Random access protocols, scheduling in routers and wireless networks

Laurent Massoulié

Inria

January 11, 2021

Aloha with finitely many stations

Stations $s \in \mathcal{S}$, $|\mathcal{S}| < \infty$

- New arrivals at station s in slot n: $A_{n,s} \in \mathbb{N}$, $\{A_{n,s}\}_{n\geq 0}$ i.i.d.
- Probability of transmission by s if message in queue: p_s
- Source of randomness: $\{B_{n,s}\}_{n\geq 0}$ i.i.d., Bernoulli(p_s)
- Transmits iff $B'_{n,s} = 1$ where $B'_{n,s} = B_{n,s}\mathbb{I}_{L_{n,s}>0}$

Queue dynamics

$$L_{n+1,s} = L_{n,s} + A_{n,s} - B'_{n,s}\prod_{s'\neq s}(1 - B'_{n,s'})$$

Aloha with finitely many stations

Assume $\forall s,\ 0 < \mathbb{P}(A_{n,s} = 0) < 1$ and $\forall s,\ 0 < p_s < 1$ Then chain is irreducible and aperiodic

Sufficient condition for ergodicity

$$\forall s, \mathbb{E}(A_{n,s}^2) < +\infty \text{ and } \lambda_s := \mathbb{E}(A_{n,s}) < p_s \prod_{s' \neq s} (1 - p_{s'})$$

Sufficient condition for transience

$$\forall s, \lambda_s > p_s \prod_{s' \neq s} (1 - p_{s'})$$

Aloha with finitely many stations

Symmetric case $\lambda_s = \lambda/|\mathcal{S}|$, $p_s \equiv p$:

Recurrence if $\lambda < |\mathcal{S}|p(1-p)^{|\mathcal{S}|-1}$

Transience if $\lambda > |\mathcal{S}|p(1-p)^{|\mathcal{S}|-1}$

$\Rightarrow$ To achieve stability (ergodicity) for fixed λ, need $p \to 0$ as $|\mathcal{S}| \to \infty$

Impractical! (Collisions take forever to be resolved)

Aloha with infinitely many stations

Many stations, very rarely active (just one message)

- A_n new messages in interval n, $\{A_n\}_{n\geq 0}$ i.i.d.
- Source of randomness $\{B_{n,i}\}_{n,i\geq 0}$ i.i.d., Bernoulli (p)
- Queue evolution

$$L_{n+1} = L_n + A_n - \mathbb{I}_{\sum_{i=1}^{L_n} B_{n,i}=1}$$

- Assumption $0 < \mathbb{P}(A_n = 0) < 1$ ensures irreducibility (and aperiodicity)

Aloha with infinitely many stations

Abramson's heuristic argument

For $A_n \sim$ Poisson(λ), Nb of attempts per slot $\approx$ Poisson(G) for unknown G

Hence successful transmission with probability Ge^{-G} per slot

Solution to $\lambda = Ge^{-G}$ exists for all $\lambda < 1/e$

Hence "Aloha should be stable (ergodic) whenever $\lambda < 1/e$"

Theorem: Instability of Aloha

With probability 1, channel jammed forever ($\sum_{i=1}^{L_n} B_{n,i} > 1$) after finite time. Hence only finite number of messages ever transmitted.

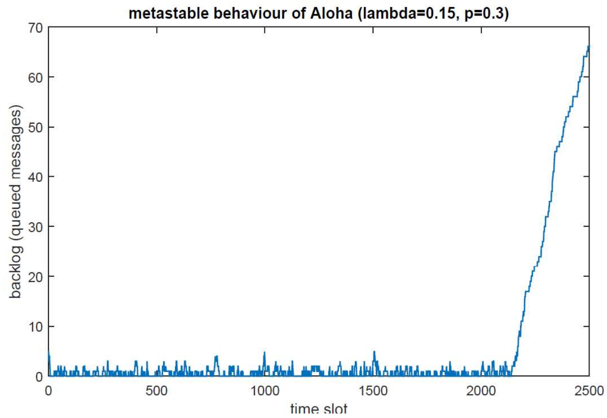

Fixing Aloha: richer feedback

Assumption: L_n known

Backlog-dependent retransmission probability $p_n = 1/L_n$
Then system ergodic if $\lambda := \mathbb{E}(A_n) < \frac{1}{e} \approx 0.368$

Denote $J_n = \{0, 1, *\}$ outcome of n-th channel use
(0: no transmission. 1: single successful transmission. $*$: collision)

