Permutation Routing

Andreas Klappenecker

Texas A&M University

The Model

Consider a directed graph $G = (V, E)$ of n nodes that models a communication network.

A directed edge (u, v) means that u can send packets to v.

The goal is to transmit a set of packets through the network, where each packet has a start node and a destination node.

The route of a packet is a path in the graph G.

Packet Routing

The Model

1. A packet can travel at most one edge per timestep.
2. At most one packet can travel along any single edge.
3. Each node has sufficient buffer to store packets.

If the in-degree of a node v is larger than 1, then (2) implies that a packet might get delayed at v.

Packet Routing

Permutation Routing

We want to consider how the network performs under high but fair load.

We assume that each node has one packet starting at it, and one addressed to it. In other words, the routing problem is a permutation π on the set of nodes V.

Hypercube

As an example we consider the hypercube.

Definition

The n-**dimensional hypercube** $Q_n = (V, E)$ has 2^n nodes,

$$V = \{0, 1\}^n.$$

Two nodes u and v are connected by an edge if and only if the Hamming distance between their labels $d(u, v) = 1$, so

$$E = \{(u, v) \in V \times V \mid d(u, v) = 1\}.$$

There is an edge between two nodes if and only if their labels differ in exactly a single bit.

Hypercube

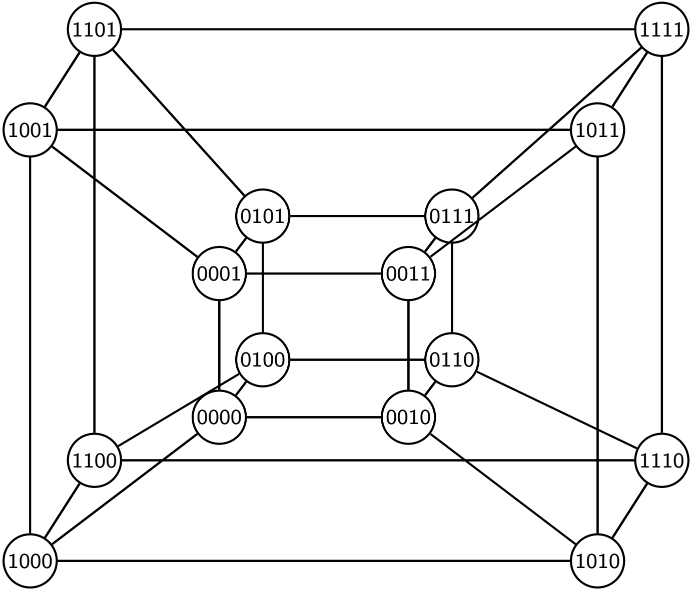
1101
1111
1001
1011
0101
0111
0001
0011
0100
0110
0000
0010
1100
1110
1000
1010

Properties

The hypercube Q_n is a sparse graph, since it has $N = 2^n$ vertices, but just

$$\Theta(N \log N)$$

edges.

Bit-Fixing Routing

The starting point for routing in an n-cube is the bit-fixing algorithm. A packet starting from a node $u = (a_1, a_2, \ldots, a_n)$ with destination $b = (b_1, b_2, \ldots, b_n)$ is routed through

$$
\begin{aligned}
&(a_1, a_2, a_3, \ldots, a_n) \\
\rightsquigarrow&(b_1, a_2, a_3, \ldots, a_n) \\
\rightsquigarrow&(b_1, b_2, a_3, \ldots, a_n) \\
&\vdots \\
\rightsquigarrow&(b_1, b_2, b_3, \ldots, b_n)
\end{aligned}
$$

The actual path is obtained by removing from this the repetitions that occur when $a_i = b_i$.

Congestion

Swap-Bottleneck

Consider the hypercube Q_{2n}. Define the permutation

$$\pi(a_1, a_2, \ldots, a_n, c_1, c_2, \ldots, c_n) = (c_1, c_2, \ldots, c_n, a_1, a_2, \ldots, a_n).$$

Then every routing path reaches a node of the form

$$C = (c_1, c_2, \ldots, c_n, c_1, c_2, \ldots, c_n)$$

with two repeated bit-patterns. There are $N = 2^{2n}$ nodes overall, but just $\sqrt{N} = 2^n$ nodes with repeated address labels.

