Popular Communication Operations

- One-one sending/one-all broadcasting
- All-all broadcasting

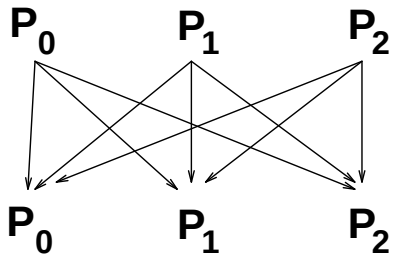

- Accumulation (or gather)

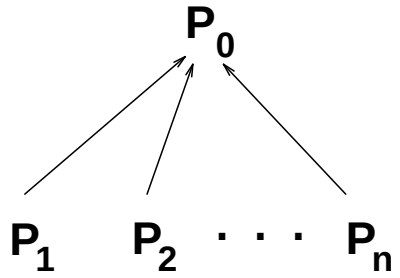

- Reduction

$$\mathrm{I}nitially: \quad V_0 \;\; V_1 \;\; \ldots \;\; V_{p-1}$$

$$\mathrm{T}hen: \quad V_0 \oplus V_1 \oplus \ldots \oplus V_{p-1}$$

to all processors.

- One-all personalized communication (single node scatter).
- All to all personalized communication.

Implementation of One-to-All Broadcasting

Store-Forward Routing

Ring

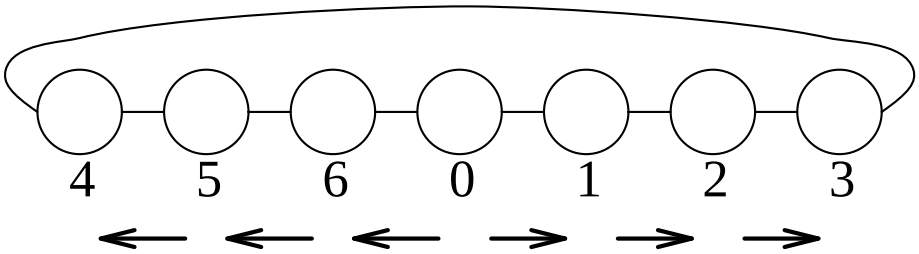

α - startup time

β - transmission speed.

$$\text{The cost is: } \frac{p}{2} \times (\alpha + \beta m)$$

Linear array: the worst cost is $p(\alpha + \beta m)$.

Broadcasting on Mesh

MESH

Stage 1

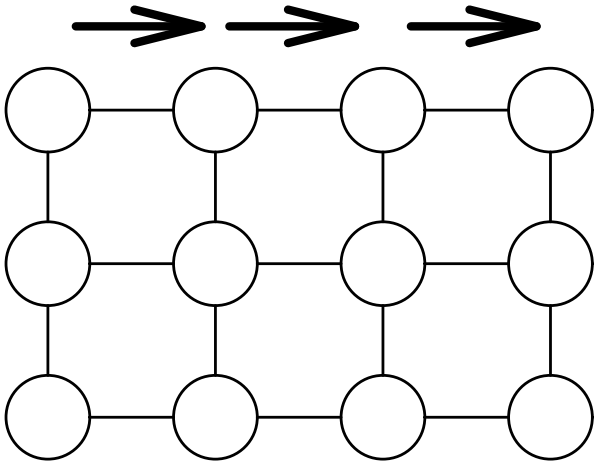

Stage 2

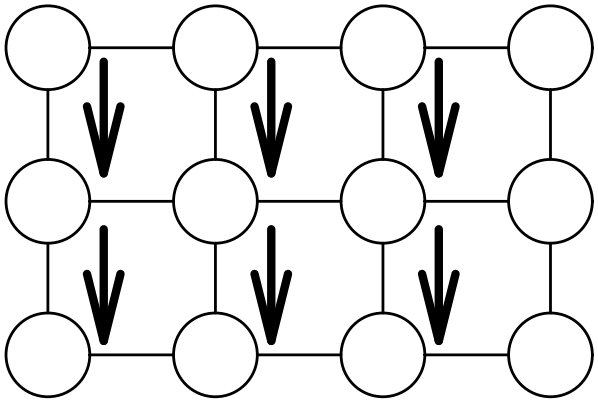

The cost

$$2\sqrt{p}(\alpha + \beta m)$$

Broadcast with wormhole routing

Wormhole routing: fast pipelined message sending between two nodes even they are not not directly connected. Node-to-node communication cost $\approx \alpha + \beta m$ and it does not depend on the node distance (i.e., the number of hops).

