The critical greedy server on the integers

Andrew Wade
Department of Mathematical Sciences
Durham University

Joint work with
James Cruise (Heriot-Watt)

Mikhail Menshikov 70th Birthday Conference
March 2018

The model

Continuous-time model introduced by KURKOVA AND MENSHIKOV (1997); earlier work on related models includes COFFMAN AND GILBERT (1987) and FOSS AND LAST (1996).

Markov Processes Relat. Fields **3**, 243–259 (1997)

Markov Processes and Related Fields

Greedy algorithm, Z^1 case

I.A. Kurkova* and M.V. Menshikov*

Laboratory of Large Random Systems, Faculty of Mathematics and Mechanics, Moscow State University, 119899, Moscow, Russia

Received November 15, 1996

Abstract. We consider a single server queueing system with stations in all integer points of the real line. The customers arrival streams at the different stations are independent Poisson processes. The service times of customers are mutually independent and exponentially distributed. The server serves each station exhaustively, i.e. till the station is empty. The next station to be served, is selected using the greedy algorithm: the server goes to the neighbouring station with the maximum number of customers. We study the trajectories of the server and his asymptotic position, as time tends to infinity.

The model

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.

The model

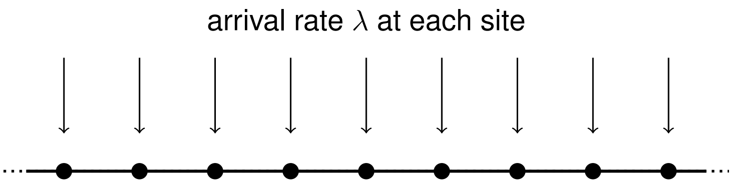

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.

The model

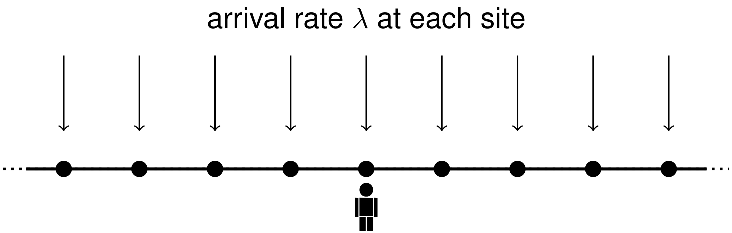

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.
- A single server, who starts at $0 \in \mathbb{Z}$.

The model

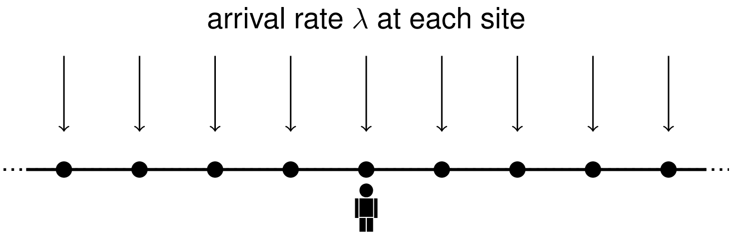

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.
- A single server, who starts at $0 \in \mathbb{Z}$.
- The server serves the current queue until empty, then the server picks the largest neighbouring queue and moves there

The model

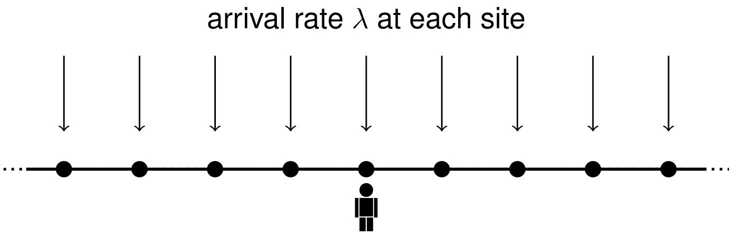

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.
- A single server, who starts at $0 \in \mathbb{Z}$.
- The server serves the current queue until empty, then the server picks the largest neighbouring queue and moves there (randomly break ties)

The model

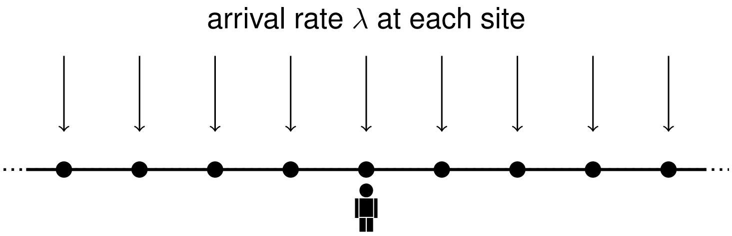

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.
- A single server, who starts at $0 \in \mathbb{Z}$.
- The server serves the current queue until empty, then the server picks the largest neighbouring queue and moves there (randomly break ties), taking time $= 1$ unit to move.

