Omnithermal Perfect Simulation

Stephen Connor
stephen.connor@york.ac.uk

University of York

61st World Statistics Congress
Marrakech, July 2017

Outline

Introduction

CFTP/domCFTP in a nutshell

Suppose that we're interested in simulating from the equilibrium distribution of some ergodic Markov chain X.

Think of a (hypothetical) version of the chain, $\tilde{X}$, which was started by your (presumably distant) ancestor from some state x at time $-\infty$:

- at time zero this chain is in equilibrium: $\tilde{X}_0 \sim \pi$;
- CFTP/domCFTP tries to determine the value of $\tilde{X}_0$ by looking into the past only a *finite* number of steps;
- do this by identifying a time in the past such that *all earlier starts from x lead to the same result at time zero*.

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time

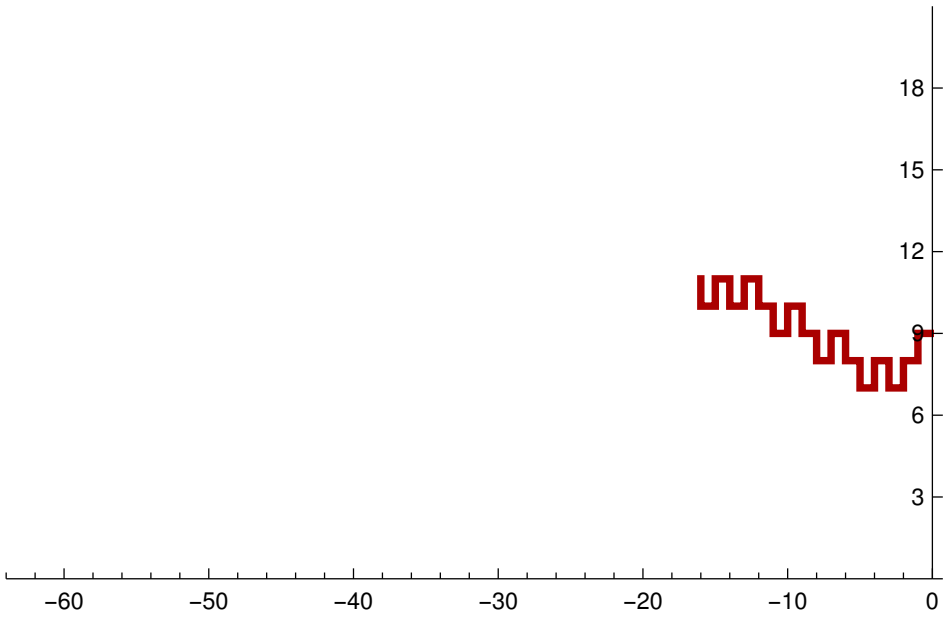

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time
- *sandwiching*

$$\text{Lower}_{\text{late}} \preccurlyeq \text{Lower}_{\text{early}} \preccurlyeq \ldots \preccurlyeq \text{Target} \preccurlyeq \ldots \preccurlyeq \text{Upper}_{\text{early}} \preccurlyeq \text{Upper}_{\text{late}}$$

