An Exact Analysis of a Class of Markovian Bitcoin Models

Kayla Javier and Brian Fralix
School of Mathematical and Statistical Sciences, and CORI
Clemson University
Clemson, South Carolina, USA

10th Annual MAM Conference
University of Tasmania
15 February 2019

Overview of Talk

We consider two different continuous-time Markov chain (CTMC) models recently proposed in the work of Göbel et al (2016) to model a type of interaction between two pools of miners in the Bitcoin blockchain.

- One pool is a smaller pool, which mines blocks (containing transaction data) and adds them to the blockchain at a rate λ_1.
- Another pool is a larger group that mines blocks at a rate $\lambda_2 > \lambda_1$
- Miners within a pool can communicate with one another instantaneously, yet pools can only communicate with each other after an exponentially distributed amount of time with rate μ.

What happens when the smaller pool tries to withhold information from the larger pool on the number of new blocks it has created since both pools last had identical knowledge of the blockchain?

Our Contribution: We perform a further analysis of the CTMCs introduced in Göbel et al (2016), with 'matrix-analytic methods' (extensions discussed as well, still a work-in-progress).

Literature Review

The following papers contain applied probability models used to describe aspects of the Bitcoin blockchain:

- Bowden. R., Keeler, H.P., Krzesinski, A.E., and Taylor, P.G. (arXiv, 2018)
- Frolkova, M., and Mandjes, M. (*Stochastic Models*, to appear).
- Göbel, J., Keeler, H.P., Krzesinski, A.E., and Taylor, P.G. (*Performance Evaluation*, 2016)
- Huberman, G., Leshno, J., and Moallemi, C. (*Technical Report*, Columbia, 2018)
- Kasahara, S., and Kawahara, J. (arXiv, 2017).
- Li, Q., Ma, J., and Chang, Y. (arXiv, 2018).

Model 1: When Everyone is Honest

Suppose we have two groups of miners working on adding blocks to the blockchain: a smaller pool, and a larger group.

- The smaller pool discovers new blocks in accordance to a Poisson process with rate λ_1.
- The larger group discovers new blocks in accordance to a Poisson process with rate λ_2: since this pool is larger, we assume $\lambda_1 < \lambda_2$ (not necessary from a purely mathematical point of view).
- 'Communication instants' occur in accordance to a Poisson process with rate μ: if the number of newly-discovered blocks by the pool is equal to the number of such blocks discovered by the larger group, both groups continue mining. Otherwise,
- both groups begin using the larger collection of new blocks.

Model 1: When Everyone is Honest

This process can be modeled as a CTMC $\{X(t); t \geq 0\}$.

- The state space of this CTMC is given by S, where

$$S := \{(i, j) : i \geq 0, j \geq 0, i, j \in \mathbb{Z}\}.$$

- It will help to express S as a partition $\{D_k\}_{k \in \mathbb{Z}}$, where for each $k \in \mathbb{Z}$,

$$D_k := \{(i, j) : j - i = k\}.$$

Model 1: When Everyone is Honest

- The transition rate matrix (generator matrix) $\mathbf{Q} := [q(x,y)]_{x,y \in S}$ has elements that are defined as follows:

$$q((i,j),(k,\ell)) := \begin{cases} \lambda_1, & k = i+1,\ \ell = j; \\ \lambda_2, & k = i,\ \ell = j+1; \\ \mu, & k = \ell = 0,\ i \neq j; \end{cases}$$

  with all other off-diagonal entries of $\mathbf{Q}$ set equal to zero.

- $\{X(t); t \geq 0\}$ is known to be ergodic when λ_1, λ_2, and μ are all (strictly) positive.
- First Question: How much is known about the (unique) stationary distribution $\mathbf{p} := [p_x]_{x \in S}$ of this CTMC?

