Creating The Internet Technology

Leonard Kleinrock

Professor, UCLA Computer Science Dept

Founder & Chairman, Nomadix Inc

SIGCOMM Tutorial

August 31, 1999

My Early Years at MIT

by

Leonard Kleinrock

- **1959 Decided to pursue PhD, but decided NOT to work in Coding Theory, but rather set out to uncover the principles of data networks**
- **1961 Published PhD Proposal : 1st paper on modern data networking**
- **1962 Filed PhD Dissertation; MIT + McGraw-Hill decide to publish it as a book**
- **1963 Joined UCLA faculty**
- **1960's Telecom industry could care less!**
- **1966 ARPA gets interested**
- **1969+ The network locomotive starts its wild ride**

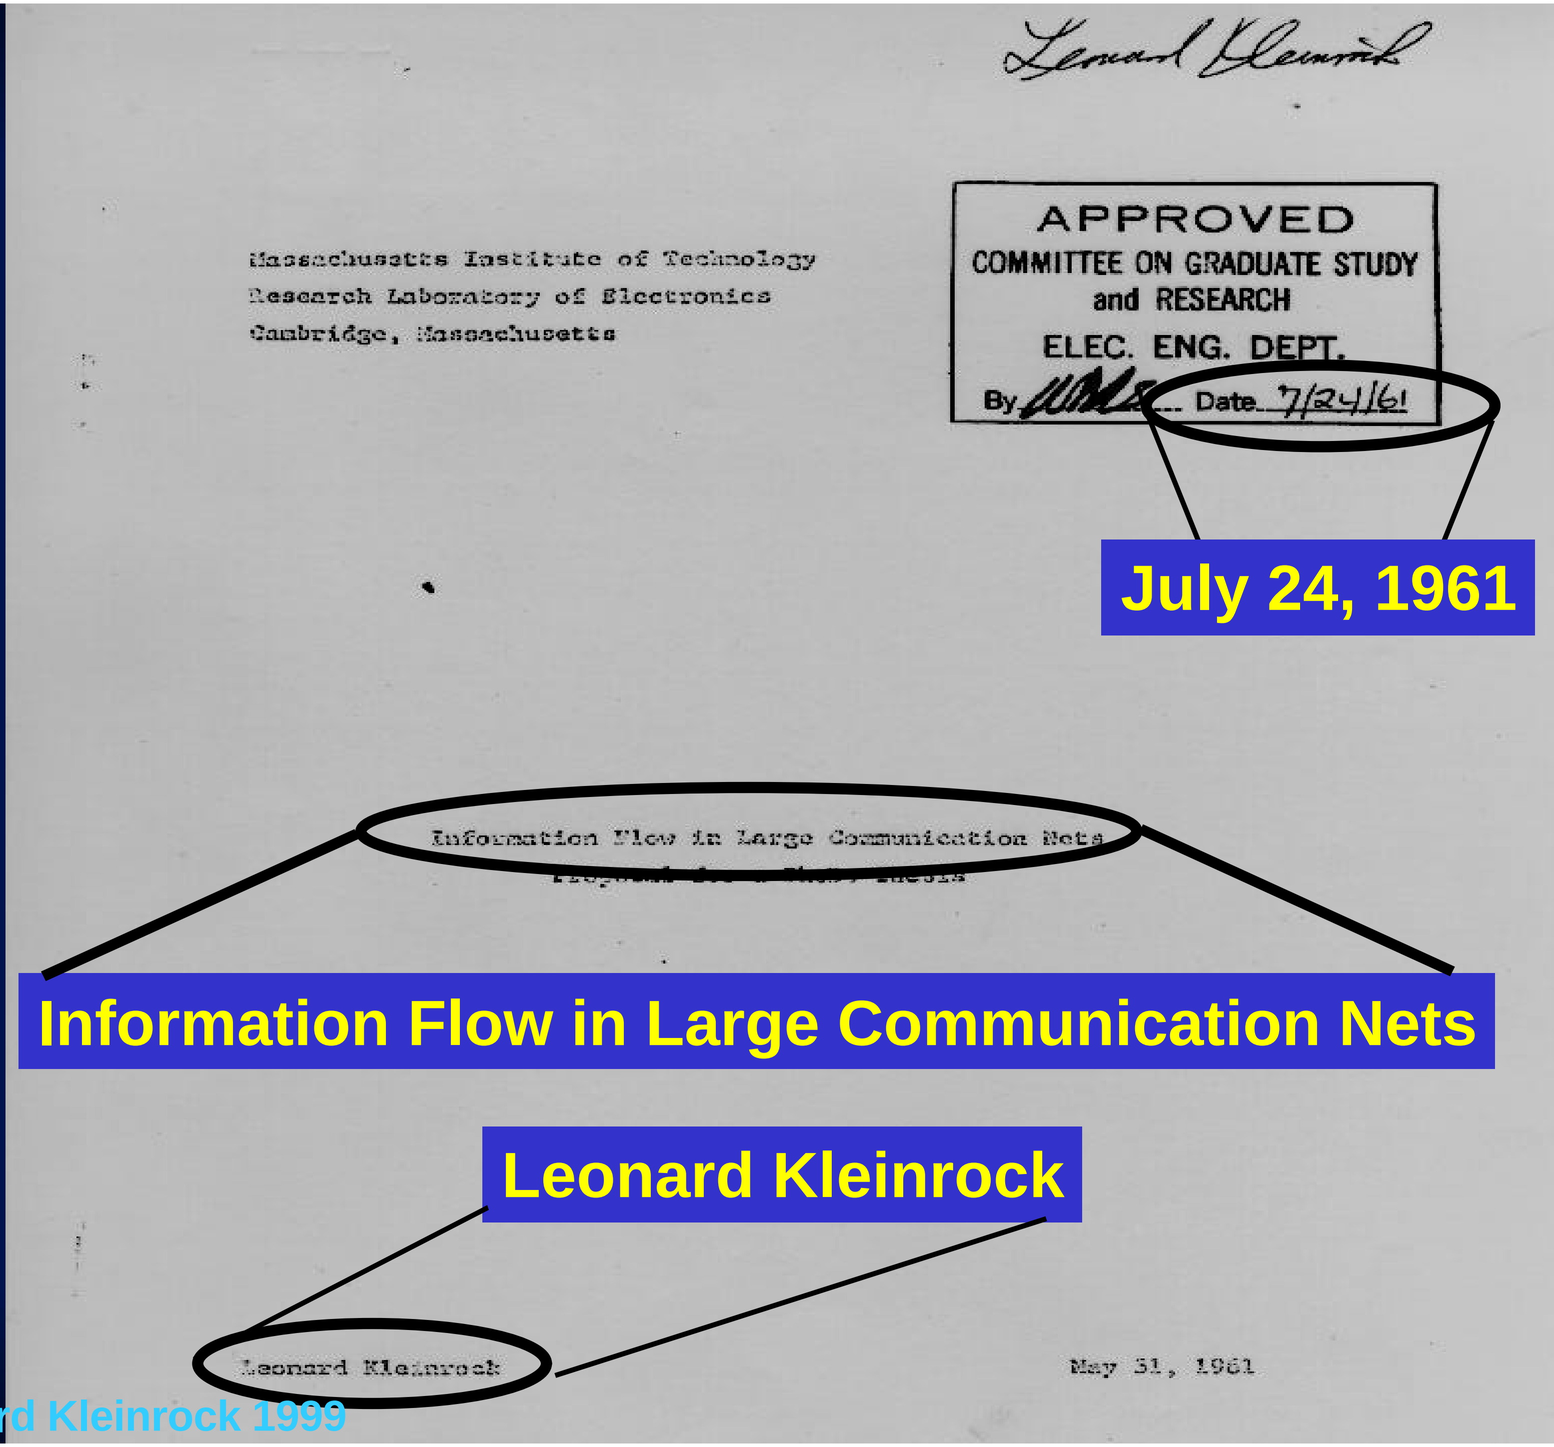

Leonard Kleinrock

APPROVED
COMMITTEE ON GRADUATE STUDY
and RESEARCH
ELEC. ENG. DEPT.
By ______ Date 7/24/61

Massachusetts Institute of Technology
Research Laboratory of Electronics
Cambridge, Massachusetts

Information Flow in Large Communication Nets

Leonard Kleinrock

May 31, 1961

Information Flow in Large Communication Nets

Proposal for a Ph.D. Thesis

Leonard Kleinrock

I. <u>Statement of the Problem</u>:

The purpose of this thesis is to investigate the problems associated with information flow in large communication nets. These problems appear to have wide application, and yet, little serious research has been conducted in this field. The nets under consideration consist of nodes, connected to each other by links. The nodes receive, sort, store, and transmit messages that enter and leave via the links. The links consist of one-way channels, with fixed capacities. Among the typical systems which fit this description are the Post Office System, telegraph systems, and satellite communication systems.

A number of interesting and important questions can be asked about this system, and it is the purpose of this research to investigate the answers to some of these questions. A partial list of such questions might be as follows:

(1) What is the probability density distribution for the total time lapse between the initiation and reception of a message between any two nodes? In particular, what is the expected value of this distribution?

(2) Can one discuss the effective channel capacity between any two nodes?

(3) Is it possible to predict the transient behavior and recovery time of the net under sudden changes in the traffic statistics?

(4) How large should the storage capacity be at each node?

(5) In what way does one arrive at a routing doctrine for incoming messages in different nets? In fact, can one state some bounds on the optimum performance of the net, independent of the routing doctrine (under some constraint on the set of allowable doctrines)?

Information Flow in Large Communication Nets

"The purpose of this thesis is to investigate the problems associated with information flow in large communication nets. …."

"…The nets under consideration consist of nodes, connected to each other by links. The nodes receive, sort, store, and transmit messages that enter and leave via the links…."

- Time lapse between initiation and reception
- Channel capacity
- Storage capacity size
- Routing doctrine
- Transient behavior and recovery time

(6) Under what conditions does the net jam up, i.e., present an excessive delay in transmitting messages through the net? The solution to this problem will dictate the extent to which the capacity of each link can be used (i.e., the ratio of rate to channel capacity, which is commonly known as the utilization factor).

(7) What are the effects of such things as additional intra-node delays, and priority messages?

One other variable in the system is the amount of information that each node has about the state of the system (i.e., how long the queues are in each other node). It is clear that these are critical questions which need answers, and it is the intent of this research to answer some of them.

In attempting the solution of some of these problems, it may well be that the study of a specific system or application will expose the basis for an understanding of the problem. It is anticipated that such a study, as well as a simulation of the system on a digital computer, will be undertaken in the course of this research.

II. History of the Problem

The application of Probability Theory to problems of telephone traffic represents the earliest area of investigation related to the present communication network problem. The first work in this direction dates back to 1907 and 1908 when E. Johannsen [1][1,2] published two essays, the one dealing with delays to incoming calls in a manual telephone exchange, and the other being an investigation as to how often subscribers with one or more lines are reported "busy." It was Dr. Johannsen who encouraged A. K. Erlang to investigate problems of this nature. Erlang, an engineer with the Copenhagen telephone exchange, made a number of major contributions to the theory of telephone traffic, all of which are translated and reported in [1]; his first paper (on the Poissonian distribution of incoming calls) appeared in 1909 and the paper containing the results of his main work was published in 1917, in which

---

[1] Numerals in square brackets refer to the bibliography.

