Deadlock & Livelock Avoidance

In this document we will proof that the Rotary router allows deadlock and livelock freedom in a generic network. In order to demonstrate it we will consider the worst-case scenario. Any other situation can be proved in an analogous way.

Initially, we will take the following assumptions: Firstly, a rotary router is modeled as a circular queue in which packets can move continuously. Secondly, in order to represent the two rings of the Rotary router, we consider that every time a hole enters a router, it can choose each turning way with the same probability.

Let us assume that we have a network with N nodes. Each node in the network is connected to C neighbors, where C is the number of different positions on each circular queue. The worst-case situation would be the one in which the whole network has only one extra hole in one node. As Figure 1 describes, this hole allows the advance of a packet between two nodes. The rest of the nodes have no available buffering space.

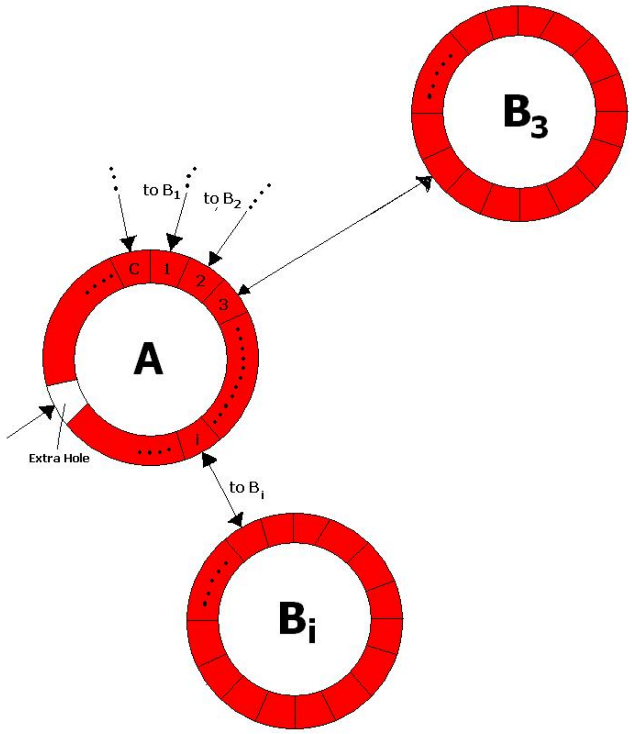

Figure 1: Three nodes of the generic network

Lemma 1

For every position of the hole H in a circular queue, there exist $0<\varepsilon\leq 1$, such that H will leave the queue with probability $p\geq \varepsilon$ and keep on circulating with probability $q=1-p$.

Proof

Let B_i be a node in which none of its packets is able to leave the queue. In a finite amount of time the probability of requesting a port x at A from B_i is different to zero. This probability is equal to the probability of a hole at port x moves from node A to B_i.

On one side, if there exist X packets in B_i that will be routed through A, the probability that a packet requests port x can be expressed as $P(x) = \frac{X}{C} > 0$. On the other side, if none of the packets in node B_i request port x, in a finite amount of time some of the packets in B_i will be marked for misrouting. This packets request every link of the node, including x. Therefore, $P(x) = \frac{Miss}{C} > 0$, where Miss is the number of misrouted packets.

This situation is the same at every node of the network. Thus, every neighbor of node A will have packets requesting the hole H. Consequently, the probability that the hole at node A leaves the queue is different to zero for any position in the queue.

Lemma 2

An extra hole H present in node A, eventually will move to another neighbor node B_i.

Proof

Let us define the advance of the hole in the circular queue as a sequence of moves. By Lemma 1, at each move hole H either moves to a neighbor node with probability $p \geq \varepsilon$ or it keeps on circulating inside the queue with probability $q=1-p$.

