Networks — an OR perspective

Michel Mandjes[1,2,3]

[1]Korteweg-de Vries Institute for Mathematics,
University of Amsterdam

[2]CWI, Amsterdam

[3]Eurandom, Eindhoven

Back-to-School day

January 14th 2016

Lunteren, the Netherlands

TODAY

This *Back-to-School day* is on *networks*.

TODAY

This *Back-to-School day* is on *networks*.

- Networks arguably form the main concept in operations research,

TODAY

This *Back-to-School day* is on *networks*.

- Networks arguably form the main concept in operations research,
- with the ultimate goal to devise procedures for optimal design (long time scale) and operations (short time scale).

OPERATIONS RESEARCH & NETWORKS

Virtually any networks-related problem has *combinatorial/deterministic* and *stochastic* elements.

OPERATIONS RESEARCH & NETWORKS

Virtually any networks-related problem has *combinatorial/deterministic* and *stochastic* elements.

Often:

OPERATIONS RESEARCH & NETWORKS

Virtually any networks-related problem has *combinatorial/deterministic* and *stochastic* elements.

Often:

- Combinatorial/deterministic techniques needed to design network such that demand is met on the longer term;
- Stochastic techniques needed to assess whether random (short term) fluctuations are adequately dealt with.

OPERATIONS RESEARCH & NETWORKS

. . . but in reality things are more subtle:

OPERATIONS RESEARCH & NETWORKS

. . . but in reality things are more subtle:

- setting up algorithms to deal with short-term fluctuations often requires combinatorial techniques (think of scheduling),

OPERATIONS RESEARCH & NETWORKS

... but in reality things are more subtle:

- setting up algorithms to deal with short-term fluctuations often requires combinatorial techniques (think of scheduling),
- and also in the design phase uncertainty is typically already included.

THE POWER OF MATH

At a ‘stylized’ level, OR has developed a set of generic models and concepts that fit in many situations — think of

THE POWER OF MATH

At a ‘stylized’ level, OR has developed a set of generic models and concepts that fit in many situations — think of

- various sorts of standard queueing models (e.g., GI/G/c, fluid queues),

THE POWER OF MATH

At a ‘stylized’ level, OR has developed a set of generic models and concepts that fit in many situations — think of

- various sorts of standard queueing models (e.g., GI/G/c, fluid queues),
- various sorts of standard combinatorial structures (graphs, matchings, travelling salesman),

THE POWER OF MATH

At a ‘stylized’ level, OR has developed a set of generic models and concepts that fit in many situations — think of

- various sorts of standard queueing models (e.g., GI/G/c, fluid queues),
- various sorts of standard combinatorial structures (graphs, matchings, travelling salesman),
- various sorts of scheduling models,

etc.

THE POWER OF MATH

but ...

THE POWER OF MATH

but ...

- stylized models are limited: new (variants of existing) models keep on being studied, called for by new developments in the application domains.

THE POWER OF MATH

but . . .

- stylized models are limited: new (variants of existing) models keep on being studied, called for by new developments in the application domains.
- size: networks become larger (increasingly connected world, with more data available). Requires structural understanding of such large networks, and powerful computational techniques.

THE POWER OF MATH

but ...

- stylized models are limited: new (variants of existing) models keep on being studied, called for by new developments in the application domains.
- size: networks become larger (increasingly connected world, with more data available). Requires structural understanding of such large networks, and powerful computational techniques.
- not stand-alone: increasing awareness that stochastic and deterministic issues should not be dealt with separately. Along the same lines: link with theoretical computer science, data science, statistics (think of simultaneous estimation and optimization), etc.

Networks

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks.

Networks

NETWORKS

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks. Runs from 2014 till 2024.

Networks

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks.
Runs from 2014 till 2024.
Focuses on understanding and optimizing large complex networks.

NETWORKS

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks.
Runs from 2014 till 2024.
Focuses on understanding and optimizing large complex networks.
Math/cs teams Univ. of Amsterdam, CWI, Eindhoven, Leiden.

NETWORKS

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks. Runs from 2014 till 2024.

Focuses on understanding and optimizing large complex networks.

Math/cs teams Univ. of Amsterdam, CWI, Eindhoven, Leiden.

Big volume (7 TTers, $\pm$ 50 PhD students, $\pm$ 20 postdocs, $\pm$ 20 local staff involved).

Networks

December 2013: large grant for fundamental research was awarded for research on (stochastic and combinatorial) analysis of networks.
Runs from 2014 till 2024.
Focuses on understanding and optimizing large complex networks.
Math/cs teams Univ. of Amsterdam, CWI, Eindhoven, Leiden.
Big volume (7 TTers, $\pm$ 50 PhD students, $\pm$ 20 postdocs, $\pm$ 20 local staff involved).
Surrounded by practical domain-specific projects.

Have a look at http://www.thenetworkcenter.nl

TODAY

Content of today very much in line with mentioned trends. Six talks with new developments, seen from theoretical as well as practical angle.

10.40 - 11.20 The power of social network analysis – *Ana Barros* (TNO).

11.35 - 12.15 The hub-network of KLM: importance and succesfactors – *Pieter Cornelisse* (KLM).

13.30 - 14.10 Design and operational challenges of communication networks – *Richa Malhotra* (SURFnet).

14.10 - 14.50 Supply Network Analytics - Operations Research in the Supply Chain – *Jan van Doremalen* (CQM).

15.05 - 15.45 Design and analysis of container liner shipping networks – *Rommert Dekker* (Erasmus University Rotterdam).

15.45 - 16.25 Towards data-driven models for the mobility system – *Maaike Snelders* (TNO and TU Delft).

A STOCHASTIC NETWORK...

At an abstract level a stochastic network can be seen as

A STOCHASTIC NETWORK...

At an abstract level a stochastic network can be seen as

- a set of *connected resources*,

A STOCHASTIC NETWORK...

At an abstract level a stochastic network can be seen as

- a set of *connected resources*,
- used by customers, imposing a randomly fluctuating demand.

A STOCHASTIC NETWORK: A TOY EXAMPLE

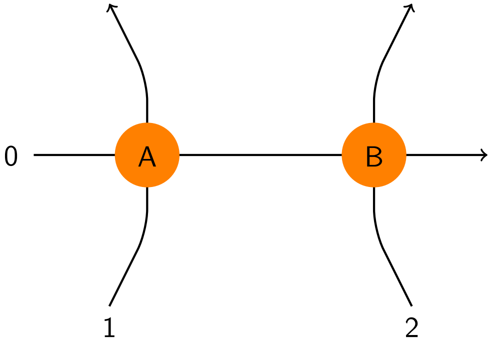

A STOCHASTIC NETWORK: A TOY EXAMPLE

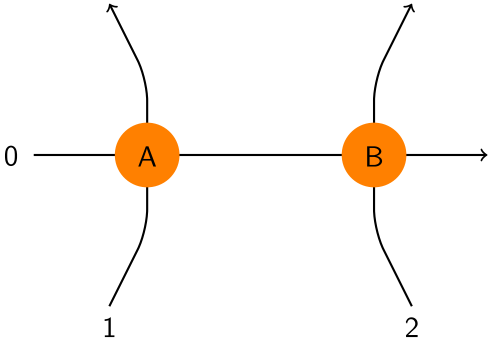

Class 0 uses node A and B,

A STOCHASTIC NETWORK: A TOY EXAMPLE

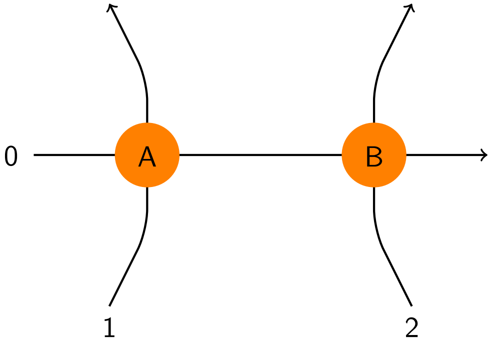

Class 0 uses node A and B, class 1 node A,

A STOCHASTIC NETWORK: A TOY EXAMPLE

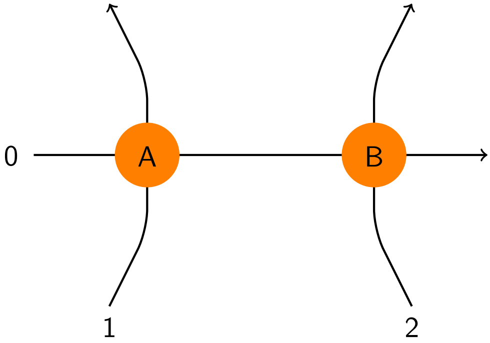

Class 0 uses node A and B, class 1 node A, and class 2 node B.

