Recent Development in Multicast Switching Networks

Yuanyuan Yang

Department of Electrical & Computer Engineering
State University of New York at Stony Brook
U.S.A

Overview

- Introduction and background
- Previous related results
- Nonblocking multicast network
- Routing algorithm and network controller
- Necessary nonblocking condition
- Blocking probability analysis
- Experimental simulations
- Summary

Introduction

- Multicast (one-to-many) communication
  - transmitting information from a single source to multiple destinations in a network
  - a requirement in high-performance networks
  - increasingly used to support various applications
    - ∗ audio and video multimedia conferencing
    - ∗ web servers
    - ∗ E-commerce on the Internet
    - ∗ distributed database updates
    - ∗ cache coherence protocols

Introduction

- Many multicast applications require
  - not only multicast capability
  - but also predictable communication performance: i.e. guaranteed quality-of-service (QoS)
    - guaranteed multicast latency
    - guaranteed multicast bandwidth
- The combination of the non-uniform nature of multicast traffic and the requirement of QoS guarantees makes the problem very challenging.

Introduction

- **Performance of multicast communication is mainly measured in terms of its latency in delivering a message to all destinations**

- **By far most of work aims to**
  - **minimize multicast latency**
  - **design deadlock-free multicast routing algorithms**
  - **provide best-effort services**

- **Software approach: supporting multicast in software (unicast-based multicast)**

- **Hardware approach: providing hardware support for multicast at the network level**
  - **Router-based networks**
    - **how to design a deadlock-free routing algorithm is a critical and difficult issue**
  - **Switch-based (switching) networks**
    - **easily achieve deadlock-free routing**
    - **equal communication latency between any source and destination**
    - **good candidate for a QoS capable multicast architecture**

Objective of This Work:

Design nonblocking multicast switching networks with low hardware cost and fast routing algorithm

Terminologies

- **A switching network is an $N \times M$ switch with N inputs and M outputs which provides connection paths between the input ports and output ports.**

- **A one-to-one connection is a connection of an input port to one output port. A maximal set of concurrent one-to-one connections is called a permutation assignment.**

- **A multicast connection is a connection where an input port can be simultaneously connected to more than one output port (but an output port can be connected to at most one input port at a time). A maximal set of concurrent multicast connections is called a multicast assignment.**

- **A connection request means an idle input port requests connection path(s) to idle output port(s).**

Terminologies

- **A multicast switching network is a network which can realize all possible multicast assignments.**

- **A rearrangeable network can satisfy all connection requests but sometimes requires rearranging the connection paths of existing connections.**

- **A nonblocking network can satisfy all connection requests and rearrangement is never required.**

Motivation

- **Multicast communication is a fundamental communication pattern in both telecommunication networks and scalable parallel and distributed computing systems**
- **A permutation network cannot support arbitrary multicast.**
  For an $N \times N$ network,
  Number of permutation assignments: $N!$
  Number of multicast assignments: N^N
- **Efficient implementation for multicast is critical to system performance.**
- **Support multicast at interconnection network level.**
- **Many applications require nonblocking capability.**
  - **Reduce overhead associated with rearrangements**
  - **Avoid disturbances of existing connections in the network**
  - **Especially important in real-time applications.**

Outline of this work

- Design of nonblocking multicast networks
- Routing algorithm
- Parallel network controller
- Necessary nonblocking condition
- Analytical model for blocking probability
- Experimental simulations

Definition of the Clos network or $v(m, n, r)$ network

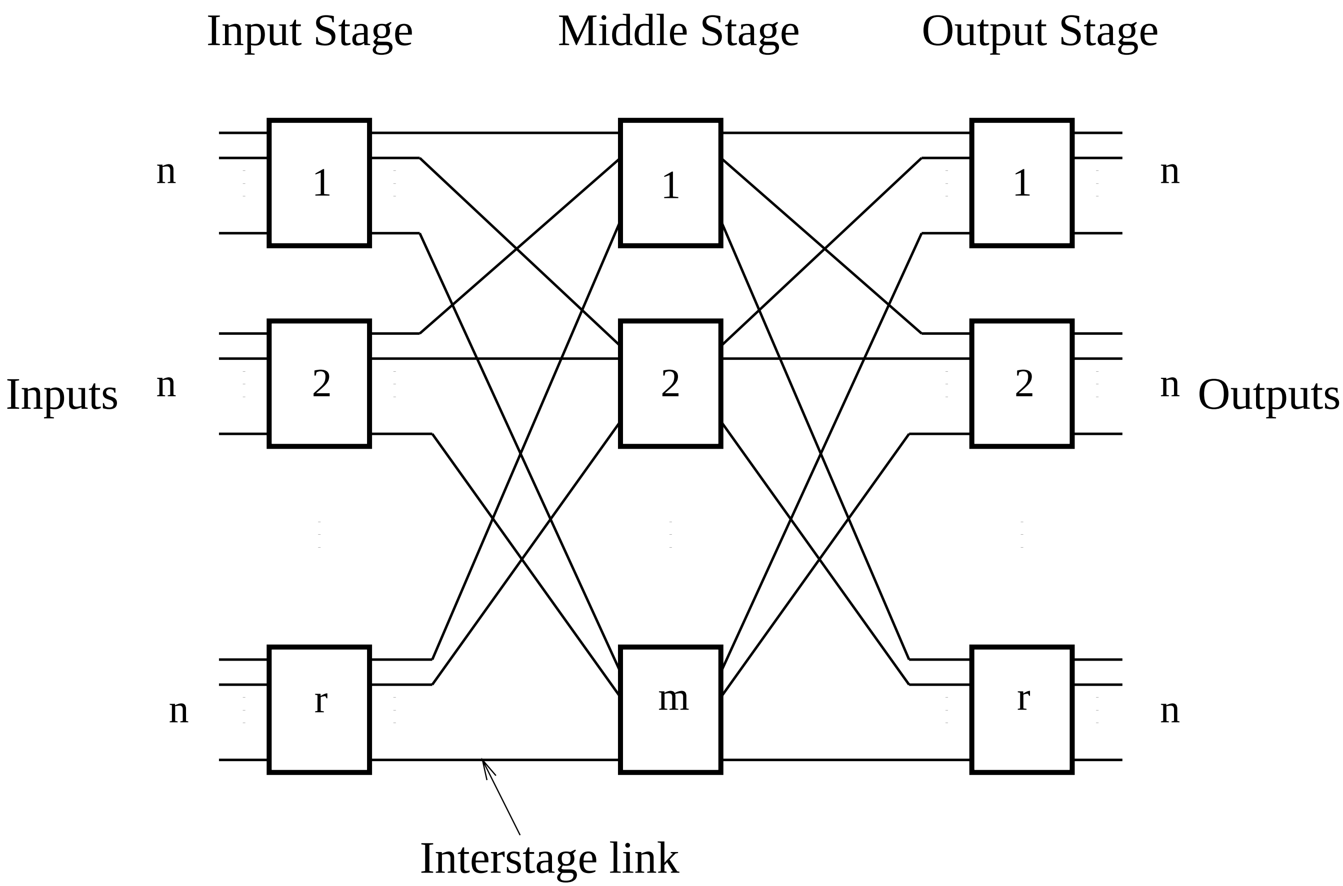

Previous results on nonblocking conditions of the Clos network

- **A $v(m, n, r)$ network is nonblocking for permutation assignments [Clos, *BSTJ*, '53], if the number of middle stage switches**

$$m \geq 2n - 1$$

- **A $v(m, n, r)$ network is nonblocking for multicast assignments [Masson, *Networks*, '72; Hwang and Jajszczyk, *IEEE Trans. Comm.*, '86], if the number of middle stage switches**

$$m \geq c(nr)$$

More terminologies

Fanout: A multicast connection from an input port to output ports on r' output switches is said to have fanout r'.

