ON RECURRENCE AND TRANSIENCE IN HEAVY-TAILED GSMP'S

Peter J. Haas[1] Peter W. Glynn[2]

[1]IBM Almaden Research Center
San Jose, CA

[2]Dept. of Management Science and Engineering
Stanford University

ETH April 2009

Discrete-Event Stochastic Systems

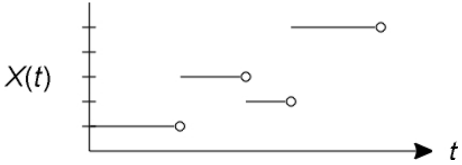

- System changes state when events occur
  - Stochastic state changes
  - At strictly increasing sequence of random times
- Underlying stochastic process $\{ X(t) \colon t \geq 0 \}$
  - $X(t)$ = state of system at time t (a random variable)
  - Piecewise-constant sample paths
  - Typically not a continuous-time Markov chain

Generalized Semi-Markov Processes [Matthes62, Whitt80]

- Classical model for discrete-event stochastic systems
  - Subsumes: queueing networks, SMPs, CTMCs, SPNs
  - Central to simulation theory
- Building blocks
  - S = set of states (finite or countably infinite)
  - E = set of events (finite)
  - $E(s)$ = active events in state s
  - $p(s'; s, e^*)$ = state-transition probability
- One clock per event: records remaining time until occurrence

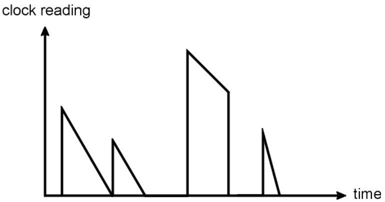

Clocks (Event Scheduling)

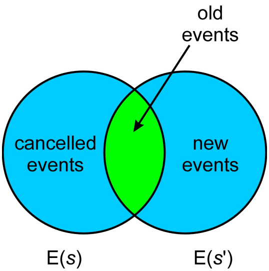

- Active events compete to trigger state transition
  - The clock that runs down to 0 first is the "winner"
  - Can have simultaneous event occurrence: $p(s'; s, E^*)$
- At a state transition $s \xrightarrow{e^*} s'$: three kinds of events
  - New events: Clock for e' set according to $F(x; e')$
  - Old events: Clocks continue to run down
  - Cancelled events: Clock readings are discarded

The GSMP Process

- The continuous-time process: $\{ X(t) : t \geq 0 \}$
  - $X(t) =$ the state at time t
  - A very complicated process
- Defined in terms of Markov chain $\{ (S_n, C_n) : n \geq 0 \}$
  - System observed after the nth state transition
  - $S_n =$ the state
  - $C_n = (C_{n,1}, \ldots, C_{n,M}) =$ the clock-reading vector
  - Chain defined via GSMP building blocks
  - initial distribution μ on state and clocks

