Queue-based activation protocols in random-access networks

Matteo Sfragara
University of Leiden

Joint work with
S. Borst (Eindhoven), F. den Hollander (Leiden), F. R. Nardi (Florence)

YEQT Toulouse
3-4 December 2018

Index

Index

Random-access network

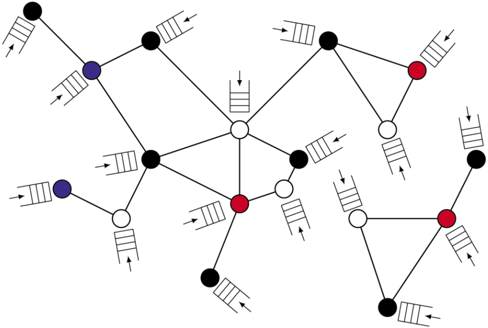

- **Servers** form a network.
- **Packets** arrive at each server, heavy load.
- Servers **interfere** with each other when "too close"
- Only active servers can process packets.

The interference graph

- Network $\rightarrow$ graph $G(N, B)$.
- Servers $\rightarrow$ set of nodes N.
- "Closeness" $\rightarrow$ set of bonds B.

Network model

- **bipartite graph** $G(U \cup V, B)$
- red nodes in U, blue nodes in V
- state of node i at time t is
  $$X_i(t) = \begin{cases} 0, & \text{inactive} \\ 1, & \text{active} \end{cases}$$
- nodes connected by bonds cannot be active at the same time

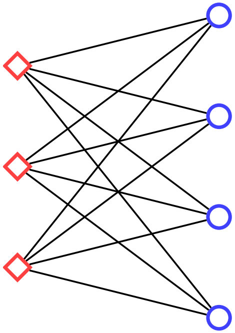

Transition time

Two **stable configurations**:

- u = all nodes in U active, all nodes in V inactive;
- v = all nodes in U inactive, all nodes in V active.

Transition time

Define the *transition time* τ_v to be the first time the system hits the configuration v, i.e.,

$$\tau_v = \min \left\{ t \geq 0 \colon \; X_i(t) = 0 \; \forall \, i \in U, \; X_i(t) = 1 \; \forall \, i \in V \right\}.$$

We are interested in the distribution of τ_v given that the initial configuration is u.

The transition from u to v represents a "global switch in the network".

Two type of models

- ON $\to$ OFF: Poisson deactivation clock ticking at rate 1.
- OFF $\to$ ON: Poisson activation clock ticking at a time-varying rate; the attempt is succesful when no neighbours are active at time t^-.

Different models

- **External**: activation rates depend on a deterministic function $f(t)$

$$r_i^{\text{ext}}(t) = \begin{cases} g_U(f(t)), & i \in U, \\ g_V(f(t)), & i \in V. \end{cases}$$

- **Internal**: activation rates depend on the actual queue length $Q_i(t)$

$$r_i^{\text{int}}(t) = \begin{cases} g_U(Q_i(t)), & i \in U, \\ g_V(Q_i(t)), & i \in V. \end{cases}$$

Acitivation rates

The choice g_U, g_V determines the transition time of the network.

Assumptions on the activation rates:

(i) let $g_U, g_V \in \mathcal{G}$, with

$$\mathcal{G} = \Big\{ g\colon \mathbb{R} \to \mathbb{R}_{\geq 0}\colon \; g \text{ non-decreasing and globally Lipschitz,} \\ g(\mathbb{R}_{\leq 0}) = 0, \lim_{x \to \infty} g(x) = \infty \Big\};$$

(ii) we want nodes in V more aggressive than nodes in U, i.e.,

$$\lim_{x \to \infty} \frac{g_V(x)}{g_U(x)} = \infty.$$

We focus on polynomial functions

$$g_U(x) = Gx^{\beta}, \qquad x \in [0, \infty),$$

with $G, \beta \in (0, \infty)$.

The queue length

Queue length

The length of the queue at node i at time t is

$$Q_i(t) = 0 \vee [Q_i(0) + Q_i^+(t) - Q_i^-(t)].$$

- Input process $Q_i^+(t) = \sum_{j=0}^{N_i(t)} Y_{ij}$: compound Poisson process with mean ρ.
- Output process $Q_i^-(t) = c \int_0^t X_i(u)du$: a server processes its packets at rate c.
- We want $c > \rho$.