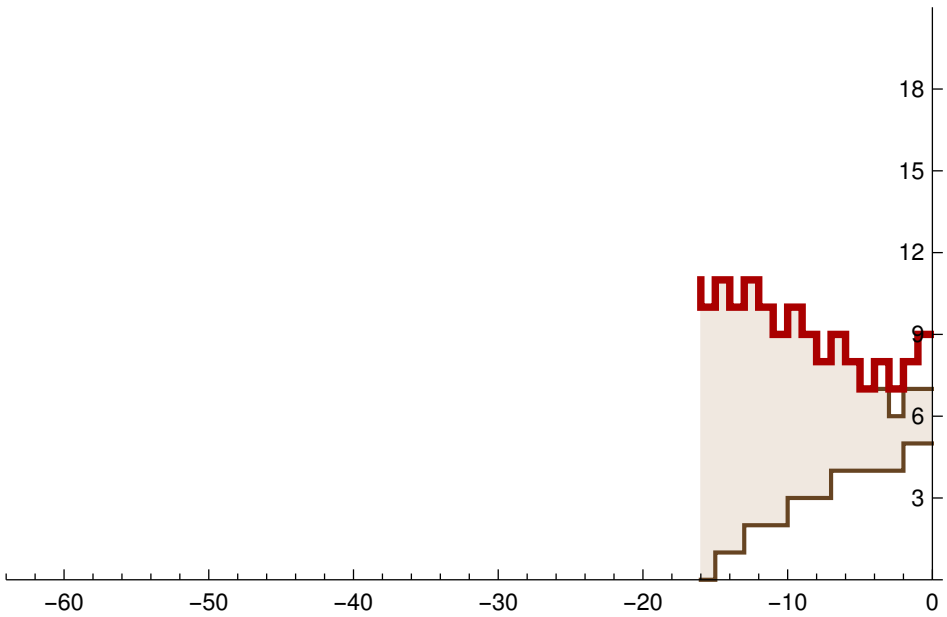

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time
- *sandwiching*

  $\text{Lower}_{\text{late}} \preccurlyeq \text{Lower}_{\text{early}} \preccurlyeq \ldots \preccurlyeq \text{Target} \preccurlyeq \ldots \preccurlyeq \text{Upper}_{\text{early}} \preccurlyeq \text{Upper}_{\text{late}}$

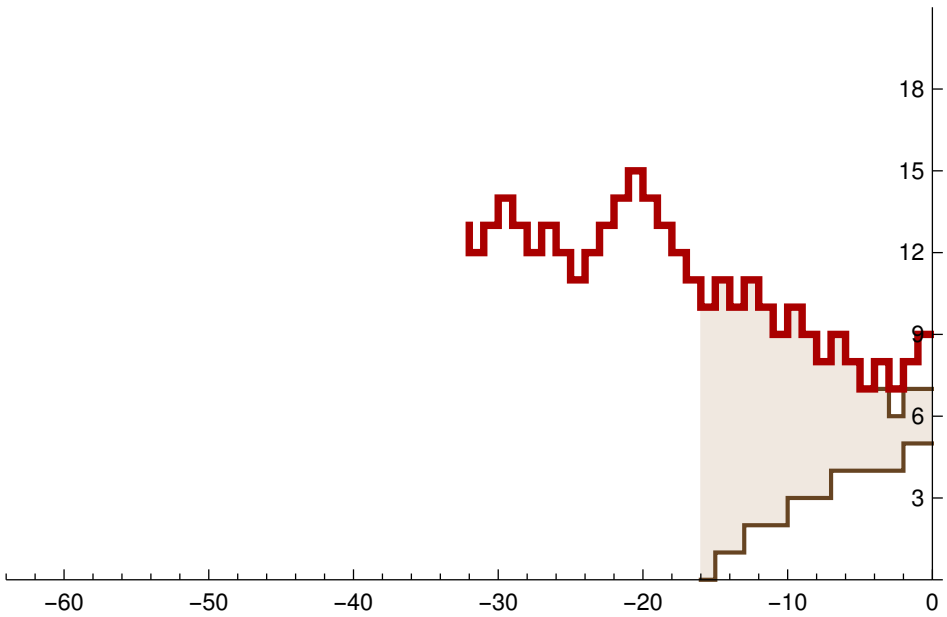

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time
- *sandwiching*

$$\text{Lower}_{\text{late}} \preccurlyeq \text{Lower}_{\text{early}} \preccurlyeq \ldots \preccurlyeq \text{Target} \preccurlyeq \ldots \preccurlyeq \text{Upper}_{\text{early}} \preccurlyeq \text{Upper}_{\text{late}}$$