Weaker assumption: channel state J_n heard by all stations

Backlog-dependent retransmission probability $p_n = 1/\hat{L}_n$, where estimate $\hat{L}_n$ computed by

$$\hat{L}_{n+1} = \max(1, \hat{L}_n + \alpha \mathbb{I}_{J_n = *} - \beta \mathbb{I}_{J_n = 0})$$

renders Markov chain $(L_n, \hat{L}_n)_{n \geq 0}$ ergodic for suitable $\alpha, \beta > 0$ if $\lambda := \mathbb{E}(A_n) < \frac{1}{e} \approx 0.368$

Fixing Aloha: richer feedback

With same ternary feedback $J_n = \{0, 1, *\}$, can stability hold for $\lambda > 1/e$?

Yes: rather intricate protocols have been invented and shown to achieve stability up to $\lambda = 0.487$

Largest λ for which some protocol based on this feedback is stable? Unknown (only bounds)

Ethernet and variants

Return to Acknowledgement-based feedback (only listen channel's state after transmission)
Variant of exponential backoff: transmit with probability 2^{-k} after k collisions
Assume $A_n \sim$ Poisson (λ)

Theorem: instability of Ethernet's variant

For any $\lambda > 0$, (modification of) Ethernet is transient.

Weaker performance guarantees

Ethernet and its modification are such that with probability 1:
For $\lambda < \ln(2) \approx 0.693$, infinite number of messages is transmitted
For $\lambda > \ln(2)$, only finitely many messages are transmitted

Unsolved conjecture

No acknowledgement-based scheme can induce a stable (ergodic) system for any $\lambda > 0$.

Conclusions on Random Access Protocols

Analysis of Aloha was useful to guide design of Ethernet.
Negative results in theory (no ergodicity), both for Aloha and Ethernet, yet...
...In practice, Ethernet and Wi-Fi's 802.11x protocols perform well

- Finite number of stations helps
- Time to instability could be huge ("metastable" behavior)
- Only small fraction of channel time used for random access collision resolution:
  Once station "wins" channel access, others wait till its transmission is over

$\rightarrow$ Alternative protocols based on ternary feedback have not been used

Scheduling in cross-bar switches

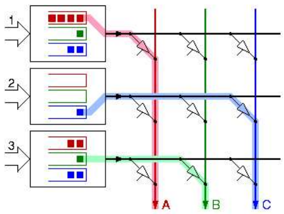

- Switch with N input and N output ports
- Time slot n: $A_n(i,j)$ packets arrive at input port i, destined to port j
- Transmission: permutation $\sigma_n \in \mathcal{S}_N$, symmetric group, matches input port i with output port $\sigma_n(i)$

$\Rightarrow$ How to choose σ_n to ensure ergodicity, i.e. stationary regime instead of queue blowup?

Scheduling downlink wireless transmissions

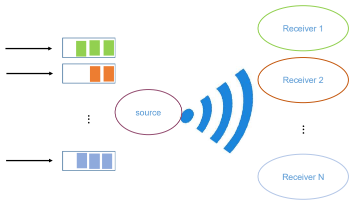

- Wireless source to send packets to wireless receivers
- Time slot n: $A_n(r)$ packets arrive at source for receiver r
- Wireless medium conditions change in each slot n: $S_n(r) =$ number of packets that could be sent to receiver r if it was chosen then

$\Rightarrow$ How to choose which receiver to schedule based on queue lengths (backlogs) and medium condition to ensure ergodicity, i.e. stationary regime instead of queue blowup?

Max-Weight scheduling

- Traffic types $r \in \mathcal{R}$, i.i.d. arrivals: $A_n(r) \in \mathbb{N}$ in slot n
- i.i.d. set $\mathcal{S}_n \subset \{0, \ldots, s_{max}\}^{\mathcal{R}}$ of feasible services in slot n
- $X_n(r)$: backlog of type r requests at end of slot n
- Evolution equation $X_{n+1}(r) = (X_n(r) - s_n(r))^+ + A_{n+1}(r)$, where $s_n \in \mathcal{S}_n$
- (w, α)-Max-weight scheduling rule for $w_r, \alpha > 0$:

  Choose $s_n \in \text{Argmax}_{s \in \mathcal{S}_n} \left\{ \sum_{r \in \mathcal{R}} w_r X_n(r)^{\alpha} s(r) \right\}$

