Two examples of queueing networks

Susceptible-Infective Epidemic propagation models

Laurent Massoulié

MSR-Inria Joint Centre

October 6, 2015

Priority queue

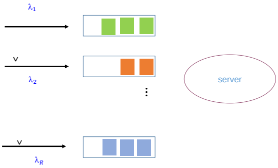

- R customer types
- Infinite queue, single server with unit capacity
- Policy: always serve customer with highest priority (lowest class index)
  Interrupt lower priority service upon higher priority arrival
  Resume interrupted service where it was stopped (FIFO per class)
- Poisson λ_r arrivals in class r; Exponential μ_r service times
  Loads: $\rho_r := \lambda_r/\mu_r$

- $X_r(t)$: number of class-r customers present at time t
- A Markov jump process with only non-zero rates

$$q_{x,x+e_r} = \lambda_r, \quad q_{x,x-e_r} = \mu_r \mathbb{I}_{x_r>0}\mathbb{I}_{x_1=\cdots=x_{r-1}=0}$$

Proposition

Process ergodic if $\rho := \sum_r \rho_r < 1$, transient if $\rho > 1$

Assume $\mu_r \equiv \mu$ and ergodicity.
Then mean number of customers at equilibrium:

$$\mathbb{E}(X_r) = \frac{\rho_r}{(1-\sum_{s<r}\rho_s)(1-\sum_{s\le r}\rho_s)}$$

Proof elements – priority queue

- Process clearly irreducible non-explosive
  Foster's criterion with $V(x) := \sum_r x_r/\mu_r \Rightarrow$ ergodic if $\rho < 1$

Proof elements – priority queue

- Process clearly irreducible non-explosive
  Foster's criterion with $V(x) := \sum_r x_r/\mu_r \Rightarrow$ ergodic if $\rho < 1$
- $\rho > 1$: with $X(0) = 0$, $X_r(t) = N_r(t) - D_r(t)$ (arrivals minus departures)
  Work spent on class r by time t: $W_r(t) \geq \sum_{m=1}^{D_r(t)} \sigma_{r,m}$
  for i.i.d. service times $\sigma_{r,m}$

Proof elements – priority queue

- Process clearly irreducible non-explosive
  Foster's criterion with $V(x) := \sum_r x_r/\mu_r \Rightarrow$ ergodic if $\rho < 1$
- $\rho > 1$: with $X(0) = 0$, $X_r(t) = N_r(t) - D_r(t)$ (arrivals minus departures)
  Work spent on class r by time t: $W_r(t) \geq \sum_{m=1}^{D_r(t)} \sigma_{r,m}$
  for i.i.d. service times $\sigma_{r,m}$
- Law of large numbers for Poisson processes: almost surely, $\lim_{t\to\infty} N_r(t)/t = \lambda_r$

Proof elements – priority queue

- Process clearly irreducible non-explosive
  Foster's criterion with $V(x) := \sum_r x_r/\mu_r \Rightarrow$ ergodic if $\rho < 1$
- $\rho > 1$: with $X(0) = 0$, $X_r(t) = N_r(t) - D_r(t)$ (arrivals minus departures)
  Work spent on class r by time t: $W_r(t) \geq \sum_{m=1}^{D_r(t)} \sigma_{r,m}$
  for i.i.d. service times $\sigma_{r,m}$
- Law of large numbers for Poisson processes: almost surely, $\lim_{t\to\infty} N_r(t)/t = \lambda_r$
- If for some r, $D_r(t) \leq \lambda_r t/2$ then $X_r(t) \geq \lambda_r t/2 + o(t)$
  Else, by Law of large numbers for $\sigma_{r,m}$,
  $\forall r, W_r(t) \geq D_r(t)/\mu_r + o(t)$

Proof elements – priority queue

- Process clearly irreducible non-explosive
  Foster's criterion with $V(x) := \sum_r x_r/\mu_r \Rightarrow$ ergodic if $\rho < 1$
- $\rho > 1$: with $X(0) = 0$, $X_r(t) = N_r(t) - D_r(t)$ (arrivals minus departures)
  Work spent on class r by time t: $W_r(t) \geq \sum_{m=1}^{D_r(t)} \sigma_{r,m}$
  for i.i.d. service times $\sigma_{r,m}$
- Law of large numbers for Poisson processes: almost surely, $\lim_{t\to\infty} N_r(t)/t = \lambda_r$
- If for some r, $D_r(t) \leq \lambda_r t/2$ then $X_r(t) \geq \lambda_r t/2 + o(t)$
  Else, by Law of large numbers for $\sigma_{r,m}$,
  $\forall r, W_r(t) \geq D_r(t)/\mu_r + o(t)$
- Since $\sum_r W_r(t) \leq t$, implies

$$\sum_r X_r(t)/\mu_r \geq \rho t - t + o(t)$$

In both cases $\max_r X_r(t) \to \infty$ almost surely

Exactly solvable bandwidth sharing network

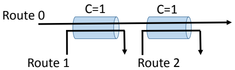

- Linear network with L unit capacity links, $L+1$ classes. Class 0 uses all links, class r uses only link r, $r \geq 1$
- Poisson ν_r arrivals, Exponential μ_r service times, loads $\rho_r = \nu_r / \mu_r$
- x_r: number of ongoing type r transfers

