DIVIT RAWAL 122

Host - contains info
Router - forwards info
Link - carries info

packet switching - flexible routing, grouped into packets (datagrams)
circuit switching - inflexible routing, selects set of links only once
virtual circuit - set of links selected once, required data rate not reserved

IPv4 - 32 bits - a.b.c.d
$0 \le a,b,c,d \le 255$
IPv6 - 128 bits
IP addresses assigned by location
- CIDR (classless interdomain routing)
- longest prefix matching

Metrics

Link Rate (Mbps) (R)

Link Bandwidth (Hz) (W)
- width of range of frequencies link can transfer over

Link Capacity (Mbps) (C)
- C = maximum reliable link rate
- Shannon Capacity
  - $C = W \log_2(1 + SNR)$
    - $SNR = \frac{pwr(signal)}{pwr(noise)}$ @ reciever

Delay
- time for packet to travel b/w two points
- end-to-end delay
  - transmission and propogation times (link)
  - queuing & processing times (node)

Throughput
- size/time
  - not same as link rate
    - includes gaps/delays
- limited by window size/RTT
- determined by slowest link

Delay Jitter
- longest delivery time - shortest delivery time
- if jitter is J, destination should store each packet for at least J seconds
- practically, start w/ J = T, then adjust as packets come in

M/M/1 Queue
- λ packets arrive per second (avg)
- μ packets sent per second (avg)
- assuming $\lambda < \mu$, average time each packet waits = $T = \frac{1}{\mu - \lambda}$
  - if $\lambda \ll 1$, $T = \frac{1}{\mu}$
    - average queuing time a packet waits = $T - \frac{1}{\mu}$
- avg # of packets in queue = $L = \frac{\lambda}{\mu - \lambda}$
- T, L grow w/o bound as $\lambda \to \mu$
- $\mu = \frac{\text{link rate}}{\text{avg packet length}}$
- utilization = $\rho = \frac{\lambda}{\mu}$

$\lambda < \mu$ is positive recurrent
$\lambda = \mu$ is null recurrent
$\lambda > \mu$ is transient

Hub - bullhorn
Switch - USPS

Transmission Error
- parity bit - extra bit to make # of 1's even
- header checksum

Congestion loss
- packets arrive faster than they can be forwarded
- solved w/ ARQ (Automatic Retransmission Request)
  - send ACK, if not recieved in some time, wait (backoff) then resend

2-Level Routing
- group nodes into domains
- each domain uses a shortest path algorithm
- use a shortest path algorithm between domains

Layers
- physical layer
- link layer
- network layer
- transport layer
- application layer

Congestion Control - hosts slow down when they miss acknowledgements

Flow Control - reciever sends back # of available buffer bits in each ACK so sender knows when to stop/slow

Multiplexing Gain - measures benefit of sharing link amongst many users
ex. $C/10$ vs $C/1000 \Rightarrow MG = 100$

outstanding packet - transmitted but no ACK recieved
- max # of outstanding bytes is window size

RTT - round trip time - time to send + recieve ACK

End-to-end principle
- tasks should not be performed by routers if they
- routers are stateless - consider 1 packet at a time, no info about connections

Multicast - one to many
Anycast - one to any

Aloha introduced Randomized Multiple Access
- device transmissions @ random, not scheduled

Aloha
- if ACK is not outstanding, transmit immediately
- if no ACK, repeat after random delay

Ethernet
- (Wait) 0 for new packet, backoff scheme used for retrans.
- (Transmit) after waiting time, when channel is idle, transmit
- (Collision Detection) if collision, abort trans. & go to (wait) for re-trans. else, go to (wait) for a new packet
- backoff scheme wait X timeslots where X is selected uniformly from $\{0, 1, \dots, 2^n - 1\}$, m = # collisions, $n = \min(m, 10)$
  - truncated binary exponential backoff - give up for $m \ge 16$
- time slots - 512 bit trans. times for 10/100 Mbps Ethernet, 4096 bit times for Gbps Ethernet
  - larger than 2PROP
- winner takes all consequence
- uses MAC (Media Access Control) Address: 48 bits
  - all 1's = broadcast addr.

Ethernet II Frame Format

Bytes	7	1	6	6	2	46-1500	4
	PRE	SFD	DA	SA	Type	Data	CRC

PRE - preamble, used for sync
SFD - start of frame delimiter
DA - destination addr.
SA - source addr.
Type - type of payload
CRC - cyclic redundancy check (used for detecting corruption)

Class A
0, 8
| 0 net | host |

Class B, 16
| 10 net | host |

Class C, 24
| 110 | net | host |

Little's Result
- L - avg # of packets in sys
- W - avg. time packet spends in system
- $L = \lambda W$

$P(n, n+1) = 1 - 2\epsilon - \mu\epsilon$
$\pi(n) = (1 - \rho)\rho^n$
$0 < \rho < 1$, $n \ge 0$

Fairness
- max-sum - give full capacity to device that needs it most
- max utility - pick dist such that each gets capacity α to need (max-min)
- proportional fairness - maximize product of utilities
- utility = need $\times$ proportion

Spanning Tree Protocol
broadcast [id | root | dist to root]
determine root as smallest ID
each node determines smallest path to root
other links are blocked
Ports
- Root
- Designated
- Blocked