Max-Weight scheduling: ergodicity properties

- Assume (to ensure irreducibility on set of states reachable from 0) $\mathbb{P}(\forall r \in \mathcal{R}, A_n(r) = 0) \in]0,1[$, $\forall r \in \mathcal{R},\ \mathbb{P}(\exists s \in \mathcal{S}_n : s(r) > 0) > 0$
- Let **schedulable region** $\mathcal{C}$ be set of vectors $x \in \mathbb{R}_+^{\mathcal{R}}$ such that $\exists z^{(\mathcal{S})} \in \text{env}(\mathcal{S}) :\ \forall r \in \mathcal{R}, x_r \le \sum_{\mathcal{S} \subset \{0,\ldots,s_{max}\}^{\mathcal{R}}} \mathbb{P}(\mathcal{S}_n = \mathcal{S}) z^{(\mathcal{S})}(r)$ where $\text{env}(\mathcal{S})$: convex hull of set $\mathcal{S}$
- Let $\rho_r := \mathbb{E}(A_n(r))$

Theorem

If $\mathbb{E}A_n(r)^{1+\alpha} < +\infty$ and for some $\epsilon > 0$, $(\rho_r + \epsilon)_{r \in \mathcal{R}} \in \mathcal{C}$, then process $\{X_n\}_{n \in \mathbb{N}}$ is ergodic.
Conversely, if $\rho \notin \mathcal{C}$, then for any strategy (max-weight or other), process $\{X_n\}_{n \in \mathbb{N}}$ is transient.

Comments

- Maximizes set of offered loads ρ for which ergodicity holds (for ρ on frontier of $\mathcal{C}$, chain at best null-recurrent)
- Does not require explicit learning of either ρ (statistics of request arrivals) or $\mathcal{S}_n$ (statistics of time varying capacity)
- Switch scheduling: convex enveloppe of permutation matrices $M_\sigma = (\mathbb{I}_{j=\sigma(i)})_{i,j \in [N]}$ = Doubly stochastic matrices, i.e. $M \in \mathbb{R}_+^{N \times N}$ such that

$$\forall i \in [N], \sum_{j \in [N]} M_{ij} = 1 = \sum_{j \in [N]} M_{ji}$$

(Birkhoff-von Neumann theorem)
Hence switch process ergodic if and only if

$$\forall i \in [N], \sum_{j \in [N]} \mathbb{E}(A(i,j)) < 1 \& \sum_{j \in [N]} \mathbb{E}(A(j,i)) < 1.$$

Proof elements

- Ergodicity: Use Foster's criterion with Lyapunov function $V(X) = \sum_{r \in \mathcal{R}} w_r \frac{X_r^{1+\alpha}}{1+\alpha}$
- Transience: for $\rho \notin \mathcal{C}$, use convex separation theorem:

$$\exists b \in \mathbb{R}^{\mathcal{R}}, \delta > 0 : \forall x \in \mathcal{C}, \sum_{r \in \mathcal{R}} b_r \rho_r \geq \delta + \sum_{r \in \mathcal{R}} b_r x_r.$$

From monotonicity of $\mathcal{C}$, can choose $b_r \geq 0,\ r \in \mathcal{R}$
$\Rightarrow$ Lower bound:

$$\sum_r b_r X_n(r) \geq \sum_{m=1}^{n} \sum_{r \in \mathcal{R}} b_r A_m(r) - \sum_{m=1}^{n} \sum_{r \in \mathcal{R}} b_r s_n(r)$$
$$\geq n \left[\sum_{r \in \mathcal{R}} b_r \left(\rho_r - \sum_{\mathcal{S} \subset \{0,\ldots,s_{max}\}^{\mathcal{R}}} \mathbb{P}(\mathcal{S}_n = \mathcal{S}) z_r(\mathcal{S}) \right) \right] + o(n)$$
$$\geq n\delta + o(n),$$

by law of large numbers and convex separation result. Hence almost surely $\lim_{n \to \infty} \sup_{r \in \mathcal{R}} X_n(r) = +\infty$

multi-hop, multipath networks

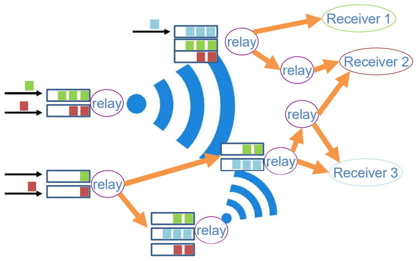

- Several traffic types, packets from each type: may be created at several network locations
- Each network location: may choose which traffic type to forward, and to which neighbor to forward it (interferences may constrain decisions at distinct locations)
- Each created packet replicated at only one location if still present; disappears when reaches its destination