A STOCHASTIC NETWORK: A TOY EXAMPLE

This model is still rather imprecise:

A STOCHASTIC NETWORK: A TOY EXAMPLE

This model is still rather imprecise:

- How do customers arrive (Poisson process?), and what is the distribution about the amount of work they bring along?

A STOCHASTIC NETWORK: A TOY EXAMPLE

This model is still rather imprecise:

- How do customers arrive (Poisson process?), and what is the distribution about the amount of work they bring along?
- How is the available capacity shared among the users?

LARGER STOCHASTIC NETWORKS. . .

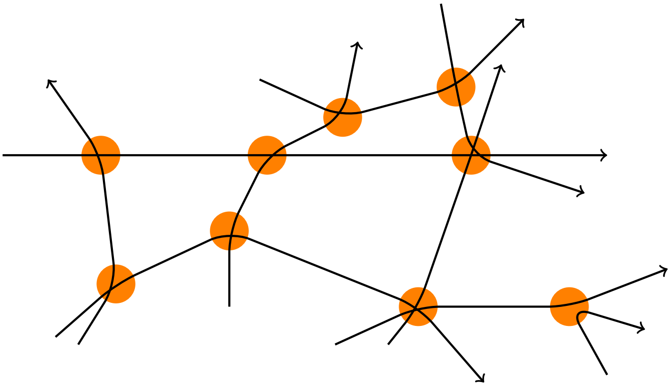

LARGER STOCHASTIC NETWORKS...

Complication: *link failures*

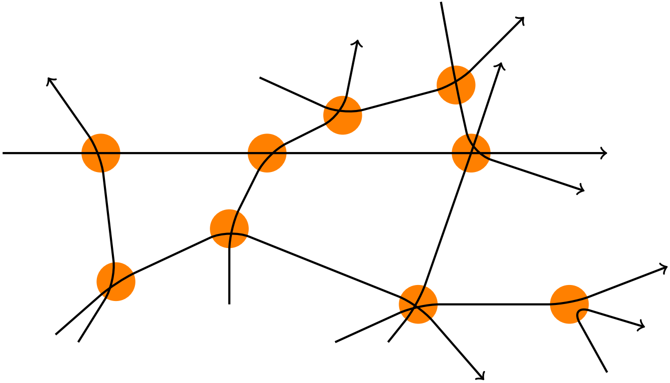

LARGER STOCHASTIC NETWORKS...

Complication: *link failures*

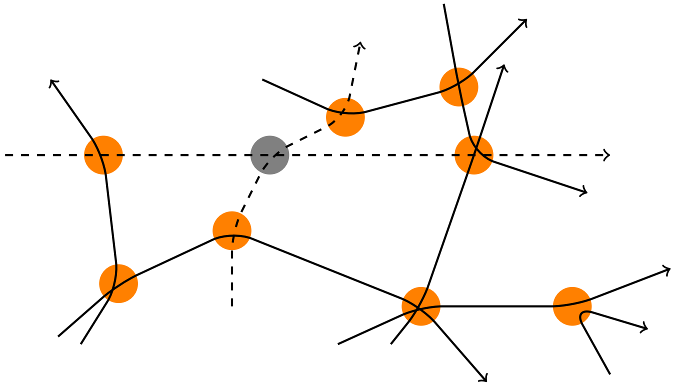

LARGER STOCHASTIC NETWORKS...

Complication: *link failures*

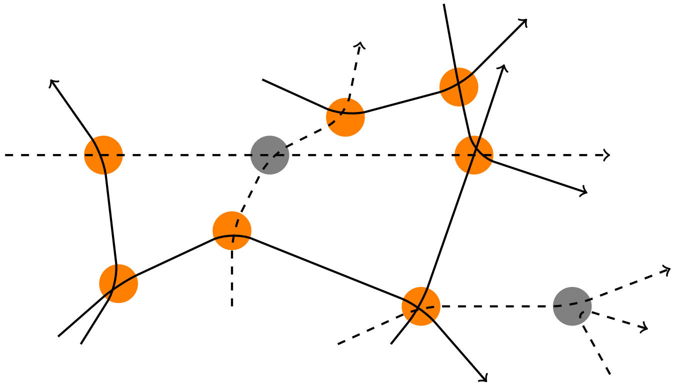

LARGER STOCHASTIC NETWORKS...

Complication: *link failures*

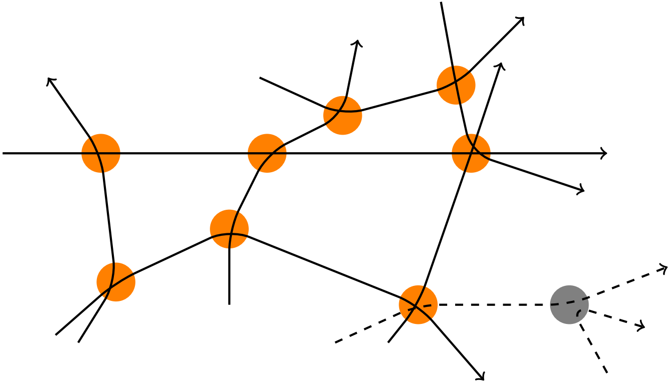

LARGER STOCHASTIC NETWORKS. . .

Observe that we consider

LARGER STOCHASTIC NETWORKS. . .

Observe that we consider

- Random dynamics on networks,

LARGER STOCHASTIC NETWORKS...

Observe that we consider

- Random dynamics on networks,
- but the structure of the network is random as well!

STOCHASTIC NETWORKS: COMMON FEATURES

- Multiple streams in network react to ‘random fluctuations in outer world’;

STOCHASTIC NETWORKS: COMMON FEATURES

- Multiple streams in network react to ‘random fluctuations in outer world’;
- this class of models can be applied in a wide variety of areas.

APPLICATION AREA 1: WIRED COMMUNICATION NETWORK

APPLICATION AREA 1: WIRED COMMUNICATION NETWORK

- links are routers in, say, the Internet;

APPLICATION AREA 1: WIRED COMMUNICATION NETWORK

- links are routers in, say, the Internet;
- links can be ‘up’ and ‘down’, affecting specific classes of users;

APPLICATION AREA 1: WIRED COMMUNICATION NETWORK

- links are routers in, say, the Internet;
- links can be ‘up’ and ‘down’, affecting specific classes of users;
- external factors can spur increased activity across user classes.

APPLICATION AREA 1: WIRED COMMUNICATION NETWORK

- links are routers in, say, the Internet;
- links can be ‘up’ and ‘down’, affecting specific classes of users;
- external factors can spur increased activity across user classes.

Likewise: wireless communication network, with multiple classes of users reacting to fluctuations in channel conditions.

APPLICATION AREA 2: ROAD TRAFFIC NETWORK

APPLICATION AREA 2: ROAD TRAFFIC NETWORK

- link failures could represent roadworks or accidents;

APPLICATION AREA 2: ROAD TRAFFIC NETWORK

- link failures could represent roadworks or accidents;
- external factors can spur increased activity across user classes (e.g. weather-related).

APPLICATION AREA 3: CHEMICAL REACTION NETWORK

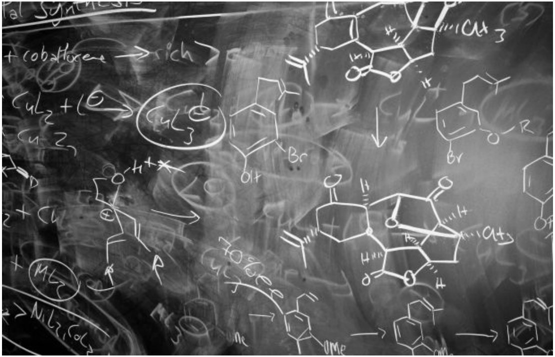

APPLICATION AREA 3: CHEMICAL REACTION NETWORK

- customers $\sim$ concentrations of different substances, that react with each other;

APPLICATION AREA 3: CHEMICAL REACTION NETWORK

- customers $\sim$ concentrations of different substances, that react with each other;
- external factors may affect reaction speed (e.g. temperature-related).