Let the event that H does not leave the circular queue during move i be denoted by l_i, and the event that H leaves the queue during move i be denoted by w_i. Let $Q(i)$ be the probability that hole H has not leaved the circular queue after i moves. Then

$$Q(i) = P(l_i l_{i-1} \ldots l_1) = P(l_i | l_{i-1} \ldots l_1) \cdot P(l_{i-1} \ldots l_1)$$

We define $F_k = l_k \ldots l_2 l_1$ for $1 \leq k \leq i$ Then

$$Q(i) = P(l_i | F_{i-1}) \cdot P(l_{i-1} | F_{i-2}) \cdots P(l_1)$$

Clearly, we have that $P(l_j | F_{j-1}) = 1 - P(w_j | F_{j-1}),\ 1 \leq j \leq i$ and by Lemma 1 we have also that $P(w_j | F_{j-1}) \geq \varepsilon > 0$. Then

$$P(w_j | F_{j-1}) \geq \varepsilon \cdot \Rightarrow$$
$$P(l_j | F_{j-1}) \leq (1-\varepsilon) \Rightarrow$$
$$Q(i) \leq (1-\varepsilon)^i$$

Thus the probability that H will not leave the circular queue after i moves, when $i \rightarrow \infty$ is:

$$\lim_{i\to\infty} Q(i) = (1-\varepsilon)^i = 0$$

The probability $P(i)$ that the hole H will move to a neighbor node after i moves, when $i \to \infty$ is:

$$\lim_{i\to\infty} P(i) = 1$$

Theorem 1 (deadlock)

The probability of the extra hole *H* not traveling through a node *A* diminishes as the number of hops increases

Proof

Let us define the path of the hole in the network as a sequence of hops. By Lemma 2, at each hop, hole *H* either moves closer to *A* with probability $p \geq \lambda$ or further from *A* with probability $q = 1 - p$.

We define a game as a sequence of *N* moves. We denote the event that *H* did not travel through *A* during game *i* by l_i, and the event that *H* traveled through *A* during game *i* by w_i. Let *S(i)* be the probability that hole *H* has not visited *A* after *i* games. Then

$$S(i) = P(l_i l_{i-1} \ldots l_1) = P(l_i | l_{i-1} \ldots l_1) \cdot P(l_{i-1} \ldots l_1)$$

We define F_k as $l_k \ldots l_2 l_1$, $1 \leq k \leq i$. Then

$$S(i) = P(l_i | F_{i-1}) \cdot P(l_{i-1} | F_{i-2}) \cdots P(l_1)$$

Clearly $P(l_j | F_{j-1}) = 1 - P(w_j | F_{j-1})$, $1 \leq j \leq i$. In the following, we are going to determine $P(w_j | F_{j-1})$. We denote the event that hole *H* starts game *j* at distance k from *A* by $S_{j,k}$. We denote the maximum distance between two nodes of the network by D. Events $S_{j,k}$ are mutually exclusive and one of them necessarily happens. Thus,

$$P(w_j | F_{j-1}) = P(w_j S_{j,1} \cup \ldots \cup w_j S_{j,D} | F_{j-1}) \Rightarrow$$

$$P(w_j | F_{j-1}) = \sum_{k=1}^{D} P(w_j S_{j,k} | F_{j-1}) \Rightarrow$$

$$P(w_j | F_{j-1}) = \sum_{k=1}^{D} P(w_j | S_{j,k} F_{j-1}) \cdot P(S_{j,k} | F_{j-1})$$

By Lemma 2 we have that $P(w_j | S_{j,k} F_{j-1}) \geq \lambda^D$ which implies that

$$P(w_j | F_{j-1}) \geq \lambda^D \cdot \sum_{k=1}^{D} P(S_{j,k} | F_{j-1})$$

Since we have that $\cdot \sum_{k=1}^{D} P\left(S_{j,k} \middle| F_{j-1}\right) = 1$ then,

$$P\left(w_j \middle| F_{j-1}\right) \geq \lambda^D \cdot \Rightarrow$$
$$P\left(l_j \middle| F_{j-1}\right) \leq \left(1 - \lambda^D\right) \Rightarrow$$
$$S(i) \leq \left(1 - \lambda^D\right)^i$$

Thus the probability that *H* will not have visited *A* after i games, when $i \rightarrow \infty$ is:

$$\lim_{i \rightarrow \infty} S(i) = \left(1 - \lambda^D\right)^i = 0$$

The probability $T(i)$ that *H* travels through *A* after i games, when $i \rightarrow \infty$ is:

$$\lim_{i \rightarrow \infty} T(i) = 1$$

Lemma 3

There exists $0 < \eta < 1$ such that, for every message *M* in the circular queue of a node *A*, *M* will be routed with probability $p \geq \eta$ and derouted with probability $q=1-p$.