Input connection request: An input connection request I_j is the set of output switches to which input port j is to be connected.

Destination sets: Destination set M_j is the set of output switches to which the middle switch j is providing connection paths from the input ports.

Available middle switches: The available middle switches of input port j is the set of middle switches with currently unused links to the input switch associated with input j.

Design of nonblocking multicast switching networks

- Specify an "intelligent" control strategy for satisfying each multicast connection request: choose no more than a certain number of middle switches, say, x, whose destination set intersections are empty from available middle switches.
- Determine how many available middle switches can guarantee that these x middle switches can always be chosen. We are interested in finding as few as possible available middle switches with this property.
- Find the optimal value of x to minimize the number of middle switches.
- Develop an efficient control algorithm to actually find these x middle switches.
- Hardware implementation of the control algorithm to further speed up the routing process.

Nonblocking condition for multicast switching networks

Theorem 1. *We can satisfy a new connection request* $I_i, i \in \{1, 2, \ldots, nr\}$*, in a* $v(m, n, r)$ *network using some* x $(x \geq 1)$ *middle switches, say,* $j_1, \ldots, j_x$*, from among the available middle switches if and only if* I_i *and the current destination sets of these* x *middle switches satisfy*

$$I_i \cap (\bigcap_{k=1}^{x} M_{j_k}) = \phi. \tag{1}$$

Nonblocking condition for multicast switching networks

Theorem 2. *For all n', $1 \leq n' \leq n$, and for all x, $1 \leq x \leq \min\{n', r\}$, let m' be the maximum number of middle switches whose destination sets have the following properties:*

1. *there are at most n' 1's, n' 2's, ... , n' r's distributed among the destination sets;*
2. *the intersection of any x of the destination sets is nonempty.*

Then

$$m' \leq n' r^{\frac{1}{x}}$$

Proof sketch of Theorem 2

WLOG, suppose these m' middle switches are $1, 2, \ldots, m'$ with destination sets $M_1, M_2, \ldots, M_{m'}$. Clearly,

$$\sum_{i=1}^{m'} |M_i| \leq n'r$$

Let

$$c_1 = \min_i\{|M_i|\}.$$

Then

$$m' \leq \frac{n'r}{c_1} \tag{2}$$

WLOG, suppose that the destination set of middle switch 1 has cardinality c_1, and $M_1 = \{1, 2, \ldots, c_1\}$.

Proof sketch of Theorem 2

Intersect each of the destination sets M_1, M_2, . . . , $M_{m'}$ with M_1 and obtain m' sets M_1^1, M_2^1, . . . , $M_{m'}^1$ which consist of only elements in $\{1, 2, \ldots, c_1\}$, and distributed among the M_i^1 's are at most n' 1's, n' 2's, . . . , n' c_1's.
Let

$$c_2 = \min_i \{|M_i^1|\}$$

WLOG, suppose that M_2^1 has cardinality c_2, and $M_2^1 = \{1, 2, \ldots, c_2\}$.
Then

$$m' \leq \frac{n' c_1}{c_2} \tag{3}$$

Then intersect each of $M_1^1, M_2^1, \ldots, M_{m'}^1$ with M_2^1 and obtain m' sets $M_1^2, M_2^2, \ldots, M_{m'}^2$, which consist of only elements in $\{1, 2, \ldots, c_2\}$.

Proof sketch of Theorem 2

After repeating the above process $x-1$ times, we have

$$m' \leq \textbf{min}\{\frac{n'r}{c_1}, \frac{n'c_1}{c_2}, \ldots, \frac{n'c_{x-2}}{c_{x-1}}\} \quad (4)$$

We now have a set of m' intersected destination sets M_1^{x-1}, M_2^{x-1}, $\ldots$, $M_{m'}^{x-1}$. Moreover, each M_k^{x-1} consists of only elements in $\{1, 2, \ldots,$ $c_{x-1}\}$, and there are at most n' 1's, n' 2's, $\ldots$, and n' c_{x-1}'s distributed among these m' sets. Thus, we have

$$m' \leq n'c_{x-1} \quad (5)$$

Therefore, from (3) and (4),

$$m' \leq \textbf{min}\{\frac{n'r}{c_1}, \frac{n'c_1}{c_2}, \ldots, \frac{n'c_{x-2}}{c_{x-1}}, n'c_{x-1}\}$$

It can be shown that

$$\max_{c_1, c_2, \cdots, c_{x-1}} \textbf{min}\{\frac{n'r}{c_1}, \frac{n'c_1}{c_2}, \ldots, \frac{n'c_{x-2}}{c_{x-1}}, n'c_{x-1}\} = n'r^{\frac{1}{x}}$$