APPLICATION AREA 4: ECONOMICS

APPLICATION AREA 4: ECONOMICS

- prices of correlated economic assets randomly fluctuate;

APPLICATION AREA 4: ECONOMICS

- prices of correlated economic assets randomly fluctuate;
- and do so by reacting to the same ‘outer world’ (i.e., state of the economy);

APPLICATION AREA 4: ECONOMICS

- prices of correlated economic assets randomly fluctuate;
- and do so by reacting to the same ‘outer world’ (i.e., state of the economy);
- (but in this case one rather uses continuous state-space stochastic process, rather than discrete state-space — e.g. stochastic differential equations.)

GENERIC FRAMEWORK: NETWORK OF QUEUES, OPERATING UNDER MARKOV MODULATION

- particles ('customers' in queueing lingo) move through a network;

GENERIC FRAMEWORK: NETWORK OF QUEUES, OPERATING UNDER MARKOV MODULATION

- particles (‘customers’ in queueing lingo) move through a network;
- arrival rates and service rates are affected by an external process (‘background process’);

GENERIC FRAMEWORK: NETWORK OF QUEUES, OPERATING UNDER MARKOV MODULATION

- particles (‘customers’ in queueing lingo) move through a network;
- arrival rates and service rates are affected by an external process (‘background process’);
- queues are ‘coupled’ because they react to *common* background process.

QUEUEING NETWORK

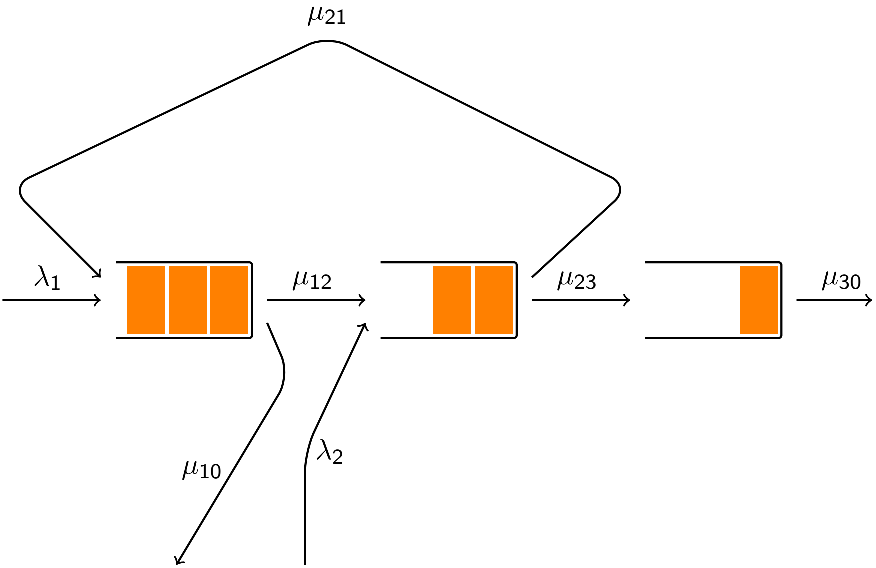

QUEUEING NETWORK, WITH MARKOV MODULATION

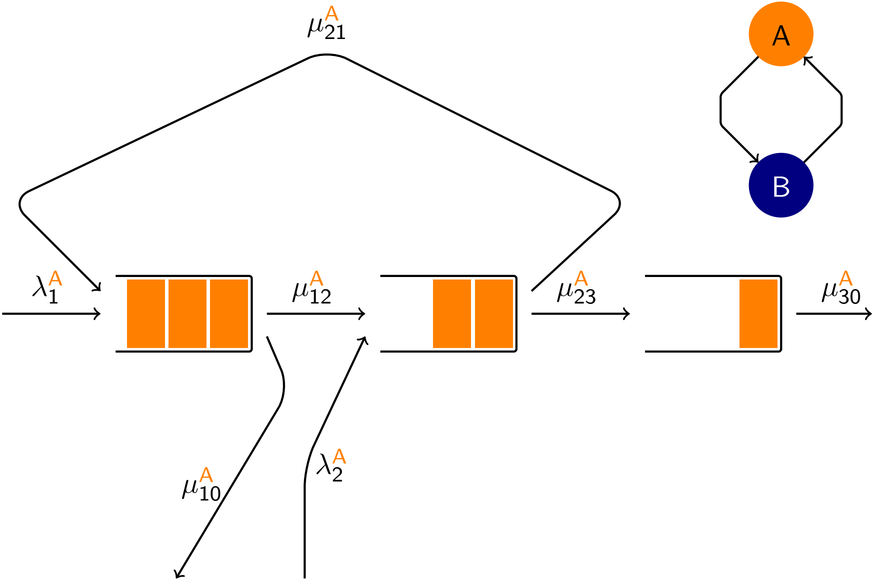

QUEUEING NETWORK, WITH MARKOV MODULATION

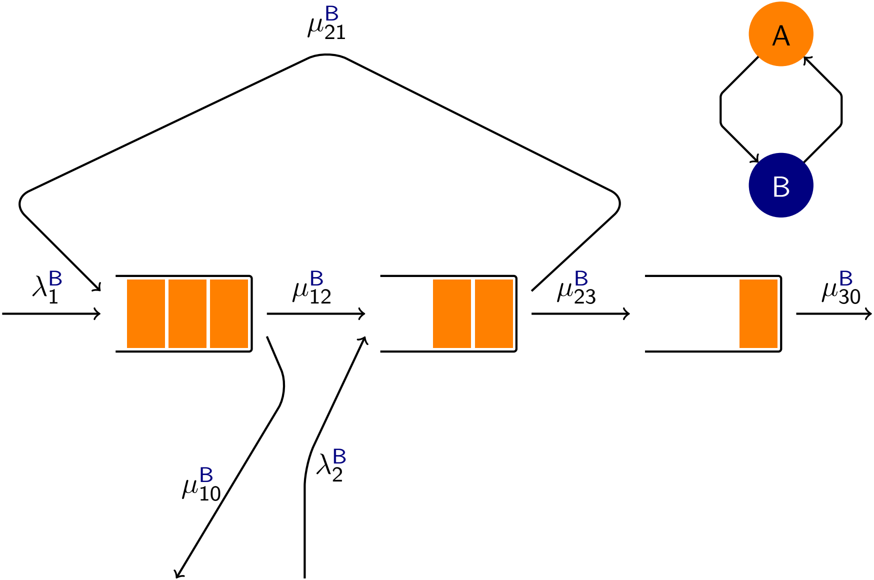

QUEUEING NETWORK, WITH MARKOV MODULATION

Framework is *rich*; covers also failure/repair system.

QUEUEING NETWORK, WITH MARKOV MODULATION

Framework is *rich*; covers also failure/repair system.

Background process corresponds to which links are ‘up’/‘down’.

QUEUEING NETWORK, WITH MARKOV MODULATION

Framework is *rich*; covers also failure/repair system.

Background process corresponds to which links are 'up'/'down'.

When link is down, all classes using this link have arrival rate 0 and departure rate ∞.