Proof

To prove the Lemma we are going to use Figure 2 as a reference. Given the situation of Figure 2, in which the network presents a hole in node *A* and a packet *M* in node *B* that needs to move to node *D* in order to approach destination, the packet will be routed to node *D* if the hole *H* is at position *x* when packet *M* is at position *y*.

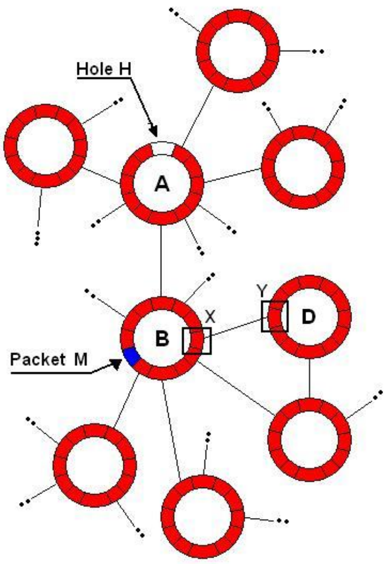

Figure 2: Generic situation of packet movement

By Theorem 1 we know that the hole is able to move through the whole network, therefore

P(H at position y)=ρ with 0< ρ <1

And since the packets move continuously in the circular queue we have that

P(M at position x)=1/C

being *C* the number of positions of the circular queue. Then, the probability of both events happening simultaneously is

*P(M routed)= P(H at position x) *P(M at position y)= ρ /C= τ 0< τ <1*

Theorem 2 (livelock)

The probability of long paths in the network diminishes as their length increases.

Proof

Let us define the path of a message in the network as a sequence of hops. By Lemma 3, at each hop, message *M* either moves closer to its destination with probability $p \geq \tau$ or further from its destination with probability $q = 1 - p$.

Let us define a game as a sequence of *N* moves. The event that *M* was not delivered during game *i* is denoted by l_i, and the event that *M* was delivered during game *i* is denoted by w_i. Let *S(i)* be the probability that message *M* has not been delivered after i games. Then

$$S(i) = P(l_i l_{i-1} \ldots l_1) = P(l_i | l_{i-1} \ldots l_1) \cdot P(l_{i-1} \ldots l_1)$$

We define F_k for $l_k \ldots l_2 l_1$, $1 \leq k \leq i$. Then

$$S(i) = P(l_i | F_{i-1}) \cdot P(l_{i-1} | F_{i-2}) \cdots P(l_1)$$

Clearly $P(l_j | F_{j-1}) = 1 - P(w_j | F_{j-1})$, $1 \leq j \leq i$. In the following we are going to determine $P(w_j | F_{j-1})$. We denote the event that message *M* starts game *j* at distance k from its destination by $S_{j,k}$ and the maximum distance between two nodes of the network by *D*. The events $S_{j,k}$ are mutually exclusive and one of them necessarily happens. Thus

$$P(w_j | F_{j-1}) = P(w_j S_{j,1} \cup \ldots \cup w_j S_{j,D} | F_{j-1}) \Rightarrow$$

$$P(w_j | F_{j-1}) = \sum_{k=1}^{D} P(w_j S_{j,k} | F_{j-1}) \Rightarrow$$

$$P(w_j | F_{j-1}) = \sum_{k=1}^{D} P(w_j | S_{j,k} F_{j-1}) \cdot P(S_{j,k} | F_{j-1})$$

We know that $P\left(w_j \middle| S_{j,k} F_{j-1}\right) \geq \tau^D$, which implies

$$P\left(w_j \middle| F_{j-1}\right) \geq \tau^D \cdot \sum_{k=1}^{D} P\left(S_{j,k} \middle| F_{j-1}\right)$$

Since we have that $\cdot \sum_{k=1}^{D} P\left(S_{j,k} \middle| F_{j-1}\right) = 1$ then,

$$P\left(w_j \middle| F_{j-1}\right) \geq \tau^D \cdot \Rightarrow$$
$$P\left(l_j \middle| F_{j-1}\right) \leq \left(1 - \tau^D\right) \Rightarrow$$
$$S(i) \leq \left(1 - \tau^D\right)^i$$

Thus, the probability that M will not have been delivered after i games, when $i \to \infty$ is:

$$\lim_{i \to \infty} S(i) = \left(1 - \tau^D\right)^i = 0$$

The probability $T(i)$ that M will be delivered after i games, when $i \to \infty$ is:

$$\lim_{i \to \infty} T(i) = 1$$