The model

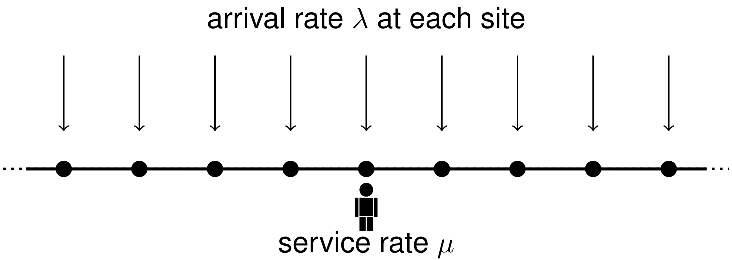

- There is a queue at each site of $\mathbb{Z}$; all empty at time 0.
- Independent Poisson arrivals at rate λ at each queue.
- A single server, who starts at $0 \in \mathbb{Z}$.
- The server serves the current queue until empty, then the server picks the largest neighbouring queue and moves there (randomly break ties), taking time $= 1$ unit to move.
- The service rate when the server is at a queue is μ.

The model

Let $S(t)$ = position of server at time t.

The model

Let $S(t) =$ position of server at time t.

Question: How does $S(t)$ behave as $t \to \infty$?

The model

Let $S(t)$ = position of server at time t.

Question: How does $S(t)$ behave as $t \to \infty$?

Starting point: Under service, a queue is an M/M/1 queue, i.e., a continuous-time random walk on $\mathbb{Z}_+$, with positive jumps at rate λ and (from positive sites) negative jumps at rate μ.

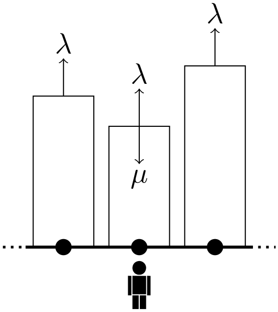

Subcritical case: $\mu < \lambda$

- If $\mu < \lambda$, the queue under service is transient.

Subcritical case: $\mu < \lambda$

- If $\mu < \lambda$, the queue under service is transient.
- So, uniformly over non-empty queues, there is positive probability that service is never completed.

Subcritical case: $\mu < \lambda$

- If $\mu < \lambda$, the queue under service is transient.
- So, uniformly over non-empty queues, there is positive probability that service is never completed.
- So the server empties (at most) finitely many queues, and then gets stuck.

Proposition (KM97)

If $\mu < \lambda$, then $S(t)$ converges to a finite limit in $\mathbb{Z}$, a.s.

Supercritical case: $\mu > \lambda$

- If $\mu > \lambda$, the queue under service is positive recurrent, so will certainly empty.

Supercritical case: $\mu > \lambda$

- If $\mu > \lambda$, the queue under service is positive recurrent, so will certainly empty.
- Moreover, the negative drift ensures that the queue empties in linear time, i.e., a queue of length ℓ takes time about $c\ell$ to empty.

Supercritical case: $\mu > \lambda$

- If $\mu > \lambda$, the queue under service is positive recurrent, so will certainly empty.
- Moreover, the negative drift ensures that the queue empties in linear time, i.e., a queue of length ℓ takes time about $c\ell$ to empty.