Intensity parameter $r \to \infty$. Given $\gamma_U \geq \gamma_V > 0$, the initial queue lengths are assumed to be

$$Q_i(0) = \begin{cases} \gamma_U r, & i \in U, \\ \gamma_V r, & i \in V. \end{cases}$$

Index

Complete bipartite graphs

Define the **pre-transition time** $\bar{\tau}_v$ to be the first time a node in V activates.

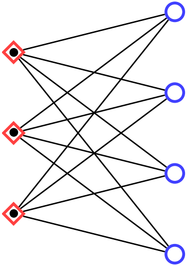

Initial configuration u.

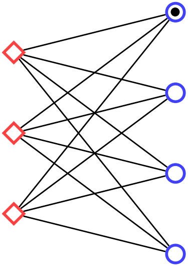

Pre-transition.

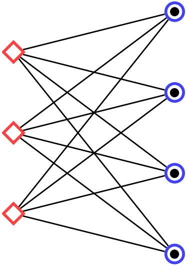

Transition, v.

Ideas

- Borst, den Hollander, Nardi, Taati (2017): results on the transition time τ_v^{ext} for external models.
- The internal model is expected to be more efficient, since it looks at the actual queue lengths.
- We **compare** the internal model with the external model via the mean queue length: activation rates with $f_i(t) = \mathbb{E}[Q_i(t)]$.
- We construct two auxiliary external models (lower and upper) by perturbing the mean queue length, hence the activation rates of the external model. We know how to deal with them.
- Results are obtained by **coupling** the three models and sandwiching τ_v^{int} between τ_v^{low} and τ_v^{upp}.

Mean transition time

Recall that $g_U(x) = Gx^{\beta}$, $x \in [0, \infty)$.

- Subcritical regime: $\beta \in \left(0, \frac{1}{|U|-1}\right)$.
- Critical regime: $\beta = \frac{1}{|U|-1}$.
- Supercritical regime: $\beta \in \left(\frac{1}{|U|-1}, \infty\right)$.

Theorem 1: Mean transition time in the external model

$$\mathbb{E}[\tau_v^{\text{ext}}] = F r^{1 \vee \beta(|U|-1)} \, [1 + o(1)], \qquad r \to \infty,$$

with

$$F = \begin{cases} \frac{\gamma_U^{\beta(|U|-1)}}{|U| G^{-(|U|-1)}}, & \text{if } \beta \in (0, \frac{1}{|U|-1}), \\ \frac{\gamma_U}{|U| G^{-(|U|-1)} + (c - \rho_U)}, & \text{if } \beta = \frac{1}{|U|-1} \\ \frac{\gamma_U}{c - \rho_U}, & \text{if } \beta = (\frac{1}{|U|-1}, \infty). \end{cases}$$

Law of the transition time

Theorem 2: Law in the external model

$$\lim_{r\to\infty} \mathbb{P}\left(\frac{\tau_v^{\text{ext}}}{\mathbb{E}[\tau_v^{\text{ext}}]} > x\right) = \mathcal{P}(x), \qquad x \in [0,\infty).$$

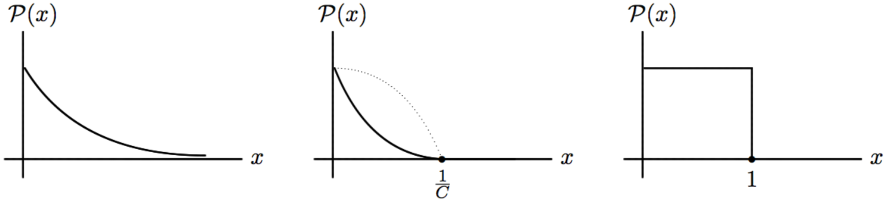

Trichotomy for $x \mapsto \mathcal{P}(x)$.