Aloha

↳ Time Slotted
- ↳ n devices
- ↳ probability p of transmitting in a given time slot
- ↳ chances of exactly one transmission in slot $= R(p) = np(1-p)^{n-1}$
- ↳ optimal $p = 1/n$
  - ↳ $R(\frac{1}{n}) = (1-\frac{1}{n})^{n-1} = \frac{1}{e} = .37$
    - ↳ 37% efficiency @ most

↳ Non Slotted
- ↳ 18% efficiency @ most

Hub Ethernet

- ↳ max time to detect collision = 2PROP
- ↳ fraction of time where stations transmit successfully $= \frac{1}{1+3.4A}$

$A = PROP/TRANS$

simulations show $\frac{1}{1+5A}$

WiFi

- ↳ uses Distributed Coordination Fn (DCF)
- ↳ com. through Access Point (AP)
- ↳ set of devices that communicate w/ given AP is Basic Service Set (BSS)
- ↳ MAC Sublayer & Physical Layer

MAC Sublayer

- ↳ binary exponential backoff
- ↳ wait short time before sending ACK
- ↳ wait longer time + backoff
  - ↳ avoids interfering w/ ACK
- ↳ some variations
  - ↳ sender sends RTS
  - ↳ recipient replies w/ CTS
  - ↳ other stations wait for P & ACK to send

Medium Access Control (MAC)

Constant	802.11b	802.11a
Slot time	20μs	9μs
SIFS	10μs	16μs
DIFS	50μs	34μs
EIFS	364μs	94μs
CW_{min}	31	15
CW_{max}	1023	1023

- send ACK after SIFS
- wait DIFS + backoff for channel to be idle for
- backoff = z, z uniformly drawn from $\{0, 1, \ldots, CW_n\}$, $CW_n = \min(CW_{max}, (CW_{min}+1) \cdot 2^n - 1)$
  - n is # of failed tries
- every station tries to decode every packet
- if station recieves incorrect packet, waits EIFS before transmitting
  - ↳ decrements backoff counter after channel idle for EIFS
  - ↳ if a correct packet recieved during EIFS, waits for only DIFS before dec.
- Dealing w/ Collisions
  - ↳ transmitting device sets ACK Time Out (TO) to SIFS + physical layer overheads + margin of slot times (for prop + processing)
    - ↳ if no MAC frame starts during this time, can conclude collision/error
    - ↳ if MAC frame starts during this time, waits to verify valid ACK

Carrier Sense Multiple Access w/ Collision Avoidance (CSMA/CA)

- ↳ hidden terminal problem
  - ↳ doesn't realize another station is transmitting
    - ↳ solved w/ RTS/CTS
      - ↳ hears at least CTS
- ↳ exposed terminal problem
  - ↳ thinks another station is transmitting
    - ↳ not an issue if using same AP
- ↳ Network Allocation Vector (NAV)
  - ↳ indicates how long channel will be busy w/ current exchange
  - ↳ virtual carrier sensing

→ 58% of bit rate is used for data pause during EIFS

PIFS = SIFS + slot time

DIFS = SIFS + 2 · slot time

EIFS = SIFS + DIFS + (ACK Time @ 1Mbps)

Transit - ISP sells access to destinations

Peering - Transit but free

Border Gateway Protocol (BGP) uses path-vector algo
- ↳ each router advertises its prefferred path to nodes

MAC Frame Format

	Bytes	
Frame Control	2	type of frame CTS, RTS, ACK, etc
Duration or ID	2	used for NAV
addr 1	6	MAC addr of reciever
addr 2	6	MAC addr of sender
addr 3	6	for extended LAN, MAC addr of source or destination device
seq control	2	frag # & higher layer frame #
addr 4	6	only used if WiFi network used as wireless bridge
frame payload	0-2312	
FCS	4	checksum

Djikstra's Algorithm

- ↳ link state, each router provides list [neighbor | dist to neighbor]
- ↳ computes recursively set of P(k) nodes closest to root
  - ↳ P(n) = P(n-1) + closest node to P(n-1)
- ↳ one pass, steps = # nodes
- ↳ only positive weights

Steiner Tree is NP hard

Improving Reliability in multicast

$P_1, P_2 \to \{P_1, P_2, C\}$
- bit by bit addition 2 2
- header specifies which packets

$C_1, C_2, \ldots, C_m$; $P_1, P_2, \ldots, P_n$

m ≥ n

Bellman-Ford → share routing tables periodically

- ↳ nodes regularly send neighbors current estimate of shortest dist. to other routers
  - relaxation: trying to find better path thru edge
- ↳ Bad News Travels Slowly
  - ↳ takes time to update when link fails

|V|-1 iterations

Link State Advertisement

Ad Hoc Networks

- ↳ all nodes talk directly to each other
- ↳ AODV
  - ↳ on demand
  - ↳ asks neighbors recursively
- ↳ OLSR
  - ↳ proactive
- ↳ Ant Routing
- ↳ Geographic
  - ↳ shortest next hop