Consider the server's first arrival at $x > 0$ at time t. As neither x nor $x+1$ have been previously visited, the picture is:

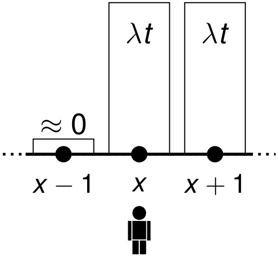

Supercritical case: $\mu > \lambda$

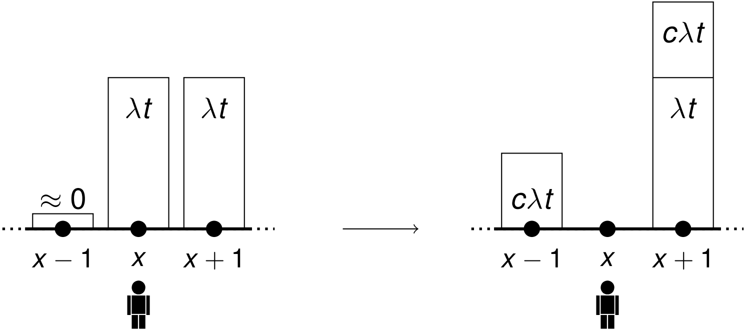

- Takes time $\sim ct$ to serve the queue at x.
- In that time $\sim c\lambda t$ new customers arrive at $x \pm 1$.

Supercritical case: $\mu > \lambda$

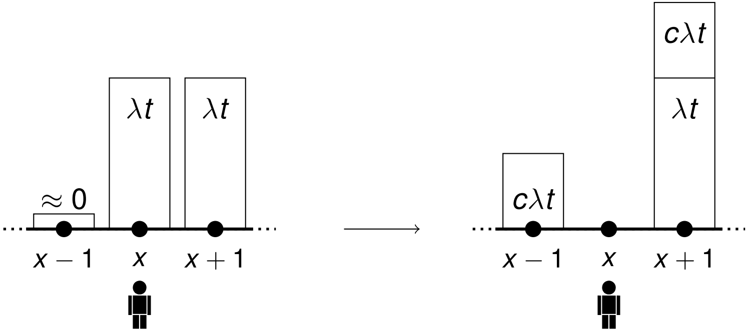

- Takes time $\sim ct$ to serve the queue at x.
- In that time $\sim c\lambda t$ new customers arrive at $x \pm 1$.
- Fluctuations in this picture are $O(\sqrt{t})$.

Supercritical case: $\mu > \lambda$

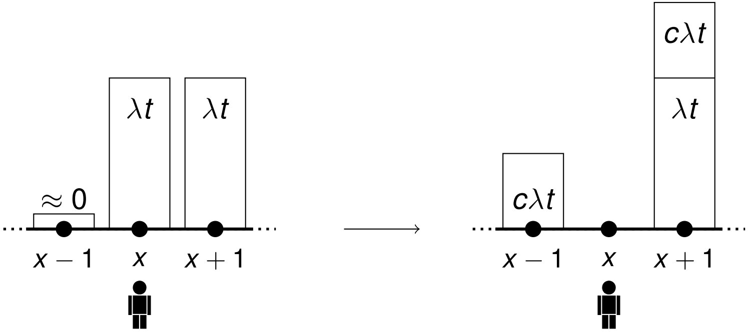

- Takes time $\sim ct$ to serve the queue at x.
- In that time $\sim c\lambda t$ new customers arrive at $x \pm 1$.
- Fluctuations in this picture are $O(\sqrt{t})$.
- With very high probability, the server moves to $x + 1$ next.

Supercritical case: $\mu > \lambda$

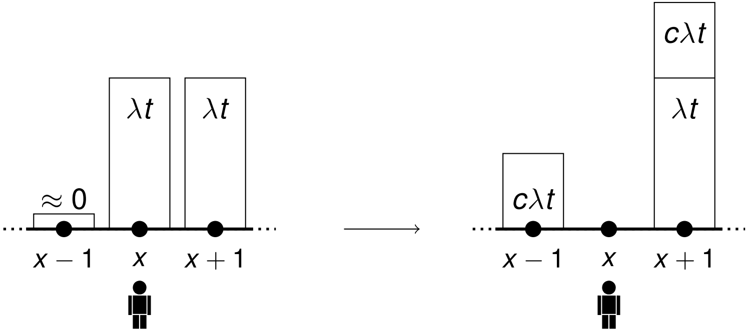

- Takes time $\sim ct$ to serve the queue at x.
- In that time $\sim c\lambda t$ new customers arrive at $x \pm 1$.
- Fluctuations in this picture are $O(\sqrt{t})$.
- With very high probability, the server moves to $x + 1$ next.