QUEUEING NETWORK, WITH MARKOV MODULATION

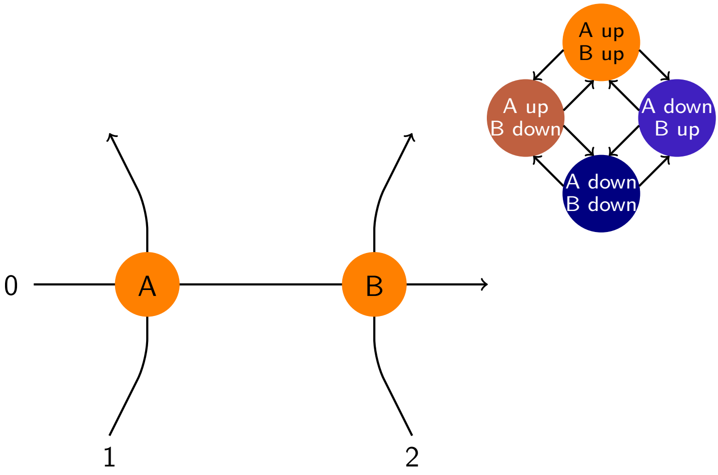

QUEUEING NETWORK, WITH MARKOV MODULATION

Class 0:

$\lambda_0^{\color{orange}\bullet} = \lambda_0, \quad \lambda_0^{\color{brown}\bullet} = \lambda_0^{\color{violet}\bullet} = \lambda_0^{\color{navy}\bullet} = 0$

$\mu_0^{\color{orange}\bullet} = \mu_0, \quad \mu_0^{\color{brown}\bullet} = \mu_0^{\color{violet}\bullet} = \mu_0^{\color{navy}\bullet} = \infty$

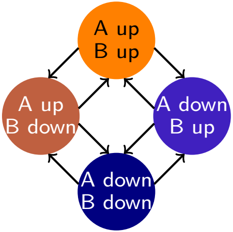

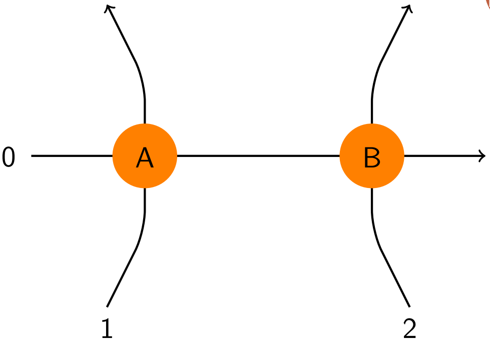

QUEUEING NETWORK, WITH MARKOV MODULATION

Class 1: (Class 2 dealt with analogously)

$$\lambda_1^{\bullet} = \lambda_1^{\bullet} = \lambda_1, \quad \lambda_1^{\bullet} = \lambda_1^{\bullet} = 0$$

$$\mu_1^{\bullet} = \mu_1^{\bullet} = \mu_1, \quad \mu_1^{\bullet} = \mu_1^{\bullet} = \infty$$

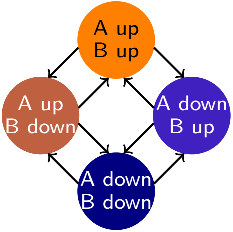

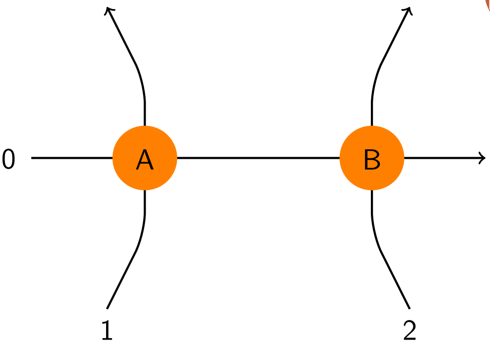

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Notoriously hard...

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Notoriously hard...
- Simplest model: two single-server queues, modulated by the same background process $X(t)$, irreducible continuous-time Markov chain, living on $\{1, \ldots, d\}$.

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Notoriously hard...
- Simplest model: two single-server queues, modulated by the same background process $X(t)$, irreducible continuous-time Markov chain, living on $\{1, \dots, d\}$.

  If $X(t) = i$, then

  - arrival rate of queue A is λ_i^{A} ($\rightarrow$ number of customers at queue A increases by 1), and of queue B it is λ_i^{B} ($\rightarrow$ number of customers at queue B increases by 1).
  - departure rate of queue A is μ_i^{A} ($\rightarrow$ number of customers at queue A decreases by 1), and of queue B it is μ_i^{B} ($\rightarrow$ number of customers at queue B decreases by 1).

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Notoriously hard...
- Simplest model: two single-server queues, modulated by the same background process $X(t)$, irreducible continuous-time Markov chain, living on $\{1, \dots, d\}$.

  If $X(t) = i$, then
  - arrival rate of queue A is λ_i^{A} ($\to$ number of customers at queue A increases by 1), and of queue B it is λ_i^{B} ($\to$ number of customers at queue B increases by 1).
  - departure rate of queue A is μ_i^{A} ($\to$ number of customers at queue A decreases by 1), and of queue B it is μ_i^{B} ($\to$ number of customers at queue B decreases by 1).
- Goal: joint distribution of stationary number in both queues:

$$\mathbb{P}(M^{\mathrm{A}} = k, M^{\mathrm{B}} = \ell).$$

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Ironically:
  marginals $\mathbb{P}(M^{\mathsf{A}} = k)$ and $\mathbb{P}(M^{\mathsf{B}} = \ell)$ can be found by elementary methods ('matrix-geometric form' — Neuts, early 1980s).

MARKOV-MODULATED SINGLE-SERVER NETWORK

- Ironically:
  marginals $\mathbb{P}(M^{\mathsf{A}} = k)$ and $\mathbb{P}(M^{\mathsf{B}} = \ell)$ can be found by elementary methods ('matrix-geometric form' — Neuts, early 1980s).
- Joint distribution has not been found (apart from trivial cases).

  Problem: discontinuity at 0 (as queue cannot become negative). Solution requires solving non-trivial boundary value problem, unless one queue systematically majorizes the other.

MARKOV-MODULATED INFINITE-SERVER NETWORK

When assuming *infinitely many* servers, analysis *is* possible.

MARKOV-MODULATED INFINITE-SERVER NETWORK

When assuming *infinitely many* servers, analysis *is* possible.

Infinite-server queue is useful proxy for model with many servers (channels in wireless network, call center, segment of a road, generation and decay of mRNA in cells, etc.).

‘CLASSICAL’ INFINITE-SERVER QUEUE

First: single infinite-server queue, in non-modulated setting.

‘CLASSICAL’ INFINITE-SERVER QUEUE

First: single infinite-server queue, in non-modulated setting.

$M(t)$ is number of customers at time t; lives on $\{0, 1, 2 \ldots\}$.

'CLASSICAL' INFINITE-SERVER QUEUE

First: single infinite-server queue, in non-modulated setting.

$M(t)$ is number of customers at time t; lives on $\{0, 1, 2 \ldots\}$.

- Rate up (from i to $i+1$) is λ,
- rate down (from i to $i-1$) is $i\,\mu$.

'CLASSICAL' INFINITE-SERVER QUEUE

First: single infinite-server queue, in non-modulated setting.

$M(t)$ is number of customers at time t; lives on $\{0, 1, 2 \ldots\}$.

- Rate up (from i to $i + 1$) is λ,
- rate down (from i to $i - 1$) is $i\,\mu$.

Fairly complete analysis is possible: steady-state, transient, various performance metrics, etc.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Now impose *Markov modulation.* As before, $X(t)$ is background process.

- $(X(t))_{t\geqslant 0}$: irreducible, Markov process on $\{1,\ldots,d\}$.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Now impose *Markov modulation.* As before, $X(t)$ is background process.

- $(X(t))_{t\geqslant 0}$: irreducible, Markov process on $\{1, \ldots, d\}$.
- Transition rates: $Q = (q_{ij})_{i,j=1}^{d}$, (unique) invariant distribution: $\boldsymbol{\pi}$.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Now impose *Markov modulation.* As before, $X(t)$ is background process.

- $(X(t))_{t\geqslant 0}$: irreducible, Markov process on $\{1,\ldots,d\}$.
- Transition rates: $Q=(q_{ij})_{i,j=1}^{d}$, (unique) invariant distribution: $\boldsymbol{\pi}$.

Let $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ be non-negative d-dimensional vectors.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Now impose *Markov modulation.* As before, $X(t)$ is background process.

- $(X(t))_{t \geqslant 0}$: irreducible, Markov process on $\{1, \ldots, d\}$.
- Transition rates: $Q = (q_{ij})_{i,j=1}^{d}$, (unique) invariant distribution: $\boldsymbol{\pi}$.

Let $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ be non-negative d-dimensional vectors.

$M(t)$ lives on $\{0, 1, 2 \ldots\}$.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Now impose *Markov modulation.* As before, $X(t)$ is background process.

- $(X(t))_{t \geqslant 0}$: irreducible, Markov process on $\{1, \ldots, d\}$.
- Transition rates: $Q = (q_{ij})_{i,j=1}^{d}$, (unique) invariant distribution: $\boldsymbol{\pi}$.