Therefore, we have $m' \leq n'r^{\frac{1}{x}}$

Nonblocking condition for multicast switching networks

Theorem 3. *A* $v(m,n,r)$ *network is nonblocking for multicast assignments if*

$$m > \min_{1 \le x \le \min\{n-1,r\}} \{(n-1)(x + r^{\frac{1}{x}})\}$$

Nonblocking condition for multicast switching networks

Curve of function $y = x + r^{\frac{1}{x}}$

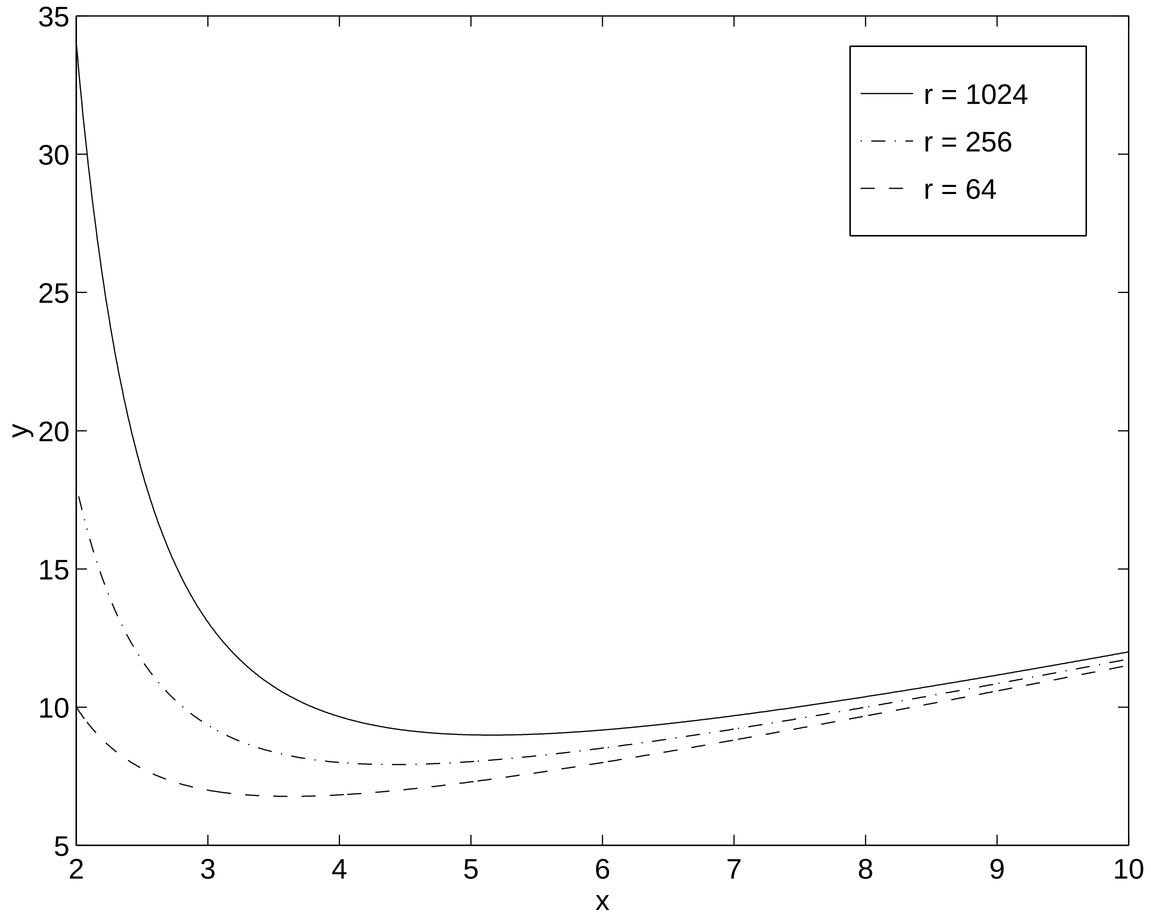

Relationship between r and coefficient of m, $\min_{1 \le x \le \min\{n-1, r\}}(x + r^{\frac{1}{x}})$

r	x	$x + r^{\frac{1}{x}}$
1	**1**	**2**
2	**1**	**3**
$4 = 2^2$	**2**	**4**
$9 = 3^2$	**2**	**5**
$27 = 3^3$	**3**	**6**
$81 = 3^4$	**4**	**7**
$256 = 4^4$	**4**	**8**
$1024 = 4^5$	**5**	**9**
$4096 = 4^6$	**6**	**10**
$16384 = 4^7$	**7**	**11**
$78125 = 5^7$	**7**	**12**
$390625 = 5^8$	**8**	**13**
$1953125 = 5^9$	**9**	**14**
$10077696 = 6^9$	**9**	**15**

Nonblocking condition for multicast switching networks

A bound on m as a function of n and r

Theorem 4. *A $v(m, n, r)$ network is nonblocking for multicast assignments if*

$$m \geq 3(n-1)\frac{\log r}{\log \log r}$$

Generalization to restricted multicast assignments:

Corollary 1. *A* $v(m,n,r)$ *network is nonblocking for restricted multicast assignments, in which each input port can be connected to at most* d $(1 \le d < r)$ *output switches, if*

$$m > \min_{1 \le x \le \min\{n-1,d\}} \{(n-1)(x + d^{\frac{1}{x}})\}$$

In particular, we have

$$m > (n-1)\left(\frac{2\log d}{\log\log d} + (\log d)^{\frac{1}{2}}\right)$$

Permutation is a special case:

Corollary 2. *Setting* $d = 1$ *in Corollary 1 yields* $m \ge 2n-1$*, which is the bound on* m *associated with the classical Clos nonblocking permutation networks.*

Generalization to $(2k+1)$-stage networks ($k > 1$)

Recursively applying the design criteria on each middle stage switch.

Theorem 5. *For each fixed integer $k \geq 1$, the minimum number of cross-points of our $(2k+1)$ stage $(N \times N)$ multicast network is*

$$O\left(N^{1+\frac{1}{k+1}}(\log N/\log\log N)^{\frac{k+2}{2}-\frac{1}{k+1}}\right)$$

Network cost comparison

- **Constructive multicast networks:**
  - **Masson's three-stage network:** $O(N^{\frac{5}{3}})$
  - **Hwang and Jajszczyk's three-stage network:** $O(N^{\frac{5}{3}})$
  - **Feldman, Friedman and Pippenger's two-stage network:** $O(N^{\frac{5}{3}})$**; three-stage network:** $O(N^{\frac{11}{7}})$
  - **Three-stage version of the new design:** $O(N^{\frac{3}{2}}(\frac{\log N}{\log\log N}))$

- **Nonconstructive multicast networks:**
  - **Feldman, Friedman and Pippenger** k**-stage network:** $O\left(N^{1+\frac{1}{k}}(\log N)^{1-\frac{1}{k}}\right)$**.** **The routing control algorithm is NP-complete.**

A linear routing control algorithm

- **Given a** $v(m, n, r)$ **network satisfying the nonblocking condition on** m **in Theorem 3.**

- **Some** x**,** $1 \leq x \leq min\{n-1, r\}$

- **A connection request** I_i **with** $|I_i| = r' \leq r$

- $m' = (n-1)r'^{\frac{1}{x}} + 1$ **available middle switches with destination sets** $M_1, M_2, \ldots, M_{m'}$**.**

A linear routing control algorithm:

Step 1: $mid\_switch \leftarrow \phi$;
for $j = 1$ **to** m' **do**
$S_j \leftarrow M_j \bigcap I_i$;

Step 2: repeat
find S_k **(**$1 \leq k \leq m'$**) such that**
$|S_k| =$ **min**$\{|S_1|, |S_2|, \ldots, |S_{m'}|\}$;
$min\_set \leftarrow S_k$;
$mid\_switch \leftarrow mid\_switch \bigcup \{k\}$;
if $min\_set \neq \phi$ **then**
for $j = 1$ **to** m' **do**
$S_j \leftarrow S_j \bigcap min\_set$;
until $min\_set = \phi$;

Step 3: connect I_i **through the middle switches in** mid_switch **and update the destination sets of these middle switches.**