Exactly solvable bandwidth sharing network

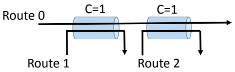

- Linear network with L unit capacity links, $L+1$ classes. Class 0 uses all links, class r uses only link r, $r \geq 1$
- Poisson ν_r arrivals, Exponential μ_r service times, loads $\rho_r = \nu_r/\mu_r$
- x_r: number of ongoing type r transfers
- Proportionally fair allocations: service rate to class $0: \ \Lambda_0 = x_0/[x_0 + y]$ where $y = \sum_{r=1}^{L} x_r$; service to class $r \geq 1 : \Lambda_r = y/(y + x_0)\mathbb{I}_{x_r > 0}$

Exactly solvable bandwidth sharing network

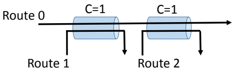

- Linear network with L unit capacity links, $L+1$ classes. Class 0 uses all links, class r uses only link r, $r \geq 1$
- Poisson ν_r arrivals, Exponential μ_r service times, loads $\rho_r = \nu_r/\mu_r$
- x_r: number of ongoing type r transfers
- Proportionally fair allocations: service rate to class $0: \ \Lambda_0 = x_0/[x_0+y]$ where $y = \sum_{r=1}^{L} x_r$; service to class $r \geq 1 : \Lambda_r = y/(y+x_0)\mathbb{I}_{x_r>0}$

See Problem 2, PC 3: proportionally fair shares at “macroscopic” (transmission) level result from simple processor sharing at “microscopic” (data packet) level

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates
  $q_{x,x+e_r} = \nu_r, \; q_{x,x-e_r} = \mu_r \Lambda_r$

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates $q_{x,x+e_r} = \nu_r, \; q_{x,x-e_r} = \mu_r \Lambda_r$
- reversible process for measure $\pi(x) := \binom{y+x_0}{y} \prod_{r \geq 0} \rho_r^{x_r}$

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates
  $q_{x,x+e_r} = \nu_r, \; q_{x,x-e_r} = \mu_r \Lambda_r$
- reversible process for measure $\pi(x) := \binom{y+x_0}{y} \prod_{r \geq 0} \rho_r^{x_r}$
  Negative binomial formula: $\sum_{x_0 \geq 0} \binom{y+x_0}{y} \rho_0^{x_0} = (1-\rho_0)^{-y-1}$

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates
  $q_{x,x+e_r} = \nu_r,\ q_{x,x-e_r} = \mu_r \Lambda_r$
- reversible process for measure $\pi(x) := \binom{y+x_0}{y} \prod_{r\geq 0} \rho_r^{x_r}$
  Negative binomial formula: $\sum_{x_0\geq 0} \binom{y+x_0}{y} \rho_0^{x_0} = (1-\rho_0)^{-y-1}$

  $\Rightarrow \pi$ summable (hence process ergodic) if and only if $\rho_0 + \rho_r < 1, r = 1, \ldots, L$ yielding stationary distribution
  $\pi(x) = \binom{y+x_0}{x_0} (1-\rho_0)^{-L+1} \prod_{r=1}^{L} (1-\rho_0-\rho_r) \prod_{r=0}^{L} \rho_r^{x_r}$

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates
  $q_{x,x+e_r} = \nu_r, \; q_{x,x-e_r} = \mu_r \Lambda_r$
- reversible process for measure $\pi(x) := \binom{y+x_0}{y} \prod_{r\geq 0} \rho_r^{x_r}$
  Negative binomial formula: $\sum_{x_0\geq 0} \binom{y+x_0}{y} \rho_0^{x_0} = (1-\rho_0)^{-y-1}$

  $\Rightarrow \pi$ summable (hence process ergodic) if and only if $\rho_0 + \rho_r < 1, r = 1, \ldots, L$ yielding stationary distribution
  $\pi(x) = \binom{y+x_0}{x_0} (1-\rho_0)^{-L+1} \prod_{r=1}^{L} (1-\rho_0-\rho_r) \prod_{r=0}^{L} \rho_r^{x_r}$
- Generating function (z-transform):
  $\mathbb{E} \prod_{r=0}^{L} z_r^{X_r} = \frac{(1-\rho_0 z_0)^{L-1}}{(1-\rho_0)^{L-1}} \prod_{r=1}^{L} \frac{1-\rho_0-\rho_r}{1-\rho_0 z_0-\rho_r z_r}$