Ring (Linear Array)

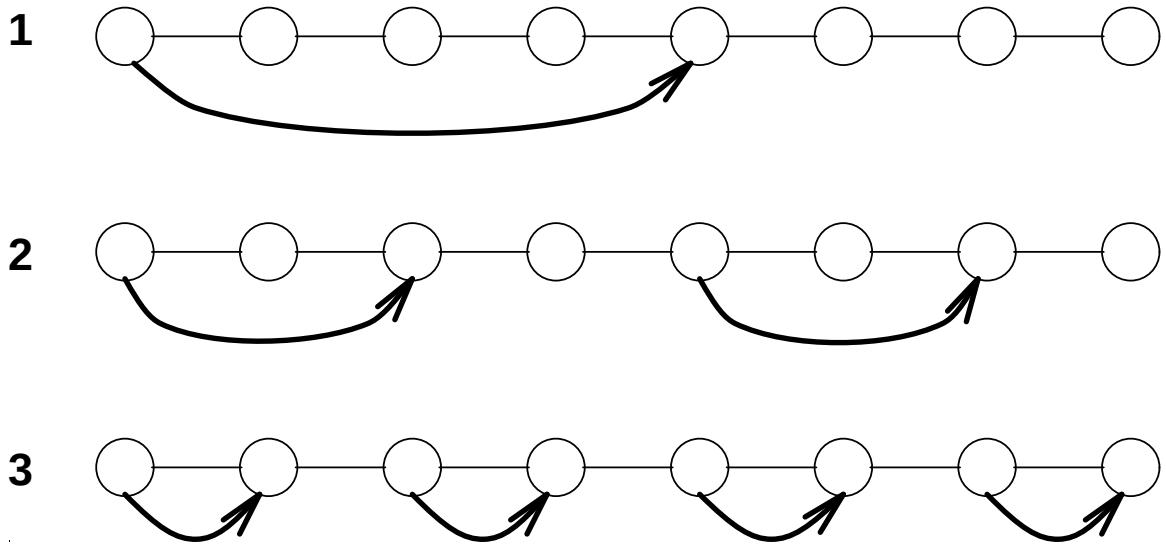

About log p steps and *there is no message pipeline contention.* Total communication cost is about $\log p(\alpha + \beta m)$.

Implementation of All-to-All Broadcast

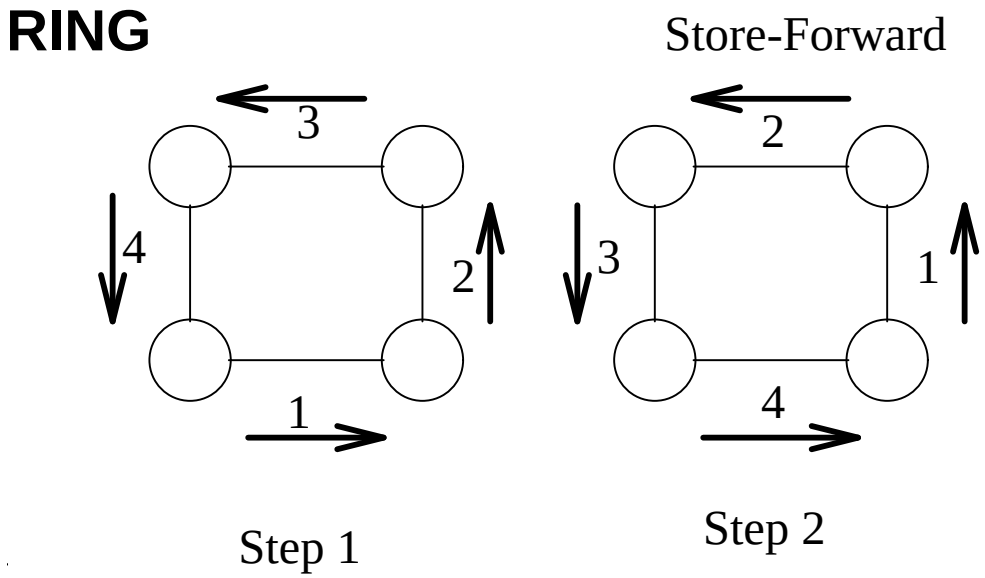

Total $p-1$ steps. Thus the cost is $(p-1)(\alpha+\beta m)$.

Using store-forward is good enough.

One-to-all personalized broadcasting

Broadcasting from the center of a linear array.

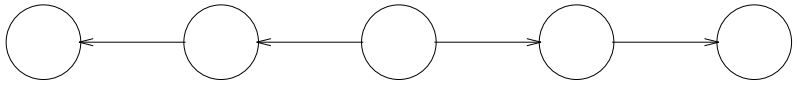

$$\begin{aligned} \alpha &+ \tfrac{p}{2} m\beta \\ \alpha &+ (\tfrac{p}{2} - 1) m\beta \\ &\ldots \\ \alpha &+ 1 m\beta \\ \texttt{Total cost} \quad &\approx \tfrac{p}{2}\alpha + \tfrac{1}{2}(\tfrac{p}{2})^2 m\beta. \end{aligned}$$

Broadcast from the left end-point of a linear array.
$(p-1)\alpha + \frac{(p-1)^2}{2} m\beta$.