Subtlety: At time t queue at x has never been visitied, but has been *inspected*. Small effect…

Supercritical case: $\mu > \lambda$

This is the intuition behind:

Theorem (KM97)

If $\mu > \lambda$, then

$$\mathbb{P}\left(\lim_{t\to\infty} S(t) = +\infty\right) = \mathbb{P}\left(\lim_{t\to\infty} S(t) = -\infty\right) = \frac{1}{2}.$$

Supercritical case: $\mu > \lambda$

This is the intuition behind:

Theorem (KM97)

If $\mu > \lambda$, then

$$\mathbb{P}\left(\lim_{t\to\infty} S(t) = +\infty\right) = \mathbb{P}\left(\lim_{t\to\infty} S(t) = -\infty\right) = \frac{1}{2}.$$

That is, the server is ‘transient’. Moreover, the rate of escape:

Theorem (KM97)

If $\mu > \lambda$, then there is a constant $\rho = \rho(\mu, \lambda) \in (0, \infty)$ such that

$$\mathbb{P}\left(\lim_{t\to\infty} \frac{|S(t)|}{\log t} = \rho\right) = 1.$$

Critical case: $\mu = \lambda$

The critical case $\mu = \lambda$ was left largely open; KM97 did show that the server never gets stuck in a finite region:

$$\limsup_{t \to \infty} |S(t)| = +\infty, \text{ a.s.}$$

New intuition:

- The queue under service now has zero drift and is null recurrent.

Critical case: $\mu = \lambda$

The critical case $\mu = \lambda$ was left largely open; KM97 did show that the server never gets stuck in a finite region:

$$\limsup_{t \to \infty} |S(t)| = +\infty, \text{ a.s.}$$

New intuition:

- The queue under service now has zero drift and is null recurrent.
- So again the queue always empties, but now the time to empty a queue of length ℓ is of order ℓ^2.

Let's try to repeat our argument from the supercritical case.

Critical case: $\mu = \lambda$

Again consider first arrival at site $x > 0$ at time t.

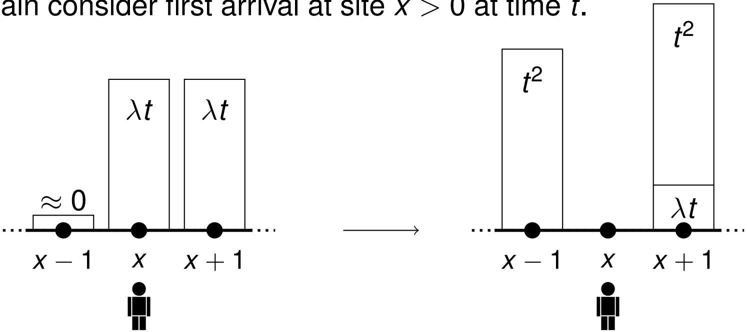

- Takes time $\approx t^2$ to serve the queue at x.
- In that time $\approx t^2$ new customers arrive at $x \pm 1$.

Critical case: $\mu = \lambda$

Again consider first arrival at site $x > 0$ at time t.

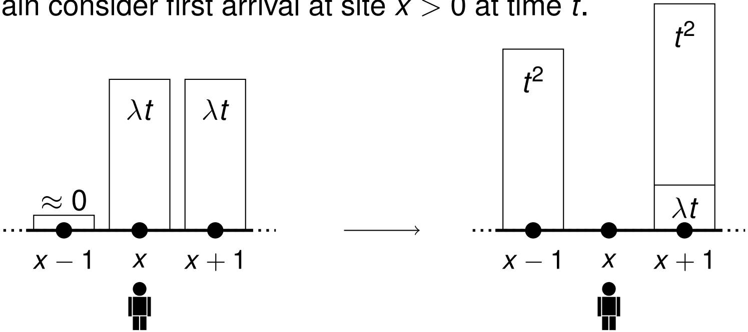

- Takes time $\approx t^2$ to serve the queue at x.
- In that time $\approx t^2$ new customers arrive at $x \pm 1$.
- Fluctuations in the arrivals are $O(t)$, the same order as the rightwards bias.

Critical case: $\mu = \lambda$

Again consider first arrival at site $x > 0$ at time t.

