Scaling Properties of Traffic in Communication Networks

—

Probabilistic Resources Allocation in Cloud Environments

Paulo Gonçalves
Inria (DANTE) - ENS Lyon

Shubhabrata Roy (PhD, 2010-2013)
Thomas Begin (Ass. Prof., UCBL Lyon 1)
Patrick Loiseau (Ass. Prof. EURECOM)

CLOSER 2014 – Barcelona (Spain) 3-5 April 2014

Historical perspective

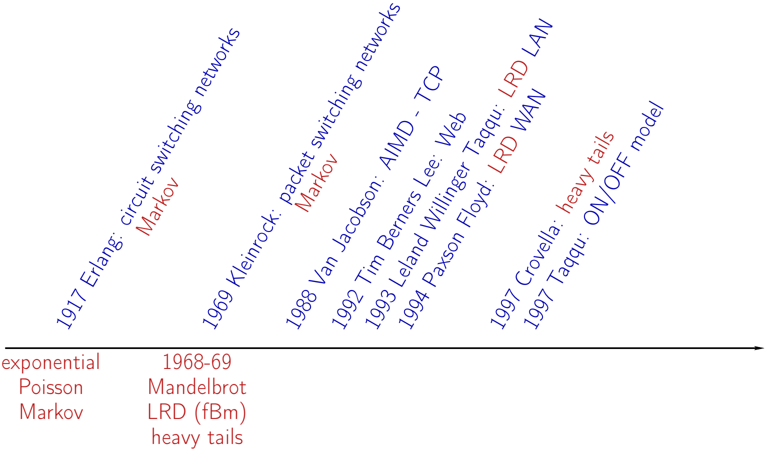

Historical perspective

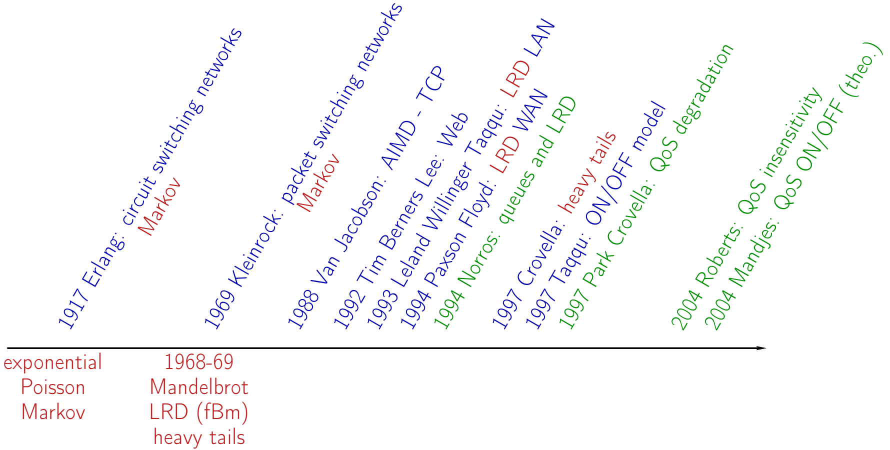

Historical perspective

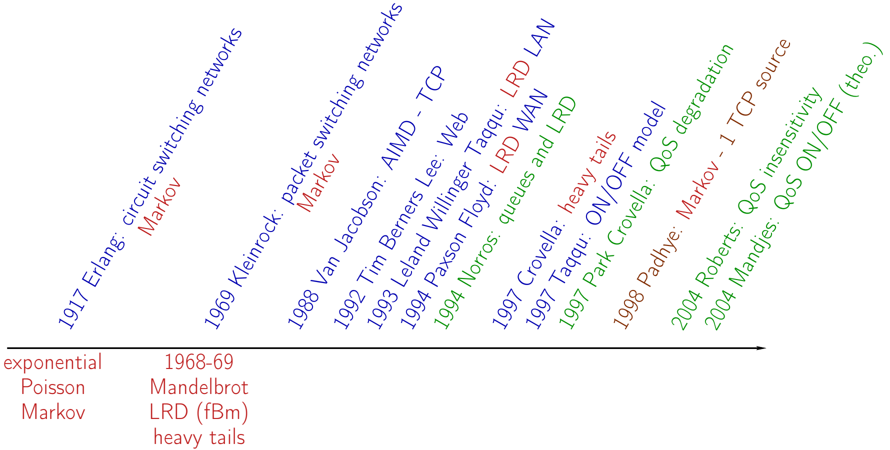

Some open questions:

- Long Range Dependence / Heavy Tailed distributions impact on QoS ?
- Existing models (e.g. Padhye) only predict mean metrics (e.g. throughput) : what about variability?

Our approach

To combine theoretical models with controlled experiments in realistic environments and real-world traffic traces

Simplified System

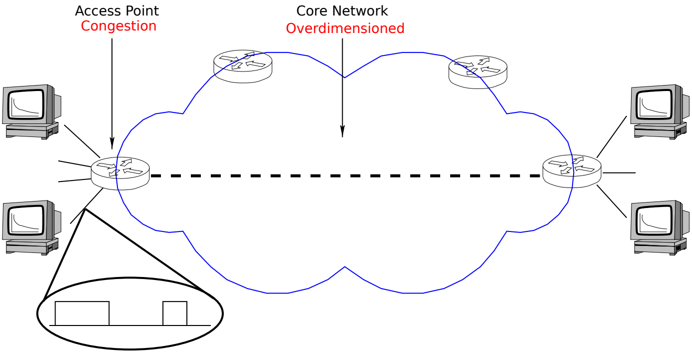

- Congestion essentially arises at the access points
  $\rightarrow$ Simplified System : single bottleneck
- Users' behavior : ON/OFF source model
- *MetroFlux* : a probe for traffic capture at packet level (O. Goga,. . .)

Long memory in aggregated traffic: the Taqqu model

- Heavy-tailed distributed ON periods: heavy tail index $\alpha_{ON} > 1$

Theorem (Taqqu, Willinger, Sherman, 1997)

In the limit of a large number of sources N_{src}, if:

- *flow throughput is constant,*
- *same throughput for all flows ;*

aggregated bandwidth $B^{(\Delta)}(t)$ is long range dependent, with parameter:

$$H = \max\left(\frac{3-\alpha_{ON}}{2}\ ,\ \frac{1}{2}\right)$$

Long memory: long range correlation $(H > 1/2)$

$$Cov_{B^{(\Delta)}}(\tau) = \mathbb{E}\left\{B^{(\Delta)}(t)B^{(\Delta)}(t+\tau)\right\} \underset{\tau\to\infty}{\sim} \tau^{(2H-2)}$$

Variance grows faster than Δ: $\mathbb{V}ar\left\{B^{(\Delta)}(t)\right\} \sim \Delta^{2H}$

Theorem validation on a realistic environment

- Controlled experiment: *MetroFlux* 1 Gbps, 100 sources, 8 hours traffic
- UDP/TCP: throughput limited to 5 Mbps (no congestion)

ON Distribution (source)

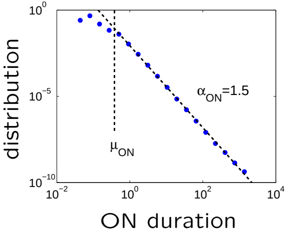

Log-diagram (aggregated traffic)