Definition of the GSMP

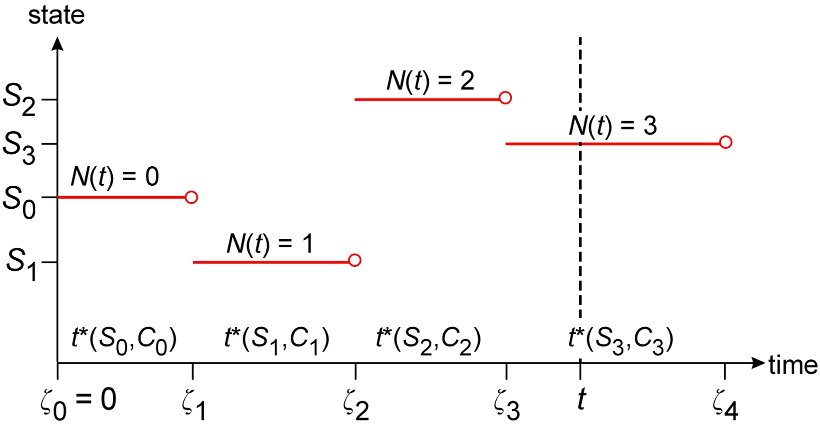

$$X(t) = S_{N(t)}$$

Example: GI/G/1 Queue

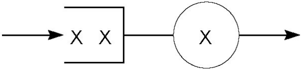

- $X(t) =$ "number of jobs in system at time t"
- $S = \{ 0, 1, 2, \ldots \}$
- $E = \{ e_1, e_2 \}$
  - $e_1 =$ "arrival"
  - $e_2 =$ "service completion"
- $E(0) = \{ e_1 \}$ and $E(s) = \{ e_1, e_2 \}$ for $s > 0$
- $p(s+1; s, e_1) = 1$ and $p(s-1; s, e_2) = 1$
- $F(\,\cdot\,; e_1) = F_{\text{interarrival}}$ and $F(\,\cdot\,; e_2) = F_{\text{service}}$

Harris Recurrence: A Basic Form of Stability

- Definition for general chain $\{ Z_n : n \geq 0 \}$ with state space Γ

$$P_z \{ Z_n \in A \text{ i.o.} \} = 1,\ z \in \Gamma \quad \text{whenever} \quad \phi(A) > 0$$

  - ϕ is a recurrence measure (often "Lebesgue-like")
  - Every "dense enough" set is hit infinitely often w.p. 1
  - No "wandering off to ∞"
  - Chain admits invariant measure π_0: $\int P(z, A)\, \pi_0(dz) = \pi_0(A)$
- Positive Harris recurrence:
  - Chain admits invariant probability measure π
  - $P_\pi \{ Z_1 \in A \} = \pi(A)$
  - Implies stationarity when initial dist'n is π
- When is $\{ (S_n, C_n) : n \geq 0 \}$ (positive) Harris recurrent?

Some Stability Conditions

- Density component g of a cdf F: $F(t) \geq \int_0^t g(u)\,du$
- $s \to s'$ iff $p(s'; s, e) > 0$ for some e
- $s \rightsquigarrow s'$: either $s \to s'$ or $s \to s^{(1)} \to \cdots \to s^{(n)} \to s'$
- Assumption PD(q):
  - State space S is finite
  - GSMP is irreducible: $s \rightsquigarrow s'$ for all $s, s' \in S$
  - There exists $\bar{x} \in (0, \infty)$ s.t. all clock-setting dist'n functions
    - Have finite qth moment
    - Have density component positive on $[0, \bar{x}]$

Positive Harris Recurrence in Light-Tailed GSMPs [Haas99]

- $\bar{\phi}(\{s\} \times A) =$ Lebesgue measure of $A \cap [0, \bar{x}]^M$
- Theorem: If Assumption PD(1) holds, then the (S_n, C_n) chain is positive Harris recurrent with recurrence measure $\bar{\phi}$
- Implies $P_\mu\{\, S_n = s \text{ i.o.}\,\} = 1$ for all $s \in S$ with finite expected (continuous-time) hitting times
- Proof:
  - Show that chain is "$\bar{\phi}$-irreducible"
  - Establish Lyapunov drift condition and apply MC machinery (Meyn and Tweedie, 1993)
- Alternate approach to recurrence: geometric-trials arguments
  - Can drop positive-density assumption
  - Use detailed analysis of specific GSMP structure

Positive Recurrence in Heavy-Tailed GSMPs: It Depends

- Example 1: Uninfluential events
  - For all $s \in S$: $e \in E(s)$ and $p(s; s, e) = 1$
  - No effect on state or other clocks
  - If all heavy-tailed events are uninfluential and PD(1) holds otherwise, then positive recurrence
- Example 2: e_1 is heavy-tailed, so no positive-recurrence