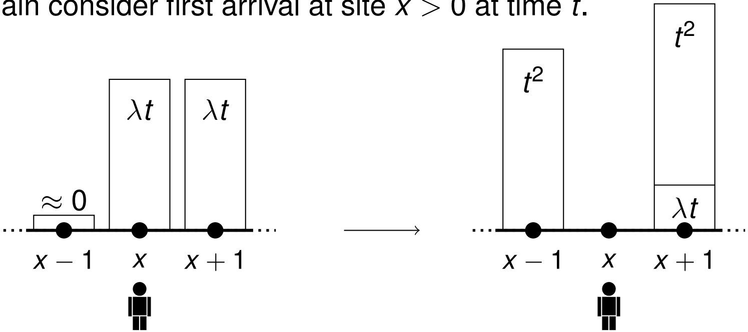

- Takes time $\approx t^2$ to serve the queue at x.
- In that time $\approx t^2$ new customers arrive at $x \pm 1$.
- Fluctuations in the arrivals are $O(t)$, the same order as the rightwards bias.
- Suggests $\mathbb{P}\,$(changes direction) $> \varepsilon > 0$?

So in this case the behaviour is more complicated.

Critical case: results

We show that the server is 'recurrent':

Theorem 1 (CW)

If $\mu = \lambda$, then a.s., for every $x \in \mathbb{R}$,

$$\{t \geq 0 : S(t) = x\} \text{ is unbounded.}$$

Critical case: results

We show that the server is 'recurrent':

Theorem 1 (CW)

If $\mu = \lambda$, then a.s., for every $x \in \mathbb{R}$,

$$\{t \geq 0 : S(t) = x\} \text{ is unbounded.}$$

Moreover, we have an iterated logarithm law for the position:

Theorem 2 (CW)

If $\mu = \lambda$, then, a.s.,

$$\limsup_{t\to\infty} \frac{S(t)}{\sqrt{\log\log t \log\log\log\log t}} = \sqrt{\frac{6}{\log 2}},$$

$$\liminf_{t\to\infty} \frac{S(t)}{\sqrt{\log\log t \log\log\log\log t}} = -\sqrt{\frac{6}{\log 2}}.$$

Ingredients in the proofs

We will outline the main ideas in the proofs.

- The times between emptying successive queues show doubly-exponential growth.

Ingredients in the proofs

We will outline the main ideas in the proofs.

- The times between emptying successive queues show doubly-exponential growth.
- The probability of turning around converges to $1/4$.

Ingredients in the proofs

We will outline the main ideas in the proofs.

- The times between emptying successive queues show doubly-exponential growth.
- The probability of turning around converges to $1/4$.
- A martingale argument.

Ingredients in the proofs

We will outline the main ideas in the proofs.

- The times between emptying successive queues show doubly-exponential growth.
- The probability of turning around converges to $1/4$.
- A martingale argument.

For convenience, take $\lambda = \mu = 1$ from now on.

The critical M/M/1 queue

Let $\zeta(k)$ be the time to serve a queue starting with $k \in \mathbb{N}$ initial customers.

The critical M/M/1 queue

Let $\zeta(k)$ be the time to serve a queue starting with $k \in \mathbb{N}$ initial customers.

In other words, $\zeta(k)$ is the time to hit 0 of a continuous-time simple random walk started from k, with rate 1 of stepping to the left and rate 1 of stepping to the right.

The critical M/M/1 queue

Let $\zeta(k)$ be the time to serve a queue starting with $k \in \mathbb{N}$ initial customers.

In other words, $\zeta(k)$ is the time to hit 0 of a continuous-time simple random walk started from k, with rate 1 of stepping to the left and rate 1 of stepping to the right.

Since the random walk $\Rightarrow$ Brownian motion:

Lemma 3

As $k \to \infty$,

$$\frac{2}{k^2}\zeta(k) \xrightarrow{d} S,$$

where $F_S(u) := \mathbb{P}(S \leq u) = 2\overline{\Phi}(u^{-1/2})$, where $\overline{\Phi}(z) = \mathbb{P}(Z > z)$ for $Z \sim \mathcal{N}(0,1)$.

The critical M/M/1 queue

Let $\zeta(k)$ be the time to serve a queue starting with $k \in \mathbb{N}$ initial customers.

In other words, $\zeta(k)$ is the time to hit 0 of a continuous-time simple random walk started from k, with rate 1 of stepping to the left and rate 1 of stepping to the right.

