A Non-stationary Service Curve Model for Performance Analysis of Transient Phases in Cellular Networks

Nico Becker Markus Fidler

LEIBNIZ UNIVERSITÄT HANNOVER

FAKULTÄT FÜR ELEKTROTECHNIK UND INFORMATIK

INSTITUT FÜR KOMMUNIKATIONSTECHNIK

Transient Backlog in LTE Networks

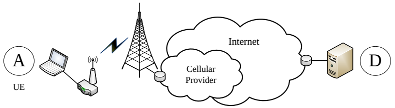

Testbed LTE

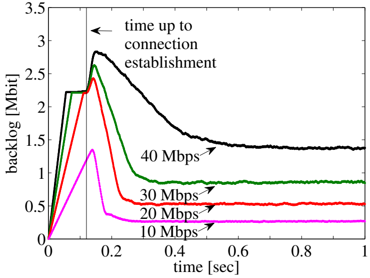

Backlog LTE

- Transient phase due to DRX mode
- In DRX mobile devices enter sleep phases to save energy
- Waking up causes additional delays
- Trade-off between energy saving and additional delay
- Relevant for safety-critical applications

Outline

Institut für Kommunikations-Technik

- Stationary vs. Non-stationary Service Curves
  - Deterministic Sleep Scheduler
  - Random Sleep Scheduler
- Measurement-Based Estimation
  - Rate Scanning
  - Burst Response
  - Minimal Probing
- Measurements in LTE

Stationarity

Consider an service process $S(t)$. The process is stationary, if

$$P[S(\tau,t) \leq x] = P[S(\tau+\delta, t+\delta) \leq x], \quad (1)$$

for any $\tau, t, \delta \geq 0$, i.e. the probability to see a certain amount of service in an interval does not depend on the time instance at which the interval starts but only on the duration of the interval.

Deterministic Sleep Scheduler

- Consider a transmitter and a receiver that if idle go to a sleep state according a defined protocol.
- Wake up is scheduled deterministically, T time units after entering sleep state.
- The transmission rate in sleep state is zero and otherwise it is R.

Backlog progression for time-variant vs. time-invariant service

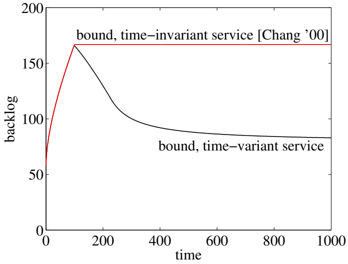

time-variant vs. time-invariant

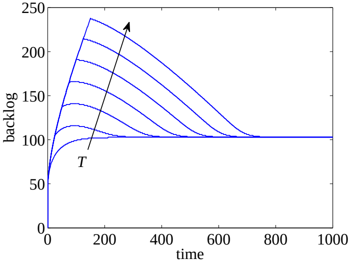

$T = \{0, 25, \ldots, 150\}$

Random Sleep Scheduler

- Consider a transmitter and a receiver that if idle go to a sleep state according a defined protocol.
- Wake up is scheduled randomly T time units after entering sleep state, i.e., T is geomatrically distributed with parameter p.
- The transmission rate in sleep state is zero and otherwise a Bernoulli increment process with parameter q.

Transient backlog and delay of random sleep scheduling

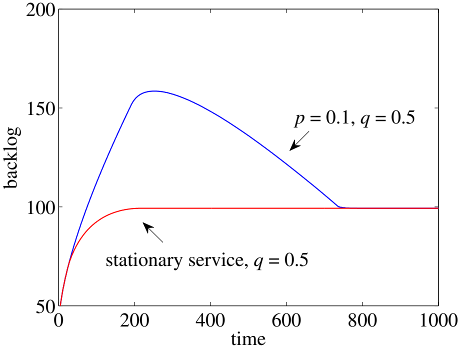

Backlog

Non-stationary Service Curves [BF '15]

Let $S(\tau, t)$ be a bivariate random service process. Then, any function $S^{\varepsilon}(\tau, t)$ that satisfies

$$P[S(\tau, t) \geq S^{\varepsilon}(\tau, t), \ \forall \tau \in [0, t]] \geq 1 - \varepsilon, \quad (2)$$

is a **non-stationary service curve** for all $t \geq 0$, where $\varepsilon \in (0, 1]$ is the underflow probability.