Exactly solvable bandwidth sharing network

- Markov jump process with non-zero rates
  $q_{x,x+e_r} = \nu_r,\ q_{x,x-e_r} = \mu_r \Lambda_r$
- reversible process for measure $\pi(x) := \binom{y+x_0}{y} \prod_{r\geq 0} \rho_r^{x_r}$
  Negative binomial formula: $\sum_{x_0\geq 0} \binom{y+x_0}{y} \rho_0^{x_0} = (1-\rho_0)^{-y-1}$

  $\Rightarrow \pi$ summable (hence process ergodic) if and only if
  $\rho_0 + \rho_r < 1, r = 1, \ldots, L$ yielding stationary distribution
  $\pi(x) = \binom{y+x_0}{x_0} (1-\rho_0)^{-L+1} \prod_{r=1}^{L} (1-\rho_0-\rho_r) \prod_{r=0}^{L} \rho_r^{x_r}$
- Generating function (z-transform):
  $\mathbb{E} \prod_{r=0}^{L} z_r^{X_r} = \frac{(1-\rho_0 z_0)^{L-1}}{(1-\rho_0)^{L-1}} \prod_{r=1}^{L} \frac{1-\rho_0-\rho_r}{1-\rho_0 z_0 - \rho_r z_r}$
- Yields explicit formulas for per class generating functions, e.g. X_r Geometric $(\rho_r/(1-\rho_0))$ for $r \geq 1$, and

$$\begin{aligned} \mathbb{E}(X_r) &= \frac{\rho_r}{1-\rho_0-\rho_r},\ r \geq 1, \\ \mathbb{E}(X_0) &= \frac{\rho_0}{1-\rho_0} \left[1 + \sum_{r=1}^{L} \frac{\rho_r}{1-\rho_0-\rho_r} \right] \end{aligned}$$

Susceptible-Infective epidemic propagation

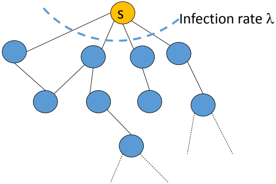

- Graph $G = (V, E)$ with n nodes ($V = [n]$)
- Infected node makes infection attempts at instants of Poisson λ process, towards graph neighbor chosen uniformly at random
- Keeps attempting forever

Susceptible-Infective epidemic propagation

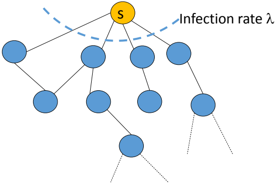

- Graph $G = (V, E)$ with n nodes ($V = [n]$)
- Infected node makes infection attempts at instants of Poisson λ process, towards graph neighbor chosen uniformly at random
- Keeps attempting forever

$\Rightarrow$ Average time to total infection? Fluctuations around average? Impact of graph topology?

Susceptible-Infective epidemic propagation

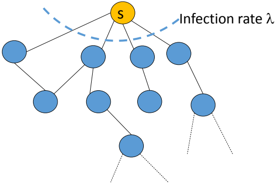

- Graph $G = (V, E)$ with n nodes ($V = [n]$)
- Infected node makes infection attempts at instants of Poisson λ process, towards graph neighbor chosen uniformly at random
- Keeps attempting forever

  $\Rightarrow$ Average time to total infection? Fluctuations around average? Impact of graph topology?
- Variant: each node = origin of its own specific epidemics; each propagation: forwards all epidemics currently held

Susceptible-Infective epidemic propagation

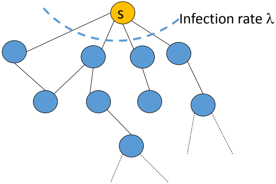

- Graph $G = (V, E)$ with n nodes ($V = [n]$)
- Infected node makes infection attempts at instants of Poisson λ process, towards graph neighbor chosen uniformly at random
- Keeps attempting forever

  $\Rightarrow$ Average time to total infection? Fluctuations around average? Impact of graph topology?
- Variant: each node = origin of its own specific epidemics; each propagation: forwards all epidemics currently held $\Rightarrow$Time till everyone heard from everyone else (“all-to-all” broadcast)?