Since the random walk $\Rightarrow$ Brownian motion:

Lemma 3

As $k \to \infty$,

$$\frac{2}{k^2}\zeta(k) \xrightarrow{d} S,$$

where $F_S(u) := \mathbb{P}(S \le u) = 2\overline{\Phi}(u^{-1/2})$, where $\overline{\Phi}(z) = \mathbb{P}(Z > z)$ for $Z \sim \mathcal{N}(0,1)$.

Note $\mathbb{P}(S > u) \sim cu^{-1/2}$ and S is $1/2$-stable; sometimes known as Lévy distribution.

Time-scale estimates

Let τ_n = time to service the nth queue served.

And let $T_n = \tau_1 + \cdots + \tau_n$.

Lemma 4

As $n \to \infty$,

$$\frac{\tau_n}{\tau_{n-1}^2} \xrightarrow{d} \frac{1}{2} S.$$

Sketch proof.

Let Q_{n-1} be the number of customers at the queue to be served at the start of the nth service.

Then $\tau_n = \zeta(Q_{n-1})$ and $Q_{n-1} \approx \lambda \tau_{n-1}$ (at least...).

By Lemma 3,

$$\frac{\zeta(\tau_{n-1})}{\tau_{n-1}^2} \xrightarrow{d} \frac{1}{2} S. \qquad \square$$

Time-scale estimates

Lemma 5

For any $\alpha \in (0, 2)$ *and* $\beta > 2$*, a.s.,*

$$e^{\alpha^n} < \tau_n < e^{\beta^n} \text{ for all but finitely many } n.$$

Sketch proof.

Up to (random) multiplicative factors, $\tau_n \approx \tau_{n-1}^2$. So if $\tau_1 \approx e^2$, we have $\tau_2 \approx e^4$, $\tau_3 \approx e^8, \ldots, \tau_n \approx e^{2^n}$. □

Time-scale estimates

Lemma 5

For any $\alpha \in (0, 2)$ *and* $\beta > 2$*, a.s.,*

$$e^{\alpha^n} < \tau_n < e^{\beta^n} \text{ for all but finitely many } n.$$

Sketch proof.

Up to (random) multiplicative factors, $\tau_n \approx \tau_{n-1}^2$. So if $\tau_1 \approx e^2$, we have $\tau_2 \approx e^4$, $\tau_3 \approx e^8$, $\ldots$, $\tau_n \approx e^{2^n}$. □

This is the doubly-exponential growth of the service times.

Time-scale estimates

Lemma 5

For any $\alpha \in (0,2)$ *and* $\beta > 2$*, a.s.,*

$$e^{\alpha^n} < \tau_n < e^{\beta^n} \text{ for all but finitely many } n.$$

Sketch proof.

Up to (random) multiplicative factors, $\tau_n \approx \tau_{n-1}^2$. So if $\tau_1 \approx e^2$, we have $\tau_2 \approx e^4$, $\tau_3 \approx e^8$, $\ldots$, $\tau_n \approx e^{2^n}$. $\square$

This is the doubly-exponential growth of the service times.

In particular, $T_n \approx \tau_n$.

Time-scale estimates

Lemma 5

For any $\alpha \in (0, 2)$ *and* $\beta > 2$*, a.s.,*

$$e^{\alpha^n} < \tau_n < e^{\beta^n} \text{ for all but finitely many } n.$$

Sketch proof.

Up to (random) multiplicative factors, $\tau_n \approx \tau_{n-1}^2$. So if $\tau_1 \approx e^2$, we have $\tau_2 \approx e^4$, $\tau_3 \approx e^8$, $\ldots$, $\tau_n \approx e^{2^n}$. □

This is the doubly-exponential growth of the service times.

In particular, $T_n \approx \tau_n$.

Corollary 6

If N_t = *number of queues emptied by time* t*,*

$$\lim_{t \to \infty} \frac{N_t}{\log \log t} = \frac{1}{\log 2}, \text{ a.s.}$$

Turning probability

Let X_n = location of nth served queue. Let $\eta_n = X_n - X_{n-1}$.

Turning probability

Let X_n = location of nth served queue. Let $\eta_n = X_n - X_{n-1}$.
Let $q_n = \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1})$.
Here $\mathcal{F}_n$ = everything up to the end of the nth service.

Turning probability

Let X_n = location of nth served queue. Let $\eta_n = X_n - X_{n-1}$.

Let $q_n = \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1})$.
Here $\mathcal{F}_n$ = everything up to the end of the nth service.