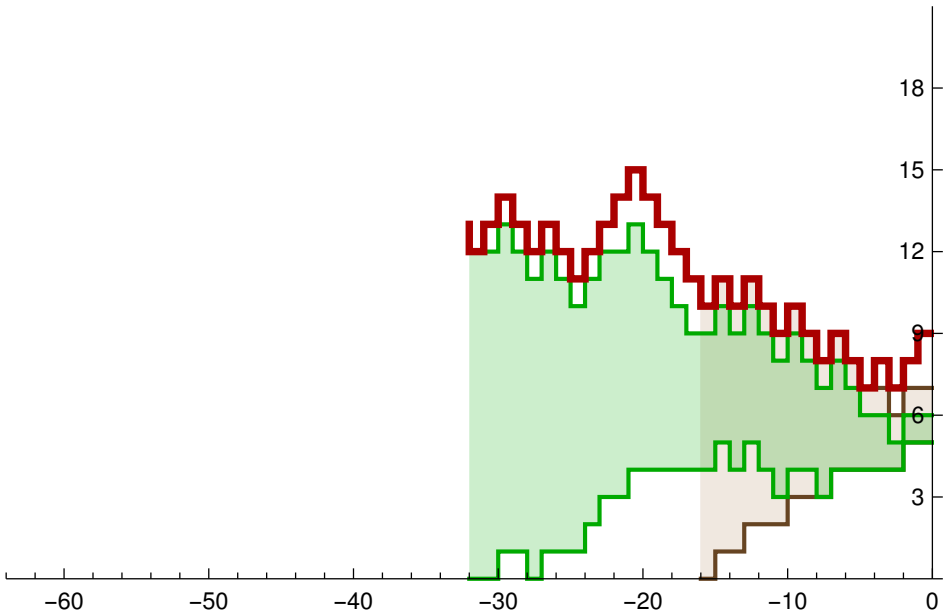

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time
- *sandwiching*

  $\text{Lower}_{\text{late}} \preccurlyeq \text{Lower}_{\text{early}} \preccurlyeq \ldots \preccurlyeq \text{Target} \preccurlyeq \ldots \preccurlyeq \text{Upper}_{\text{early}} \preccurlyeq \text{Upper}_{\text{late}}$

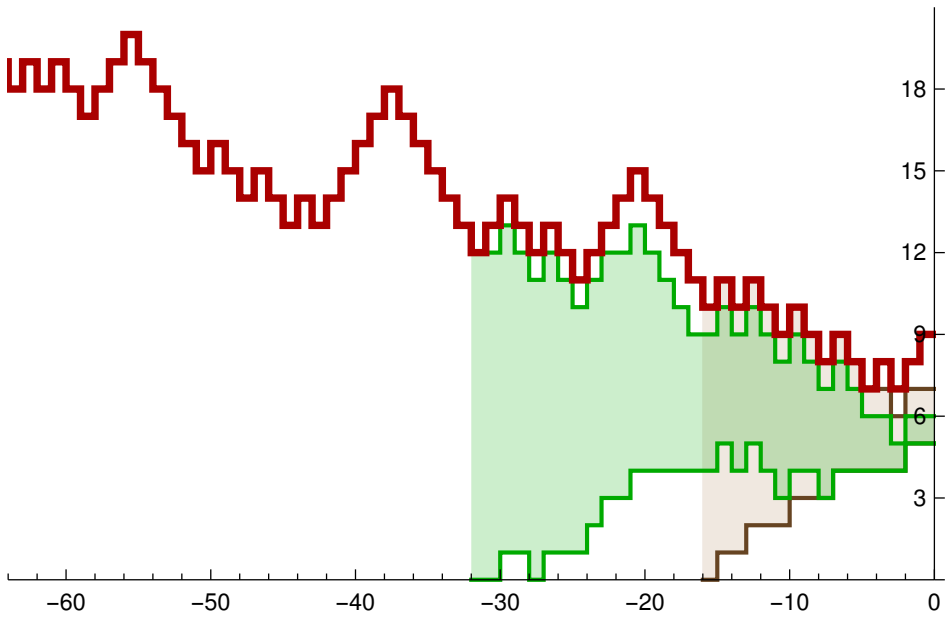

domCFTP: basic ingredients

- *dominating process Y*
  - draw from equilibrium π_Y
  - simulate backwards in time
- *sandwiching*