Susceptible-Infective epidemic propagation

- Running assumption: **complete graph**

Susceptible-Infective epidemic propagation

- Running assumption: **complete graph**
- System description: X_t= number of infected nodes at time t

Susceptible-Infective epidemic propagation

- Running assumption: **complete graph**
- System description: X_t= number of infected nodes at time t
- Rate at which new attempts made when in state x:
  superposition of x Poisson λ processes
  Success probability of infection attempt: $(n-x)/(n-1)$

Susceptible-Infective epidemic propagation

- Running assumption: **complete graph**
- System description: $X_t =$ number of infected nodes at time t
- Rate at which new attempts made when in state x:
  superposition of x Poisson λ processes
  Success probability of infection attempt: $(n-x)/(n-1)$

  $\Rightarrow$ next infection time: first time of Poisson $x\lambda$ process, thinned with probability $(n-x)/(n-1)$ of retaining points: Exponential $\lambda x(n-x)/(n-1)$ random variable

Susceptible-Infective epidemic propagation

- Running assumption: **complete graph**
- System description: X_t= number of infected nodes at time t
- Rate at which new attempts made when in state x: superposition of x Poisson λ processes
  Success probability of infection attempt: $(n-x)/(n-1)$

  $\Rightarrow$ next infection time: first time of Poisson $x\lambda$ process, thinned with probability $(n-x)/(n-1)$ of retaining points: Exponential $\lambda x(n-x)/(n-1)$ random variable

  $\Rightarrow$ X_t a Markov jump process with non-zero jump rate $q_{x,x+1} = \lambda x(n-x)/(n-1)$

Time to total infection

Let E_x: i.i.d. Exponential(1) random variables, T_n: time to total outbreak
Then $T_n = \sum_{x=1}^{n-1} \frac{1}{q_x} E_x$, with $q_x = \lambda x(n-x)/(n-1)$

Time to total infection

Let E_x: i.i.d. Exponential(1) random variables, T_n: time to total outbreak

Then $T_n = \sum_{x=1}^{n-1} \frac{1}{q_x} E_x$, with $q_x = \lambda x(n-x)/(n-1)$

$$\begin{aligned} \mathbb{E}(T_n) &= \sum_{x=1}^{n-1} \frac{1}{q_x} = \frac{n-1}{n}\frac{1}{\lambda}\sum_{x=1}^{n-1}\left(\frac{1}{x}+\frac{1}{n-x}\right) \\ &= \frac{n-1}{n}\frac{2}{\lambda}H(n-1) \\ &= \frac{2}{\lambda}[\ln(n)+\gamma+o(1)] \end{aligned}$$

where $H(k)$: k-th Harmonic number, and $\gamma \approx 0.577$: Euler's constant

Time to total infection

Let E_x: i.i.d. Exponential(1) random variables, T_n: time to total outbreak

Then $T_n = \sum_{x=1}^{n-1} \frac{1}{q_x} E_x$, with $q_x = \lambda x(n-x)/(n-1)$

$$\begin{aligned}\mathbb{E}(T_n) &= \sum_{x=1}^{n-1} \frac{1}{q_x} = \frac{n-1}{n}\frac{1}{\lambda}\sum_{x=1}^{n-1}\left(\frac{1}{x}+\frac{1}{n-x}\right)\\ &= \frac{n-1}{n}\frac{2}{\lambda}H(n-1)\\ &= \frac{2}{\lambda}[\ln(n)+\gamma+o(1)]\end{aligned}$$

where $H(k)$: k-th Harmonic number, and $\gamma \approx 0.577$: Euler's constant

Similarly, for $0 < a < b < 1$: $\mathbb{E}(T_{bn} - T_{an}) \to \frac{1}{\lambda}\ln\left(\frac{b}{1-b}\frac{1-a}{a}\right)$

Time to total infection

Let E_x: i.i.d. Exponential(1) random variables, T_n: time to total outbreak

Then $T_n = \sum_{x=1}^{n-1} \frac{1}{q_x} E_x$, with $q_x = \lambda x(n-x)/(n-1)$

$$\begin{aligned}\mathbb{E}(T_n) &= \sum_{x=1}^{n-1} \frac{1}{q_x} = \frac{n-1}{n}\frac{1}{\lambda}\sum_{x=1}^{n-1}\left(\frac{1}{x}+\frac{1}{n-x}\right)\\ &= \frac{n-1}{n}\frac{2}{\lambda}H(n-1)\\ &= \frac{2}{\lambda}[\ln(n)+\gamma+o(1)]\end{aligned}$$

where $H(k)$: k-th Harmonic number, and $\gamma \approx 0.577$: Euler's constant

Similarly, for $0 < a < b < 1$: $\mathbb{E}(T_{bn} - T_{an}) \to \frac{1}{\lambda}\ln\left(\frac{b}{1-b}\frac{1-a}{a}\right)$

Heuristic inversion: starting from $X_0 = an$, $X_t \approx n\frac{ae^{\lambda t}}{1-a+ae^{\lambda t}}$