Network needs
- ↳ connectivity - all nodes can talk to each other
- ↳ broadcast capability - router can send msg to reach all devices

subnet addr specifies leading bits common to all IPs

ARP gets MAC addr. given IP addr. Not forwarded by routers uses broadcast

DIVIT RAWAL 12

Bianchi Markov Chain Model

$\{s(k), b(k)\}$

- s - prev attempts @ sending this packet @ epoch k
- b - backoff counter @ epoch k

for cases where n devices have data to transmit (or n-1 devices + AP)

$W_i = 2^i(CW_{min}+1)$

$CW_{max}+1 = 2^m(CW_{min}+1)$

when station transmits, it is successful w/ probability $(1-p)$, collides w/ probability p

probability that MC is in state of form $(i,0)$ ~ α

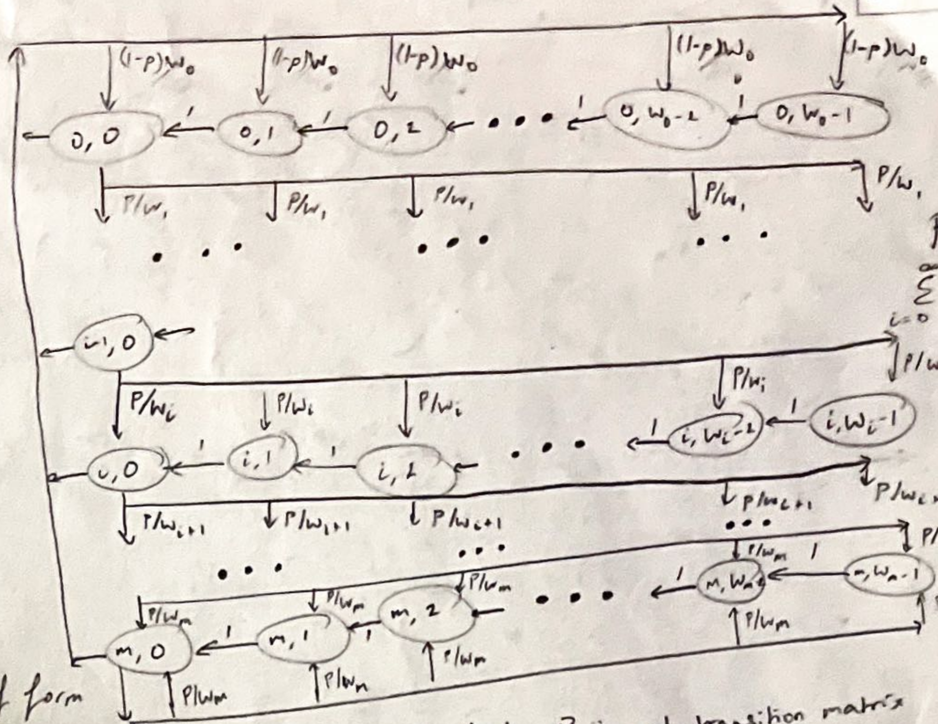

$P_{ij} = P(i \to j)$

$\sum_{i=0}^{\infty} \pi(i) = 1$

assume all n stations have prob α of transmitting ($\hat{\alpha}$ prob avg)

$1-p = (1-\alpha)^{n-1}$

function of p

$$\alpha(p) = \frac{2(1-2p)}{(1-2p)(W+1)+pW(1-(2p)^m)}$$

$W = CW_{min}+1$

successful transmission probability

$= \beta = n\alpha(1-\alpha)^{n-1}$

- length T
- bits B

Throughput $= \frac{\beta B}{T}$

Sys Throughput $= \frac{\text{\# data bits b/w consecutive epochs}}{\text{time b/w consecutive epochs}}$

let π denote invariate dist, P - prob. transition matrix

$\Rightarrow \pi P = \pi$

$\Rightarrow \pi(i,0) = p^i\pi(0,0)$ for $0 < i < m$

$\pi(i-1,0)P = \pi(i,0)$

reducible - one node is unreachable from another

irreducible - not reducible (any)

- has unique invariant dist
  - $\lim_{N\to\infty} \frac{1}{N}\sum_{n=0}^{N-1} 1\{X_n = i\} = \pi(i),\ i \in X$ — state space

recurrent iff $P(T_{ii} < \infty) = 1$, transient iff $P(T_{ii} < \infty) < 1$

positive recurrent iff $E(T_{ii}) < \infty$, null recurrent iff $E(T_{ii}) = \infty$

irreducible finite DTMC is always + recurrent

if one state of irreducible DTMC is positive/null/transient recurrent, the whole DTMC is

$d(i) = GCD\{n \geq 1 \mid P^n(i,i) > 0\}$

- periodic if $d > 1$, else if $d = 1$ aperiodic

for each node, visit unvisited neighbors

- pick next node based on smallest dist/weight

~~Djikstra's Algo~~ Djikstra's Algorithm

Internetworking

each device must

- have IP
- know subnet addr
- know gateway router IP
- IP addr of DNS server
- implement ARP

DHCP loans IP addr

NAT enables reusing of private IPs

- use port #s of transport protocol

Multiplexing

TDM

- time slots

FDM

CDM

Inf geo sum

$\frac{a}{1-r}$ — a: first term, r: multiplier

$[\pi_n] \begin{bmatrix} \alpha \\ \beta \quad P \\ \vdots \end{bmatrix} = [\pi_{n+1}]$