  $\text{Lower}_{\text{late}} \preccurlyeq \text{Lower}_{\text{early}} \preccurlyeq \ldots \preccurlyeq \text{Target} \preccurlyeq \ldots \preccurlyeq \text{Upper}_{\text{early}} \preccurlyeq \text{Upper}_{\text{late}}$
- *coalescence*

  eventually a Lower and an Upper process must coalesce

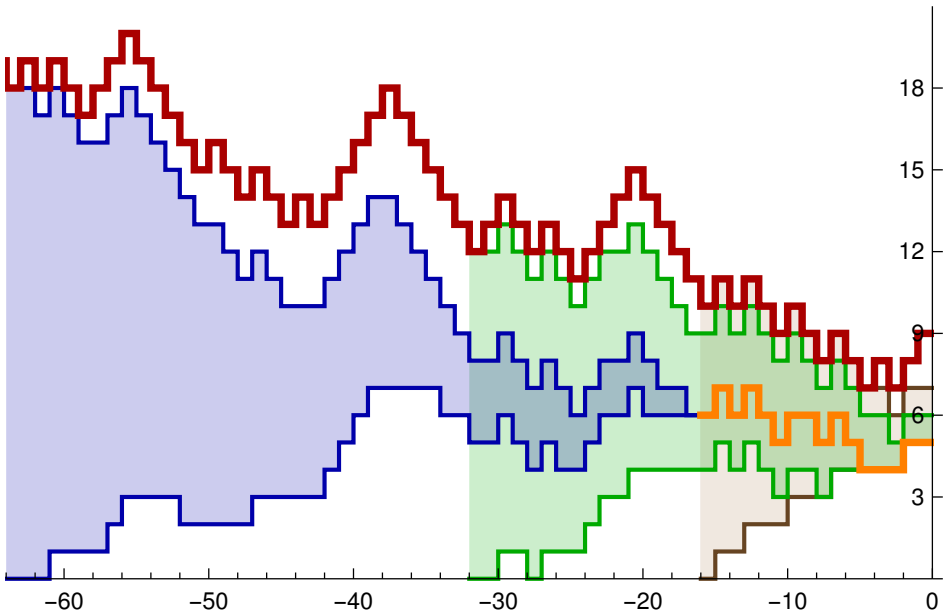

Omnithermal simulation

Suppose that the target process X has a distribution π_β that depends on some underlying parameter β.

In some situations it is possible to modify a perfect simulation algorithm so as to sample *simultaneously* from π_β for all β in some given range: call this **omnithermal simulation**.

Clearly desirable to be able to do this, particularly if it requires minimal additional computational overhead.

Let's look at some examples...

Random Cluster Model

First example of omnithermal CFTP (Propp & Wilson, 1996).
States are subsets of edges of undirected graph G, with

$$\pi_p(H) \propto \left(\prod_{e \in H} p\right) \left(\prod_{e \notin H} (1-p)\right) 2^{C(H)}, \quad H \subseteq G\,.$$

- $p \in [0,1]$
- $C(H) =$ number of connected components of H

Random Cluster Model

First example of omnithermal CFTP (Propp & Wilson, 1996).
States are subsets of edges of undirected graph G, with

$$\pi_p(H) \propto \left(\prod_{e \in H} p\right)\left(\prod_{e \notin H}(1-p)\right) 2^{C(H)}, \quad H \subseteq G\,.$$

- $p \in [0,1]$
- $C(H)$ = number of connected components of H

Assigning commom random spin (± 1) to connected vertices gives random (attractive) Ising model state

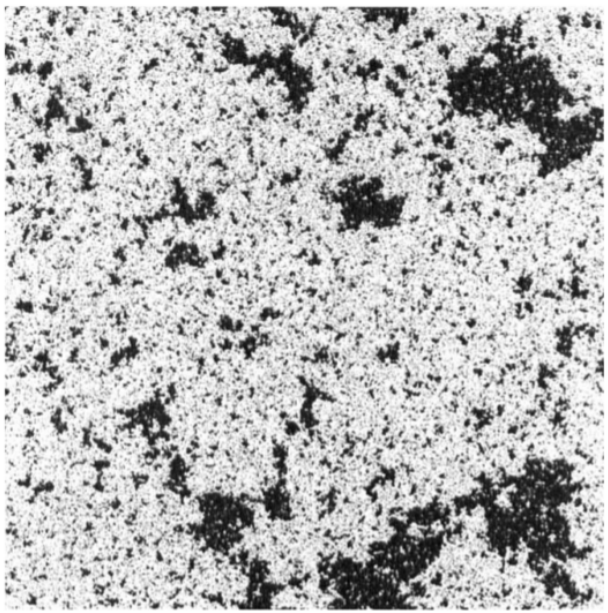

Single-bond heat-bath (Glauber dynamics) is monotone w.r.t. subgraph inclusion: allows for sampling via (monotone) CFTP. (Top state $= G$, bottom $=$ empty graph.)