$\Rightarrow$ The celebrated **logistic function**, or S-curve

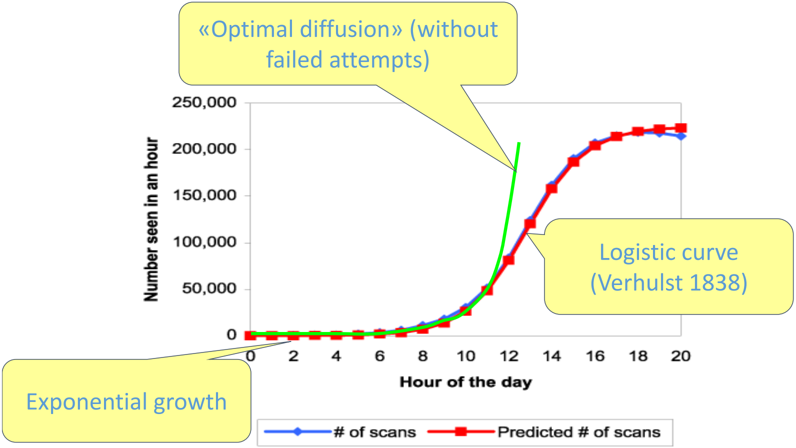

Time to total infection order-optimal
(logarithmic in number of targets) despite random target selection

Controlling fluctuations

- **Markov's inequality**: random variable $X \geq 0$, $a > 0 \Rightarrow \mathbb{P}(X \geq a) \leq \mathbb{E}(X)/a$

Controlling fluctuations

- **Markov's inequality**: random variable $X \geq 0$, $a > 0 \Rightarrow \mathbb{P}(X \geq a) \leq \mathbb{E}(X)/a$
- **Bienaymé-Tchebitchev's inequality**: random variable $X \in \mathbb{R}$: $\mathbb{P}(|X - \mathbb{E}(X)| \geq a) \leq \mathrm{Var}(X)/a^2$

Controlling fluctuations

- **Markov's inequality**: random variable $X \geq 0$, $a > 0 \Rightarrow \mathbb{P}(X \geq a) \leq \mathbb{E}(X)/a$
- **Bienaymé-Tchebitchev's inequality**: random variable $X \in \mathbb{R}$: $\mathbb{P}(|X - \mathbb{E}(X)| \geq a) \leq \mathrm{Var}(X)/a^2$
- **Exponential version**: for $\theta > 0$, $\mathbb{P}(X \geq t) \leq \mathbb{E}(e^{\theta X})e^{-\theta t}$ i.e. finite exponential moments yield exponentially decaying control of tail probabilities

Controlling fluctuations

- **Markov's inequality**: random variable $X \geq 0$, $a > 0 \Rightarrow \mathbb{P}(X \geq a) \leq \mathbb{E}(X)/a$
- **Bienaymé-Tchebitchev's inequality**: random variable $X \in \mathbb{R}$: $\mathbb{P}(|X - \mathbb{E}(X)| \geq a) \leq \mathrm{Var}(X)/a^2$
- **Exponential version**: for $\theta > 0$, $\mathbb{P}(X \geq t) \leq \mathbb{E}(e^{\theta X})e^{-\theta t}$ i.e. finite exponential moments yield exponentially decaying control of tail probabilities

Variable $S_n := \lambda(T_n - \mathbb{E}(T_n))$ satisfies for all $\theta \in [0, 1/2]$
$\mathbb{E}(\exp(\theta S_n)) \leq \exp(4\pi^2\theta^2/3) =: C_\theta < +\infty$
hence $\mathbb{P}(\lambda(T_n - \mathbb{E}(T_n)) \geq t) \leq C_\theta e^{-\theta t}$ (fluctuations small compared to mean)

Controlling fluctuations

- **Markov's inequality**: random variable $X \geq 0$, $a > 0 \Rightarrow \mathbb{P}(X \geq a) \leq \mathbb{E}(X)/a$
- **Bienaymé-Tchebitchev's inequality**: random variable $X \in \mathbb{R}$: $\mathbb{P}(|X - \mathbb{E}(X)| \geq a) \leq \mathrm{Var}(X)/a^2$
- **Exponential version**: for $\theta > 0$, $\mathbb{P}(X \geq t) \leq \mathbb{E}(e^{\theta X})e^{-\theta t}$ i.e. finite exponential moments yield exponentially decaying control of tail probabilities

Variable $S_n := \lambda(T_n - \mathbb{E}(T_n))$ satisfies for all $\theta \in [0, 1/2]$
$\mathbb{E}(\exp(\theta S_n)) \leq \exp(4\pi^2\theta^2/3) =: C_\theta < +\infty$
hence $\mathbb{P}(\lambda(T_n - \mathbb{E}(T_n)) \geq t) \leq C_\theta e^{-\theta t}$ (fluctuations small compared to mean)