Let $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ be non-negative d-dimensional vectors.

$M(t)$ lives on $\{0, 1, 2 \ldots\}$.

- Rate up (from i to $i+1$) is $\lambda_{X(t)}$,
- rate down (from i to $i-1$) is $i\,\mu_{X(t)}$.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Relatively small number of papers available (< 30...).

MARKOV-MODULATED INFINITE-SERVER QUEUE

Relatively small number of papers available (< 30...).

Remarkably sharp distinction:

- *Markov-modulated single-server queues*: single queue easy, multiple queues hard (if not impossible);
- *Markov-modulated infinite-server queues*: partial results on single queue, but whatever can be done for single queue can be done for multiple queues as well.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Even in single queue, no explicit expression for distribution of stationary number of customers M...

MARKOV-MODULATED INFINITE-SERVER QUEUE

Even in single queue, no explicit expression for distribution of stationary number of customers M...

Available results for Markov-modulated infinite-server (MMIS) queue typically in terms of

- d-dimensional system of (partial) differential equations to describe pgf of $M(t)$ and stationary counterpart, M.
- Recursive scheme to determine all moments; for transient moments in all steps non-homogeneous system of linear differential equations needs to be solved.

MARKOV-MODULATED INFINITE-SERVER QUEUE

Even in single queue, no explicit expression for distribution of stationary number of customers M...

Available results for Markov-modulated infinite-server (MMIS) queue typically in terms of

- d-dimensional system of (partial) differential equations to describe pgf of $M(t)$ and stationary counterpart, M.
- Recursive scheme to determine all moments; for transient moments in all steps non-homogeneous system of linear differential equations needs to be solved.

See papers (between $\pm$ 1990 and $\pm$ 2005) by O'Cinneide/Purdue, Keilson/Servi, Adan/Fralix, D'Auria, ...

MARKOV-MODULATED INFINITE-SERVER QUEUE

For Markov-modulated single-server queue a lot is known (Neuts): stationary number in the system follows a *matrix-geometric distribution* (generalization of M/M/1).

MARKOV-MODULATED INFINITE-SERVER QUEUE

For Markov-modulated single-server queue a lot is known (Neuts): stationary number in the system follows a *matrix-geometric distribution* (generalization of M/M/1).

Therefore in the context of MMIS queue one would naïvely expect a matrix-Poisson distribution (generalization of M/M/∞)...

MARKOV-MODULATED INFINITE-SERVER QUEUE

For Markov-modulated single-server queue a lot is known (Neuts): stationary number in the system follows a *matrix-geometric distribution* (generalization of M/M/1).

Therefore in the context of MMIS queue one would naïvely expect a matrix-Poisson distribution (generalization of $M/M/\infty$)...

but this is *not* true.

MARKOV-MODULATED INFINITE-SERVER QUEUE

MMIS queue comes in two flavors.

In above model (referred to as Model I) the transition rates depend on the current state of the background process. $M(t)$ has Poisson distribution with random parameter

$$\int_0^t \lambda_{X(s)} e^{-\int_s^t \mu_{X(r)}\,\mathrm{d}r} \mathrm{d}s.$$

MARKOV-MODULATED INFINITE-SERVER QUEUE

MMIS queue comes in two flavors.

In above model (referred to as Model I) the transition rates depend on the current state of the background process. $M(t)$ has Poisson distribution with random parameter

$$\int_0^t \lambda_{X(s)} e^{-\int_s^t \mu_{X(r)}\,\mathrm{d}r} \mathrm{d}s.$$

Alternative model (Model II): service times are sampled upon arrival. $M(t)$ has Poisson distribution with random parameter

$$\int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}\,(t-s)} \mathrm{d}s.$$

MARKOV-MODULATED INFINITE-SERVER QUEUE

MMIS queue comes in two flavors.

In above model (referred to as Model I) the transition rates depend on the current state of the background process. $M(t)$ has Poisson distribution with random parameter

$$\int_0^t \lambda_{X(s)} e^{-\int_s^t \mu_{X(r)}\,\mathrm{d}r} \mathrm{d}s.$$

Alternative model (Model II): service times are sampled upon arrival. $M(t)$ has Poisson distribution with random parameter

$$\int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}\,(t-s)} \mathrm{d}s.$$

In this talk: μ_i identical across i, so that both models coincide.

MMIS: THE LOW HANGING FRUIT...

First characterize invariant distribution $(\boldsymbol{p}_k)_{k=0}^{\infty}$, where $\boldsymbol{p}_k$ is d-dimensional row-vector, defined by

$$[\boldsymbol{p}_k]_j := \mathbb{P}(M = k, X = j).$$

The (row-vector-)pgf $\boldsymbol{p}(z)$ is then given by

$$\boldsymbol{p}(z) := \sum_{k=0}^{\infty} \boldsymbol{p}_k z^k.$$

MMIS: THE LOW HANGING FRUIT...

First characterize invariant distribution $(\boldsymbol{p}_k)_{k=0}^{\infty}$, where $\boldsymbol{p}_k$ is d-dimensional row-vector, defined by

$$[\boldsymbol{p}_k]_j := \mathbb{P}(M = k, X = j).$$

The (row-vector-)pgf $\boldsymbol{p}(z)$ is then given by

$$\boldsymbol{p}(z) := \sum_{k=0}^{\infty} \boldsymbol{p}_k z^k.$$

Elementary (from Kolmogorov equations): $\boldsymbol{p}(z)$ satisfies ODE

$$\boldsymbol{p}(z) Q = (z-1)\left(\boldsymbol{p}'(z)\mathrm{diag}\{\boldsymbol{\mu}\} - \boldsymbol{p}(z)\mathrm{diag}\{\boldsymbol{\lambda}\}\right).$$

For transient behavior we obtain similar DE (which is a PDE).

MMIS: THE LOW HANGING FRUIT...

With ODE

$$\boldsymbol{p}(z)Q = (z-1)\left(\boldsymbol{p}'(z)\mathrm{diag}\{\boldsymbol{\mu}\} - \boldsymbol{p}(z)\mathrm{diag}\{\boldsymbol{\lambda}\}\right).$$

stationary (factorial) moments can be found by differentiation and plugging in $z \uparrow 1$.

MMIS: THE LOW HANGING FRUIT...

With ODE

$$\boldsymbol{p}(z)Q = (z-1)\left(\boldsymbol{p}'(z)\mathrm{diag}\{\boldsymbol{\mu}\} - \boldsymbol{p}(z)\mathrm{diag}\{\boldsymbol{\lambda}\}\right).$$

stationary (factorial) moments can be found by differentiation and plugging in $z \uparrow 1$.

Define

$$\boldsymbol{m}_k := \mathbb{E}[M(M-1)\cdots(M-k+1)1_{\{X=i\}}] = \boldsymbol{p}^{(k)}(1).$$

Recursion (realize $\boldsymbol{m}_0 = \boldsymbol{\pi}$):

$$\boldsymbol{m}_k Q = k\boldsymbol{m}_k \,\mathrm{diag}\{\boldsymbol{\mu}\} - k\boldsymbol{m}_{k-1}\,\mathrm{diag}\{\boldsymbol{\lambda}\}.$$

MMIS: THE LOW HANGING FRUIT...

With ODE

$$\boldsymbol{p}(z)Q = (z-1)\left(\boldsymbol{p}'(z)\mathrm{diag}\{\boldsymbol{\mu}\} - \boldsymbol{p}(z)\mathrm{diag}\{\boldsymbol{\lambda}\}\right).$$

stationary (factorial) moments can be found by differentiation and plugging in $z \uparrow 1$.

Define

$$\boldsymbol{m}_k := \mathbb{E}[M(M-1)\cdots(M-k+1)1_{\{X=i\}}] = \boldsymbol{p}^{(k)}(1).$$

Recursion (realize $\boldsymbol{m}_0 = \boldsymbol{\pi}$):

$$\boldsymbol{m}_k Q = k\boldsymbol{m}_k \,\mathrm{diag}\{\boldsymbol{\mu}\} - k\boldsymbol{m}_{k-1}\,\mathrm{diag}\{\boldsymbol{\lambda}\}.$$

(Same for transient moments: then in each step of the recursion non-homogeneous system of differential equations must be solved.)

MMIS: THE LOW HANGING FRUIT...