Heat-bath dynamics also monotone w.r.t. parameter p (linked to *temperature* in Ising model) hence **omnithermal version**:

- record *set of values* $p(e)$ for which edge e belongs to H
- monotonicity ensures that

$$p \in p(e) \text{ and } p' \geq p \implies p' \in p(e)$$

- added complexity: determining limit of each interval $p(e)$.

Area Interaction Process (or Widom-Rowlinson Process)

Point process in a compact region of $\mathbb{R}^2$. Density w.r.t. unit rate Poisson process given by

$$\pi_\beta(\mathbf{x}) \propto \lambda^{n(\mathbf{x})} e^{-\beta m(U_r(\mathbf{x}))}$$

- $\lambda > 0$
- $n(\mathbf{x}) =$ number of points in $\mathbf{x}$
- $m =$ Lebesgue measure on $\mathbb{R}^2$
- $U_r(\mathbf{x}) =$ union of disks of radius r centred at points of $\mathbf{x}$

$\beta \in \mathbb{R}$ controls area-interaction:
$\beta > 0$ is *attractive* case

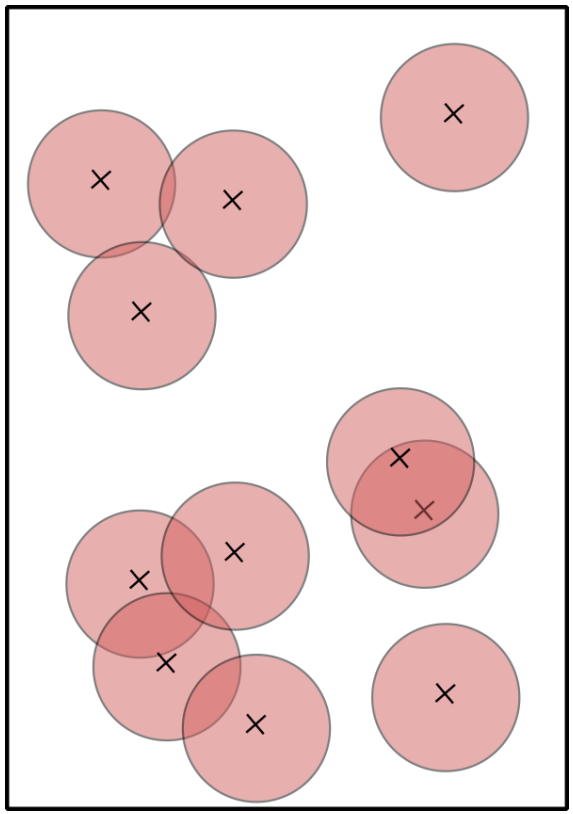

π_β can be viewed as equilibrium distribution of a **spatial birth-death process** Ψ_β: new points are born at a rate depending upon the current configuration, and die after an Exp(1) lifetime.

π_β can be viewed as equilibrium distribution of a **spatial birth-death process** Ψ_β: new points are born at a rate depending upon the current configuration, and die after an Exp(1) lifetime.

Implement by simulating a **free process** Φ, and censoring births accordingly: $\Psi_\beta(t) \subseteq \Phi(t)$

- Equilibrium of Φ is just a Poisson PP of rate λ
- Φ simple to run in reverse-time
- Censoring of births is monotonic w.r.t. set inclusion, so sandwiching holds

So we have all the necessary ingredients for domCFTP. (Kendall, 1998)

Going Omnithermal

Censoring of births is also monotonic in β:

$$\beta < \beta' \text{ and } \Psi_\beta(0) = \Psi_{\beta'}(0) \implies \Psi_\beta(t) \supseteq \Psi_{\beta'}(t) \quad \text{for all } t \geq 0$$

Going Omnithermal

Censoring of births is also monotonic in β:

$$\beta < \beta' \text{ and } \Psi_\beta(0) = \Psi_{\beta'}(0) \implies \Psi_\beta(t) \supseteq \Psi_{\beta'}(t) \quad \text{for all } t \geq 0$$

- given birth of a point ξ in free process Φ, record *set of values* $\beta(\xi)$ for which the birth is accepted in all target processes Ψ_β with $\beta \leq \beta(\xi)$, and rejected otherwise
- careful construction yields set of points of form $(\xi, \beta(\xi))$, which can be thresholded to obtain a draw from π_β
- added complexity: determining values $\beta(\xi)$

See (Shah, 2004) for details.