[2] Reference to Johannsen's work will be found in [1], page 16.

Under what conditions does the net jam up?

(6) Under what

presen

monly

ch things as additional intra-

rity messages?

riable in the system is the amount of information that each node has about the state of the system (i.e., how long the queues are in each other node). It is clear that these are critical questions which need answers, and it is the intent of this research to answer some of them.

In attempting the solution of some of these problems, it may well be that the study of a specific system or application will expose the basis for an understanding of the problem. It is anticipated that such a study, as well as a simulation of the system on a digital computer, will be undertaken in the course of this research.

II. History of the Problem

The application of Probability Theory to problems of telephone traffic represents the earliest area of investigation related to the present communication network problem. The first work in this direction dates back to 1907 and 1908 when E. Johannsen [1][1,2] published two essays, the one dealing with delays to incoming calls in a manual telephone exchange, and the other being an investigation as to how often subscribers with one or more lines are reported "busy." It was Dr. Johannsen who encouraged A. K. Erlang to investigate problems of this nature. Erlang, an engineer with the Copenhagen telephone exchange, made a number of major contributions to the theory of telephone traffic, all of which are translated and reported in [1]; his first paper (on the Poissonian distribution of incoming calls) appeared in 1909 and the paper containing the results of his main work was published in 1917, in which

---

[1] Numerals in square brackets refer to the bibliography.

[2] Reference to Johannsen's work will be found in [1], page 16.

My Early Dissertation Work

- Developed theory of stochastic flow of message traffic in connected networks of communication centers:
  - Channel capacity limited
  - Mean response time as key metric
  - Optimal assignment of channel capacity
  - Choice of priority queueing discipline
  - Choice of routing procedure
  - Design of topological structure
- Developed underlying principles of data networks

Systems of Flow

- Steady flow through a single channel
  - Trivial and deterministic
- Unsteady flow through a single channel
  - Queueing theory; stochastics get you
- Steady flow through a network of channels
  - Network flow theory; multicommodity gets you
- Unsteady flow through a network of channels
  - A New domain; everything gets you!
    - Jackson’s networks of queues (1957)
    - Kleinrock’s Independence Assumption cracks the problem wide open

Key Results in My PhD Dissertation

- Set up the model:
  - Use of queueing theory; Erlang's heritage
  - Independence assumption (critical!)
- Evaluated network performance
- Developed optimal design procedures
  - Capacity, topology, routing, message size
- Introduced and evaluated distributed adaptive routing control
- Evaluated different queueing disciplines for handling traffic in the nodes, specifically, chopping messages into smaller segments

Key Equation for Networks

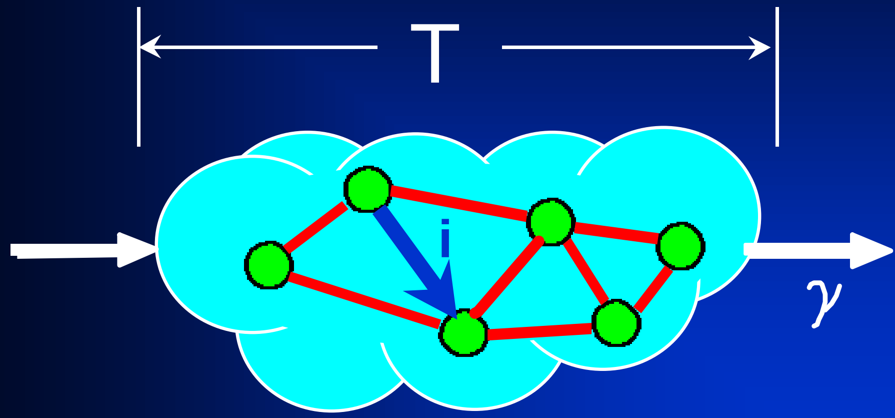

$$T = \sum_i \frac{\lambda_i}{\gamma} T_i$$

This is EXACT!!

- T = Average network delay
- λ_i = Traffic on channel i (Msg/sec)
- γ = Network throughput (Msg/sec)
- T_i = Average delay for channel i

But how do you find this term?

Key Assumption

The Independence Assumption

Each time that a message is received at a node within the net, a new length is chosen for this message independently from an exponential distribution

The Independence Assumption

- **Without the Independence Assumption, the problem is intractable.**
- **With the Independence Assumption, the problem is totally manageable!!**
- **We get:**

$$T_i = \frac{1}{\mu C_i - \lambda_i}$$

where $\mu C_i =$ **Capacity of channel i (Msg/sec)**

Response Time vs Throughput

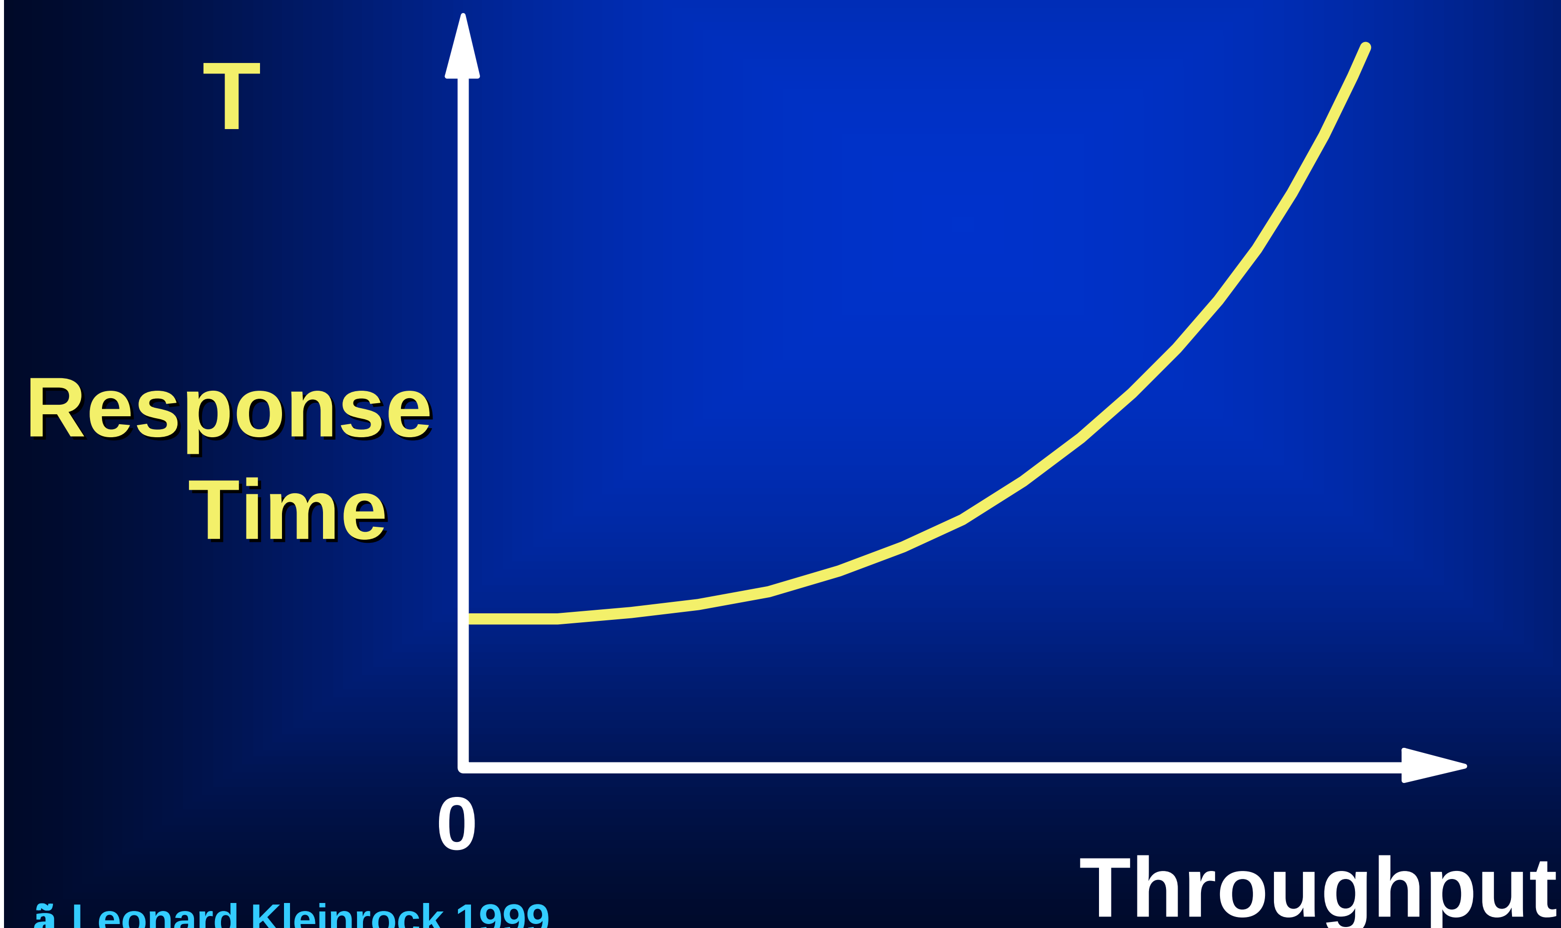

How Do Queues Form?