M/M/1

$L = \sum_{n\geq 0} n(1-\rho)\rho^n = \frac{\rho}{1-\rho} = \frac{\lambda}{\mu-\lambda}$

avg # packets

$W(\text{delay}) = \frac{1}{\mu-\lambda}$

conservation principle for inv. dist

flow in = flow out for appropriate boxes

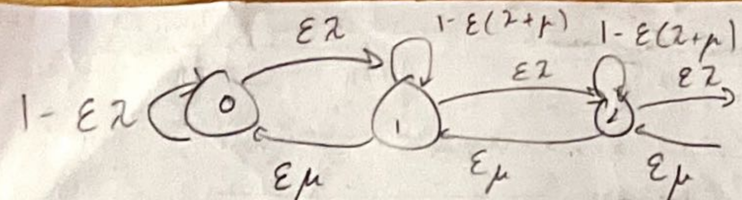

if $T \sim Expo(\lambda)$

$\rightarrow f_T(t) = \lambda e^{-\lambda t}$ exp^{ectation} $= \frac{1}{\lambda}$ variance $= \frac{1}{\lambda^2}$

if $X \sim Poisson(\mu)$

$\rightarrow P(X=k) = e^{-\mu} \frac{\mu^k}{k!}$

variance $= \mu$ expectation $= \mu$

cut 1

$\lambda \varepsilon \pi(0) = \mu \varepsilon \pi(1)$

cut $n+1$: $\lambda \varepsilon \pi(n) = \mu \varepsilon \pi(n+1)$

$\Rightarrow \pi(n+1) = \frac{\lambda}{\mu} \pi(n) = \left(\frac{\lambda}{\mu}\right)^{n+1} \pi(0)$ by Markov Prop

$\sum_{i=0}^{\infty} \left(\frac{\lambda}{\mu}\right)^{n+1} \pi(0) = 1$ law of total prob.

μT is # packets sys can serve

$\rho \mu T$ is # packets sys actually serves $= \lambda T$

sum of inf. geo. series $\left[\frac{\pi(0)}{1 - \frac{\lambda}{\mu}} = 1 \Rightarrow \pi(0) = 1 - \frac{\lambda}{\mu} \Rightarrow \pi(n) = \left(\frac{\lambda}{\mu}\right)^n \left(1 - \frac{\lambda}{\mu}\right)\right.$

avg. val of $X (X \geq 0) = \sum_{n=0}^{\infty} n P(X=n)$

for $0 < r < 1$

$\sum_{n=0}^{\infty} n r^n = \frac{r}{(1-r)^2}$

$\rho = \frac{\lambda}{\mu} \Rightarrow \pi(n) = (1-\rho)\rho^n$

$L = \sum_{i=0}^{\infty} i (1-\rho) \rho^i$

$= (1-\rho) \sum_{i=0}^{\infty} i \rho^i = \frac{\rho}{1-\rho}$

$\Rightarrow \frac{\lambda}{\mu - \lambda}$

RIP is least # of hops

DHCP returns

IP of device, IP of gateway router, IP of DNS, subnet mask

avg. Delay $= \frac{L}{\lambda}$

$= \frac{1}{\mu - \lambda}$

Transport

- Software ports
  - 1-65535
    - 1-1023
      - well-known
      - correspond to fixed processes
    - 1024-49451
      - registered
      - specific applications by companies

UDP

- delivers individual packets
- unreliable delivery

TCP

- delivers a byte stream
- reliable

Go Back N - TCP Tahoe

- @ any time, source may have sent up to N packets not ACKed yet
- dest responds to each packet w/ packet # it expects next in order
- if ACK fails to arrive after some time, retransmit that packet + next N-1 packets
- if no err. - N packets/RTT
  => N × stop & wait throughput

Selective ACKs

- indicate positive ACK for packets recieved
  - avoids retransmission unnecessary

TCP-SACK

AIMD

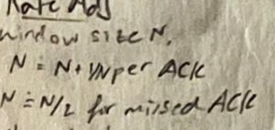

x & y increment until $x+y > C$ (packets drop), then each divides by 2 & repeats

Rate Adj

window size N,
$N = N + 1/N$ per ACK
$N \doteq N/2$ for missed ACK
slow start - start by doubling size every RTT until missed ACK or threshold
→ $N \doteq N+1$ per ACK

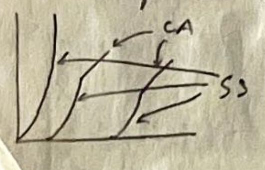

TCP States

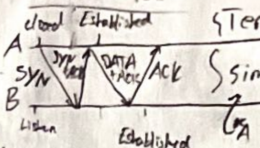

{Termination} { similar }

A waits in case B retransmits FIN & A must ACK again

Stop & Wait

→ wait for T sec for ACK, if not recieved, retransmit

rate = $\frac{\text{packet size}}{T}$

0	16	32
src port	dest port	
UDP len	checksum	
payload		

Header Info

SYN, FIN - establish/terminate TCP conn
ACK - set when Acknowledgement field is valid
URG - urgent data
→ Urgent Ptr shows where non-urgent data starts
PUSH - don't wait to fill segment
RESET - abort connection