End

Hardware implementation of the control algorithm

- **Overview structure**

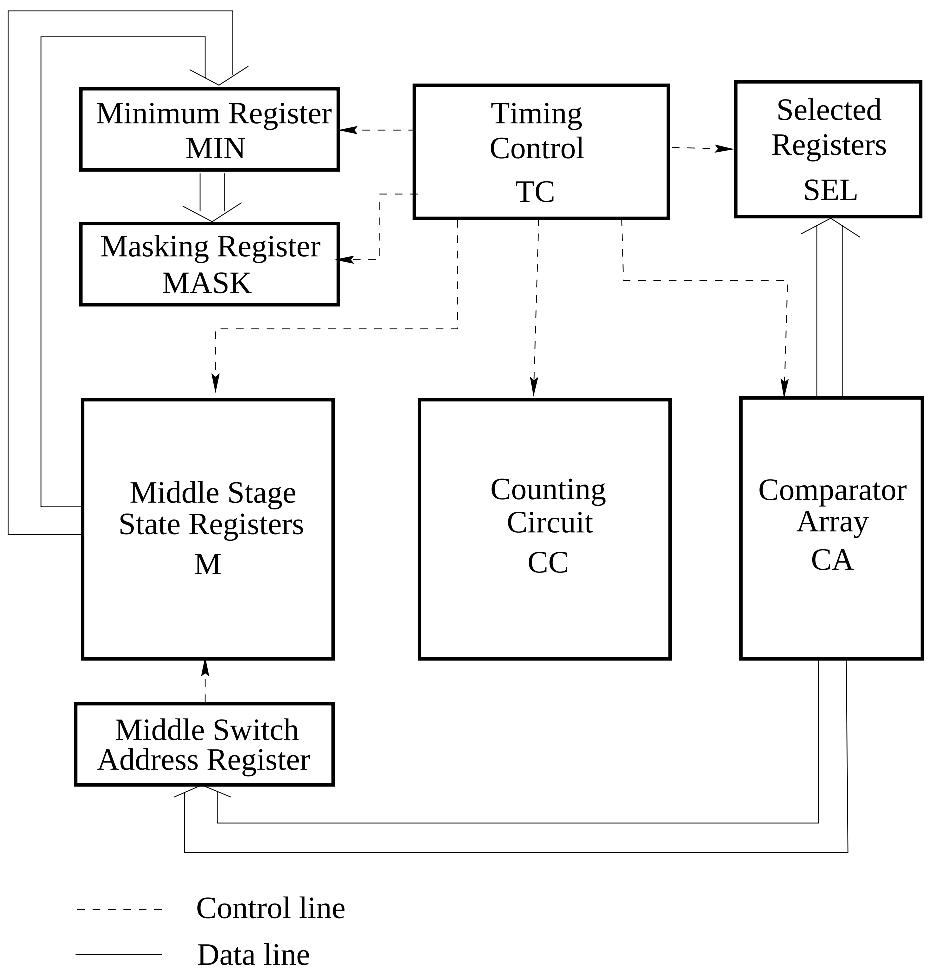

Hardware implementation of the control algorithm

- **Sequential implementation**
  - **Sequentially evaluate the cardinality of each destination set of middle switches and find the minimum cardinality set.**

- **Parallel implementation**
  - **Parallel evaluation of the cardinalities of all destination sets of middle switches by using a sequential circuit or a combinational circuit.**

Hardware implementation of the control algorithm

- **Summary of the various designs of the controller**

Design	Gates	Clocks	Gate Delays
Seq./ Counter	$O(\log r)$	$O\left(\frac{mr \log r}{\log \log r}\right)$	—
Seq./ Adder	$O(r)$	$O\left(\frac{m \log r}{\log \log r}\right)$	—
Parall./ Counter	$O(m \log r)$	$O\left(\frac{r \log r}{\log \log r}\right)$	—
Parall./ Adder	$O(mr)$	$O\left(\frac{\log r}{\log \log r}\right)$	$O\left(\frac{(\log r)^2}{\log \log r}\right)$

Necessary nonblocking condition

- **Can the sufficient condition we obtained be further reduced?**
- **What is the optimal design for this type of multicast network?**
- **Derive necessary conditions for supporting arbitrary multicast assignment under different control strategies by constructing worst case network states which force us to use a certain number of middle stage switches.**
- **Employ these necessary conditions to guide the design process of multicast networks.**

Necessary nonblocking condition

Routing control strategies used to derive necessary condition:

Strategy 1. *For each input connection request,* I_i*,* $i \in \{1, 2, \ldots, nr\}$*, in the network, always choose the middle switch with the minimum cardinality of destination sets with regard to the unsatisfied portion of* I_i *from available middle switches, until* I_i *is satisfied, that is, until all middle switches chosen satisfy condition* $I_i \cap (\bigcap_{k=1}^{x} M_{j_k}) = \phi$.

Note: This is the strategy used by the routing control algorithm we discussed earlier.

Necessary nonblocking condition

Strategy 2. *For each input connection request* I_i, $i \in \{1, 2, \ldots, nr\}$, *choose the minimum number of middle switches that satisfy condition* $I_i \cap (\bigcap_{k=1}^{x} M_{j_k}) = \phi$ *for the current network state from the available middle switches.*

Strategy 3. *For each input connection request* I_i, $i \in \{1, 2, \ldots, nr\}$, *use an empty available middle switch (i.e., middle switch with no connections) only when there is no any subset of non-empty available middle switches can satisfy condition* $I_i \cap (\bigcap_{k=1}^{x} M_{j_k}) = \phi$.

Necessary nonblocking condition

A fundamental lemma for constructing worst case network states:

Lemma 1. *For sufficiently large* n, r *and* $\frac{m}{n}(m > n)$, *there exist* $m + n$ *subsets of set* $\{1, 2, \ldots, r\}$, $\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_n, \mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_m$, *which satisfy the following conditions:*

1. *the flattened set of* $\{\mathcal{I}_1, \mathcal{I}_2, \ldots, \mathcal{I}_n, \mathcal{M}_1, \mathcal{M}_2, \ldots, \mathcal{M}_m\}$ *is a multi-set chosen from set* $\{1, 2, \ldots, r\}$ *with multiplicity of each element no more than* n*;*

2. *for some* $x = \Theta\left(\frac{\log r}{\log m - \log n}\right)$ *and for any* $\mathcal{I}_i$ $(1 \leq i \leq n)$ *and any* $\mathcal{M}_{j_1}, \mathcal{M}_{j_2}, \ldots, \mathcal{M}_{j_x}$ $(1 \leq j_1 < j_2 < \cdots < j_x \leq m)$ $\mathcal{I}_i \cap (\bigcap_{k=1}^{x} \mathcal{M}_{j_k}) \neq \phi$.

Necessary nonblocking condition

Theorem 6. *The necessary condition for a* $v(m,n,r)$ *network to be strictly nonblocking for multicast assignments is* $m \geq \Theta\left(n\frac{\log r}{\log\log r}\right)$.

Theorem 7. *The necessary condition for a* $v(m,n,r)$ *multicast network to be nonblocking under Strategy 1 is* $m \geq \Theta\left(n\frac{\log r}{\log\log r}\right)$.

Theorem 8. *The necessary condition for a* $v(m,n,r)$ *multicast network to be nonblocking under Strategy 2 is* $m \geq \Theta\left(n\frac{\log r}{\log\log r}\right)$.

Theorem 9. *The necessary condition for a* $v(m,n,r)$ *multicast network to be nonblocking under Strategy 3 is* $m \geq \Theta\left(n\frac{\log r}{\log\log r}\right)$.

Necessary nonblocking condition

Conjecture 1. *The necessary condition for a $v(m,n,r)$ multicast network to be nonblocking under any strategy is $m \geq \Theta\left(n\frac{\log r}{\log\log r}\right)$.*

Based on this conjecture, $m = O\left(n\frac{\log r}{\log\log r}\right)$ is optimal for nonblocking $v(m,n,r)$ multicast network.