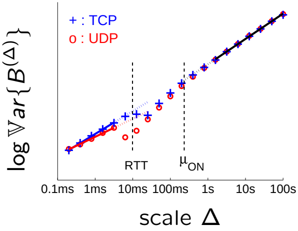

Taqqu Prediction

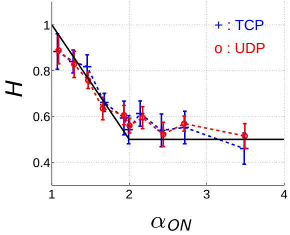

⇒ Protocol has no influence at large scales

⇒ Long memory shows up beyond scale $\Delta = \mu_{ON}$ (mean flow duration)

Influence of flow mean throughput / duration correlation

- Web traffic acquired at in2p3 (Lyon) with *MetroFlux* 10 Gbps

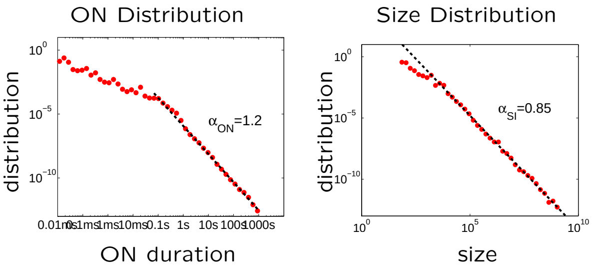

- Heavy-tailed ON periods, $\alpha_{ON} = 1.2$
- Heavy tailed flow sizes, $\alpha_{SI} = 0.85$

Influence of flow mean throughput / duration correlation

- Web traffic acquired at in2p3 (Lyon) with *MetroFlux* 10 Gbps

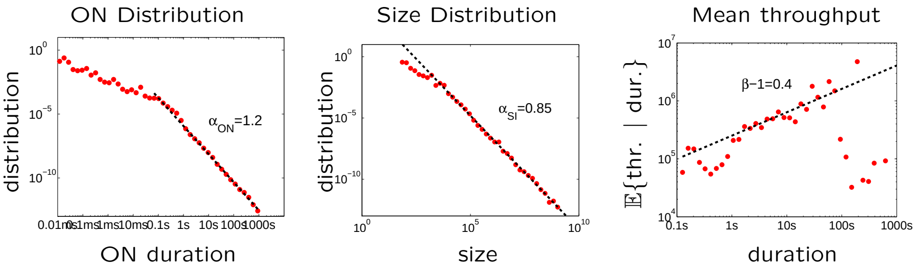

- Heavy-tailed ON periods, $\alpha_{ON} = 1.2$
- Heavy tailed flow sizes, $\alpha_{SI} = 0.85$
- Flow throughput and duration are correlated:

$$\mathbb{E}\{\text{thr.}|\text{dur.}\} \propto (\text{dur.})^{\beta-1}, \quad \beta = \alpha_{ON}/\alpha_{SI} \quad (= 1.4)$$

$\Rightarrow$ Which heavy tail index does control LRD ? $(\alpha_{ON}, \alpha_{SI})$?

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

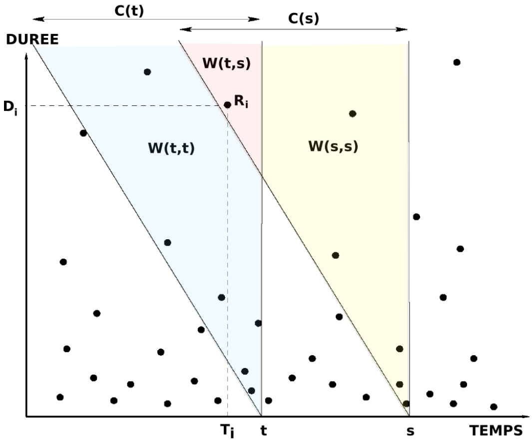

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

Proposition (LGVBP, 2009)

Model: $\mathbb{E}\{through.|dur.\} = M \cdot (dur.)^{\beta-1}$; $\mathbb{V}ar\{through.|dur.\} = V$

$$Cov_{B(\Delta)}(\tau) = CM^2\tau^{-(\alpha_{ON}-2(\beta-1))+1} + C'V\tau^{-\alpha_{ON}+1}$$

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

Proposition (LGVBP, 2009)

Model: $\mathbb{E}\{through.|dur.\} = M \cdot (dur.)^{\beta-1}$; $\mathbb{V}ar\{through.|dur.\} = V$

$$Cov_{B(\Delta)}(\tau) = CM^2\tau^{-(\alpha_{ON}-2(\beta-1))+1} + C'V\tau^{-\alpha_{ON}+1}$$

Log-diagram, $\beta > 1$

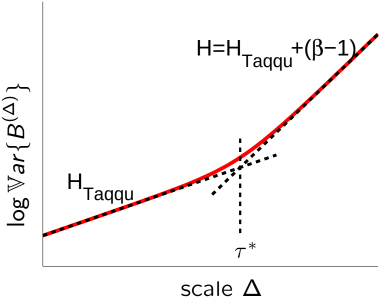

- threshold $\tau^* = \left(\frac{C'V}{CM^2}\right)^{1/(2(\beta-1))}$

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

Proposition (LGVBP, 2009)

Model: $\mathbb{E}\{through.|dur.\} = M \cdot (dur.)^{\beta-1}$; $\mathbb{V}ar\{through.|dur.\} = V$

$$Cov_{B(\Delta)}(\tau) = CM^2\tau^{-(\alpha_{ON}-2(\beta-1))+1} + C'V\tau^{-\alpha_{ON}+1}$$

Log-diagram, $\beta > 1$

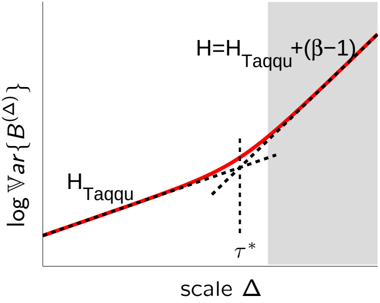

- threshold $\tau^* = \left(\frac{C'V}{CM^2}\right)^{1/(2(\beta-1))}$

  $\rightarrow$ if $\Delta \gg \tau^*$: $H = H_{\text{Taqqu}} + (\beta - 1)$

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

Proposition (LGVBP, 2009)

Model: $\mathbb{E}\{through.|dur.\} = M \cdot (dur.)^{\beta-1}$; $\mathbb{V}ar\{through.|dur.\} = V$

$$Cov_{B(\Delta)}(\tau) = CM^2\tau^{-(\alpha_{ON}-2(\beta-1))+1} + C'V\tau^{-\alpha_{ON}+1}$$

Log-diagram, $\beta > 1$

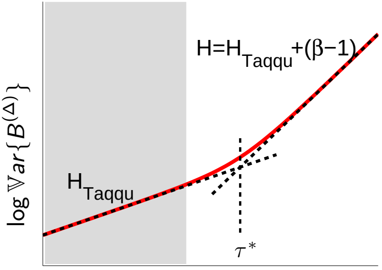

- threshold $\tau^* = \left(\frac{C'V}{CM^2}\right)^{1/(2(\beta-1))}$

  $\rightarrow$ if $\Delta \gg \tau^*$: $H = H_{\text{Taqqu}} + (\beta - 1)$

  $\rightarrow$ if $\Delta \ll \tau^*$: $H = H_{\text{Taqqu}}$

Taqqu model extension

- Planar Poisson process to describe arrival instant *vs* duration

Proposition (LGVBP, 2009)