Congestion

Swap-Bottleneck

There are 2^{2n} packets that are routed through 2^n bottleneck nodes such as C on a route that is $2n$ steps long. This means that

$$\Omega\left(\frac{2^n}{2n}\right) = \Omega\left(\frac{\sqrt{N}}{\log N}\right)$$

steps are needed in the permutation routing problem π.

Avoiding the Bottleneck

Two-Phase Randomized Routing

1. Pick for each packet from a to b a random intermediate node c. Route from a to c using bit-fixing.
2. Route from c to b using bit-fixing.

Gain

We will show that all packets can be delivered in $O(\log N)$ steps with high probability (meaning with probability $1 - O(1/N)$).

We will now give the analysis. There are a couple of difficulties that could make the analysis very difficult (or perhaps even impossible). Pay attention to the tricks that are used to circumvent problems.

Analysis

Let $T(M)$ denote the time the packet M takes to reach its destination. In each of these $T(M)$ steps,

1. the packet M crosses an edge, or
2. the packet M is in a queue (as some other packet crosses an edge that M needs to cross).

Analysis

Let $T(M)$ denote the time the packet M takes to reach its destination. In each of these $T(M)$ steps,

1. the packet M crosses an edge, or
2. the packet M is in a queue (as some other packet crosses an edge that M needs to cross).

Queuing Policy

The queuing policy has a significant impact on $T(M)$, but it is difficult to analyze.

Let $X(e)$ denote the number of packets that have edge e on their route.

Remark

If the route of packet M consists of the edges $(e_1, e_2, \ldots, e_m)$, then

$$T(M) \leqslant \sum_{k=1}^{m} X(e_k).$$

The bound seems loose, but does not depend on a particular routing policy.

Analysis

For a path $P = (e_1, e_2, \ldots, e_m)$, we define

$$T(P) = \sum_{k=1}^{m} X(e_k).$$

If R is the set of all paths used in routing, then

$$\max_{P \in R} T(P)$$

is an upper bound on the routing time for all messages.

Let us consider just phase (1) of the routing algorithm. Let T_1 denote time used to route to the intermediate node and X_1 denote the number of messages that use an edge in their route.

We will show that it is very likely that

$$T_1(P) \leqslant 30n$$

for all possible paths P.

Problem

We want to a high-probability bound for $T_1(P) = \sum_{k=1}^{m} X_1(e_k)$, but the random variables $X_1(e_k)$ are **not** independent. We cannot directly apply a Chernoff bound.

Let $P = (e_1, e_2, \ldots, e_m)$ be a path, and $e_k = (v_{k-1}, v_k)$.

Let j be the bit on which v_{k-1} and v_k differ. We say that a packet M is **active** in node v_{k-1} if and only if

1. the packet M is routed through v_{k-1}, and
2. when the packet M reaches v_{k-1}, its j-th bit has not been fixed yet, but the previous bits $1, 2, \ldots, j-1$ have been fixed or were correct to begin with.

An active packet at v_{k-1} has the potential to cross the edge e_k.

Analysis

Let H_M be the indicator random variable that packet M is active in at least one node on the path P. Let

$$H = \sum_{M \in \{0,1\}^n} H_M.$$

The random variables H_M are mutually independent, because H_M depends only on the intermediate destination of phase I, and these choices are independent for each packet.

Proposition

$$\mathrm{E}[H] \leqslant m \leqslant n.$$

Proof.

Let's consider active packets at v_{k-1}. Suppose that

$$\begin{aligned} v_{k-1} &= (b_1, b_2, \ldots, b_{j-1}, a_j, a_{j+1}, \ldots, a_n) \\ v_k &= (b_1, b_2, \ldots, b_{j-1}, b_j, a_{j+1}, \ldots, a_n) \end{aligned}$$

A packet that is active at v_{k-1} started in a node of the form $(*, \ldots, *, a_j, \ldots, a_n)$. So there are 2^{j-1} possible start nodes.

By condition (1), if a packet is active in v_{k-1}, then its destination must be of the form $(b_1, \ldots, b_{j-1}, *, \ldots, *)$. Therefore, if we consider a fixed starting node, then it will cross the edge e_k with probability $2^{-(j-1)}$.

Therefore, the expected number of active packets at v_{k-1} is 1.