Timers

timeout $= A_n + 4D_n$
A_n = avg. ACK time for packets w/o outstanding ACKs
D_n = deviation around A_n
$A_{n+1} = (1-b)A_n + bT_{n+1}$
$D_{n+1} = (1-b)D_n + b|T_{n+1} - A_{n+1}|$
$n \geq 1,\ b < 1$

$\sigma^2 = \frac{\rho}{(1-\rho)^2}$

fast retransmit - retransmits after 3 duplicate ACKs, avoiding wait for timeout } TCP Reno

fast recovery - when packets lost, lowers window size, continues transmitting, relying on fast retransmit to fill gaps & increase window size once sender's window size reaches $W/2 - 1$, can send all packets, then $W \doteq W/2$ } TCP Selective Reno

Flow Control

respond w/ RAW (Reciever Advertised Window)
OUT = outstanding bytes
sends min(RAW-OUT, W-OUT)

TCP Vegas

diff = W/minRTT - actual rate
$\alpha < \beta$
if $\alpha >$ diff || diff $> \beta$, change window linearly
clobbered in presence of Reno users

Models

arc is a pair of vertices
each arc(i,j) has capacity $C(i,j)$
cut - subset of nodes
capacity of cut is sum of arc capacities from S to S^c
$R(i,j) = R(i,k) + R(k,j)$
$R(i,j) \leq C(i,j)$ for arc(i,j)

PASTA

→ poisson arr. see time avgs

Graphs

max-flow min-cut
→ max flow from $v \to w$ = min capacity of cuts where $v \in S, w \in S^c$
graph is complete if any two nodes are connected by a link
max deg of graph = max degree of vertices
coloring - st no neighbors have same color
→ # colors - chromatic #
clique - set of vertices that is pairwise connected
→ all must have diff color
→ ⇒ chromatic # ≥ size(largest clique)

Brook's Thm

chromatic # if connected graph w/ max degree Δ is at most Δ except for odd cycles & complete graphs ⇒ = Δ+1

conflict graph - link represents two nodes cannot transmit @ same time
→ same color can transmit @ once
→ independent set is all 1 color
→ maximal indep. set if every other node is attached to one of the nodes in the set

for packets arriving @ nodes 1-J @ rate λ_i. $\{S_1 \dots S_K\}$ be set of maximal indep. sets
$\lambda = \{\lambda_1 \dots \lambda_J\}$
$p = \{p_1 \dots p_K\}$ be fraction of time each node transmits
node i transmits for fraction of time $= \sum_{k=1}^{K} p_k 1\{i \in S_k\}$
⇒ λ feasible iff $p_k \geq 0\ \forall k$ & $\sum_{k=1}^{K} p_k = 1$ st. $\lambda_i \leq \sum_{k=1}^{K} p_k 1\{i \in S_k\},\ i = 1 \dots J$

Queues

let $0 < \epsilon \ll 1$ and observe X_t packets in M/M/1 Queue for $t = 0, \epsilon, 2\epsilon \dots$
⇒ $P(n, n+1) \approx \lambda\epsilon$, $P(n, n-1) \approx \mu\epsilon$, $P(n,n) \approx 1 - \lambda\epsilon - \mu\epsilon$
⇒ $\pi(n) = (1-\rho)\rho^n,\ n > 0,\ \rho = \lambda/\mu$ ⇒ $L = \sum_{n \geq 0} n(1-\rho)\rho^n = \frac{\rho}{1-\rho} = \frac{\lambda}{\mu - \lambda}$ assume $\lambda < \mu$
⇒ $W = \frac{1}{\mu - \lambda}$

Jackson Networks

J Queues C classes
class c goes through $S(c) \subset \{1 \dots J\}$ ⇒ $\lambda = \sum_{c=1}^{C} \lambda_c$
rate of customers through queue $j = \gamma_j = \sum_{c=1}^{C} \lambda_c 1\{j \in S(c)\}$
assume $\gamma_j < \mu_j$

$W_c = \sum_{j \in S(c)} \frac{1}{\mu_j - \gamma_j}$, $L_j = \frac{\gamma_j}{\mu_j - \gamma_j}$ avg. delay $1/(\mu_j - \lambda_j)$ through j as if j was M/M/1

assuming two queues
⇒ $\min \{T = \frac{p}{\mu_1 - \lambda_1} + \frac{1-p}{\mu_2 - (1-p)\lambda_2}\}$
⇒ $p = \frac{\lambda\sqrt{\mu_1} + \mu_1\sqrt{\mu_2} - \mu_2\sqrt{\mu_1}}{\lambda(\sqrt{\mu_1} + \sqrt{\mu_2})}$

$u(x) = \begin{cases} \frac{x^{1-\alpha}}{1-\alpha} & \text{for } \alpha \geq 0,\ \alpha \neq 1 \\ \log(x) & \text{if } \alpha = 1 \end{cases}$