Model: $\mathbb{E}\{through.|dur.\} = M \cdot (dur.)^{\beta-1}$; $\mathbb{V}ar\{through.|dur.\} = V$

$$Cov_{B(\Delta)}(\tau) = CM^2\tau^{-(\alpha_{ON}-2(\beta-1))+1} + C'V\tau^{-\alpha_{ON}+1}$$

Log-diagram, $\beta > 1$

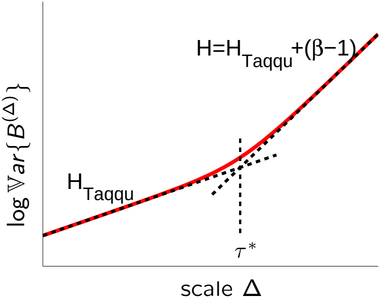

- Correlations intensify LRD ($\beta > 1$)
- Traffic evolution, future Internet: "flow-aware" control mechanisms, FTTH

LRD impact on QoS: a brief (experimental) outlook

The situation is complex. . .

LRD impact on QoS: a brief (experimental) outlook

The situation is complex. . .

- Negative on finite queues with UDP flows [cf. Mandjes, 2004 (infinite queues)]
  - LRD degrades QoS for large queue sizes (beyond some threshold)
  - but the threshold depends on the considered QoS metric (loss rate *vs* mean load)

LRD impact on QoS: a brief (experimental) outlook

The situation is complex. . .

- Negative on finite queues with UDP flows [cf. Mandjes, 2004 (infinite queues)]
  - LRD degrades QoS for large queue sizes (beyond some threshold)
  - but the threshold depends on the considered QoS metric (loss rate *vs* mean load)
- Questionable with TCP flows: [Park, 1997] against [Ben Fredj, 2001]
  - LRD has contradictory effects on QoS metrics depending on:

	with slow start	without slow start
Delay	$\searrow$	$\nearrow$
loss rate	$\searrow$	$\rightarrow$
mean throughput	$\rightarrow$	$\nearrow$

  - Heavy tailed distributions (i.e LRD) can favour QoS for large flows

LRD impact on QoS: a brief (experimental) outlook

The situation is complex. . .

- Negative on finite queues with UDP flows [cf. Mandjes, 2004 (infinite queues)]
  - LRD degrades QoS for large queue sizes (beyond some threshold)
  - but the threshold depends on the considered QoS metric (loss rate *vs* mean load)
- Questionable with TCP flows: [Park, 1997] against [Ben Fredj, 2001]
  - LRD has contradictory effects on QoS metrics depending on:

	with slow start	without slow start
Delay	$\searrow$	$\nearrow$
loss rate	$\searrow$	$\rightarrow$
mean throughput	$\rightarrow$	$\nearrow$

  - Heavy tailed distributions (i.e LRD) can favour QoS for large flows
- But in general, QOS is a complex function of multiple variables

Second level of description : single TCP source traffic

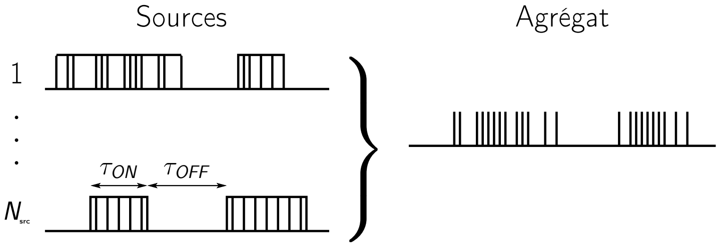

Second level of description : single TCP source traffic

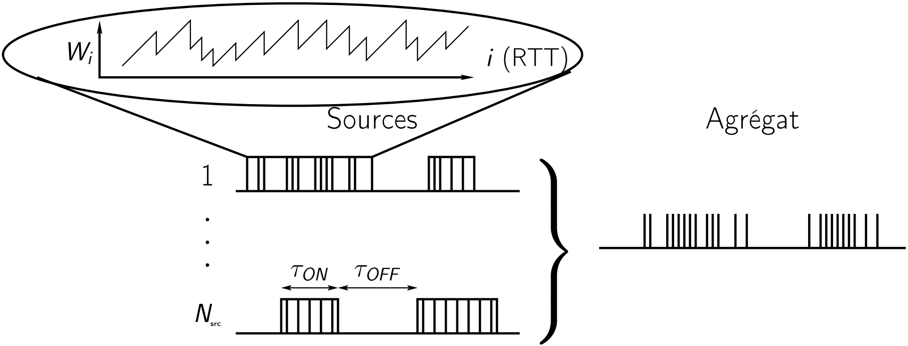

- single TCP source traffic detail
- Long-lived flow → stationary regime

⇒ How to characterize the congestion window evolution?

Markov model

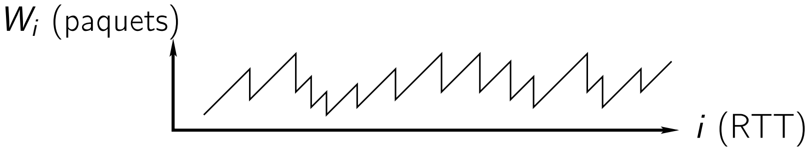

- long-lived flow stationary regime: AIMD
- model: $(W_i)_{i\geq 1}$ finite Markov chain (irreducible, aperiodic), transition matrix Q :

$$\begin{cases} Q_{w,\min(w+1,w_{\max})} &= 1-p(w), \\ Q_{w,\max(\lfloor w/2\rfloor,1)} &= p(w). \end{cases}$$

- $p(\cdot)$ loss probability of at least one packet, only depends on the current congestion window (hyp.)
- Example: [Padhye, 1998] Bernoulli loss: $p(w) = 1-(1-p_{pkt})^w$

Almost sure mean throughput

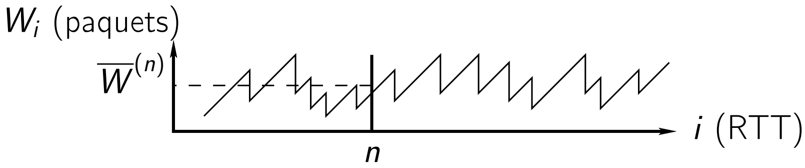

- mean throughput at scale n (RTT): $\overline{W}^{(n)} = \frac{\sum_{i=1}^{n} W_i}{n}$

Ergodic Birkhoff theorem (1931): almost sure mean

For almost all realisation, the mean throughput at scale n converges towards a value corresponding to the expectation of the invariant distribution:

$$\overline{W}^{(n)} \xrightarrow[n\to\infty]{p.s.} \overline{W}^{(\infty)} = \mathbb{E}\{W_i\}$$

- Example: [Padhye, 1998], $\overline{W}^{(\infty)} \underset{p_{pkt}\to 0}{\sim} \sqrt{\frac{3}{2p_{pkt}}}$ (RTT=1, MSS=1)