This is nice – one could numerically analyze the model now. However, we'd like to have 'structural insight' into the model...

MMIS: THE LOW HANGING FRUIT...

This is nice – one could numerically analyze the model now. However, we'd like to have 'structural insight' into the model...

Therefore: consider *scaling limits*.

MMIS: THE LOW HANGING FRUIT...

This is nice – one could numerically analyze the model now. However, we'd like to have 'structural insight' into the model...

Therefore: consider *scaling limits*.

We let some of the parameters of the model (viz. $\boldsymbol{\lambda}$, $\boldsymbol{\mu}$, and Q) grow large of small, in a 'coordinated manner', and see whether we obtain any explicit results...

MMIS: SCALING LIMITS...

'Black magic': what is the right scaling?
To provide intuition, let's explicitly compute the mean and variance of $M(t)$.

MMIS: MEAN AND VARIANCE

Straightforward (for instance from Poisson-with-random-mean representation):

$$\mathbb{E}M(t) = \sum_{i=1}^{d} \pi_i \frac{\lambda_i}{\mu}(1 - e^{-\mu t}).$$

MMIS: MEAN AND VARIANCE

Straightforward (for instance from Poisson-with-random-mean representation):

$$\mathbb{E}M(t) = \sum_{i=1}^{d} \pi_i \frac{\lambda_i}{\mu}(1 - e^{-\mu t}).$$

Scaling the λ_is by N blows up scale of process by a factor N...

MMIS: MEAN AND VARIANCE

Variance can be computed with law of total variance:

$$\mathbb{V}\mathrm{ar} M(t) = \mathbb{E}(\mathbb{V}\mathrm{ar}(M(t) \,|\, X)) + \mathbb{V}\mathrm{ar}(\mathbb{E}(M(t) \,|\, X)),$$

with $X \equiv (X(s))_{s \in [0,t]}$.

MMIS: MEAN AND VARIANCE

Variance can be computed with law of total variance:

$$\mathbb{V}\mathrm{ar} M(t) = \mathbb{E}(\mathbb{V}\mathrm{ar}(M(t) \mid X)) + \mathbb{V}\mathrm{ar}(\mathbb{E}(M(t) \mid X)),$$

with $X \equiv (X(s))_{s \in [0,t]}$.

Clearly,

$$\mathbb{E}(\mathbb{V}\mathrm{ar}(M(t) \mid X)) = \mathbb{E} M(t) = \sum_{i=1}^{d} \pi_i \frac{\lambda_i}{\mu}(1 - e^{-\mu t}).$$

MMIS: MEAN AND VARIANCE

$$\mathbb{V}\mathrm{ar}(\mathbb{E}(M(t)\,|\,X)) = \mathbb{V}\mathrm{ar}\left(\int_0^t \lambda_{X(s)} e^{-\mu(t-s)} \mathrm{d}s\right)$$

$$= \int_0^t \int_0^t \mathbb{C}\mathrm{ov}\left(\lambda_{X(s)} e^{-\mu(t-s)}, \lambda_{X(u)} e^{-\mu(t-u)}\right) \mathrm{d}s\mathrm{d}u$$

$$= \sum_{i,j=1}^{d} \lambda_i \lambda_j \int_0^t \int_0^t e^{-\mu(t-s)} e^{-\mu(t-u)} \mathbb{C}\mathrm{ov}\left(1_{\{X(s)=i\}}, 1_{\{X(u)=j\}}\right) \mathrm{d}s\mathrm{d}u.$$

MMIS: MEAN AND VARIANCE

$$\mathbb{V}\mathrm{ar}(\mathbb{E}(M(t)\,|\,X)) = \mathbb{V}\mathrm{ar}\left(\int_0^t \lambda_{X(s)} e^{-\mu(t-s)} \mathrm{d}s\right)$$

$$= \int_0^t \int_0^t \mathbb{C}\mathrm{ov}\left(\lambda_{X(s)} e^{-\mu(t-s)}, \lambda_{X(u)} e^{-\mu(t-u)}\right) \mathrm{d}s\mathrm{d}u$$

$$= \sum_{i,j=1}^{d} \lambda_i \lambda_j \int_0^t \int_0^t e^{-\mu(t-s)} e^{-\mu(t-u)} \mathbb{C}\mathrm{ov}\left(1_{\{X(s)=i\}}, 1_{\{X(u)=j\}}\right) \mathrm{d}s\mathrm{d}u.$$

Reduces to:

$$\sum_{i,j=1}^{d} \lambda_i \lambda_j \int_0^t \int_0^u e^{-\mu(t-s)} e^{-\mu(t-u)} \pi_i (p_{ij}(u-s) - \pi_j) \mathrm{d}s\mathrm{d}u$$

$$+ \sum_{i,j=1}^{d} \lambda_i \lambda_j \int_0^t \int_u^t e^{-\mu(t-s)} e^{-\mu(t-u)} \pi_i (p_{ij}(u-s) - \pi_j) \mathrm{d}s\mathrm{d}u.$$

MMIS: MEAN AND VARIANCE

Deviation matrix:

$$D_{ij} := \int_0^\infty (p_{ij}(t) - \pi_j)\mathrm{d}t.$$

MMIS: MEAN AND VARIANCE

Deviation matrix:

$$D_{ij} := \int_0^\infty (p_{ij}(t) - \pi_j)\mathrm{d}t.$$

Perform parameter scaling $\boldsymbol{\lambda} \mapsto \boldsymbol{\lambda} N$, and $Q \mapsto QN^f$, for some $f > 0$.

MMIS: MEAN AND VARIANCE

Deviation matrix:

$$D_{ij} := \int_0^\infty (p_{ij}(t) - \pi_j)\mathrm{d}t.$$

Perform parameter scaling $\boldsymbol{\lambda} \mapsto \boldsymbol{\lambda} N$, and $Q \mapsto QN^f$, for some $f > 0$.

Elementary calculations for stationary number in system:

$$\mathbb{V}\mathrm{ar} M^{(N)} \sim N \frac{\lambda_\infty}{\mu} + N^{2-f} \sum_{i,j=1}^{d} \pi_i \frac{\lambda_i \lambda_j}{\mu} D_{ij},$$

with $\lambda_\infty := \sum_{i=1}^{d} \pi_i \lambda_i$.

MMIS: DICHOTOMY

Interesting dichotomy:

MMIS: DICHOTOMY

Interesting dichotomy:

- If $f > 1$ the variance essentially equals

$$\mathbb{V}\mathrm{ar}\, M^{(N)} \sim N\varrho, \quad \text{where} \quad \varrho := \frac{\lambda_\infty}{\mu}.$$

The system behaves 'Poissonian': background process moves faster than arrival process.
Limiting system is effectively a *non-modulated* infinite-server queue.

MMIS: DICHOTOMY

Interesting dichotomy:

- If $f > 1$ the variance essentially equals

$$\mathbb{V}\mathrm{ar}\, M^{(N)} \sim N\varrho, \quad \text{where} \quad \varrho := \frac{\lambda_\infty}{\mu}.$$

  The system behaves 'Poissonian': background process moves faster than arrival process.
  Limiting system is effectively a *non-modulated* infinite-server queue.

- If $f < 1$ the variance essentially equals

$$\mathbb{V}\mathrm{ar}\, M^{(N)} \sim N^{2-f} \sum_{i,j=1}^{d} \pi_i \frac{\lambda_i \lambda_j}{\mu} D_{ij}.$$

  'Local equilibria'.

MMIS: DICHOTOMY

We consider two types of limit results:

MMIS: DICHOTOMY

We consider two types of limit results:

- Behavior 'around the mean': *central limit theorems*. Crucially different behavior for $f < 1$, $f = 1$, and $f > 1$: apparently the right CLT scaling is N^γ, with

$$\gamma := \max\left\{\frac{1}{2}, 1 - \frac{f}{2}\right\}.$$

MMIS: DICHOTOMY

We consider two types of limit results:

- Behavior 'around the mean': *central limit theorems*.
  Crucially different behavior for $f < 1$, $f = 1$, and $f > 1$: apparently the right CLT scaling is N^γ, with

$$\gamma := \max\left\{\frac{1}{2}, 1 - \frac{f}{2}\right\}.$$

- Rare-event behavior, 'far away from the mean': *large deviations*.
  Again crucially different behavior for $f < 1$, $f = 1$, and $f > 1$.

