof. The proof is by a charging argument. Let p_v be the packet sent from v to $\pi(v)$. Assume that at some time step t, it is supposed to traverse an edge e_i, but instead for some $v' \in S$ the packet $p_{v'}$ is sent over e_i. In this case, the packet p_v generates a "token" and places it on the packet $p_{v'}$. Moreover, whenever a packet $p_{v'}$ with a token on it is supposed to traverse an edge e_j, but instead some other packet $p_{v''}$ traverse it (with $v'' \in S$ necessarily), the token on $p_{v'}$ is moved to $p_{v''}$.

We claim that any packet $p_{v'}$ for $v' \in S$ can have at most one token on it at any given moment. This will show that at most $|S|$ tokens are generated by p_v, and hence it it delayed at most $|S|$ steps, and it reaches its destination after at most $k + |S|$ steps.

To see that, note that tokens always move forward along $e_1, \ldots, e_k$. That is, if we follow a specific token, it starts at some edge e_i, follows a path $e_i, \ldots, e_j$ for some $j \geq i$, and then traverses an edge outside $e_1, \ldots, e_k$. At this point, it can never get back to it, by the claim we just proved. So, we can never have two tokens which traverse the same edge at the same time, and hence two tokens can never be on the same packet. □

Lemma 3.4. *Let $d(v) \in \{0,1\}^n$ be uniformly and independently chosen for each $v \in \{0,1\}^n$. Let $P(v)$ be the path between v and $d(v)$ obtained by the bit-fixing routing scheme. Then with probability at least $1 - 2^{-n}$, any path $P(v)$ intersects with at most $6n$ other pathes $P(w)$.*

Proof. For $v, w \in \{0,1\}^n$, let $X_{v,w} \in \{0,1\}$ be the indicator random variable for the event that $P(v)$ and $P(w)$ intersect. Our goal is to upper bound $X_v = \sum_w X_{v,w}$ for all $v \in \{0,1\}^n$. Before analyzing it, we first analyze a simpler random variable.

Fix an edge $e = (u, u \oplus e_a)$. For $v \in \{0,1\}^n$, let $Y_{e,v} \in \{0,1\}$ be an indicator variable for the event that the edge e belongs to the path $P(v)$ between v and $d(v)$, obtained by the bit-fixing routing scheme. The number of pathes which pass through e is then $\sum_v Y_{e,v}$. Now, the path between v and $d(v)$ pathes through e iff $v_i = u_i$ for all $i > a$, and $d(v)_i = u_i$ for all $i < a$. Let $A_e = \{v \in \{0,1\}^n : v_i = u_i \ \forall i > a\}$. Only v with $v \in A_e$ has a nonzero probability for the path from v to $d(v)$ to go through e. Note that $|A_e| = 2^a$. Moreover, for any $v \in A_e$, the probability that $P(v)$ indeed passes through e is given by

$$\Pr[Y_{e,v} = 1] = \Pr[d(v)_i = u_i \ \forall i < a] = 2^{-(a-1)}.$$

Hence, the expected number of pathes $P(v)$ which go through any edge e is

$$\mathbb{E}[\sum_v Y_{e,v}] = \sum_{v \in A_e} \mathbb{E}[X_v] = 2^a \cdot 2^{-(a-1)} = 2.$$

Now, the pathes $P(v), P(w)$ intersect if some $e \in P(v)$ belongs to $P(w)$. So,

$$X_{v,w} \leq \sum_{e \in P(v)} Y_{e,w}$$

and hence

$$\mathbb{E}[X_v] = \mathbb{E}[\sum_w X_{v,w}] \leq 2n.$$

This means that on average, any fixed path $P(v)$ intersects at most $2n$ other pathes. We would like to show that this happens for all v, with possible worse bounds. To do so, we need a tail bound. We will use a multiplicative version of the Chernoff inequality: if $Z_1, \ldots, Z_N \in \{0,1\}$ are independent random variables, $Z = Z_1 + \ldots + Z_N$ with $\mathbb{E}[Z] = \mu$, then for any $\delta > 0$,

$$\Pr[Z \geq \mu(1+\lambda)] \leq \exp(-\frac{\delta^2}{2+\delta}\mu).$$

In our case, $Z_1, \ldots, Z_N$ are the random variables $\{X_{v,w} : w \in \{0,1\}^n\}$ and $\mu = \mathbb{E}[X_v] \leq 2n$. So

$$\Pr[X_v \geq 6n] \leq \exp(-2n).$$

By the union bound, the probability that $X_v \geq 6n$ for some $v \in \{0,1\}^n$ is bounded by

$$\Pr[\exists v \in \{0,1\}^n, X_v \geq 6n] \leq 2^n \exp(-2n) \leq 2^{-n}.$$

□

So, with very high probability $(1 - 2^{-n})$, we can send all the packets from v to $d(v)$ in at most $6n + n = 7n$ time steps. Similarly, we can send the packets from $d(v)$ to $\pi(v)$ in at most $7n$ time steps (note that this is exactly the same argument, except that the starting point is now randomized, instead of the end point). So, with very high probability, all packets go from v to $\pi(v)$ in at most $14n$ time steps.