Throughput variability: Large Deviations

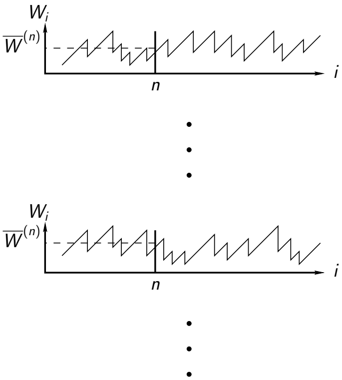

- $\overline{W}^{(n)} \simeq \alpha \neq \overline{W}^{(\infty)}$ Rare events

Large Deviations theorem (Ellis, 84)

$$\mathbb{P}(\overline{W}^{(n)} \simeq \alpha) \underset{n\to\infty}{\sim} \exp(n \cdot f(\alpha))$$

- $f(\alpha)$ Large Deviation spectrum

$\rightarrow$ Scale invariant quantity

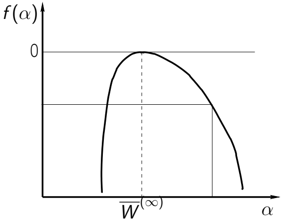

Throughput variability: Large Deviations

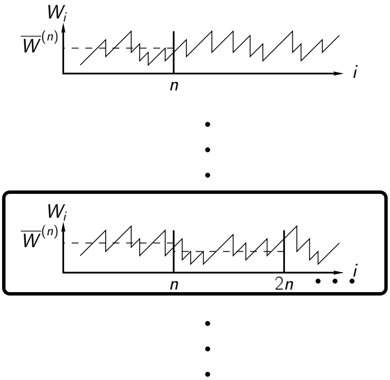

- $\overline{W}^{(n)} \simeq \alpha \neq \overline{W}^{(\infty)}$ Rare events

Large Deviations theorem (Ellis, 84)

$$\mathbb{P}(\overline{W}^{(n)} \simeq \alpha) \underset{n\to\infty}{\sim} \exp(n \cdot f(\alpha))$$

- $f(\alpha)$ Large Deviation spectrum

$\rightarrow$ Scale invariant quantity

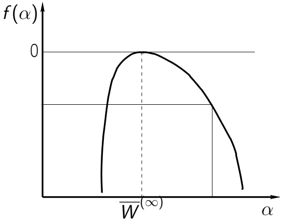

$\Rightarrow$ Does a similar theorem exist for a single realization?

Large Deviation on almost all realizations

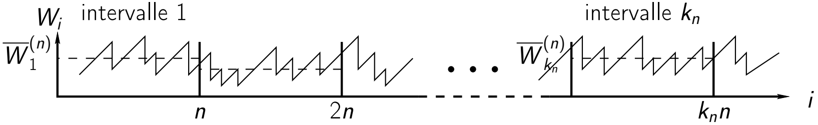

Large Deviation theorem on almost all realisations (Loiseau et al., 2010)

For a given α, if $k_n \geq e^{nR(\alpha)}$, then a.s.

$$\frac{\#\left\{j \in \{1, \cdots, k_n\} : \overline{W}_j^{(n)} \simeq \alpha\right\}}{k_n} \underset{n\to\infty}{\sim} \exp(n \cdot f(\alpha))$$

- "Price to pay": exponential increase of the number of intervals
- Finite realization (of size N): $nk_n = N$

$\Rightarrow$ $[\alpha_{\min}(n), \alpha_{\max}(n)]$ support of observable spectrum at scale n

Large Deviation on almost all realizations

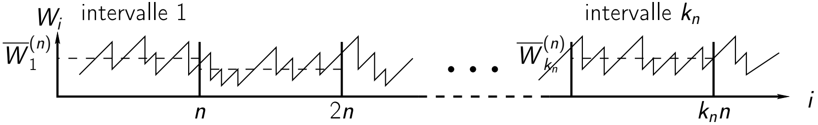

Large Deviation theorem on almost all realisations (Loiseau et al., 2010)

For a given α, if $k_n \geq e^{nR(\alpha)}$, then a.s.

$$\frac{\#\left\{j \in \{1, \cdots, k_n\} : \overline{W}_j^{(n)} \simeq \alpha\right\}}{k_n} \underset{n\to\infty}{\sim} \exp(n \cdot f(\alpha))$$

- "Price to pay": exponential increase of the number of intervals
- Finite realization (of size N): $nk_n = N$

$\Rightarrow$ $[\alpha_{\min}(n), \alpha_{\max}(n)]$ support of observable spectrum at scale n

- Theory: $p(\cdot) \to Q \to f(\alpha), R(\alpha), \alpha_{\min}, \alpha_{\max}$
- Practice: $(W_i)_{i \leq N} \to$ observed distribution

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

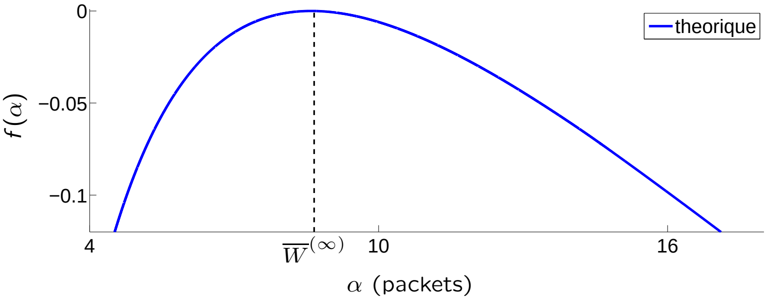

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

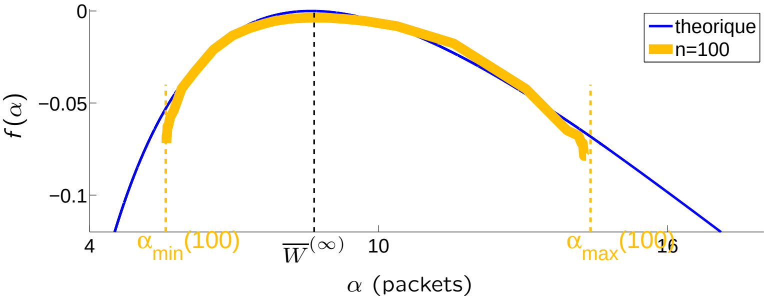

- Apex: almost sure mean: 8.6 packets (Padhye: $\sqrt{\frac{3}{2p_{pkt}}} = 8.66$)

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

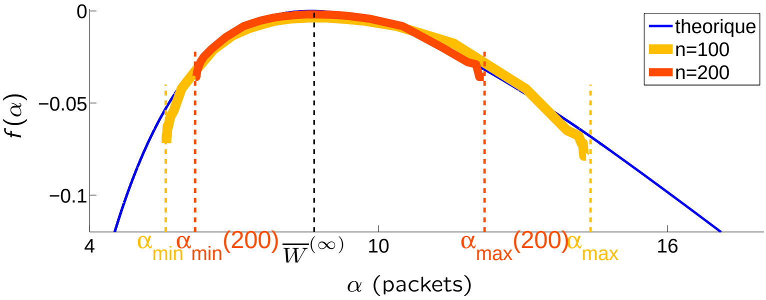

- Apex: almost sure mean: 8.6 packets (Padhye: $\sqrt{\frac{3}{2p_{pkt}}} = 8.66$)