Note that

$$\eta_n = \begin{cases} +1 & \text{if } Q_{n-1}(X_{n-1}+1) > Q_{n-1}(X_{n-1}-1) \\ -1 & \text{if } Q_{n-1}(X_{n-1}+1) < Q_{n-1}(X_{n-1}-1) \end{cases}$$

where $Q_n(x)$ = length of queue at x on completion of nth service.

Turning probability

Let X_n = location of nth served queue. Let $\eta_n = X_n - X_{n-1}$.

Let $q_n = \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1})$.
Here $\mathcal{F}_n$ = everything up to the end of the nth service.

Note that

$$\eta_n = \begin{cases} +1 & \text{if } Q_{n-1}(X_{n-1}+1) > Q_{n-1}(X_{n-1}-1) \\ -1 & \text{if } Q_{n-1}(X_{n-1}+1) < Q_{n-1}(X_{n-1}-1) \end{cases}$$

where $Q_n(x)$ = length of queue at x on completion of nth service.

In particular, both η_n and X_n are $\mathcal{F}_{n-1}$-measurable.

Turning probability

Let X_n = location of nth served queue. Let $\eta_n = X_n - X_{n-1}$.

Let $q_n = \mathbb{P}(\eta_{n+1} \neq \eta_n \mid \mathcal{F}_{n-1})$.
Here $\mathcal{F}_n$ = everything up to the end of the nth service.

Note that

$$\eta_n = \begin{cases} +1 & \text{if } Q_{n-1}(X_{n-1}+1) > Q_{n-1}(X_{n-1}-1) \\ -1 & \text{if } Q_{n-1}(X_{n-1}+1) < Q_{n-1}(X_{n-1}-1) \end{cases}$$

where $Q_n(x)$ = length of queue at x on completion of nth service.

In particular, both η_n and X_n are $\mathcal{F}_{n-1}$-measurable.

Lemma 7

$q_n \to 1/4$, a.s.

Turning probability

Lemma 7

$q_n \to 1/4$, a.s.

Sketch proof.

Suppose the nth queue has just been emptied; the service time was τ_n.

Turning probability

Lemma 7

$q_n \to 1/4$, a.s.

Sketch proof.

Suppose the *n*th queue has just been emptied; the service time was τ_n.

At this time, the *previous* queue has about Po(τ_n) customers, while the other neighbouring queue has about Po($\tau_{n-1} + \tau_n$) (ignoring the much smaller number of prior customers).

Turning probability

Lemma 7

$q_n \to 1/4$, a.s.

Sketch proof.

Suppose the *n*th queue has just been emptied; the service time was τ_n.

At this time, the *previous* queue has about $\mathrm{Po}(\tau_n)$ customers, while the other neighbouring queue has about $\mathrm{Po}(\tau_{n-1} + \tau_n)$ (ignoring the much smaller number of prior customers).

So

$$\mathbb{P}(\eta_{n+1} \neq \eta_n) \approx \mathbb{P}(\mathrm{Po}(\tau_n) > \mathrm{Po}(\tau_{n-1} + \tau_n)).$$

Turning probability

Lemma 7

$q_n \to 1/4$, a.s.

Sketch proof.

Suppose the *n*th queue has just been emptied; the service time was τ_n.

At this time, the *previous* queue has about $\text{Po}(\tau_n)$ customers, while the other neighbouring queue has about $\text{Po}(\tau_{n-1} + \tau_n)$ (ignoring the much smaller number of prior customers).

So

$$\mathbb{P}(\eta_{n+1} \neq \eta_n) \approx \mathbb{P}(\text{Po}(\tau_n) > \text{Po}(\tau_{n-1} + \tau_n)).$$

By the CLT, and using the fact that $\tau_n \gg \tau_{n-1}$,

$$\mathbb{P}(\eta_{n+1} \neq \eta_n) \approx \mathbb{P}(\tau_n + Z\tau_n^{1/2} > \tau_{n-1} + \tau_n + Z'\tau_n^{1/2}),$$

where Z, Z' are independent $\mathcal{N}(0, 1)$.

Turning probability

We have shown that

$$\mathbb{P}(\eta_{n+1} \neq \eta_n) \approx \mathbb{P}(\tau_n + Z\tau_n^{1/2} > \tau_{n-1} + \tau_n + Z'\tau_n^{1/2}),$$

where Z, Z' are independent $\mathcal{N}(0, 1)$.