Proof: For $r_x = x(n-x)/(n-1) = q_x/\lambda$,
$\mathbb{E}e^{\theta S_n} = \prod_{x=1}^{n-1} \frac{r_x}{r_x - \theta} e^{-\theta/r_x}$
For $u \in (0, 1/2]$, $\frac{e^{-u}}{1-u} \leq 1 + 2u^2$, hence:
$\mathbb{E}e^{\theta S_n} \leq \prod_{x=1}^{n-1}[1 + 2(\theta/r_x)^2] \leq e^{\sum_{x=1}^{n-1} 2(\theta/r_x)^2} \leq e^{8\theta^2 \sum_{x\geq 1} x^{-2}}$

Application: All-to-all scenario (one epidemic per user)

Lemma

Let random variables $S^1, \dots S^n$ *be such that for some* $a, b > 0$:
$\forall t > 0, \forall i \in [n], \mathbb{P}(S^i \geq t) \leq ae^{-bt}$
Then $\mathbb{E}(\sup_i S^i) \leq \mathbb{E}((\sup_i S^i)^+) \leq \frac{\ln(an)+1}{b}$

Application: All-to-all scenario (one epidemic per user)

Lemma

Let random variables $S^1, \dots S^n$ *be such that for some* $a, b > 0$:
$\forall t > 0, \forall i \in [n], \mathbb{P}(S^i \geq t) \leq a e^{-bt}$
Then $\mathbb{E}(\sup_i S^i) \leq \mathbb{E}((\sup_i S^i)^+) \leq \frac{\ln(an)+1}{b}$

Proof: Write $\mathbb{E}((\sup_i S^i)^+) = \int_0^\infty \mathbb{P}(\sup_i S^i \geq t) dt$
Then upper-bound $\mathbb{P}(\sup_i S^i \geq t)$ by nae^{-bt} for $t \geq \ln(an)/b$ (**union bound**) and by 1 otherwise.

Lemma

Let random variables $S^1, \dots S^n$ *be such that for some* $a, b > 0$:
$\forall t > 0, \forall i \in [n], \mathbb{P}(S^i \geq t) \leq ae^{-bt}$
Then $\mathbb{E}(\sup_i S^i) \leq \mathbb{E}((\sup_i S^i)^+) \leq \frac{\ln(an)+1}{b}$

Proof: Write $\mathbb{E}((\sup_i S^i)^+) = \int_0^\infty \mathbb{P}(\sup_i S^i \geq t)dt$
Then upper-bound $\mathbb{P}(\sup_i S^i \geq t)$ by nae^{-bt} for $t \geq \ln(an)/b$ (**union bound**) and by 1 otherwise.

Corollary

All-to-all propagation time T satisfies for all $\theta \in (0, 1/2]$
$\mathbb{E}T \leq \frac{1}{\lambda}\left[2(\ln(n)+\gamma)+o(1)+\frac{\ln(C_\theta n)+1}{\theta}\right] = O(\ln(n))$,
same order still

Application: All-to-all scenario (one epidemic per user)

Lemma

Let random variables $S^1, \ldots S^n$ *be such that for some* $a, b > 0$:
$\forall t > 0, \forall i \in [n], \mathbb{P}(S^i \geq t) \leq ae^{-bt}$
Then $\mathbb{E}(\sup_i S^i) \leq \mathbb{E}((\sup_i S^i)^+) \leq \frac{\ln(an)+1}{b}$

Proof: Write $\mathbb{E}((\sup_i S^i)^+) = \int_0^\infty \mathbb{P}(\sup_i S^i \geq t)dt$
Then upper-bound $\mathbb{P}(\sup_i S^i \geq t)$ by nae^{-bt} for $t \geq \ln(an)/b$ (**union bound**) and by 1 otherwise.

Corollary

All-to-all propagation time T satisfies for all $\theta \in (0, 1/2]$
$\mathbb{E}T \leq \frac{1}{\lambda}\left[2(\ln(n)+\gamma)+o(1)+\frac{\ln(C_\theta n)+1}{\theta}\right] = O(\ln(n))$,
same order still

Indeed: T = supremum of n propagation times corresponding each to single epidemic propagation

Towards Susceptible-Infective-Removed (SIR) epidemics: Galton-Watson branching process (1873)

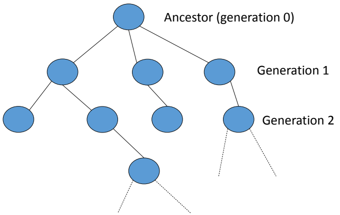

Offspring distribution $\{p_k\}_{k\in\mathbb{N}}$

Z_k number of individuals per generation:

$Z_0 = 1, Z_k = \sum_{m=1}^{Z_{k-1}} X_{m,k}$ where $\{X_{m,k}\}_{m,k\geq 0}$: i.i.d., $\sim \{p_k\}_{k\in\mathbb{N}}$