Model 1: When Everyone is Honest

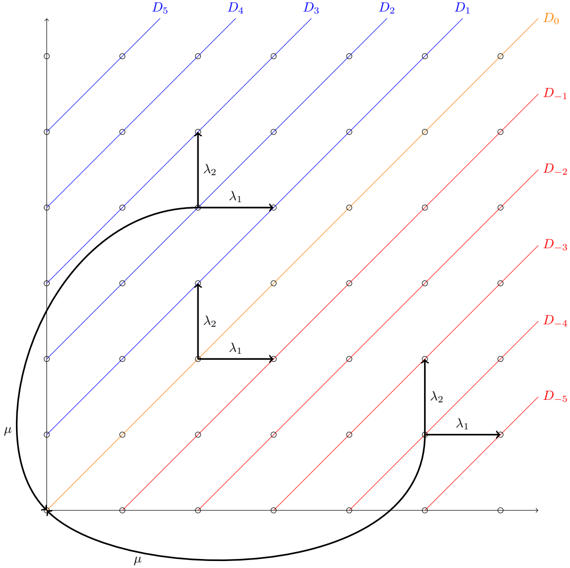

The ‘most important’ diagonal in this picture is D_0.

Model 1: When Everyone is Honest

- Second Question: How much is known about the time-dependent (transient) behavior of this CTMC? **We will have to assume that $X(0) = (0,0)$ w.p.1**.
- Given each state $(i,j) \in S$, we define the function $p_{(0,0),(i,j)} : [0,\infty) \to [0,1]$ as

$$p_{(0,0),(i,j)}(t) := \mathbb{P}(X(t) = (i,j) \mid X(0) = (0,0)).$$

- Further associated with the function $p_{(0,0),(i,j)}$ is its Laplace transform $\pi_{(i,j)}$, defined as

$$\pi_{(i,j)}(\alpha) := \int_0^\infty e^{-\alpha t} p_{(0,0),(i,j)}(t)dt, \qquad \alpha \in \mathbb{C}_+$$

where

$$\mathbb{C}_+ := \{\alpha \in \mathbb{C} : Re(\alpha) > 0\}.$$

Model 1: When Everyone is Honest

Theorem

The stationary distribution of the honest-mining CTMC is as follows: for $(i,j) \neq (0,0)$, $p_{(i,j)}$ is simply

$$p_{(0,0)} \sum_{x=0}^{\min(i,j)} \left[\frac{2^x (x + |i-j|)}{i+j-x} \binom{i+j-x}{j} \right] \frac{\lambda_1^i \lambda_2^j}{(\lambda_1+\lambda_2)^x (\lambda_1+\lambda_2+\mu)^{i+j-x}}.$$

Furthermore,

$$p_{(0,0)} = \frac{1 - \frac{2\lambda_1}{\lambda_1+\lambda_2}\phi_1(\mu)}{1 + \frac{\lambda_1+\lambda_2}{\mu} - \frac{2\lambda_1}{\mu}\phi_1(\mu)}$$

where ϕ_1 is the Laplace-Stieltjes transform of the busy period of an $M/M/1$ queue, having arrival rate λ_1 and service rate λ_2.

Model 1: When Everyone is Honest

Theorem

The Laplace-Stieltjes transforms of this CTMC is as follows: for $(i,j) \neq (0,0)$, $\pi_{(i,j)}$ satisfies

$$\pi_{(i,j)}(\alpha) = \pi_{(0,0)}(\alpha) \sum_{x=0}^{\min(i,j)} d_x(i,j) \frac{\lambda_1^i \lambda_2^j}{(\lambda_1 + \lambda_2 + \alpha)^x (\lambda_1 + \lambda_2 + \mu + \alpha)^{i+j-x}}.$$

Furthermore,

$$\pi_{(0,0)}(\alpha) = \frac{(\mu + \alpha)\left[1 - \frac{2\lambda_1}{\lambda_1+\lambda_2+\alpha}\phi_1(\mu + \alpha)\right]}{\alpha\mu\left[1 + \frac{\lambda_1+\lambda_2+\alpha}{\mu} - \frac{2\lambda_1\phi_1(\mu+\alpha)}{\mu}\right]}.$$

Here

$$d_x(i,j) = \frac{2^x(x + |i - j|)}{i + j - x}\binom{i + j - x}{j}.$$

Model 1: When Everyone is Honest

A few remarks need to be made:

- The expression previously given for $p_{(i,j)}$ (in terms of $p_{(0,0)}$) can also be found in the work of Göbel et al (2016).
- We give an alternative derivation of this formula, which involves making use of the *random-product technique* (brief description given shortly).
- Our expression for $p_{(0,0)}$ can be derived by making use of an M/M/1-like queueing structure hidden within the model.
- The Laplace transforms of the transition functions can be derived in an analogous manner, although the analysis is slightly more difficult.
- The analysis of our next model does involve ideas more explicitly related to Matrix-Analytic Methods.