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

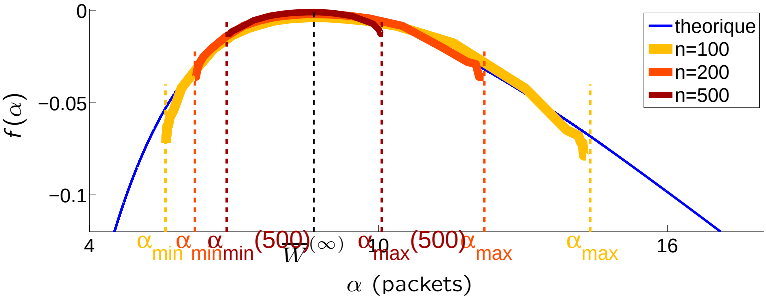

- Apex: almost sure mean: 8.6 packets (Padhye: $\sqrt{\frac{3}{2p_{pkt}}} = 8.66$)

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

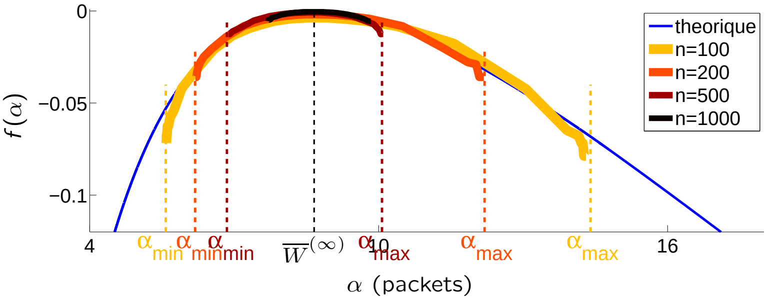

- Apex: almost sure mean: 8.6 packets (Padhye: $\sqrt{\frac{3}{2p_{pkt}}} = 8.66$)
- Superimposition at different scales → scale invariance

Results: example of Bernoulli losses ($p_{\text{pkt}} = 0.02$)

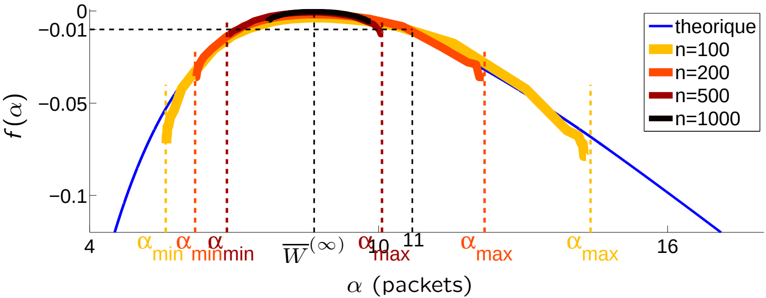

- Apex: almost sure mean: 8.6 packets (Padhye: $\sqrt{\frac{3}{2p_{pkt}}} = 8.66$)
- Superimposition at different scales $\rightarrow$ scale invariance
- beyond $n = 100$: variability

  $n = 100$, portion of intervals with mean ~ 11: $e^{-100\times 0.01} = 0.37$

  $n = 200$, portion of intervals with mean ~ 11: $e^{-200\times 0.01} = 0.14$

$\Rightarrow$ More accurate information than the almost sure mean

Results II: case of a long-lived flow

- losses: not Bernoulli
- empirical losses

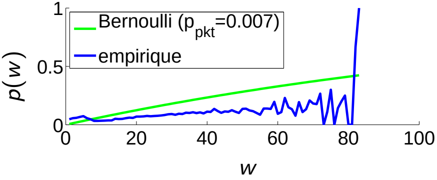

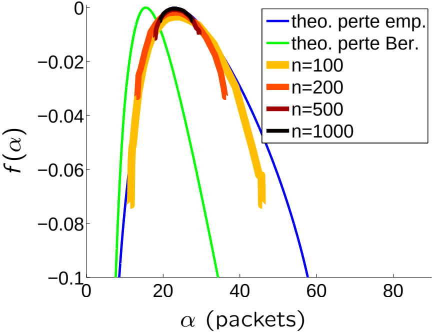

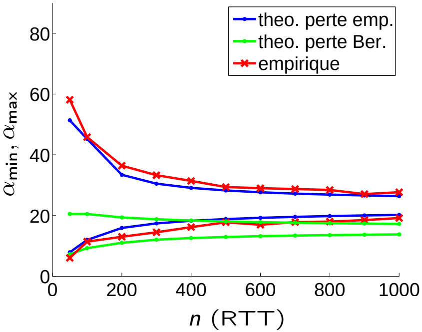

Two important assets for Large Deviations Utility

General result ("Large deviations for the local fluctuations of random walks", J. Barral, P. Loiseau, *Stochastic Processes and their Applications*, 2011)

A wide class of processes (stationary & mixing) verifies an empirical large deviation principle. In particular, this results holds true any time series that can reliably be modelled by an irreducible, aperiodic Markov process.

Theorem ("On the estimation of the Large Deviations spectrum", J. Barral, P. G., *J. stat. Phys.*, 2011)

We derived a consistent estimator of the large deviation spectrum from a finite size time series (observation samples). We proved convergence on mathematical objects with scale invariance properties (multifractal measures and processes).

Empirical estimation from a finite length trace

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

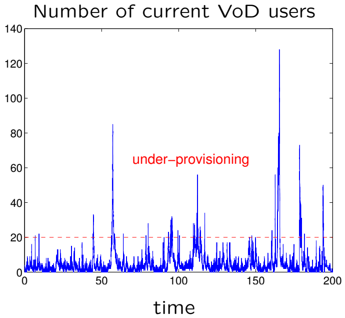

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

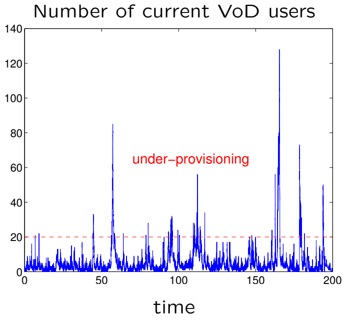

Approach Combine the three ingredients:

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

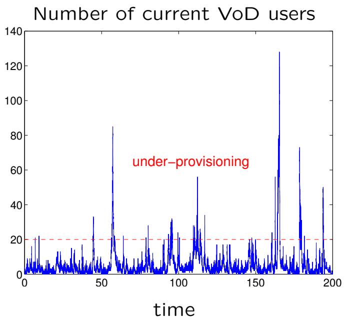

Approach Combine the three ingredients:

- A sensible (epidemic) model to catch the burstiness and the dynamics of the workload

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

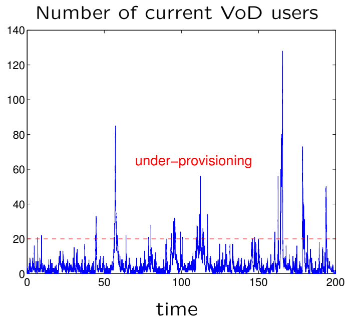

Approach Combine the three ingredients:

- A sensible (epidemic) model to catch the burstiness and the dynamics of the workload
- A (Markov) model that verifies a large deviations principle