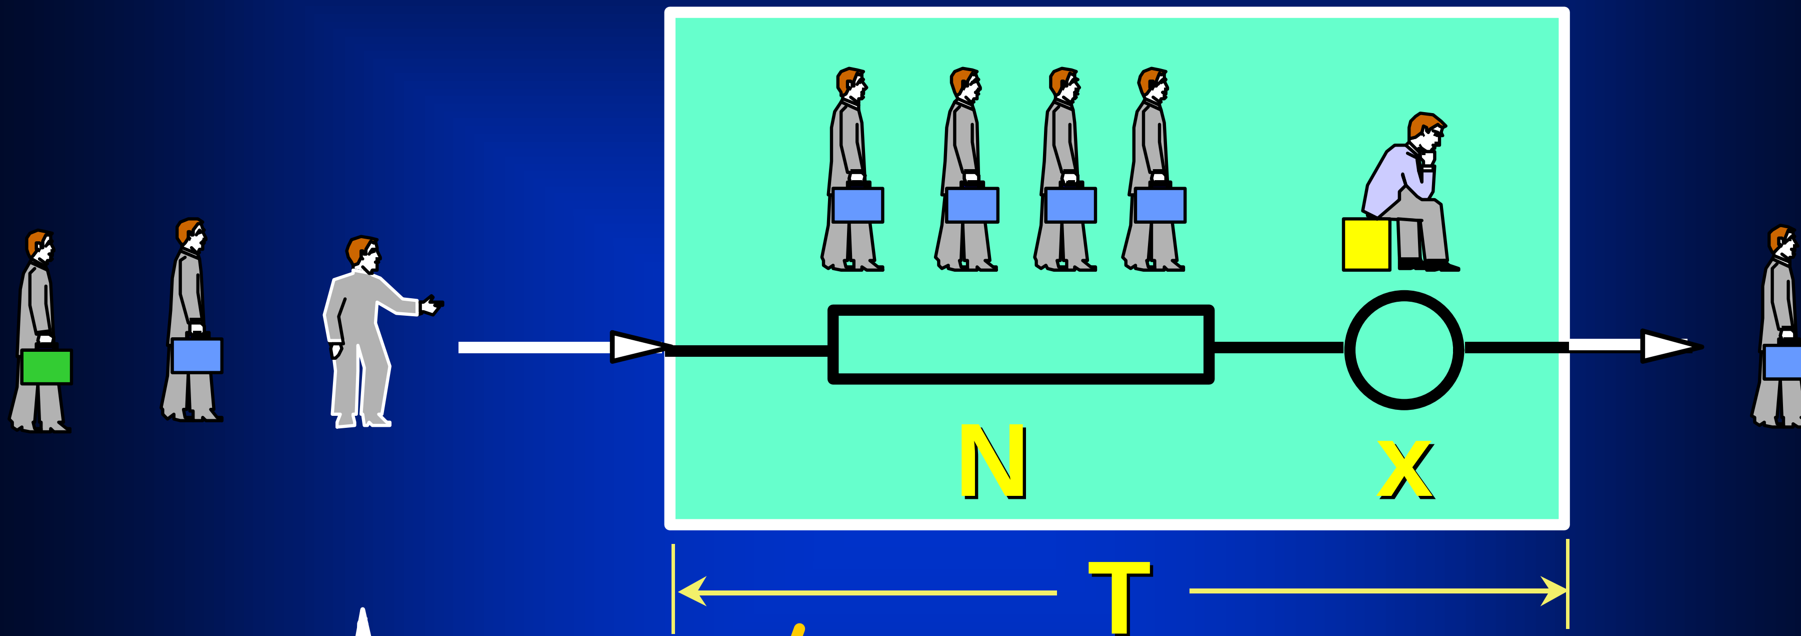

$$T = Nx + x$$

$$N = \gamma T \quad \text{(Little's Law)}$$

$$T = T\gamma x + x$$

$$T = x \,/\, (1 - \gamma x)$$

Simple 2-parameter Model For Delay

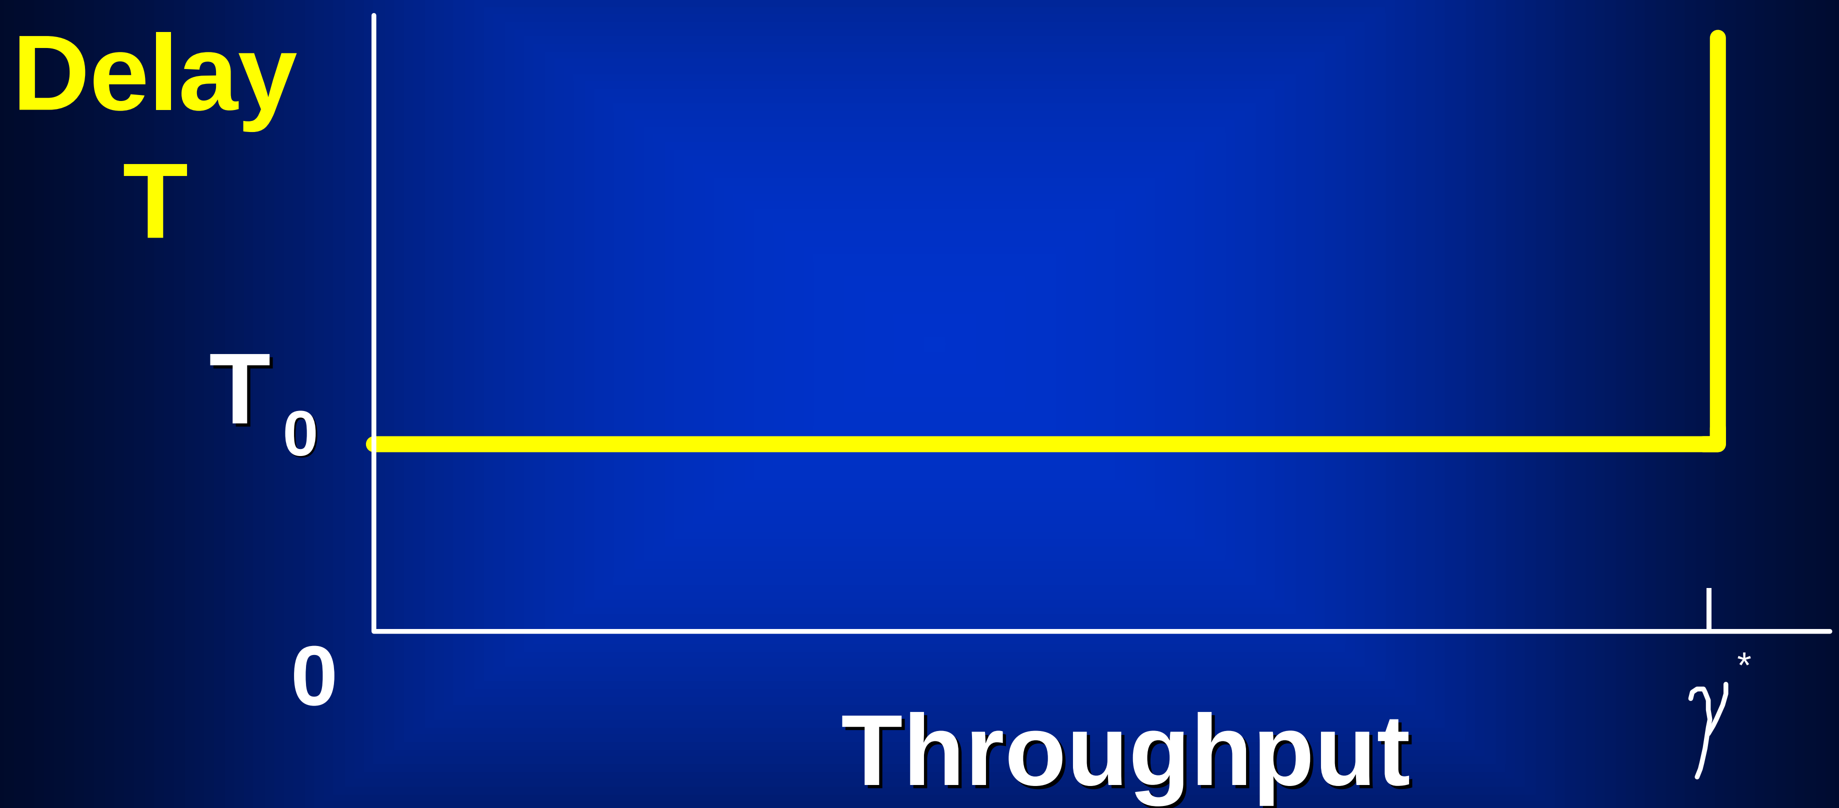

The General Optimization Problem

- **Minimize** $T = \sum_i \frac{\lambda_i}{\gamma} T_i$

{ Channel Capacity Assignment
Routing Procedure
Message queueing discipline
Topology }

- **Subject to:** $D = \sum_i d_i C_i$

Where

C_i = Channel capacity of i^{th} channel
d_i = Cost to supply 1 unit of capacity to i^{th} channel
D = Total dollars available for design

Solution to the Problem

- Exact solution for $d_i = 1$
- Exact solution for arbitrary d_i
- Implications for topology
- Implications for routing procedure
- Implications for message sizes

The Underlying Principles

- Resource Sharing (demand access)
  - Only assign a resource to data that is present
  - Examples are:
    - Message switching
    - Packet switching
    - Polling
    - ATDM
- Economy of Scale in Networks
- Distributed control
  - It is efficient, stable, robust, fault-tolerant and WORKS!

Resources

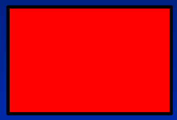

- A Resource is a device that can do work for you at a finite rate
- Examples:
  - A Communication Channel
  - A Computer

Resources

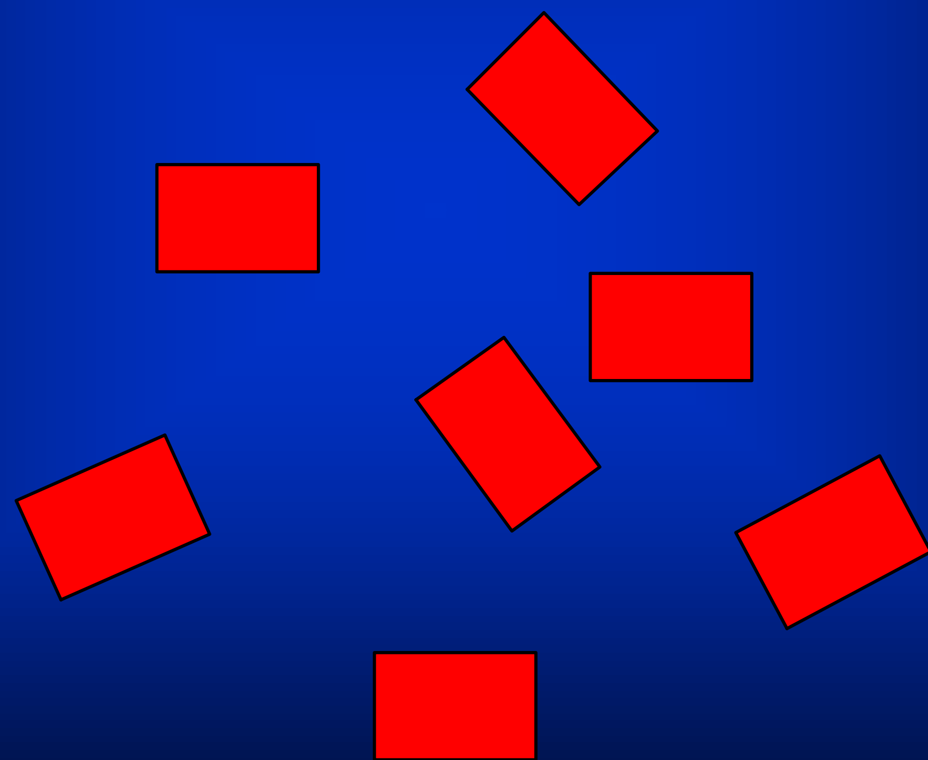

Demands

- A **Demand** requires work from resources

- Examples:
  - Packets (require transmission)
  - Jobs (require processing)