Model 1: Calculating $p_{(0,0)}$

We can calculate $p_{(0,0)}$ by making use of the well-known fact:

$$p_{(0,0)} = \frac{1}{q((0,0))\mathbb{E}_{(0,0)}[\tau_{(0,0)}]} = \frac{1}{(\lambda_1 + \lambda_2)\mathbb{E}_{(0,0)}[\tau_{(0,0)}]}$$

where

$$\tau_{(0,0)} := \inf\{t \geq 0 : X(t-) \neq X(t) = (0,0)\}.$$

- First step: Show that $\mathbb{E}_{(i,i)}[\tau_{(0,0)}] = \mathbb{E}_{(0,0)}[\tau_{(0,0)}]$, for each $i \geq 0$.
- Next step: Show that $\mathbb{E}_{(0,1)}[\tau_{(0,0)}]$ can be expressed in terms of $\mathbb{E}_{(0,0)}[\tau_{(0,0)}]$, plus the expected busy period length of an $M/M/1$ model with exponential clearing.
- Next step: Show that $\mathbb{E}_{(1,0)}[\tau_{(0,0)}]$ can be expressed in a similar way.
- Finally: Solve for $\mathbb{E}_{(0,0)}[\tau_{(0,0)}]$.

Model 1: When Everyone is Honest

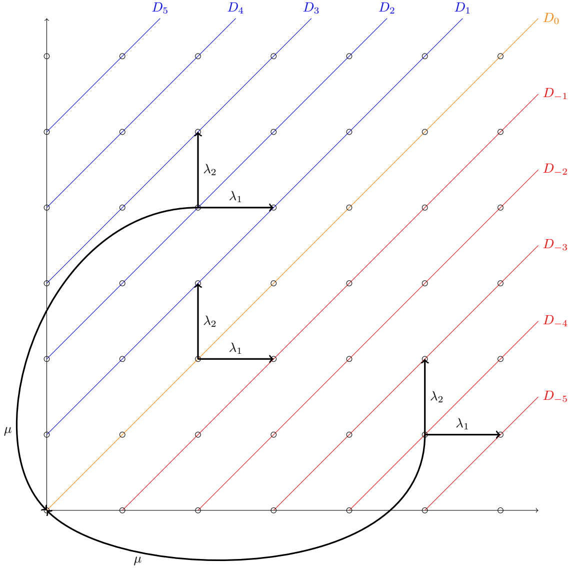

The ‘most important’ diagonal in this picture is D_0.

The Random Product Technique

Given $\{X(t); t \geq 0\}$, construct another CTMC $\{\tilde{X}(t); t \geq 0\}$ whose transition rate matrix $\tilde{\mathbf{Q}}$ satisfies two properties:

- $\tilde{q}(x, y) > 0$ if and only if $q(y, x) > 0$ for each $x, y \in E$.
- $\sum_{y \neq x} \tilde{q}(x, y) = \sum_{y \neq x} q(x, y)$ for each $x \in E$.

In other words, $\tilde{X}$ can make a one-step transition from state x to state y if and only if X can make a one-step transition from state y to state x.

Furthermore, each sojourn in state x is exponentially distributed with rate $q(x) := \sum_{y \neq x} q(x, y)$ under both X and $\tilde{X}$.

The Random Product Technique

Let $\{\tilde{T}_n\}_{n\geq 0}$ denote the transition times of $\{\tilde{X}(t); t \geq 0\}$, where $\tilde{T}_0 := 0$, and for each integer $n \geq 0$, define

$$\tilde{X}_n := \tilde{X}(T_n).$$

Furthermore, define the hitting times

$$\tilde{\tau}_x := \inf\{t \geq 0 : \tilde{X}(t) = x\}, \qquad \tilde{\eta}_x := \inf\{n \geq 0 : \tilde{X}_n = x\}.$$

Finally, for each $x, y \in E$ we define $w_{x,y} : \mathbb{C}_+^0 \to \mathbb{C}$ on the set $\mathbb{C}_+^0 := \{0\} \cup \mathbb{C}_+$ as

$$w_{x,y}(\alpha) := \mathbb{E}_y \left[\mathbf{1}(\tilde{\eta}_x < \infty) e^{-\alpha \tilde{\tau}_x} \prod_{\ell=1}^{\tilde{\eta}_x} \frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)} \right].$$

Notice that $w_{x,x}(0) = 1$ for each $x \in E$.