$M/G/c$ Queue

- Customers arrive at times of a Poisson process: interarrival times $T_n \sim \text{Exp}(\lambda)$
- Service durations S_n are i.i.d. with $\mathbb{E}[S] = 1/\mu$ (and we assume that $\mathbb{E}\left[S^2\right] < \infty$)
- Customers are served by c servers, on a First Come First Served (FCFS) basis

Queue is *stable* iff $\rho := \dfrac{\lambda}{\mu c} < 1$.

Interested in equilibrium distribution of (ordered) workload vector.

DomCFTP Algorithm (C. & Kendall, 2015)

- Dominating process Y is stationary $M/G/c$ **[RA]** queue
- Check for coalescence of **sandwiching processes**, U^c and L^c:
  - these are workload vectors of $M/G/c$ $[FCFS]$ queues
  - L^c starts from empty
  - U^c is instantiated using **residual workloads** from Y

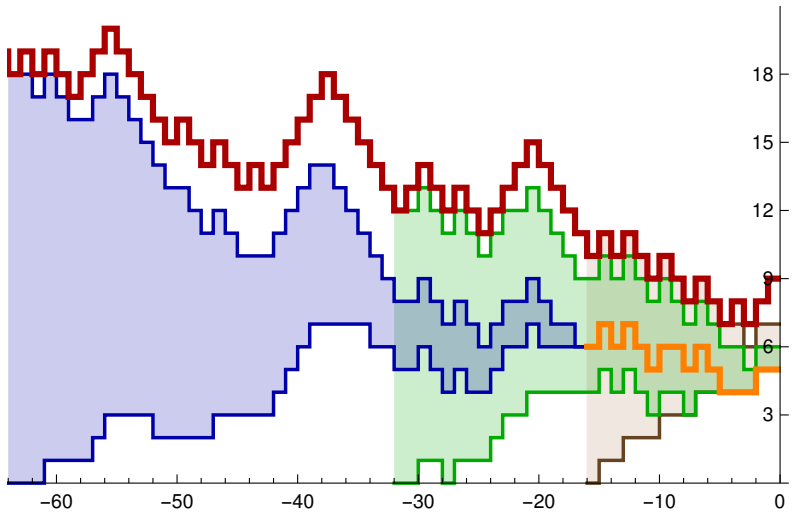

Going Omnithermal

Dynamics for workload vectors with **different numbers of servers** are monotonic w.r.t. a natural partial order:

for $V^c \in \mathbb{R}^c$ and $V^{c+m} \in \mathbb{R}^{c+m}$, write $V^{c+m} \preceq V^c$ if and only if

$$V^{c+m}(k+m) \leq V^c(k)\,, \quad k = 1, \ldots, c\,.$$

("Busiest c servers in V^{c+m} each no busier than corresponding server in V^c".)

So we can produce processes U^{c+m} and L^{c+m} over $[T, 0]$, coupled to our c-server dominating process Y, such that:

- U^{c+m} and L^{c+m} sandwich our $M/G/(c+m)$ FCFS process of interest
- $U_t^{c+m} \preceq U_t^c$ and $L_t^{c+m} \preceq L_t^c$

So we can produce processes U^{c+m} and L^{c+m} over $[T, 0]$, coupled to our c-server dominating process Y, such that:

- U^{c+m} and L^{c+m} sandwich our $M/G/(c+m)$ FCFS process of interest
- $U_t^{c+m} \preceq U_t^c$ and $L_t^{c+m} \preceq L_t^c$

But U^{c+m} and L^{c+m} **won't** necessarily coalesce before time 0!