Congestion Control

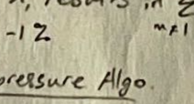

$g_a(x) = x_1 + x_2 - 1 \leq 0$ & $g_b(x) = x_1 + x_3 - 1 \leq 0$
total throughput $= x_1 + x_2 + x_3$
utility $u(x)$ is concave & increasing for of $x > 0$
choose x s.t. $\underbrace{u(x_1) + u(x_2) \dots}_{f(x)}$ maximize subject to constraints

max-min fair rate - x s.t. not possible to increase x_i w/o decreasing x_j where $x_j \leq x_i$ & max $f(x)$

proportionally fair - change of 1% in x_1 results in $\sum_{m \neq 1} x_m$ change of $\leq -1\%$

Backpressure Algo

pressure for each node based on size of packet queues
↑pressure ⇒ recieving > sending
↓pressure ⇒ sending > recieving
route to neighbor w/ highest positive pressure
max-sum fairness - maximize sum of individual rate

Primal Problem

let $f(x)$ be concave & $g_j(x)$ be convex for $j = a, b$, maximize $f(x)$ subject to $g_j(x) \leq 0$
iff x^* ~~...~~ satisfies constraint & x^* maximizes $L(x, \lambda^*) = f(x) - \lambda_a^* g_a(x) - \lambda_b^* g_b(x)$
for $\lambda_a^*, \lambda_b^* \geq 0$ s.t. $\lambda_j^* g_j(x^*) = 0$

dual of
& if $L(x(\lambda), \lambda) = \max_x L(x, \lambda)$
⇒ λ_a^*, λ_b^* minimize $L(x(\lambda), \lambda)$

Primal Problem: $\max_{x \in S} f(x)$

subject to

$h_i(x) \geq 0,\ i = 1, 2, \ldots, I$

$g_j(x) = 0,\ j = 1, 2, \ldots, J$

TDMA

FDMA

CDMA code

↳ Lagrangian

$L(x, \lambda, \mu) = f(x) + \sum_{i=1}^{I} \lambda_i h_i(x) + \sum_{j=1}^{J} \mu_j g_j(x),\ \lambda_i \geq 0\ \forall i$

Lagrangian Dual Fn ← convex, $D(\lambda, \mu) \geq f^*$ where $f^* = \max_{x \in S} f(x)$

$D(\lambda, \mu) = \sup_{x \in S} L(x, \lambda, \mu)$

upper bound on max of problem

Lagrange Dual Problem ← gives best upper bound

$\inf_{\lambda \geq 0, \mu} D(\lambda, \mu) = d^*$ ← let

generally $d^* \geq f^*$, $d^* - f^*$ → duality gap

strong duality ⇒ $d^* = f^*$

Slater Conditions

↳ strong duality holds if

- $f(x)$ is concave & $h_i(x)$ are convex
- $g_j(x)$ are affine - of form $ax+b$
- $\exists\, x \in$ relative interior of S s.t. $h_i(x) < 0\ \forall i$ ∧ $g_j(x) = 0\ \forall j$

KKT Conditions

assume f, h_i, g_j are differentiable & Slater conditions are satisfied

let x^* be a point that satisfies constraints

x^* is a global maximizer iff $\exists\ \lambda_i^* \geq 0, \mu_j^*$ s.t.

$$\frac{\partial f}{\partial x_k}(x^*) - \sum_i \lambda_i^* \frac{\partial h_i}{\partial x_k}(x^*) + \sum_j \mu_j^* \frac{\partial g_j}{\partial x_k}(x^*) = 0 \quad \forall k$$

$\lambda_i^* h_i(x^*) = 0\ \forall i$

complementary slackness condition

Cellular

- divide geographical area into cells
- users in the cell communicate w/ its base station
- neighboring cells use diff freqs to minimize interference
- user picks BS w/ strongest signal
- user picks a timeslot to try to connect to BS
  - ↳ if req does not collide, allocate slots
- home registry - records loc
  - ↳ periodically send to BS
- modulation scheme determines rate - shared by all users in cell

LTE

FDD (freq Division Duplex) ← more common

TDD (time Div. Duplex) ← used when less available freq

open-loop pwr control - estimate necessary signal strength by BS signal strength

closed-loop pwr control - feedback from BS on sig. strength

LTE (cont.)

- eNBs (Base Stations) connected by wired IP network (EPC) ← control plane & data plane
- Orthogonal Freq Div. Multiplexing (OFDM)

Slow fading - $\frac{\Delta \text{channel}}{\Delta \text{signal}}$ is small

fast fading - $\frac{\Delta \text{channel}}{\Delta \text{signal}}$ is large

TTI - subframe ← trans. time int. 2 slots

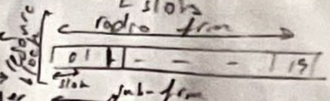

smallest unit allocated to user

MIMO

- beamforming - direct main beams b/w trans-recieve pair for higher SNR
- spatial diversity - multiple copies for redundancy (used in low SNR)
- spatial multiplexing - diff data b/w trans-r pairs for higher throughput (used in high SNR)

1 slot - 7 modulated symbols, 12 subcarriers } resource block ← smallest unit of allocation to usr