MMIS: COAUTHORS

- At University of Amsterdam: *Peter Spreij* and *Gang Huang*.
- At CWI: *Joke Blom* and *Halldóra Þórsdottir*.

 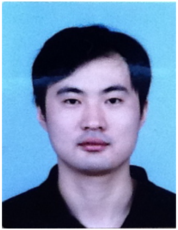  

MMIS: COAUTHORS

- At University of Melbourne: *Peter Taylor*.
- At Hebrew University: *Offer Kella*.
- At Supélec Paris: *Koen de Turck*.
- At University Ghent: *Marijn Jansen*.

 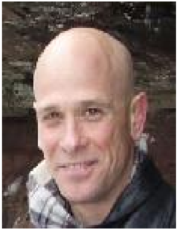  

REST OF THE TALK

- Central limit theorems,
- Large deviations (very brief, time permitting!).

REST OF THE TALK

- Central limit theorems,
- Large deviations (very brief, time permitting!).

For both I'll present the main ideas and underlying reasoning, state the result in its basic form. Many extensions, generalizations, and ramifications are possible.

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

- Set up a DE for the PGF of $M^{(N)}$.

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

- Set up a DE for the PGF of $M^{(N)}$.
- Transform this is into a DE for the MGF of

$$\tilde{M}^{(N)} := \frac{M^{(N)} - N\varrho}{N^{\gamma}}.$$

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

- Set up a DE for the PGF of $M^{(N)}$.
- Transform this is into a DE for the MGF of

$$\tilde{M}^{(N)} := \frac{M^{(N)} - N\varrho}{N^{\gamma}}.$$

- Manipulate this expression and let $N \to \infty$.

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

- Set up a DE for the PGF of $M^{(N)}$.
- Transform this is into a DE for the MGF of

$$\tilde{M}^{(N)} := \frac{M^{(N)} - N\varrho}{N^{\gamma}}.$$

- Manipulate this expression and let $N \to \infty$.
- Observe that we obtain a Gaussian limit.

CENTRAL LIMIT THEOREM

Basic form: single MMIS queue, stationary behavior.

- Set up a DE for the PGF of $M^{(N)}$.
- Transform this is into a DE for the MGF of

$$\tilde{M}^{(N)} := \frac{M^{(N)} - N\varrho}{N^{\gamma}}.$$

- Manipulate this expression and let $N \to \infty$.
- Observe that we obtain a Gaussian limit.

By now we have various alternative techniques (generator-based; martingale-based); this one is most insightful.

CENTRAL LIMIT THEOREM, ctd.

First characterize invariant distribution $(\boldsymbol{p}_k^{(N)})_{k=0}^{\infty}$, where $\boldsymbol{p}_k^{(N)}$ is d-dimensional row-vector, defined by

$$[\boldsymbol{p}_k^{(N)}]_j := \mathbb{P}(M^{(N)} = k, X^{(N)} = j).$$

The (row-vector-)pgf $\boldsymbol{p}^{(N)}(z)$ is then given by

$$\boldsymbol{p}^{(N)}(z) := \sum_{k=0}^{\infty} \boldsymbol{p}_k^{(N)} z^k.$$

CENTRAL LIMIT THEOREM, ctd.

First characterize invariant distribution $(\boldsymbol{p}_k^{(N)})_{k=0}^{\infty}$, where $\boldsymbol{p}_k^{(N)}$ is d-dimensional row-vector, defined by

$$[\boldsymbol{p}_k^{(N)}]_j := \mathbb{P}(M^{(N)} = k, X^{(N)} = j).$$

The (row-vector-)pgf $\boldsymbol{p}^{(N)}(z)$ is then given by

$$\boldsymbol{p}^{(N)}(z) := \sum_{k=0}^{\infty} \boldsymbol{p}_k^{(N)} z^k.$$

Kolmogorov equations are now given by

$$\boldsymbol{p}^{(N)}(z)Q = \frac{(z-1)}{N^f}\left((\boldsymbol{p}^{(N)})'(z)\mathrm{diag}\{\boldsymbol{\mu}\} - N\boldsymbol{p}^{(N)}(z)\mathrm{diag}\{\boldsymbol{\lambda}\}\right).$$

CENTRAL LIMIT THEOREM, ctd.

Translate into mgf of $\tilde{M}^{(N)}$:

$$\begin{aligned} \tilde{\boldsymbol{p}}^{(N)}(\vartheta) &:= \mathbb{E} e^{\vartheta \tilde{M}^{(N)}} = \mathbb{E} \exp\left(\vartheta \frac{M^{(N)} - N\varrho}{N^{\gamma}}\right) \\ &= e^{-\vartheta N^{1-\gamma}\varrho} \boldsymbol{p}^{(N)}\left(e^{\vartheta N^{-\gamma}}\right). \end{aligned}$$

CENTRAL LIMIT THEOREM, ctd.

Manipulate resulting DE.

CENTRAL LIMIT THEOREM, ctd.

Manipulate resulting DE.

- $\Pi := \mathbf{1}\boldsymbol{\pi}^{\mathrm{T}}$.

CENTRAL LIMIT THEOREM, ctd.

Manipulate resulting DE.

- $\Pi := \mathbf{1}\boldsymbol{\pi}^{\mathrm{T}}$.
- $F := D + \Pi$ (fundamental matrix).

CENTRAL LIMIT THEOREM, ctd.

Manipulate resulting DE.

- $\Pi := \mathbf{1}\boldsymbol{\pi}^{\mathrm{T}}$.
- $F := D + \Pi$ (fundamental matrix).
- Standard properties: $QF = FQ = \Pi - I$, $F\mathbf{1} = \mathbf{1}$, and $\Pi D = D\Pi = 0$.

CENTRAL LIMIT THEOREM, ctd.

Manipulate resulting DE.

- $\Pi := \mathbf{1}\boldsymbol{\pi}^{\mathrm{T}}$.
- $F := D + \Pi$ (fundamental matrix).
- Standard properties: $QF = FQ = \Pi - I$, $F\mathbf{1} = \mathbf{1}$, and $\Pi D = D\Pi = 0$.

Postmultiply DE by F.

CENTRAL LIMIT THEOREM, ctd.

When the dust has settled...

$$\begin{aligned}\tilde{\boldsymbol{p}}^{(N)}(\vartheta) &= \tilde{\boldsymbol{p}}^{(N)}(\vartheta)\Pi + N^{1-f}\left(z^{(N)}(\vartheta)-1\right)\tilde{\boldsymbol{p}}^{(N)}(\vartheta)\mathrm{diag}\{\boldsymbol{\lambda}\}F \\ &\quad - N^{1-f}\left(1-\frac{1}{z^{(N)}(\vartheta)}\right)\varrho\,\tilde{\boldsymbol{p}}^{(N)}(\vartheta)\mathrm{diag}\{\boldsymbol{\mu}\}F \\ &\quad - N^{1-f-\beta/2}\left(1-\frac{1}{z^{(N)}(\vartheta)}\right)(\tilde{\boldsymbol{p}}^{(N)})'(\vartheta)\mathrm{diag}\{\boldsymbol{\mu}\}F.\end{aligned}$$

Here: $\beta := \min\{f,1\}$, and $z \equiv z^{(N)}(\vartheta) := \exp(\vartheta N^{-1+\beta/2})$.

CENTRAL LIMIT THEOREM, ctd.

Then

CENTRAL LIMIT THEOREM, ctd.

Then

- 'Taylor' the z, and iterate the equation to get rid of all terms that are $o(N^{-f})$:

CENTRAL LIMIT THEOREM, ctd.

Then

- 'Taylor' the z, and iterate the equation to get rid of all terms that are $o(N^{-f})$:
- Goal: transform the coupled system of ODE's in $\tilde{\boldsymbol{p}}^{(N)}(\vartheta)$ into a single-dimensional ODE in terms of $\phi^{(N)}(\vartheta) := \tilde{\boldsymbol{p}}^{(N)}(\vartheta)\mathbf{1}$. Postmultiply by $\mathbf{1}\, N^f/\vartheta$; realize that $\Pi\mathbf{1} = \mathbf{1}$ and $F\mathbf{1} = \mathbf{1}$.