Proof. (Continued)

We only need to consider the m nodes on the path $P = (e_1, e_2, \ldots, e_m)$, namely

$$v_0, v_1, \ldots, v_{m-1}.$$

Consequently, by linearity of expectation, we have

$$\mathrm{E}[H] \leqslant m \cdot 1 \leqslant n. \ \square$$

Remark.

Since the random variables H_k are mutually independent, we may apply Chernoff bounds

$$Pr(H \geqslant 6n) \leqslant 2^{-6n},$$

since $6n \geqslant 6\mathrm{E}[H]$.

Strategy

We choose the events $A = \{T_1(P) \geqslant 30n\}$ and $B = \{H \geqslant 6n\}$ in the estimate

$$\begin{aligned}\Pr[A] &= \Pr[A|B]Pr[B] + Pr[A|\overline{B}]Pr[\overline{B}] \\ &\leqslant \Pr[B] + Pr[A|\overline{B}].\end{aligned}$$

Corollary

For the events $A = \{T_1(P) \geqslant 30n\}$ and $B = \{H \geqslant 6n\}$, we get the estimate

$$\begin{aligned}\Pr[T_1(P) \geqslant 30n] &\leqslant \Pr[B] + Pr[A|\overline{B}] \\ &\leqslant 2^{-6n} + \Pr[T_1(P) \geqslant 30n | H < 6n].\end{aligned}$$

We will next estimate the latter conditional probability. In other words, given that there are less than $6n$ active packets on the path P, we need to bound the number of transitions that they make through the path P.

Analysis

Assume that packet M is active in v_{k-1}.

For M to actually cross the edge (v_{k-1}, v_k), we require its j-th address bit to be b_j. This has probability $1/2$.

However, the packet should not fix any earlier bits $1, \ldots, j-1$. Thus, the actual probability for a packet that is active in v_{k-1} to actually cross (v_{k-1}, v_k) is at most $1/2$.

Analysis

Assume that there are a total of h active packets for the nodes of P. What is the probability that together they make a total of at least $30n$ steps along path P?

Consider as an individual trial a situation where a given active packet is in some given node on P. With probability at most $1/2$ we get **success**, meaning that the packet proceeds along an edge on P. At least with probability $1/2$ we get **failure**, so the packet leaves path P and never returns. When a failure occurs, we move to considering the next active packet. Hence, each success contributes one transition along P, but each failure removes one packet from consideration. To get $30n$ transitions, the first $30n + h$ trials may have at most h failures.

Analysis

The desired conditional probability

$$\Pr(T_1(P) \geqslant 30n | H \leqslant 6n]$$

is therefore the probability that in the repeated trials we get at most $6n$ failures in $36n$ iterations. Since each success probability is at most $1/2$, we easily see that

$$\Pr[T_1(P) \geqslant 30n \mid H \leqslant 6n] \leqslant \Pr[Z \leqslant 6n],$$

where $Z \sim \mathrm{Bin}(36n, 1/2)$. We have $\mathrm{E}[Z] = 18n$.

By applying the Chernoff bound, we get

$$
\begin{aligned}
Pr[T_1(P) \geqslant 30n | H \geqslant 6n] \leqslant \Pr[Z \leqslant \underbrace{(1-2/3)18n}_{=6n}] \\
&\leqslant \exp\left(-18n\left(\frac{2}{3}\right)^2 \Big/ 2\right) \\
&= e^{-4n} \leqslant 2^{-3n-1} \quad \text{since } 2^{3n+1} \leqslant e^{4n}
\end{aligned}
$$

Hence,

$$
\begin{aligned}
\Pr[T_1(P) \geqslant 30n] &\leqslant 2^{-6n} + \Pr(T_1(P) \geqslant 30n | H < 6n) \\
&\leqslant 2^{-6n} + 2^{-3n-1} \\
&\leqslant 2^{-3n}.
\end{aligned}
$$

Conclusions

There are at most 2^{2n} possible packet paths in the hypercube Q_n. Thus, the probability that there is any possible packet path P with $T_1(P) \geqslant 30n$ is bounded by

$$2^{2n}2^{-3n} = 2^{-n} = O(1/|V|),$$

where $V = \{0, 1\}^n$.

We can conclude that both phase I and Phase II will take at most $O(\log|V|)$ steps with high probability.