- Subcritical: exponential decay, $\mathcal{P}_1(x) = e^{-x}$.
- Critical: polynomial decay, $\mathcal{P}_2(x) = (1 - Cx)^{\frac{1-C}{C}}$.
- Supercritical: $\mathcal{P}_3(x)$, cut-off.

Internal model

For any perturbation $\delta > 0$ small enough, there exists a **coupling** such that

$$\lim_{r\to\infty} \hat{\mathbb{P}}\big(\tau_v^{\mathrm{low}} \le \tau_v^{\mathrm{int}} \le \tau_v^{\mathrm{upp}}\big) = 1,$$

where $\hat{\mathbb{P}}$ is the joint law induced by the coupling, with all three models starting from the configuration u.

Theorem 3: Mean transition time and law in the internal model

$$\mathbb{E}[\tau_v^{\mathrm{int}}] = F r^{1\vee\beta(|U|-1)}\,[1+o(1)], \qquad r\to\infty,$$

and

$$\lim_{r\to\infty} \mathbb{P}\left(\frac{\tau_v^{\mathrm{int}}}{\mathbb{E}[\tau_v^{\mathrm{int}}]} > x\right) = \mathcal{P}(x), \qquad x\in[0,\infty).$$

Large deviations

Take any small $\delta > 0$.

- With high probabiliy the input process $Q^+(\cdot)$ follows a path close to its mean $\mathbb{E}[Q^+(t)] = \rho t$,

$$\lim_{r\to\infty} \mathbb{P}\big(\rho t - \delta r \leq Q^+(t) \leq \rho t + \delta r \;\; \forall t \geq 0\big) = 1.$$

- With high probability the output process $Q^-(\cdot)$ follows a path close to the deterministic path ct,

$$\lim_{r\to\infty} \mathbb{P}\big(ct - \delta r \leq Q^-(t) \leq ct \;\; \forall t \in [0, T_U]\big) = 1.$$

We use **large deviations** theorems such as Mogulskii's Theorem and Cramér's Theorem.

Bounds and auxiliary models

We can control the queue length process via a **tube** around its path: we have lower and upper bounds for $Q(t)$.

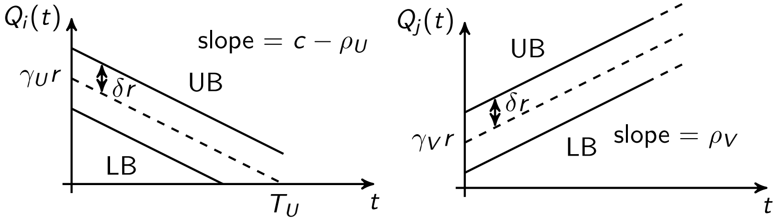

Two auxiliary external models where the activation rates depend on the above bounds:

- **Lower external model**: U less aggressive, V more aggressive.
- **Upper external model**: U more aggressive, V less aggressive.

Coupling results and sandwich

- Coupling the internal with the lower external model, we get

$$\lim_{r\to\infty} \hat{\mathbb{P}}(\tau_v^{\mathrm{low}} \leq \tau_v^{\mathrm{int}}) = 1.$$

- In a similar way, coupling with the upper external model, we get

$$\lim_{r\to\infty} \hat{\mathbb{P}}(\bar{\tau}_v^{\mathrm{int}} \leq \tau_v^{\mathrm{upp}}) = 1.$$

- **Negligible gap** between pre-transition and transition time:

$$\lim_{r\to\infty} \mathbb{P}\left(\tau_v^{\mathrm{int}} - \bar{\tau}_v^{\mathrm{int}} = o\left(\frac{1}{g_V(r)}\right)\right) = 1.$$

$\Rightarrow$ **Sandwich.** For $\delta > 0$ small enough, there exists a coupling such that

$$\lim_{r\to\infty} \hat{\mathbb{P}}\left(\tau_v^{\mathrm{low}} \leq \tau_v^{\mathrm{int}} \leq \tau_v^{\mathrm{upp}}\right) = 1.$$

Index

Aribitrary bipartite graphs

Network extension: more general bipartite graphs, like a cyclic ladder, an hypercube, an even torus...