Demands

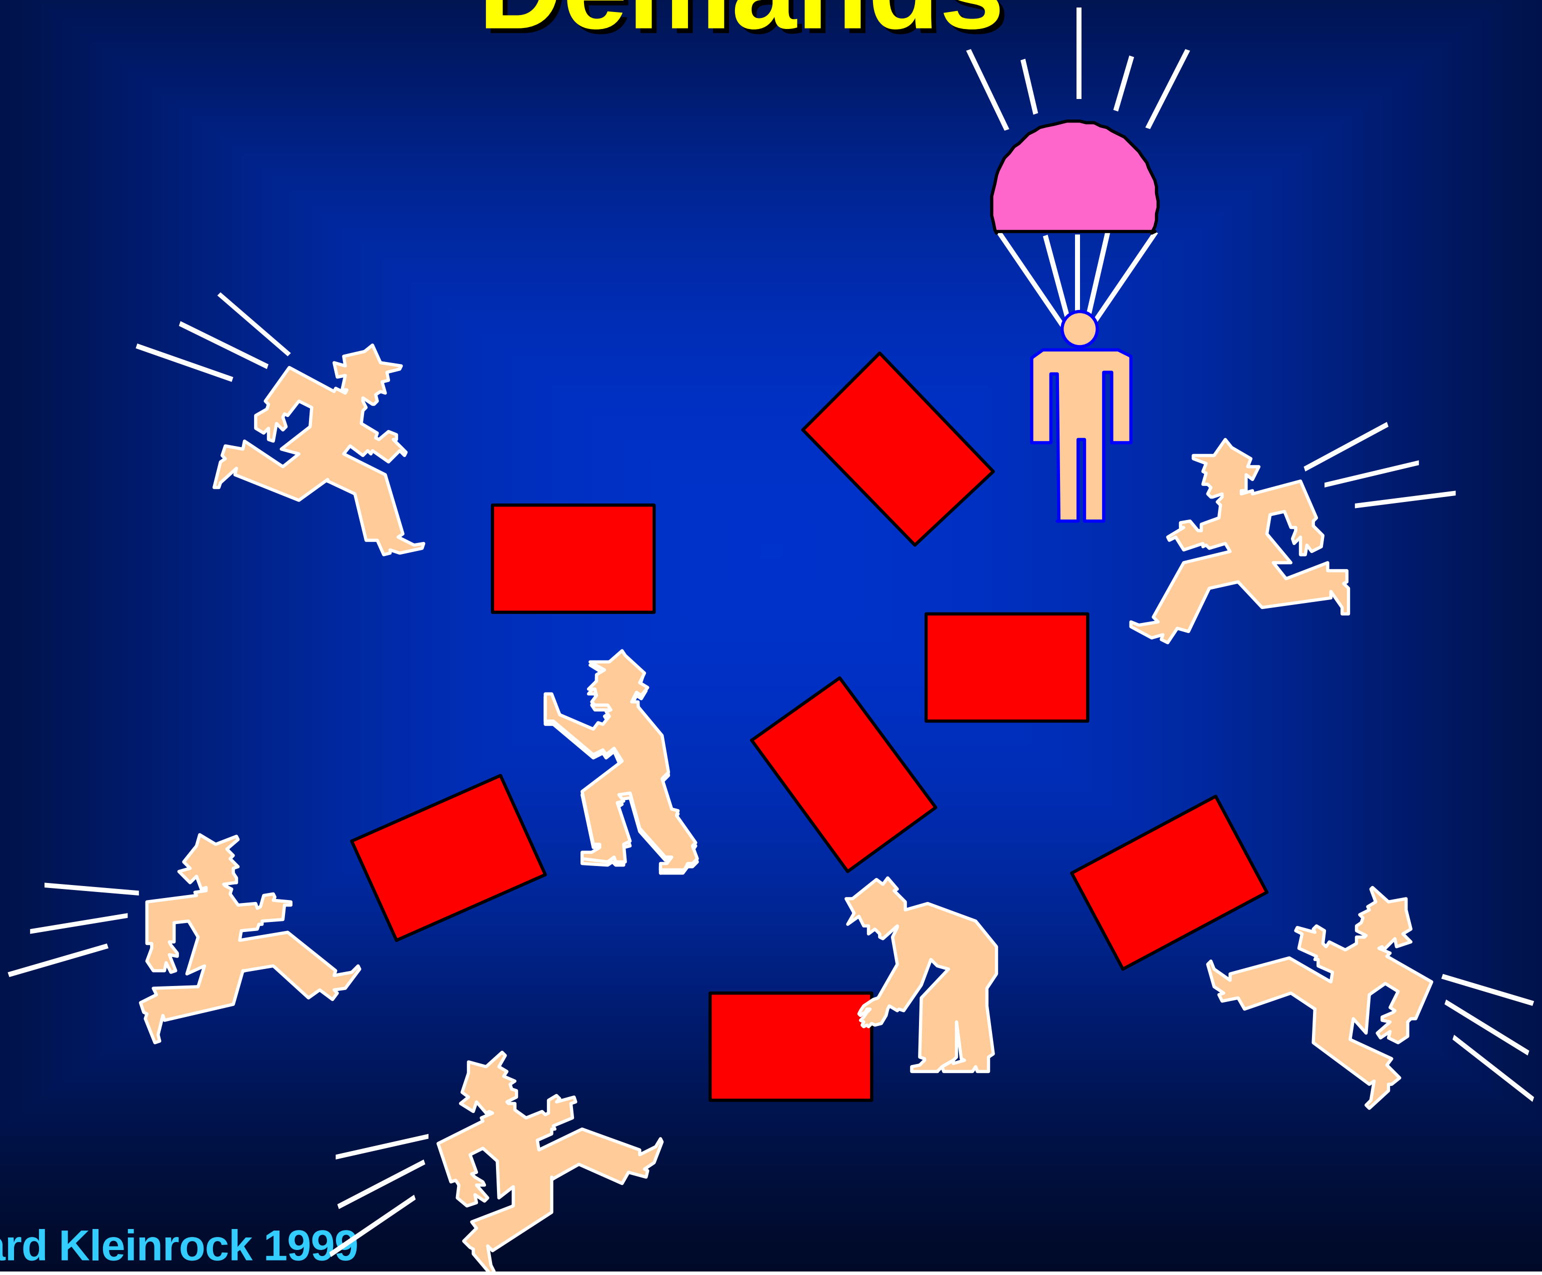

Bursty Asynchronous Demands

- You cannot predict exactly **when** they will demand access
- You cannot predict **how much** they will demand
- Most of the time they **do not need** access
- When they ask for it, they want **immediate** access!!

Resource Sharing

Type 0

Resource Sharing

Type 1
Dedicated Resources

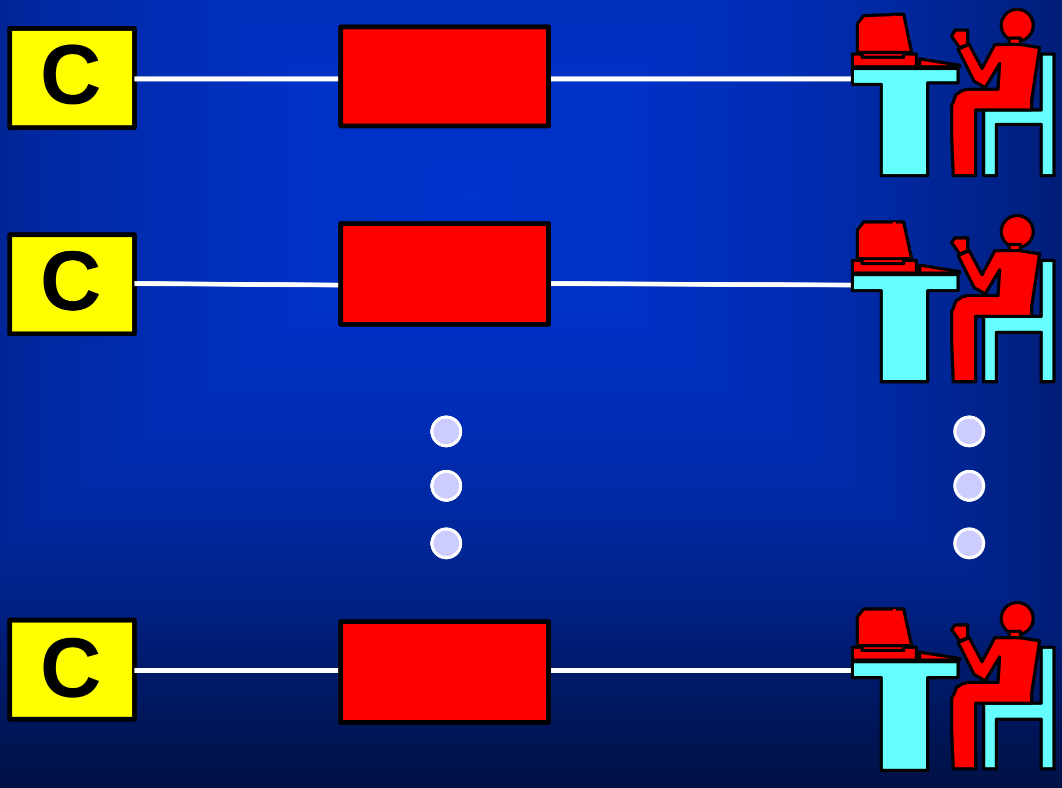

Resource Sharing

Type 2
Shared Resources

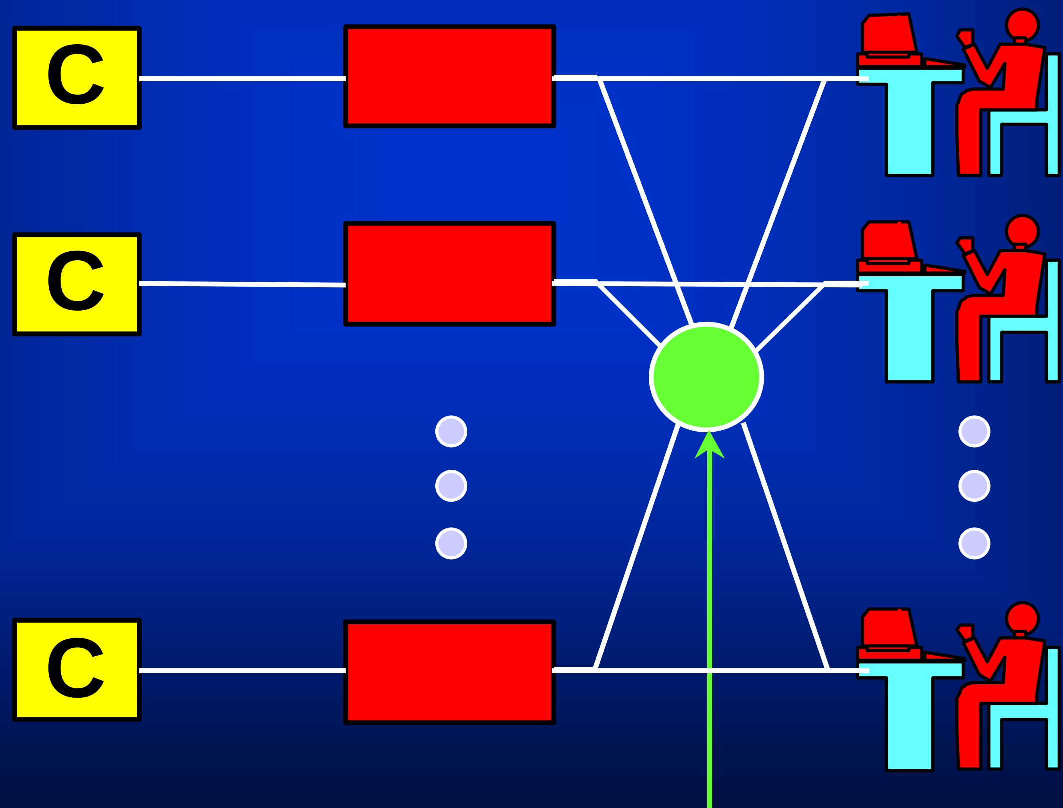

The Law of Large Numbers
(The First Resource Sharing Principle)

- Although each member of a large population may behave in a random fashion, the population as a whole behaves in a predictable fashion.
- This predictable fashion presents a total demand equal to the sum of the average demands of each member.