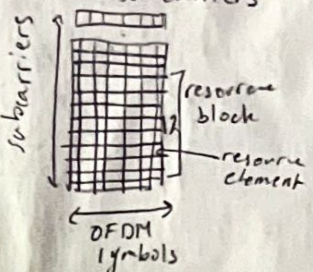

DL: OFDMA

UL: SC-FDMA

↳ lower peak-to-avg pwr ratio

symbol time = $\frac{1}{\text{subcarrier spacing}}$

LTE Sys Capacity

ex: 10 MHz channel, 600 subcarrier, ⇒ 50 PRBs per slot

QAM 64 → 6 bits per RE

84 RE in a PRB

Popular Scheduling Schemes

Priority = T^a / R^b

T - data rate potentially available by this station in curr slot

R - historic avg. data rate

Alternatives:

- a = 0, b = 1 (round-robin) (eq. throughput to each usr)
- a = 1, b = 0 (best condition first) (maximizes sys. throughput)
- a = 1, b = 1 (proportionally fair) (equal share of resources to each usr)

Clos Networks

- 3 layers of nodes
- each node in i connected to each node in i+1
- can build larger switches using smaller recursively

Clos (IN, N_1, N_2, N_3, OUT)

consider as circuit switch network

- switch is nonblocking if it can have all pt-pt connections b/w free (unused) inputs & free outputs
- SNB (strictly nonblocking): if any new connection from free input to free output can be made w/o having to modify ongoing connections
- nonblocking but not SNB is RNB (rearrangeably NB)

Thms (Clos)

(1) if $N_2 \geq$ IN + OUT − 1 ⇒ SNB

if SNB ⇒ $N_1 \geq$ OUT & $N_3 \geq$ IN

(2) consider NW built from RNB modules

NW is RNB iff $N_2 \geq \max(\text{IN}, \text{OUT})$

$\Delta f = \frac{\Delta v}{c} f_0$

Δf: freq shift; Δv: reciever rel. to src; c: signal speed; f_0: original

doppler spread related

coherence time - how long channel can be considered to remain unchanged

Path Loss: $P_r = P_t P_0 \left(\frac{d_0}{d}\right)^{\alpha},\ d \geq d_0$

P_r: recieved pwr; P_t: transmitted pwr; P_0: pwr @ dist d_0 w/ unit transmit pwr; d: dist; α: path loss exp

$PL(d) = PL(d_0) + 10\alpha \log\left(\frac{d}{d_0}\right) + X$

lognormal shadowing $N(0, \sigma)$

Input Queueing: Head-of Line Blocking

↳ fixed size packets
↳ no speed up (max 1 pack in 1 pack time)
↳ at most 1 pack to each output port in 1 pack time

HOL blocking reduces sys throughput to max val .586 under symmetric traffic conditions

Overcoming HOL Blocking

- higher depth of look-ahead
- 1 pack / pack time
- VOQ (Virtual Output Queueing)
  - ↳ separate input queues by class/output queue
  - ↳ for $N \times N$ switch, λ is $N \times N$ matrix
    - ↳ λ_{ij} = arrival rate @ i destined for j
    - ↳ λ is supportable if $\exists$ switch scheduling algo. s.t. $\lim_{L\to\infty} \lim_{t\to\infty} P(|q(t)| \geq L) = 0$
      sum of all queue lens
    - → let $C = \{\lambda : \lambda \geq 0, \sum_{i=1}^{N} \lambda_{ij} \leq 1 \ \forall j, \ \& \ \sum_{j=1}^{N} \lambda_{ij} \leq 1 \ \forall i\}$
    - → Thm. if $\lambda \notin C$, no switch scheduling algo can support it

Max Weight Scheduling

M^h = h^{th} matching & $M^h = 1$ if i matched to j, 0 otherwise

Switch finds matching $M(t)$ @ t s.t.
$M(t) \in \arg\max_{M^h} \sum_{ij} q_{ij}(t) \cdot M^h(ij)$ (# of pack in queue)
& transfers pack from VOQ(i,j) to j if there is a pack in this queue

Thm. MaxWeight Scheduling can support any arr. rate matrix λ s.t. $(1+\epsilon) * \lambda \in C$ for some $\epsilon > 0$

QoS

Leaky Bucket

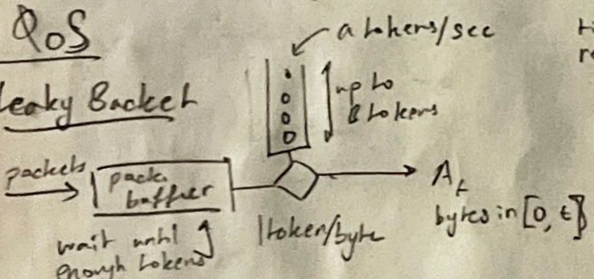

- Traffic shaping by limiting rate & burst size

$A_{s+t} - A_s \leq B + A_t$ bytes for $s, t \geq 0$

Delay Bounds

- N input streams for a buffer, each leaky bucket shaped
- each stream carries at most P bytes, and packets enter contiguously

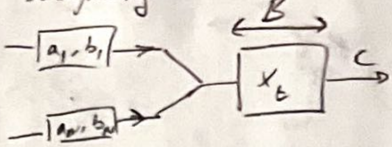

Thm assume $a_1 + \dots + a_n \leq C$,
(a) $x_t \leq b_1 + \dots + b_n + NP$, $t \geq 0$
(b) delay of every packet is bounded by $(b_1 + \dots + b_n + NP + P)/C$