The Random Product Technique

For each $x, y \in E$, we have that for $\alpha \in \mathbb{C}_+^0$,

$$\pi_{x,y}(\alpha) = \pi_{x,x}(\alpha) w_{x,y}(\alpha).$$

Furthermore, if $\{X(t); t \geq 0\}$ is both irreducible and positive recurrent, then the elements of its stationary distribution $\mathbf{p}$ satisfy

$$p_y = p_x w_{x,y}(0).$$

Once each $w_{x,y}(\alpha)$ term is known, then the remaining $\pi_{x,x}(\alpha)$ term is known as well, since for $\alpha \in \mathbb{C}_+$,

$$\sum_{y \in E} \pi_{x,y}(\alpha) = \frac{1}{\alpha}.$$

A similar statement can be made about p_x when each $w_{x,y}(0)$ term is known: from here onward we express each $w_{(i,j)}(0)$ term simply as $w_{(i,j)}$.

Model 1: Calculating $p_{(i,j)}$

Now that we have $p_{(0,0)}$, calculating $p_{(i,j)}$ exactly amounts to calculating $w_{(i,j)}$. Here

$$\begin{aligned} w_{(i,j)} &:= \mathbb{E}_{(i,j)}\left[\mathbf{1}(\tilde{\eta}_{(0,0)} < \infty) \prod_{\ell=1}^{\tilde{\eta}_{(0,0)}} \frac{q(\tilde{X}_\ell, \tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1}, \tilde{X}_\ell)}\right] \\ &= \sum_{x=0}^{\min(i,j)} d_x(i,j) \frac{\lambda_1^i \lambda_2^j}{(\lambda_1+\lambda_2)^x (\lambda_1+\lambda_2+\mu)^{i+j-x}} \end{aligned}$$

where $d_x(i,j)$ represents the number of paths from (i,j) to $(0,0)$ that make transitions from D_0 exactly x times.

Idea: Sum over paths, cancel ‘tilde-terms’ and carefully study the remaining product associated with each feasible path of $\tilde{X}$ from (i,j) to $(0,0)$.

Model 1: Calculating $p_{(i,j)}$

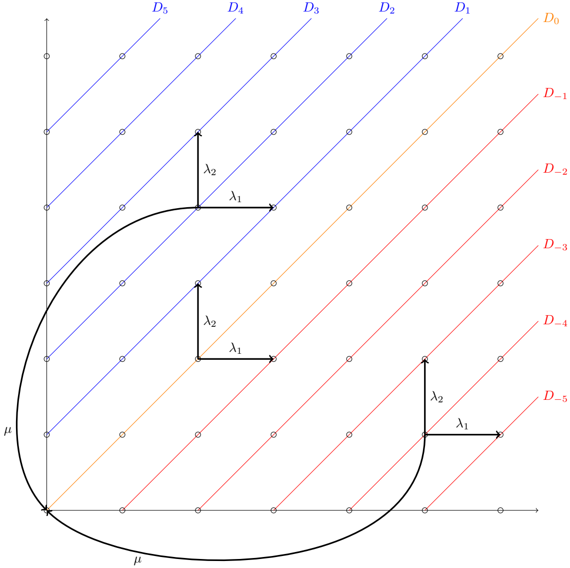

The ‘most important’ diagonal in this picture is D_0.

Model 1: Calculating $p_{(i,j)}$

How do we calculate $d_x(i,j)$? First, observe that from the well-known ballot theorem (see e.g. Renault (2007)) we have

$$d_0(i,j) = \frac{|i-j|}{j+i}\binom{j+i}{i}$$

and since, for each $x \geq 1$,

$$d_x(i,j) = \sum_{\ell=x}^{\min(i,j)} d_x(\ell,\ell)d_0(i-\ell,j-\ell)$$

it suffices to find $d_x(\ell,\ell)$ for each $\ell \geq 1$.