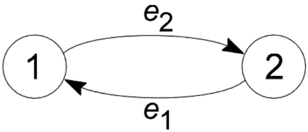

Recurrence in Heavy-Tailed GSMPs

- S_n = state just after nth state transition
- Conjecture: $P\{\, S_n = s \text{ i.o.} \,\} = 1$ for each s under PD(0)
  - State space S is finite
  - GSMP is irreducible
  - $\exists\, \bar{x} > 0$ s.t. each $F(\,\cdot\,; e)$ has positive density on $(0, \bar{x})$
- Certainly true for CTMCs
- CONJECTURE IS FALSE for GSMPs!
  - In the presence of heavy-tailed clock-setting dist'ns

A Counterexample

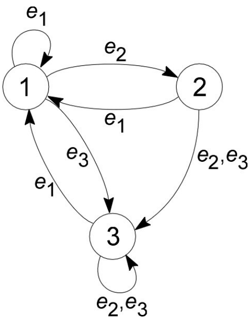

- $S = \{1, 2, 3\}$ and $E = \{e_1, e_2, e_3\}$
- Event sets: $E(s) = \{e_1, e_2, e_3\}$ for all s (renewal processes)
- $p(s'; s, e^*) = 0$ or 1
- Clock-setting distributions:
  - $F(t; e_1) = 1 - (1 + t)^{-\alpha}$
  - $F(t; e_2) = 1 - (1 + t)^{-\beta}$
  - $F(\cdot\,; e_3)$ is Uniform$[0, a]$

  with $\beta > 1/2$ and $\alpha + \beta < 1$
- GSMP hits state $s = 2$ only if:
  - e_1 occurs and then e_2 occurs with no intervening occurrence of e_3
- Claim: $P\{S_n = 2 \text{ i.o.}\} = 0$
- Intuition: heavy clocks are rarely small simultaneously

Proof

- Observe: $P\{ S_n = 2 \text{ i.o.} \} = 0$ iff $P\{ B_n \text{ i.o.} \} = 0$
  - $B_n = \{ C_2(T_n) \le C_3(T_n) \}$
  - $T_n = n$th occurrence time of e_1
  - $C_i(t) =$ clock reading for e_i at time t
- Borel-Cantelli: suffices to show that $\sum_{n=1}^{\infty} P\{ B_n \} < \infty$
- Bound $P\{ B_n \}$ by an integral:

$$P\{ B_n \} \le \int_0^{\infty} P\{ C_2(t) \le a \} f_1^{*n}(t)\, dt$$

  - $f_1^{*n} =$ density of $T_n = n$-fold convolution of $f(\,\cdot\,; e_1)$
- Sum over n:

$$\sum_{n=1}^{\infty} P\{ B_n \} \le \int_0^{\infty} h(t) u_1(t)\, dt$$

  - $h(t) = P\{ C_2(t) \le a \}$
  - $u_1 =$ renewal density function for $F(\,\cdot\,; e_1)$
- So suffices to show that $\int_0^{\infty} h(t) u_1(t)\, dt < \infty$

Proof–Continued

- To show: $\int_0^\infty h(t)u_1(t)\,dt < \infty$ where $h(t) = P\{\,C_2(t) \le a\,\}$
- Renewal argument: $h = U_2 * Q$
  - $U_2 =$ renewal function for $F(\,\cdot\,; e_2)$
  - $Q(t) = F(t+a; e_2) - F(t; e_2)$
- Heavy-tail key renewal theorem [Erickson70]:
  - Light-tail KRT: $h(t) \to m_2(a)/m_2(\infty)$
    - where $m_2(t) = \int_0^t \bar{F}(u; e_2)\,du = (1+t)^{1-\beta}/(1-\beta)$
    - $\bar{F} = 1 - F$
  - Regular variation: $G \in \mathrm{RV}_\lambda$ iff $\lim_{t\to\infty} G(tx)/G(t) = x^\lambda$
  - Erickson: $h(t) = O\big(1/m_2(t)\big)$
  - Thus $h(t) = O(t^{\beta-1})$