- This is the “smoothing effect” of large populations.

Resource Sharing:

Telephone Trunks

Buildings @ 2/3 Erlangs   Trunks   Blocking

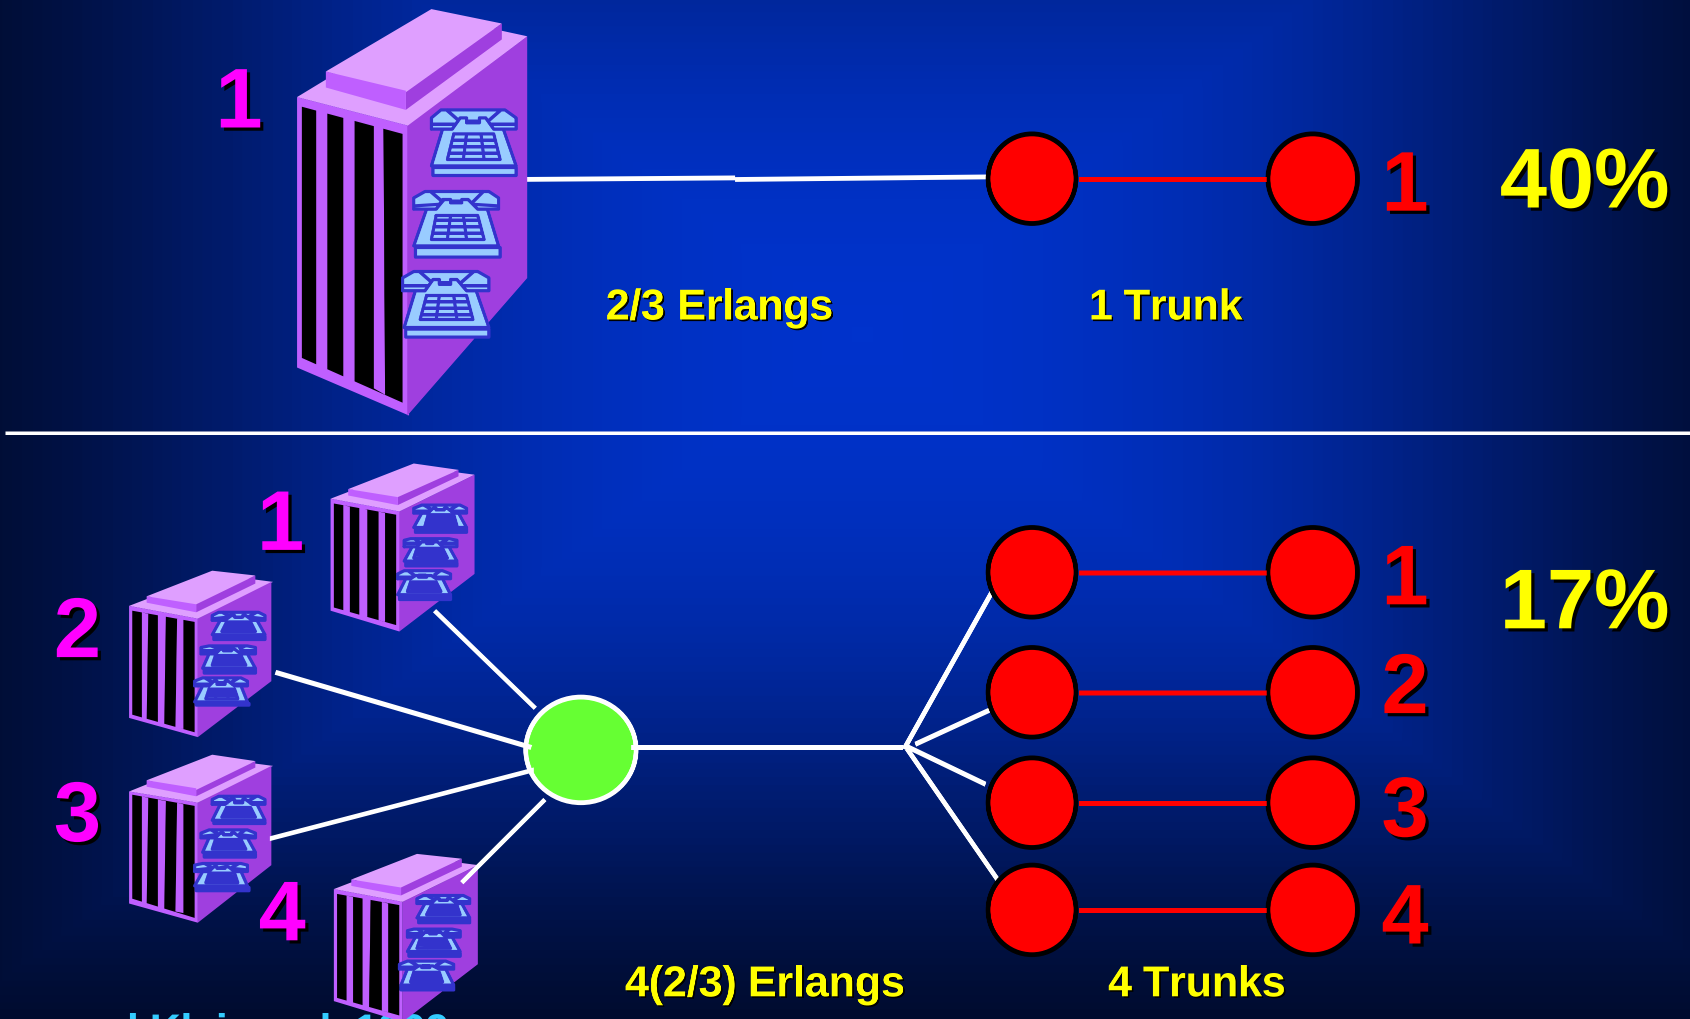

Resource Sharing:

Telephone Trunks

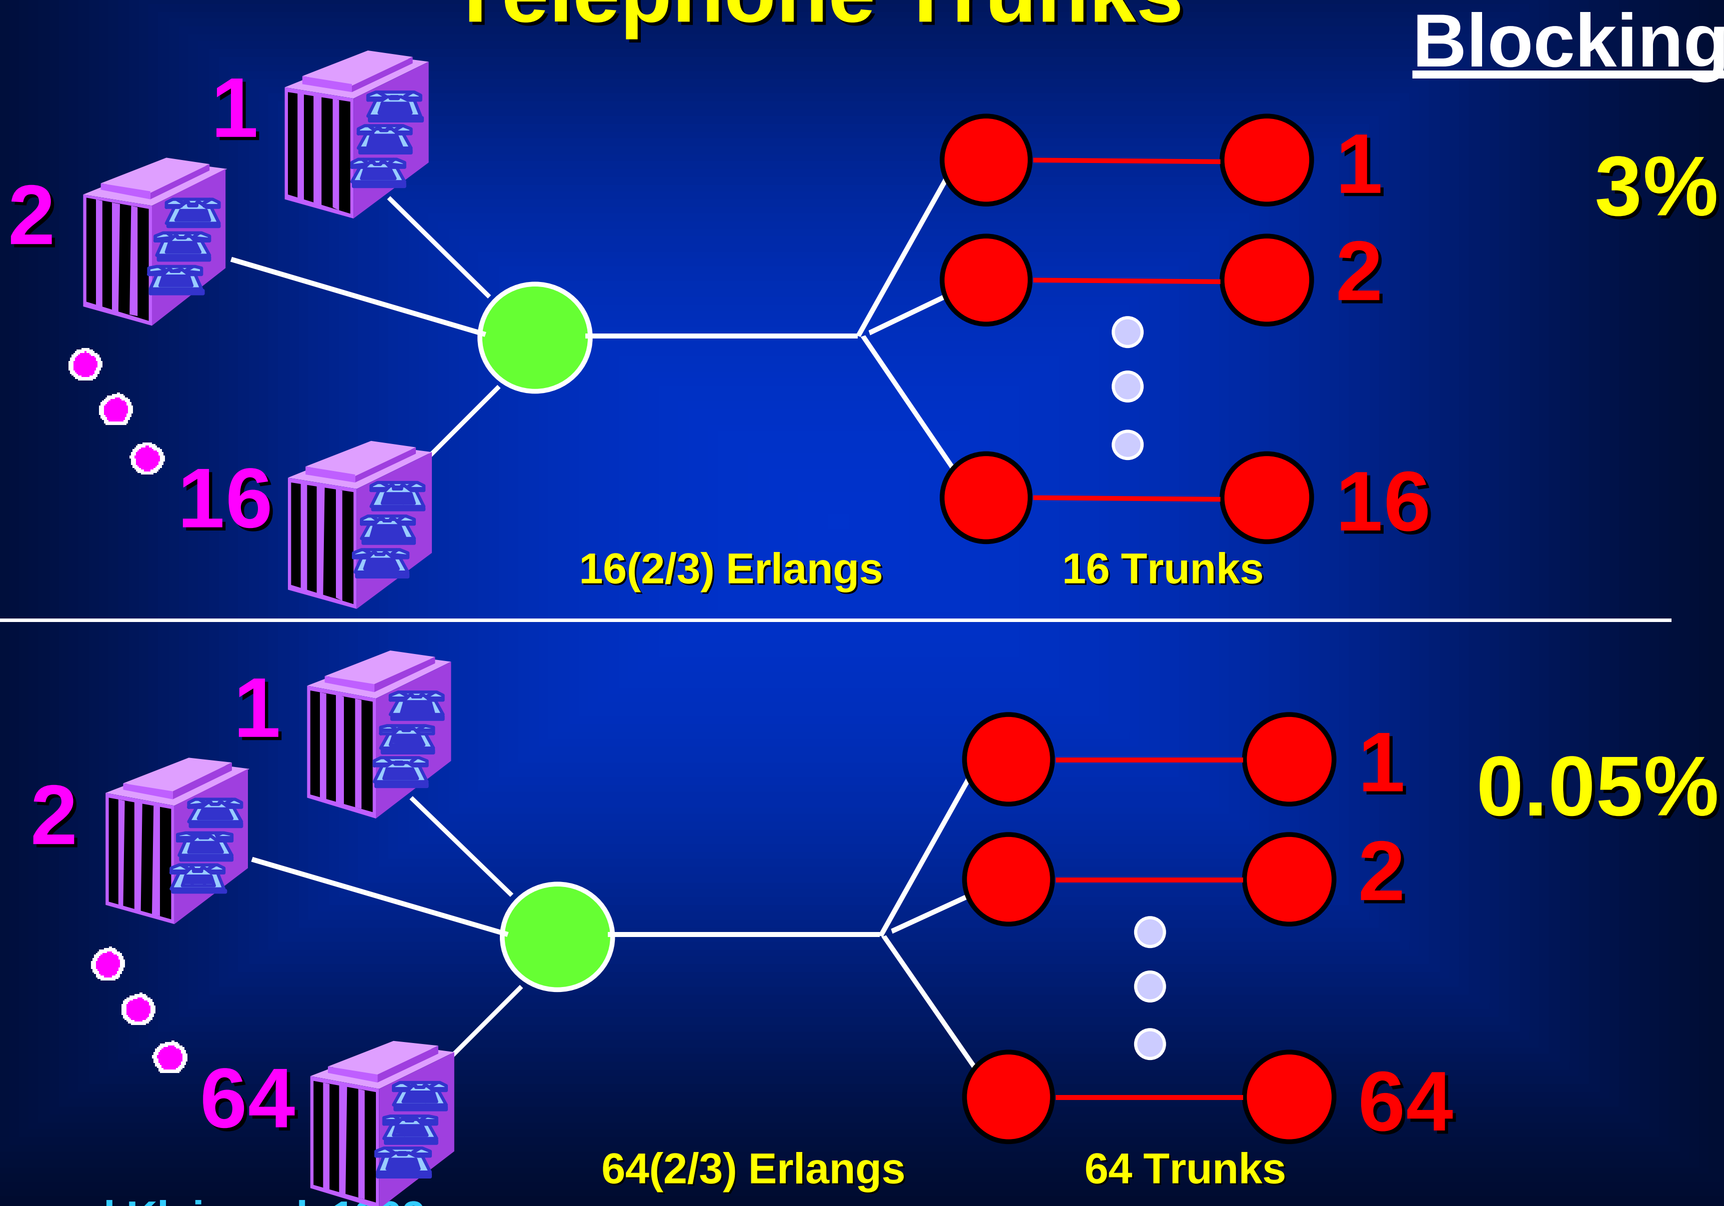

Resource Sharing

Type 2
Shared Resources

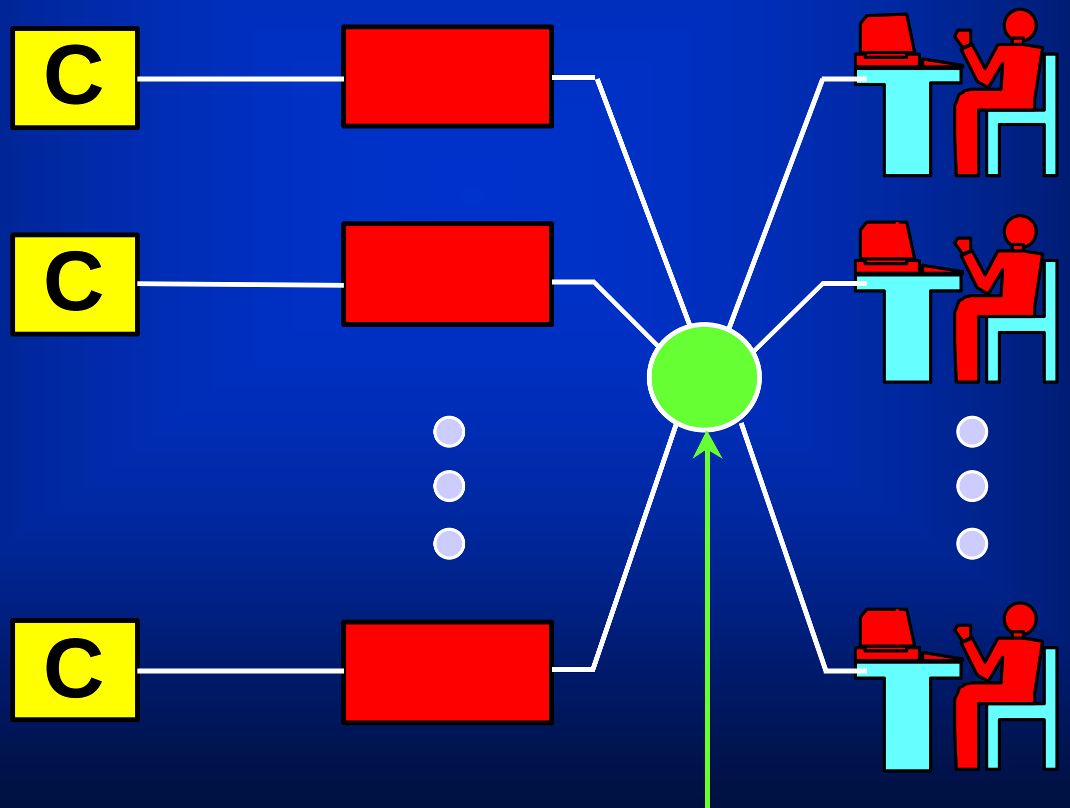

A Fancy Green Switch

Resource Sharing
Type 3
LARGE Shared Resources

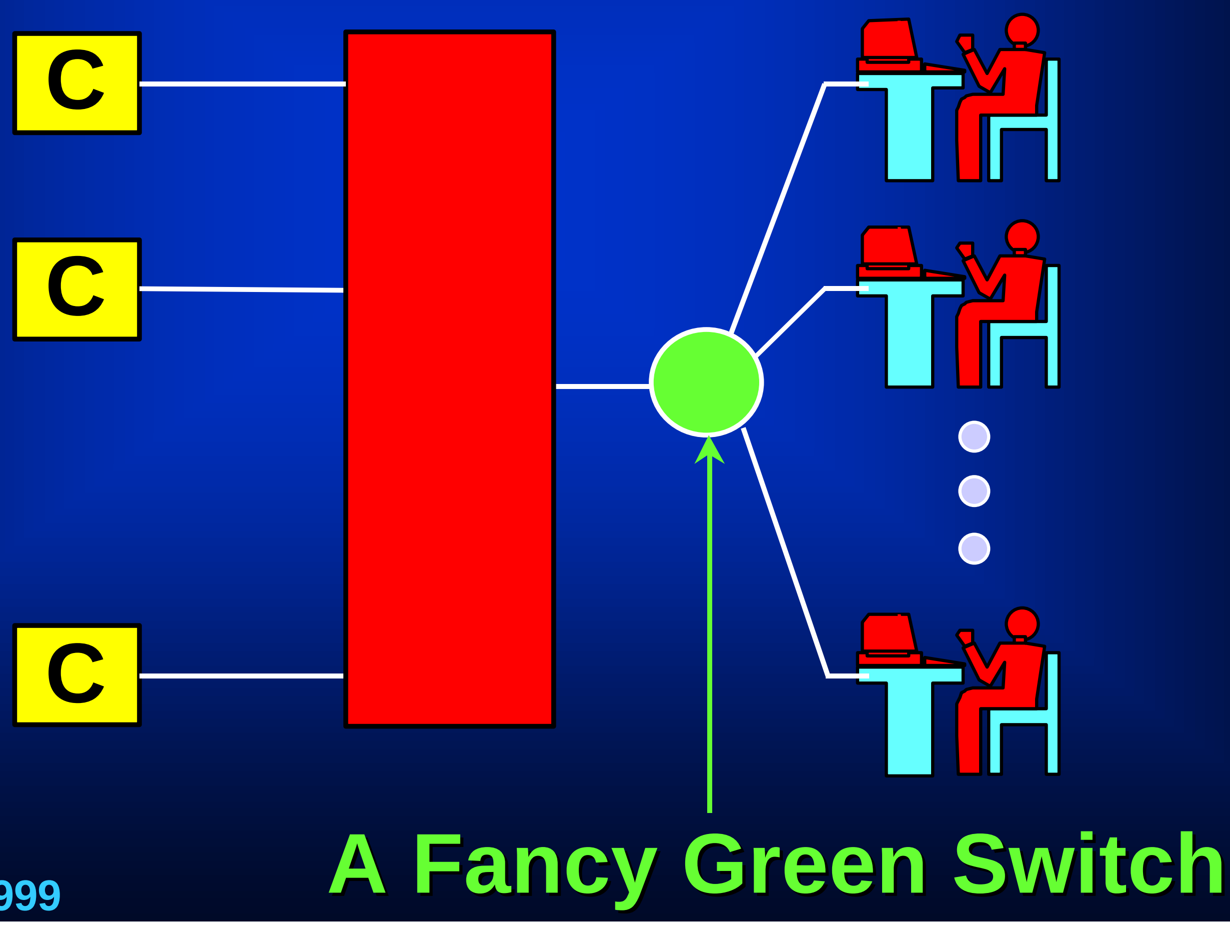

Conflict Resolution

- **Queueing:**
  - One gets served
  - All others wait
- **Splitting:**
  - Each gets a piece of the resource
- **Blocking:**
  - One gets served
  - All others are refused
- **Smashing:**
  - Nobody gets served !

The Economy of Scale
(The Second Resource Sharing Principle)

- If you scale up throughput and capacity by some factor F, then you reduce response time by that same factor.
- If you scale capacity more slowly than throughput while holding response time constant, then efficiency will increase (and can approach 100%).

Resource Sharing:

Data Communications

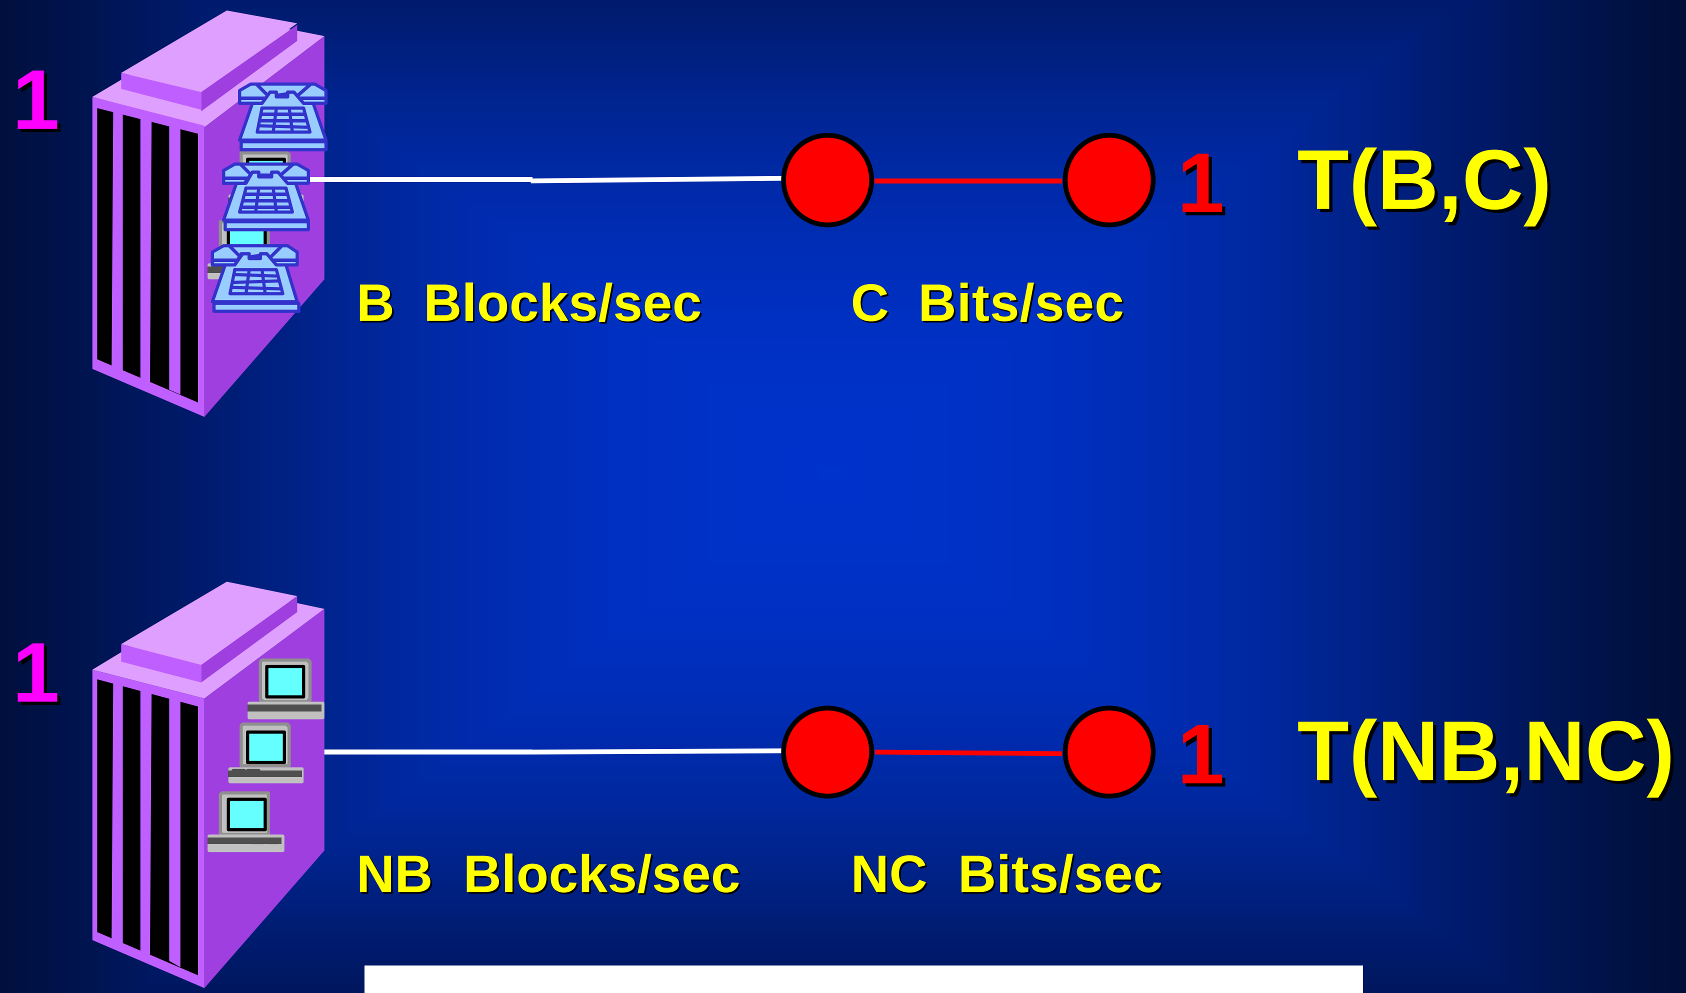

T(NB,NC) = T(B,C) / N

Key Tradeoff:
Response Time, Throughput, Efficiency

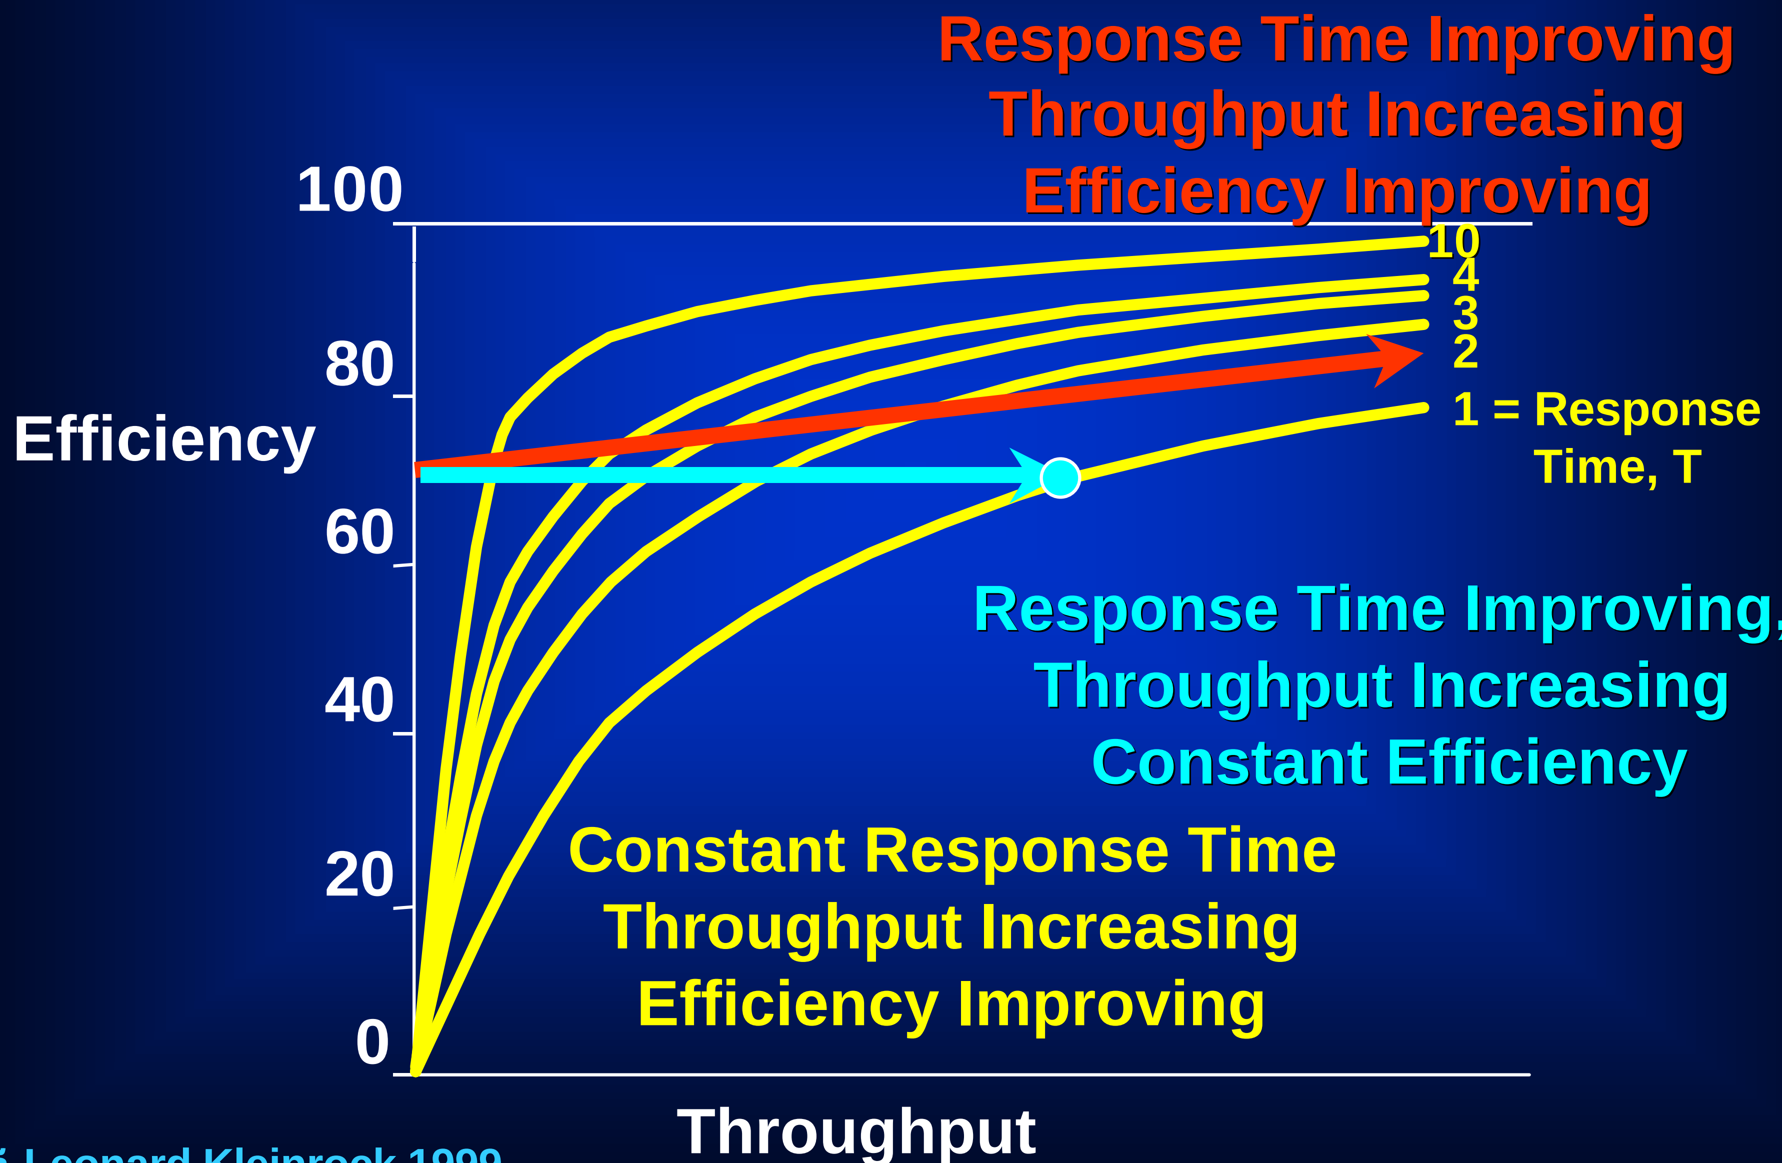

Resource Sharing

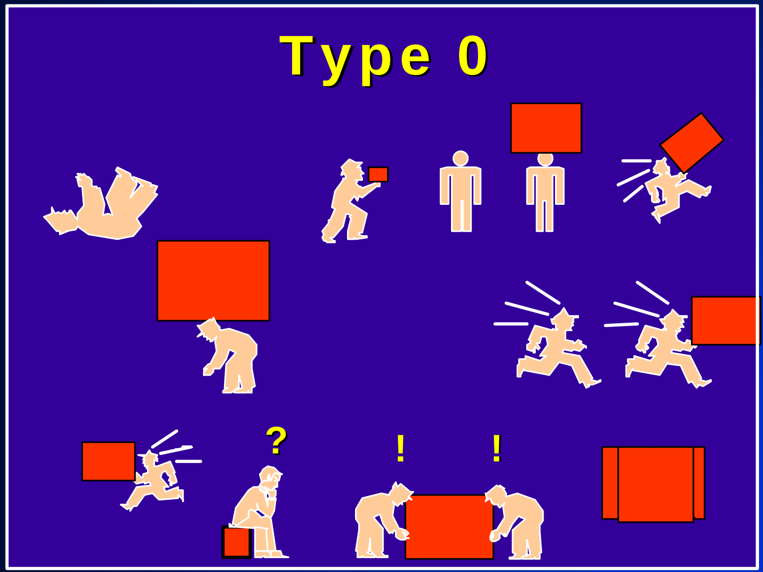