Blocking probability analysis

Analytical model for the blocking probability of the networks with smaller m

- **The necessary and sufficient nonblocking condition we obtained suggests that there is little room for further improvement on the multicast nonblocking condition.**
- **What is the blocking behavior of the multicast network with smaller number of middle stage switches?**
  For example, a network with only the same number of middle stage switches as a nonblocking permutation network, i.e.
  $m = 2n - 1$**.**
- **Develop an analytical model for the blocking probability of $v(m, n, r)$ multicast network.**
- **Look into the blocking behavior of the networks under various routing control strategies through simulations to validate our model.**

Lee's model for permutation Clos network

- **The m paths between a given input/output pair:**

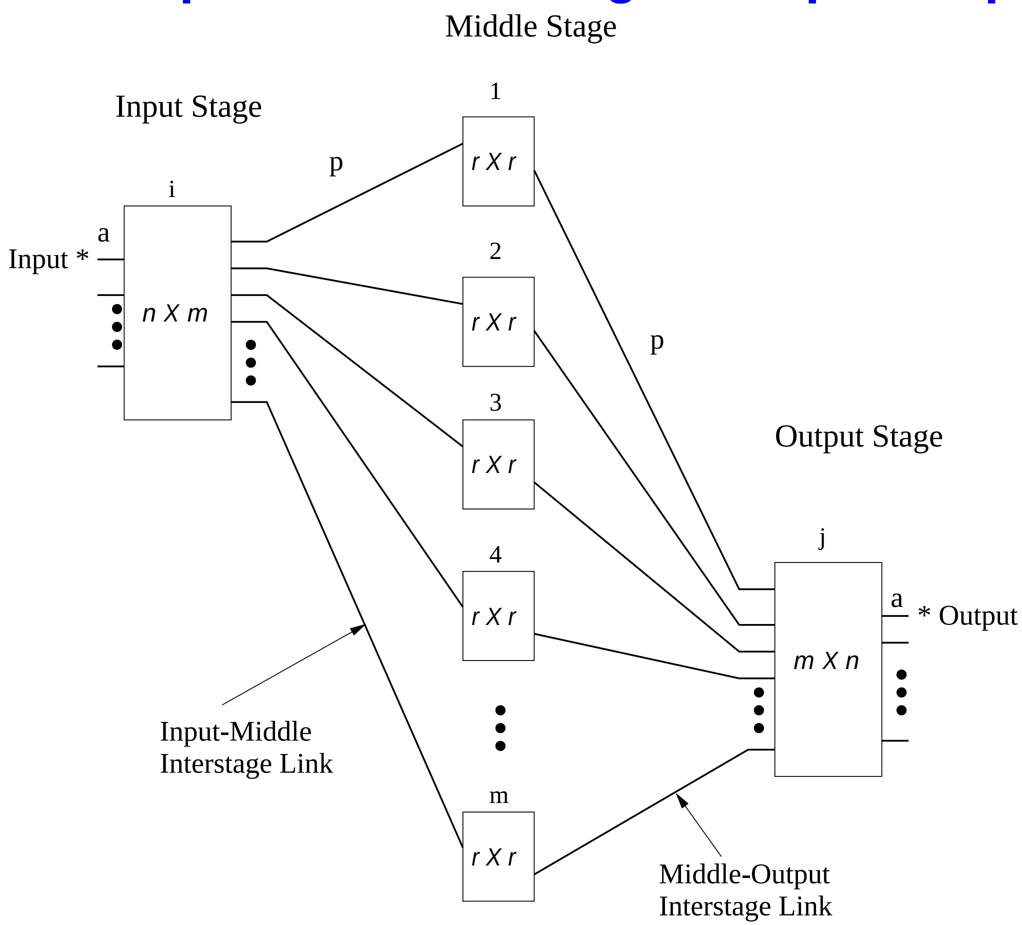

- **a: the probability that a typical input (or output) link is busy**
- **p: the probability that an interstage link is busy.**

Lee's model for permutation Clos network

- **Random routing strategy:**
  assume incoming traffic is uniformly distributed over m inter-stage links and the events that individual links in the networks are busy are independent.

- **The probability that an interstage link is busy is $p = \frac{an}{m}$.**

- **The probability that an interstage link is idle is $q = 1 - p$.**

- **The probability that a path (consisting of two interstage links) cannot be used for a connection is $1 - q^2$.**

- **The blocking probability (i.e., all m paths cannot be used)**

$$P_B = [1 - q^2]^m$$

Apply Lee's model to multicast communication

- **Different ways to realize a multicast connection with fanout f.**

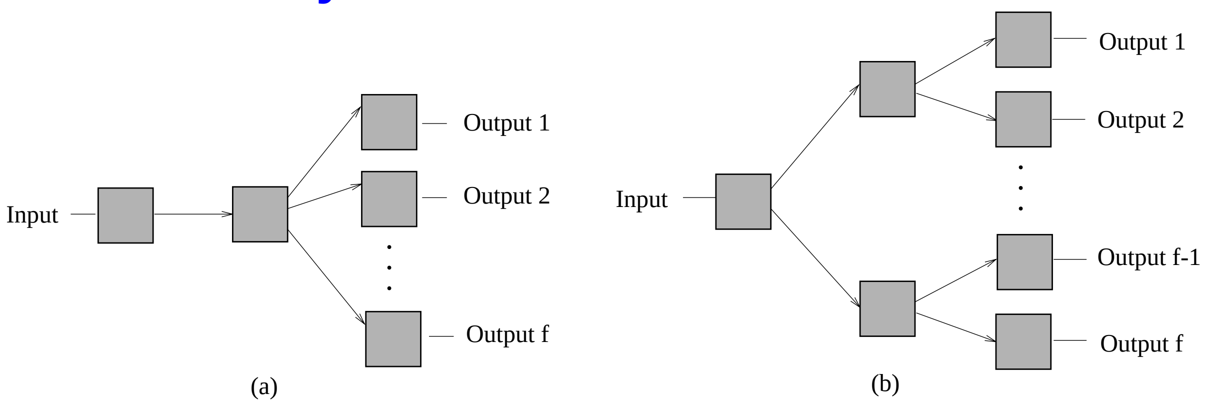

- **The total number of ways to realize a multicast connection with fanout f $(1 \le f \le r)$ is**

$$\sum_{j=1}^{f} \binom{m}{j} S(f,j) j!,$$

where $S(f,j)$ is the Stirling number of the second kind.

- **The dependencies among multicast trees make the problem intractable.**

New analytical model for multicast communication

Consider a subnetwork associated with a multicast connection with fanout f, where k input-middle interstage links are idle.

