# Two Large-scale Particle Systems With Mean-field Interaction

Alexander Stolyar (UIUC)
In part, joint work with Seva Shneer (Heriot Watt Univ.)

June 8, 2021. YEQT-XIV

# Underlying particle system

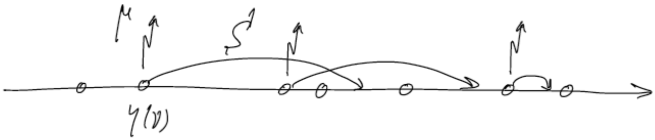

- $n$ particles moving "right" on real line
- Each gets "urges" to jump forward as Poisson process, rate $\mu$. Actually jumps with probability $\eta(\nu)$, $\nu =$ particle location quantile; $\eta(\nu)$, $0 \leq \nu \leq 1$, continuous strictly decreasing from 1 to 0
- $S \geq 0$ random jump size; CDF $J(\cdot)$; $\bar{J}(\cdot) = 1 - J(\cdot)$; $\mathbb{E}S^2 < \infty$
- WLOG: $\mu = 1, \mathbb{E}S = 1$

# Motivation. Key questions

- Parallel simulation
- Blockchain evolution
- Other synchronization models

$f^n(x,t)$ = fraction of particles located at $\leq x$ at time $t$

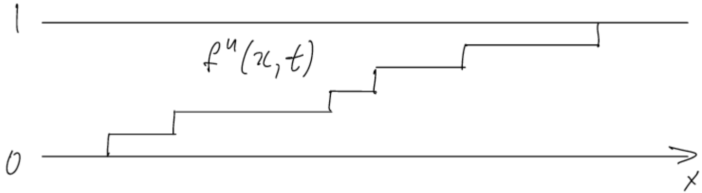


When $n$ becomes large,

- Does (recentered) $f^n(\cdot, t)$ remains "tight"?
- Does it converge to a traveling wave, as $t \to \infty$?

# Mean field limit

$\lim_{n\to\infty} f^n(x,t) = f(x,t) =$ Mean-field limit, which satisfies:

$$\frac{\partial}{\partial t} f(x,t) = -\int_{-\infty}^{x} d_y f(y,t)\eta(f(y,t))\bar{J}(x-y) \tag{1}$$

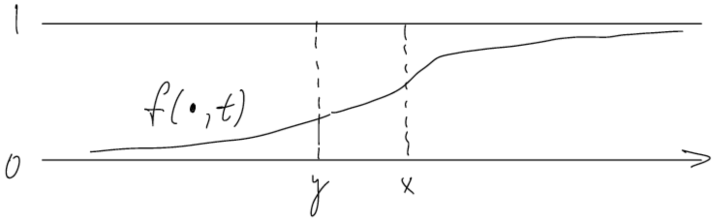

# Mean field model

Mean-field model (MFM) $=$ Any solution to (1)
Mean $\bar{f}(t) = \int_{-\infty}^{\infty} x df_x(x,t)$.
If $\bar{f}(t)$ well-defined finite, it has constant speed $v$:

$$\frac{d}{dt}\bar{f}(t) = \int_0^1 \eta(\nu) d\nu = v$$

# Traveling wave

Traveling wave: $f(x,t) = \phi(x - vt)$

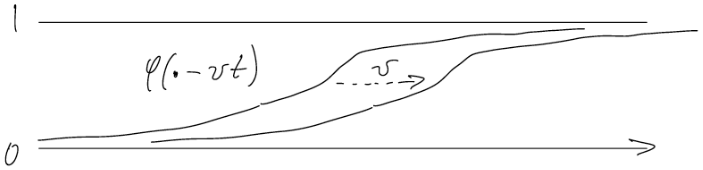


$$v\phi'(x) = \int_{-\infty}^{x} \phi'(y)\eta(\phi(y))\bar{J}(x-y)dy \quad (2)$$

Solution $\phi(\cdot)$ to (2) is a traveling wave shape (TWS).
A TWS is defined up to a shift