Probabilistic Resources Allocation Based on a LDP

Workload Volatility

Context Applications that undergo highly time-varying (elastic) workloads (e.g. Buzz demand in a VoD system)

Goal Dynamic resource allocation yielding a good compromise between capex and opex costs

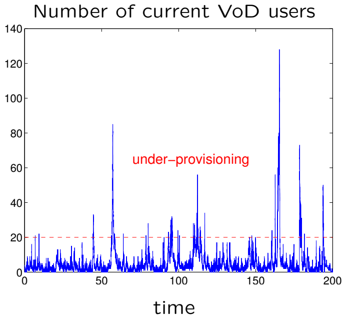

Approach Combine the three ingredients:

- A sensible (epidemic) model to catch the burstiness and the dynamics of the workload
- A (Markov) model that verifies a large deviations principle
- A probabilistic management policy based on the large deviation characterisation

An epidemic based model for volatile workload

A hidden state Markov process with memory effect [IEICE 2012, TRAC 2013]

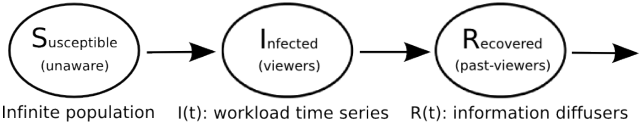

An epidemic based model for volatile workload

A hidden state Markov process with memory effect [IEICE 2012, TRAC 2013]

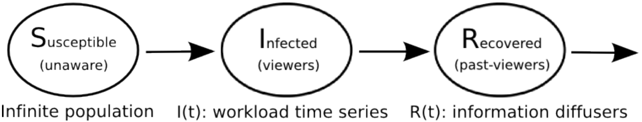

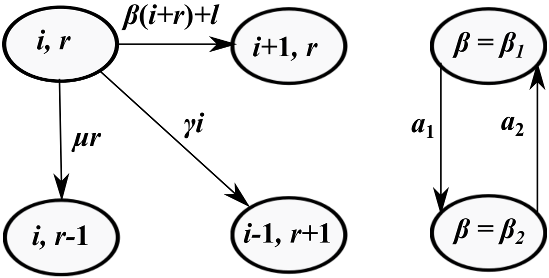

i: current # of viewers / r: current # of infected

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

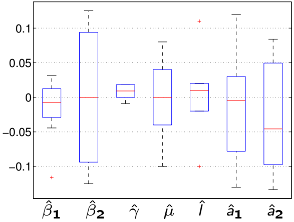

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

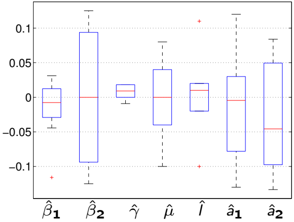

VoD workload - trace I

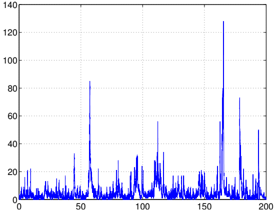

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

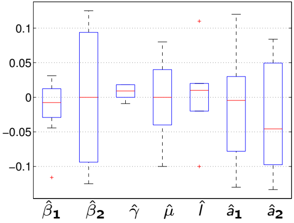

VoD workload - trace I

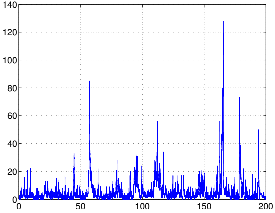

Proposed Model

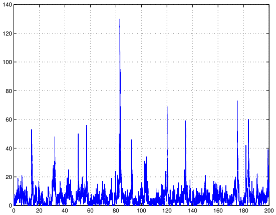

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

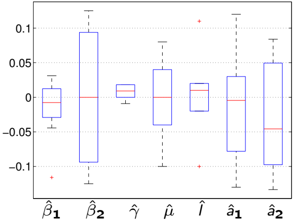

VoD workload - trace I

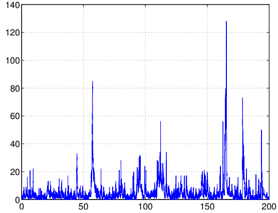

Proposed Model

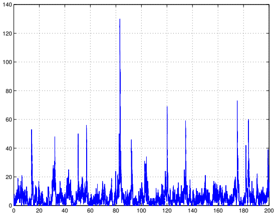

MMPP/M/1

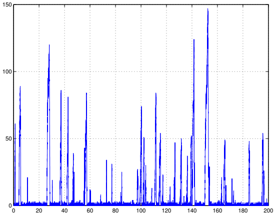

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

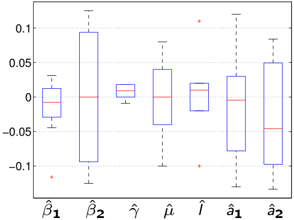

Steady state distribution

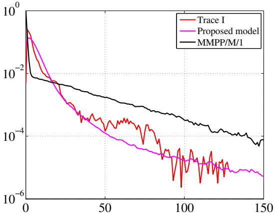

VoD workload - trace I

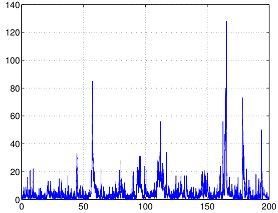

Proposed Model

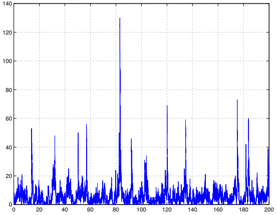

MMPP/M/1

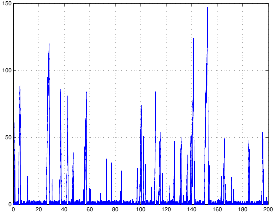

Model identification and evaluation

A MCMC based estimation procedure for the model's parameters [Gretsi 2013]

Param. estimation precision

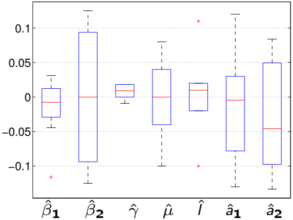

Steady state distribution

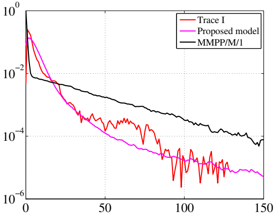

Autocorrelation function

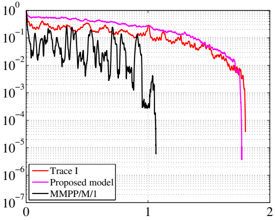

VoD workload - trace I

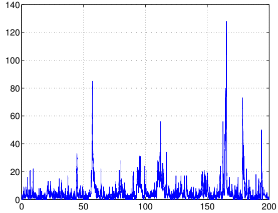

Proposed Model

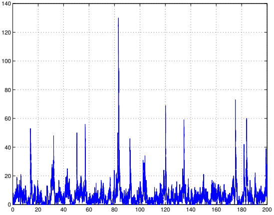

MMPP/M/1

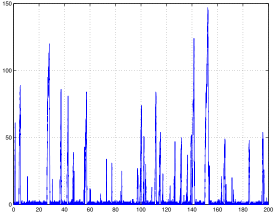

Large Deviations Principle

A process I_t verifies a large deviations principle:

$$\mathbb{P}\{\langle I_t \rangle_\tau \in [\alpha - \varepsilon_\tau, \alpha + \varepsilon_\tau]\} \sim \exp(\tau \cdot f(\alpha)), \quad \tau \to \infty$$