The key identity needed here is the Rothe-Hagen Identity, which is Identity (5.63) on page 202 of *Concrete Mathematics* by Graham, Knuth, and Patashnik.

Model 1: Calculating $p_{(i,j)}$

From this identity, combined with an induction argument, one can see that for each $x \geq 1$, and each $i \geq x$,

$$d_x(i,i) = \frac{x2^x}{2i-x}\binom{2i-x}{i}.$$

This observation, combined again with the Rothe-Hagen identity gives

$$d_x(i,j) = \frac{2^x(x+|i-j|)}{i+j-x}\binom{i+j-x}{j}.$$

This coincides with the expression already found in Göbel et al (2016).

Model 2: When some miners are selfish

Göbel et al also consider another CTMC $\{X(t); t \geq 0\}$ defined on S, where the off-diagonal elements of $\mathbf{Q}$ are now as follows:

$$q((i,j),(k,l)) := \begin{cases} \lambda_1, & i \geq 0,\ j \geq 0,\ k = i+1,\ l = k; \\ \lambda_2, & i \geq 0,\ j \geq 0,\ k = i,\ l = k+1; \\ \mu, & k = l = 0,\ (i,j) \in [\bigcup_{k=1}^{\infty} D_k] \cup [D_{-1} \setminus \{(1,0)\}]; \\ 0, & \text{otherwise.} \end{cases}$$

The next slide illustrates the transition structure.

Model 2: When some miners are selfish

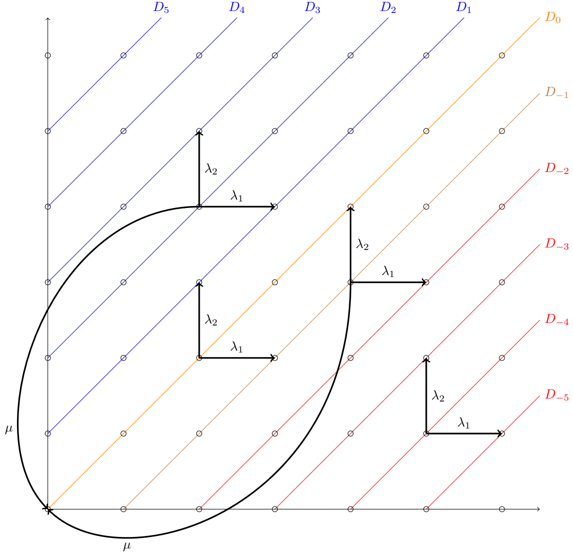

The important diagonals are D_0 and D_{-1}: state $(1, 0)$ is special.

Model 2: When some miners are selfish

The stationary distribution of this CTMC exists whenever $\mu > 0$, and $\lambda_1 < \lambda_2$. While it is not as tractable as the stationary distribution of Model 1, we can say the following:

- An explicit expression can be found for $p_{(0,0)}$.
- All other probabilities $p_{(i,j)}$, for $(i,j) \in (D_0 \cup D_{-1})^c$, can be expressed entirely in terms of a finite number of stationary probabilities associated with states in $D_0 \cup D_{-1}$.
- The stationary probabilities corresponding to states in $D_0 \cup D_{-1}$ satisfy a simple recursion that is analogous to Ramaswami's formula.

Model 2: When some miners are selfish

Theorem

The stationary probability $p_{(0,0)}$ is given by

$$p_{(0,0)} = (\lambda_1 + \lambda_2) \left[\begin{aligned} & \frac{\mu(\lambda_2 + \lambda_1(\lambda_2 - \lambda_1)) + \lambda_2(1 - \phi_2(\mu)(\lambda_2 - \lambda_1))}{\mu(\lambda_2 - \lambda_1)} \\ + \quad & \lambda_1^2 \mathbb{E}_{(2,1)}[\tau_{(0,0)}] \\ + \quad & ((\lambda_1 + \lambda_2)\lambda_2\phi_2(\mu) + \lambda_1\lambda_2)\mathbb{E}_{(1,1)}[\tau_{(0,0)}] \end{aligned} \right]^{-1}$$

where both $\mathbb{E}_{(2,1)}[\tau_{(0,0)}]$ and $\mathbb{E}_{(1,1)}[\tau_{(0,0)}]$ can be expressed in closed-form.