# Closely related previous work

- A. Greenberg, V. Malyshev, and S. Popov, Stochastic model of massively parallel computation. Markov Processes and Related Fields, 2:473-490, 1997.
  - Proves (under some conditions) convergence to MFL/MFM: $f^n(\cdot, t) \to f(\cdot, t)$ as $n \to \infty$.
- A. Greenberg, S. Shenker, and A. S., Asynchronous updates in large parallel systems. SIGMETRICS-1996, 91-103.
  - Assuming $J(\cdot)$ has positive density: for any MFM $f_1(\cdot.t)$ and $f_2(\cdot.t)$ with $\bar{f}_1(0) = \bar{f}_2(0)$, $\|f_1(\cdot.t) - f_2(\cdot.t)\|_{L_1} \to 0$ as $t \to \infty$.
  - As a corollary, if TWS $\phi(\cdot)$ exists, then it is unique (up to a shift) and any MFM $f(\cdot, t)$ with $\bar{f}(0) = \bar{\phi}$ converges to it, $f(\cdot + vt, x) \to \phi(\cdot)$.
  - Existence of TWS is shown only for $J(x) = 1 - e^{-x}$ and $\eta(\nu) = (1 - \nu)^K$.

# Our main results

**Theorem (TWS existence)**

*If $\mathbb{E}S^2 < \infty$, a TWS $\phi(\cdot)$ exists for any $J(\cdot)$ and any $\eta(\cdot)$.*
*If $\mathbb{E}S^\ell < \infty$ for some $\ell \geq 2$, $\int |x|^{\ell-1} d\phi(x) < \infty$.*

**Corollary (From the Theorem and SIGMERTICS-96 result)**

*If, in addition, $J(\cdot)$ has positive density, then the TWS $\phi(\cdot)$ is unique (up to a shift) and any MFM $f(\cdot, t)$ with $\bar{f}(0) = \bar{\phi}$ converges to it, $f(\cdot + vt, t) \to \phi(\cdot)$.*

Remarks:
1. The positive density assumption likely can be much relaxed.
2. Our results are not probabilistic – they are about properties of MFM. The methods are are almost exclusively probabilistic.

# Some other previous work on 'synchronization' systems


- A. Manita and V. Shcherbakov. Asymptotic analysis of a particle system with mean-field interaction. Markov Process. Relat. Fields, 11:489-518, 2005.
- V. Malyshev and A. Manita. Phase transitions in the time synchronization model. Theory Probab. Appl., 50:134-141, 2006.
- Marton Balazs, Miklos Racz and Balint Toth. Modeling Flocks and Prices: Jumping Particles with an Attractive Interaction. Annales de l'Institut Henri Poincare - Probabilites et Statistiques, 50(2): 425-454, 2014.

# Characterization of a TWS

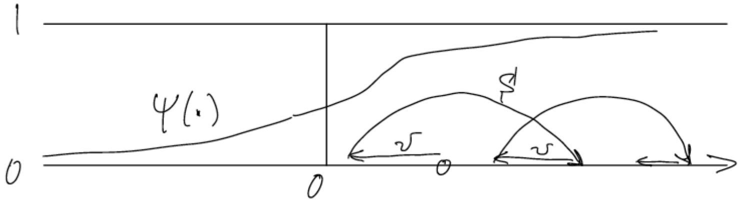


- $\psi$ = distribution ("environment"). Determines jump probabilities
- $v$ = speed. Particle moves left at constant speed $v$ when not jumping
- $A\psi$ = stationary distribution of a single particle ("at cavity"), given environment $\psi$
- $\phi$ is a TWS if and only if $\phi = A\phi$

# Finite-frame TWS

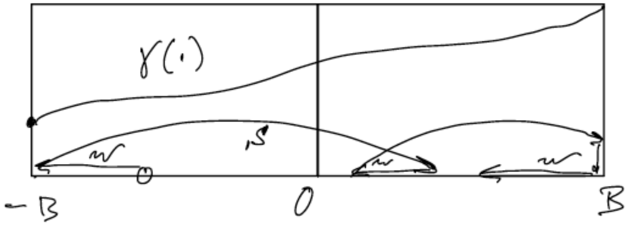


- $\gamma$ = distribution ("environment") on "finite frame" $[-B, B]$. Determines jump probabilities
- $-B$ and $B$ are regulating boundaries
- $w$ = speed. Particle moves left at constant speed $w$ when away from boundaries and not jumping
- $A^{w,B}\gamma$ = stationary distribution of a single particle ("at cavity"), given environment $\gamma$

# Finite-frame TWS

- Fixed point $\gamma^{w,B}$ of $A^{w,B}$ we call a TWS within finite frame $[-B, B]$
- Finite-frame TWS $\gamma^{w,B}$ exists and is unique
- Existence: Brower fixed point
- Uniqueness: Equation for a fixed point $\Rightarrow$ "Horizontal distance" $\|\gamma_1^{-1} - \gamma_2^{-1}\|$ between two fixed point cannot be positive