CENTRAL LIMIT THEOREM, ctd.

We thus obtain

$$(\phi^{(N)})'(\vartheta) = \vartheta N^{\beta-f}\kappa\phi^{(N)}(\vartheta) + \vartheta N^{\beta-1}\varrho\phi^{(N)}(\vartheta) + o(1), \text{ with}$$

$$\kappa := \frac{\boldsymbol{\pi}^{\mathrm{T}}(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})F(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})\mathbf{1}}{\mu}.$$

CENTRAL LIMIT THEOREM, ctd.

We thus obtain

$$(\phi^{(N)})'(\vartheta) = \vartheta N^{\beta-f}\kappa\phi^{(N)}(\vartheta) + \vartheta N^{\beta-1}\varrho\phi^{(N)}(\vartheta) + o(1), \text{ with}$$

$$\kappa := \frac{\boldsymbol{\pi}^{\mathrm{T}}(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})F(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})\mathbf{1}}{\mu}.$$

Conclude, recalling that $\beta = \min\{f, 1\}$,

- $f < 1$: only first term RHS matters $\rightarrow$ Normal distribution with variance
$$\sum_{i,j=1}^{d} \pi_i \frac{\lambda_i \lambda_j}{\mu} D_{ij}.$$

CENTRAL LIMIT THEOREM, ctd.

We thus obtain

$$(\phi^{(N)})'(\vartheta) = \vartheta N^{\beta - f}\kappa\phi^{(N)}(\vartheta) + \vartheta N^{\beta-1}\varrho\phi^{(N)}(\vartheta) + o(1), \text{ with}$$

$$\kappa := \frac{\boldsymbol{\pi}^{\mathrm{T}}(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})F(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})\mathbf{1}}{\mu}.$$

Conclude, recalling that $\beta = \min\{f, 1\}$,

- $f < 1$: only first term RHS matters $\rightarrow$ Normal distribution with variance
$$\sum_{i,j=1}^{d} \pi_i \frac{\lambda_i \lambda_j}{\mu} D_{ij}.$$
- $f > 1$: only second term RHS matters $\rightarrow$ Normal distribution with variance
$$\sum_{i=1}^{d} \pi_i \frac{\lambda_i}{\mu} = \frac{\lambda_\infty}{\mu}.$$

CENTRAL LIMIT THEOREM, ctd.

We thus obtain

$$(\phi^{(N)})'(\vartheta) = \vartheta N^{\beta - f}\kappa\phi^{(N)}(\vartheta) + \vartheta N^{\beta-1}\varrho\phi^{(N)}(\vartheta) + o(1), \text{ with}$$

$$\kappa := \frac{\boldsymbol{\pi}^{\mathrm{T}}(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})F(\mathrm{diag}\{\boldsymbol{\lambda}\} - \varrho\,\mathrm{diag}\{\boldsymbol{\mu}\})\mathbf{1}}{\mu}.$$

Conclude, recalling that $\beta = \min\{f, 1\}$,

- $f < 1$: only first term RHS matters $\rightarrow$ Normal distribution with variance
$$\sum_{i,j=1}^{d} \pi_i \frac{\lambda_i \lambda_j}{\mu} D_{ij}.$$
- $f > 1$: only second term RHS matters $\rightarrow$ Normal distribution with variance
$$\sum_{i=1}^{d} \pi_i \frac{\lambda_i}{\mu} = \frac{\lambda_\infty}{\mu}.$$
- $f = 1$: *both* terms matter.

LARGE DEVIATIONS

Under the same scaling, large deviations can be examined.
Objective in transient case:

$$\lim_{N\to\infty} \frac{1}{N} \log \mathbb{P}\left(\frac{M^{(N)}(t)}{N} \geq a\right);$$

for stationary case, replace $M^{(N)}(t)$ by $M^{(N)}$.

Again crucially different behavior for $f > 1$ and $f < 1$.

LARGE DEVIATIONS, ctd.

First concentrate on $f > 1$.

LARGE DEVIATIONS, ctd.

First concentrate on $f > 1$.

- Stationary case: rate function looks like that of Poisson random variable with parameter

$$\varrho := \frac{\lambda_\infty}{\mu}.$$

LARGE DEVIATIONS, ctd.

First concentrate on $f > 1$.

- Stationary case: rate function looks like that of Poisson random variable with parameter

$$\varrho := \frac{\lambda_\infty}{\mu}.$$

- Transient case: same result, but then with parameter

$$\varrho_t := \frac{\lambda_\infty}{\mu}(1 - e^{-\mu t}).$$

LARGE DEVIATIONS, ctd.

Second regime: $f < 1$.

Take for ease $f = 0$ (that is, background process is unscaled) and Model II (for Model I analysis is similar). Recall: $M^{(N)}(t)$ has a Poisson distribution with parameter

$$N \int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}(t-s)} \mathrm{d}s.$$

LARGE DEVIATIONS, ctd.

Second regime: $f < 1$.

Take for ease $f = 0$ (that is, background process is unscaled) and Model II (for Model I analysis is similar). Recall: $M^{(N)}(t)$ has a Poisson distribution with parameter

$$N \int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}(t-s)} \mathrm{d}s.$$

- A *single* path $f(s)$ of $X(s)$ $(s \in [0, t])$ determines asymptotics.

LARGE DEVIATIONS, ctd.

Second regime: $f < 1$.

Take for ease $f = 0$ (that is, background process is unscaled) and Model II (for Model I analysis is similar). Recall: $M^{(N)}(t)$ has a Poisson distribution with parameter

$$N \int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}(t-s)} \mathrm{d}s.$$

- A *single* path $f(s)$ of $X(s)$ ($s \in [0, t]$) determines asymptotics.
- Naïve first thought: background process (essentially) stays in state i that maximizes λ_i / μ_i.

LARGE DEVIATIONS, ctd.

Second regime: $f < 1$.

Take for ease $f = 0$ (that is, background process is unscaled) and Model II (for Model I analysis is similar). Recall: $M^{(N)}(t)$ has a Poisson distribution with parameter

$$N \int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}(t-s)} \mathrm{d}s.$$

- A *single* path $f(s)$ of $X(s)$ ($s \in [0, t]$) determines asymptotics.
- Naïve first thought: background process (essentially) stays in state i that maximizes λ_i/μ_i.
  Wrong! Result: $X(s)$ close to path $f^\star(s)$, defined by

  $$\arg\max_{f(s)} \lambda_{f(s)}\, e^{-\mu_{f(s)}(t-s)}.$$

  Idea: maximize parameter of Poisson distribution.

LARGE DEVIATIONS, ctd.

Second regime: $f < 1$.

Take for ease $f = 0$ (that is, background process is unscaled) and Model II (for Model I analysis is similar). Recall: $M^{(N)}(t)$ has a Poisson distribution with parameter

$$N \int_0^t \lambda_{X(s)} e^{-\mu_{X(s)}(t-s)} \mathrm{d}s.$$

- A *single* path $f(s)$ of $X(s)$ ($s \in [0, t]$) determines asymptotics.
- Naïve first thought: background process (essentially) stays in state i that maximizes λ_i / μ_i.
  Wrong! Result: $X(s)$ close to path $f^\star(s)$, defined by

  $$\arg\max_{f(s)} \lambda_{f(s)}\, e^{-\mu_{f(s)}(t-s)}.$$

  Idea: maximize parameter of Poisson distribution.

Again, this was result in its *basic form*. Many extensions possible!

EPILOGUE

- Area of stochastic networks highly relevant, and mathematically extremely rich,

EPILOGUE

- Area of stochastic networks highly relevant, and mathematically extremely rich,
- with many challenges for the years to come,

EPILOGUE

- Area of stochastic networks highly relevant, and mathematically extremely rich,
- with many challenges for the years to come,
- particularly at the interface with algorithmics/combinatorics and statistics,

EPILOGUE

- Area of stochastic networks highly relevant, and mathematically extremely rich,
- with many challenges for the years to come,
- particularly at the interface with algorithmics/combinatorics and statistics,
- examples in talk illustrate how they complement each other.

EPILOGUE

- Area of stochastic networks highly relevant, and mathematically extremely rich,
- with many challenges for the years to come,
- particularly at the interface with algorithmics/combinatorics and statistics,
- examples in talk illustrate how they complement each other.

Thanks for your attention!