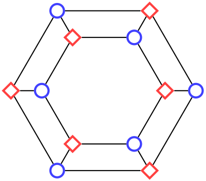

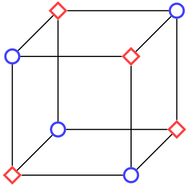

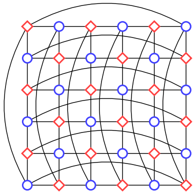

The pre-transition time does not play a key role anymore: the time between the first activation of a node in V and the transition time can be very large.

$\Longrightarrow$ How does the system behave after the first activation?

Ideas

- Nodes in V activates one by one: path according to first activations. When a node in V is active, its neighbors are blocked forever.
- We define a **greedy algorithm** that describes the most likely paths the system follows.
- The study of the transition time of the system can be reduced to the study of the transition time along a fixed path generated by the algorithm. The transition time will be given by the sum of nucleation times of a **sequence of complete bipartite subgraphs**.
- By analyzing the queue length behavior after each activation, we are able to understand how the pre-factor of each nucleation time changes.
- We are able to give explicit asymptotics for the mean transition time when $r \to \infty$ and to describe its law.

The algorithm

Bipartite graph $G = ((U, V), E)$ with $|U| = 6$ and $|V| = 4$.

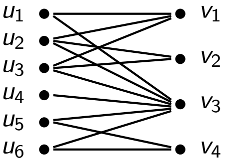

Generated path: v_2, v_1, v_4, v_3.

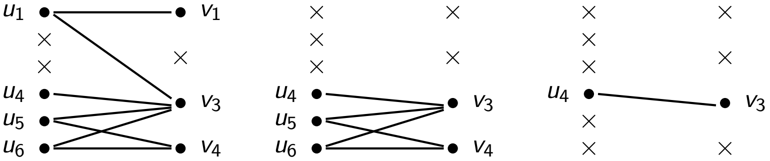

Iteration procedure

We start with $G = G_1 = ((U_1, V_1), E_1)$ and iterate the following **procedure** until V_{k+1} is empty.

1. Consider the graph G_k.
2. Look at the nodes in V_k and at the minimum degree $\bar{d}_k$ in G_k.
3. Pick a node uniformly at random from the ones with minimum degree. Denote the chosen node by a_k.
4. Eliminate the node a_k, all its neighbors in U_k and all their edges. Denote the new bipartite graph by G_{k+1}.

To each iteration $k = 1, \ldots, N$ is associated a **mean nucleation time**, given by

$$\mathbb{E}[\mathcal{T}_{a_k}^{Q^{k-1}}] = F_c(Q^{k-1}, \bar{d}_k) r^{1 \vee \beta(\bar{d}_k - 1)} [1 + o(1)], \qquad r \to \infty.$$

Most likely pahts

$\mathcal{A} \to$ set of all possible paths generated by the algorithm.
$a^* = (a_1^*, \ldots, a_N^*) \to$ path that the system follows.

Theorem 4: Most likely paths

(i) With high probability when $r \to \infty$ the system follows one of the paths generated by the algorithm, i.e.,

$$\lim_{r \to \infty} \mathbb{P}(a^* \in \mathcal{A}) = 1.$$

(ii) Given the path $a = (a_1, \ldots, a_N) \in \mathcal{A}$ and the event $A = \{a^* = a\}$,

$$\mathbb{E}[\mathcal{T}_G^{Q^0} | A] = \sum_{k=1}^{N} \mathbb{E}[\mathcal{T}_{a_k}^{Q^{k-1}}], \qquad r \to \infty.$$

Motivation

- Let $\mathcal{E} = \{\text{The system follows the algorithm}\}$.
- The mean transition time of the graph G can be written as

$$\mathbb{E}_r[\mathcal{T}_G^{Q^0}] = \mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_{\mathcal{E}}] + \mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_{\mathcal{E}^C}],$$

where

$$\mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_{\mathcal{E}^C}] = \mathbb{E}_r[\mathcal{T}_G^{Q^0}|\mathcal{E}^C]\,\mathbb{P}_r(\mathcal{E}^C). \tag{1}$$

- Since $\mathbb{P}(\mathcal{E}^C) \to 0$ as $r \to \infty$, we are interested in the term $\mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_{\mathcal{E}}]$ and we can split it as

$$\mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_{\mathcal{E}}] = \sum_{a\in\mathcal{A}} \mathbb{E}_r[\mathcal{T}_G^{Q^0}\mathbb{1}_A] = \sum_{a\in\mathcal{A}} \mathbb{E}_r[\mathcal{T}_G^{Q^0}|A]\mathbb{P}(A). \tag{2}$$

Note that we can easily recover the probability $\mathbb{P}(A)$ of each path from the algorithm.