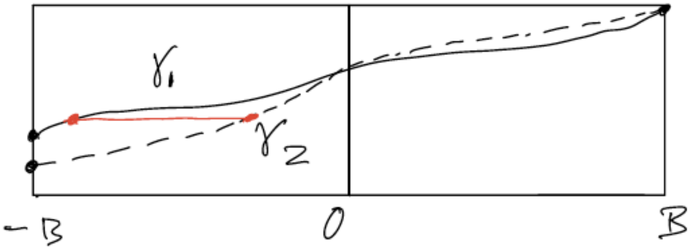

# TWS is a limit of finite-frame TWS, as $B \to \infty$

- $B \to \infty$
- For each $B$, there is unique speed $w_B$, such that the median of $\gamma^{w_B,B}$ is at 0
- Show that $w_B \to v$
- Show that $\{\gamma^{w_B,B}\}$ is tight
  - Use rescaling and quadratic Lyapunov function
- Show that any weak limit $\phi$ of $\gamma^{w_B,B}$ is a TWS

# $w_B \to v$

- Suppose, say, $w_B \to w_\infty > v$
- Show $\gamma^{w_B,B}(\cdot - B) \to \gamma^*(\cdot)$ with $\gamma^*(0) > 0$
- Show $\gamma^*$ is proper $\Rightarrow$ Contradiction with the median of $\gamma^{w_B,B}$ staying at 0

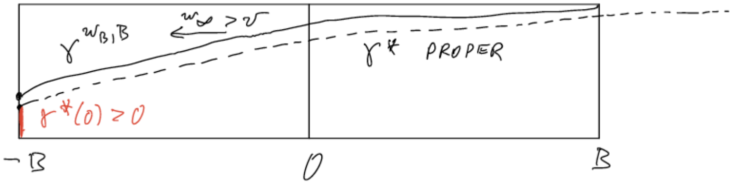

# $\{\gamma^{w_B,B}\}$ is tight

- Suppose, $C = C(B) = \max\{-q_\delta(\gamma^{w_B,B}), q_{1-\delta}(\gamma^{w_B,B})\} \to \infty$; $q_\nu$ - quantile
- Consider rescaled distribution and corresponding (cavity) particle movement process:

$$\gamma_{(C)}(y) = \gamma^{w_B,B}(Cy), \quad X_{(C)}(t) = (1/C)X_B(Ct).$$

- Show tightness of $\{\gamma_{(C)}\}$, using quadratic Lyapunov function

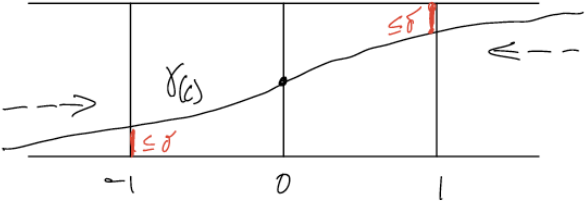

# $\{\gamma^{w_B,B}\}$ is tight

- Consider limit $\gamma_{(C)} \to \gamma_{(\infty)}$
- Process $X_{(C)}(t)$ is non-random in the limit $\Rightarrow \gamma_{(\infty)}$ is degenerate with one or two atoms
- Obtain contradiction – negative steady-state drift – using quadratic Lyapunov function

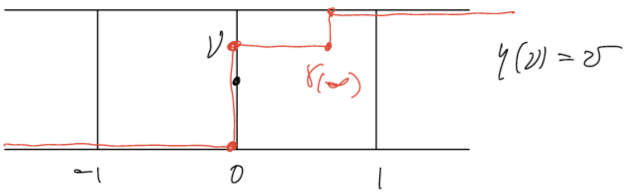

# At a high level ...

Our approach relies, almost exclusively, on three fundamental properties of the system:

- Monotonicity
- "Average jump size is non-increasing in a particle location"
- Work conservation

# Simple – but useful – generalization

- There is additional type of jumps:
  - "Urge" rate is $\mu_2$; jump occurs with prob. 1.
  - Jump size distribution: $J_2$
- Same results hold, under analogous assumptions