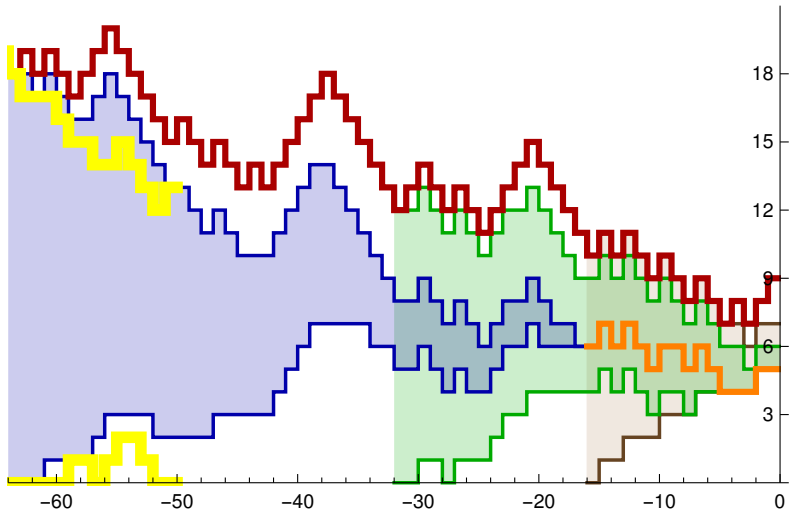

Write $T^c \leq 0$ for the coalescence time of U^c and L^c.

Condition A

$\forall$ arrival times $\tau \in [T, T^c]$, if $L^c_{\tau-}(1) = U^c_{\tau-}(1)$ then $U^c_{\tau-}(1) = 0$

Write $T^c \leq 0$ for the coalescence time of U^c and L^c.

Condition A

$\forall$ arrival times $\tau \in [T, T^c]$, if $L^c_{\tau-}(1) = U^c_{\tau-}(1)$ then $U^c_{\tau-}(1) = 0$

Theorem (C., 2016)

If Condition A holds then $T^{c+m} \leq T^c$ for any $m \in \mathbb{N}$.

Write $T^c \leq 0$ for the coalescence time of U^c and L^c.

Condition A

$\forall$ arrival times $\tau \in [T, T^c]$, if $L^c_{\tau-}(1) = U^c_{\tau-}(1)$ then $U^c_{\tau-}(1) = 0$

Theorem (C., 2016)

If Condition A holds then $T^{c+m} \leq T^c$ for any $m \in \mathbb{N}$.

This gives us a method for performing omnithermal domCFTP:

1. for a given run of the c-server domCFTP algorithm, check to see whether Condition A holds. If not, repeatedly backoff ($T \leftarrow 2T$) until Condition A is satisfied;
2. run L^{c+m} (for any $m \in \mathbb{N}$) over $[T, 0]$, and return $L^{c+m}(0)$.

Example output

Simulation results from 5,000 runs for $M/M/c$ with $\lambda = 2.85$, $\mu = 1$ and $c = 3$ ($\rho = 0.95$)

- 333 (7%) runs needed extending
- only 2 runs needed more than 2 additional backoffs

Mean workload at each server, for $c = 3$ and $m \in \{0, 1, 2, 3\}$:

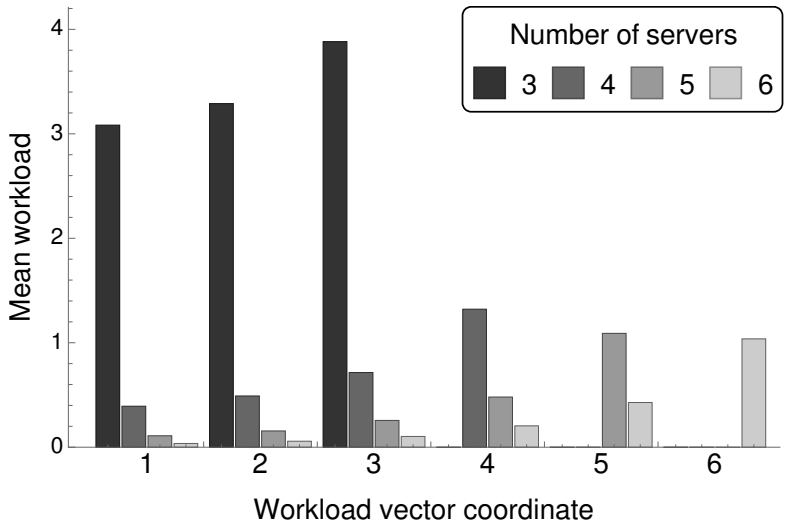

Example output

Distribution functions for workload at (a) first and (b) last coordinates of the workload vector:

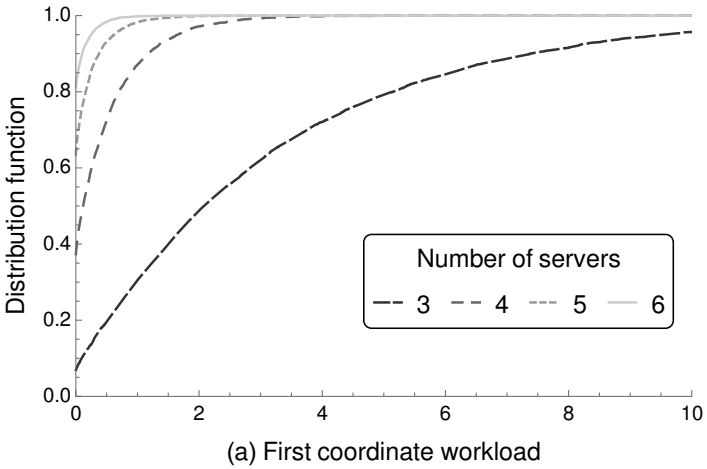

(a) First coordinate workload

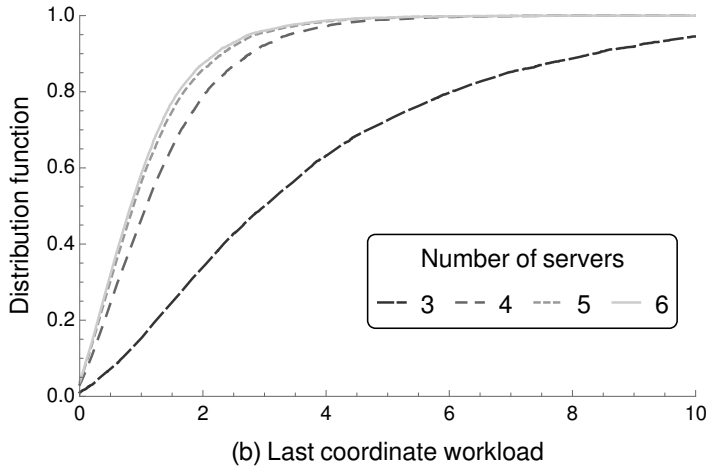

(b) Last coordinate workload

How expensive is this in practice?

Not very!

- Simulations indicate that Condition A is satisfied (with no need for further backoffs) $> 90\%$ of the time when $\rho \leq 0.75$, and in $> 70\%$ of cases when $\rho = 0.85$
- In addition, runs in which Condition A initially fails typically don't require significant extension
- Theoretical analysis of run-time would be nice, but hard!

Conclusions

Perfect simulation algorithms are of practical use (and theoretical interest!) in a growing number of applications.

Some of these algorithms can be modified to provide omnithermal samples, often with relatively little additional computational effort

There are plenty of other perfect simulation algorithms out there, e.g. gradient simulation for fork-join networks (Chen & Shi, 2016), for which it may be possible to "go omnithermal"!

References

- Chen, X & Shi, X. (2016). Perfect Sampling and Gradient Simulation for Fork-Join Networks. [arXiv]
- Connor, S.B. (2016). Omnithermal perfect simulation for multi-server queues. *Submitted.*
- Connor, S.B., & Kendall, W.S. (2015). Perfect simulation of $M/G/c$ queues. *Advances in Applied Probability*, **47**(4), 1039–1063.
- Kendall, W.S. (1998). Perfect simulation for the area-interaction point process. In *Probability Towards 2000* eds L. Accardi and C. C. Heyde. Springer, New York, pp. 218–234.
- Propp, J.G., & Wilson, D.B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures and Algorithms*, **9**, 223–252.
- Shah, S.R. (2004). *Perfect simulation of conditional & weighted models*. Ph.D. thesis, Department of Statistics, University of Warwick.