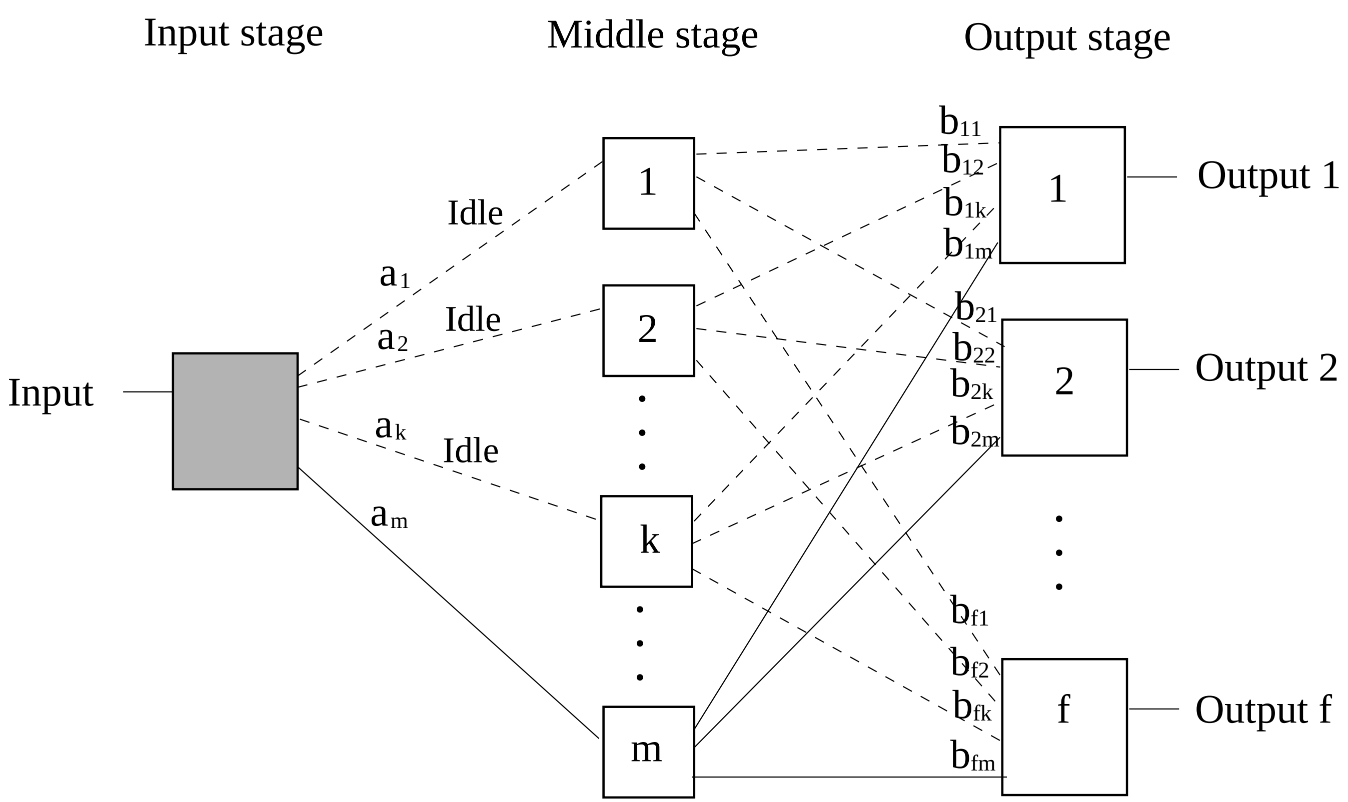

New analytical model for multicast communication

Notations and assumptions

- a_i: **the event that the input-middle interstage link** a_i **is busy.**
- b_{ij}: **the event that the middle-output interstage link** b_{ij} **is busy.**
- ε: **the event that the connection request with fanout** f **cannot be realized.**
- σ: **the state of the input-middle interstage links** $a_1, a_2, \ldots, a_m$.
- $P(\varepsilon|\sigma)$: **the conditional blocking probability in this state.**
- $P(\sigma)$: **the probability of being in state** σ.

$$P(\sigma) = q^k p^{m-k}$$

- **Still follow Lee's assumption that the events that individual links are busy are independent.**

New analytical model for multicast communication

Blocking probability for a multicast connection with fanout f

$$P_B(f) = P(\varepsilon) = \sum_{\sigma} P(\sigma) P(\varepsilon|\sigma)$$

$$= \sum_{k=0}^{m} \binom{m}{k} q^k p^{m-k} P(\varepsilon|\bar{\mathbf{a}}_1, \ldots, \bar{\mathbf{a}}_k, \mathbf{a}_{k+1}, \ldots, \mathbf{a}_m)$$

New analytical model for multicast communication

Blocking property of the subnetwork

Lemma 2. *Assume that the interstage links* $a_1, a_2, \ldots, a_k$ *in the subnetwork are idle. A multicast connection from an input of the input switch to* f *distinct output switches cannot be realized if and only if there exists an output switch whose first* k *inputs are busy.*

New analytical model for multicast communication

- **Let ε' be the event that the connection request with fanout f cannot be realized given links $a_1, a_2, \ldots, a_k$ are idle.**

$$P(\varepsilon') = P(\varepsilon|\bar{\mathbf{a}}_1, \ldots, \bar{\mathbf{a}}_k, \mathbf{a}_{k+1}, \ldots, \mathbf{a}_m).$$

- **From Lemma 2, event ε' can be expressed in terms of events $\mathbf{b}_{ij}$'s:**

$$\begin{aligned} \varepsilon' &= (\mathbf{b}_{11} \cap \mathbf{b}_{12} \cap \cdots \cap \mathbf{b}_{1k}) \\ &\quad \cup (\mathbf{b}_{21} \cap \mathbf{b}_{22} \cap \cdots \cap \mathbf{b}_{2k}) \cup \cdots \\ &\quad \cup (\mathbf{b}_{f1} \cap \mathbf{b}_{f2} \cap \cdots \cap \mathbf{b}_{fk}). \end{aligned}$$

- **The probability of event ε'**

$$\begin{aligned} P(\varepsilon') &= 1 - \prod_{i=1}^{f} P(\overline{\mathbf{b}_{i1} \cap \mathbf{b}_{i2} \cap \cdots \cap \mathbf{b}_{ik}}) \\ &= 1 - \prod_{i=1}^{f} (1 - p^k) = 1 - (1 - p^k)^f \end{aligned}$$

New analytical model for multicast communication

Blocking probability for a multicast connection with fanout f

$$P_B(f) = \sum_{k=0}^{m} \binom{m}{k} q^k p^{m-k} [1 - (1 - p^k)^f].$$

Unicast special case ($f = 1$):

$$\begin{aligned} P_B(1) &= \sum_{k=0}^{m} \binom{m}{k} q^k p^{m-k} [1 - (1 - p^k)] \\ &= (1 - q^2)^m. \end{aligned}$$

This is exactly Lee's blocking probability for the $v(m, n, r)$ permutation network.

Blocking probabilities for $v(m, 32, 32)$ network with fanouts between 1 and 32:

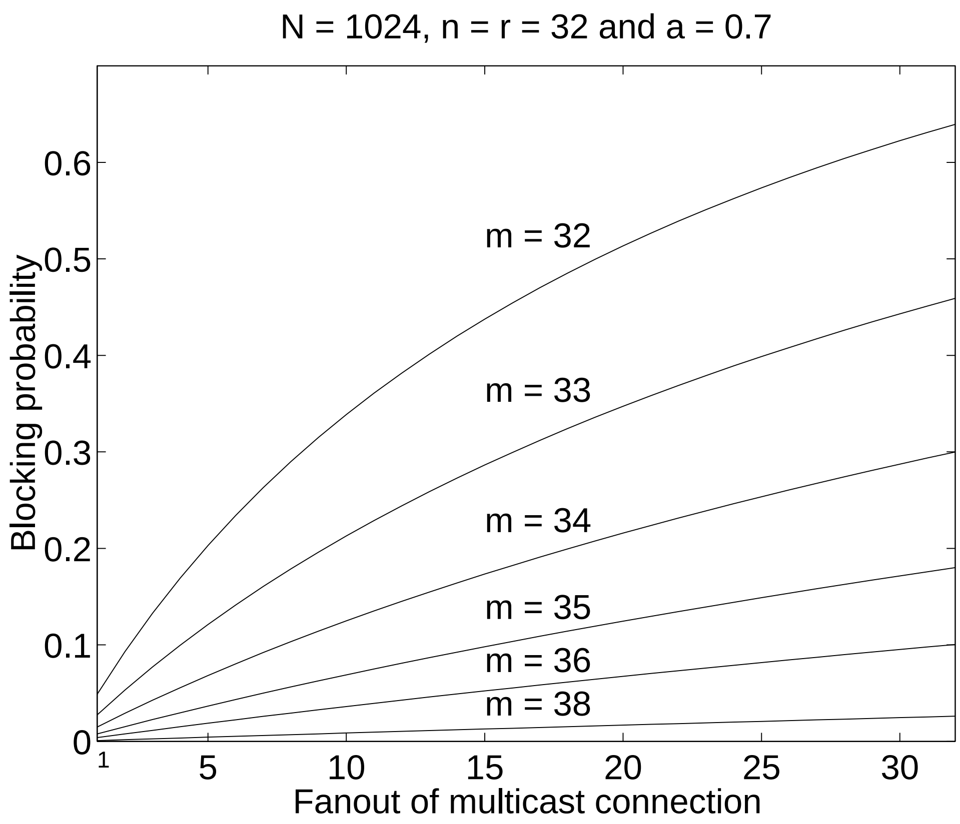

The blocking probability $P_B(f)$ is an increasing sequence of fanout f.