# Ongoing work

- Limit interchange. $f_0^n(\cdot,t)$ is $f^n(\cdot,t)$ recentered so that median at 0. $f_0^n(\cdot,\infty)$ is $f_0^n(\cdot,t)$ in steady-state.

$$\|f_0^n(\cdot,\infty)-\phi(\cdot)\| \xrightarrow{P} 0$$

# Parallel server system with multi-component jobs

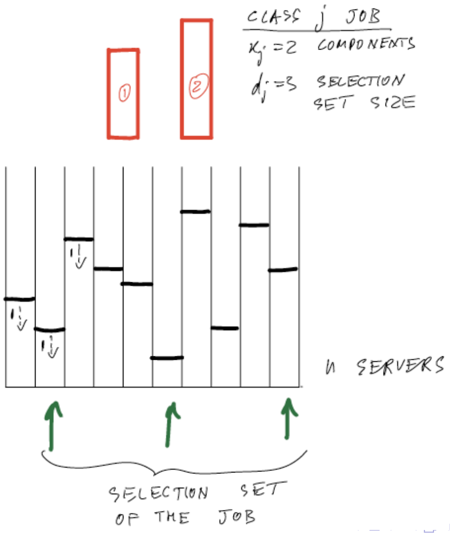

# Parallel server system with multi-component jobs

- $n$ servers processing work at unit rate
- Multiple job classes
- Class $j$ jobs arrive as Poisson process, rate $\lambda_j n$
- Class $j$ job: $k_j$ components; $d_j \geq k_j$ selection set size; component sizes $(\xi_1^{(j)}, \ldots, \xi_{k_j}^{(j)})$ drawn from an exchangeable distribution $F_j$ (for example, i.i.d.)
- A job class may be of one of two types: Least-load or Water-filling. Type determines the rule by which job adds workload to selected servers.
- Subcritical load: $\rho = \sum_j \lambda_j k_j \mathbb{E}\xi_1^{(j)} < 1$.
  WLOG, $\lambda = \sum_j \lambda_j = \rho < 1$
- Asymptotic regime: $n \to \infty$. Main question:
  Steady-state asymptotic independence of server workloads?

# Parallel server system with multi-component jobs

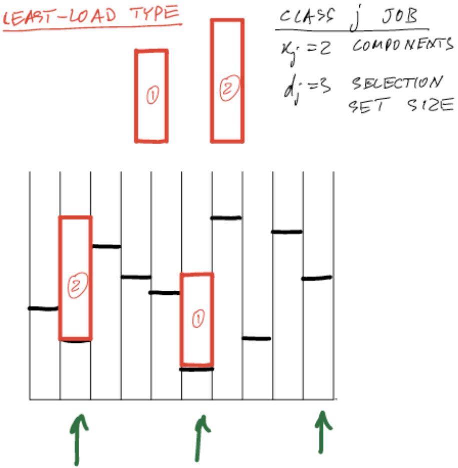


Covers cancel-on-start redundancy

# Parallel server system with multi-component jobs

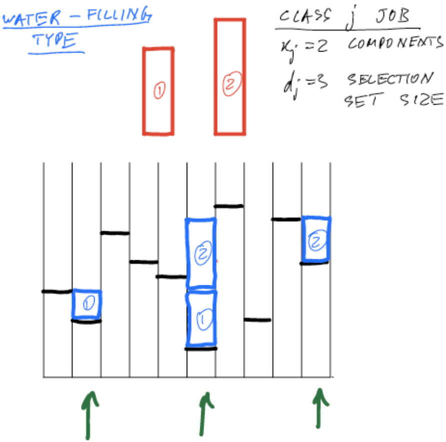


Covers cancel-on-completion redundancy, but only the special case of i.i.d. exponential components

# Particle system interpretation. Work conservation

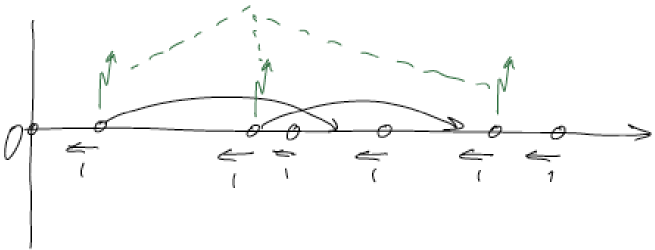

- Class $j$ job arrival:
  - $k_j$ selected "particles" get an "urge" to jump
  - Jump sizes are actual workloads added
- Work conservation: Average total of all jumps upon a class $j$ arrival is fixed