Quantities of interest: probability of extinction; in case of extinction, total population size

Theorem

Extinction probability p_{ext}*: smallest root in* $[0, 1]$ *of* $z = \phi(z)$
where $\phi(z) = \mathbb{E}(z^X) = \sum_{k \geq 0} p_k z^k$
If $\mu := \mathbb{E}(X) < 1$ *then* $p_{ext} = 1$
If $\mu = 1$ *and* $p_0 > 0$ *then* $p_{ext} = 1$
If $\mu > 1$ *then* $p_{ext} < 1$

Theorem

Extinction probability p_{ext}*: smallest root in* $[0,1]$ *of* $z = \phi(z)$
where $\phi(z) = \mathbb{E}(z^X) = \sum_{k\geq 0} p_k z^k$
If $\mu := \mathbb{E}(X) < 1$ *then* $p_{ext} = 1$
If $\mu = 1$ *and* $p_0 > 0$ *then* $p_{ext} = 1$
If $\mu > 1$ *then* $p_{ext} < 1$

Proof: $\{Z_k = 0\} \nearrow \{\text{Extinction}\}$; $\mathbb{P}(Z_k = 0) = \phi_k(0)$ where $\phi_k(z) = \mathbb{E}(z^{Z_k})$
By induction $\phi_k(z) = \phi \circ \phi_{k-1}(z)$ hence
$\mathbb{P}(Z_k = 0) = \phi(\mathbb{P}(Z_{k-1} = 0))$
$\Rightarrow$ by monotonicity of ϕ and $\mathbb{P}(Z_0 = 0) = 0$, sequence increases to (necessarily smallest) fixed point.

Theorem

Extinction probability p_{ext}*: smallest root in* $[0,1]$ *of* $z=\phi(z)$
where $\phi(z)=\mathbb{E}(z^X)=\sum_{k\geq 0} p_k z^k$
If $\mu:=\mathbb{E}(X)<1$ *then* $p_{ext}=1$
If $\mu=1$ *and* $p_0>0$ *then* $p_{ext}=1$
If $\mu>1$ *then* $p_{ext}<1$

Proof: $\{Z_k=0\}\nearrow\{\text{Extinction}\}$; $\mathbb{P}(Z_k=0)=\phi_k(0)$ where $\phi_k(z)=\mathbb{E}(z^{Z_k})$

By induction $\phi_k(z)=\phi\circ\phi_{k-1}(z)$ hence
$\mathbb{P}(Z_k=0)=\phi(\mathbb{P}(Z_{k-1}=0))$

$\Rightarrow$ by monotonicity of ϕ and $\mathbb{P}(Z_0=0)=0$, sequence increases to (necessarily smallest) fixed point.

μ: slope of ϕ at 1^-. By convexity of ϕ, only fixed point: 1 if $\mu<1$

By continuity of ϕ, $\exists$ fixed point <1 if $\mu>1$

For $\mu=1$, if $p_0>0$ then ϕ strictly convex hence only fixed point: 1; if $p_0=0$ then $p_{ext}=0$

Theorem

Extinction probability p_{ext}*: smallest root in* $[0,1]$ *of* $z = \phi(z)$
where $\phi(z) = \mathbb{E}(z^X) = \sum_{k\geq 0} p_k z^k$
If $\mu := \mathbb{E}(X) < 1$ *then* $p_{ext} = 1$
If $\mu = 1$ *and* $p_0 > 0$ *then* $p_{ext} = 1$
If $\mu > 1$ *then* $p_{ext} < 1$

Proof: $\{Z_k = 0\} \nearrow \{\text{Extinction}\}$; $\mathbb{P}(Z_k = 0) = \phi_k(0)$ where $\phi_k(z) = \mathbb{E}(z^{Z_k})$
By induction $\phi_k(z) = \phi \circ \phi_{k-1}(z)$ hence
$\mathbb{P}(Z_k = 0) = \phi(\mathbb{P}(Z_{k-1} = 0))$
$\Rightarrow$ by monotonicity of ϕ and $\mathbb{P}(Z_0 = 0) = 0$, sequence increases to (necessarily smallest) fixed point.
μ: slope of ϕ at 1^-. By convexity of ϕ, only fixed point: 1 if $\mu < 1$
By continuity of ϕ, $\exists$ fixed point < 1 if $\mu > 1$
For $\mu = 1$, if $p_0 > 0$ then ϕ strictly convex hence only fixed point: 1; if $p_0 = 0$ then $p_{ext} = 0$
Fundamental example of **phase transition**
Special case $X \sim \text{Poisson}(\mu)$: $p_{ext} = e^{-\mu(1-p_{ext})}$

Random walk exploration of Galton-Watson tree

Sequentially pick *active* node (whose children have not yet been sampled)
De-activate it and add its children to active set
Stop when active set empty (tree exploration complete)