τ : average time scale

$f(\alpha)$: large deviations spectrum of I_t

Large Deviations Principle

A process I_t verifies a large deviations principle:

$$\mathbb{P}\{\langle I_t \rangle_\tau \in [\alpha - \varepsilon_\tau, \alpha + \varepsilon_\tau]\} \sim \exp(\tau \cdot f(\alpha)), \quad \tau \to \infty$$

τ : average time scale

$f(\alpha)$: large deviations spectrum of I_t

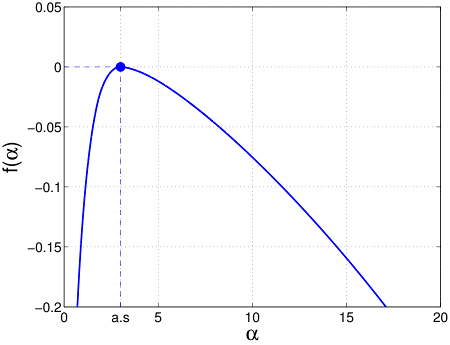

Large Deviations Principle

A process I_t verifies a large deviations principle:

$$\mathbb{P}\{\langle I_t\rangle_\tau \in [\alpha - \varepsilon_\tau, \alpha + \varepsilon_\tau]\} \sim \exp\left(\tau \cdot f(\alpha)\right), \quad \tau \to \infty$$

τ : average time scale

$f(\alpha)$: large deviations spectrum of I_t

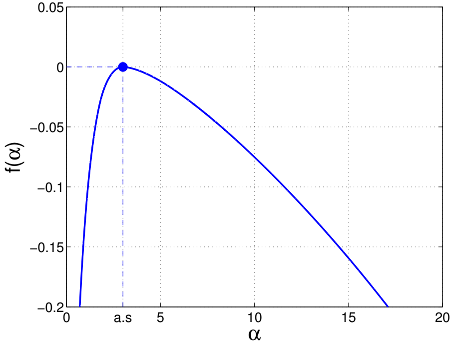

"Dynamic" implies time scale: a notion that is explicit in large deviations principle