# Mean-field scaling

$x_w^n(t)$ = fraction of particles located at $> w$ at time $t$
$x^n(t) = (x_w^n(t),\ w \geq 0)$

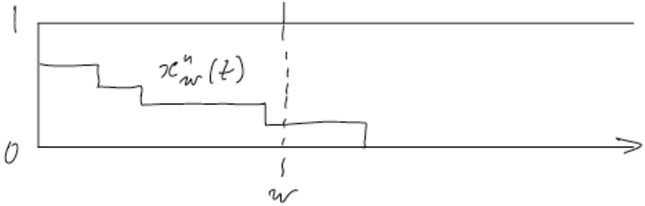


$x^n(t),\ t \geq 0,$ Markov process
$x^n(\infty)$ Process steady-state

# Main result

## Theorem (Steady-state asymptotic independence)

*There exists a fixed element $x^*$, with $x_0^* = \lambda$ and $\lim_{w\to\infty} x_w^* \to 0$, such that*

$$x^n(\infty) \Rightarrow x^*, \quad n \to \infty.$$

*Consequently, steady-state workloads of a finite set of servers are asymptotically independent.*

# Discussion

- There are no previous proofs of the steady-state asymptotic independence for such multi-component systems
  - Except the special case of single least-load class with $k = 1$. This is the LL(d) discipline in
    M. Bramson, Y. Lu, and B. Prabhakar. Asymptotic independence of queues under randomized load balancing. Queueing Systems, 71:247-292, 2012

# Discussion

- However, the conjecture of the steady-state asymptotic independence is often used for such systems, to derive useful performance estimates for large-scale systems. Our result formally proves the conjecture in some cases. Examples:
  - K. Gardner, M. Harchol-Balter, A. Scheller-Wolf, M. Velednitsky, and S. Zbarsky. Redundancy-d: The power of d choices for redundancy. Operations Research, 65(4):1078-1094, 2017. [Model within our framework.]
  - T. Hellemans, T. Bodas, and B. Van Houdt. Performance analysis of workload dependent load balancing policies. Proceedings of the ACM on Measurement and Analysis of Computing Systems, 3(2), 2019. [Some of the models within our framework.]

# Once again, at a high level ...

Our approach will rely, almost exclusively, on three fundamental properties of the system:

- Monotonicity
- "Average jump size is non-increasing in a particle location" = "On average, servers with lower workload receive more new workload"
- Work conservation

# Mean-field limit

$x(t) = \lim_{n\to\infty} x^n(t)$
$x(t),\ \ t \geq 0$, Mean-field limit process

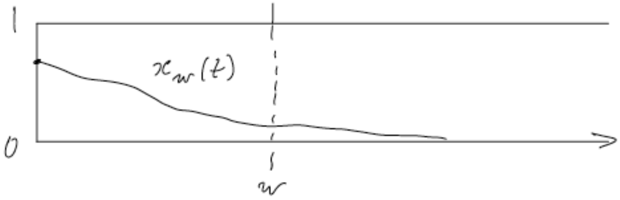

# Fixed point

$x^{\emptyset}(t) \uparrow x^*$
$x^*$ is the fixed point

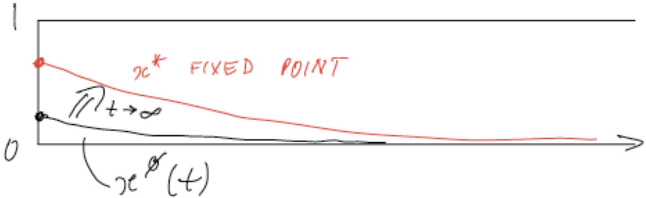


$x^*$ is proper $(x^*_\infty = 0)$
$x^*_0 \leq \lambda$ by work conservation

# $x_0^* = \lambda$

$x^{*,c}$ is the fixed point for the (truncated) system with frame $[0, c]$
$x^{*,c} \uparrow x^*$ as $c \uparrow \infty$

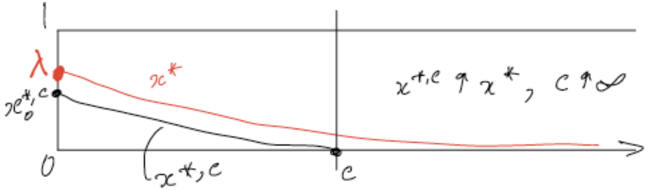


When $c$ is large, the amount of work (per job) lost due to truncation is small
$x_0^{*,c}$ is close to $\lambda$ by work conservation
$x_0^* = \lambda$