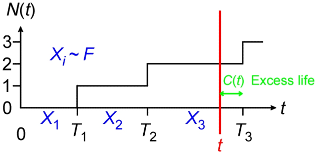

$U(t) = E[N(t)] = \sum_{n=1}^{\infty} F^{*n}(t)$

$u(t) = dU(t)/dt = \sum_{n=1}^{\infty} f^{*n}(t)$

Proof–Continued

- To show: $\int_0^\infty h(t)u_1(t)\,dt < \infty$
  where $h(t) = O(t^{\beta-1})$
- Tauberian theorems: $\bar{F}(\,\cdot\,; e_1) \in \mathrm{RV}_{-\alpha} \Rightarrow U_1 \in \mathrm{RV}_\alpha$
  - $U_1 =$ renewal function for $F(\,\cdot\,; e_1)$
- Landau's theorem: u_1 ultimately monotone $\Rightarrow u_1 \in \mathrm{RV}_{\alpha-1}$
  - Temporarily assume ultimate monotonicity
  - Result in Feller $\Rightarrow u_1(t) = O(t^{\alpha-1+\epsilon})$
- Combine: $\int_0^\infty h(t)u_1(t)\,dt = \int_0^\infty O(t^{\alpha+\beta+\epsilon-2})\,dt < \infty$
  since $\alpha + \beta + \epsilon < 1$

Ultimate-Monotonicity Proof (Sketch)

- $\mathcal{L}(f_1)(s) = \alpha s^\alpha e^s \Gamma(-\alpha, s)$, where

$$\Gamma(a, s) = \int_s^\infty e^{-t} t^{a-1}\, dt$$

- $\mathcal{L}(u_1)(s) = \dfrac{\mathcal{L}(f_1)(s)}{1 - \mathcal{L}(f_1)(s)}$

- $\mathcal{L}(u_1')(s) = \dfrac{s\mathcal{L}(u_1)(s)}{1 - \mathcal{L}(u_1)(s)} - u_1(0+) = g(s)$, where

$$g(s) = \frac{\alpha s^{\alpha+1} e^s \Gamma(-\alpha, s)}{1 - \alpha s^\alpha e^s \Gamma(-\alpha, s)} - \alpha$$

- $g(s)$ is analytic except at origin

Ultimate-Monotonicity Proof — Continued

- Inversion formula: $u_1'(t) = \dfrac{1}{2\pi i}\displaystyle\int_{1-i\infty}^{1+i\infty} e^{st} g(s)\, ds$
- Apply Cauchy's integral theorem:

$$u_1'(t) \approx -\frac{ct^{\alpha-2}\sin\bigl(\pi(1-\alpha)\bigr)}{\pi} < 0$$

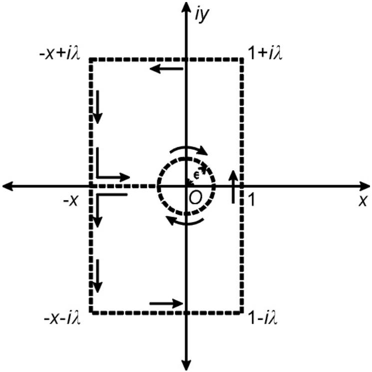

Summary and Conjecture

# HT Clocks	RW Equiv.	Recurrent?	"Positive Recurrent"?
0	–	Yes	Yes
1	2	Yes	Depends
≥ 2	≥ 4	Depends	Depends

A Special Case

- Hazard rate: $h(x;e) = f(x;e)/\bar{F}(x;e)$
- Theorem: $P\{\, S_n = s \text{ i.o.} \,\} = 1$ for each s if
  - State space S is finite
  - GSMP is irreducible
  - $\exists\, \bar{x} > 0$ s.t. each $F(\,\cdot\,;e)$ has positive density on $(0,\bar{x})$
  - At most one active event with heavy-tailed clock-setting dist'n
  - $\alpha(e) \leq h(x;e) \leq \beta(e)$ for each light-tailed event e
- Proof uses regenerative structure [Glynn89] + geometric trials

Questions?

www.almaden.ibm.com/cs/people/peterh