Note that ϕ_2 is the Laplace-Stieltjes transform of the busy period of an M/M/1 queue, having arrival rate λ_2 and service rate λ_1.

Model 2: When some miners are selfish

The key to deriving a closed-form expression for $p_{(0,0)}$ involves observing that we can calculate both $\mathbb{E}_{(1,1)}[\tau_{(0,0)}]$ and $\mathbb{E}_{(2,1)}[\tau_{(0,0)}]$ explicitly.

Indeed, we show that these two expectations satisfy the following two equalities:

$$\begin{aligned}
\mathbb{E}_{(2,1)}[\tau_{(0,0)}] &= \frac{1}{\lambda_2+\mu}\left[\frac{\lambda_2}{\lambda_2-\lambda_1}\right]+\frac{\lambda_2}{\lambda_2+\mu}\mathbb{E}_{(1,1)}[\tau_{(0,0)}] \\
\mathbb{E}_{(1,1)}[\tau_{(0,0)}] &= \frac{1}{\mu}\left[\frac{\mu+\lambda_2(1-\phi_2(\mu))}{\lambda_1+\lambda_2(1-\phi_2(\mu))}\right]+\frac{\lambda_1}{\lambda_1+\lambda_2(1-\phi_2(\mu))}\mathbb{E}_{(2,1)}[\tau_{(0,0)}].
\end{aligned}$$

This system can always be solved, when $\mu > 0$, and $\lambda_1 < \lambda_2$.

Model 2: When some miners are selfish

Here are a number of important observations that can be made about this model:

- For each integer $i \geq 1$, the law of $\tau_{(0,0)}$ under the measure $\mathbb{P}_{(i,i)}$ is the same as the law of $\tau_{(0,0)}$ under the measure $\mathbb{P}_{(1,1)}$.
- For each integer $i \geq 1$, the law of $\tau_{(0,0)}$ under the measure $\mathbb{P}_{(i+1,i)}$ is the same as the law of $\tau_{(0,0)}$ under the measure $\mathbb{P}_{(2,1)}$.
- We use first-step analysis to express $\mathbb{E}_{(0,0)}[\tau_{(0,0)}]$ in terms of $\mathbb{E}_{(0,1)}[\tau_{(0,0)}]$ and $\mathbb{E}_{(1,0)}[\tau_{(0,0)}]$.
- $\mathbb{E}_{(0,1)}[\tau_{(0,0)}]$ can be expressed in terms of $\mathbb{E}_{(1,1)}[\tau_{(0,0)}]$.
- $\mathbb{E}_{(1,0)}[\tau_{(1,0)}]$ can be expressed in terms of $\mathbb{E}_{(2,1)}[\tau_{(0,0)}]$.

Model 2: When some miners are selfish

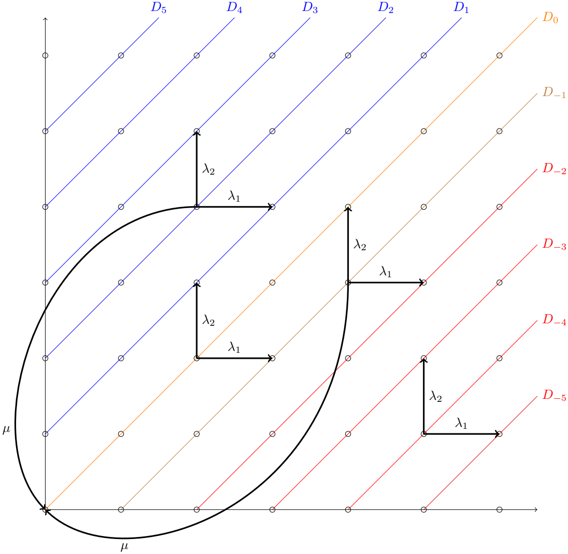

The important diagonals are D_0 and D_{-1}: state $(1, 0)$ is special.