Turning probability

We have shown that

$$\mathbb{P}(\eta_{n+1} \neq \eta_n) \approx \mathbb{P}(\tau_n + Z\tau_n^{1/2} > \tau_{n-1} + \tau_n + Z'\tau_n^{1/2}),$$

where Z, Z' are independent $\mathcal{N}(0, 1)$.

Hence

$$\begin{aligned}\mathbb{P}(\eta_{n+1} \neq \eta_n) &\approx \mathbb{P}\left(Z - Z' > \frac{\tau_{n-1}}{\tau_n^{1/2}}\right) \\ &\approx \mathbb{P}\left(\sqrt{2}Z > \left(\tfrac{1}{2}S\right)^{-1/2}\right),\end{aligned}$$

by Lemma 4.

Turning probability

We have shown that

$$\mathbb{P}\left(\eta_{n+1} \neq \eta_n\right) \approx \mathbb{P}\left(\tau_n + Z\tau_n^{1/2} > \tau_{n-1} + \tau_n + Z'\tau_n^{1/2}\right),$$

where Z, Z' are independent $\mathcal{N}(0,1)$.

Hence

$$\begin{aligned}\mathbb{P}\left(\eta_{n+1} \neq \eta_n\right) &\approx \mathbb{P}\left(Z - Z' > \frac{\tau_{n-1}}{\tau_n^{1/2}}\right) \\ &\approx \mathbb{P}\left(\sqrt{2}Z > \left(\tfrac{1}{2}S\right)^{-1/2}\right),\end{aligned}$$

by Lemma 4.

By the particular compatibility of the distribution of S (Lemma 3) with the normal distribution, this last probability is $1/4$! □

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

Why? Let $f(x,i) = x + 2\mathbf{1}\{i = 1\}$. Then

$$f(x+i,i) - f(x,i) = i$$
$$f(x-i,-i) - f(x,i) = -3i$$

so

$$\begin{aligned}\mathbb{E}\left(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1}\right) &= \eta_n(1-q_n) - 3\eta_n q_n \\ &= \eta_n(1-4q_n) \\ &\approx 0.\end{aligned}$$

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

Why? Let $f(x, i) = x + 2\mathbf{1}\{i = 1\}$. Then

$$f(x + i, i) - f(x, i) = i$$
$$f(x - i, -i) - f(x, i) = -3i$$

so

$$\begin{aligned}\mathbb{E}\left(Y_{n+1} - Y_n \mid \mathcal{F}_{n-1}\right) &= \eta_n(1 - q_n) - 3\eta_n q_n \\ &= \eta_n(1 - 4q_n) \\ &\approx 0.\end{aligned}$$

The main technical work is getting rates of convergence everywhere to quanitfy this.

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

The main technical work is getting rates of convergence everywhere to quanitfy this.

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

The main technical work is getting rates of convergence everywhere to quanitfy this.

Theorem 1 now follows from the fact that for a martingale M_n with bounded increments,

$$\liminf_{n\to\infty} M_n = -\infty \text{ and } \limsup_{n\to\infty} M_n = +\infty.$$

Final ingredient: martingale

Claim: $Y_n = X_n + 2\mathbf{1}\{\eta_n = +1\}$ is almost a martingale.

The main technical work is getting rates of convergence everywhere to quanitfy this.

Theorem 1 now follows from the fact that for a martingale M_n with bounded increments,

$$\liminf_{n\to\infty} M_n = -\infty \text{ and } \limsup_{n\to\infty} M_n = +\infty.$$

Theorem 2 follows from STOUT's martingale LIL and Corollary 6.

References

- E.G. Coffman & E.N. Gilbert, Polling and greedy servers on a line, *Queueing Syst.* **2** (1987) 115–145.
- J.R. Cruise & A.R. Wade, The critical greedy server on the integers is recurrent. arXiv:1712.03026.
- S. Foss & G. Last, Stability of polling systems with exhaustive service policies and state-dependent routing, *Ann. Appl. Probab.* **6** (1996) 116–137.
- I.A. Kurkova & M.V. Menshikov, Greedy algorithm: Z^1 case, *Markov Process. Relat. Fields* **3** (1997) 243–259.