Random walk exploration of Galton-Watson tree

Sequentially pick *active* node (whose children have not yet been sampled)
De-activate it and add its children to active set
Stop when active set empty (tree exploration complete)

- Dynamics of A_t, number of active nodes at step t:
  Random walk $A_t = A_{t-1} - 1 + X_t$ where X_t independent of past exploration $\{A_s, X_s, s < t\}$ and distributed according to $\{p_k\}_{k\geq 0}$
- Time T at which exploration stops, i.e. $A_T = 0$ gives size of tree. Indeed $A_t = 1 - t + X_1 + \ldots + X_t$ and $A_T = 0$ yield $T = 1 + X_1 + \ldots + X_T$.
- Random walk can be pursued after time T

Random walk exploration of Galton-Watson tree

Sequentially pick *active* node (whose children have not yet been sampled)
De-activate it and add its children to active set
Stop when active set empty (tree exploration complete)

- Dynamics of A_t, number of active nodes at step t:
  Random walk $A_t = A_{t-1} - 1 + X_t$ where X_t independent of past exploration $\{A_s, X_s, s < t\}$ and distributed according to $\{p_k\}_{k \geq 0}$
- Time T at which exploration stops, i.e. $A_T = 0$ gives size of tree. Indeed $A_t = 1 - t + X_1 + \ldots + X_t$ and $A_T = 0$ yield $T = 1 + X_1 + \ldots + X_T$.
- Random walk can be pursued after time T

$\Rightarrow$ Bound on population size: for continued RW $\{A_t\}_{t \geq 0}$,
$\mathbb{P}(T > t) = \mathbb{P}(A_0, \ldots, A_t > 0) \leq \mathbb{P}(A_t > 0) = \mathbb{P}(\sum_{s=1}^{t}(X_s - 1) \geq 0)$

Chernoff's inequality and bounds on population size

Theorem

For i.i.d. X_s, $\mathbb{P}(\sum_{s=1}^{t} X_s \geq at) \leq e^{-th(a)}$ *where*
$h(a) := \sup_{\theta>0}[\theta a - \ln(\mathbb{E}(e^{\theta X_1}))]$

Chernoff's inequality and bounds on population size

Theorem

For i.i.d. X_s, $\mathbb{P}(\sum_{s=1}^{t} X_s \geq at) \leq e^{-th(a)}$ *where*
$h(a) := \sup_{\theta>0}[\theta a - \ln(\mathbb{E}(e^{\theta X_1}))]$

Non-trivial exponential bound when $a > \mathbb{E}(X_1)$ and
$\exists \epsilon > 0 : \mathbb{E} e^{\epsilon X_1} < +\infty$

Chernoff's inequality and bounds on population size

Theorem

For i.i.d. X_s, $\mathbb{P}(\sum_{s=1}^{t} X_s \geq at) \leq e^{-th(a)}$ *where* $h(a) := \sup_{\theta>0}[\theta a - \ln(\mathbb{E}(e^{\theta X_1}))]$

Non-trivial exponential bound when $a > \mathbb{E}(X_1)$ and $\exists \epsilon > 0 : \mathbb{E}e^{\epsilon X_1} < +\infty$

Application to Galton-Watson process:
$\mathbb{P}(T > t) \leq e^{-th(1)}$ exponentially decaying if $\mathbb{E}(X_1) < 1$ and X_1 admits finite exponential moments.

Theorem

For i.i.d. X_s, $\mathbb{P}(\sum_{s=1}^{t} X_s \geq at) \leq e^{-th(a)}$ *where* $h(a) := \sup_{\theta>0}[\theta a - \ln(\mathbb{E}(e^{\theta X_1}))]$

Non-trivial exponential bound when $a > \mathbb{E}(X_1)$ and $\exists \epsilon > 0 : \mathbb{E}e^{\epsilon X_1} < +\infty$

Application to Galton-Watson process:
$\mathbb{P}(T > t) \leq e^{-th(1)}$ exponentially decaying if $\mathbb{E}(X_1) < 1$ and X_1 admits finite exponential moments.

Case of Poisson random variables, parameter $\mu > 0$, $a > \mu$:
$h_\mu(a) = \sup_{\theta>0}[\theta a - \mu(e^\theta - 1)]$
Gives $\theta = \ln(a/\mu)$, $h_\mu(a) = \mu h_1(a/\mu)$
with $h_1(x) = x\ln(x) - x + 1$

Takeaway messages

- Epidemic spread in logarithmic time for single propagation and for all-to-all propagation
- Same order as if infection attempts were optimized
- Motivates "epidemic algorithms" for information dissemination
- Exponential versions of Markov's inequality (in particular Chernoff's inequality): powerful tool, will be used in analysis of SIR epidemics and random graphs