Model 2: When some miners are selfish

Theorem

For each state (i, j) satisfying $i < j$,

$$p_{(i,j)} = \sum_{k=0}^{i} \frac{j-i}{i+j-2k} \binom{i+j-2k}{j-k} \frac{\lambda_1^{i-k}\lambda_2^{j-k}}{(\lambda_1+\lambda_2+\mu)^{i+j-2k}} p_{(k,k)}.$$

For each state (i, j) satisfying $i > j+1$, $j \geq 0$,

$$p_{(i,j)} = \sum_{k=0}^{j} \frac{i-(j+1)}{i+j-2k-1} \binom{i+j-2k-1}{j-k} \frac{\lambda_1^{i-(k+1)}\lambda_2^{j-k}}{(\lambda_1+\lambda_2)^{i+j-2k-1}} p_{(k+1,k)}.$$

We give a quick sketch of the derivation for the case where $i < j$: the other case can be handled in a similar manner.

Model 2: When some miners are selfish

Using the Strong Markov property at $\tilde{\eta}_{D_0}$ gives

$$\begin{aligned} w_{(i,j)} &= \mathbb{E}_{(i,j)}\left[\mathbf{1}(\tilde{\eta}_{(0,0)} < \infty)\prod_{\ell=1}^{\tilde{\eta}_{(0,0)}}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)}\right] \\ &= \sum_{k=0}^{i}\mathbb{E}_{(i,j)}\left[\mathbf{1}(\tilde{\eta}_{D_0} < \infty)\mathbf{1}(\tilde{X}_{\tilde{\eta}_{D_0}} = (k,k))\prod_{\ell=1}^{\tilde{\eta}_{D_0}}\frac{q(\tilde{X}_\ell,\tilde{X}_{\ell-1})}{\tilde{q}(\tilde{X}_{\ell-1},\tilde{X}_\ell)}\right]w_{(k,k)} \end{aligned}$$

where each remaining random-product term can be calculated via a lattice-path counting argument. Next, recall that

$$p_{(i,j)} = p_{(0,0)}w_{(i,j)}.$$

Note: This is a ‘random-product’ interpretation of the familiar $\mathbf{p}_{n+1} = \mathbf{p}_n\mathbf{R}$ equality from the theory of QBD processes.

Model 2: When some miners are selfish

It remains to calculate the $p_{(i,i)}$ and $p_{(i+1,i)}$ stationary probabilities: clearly

$$p_{(1,0)} = \frac{\lambda_1}{\lambda_1 + \lambda_2} p_{(0,0)}$$

and for each $i \geq 1$,

$$\begin{aligned} p_{(i,i)} &= \frac{\lambda_1}{\lambda_1+\lambda_2} \sum_{k=0}^{i-1} \frac{1}{2i-1-2k} \binom{2i-1-2k}{i-k} \frac{\lambda_1^{i-1-k} \lambda_2^{i-k}}{(\lambda_1+\lambda_2+\mu)^{2i-1-2k}} p_{(k,k)} \\ &+ \frac{\lambda_2}{\lambda_1+\lambda_2} p_{(i,i-1)} \end{aligned}$$

and $p_{(i+1,i)}$ is simply

$$\begin{aligned} &= \frac{\lambda_2}{\lambda_1+\lambda_2+\mu} \sum_{k=0}^{i-1} \frac{1}{2i-2k-l} \binom{2i-2k-1}{i-1-k} \frac{\lambda_1^{i-k} \lambda_2^{i-1-k}}{(\lambda_1+\lambda_2)^{2i-2k-1}} p_{(k+1,k)} \\ &+ \frac{\lambda_1}{\lambda_1+\lambda_2+\mu} p_{(i,i)}. \end{aligned}$$

Model 2: When some miners are selfish

This gives us a clear exact, numerical procedure: in order to calculate $p_{(i,j)}$ (for the case where $i < j$; the other case works in a similar manner)

- First calculate $p_{(0,0)}$.
- Next, use the recursion to calculate $p_{(1,0)}, p_{(1,1)}, p_{(2,1)}, \ldots$, up to $p_{(i,i)}$ (this recursion is analogous to Ramaswami's well-known recursion from the theory of Markov processes of M/G/1-type).
- Finally, calculate $p_{(i,j)}$ (which can be expressed in terms of $p_{(0,0)}, p_{(1,1)}, \ldots, p_{(i,i)}$.

A very similar procedure can be used to calculate the Laplace transforms of the transition functions of $\{X(t); t \geq 0\}$, if again, we further assume $X(0) = (0, 0)$.

THANK YOU!