Generalized Processor Sharing (GPS)

↳ packets classified into K classes & wait in corresponding FIFO queues until transmission
↳ each class has weight w_k
↳ scheduler serves HOL packets @ rate prop. to weight
↳ service rate of k is $\frac{w_k C}{W}$ (line rate out; sum of weights of queues currently backlogged)

Thm assume traffic of k is regulated w/ (a_k, B_k) s.t. $\rho_k = \frac{w_k C}{\sum_j w_j} \geq a_k$
⇒ backlog of k never exceeds B_k, queueing delay never exceeds B_k / ρ_k

Weighted Fair Queueing (WFQ)

↳ approximates GPS
↳ classified & queued as in GPS
↳ 1 pack at a time @ rate C
↳ when completes a transmission, starts transmitting GPS selected next pack.

Thm let F_k & G_k be completion times of pack P_k under WFQ & GPS respectively assume trans. time of pack $\leq T$
⇒ $F_k \leq G_k + T \quad k \geq 1$

Optimization

sup - least upper bnd
inf - greatest lower bnd

Utility Fns & Fairness

max-sum
↳ w/ $\alpha = 0$ maximize sum of indv. rate

max-min
↳ maximize min of allocations
$\alpha \to \infty$

Jackson Network

avg delay = avg pack / ext arr. rate

total arr. rate into i
$\lambda_i = \gamma_i + \sum_{j=1}^{J} \lambda_j r(j,i)$ for $i = 1, 2 \dots J$
γ_i external arrival rate into i

when customer leaves i, he joins j w/ prob $r(i,j)$ & leaves nw w/ prob $1 - \sum_{j=1}^{J} r(i,j)$

for $t \geq 0$, $X_t = (X_{1t} \dots X_{Jt})$
X_{it} = # cust @ node i (total)
multidim. CTMC

Thm^{sol} $(\lambda_1 \dots \lambda_J)$ s.t. $\lambda_i \leq \mu_i$ for $i = 1, \dots, J$; then CTMC X_t admits

$\pi(x_1 \dots x_J) = \pi(x_1) \dots \pi(x_J)$

$\forall j$, $\pi_j(n) = (1 - \rho_j)\rho_j^n$ for $n \geq 0$ & $\rho_j = \frac{\lambda_j}{\mu_j}$

Min Potential Delay Fairness

$U_r(x_r) = -1/x_r$

w/ $\alpha = 2$, minimize $\sum_r 1/x_r$
delay of pack of unit size

Graphs

any 2 imply 3rd
- G is connected
- G is acyclic/does not contain cycle
- G has n-1 edges

CTMC

X is a countable set

blocking prob $\pi(s) = \dfrac{\rho^s/s!}{\sum_{n=0}^{s} \rho^n/n!}$, $\rho = \dfrac{\lambda}{\mu}$

rate matrix

$Q = \{q(i,j) ; i,j \in X\} \in \mathbb{R}$

is s.t.

$0 \le q(i,j) < \infty$ for $i \ne j \in X$

$q(i) = -q(i,i) = \sum_{j \ne i} q(i,j) < \infty \ \forall i \in X$

let π be init dist.

if $X_0 = i$, select random time τ from exp dist w/ rate $q(i)$

let $X_t = i$ for $0 \le t \le \tau$

@ $t = \tau$, X_t jumps to j indep. s.t.

$P(X_\tau = j \mid X_0 = i \text{ and } \tau) = \dfrac{q(i,j)}{q(i)} = \Gamma(i,j) \ \forall j \ne i$

construction resumes @ $X_\tau = j$ same way as $X_0 = i$ indep. of past

Q on X is irreducible iff $q(i) > 0 \ \forall i \in X$

& transition prob matrix $\Gamma(i,j) = \dfrac{q(i,j)}{q(i)}$ iff $i \ne j$, 0 if $i = j$

is irreducible

Thm

X_t is irreducible CTMC over X w/ rate matrix Q & init dist π

π is inv (ie $P(X_t = i) = \pi(i)$)

iff

$\sum_{i \in X} \pi(i) q(i,j) = 0 \ \forall j \in X$ (or $\pi Q = 0$)

Thm

an irreducible CTMC has either no or exactly 1 inv. dist

↳ if finite, then exact 1

must be positive recurrent

Thm if CTMC has exactly one inv. dist π.

then for any init dist

$\lim_{t \to \infty} P(X_t = i) = \pi(i) \ \forall i \in X$

&

$\lim_{T \to \infty} \frac{1}{T} \int_0^T 1\{X_s = i\} ds = \pi(i) \ \forall i \in X$

Poisson

arr rate λ, let $X_0 = 0$, let $X_t = \#$ arr. in $[0, t]$

$\Rightarrow X$ is a CTMC w/ $q(i,j) = -\lambda$ & $q(i,i+1) = \lambda$ for $i \ge 0$

$P(\# \text{arr} = k) = \dfrac{e^{-\lambda t} (\lambda t)^k}{k!}$, $k = 0, 1, \ldots$