Mean transition time

Recall that $g_U(x) = Gx^{\beta}$.
Order: $d^* = \max_k \bar{d}_k$.
Pre-factor: $k_1, k_2 =$ number of nucleations such that $\bar{d}_k = d^*$.

Theorem 5: Mean transition time

$$\mathbb{E}[\mathcal{T}_G^{Q^0}] = \begin{cases} \frac{k_1 \gamma_U^{\beta(d^*-1)}}{d^* G^{-(d^*-1)}} r^{\beta(d^*-1)} [1+o(1)], & \text{if } \beta \in (0, \frac{1}{d^*-1}), \\ \sum_{k:\, \bar{d}_k = d^*} \frac{\gamma_U^{(h_k)}}{d^* G^{-(d^*-1)} + (c-\rho_U)} r\, [1+o(1)], & \text{if } \beta = \frac{1}{d^*-1}, \\ \frac{\gamma_U}{c-\rho_U} r\, [1+o(1)], & \text{if } \beta \in (\frac{1}{d^*-1}, \infty). \end{cases}$$

Recall that we are conditioning the system on the event A. In the subcritical and supercritical regimes, the result is actually **independent** from which path we condition on.

Law of the transition time

Let $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ be probability distributions as in Theorem 2. The law of the transition time is given by a **convolution** of these laws.

Theorem 6: Law of the transition time

$$\lim_{r\to\infty} \mathbb{P}\left(\frac{\mathcal{T}_G^{Q^0}}{\mathbb{E}[\mathcal{T}_G^{Q^0}]} > x\right) = \mathcal{P}^c(x), \qquad x \in [0,\infty), \tag{3}$$

with

$$\mathcal{P}^c(x) = \begin{cases} (\underbrace{\mathcal{P}_1 * \cdots * \mathcal{P}_1}_{k_1 \text{times}})(x), & \text{if } \beta \in (0, \frac{1}{d^*-1}), \\ (\underbrace{\mathcal{P}_2 * \cdots * \mathcal{P}_2}_{k_2 \text{times}})(x), & \text{if } \beta = \frac{1}{d^*-1}, \\ \mathcal{P}_3(x), & \text{if } \beta \in (\frac{1}{d^*-1}, \infty). \end{cases} \tag{4}$$

Updated queue lengths

- If $\beta \in (0, \frac{1}{d^*-1})$, at any step k, the queue length at a node in U is
$$Q_U^k = \gamma_U r\,[1 + o(1)], \qquad r \to \infty.$$
- If $\beta = \frac{1}{d^*-1}$, at any step k, the queue length at a node in U after activating h_k critical nodes in V is
$$Q_U^k = \gamma_U^{(h_k)} r\,[1 + o(1)], \qquad r \to \infty,$$
with
$$\gamma_U^{(h_k)} = \left(\gamma_U - (c - \rho_U)\sum_{i=1}^{h_k} Z_i\right),$$
where $(Z_i)_{i=1}^N$ is a family of random variables.
- If $\beta \in (\frac{1}{d^*-1}, \infty)$, at any step k, the queue length at a node in U after activating the first supercritical node in V is
$$Q_U^k = o(1), \qquad r \to \infty.$$

Index

Future works

- **Protocol extension**: activation rates depend also on the queue lengths of neighbouring nodes.

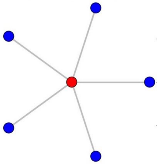

- **Network extension**: focus on different types of interference graph, not necessary bipartite.
- **Activation rates**: consider more general activation rates g_U, g_V or relax the aggressiveness assumption.

Thank you for your attention.