Average blocking probability over all fanouts

- **Suppose the fanout of a multicast connection is uniformly distributed over 1 to r. The average value of the blocking probability, simply referred to as the** *blocking probability of the $v(m, n, r)$ multicast network***:**

$$P_B \;=\; \frac{1}{r}\sum_{f=1}^{r}\sum_{k=0}^{m}\binom{m}{k}q^k p^{m-k}[1-(1-p^k)^f].$$

Average blocking probability over all fanouts

Asymptotic bound on the blocking probability

When $m = n + c$ or $m = dn$, where c and d are some constants, if $r = O(n)$ we can obtain

$$P_B = O(e^{-\epsilon n})$$

where ϵ is a constant > 0.

This suggests that the blocking probability tends to zero very quickly as n increases.

Experimental simulations

- **Routing control strategies used in the simulation**

1. **Smallest Absolute Cardinality Strategy**
2. **Largest Absolute Cardinality Strategy**
3. **Average Absolute Cardinality Strategy**
4. **Smallest Relative Cardinality Strategy**
5. **Largest Relative Cardinality Strategy**
6. **Average Relative Cardinality Strategy**
7. **Random Strategy**

Experimental simulations

- **Two network configurations considered:**
  $N = 1024, n = r = 32$**, and** $32 \leq m \leq 48$**.**
  $N = 4096, n = r = 64$**, and** $64 \leq m \leq 88$**.**

- **Seven routing control strategies**

- **Three types of traffic: uniform, uniform/constant, and Poisson**

- **Initial network utilization** $= 90\%$

- $25,000$ **connection requests processed per configuration per s-trategy**

Simulation results

The blocking probability of the $v(m, n, r)$ multicast network under seven routing control strategies:

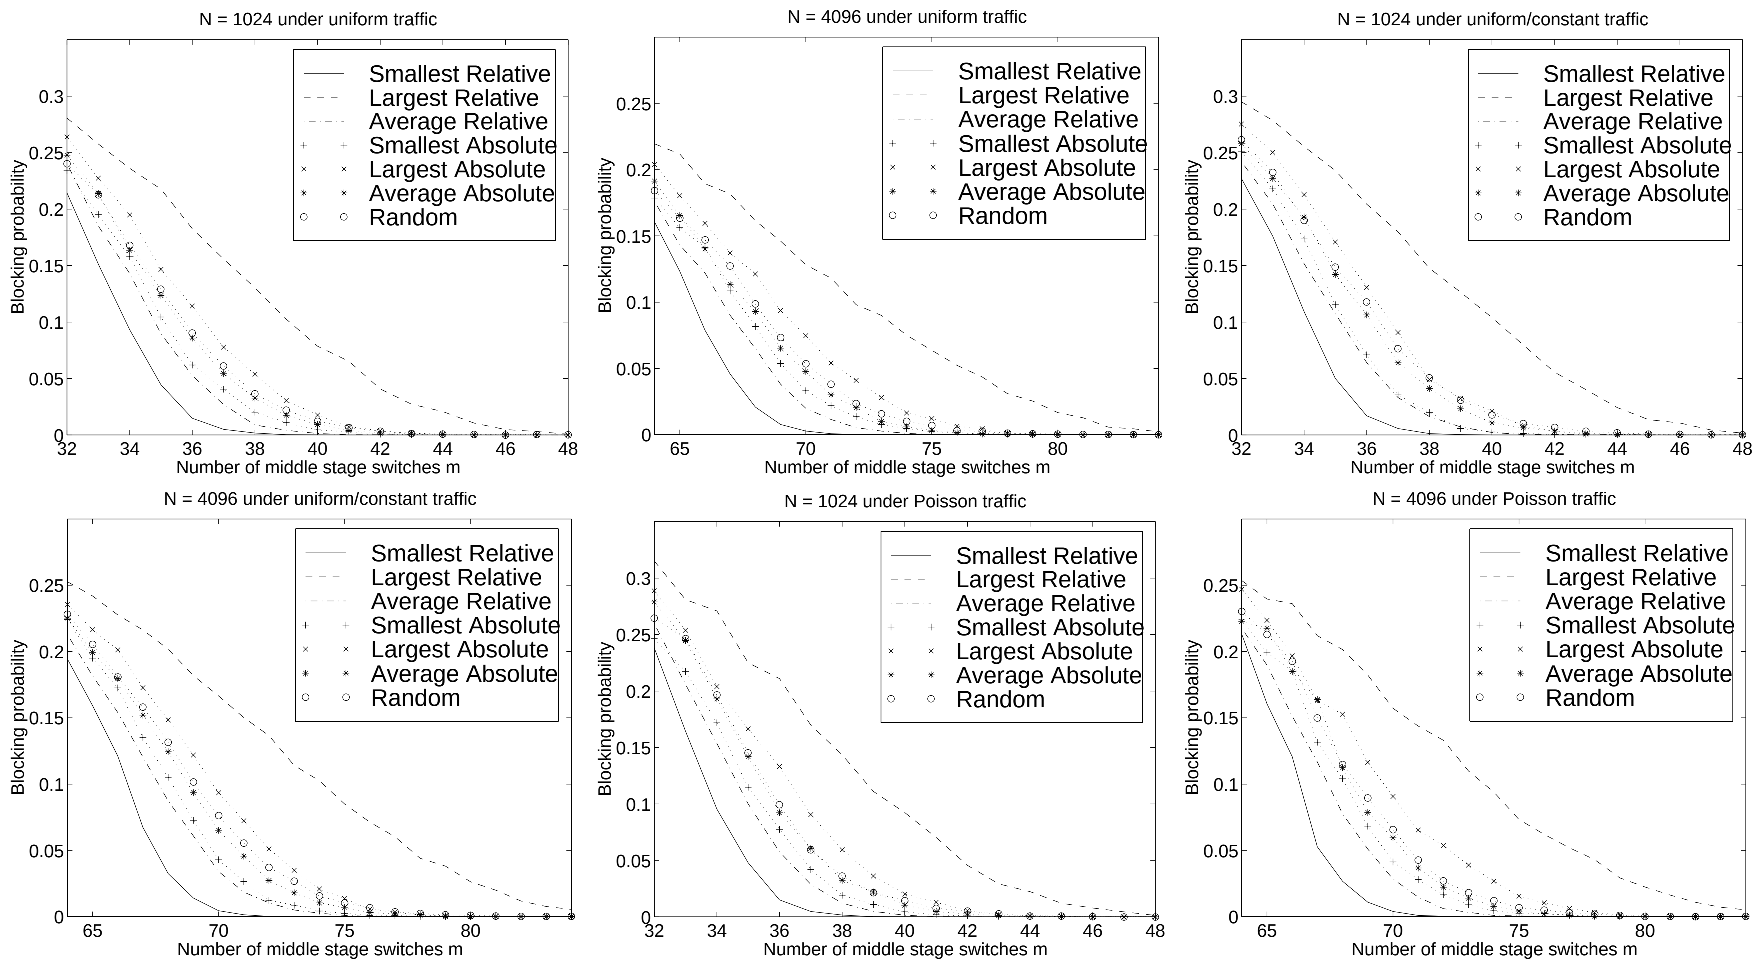

Comparison between the analytical model and the simulation results

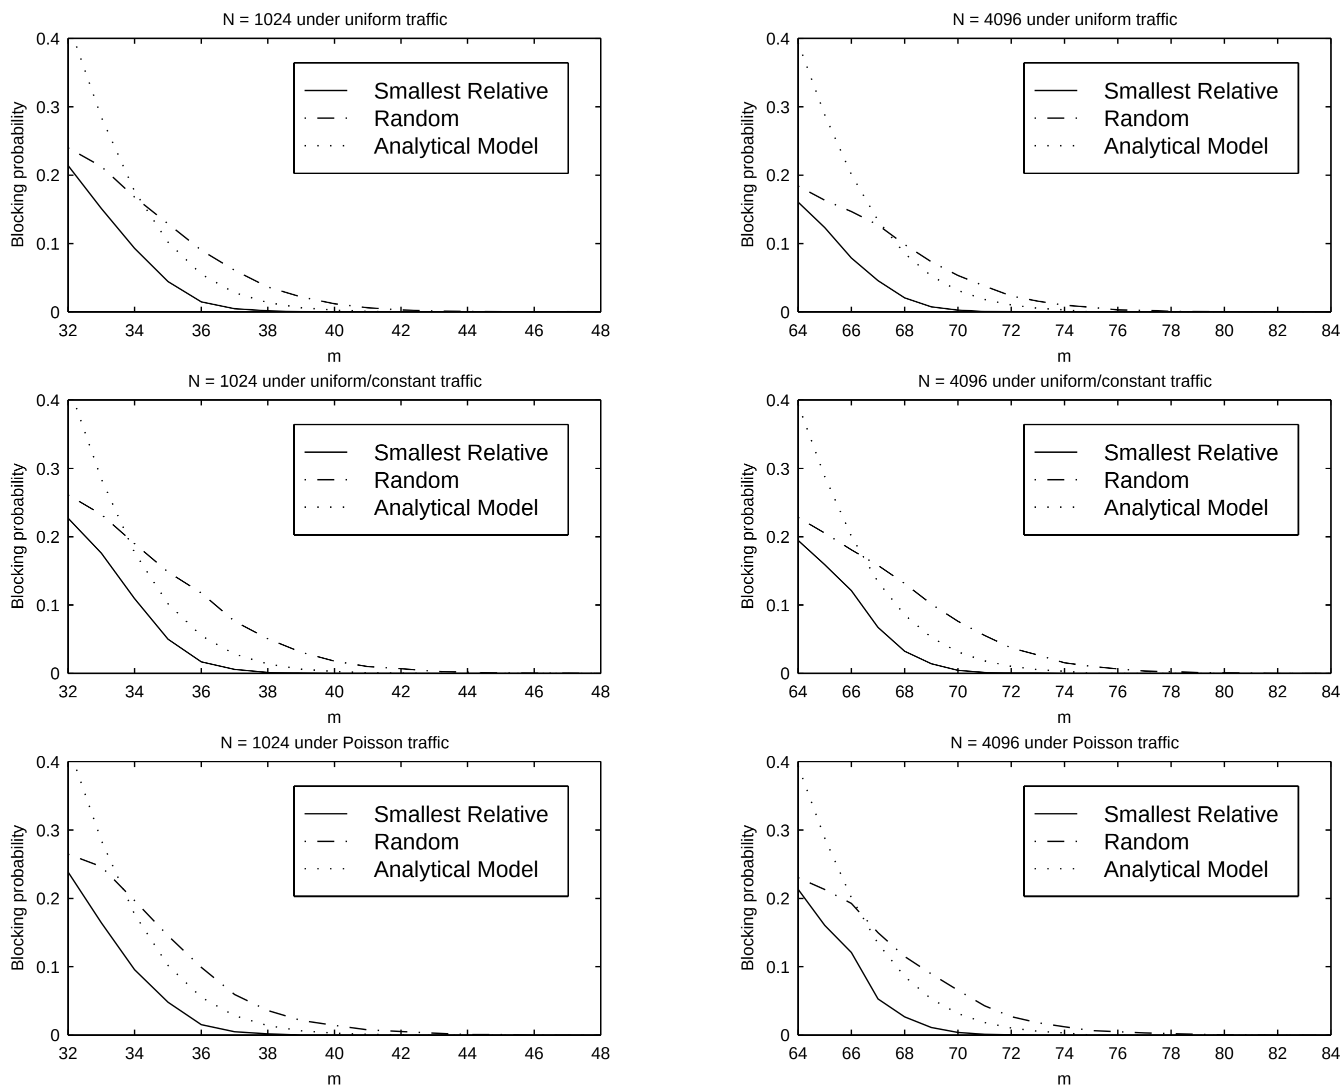

Summary

Designed currently best available explicit construction of nonblocking multicast networks.

- **Reduced the number of middle switches from $O(nr)$ to $O(n\frac{\log r}{\log\log r})$.**

- **Provided a linear time network control algorithm for satisfying connection requests.**

- **The hardware implementations of the controller provide fast path routings and require only a small amount of hardware compared with the switching hardware.**

- **Derived necessary conditions for the nonblocking multicast networks under several typical control strategies:**

$$m \geq \Theta\left(n \frac{\log r}{\log \log r}\right)$$

- **The necessary conditions obtained match the sufficient nonblocking condition under Strategy 1.**

- **Proposed an analytical model for the blocking probability of the multicast networks.**

- **Conducted extensive simulations to validate the model.**

- **The analytical and simulation results indicate that a network with a small m, such as $m = n + c$ or dn, is almost nonblocking for multicast connections.**

Related publications

1. Y. Yang and G.M. Masson, "Nonblocking broadcast switching networks," *IEEE Transactions on Computers*, vol. 40, no. 9, pp. 1005-1015, 1991.

2. Y. Yang and G.M. Masson, "The necessary conditions for Clos-type nonblocking multicast networks," *IEEE Transactions on Computers*, vol. 48, no. 11, pp. 1214-1227, November 1999.

3. Y. Yang and J. Wang, "On blocking probability of multicast networks," *IEEE Transactions on Communications*, vol. 46, no. 7, pp. 957-968, July 1998.

4. Y. Yang and J. Wang, "A more accurate analytical model on blocking probability of multicast networks," *IEEE Transactions on Communications*, vol. 48, no. 11, pp. 1930-1936, November 2000.

5. Y. Yang and G.M. Masson, “A fast network controller for non-blocking multicast networks,” *International Journal of Parallel and Distributed Systems and Networks*, vol. 1, no. 3, 1998, pp. 149-156.

6. Y. Yang, “A class of interconnection networks for multicasting,” *IEEE Transactions on Computers*, vol. 47, no. 8, pp. 899-906, August 1998.

7. Y. Yang, “An analysis model on nonblocking multirate broadcast networks,” *Proc. of the 8th ACM International Conference on Supercomputing (ICS ’94)*, pp. 256-263, 1994.

8. Y. Yang and G.M. Masson, “Non-blocking Broadcast Network,” United States Patent Number 5,451,936, issued 1995.

9. Y. Yang and G.M. Masson, “Controller for a Non-Blocking Broadcast Network,” United States Patent Number 5,801,641, issued September 1998.