Autoregressive Cascades on Random Networks

Srikanth K Iyer

(Joint work with Rahul Vaze)

DEPARTMENT OF MATHEMATICS
INDIAN INSTITUTE OF SCIENCE, BANGALORE

September 11, 2014

Model Description

- Consider a (possibly) infinite random tree where each node has out degree which is an independent copy of N with $E[N] > 1$ and

$$w_k = P[N = k], \qquad k = 0, 1, \ldots.$$

- At time $t = 0$, each node has a load that is an independent copy of $0 < L < c$ with d.f. F.
- The capacity of each node is a fixed constant $c > 0$.
- At time $t = 0$, the load at the root r increases to $\ell > c$. This causes r to fail.
- If the out-degree $N(r) = 0$, then process terminates. Else given $N(r) = k$, let $u_1, \ldots, u_k$ be the nodes connected to r. Let $\{p_{ru_i}, i = 1, \ldots, k\}$ be a random exchangeable pmf.
- The load at time $t = 1$ at u_i is given by

$$L_{u_i}^{(1)} = p_{ru_i}\ell + L_{u_i}^{(0)}.$$

Going Forward

- If $L_{u_i}^{(1)} < c$ for all $i = 1, \ldots, k$, then cascade terminates at time $t = 1$.
- Node u_i fails at time 1 if the resultant load $L_{u_i}^{(1)} \geq c$.
- The process now proceeds as above from each of the failed node with input $L_{\cdot}^{(1)}$.
- Once a node fails it takes no further part in the process.
- The process terminates at time $T + 1$ if some node at a distance T from the root has failed and none of the nodes at distance $T + 1$ fail at time $T + 1$.
- We will refer to this process as the Autoregressive Cascade or ARC model.

Cascades on Random Graphs

- Cascades in random networks was introduced by D.J. Watts (2002), to study the spread of ideas, opinion, technology etc.
- Agents are the nodes of a network
- Interactions between agents - links in the network.
- Agent initially in state 0 will adopt a new idea (state 1) as soon as a fraction of its neighbors who have adopted the new idea exceeds a threshold.
- Showed the existence of a phase transition (as a function of the threshold)
- A large number of generalizations of this model have been studied.

Cascades on Random Graphs

- Gleeson, Hurd, Melnik and Hackett (2013): Systemic risks in financial networks.
- Links have weights that may depend on the degree of the nodes that the link connects. Node changes state if the sum of the weights of links that have adopted the idea exceeds a threshold.
- Kempe, Kleinberg and Tardos (2003) and Venkatraman and Kumar (2011): Random thresholds.
- Models for epidemic (Newman (2010), Lelarge (2012)): Nodes fail or are infected based on some probabilistic or deterministic mechanism that depends only on the number of failed/infected neighbors but not the severity of the infection.
- ARC model, failure is governed by a mechanism of load transfer whose effect can persist over several generations.

Cascades on Random Graphs

- The ARC model is suited to study the behavior of outages in electrical power networks.
- A single node failure can lead to catastrophically wide-spread outages (Santhi (2010); simulation studies with exponential load distributions).
- Dobson, Carreras, Newman (2004): Each failing node results in the failure of k nodes.
- Dobson, Carreras, Newman (2005): Failure at a node results in a uniform increase in the loads at all the active nodes.
- Sand pile models.

Super-Critical Regime

$$p(k) \stackrel{d}{=} p_{ru_i}(k), \qquad \bar{F} = 1 - F.$$

- The ARC process survives indefinitely with positive probability if c satisfies the following condition:

$$\sum_{k=1}^{\infty} w_k \, k \, \mathbb{E}\left[\bar{F}(c(1-p(k)))\right] > 1.$$

- Idea of proof: Couple with a Galton-Watson branching process.

The Coupled GW process

- Gen 0: Start with an individual labelled r.
- Given $N(r) = k > 0$, let $\{u_i, i = 1, 2, \ldots, k\}$ be neighbors of r and let $\{p_{ru_i}, i = 1, 2, \ldots, k\}$ be the random allocation in the tree.
- A child labelled u_i is born to r in the GW process if

$$p_{ru_i} c + L_{u_i}^{(0)} \geq c.$$

- Note that node u_i fails in the ARC process if

$$L_{u_i}^{(1)} = p_{ru_i} L_r^{(0)} + L_{u_i}^{(0)} \geq c.$$

- Since $L_r^{(0)} = \ell > c$ node u_i fails in the ARC process if it is born in the GW process.
- The same procedure holds for generating offspring in subsequent generations.