Max-weight backpressure algorithm: general setup

- Abstract data types $r \in \mathcal{R}$, i.i.d. arrivals $A_n(r)$ in slot n .
  Also, let $\mathcal{R}' := \mathcal{R} \cup \{ext\}$
- Set of potential transmissions per time slot:
  $\mathcal{S} \subset \{0, 1, \ldots, s_{max}\}^{\mathcal{R} \times \mathcal{R}'}$,
  assumed decreasing, i.e. $s \le s', s' \in \mathcal{S} \Rightarrow s \in \mathcal{S}$
- $X_n(r)$: backlog of type r-data in time slot n
- Evolution equation

$$X_{n+1}(r) = X_n(r) + \sum_{r' \in \mathcal{R}} s'_n(r', r) - \sum_{r' \in \mathcal{R}'} s'_n(r, r') + A_{n+1}(r),$$

where $\{s'_n(r, r')\}_{(r,r') \in \mathcal{R} \times \mathcal{R}'}$: $s'_n(r, r') \le s_n(r, r')$ for some $s_n \in \mathcal{S}$, and:

$$X_n(r) - \sum_{r' \in \mathcal{R}'} s'_n(r, r') = \left(X_n(r) - \sum_{r' \in \mathcal{R}'} s_n(r, r') \right)^+$$

Max-weight backpressure: policy

- (w, α)-max-weight backpresssure policy, for $w_r > 0, \alpha > 0$, selects $s_n \in \mathcal{S}$ achieving

$$\mathrm{Max}_{s \in \mathcal{S}} \left\{ \sum_{(r,r') \in \mathcal{R} \times \mathcal{R}'} s(r, r')[w_r X_n(r)^\alpha - w_{r'} X_n(r')^\alpha] \right\}$$

  Backpressure from r to r': $w_r X_n(r)^\alpha - w_{r'} X_n(r')^\alpha$.
  Schedule transfers $r \to r'$ only if backpressure positive.
  By convention, $X_n(ext) = 0$.

- **Schedulable region** $\mathcal{C}$ = set of vectors $x \in \mathbb{R}_+^{\mathcal{R}}$ such that

$$\exists c \in \mathrm{env}(\mathcal{S}) : \; \forall r \in \mathcal{R}, x_r + \sum_{r' \in \mathcal{R}} c(r', r) \le \sum_{r' \in \mathcal{R}'} c(r, r').$$

Ergodicity properties

Denote $\rho_r := \mathbb{E}(A_n(r))$. Then

Theorem

If $\{X_n\}_{n\in\mathbb{N}}$ *is irreducible,* $\mathbb{E}A_n(r)^{1+\alpha} < +\infty$ *and for some* $\epsilon > 0, (\rho_r + \epsilon)_{r\in\mathcal{R}} \in \mathcal{C}$*, then* $\{X_n\}_{n\in\mathbb{N}}$ *is ergodic.*
Conversely, if $\rho \notin \mathcal{C}$*, then for any strategy (max-weight backpressure or other)* $\{X_n\}_{n\in\mathbb{N}}$ *is transient.*

Proof elements: parallel proof for Max-weight, showing ergodicity with same Lyapunov function $V(x) = \sum_r w_r \frac{x(r)^{1+\alpha}}{1+\alpha}$

Comments

- Enjoys same optimal ergodicity properties as Max-weight, in multi-hop setting with varieties of network paths to choose from
- No need to explicitly estimate traffic parameters
- Extends to case of i.i.d., rather than constant sets $\mathcal{S}_n$ of feasible transmissions
- Proposed in '93 as a practical way to schedule transmissions in wireless networks (Tassiulas-Ephremides), and as an algorithm to determine approximate solutions to multicommodity flow problems (Awerbuch-Leighton). Max-weight special case rediscovered later for switches

Takeaway messages

- Markov chain theory: framework for system and algorithm performance analysis
- Ergodicity (stability) analysis:
  $\rightarrow$ Determines for what demands system stabilizes into steady state
  $\rightarrow$ A "first order" performance index (know when delays remain stable, not their magnitude)
- Foster-Lyapunov criterion to prove ergodicity with adequate Lyapunov function when stationary distribution not known explicitly
- Several models for which **schedulable region** characterizes set of traffic parameters (loads per class) which make system ergodic, and for which known simple policy achieves ergodicity whenever possible with no explicit inference of traffic parameters