Overflow propability

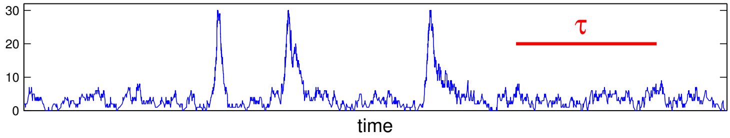

Overflow propability

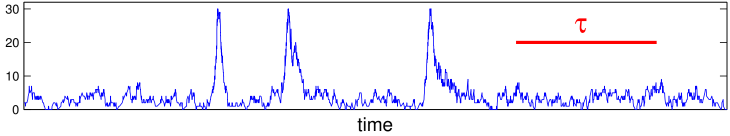

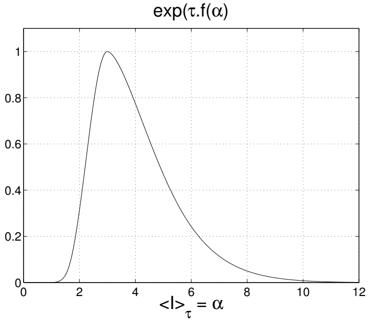

Overflow propability

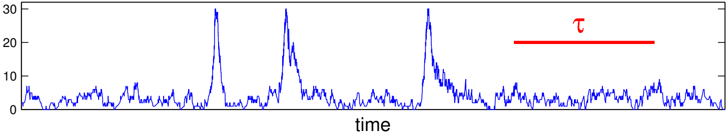

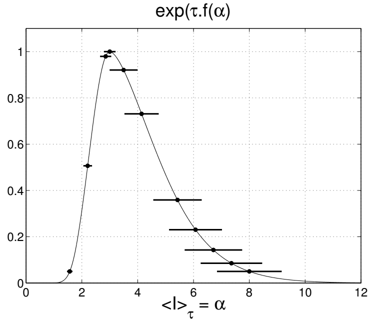

Overflow propability

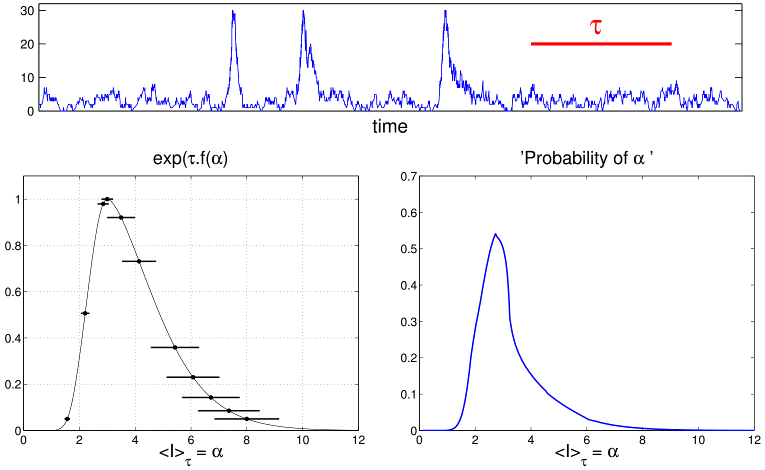

Overflow propability

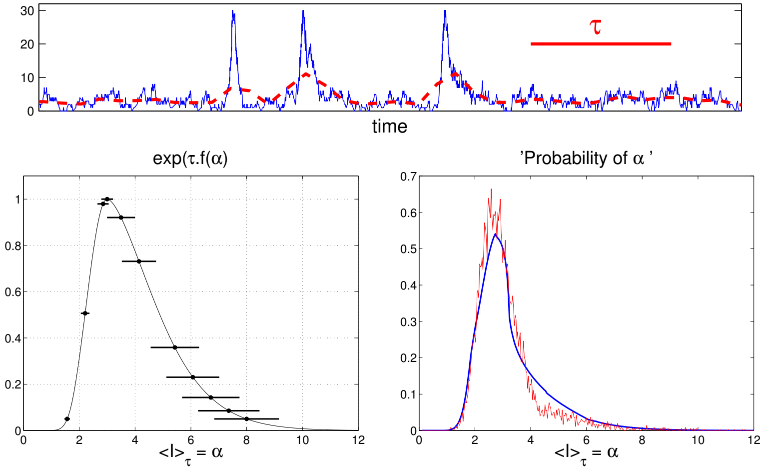

Probabilistic resource provisioning

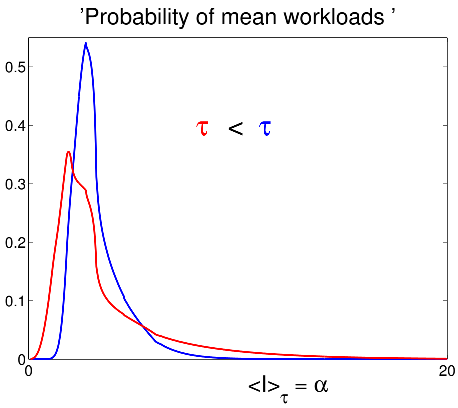

Probabilistic resource provisioning

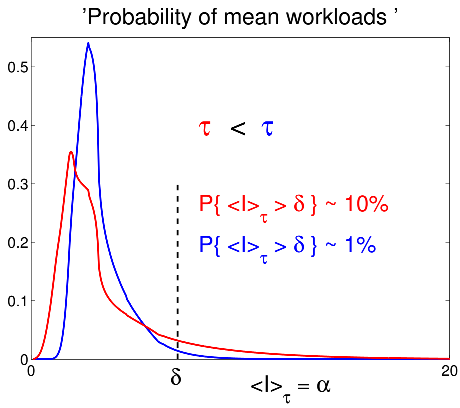

Probabilistic resource provisioning

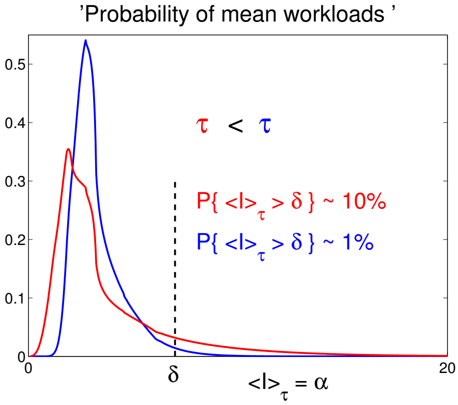

→ Resource provisioning based on a time scale dependent performance evaluation

Probabilistic resource provisioning

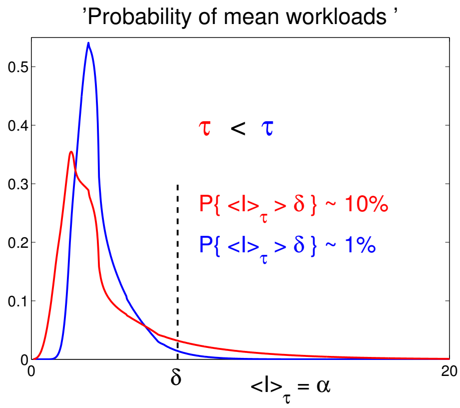

→ Resource provisioning based on a time scale dependent performance evaluation

⟶ Dynamic management

Optimal reactive time scale for reconfiguration

Optimal reactive time scale for reconfiguration

Reactivity scale for reconfiguring resource allocation is a compromise between:

- the level of congestion (or losses) yielding tolerable performance degradation
- the affordable price for a frequent reconfiguration of infrastructures

Optimal reactive time scale for reconfiguration

Reactivity scale for reconfiguring resource allocation is a compromise between:

- the level of congestion (or losses) yielding tolerable performance degradation
- the affordable price for a frequent reconfiguration of infrastructures

Assume admissible bounds for these 2 competing factors:

$\alpha^* > \alpha_{\text{a.s.}}$ beyond, it is mandatory (or profitable) reallocating resources
$\leftarrow$ capex performance concern

σ^* acceptable probability of occurrence of overflows
$\leftarrow$ opex cost

Optimal reactive time scale for reconfiguration

Reactivity scale for reconfiguring resource allocation is a compromise between:

- the level of congestion (or losses) yielding tolerable performance degradation
- the affordable price for a frequent reconfiguration of infrastructures

Assume admissible bounds for these 2 competing factors:

$\alpha^* > \alpha_{\text{a.s.}}$ beyond, it is mandatory (or profitable) reallocating resources

$\leftarrow$ capex performance concern

σ^* acceptable probability of occurrence of overflows

$\leftarrow$ opex cost

and $f(\alpha)$ is identifiable

Optimal reactive time scale for reconfiguration

Reactivity scale for reconfiguring resource allocation is a compromise between:

- the level of congestion (or losses) yielding tolerable performance degradation
- the affordable price for a frequent reconfiguration of infrastructures

Assume admissible bounds for these 2 competing factors:

$\alpha^* > \alpha_{\mathrm{a.s.}}$ beyond, it is mandatory (or profitable) reallocating resources
← capex performance concern

σ^* acceptable probability of occurrence of overflows
← opex cost

and $f(\alpha)$ is identifiable

Optimal reconfiguration time scale for dynamic resource provisioning:

$$\tau^* \; : \; \mathbf{Pr}\{\langle I \rangle_{\tau^*} \geq \alpha^*\} \approx \int_{\alpha^*}^{\infty} P_{\tau^*}(\alpha)\, \mathrm{d}\alpha > \sigma^*$$

Elastic link capacity dimensioning

Elastic link capacity dimensioning

The Service Level Agreement fixes:

- . . .
- an admissible level of losses due to link congestion

Elastic link capacity dimensioning

The Service Level Agreement fixes:

- ...
- an admissible level of losses due to link congestion

Assume $f(\alpha)$ is identifiable

Elastic link capacity dimensioning

The Service Level Agreement fixes:

- ...
- an admissible level of losses due to link congestion

Assume $f(\alpha)$ is identifiable

$C_0 = \alpha_{a.s.}$ The dedicated link capacity (nominal functioning)

Elastic link capacity dimensioning

The Service Level Agreement fixes:

- ...
- an admissible level of losses due to link congestion

Assume $f(\alpha)$ is identifiable

$C_0 = \alpha_{a.s.}$ The dedicated link capacity (nominal functioning)

$\widetilde{C}_{\tau_{\min}}$ The shared bandwidth needed to absorb bursty overflows, while guaranteeing QoS (loss rate) conformed to SLA:

$$\widetilde{C}_{\tau_{\min}} = \int_{\alpha_{\mathrm{a.s.}}}^{\infty} (\alpha - \alpha_{\mathrm{a.s.}}) P_{\tau_{\min}}(\alpha)\, \mathrm{d}\alpha$$

Elastic link capacity dimensioning

The Service Level Agreement fixes:

- ...
- an admissible level of losses due to link congestion

Assume $f(\alpha)$ is identifiable

$C_0 = \alpha_{a.s.}$ The dedicated link capacity (nominal functioning)

$\widetilde{C}_{\tau_{\min}}$ The shared bandwidth needed to absorb bursty overflows, while guaranteeing QoS (loss rate) conformed to SLA:

$$\widetilde{C}_{\tau_{\min}} = \int_{\alpha_{\mathrm{a.s.}}}^{\infty} (\alpha - \alpha_{\mathrm{a.s.}}) P_{\tau_{\min}}(\alpha)\, \mathrm{d}\alpha$$

$\tau_{\min}$ Determined by the buffer size provisioned to dampen traffic volatility

Concluding remarks

Concluding remarks

Scaling laws	Present in many (complex) systems
	Likely to become ever more ubiquitous (big data sets, heterogeneity, traffic awareness. . .)
	Impact (on performance) are still little known

Concluding remarks

Scaling laws	Present in many (complex) systems
	Likely to become ever more ubiquitous (big data sets, heterogeneity, traffic awareness. . .)
	Impact (on performance) are still little known
Large Dev. Princ.	Insufficiently exploited so far
	Holds true for a large class of modelling processes
	Takes explicitly into account the role of time scale
	Conveys information about the dynamics of the process