Supercriticality of the Coupled GW process

- Probability a child is born to a parent with out-degree k in the graph is

$$\mathbb{E}\left[\bar{F}(c(1-p(k)))\right].$$

- So expected offspring size will be

$$\sum_{k=1}^{\infty} w_k \, k \, \mathbb{E}\left[\bar{F}(c(1-p(k)))\right]$$

- The ARC process survives if the GW does, which happens if the expected number of offspring exceeds one.

Sub-Critical Regime

$$h := \inf_{\theta \geq 0} \mathbb{E}\left[N e^{\theta(L-(1-p)c)} \right]$$

- If $h < 1$ then the cascade will terminate in finite time with probability one.
- Idea: dominate the ARC process with a branching random walk

Branching Random Walks (BRW)

- A BRW is a process that starts with a single individual labelled r located at $x \in \mathbb{R}$.
- An individual labelled u born at location at $y \in \mathbb{R}$ lives for unit time ...
- at the end of which gives birth to offspring located according to the point process $y + Z_u$, where Z_u is an independent copy of a point process Z.

$$m(\theta) = E\left[\int_{-\infty}^{\infty} e^{-\theta t} dZ(t)\right]$$

$$\eta = \inf\{m(\theta) : \theta \geq 0\}.$$

Theorem [Biggins, 1977] Let $Z^{(n)}$ be the number of individuals of the BRW in generation n located in the interval $(-\infty, 0]$. Then if $\eta < 1$, then, almost surely, $Z^{(n)} = 0$ for all but finitely many n.

The Drift in the Load

- If node u fails then its load increases to $L_u^{(1)} \geq c$.
- The load at its child node v increases to

$$L_v^{(1)} = p_{uv} L_u^{(1)} + L_v^{(0)}$$

- The difference in the loads at nodes v and u, or the "drift" of the load increase satisfies

$$L_v^{(1)} - L_u^{(1)} = -(1 - p_{uv}) L_u^{(1)} + L_v^{(0)} \leq -(1 - p_{uv}) c + L_v^{(0)}$$

Constructing the Coupled BRW

- The process X starts with a particle labelled r located at $\ell > c$.
- Suppose node u fails at time t in the ARC process.
- If the number of non-failed neighbors of u , $N(u) = 0$, then no child is born in the BRW process.
- Else given $N(u) = k$, then for each neighbor v_i of u, an individual with the same label is born in the BRW process at time $t + 1$ if it is located in (c, ∞).
- The new load at node v_i in the ARC process is

$$L_{v_i}^{(1)} = p_{uv_i}(k) L_u^{(1)} + L_{v_i}^{(0)}.$$

- The location of v_i in the BRW process will

$$x_{v_i}^{(t+1)} = x_u^{(t)} - (1 - p_{uv_i}(k))c + L_{v_i}^{(0)} \geq L_{v_i}^{(1)},$$

- If X terminates then so does the ARC.

Two more tricks

- Consider another BRW Y coupled to X as follows. The process Y starts off with a single individual located at $-\ell$
- if a child is born to a parent located at x in the BRW X such that the location of the child is $x+d$, then a child is born to a corresponding parent located at y in the process Y and the child is located at $y-d$.
- Individuals in Y produce offspring whose numbers are distributed as N with displacements distributed as $-(L-(1-p)c)$. If X drifts to the left, then Y drifts to the right.
- No barrier in the Y process.
- Shift the origin to $-c$.
- The BRW X terminates if for some n, $X_n([2c,\infty))=0$ and this happens if $Y_n((-\infty,0]=0$.

Extensions

- Can extend to graphs that are "locally tree-like" such as the Newman, Strogatz and Watts (2001) random graphs with a given degree distribution.
- Size-biased distribution

$$\tilde{w}_{k-1} = \frac{k w_k}{\mu}, \qquad k \geq 1,$$

$$\mu = \sum_{k=1}^{\infty} k w_k.$$

- Random capacities and allocation depending on load at neighboring nodes?