# Limit of stationary distributions

$x^n(\infty) \Rightarrow x^{**}$ = Any limit along a subsequence, a random element
Must have $x^{**} \geq x^*$

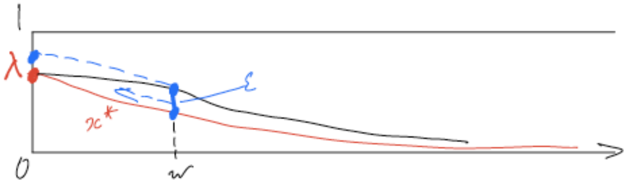


$\mathbb{P}\{x_w^{**} \geq x_w^* + \epsilon\} > 0$ is impossible; otherwise,
$\mathbb{E}x_0^{**} > \lambda$, a contradiction with work conservation
So, $x^{**} = x^*$

# Multi-component jobs. Cancel-on-completion

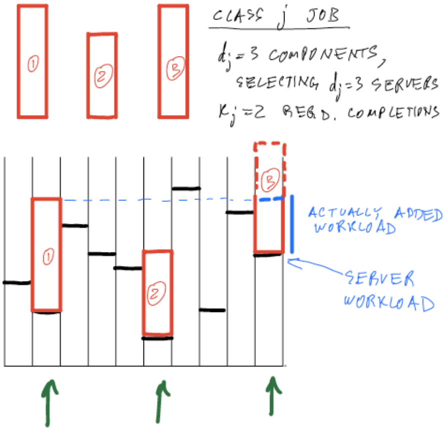


Non-work-conserving: Average added workload depends on the relative workloads of selected servers

# Multi-component jobs. Cancel-on-completion

- $n$ servers processing work at unit rate
- Multiple job classes
- Class $j$ jobs arrive as Poisson process, rate $\lambda\pi_j n$, $\sum \pi_j = 1$
- Class $j$ job: $d_j$ components; $d_j$ selection set size; component sizes $(\xi_1^{(j)}, \ldots, \xi_{d_j}^{(j)})$ drawn from an exchangeable distribution $F_j$ (for example, i.i.d.); $\mathbb{E}[\xi_1^{(j)}]^2 < \infty$, $\forall j$
- Cancel-on-completion: When $k_j(\leq d_j)$ components complete service, the remaining components' service is immediately canceled. Non-work-conserving
- Asymptotic regime: $n \to \infty$ Main question:
  Steady-state asymptotic independence of server workloads?

# The most intuitive behavior

Steady-state asymptotic independence for the full range of loads:

There exists $\bar{\lambda}$ and a continuous strictly increasing in $\lambda \in [0, \bar{\lambda})$ element $x^{*,\lambda}$, with load $x_0^{*,\lambda} \uparrow 1$ as $\lambda \uparrow \bar{\lambda}$, such that

$$x^n(\infty) \Rightarrow x^{*,\lambda}$$

# Some of the main results

FP-uniqueness condition (see the paper for exact condition):
For any $\lambda$, a fixed point $x^{*,\lambda}$ of the mean-field limit is unique (if exists).

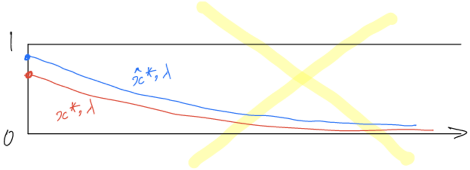


**Theorem (General sufficient condition)**

*Assuming the FP-uniqueness, the steady-state asymptotic independence for the full range of loads holds.*

# Some of the main results

## Theorem (I.i.d. IHR case)

*The FP-uniqueness (and then the steady-state asymptotic independence for the full range of loads) holds if for each class $j$ the component sizes are i.i.d. with increasing hazard rate.*
*(Or, more generally, the component size distribution for each class can be a mixture of those.)*

Covers, in particular, random identical components

# Papers

- A. S., Large-scale behavior of a particle system with mean-field interaction: Traveling wave solutions. arXiv:2004.00177
- S. Shneer, A. S., Large-scale parallel server system with multi-component jobs. Queueing Systems, 98: 21-48, 2021. arXiv:2006.11256
- A. S., Parallel server systems with cancel-on-completion redundancy. arXiv:2105.14143