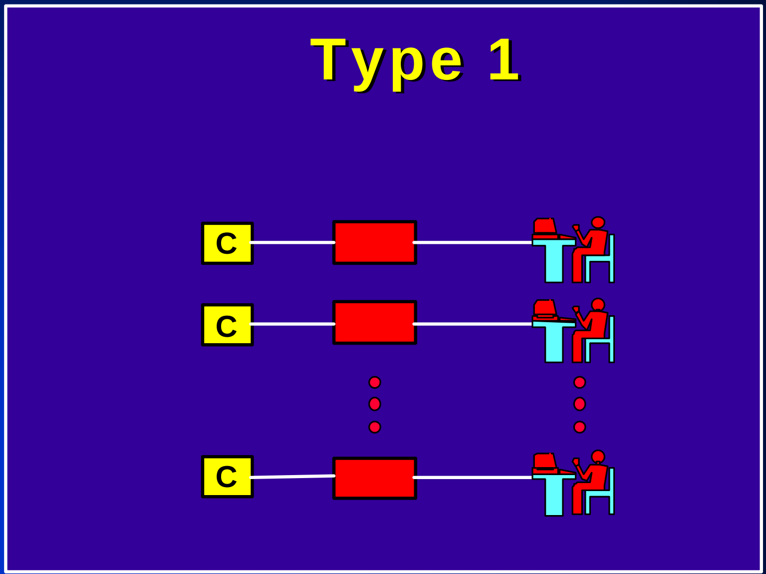

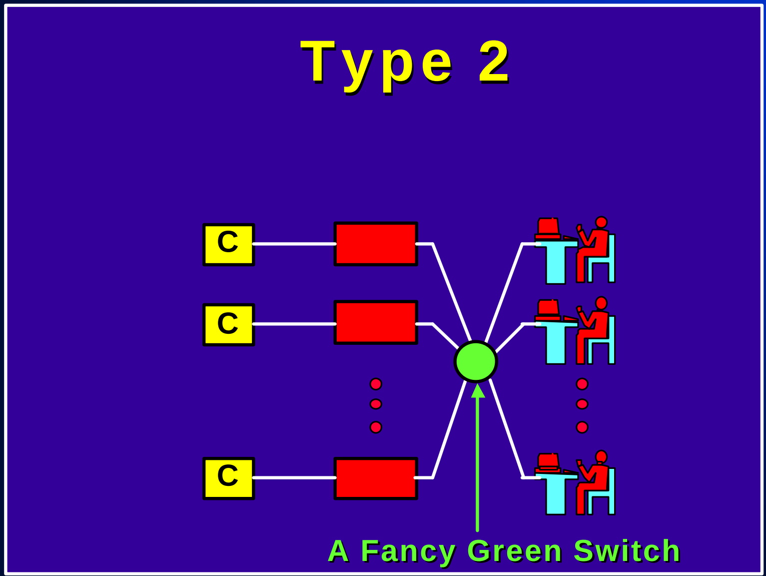

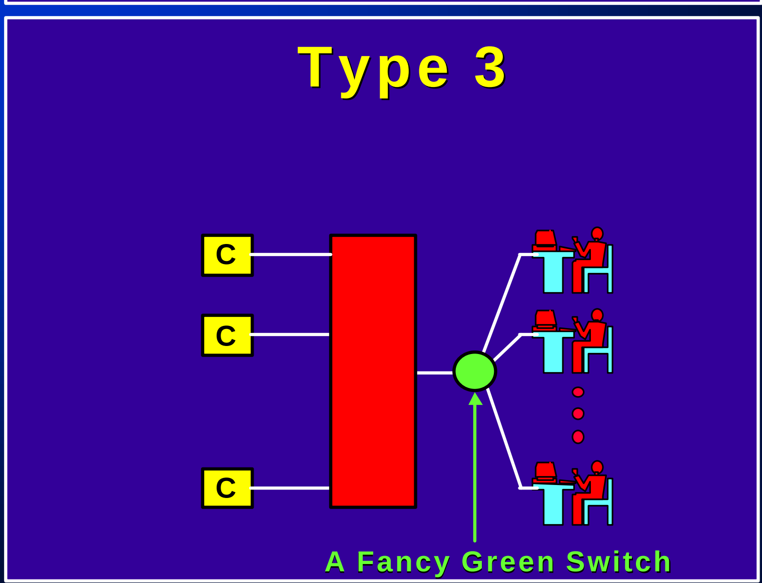

Economy of Scale in Networks

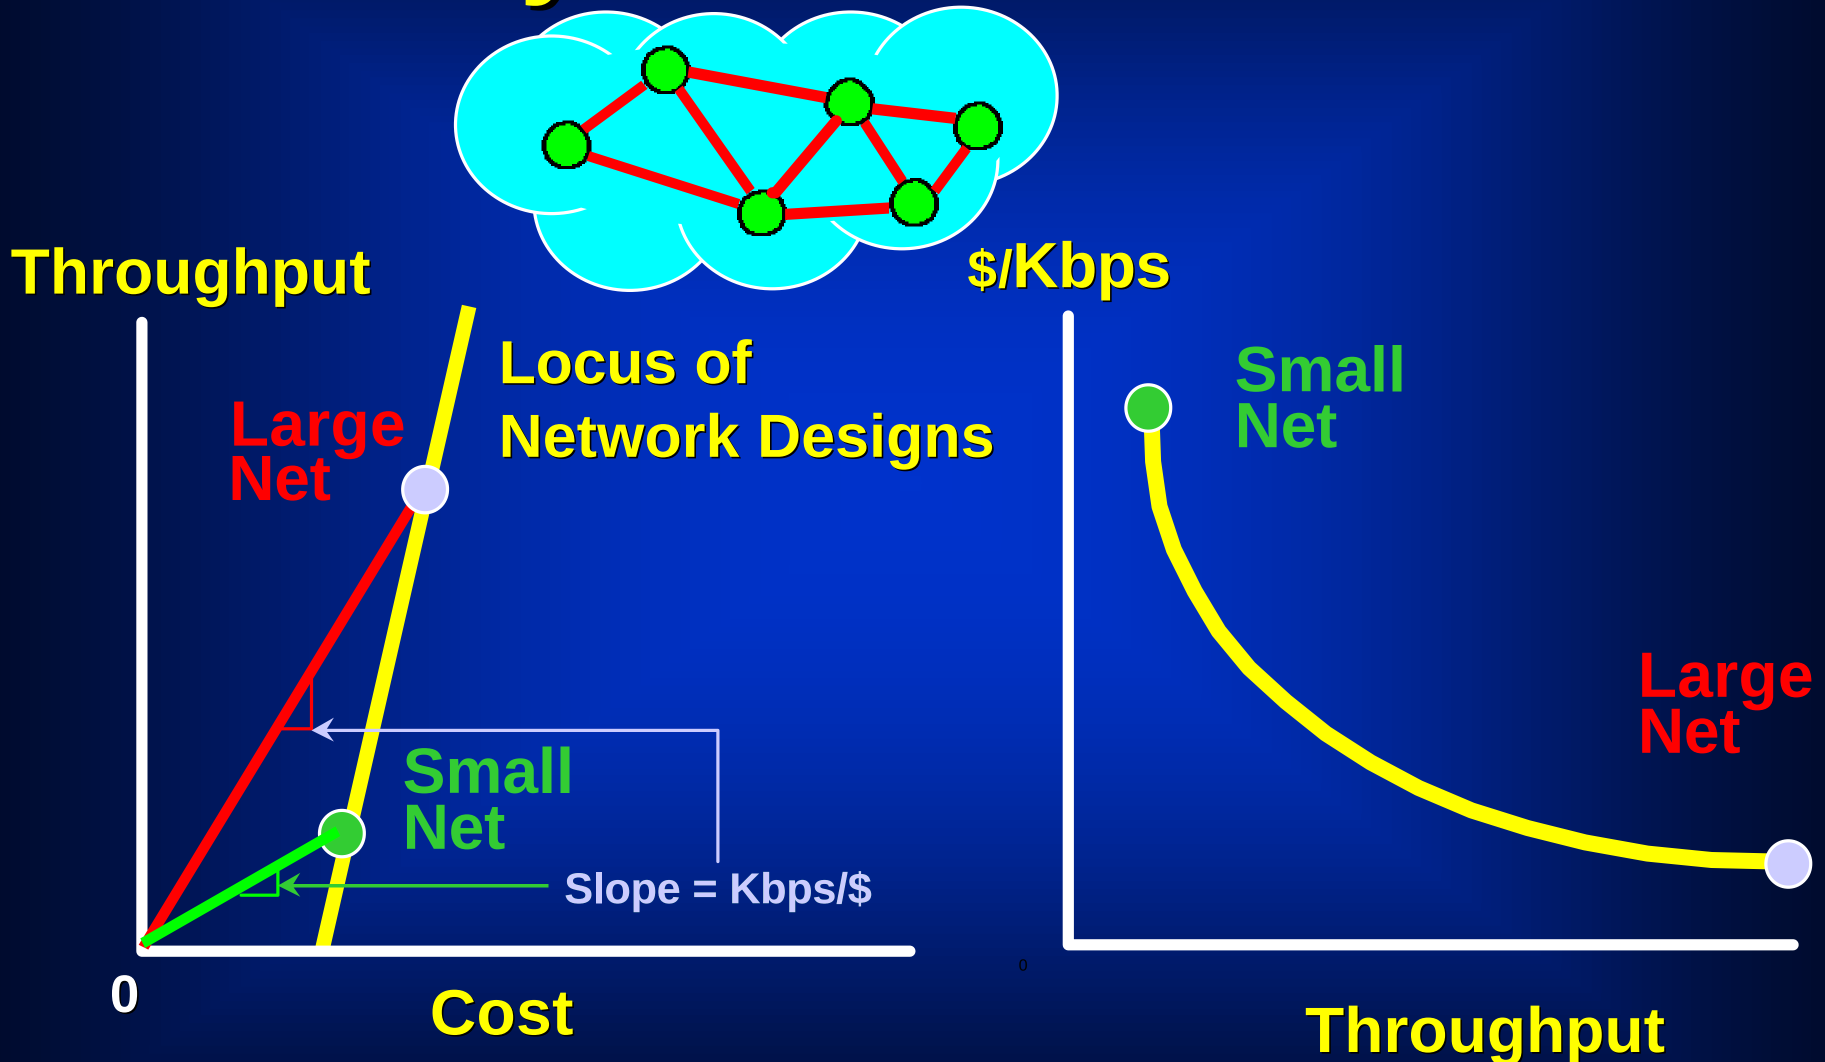

Thank You

www.lk.cs.ucla.edu