Non-stationary service curves of random sleep scheduling

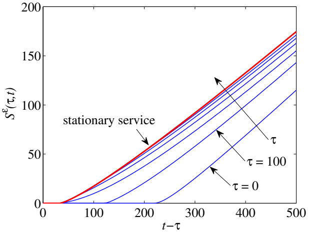

$\tau = \{0, 100, \ldots\}$
(fixing τ)

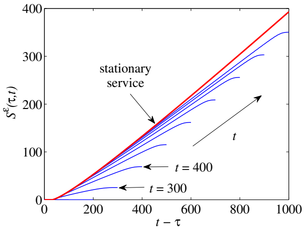

$t = \{0, 100, \ldots\}$
(fixing t)

Rate Scanning [LFV '10, LFL '14]

- Uses constant rate probes $A(t) = rt$, for a set of probes $r \in \mathbb{R}$
- $S(\tau, t) \geq \max_{r \in \mathbb{R}} \{r(t - \tau) - B(r, t)\}$
- Repeat measurements and take backlog quantile $B^{\xi}(r, t)$
- $S^{\varepsilon}(\tau, t) = \max_{r \in \mathbb{R}} \left\{r(t - \tau) - B^{\xi}(r, t)\right\}$
- $\varepsilon = \sum_{r \in \mathbb{R}} \xi$ (Union bound)

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

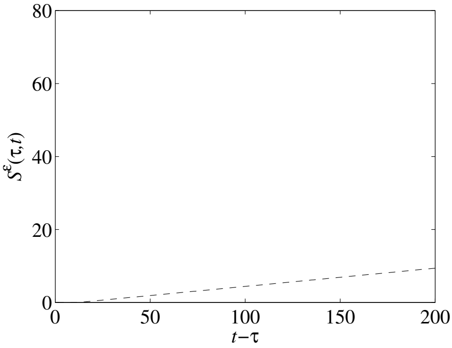

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

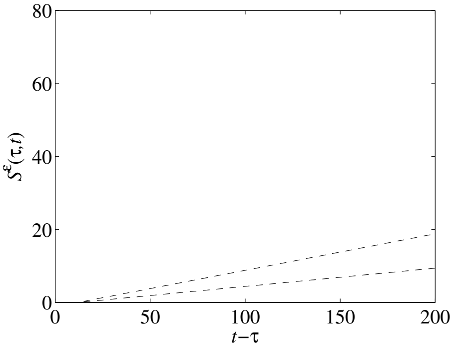

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

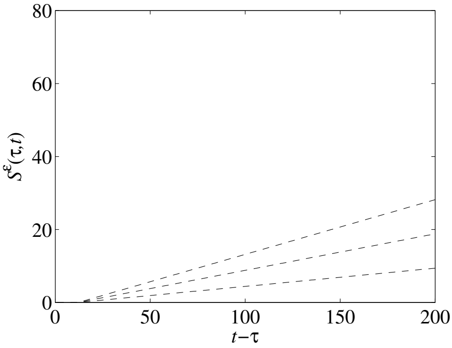

Rate Scanning [LFV ’10, LFL ’14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

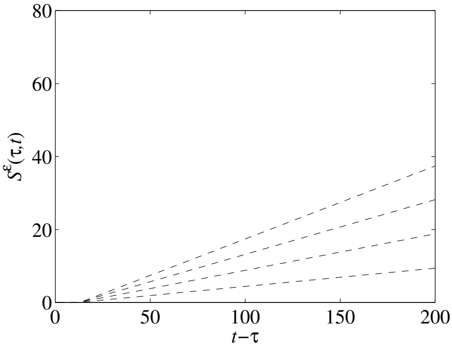

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

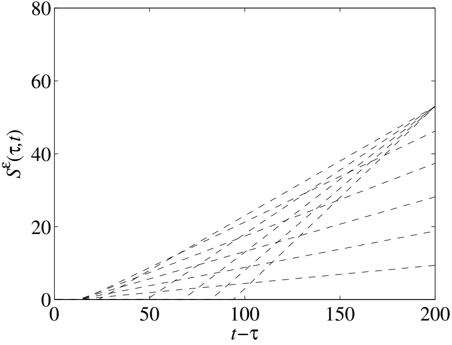

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

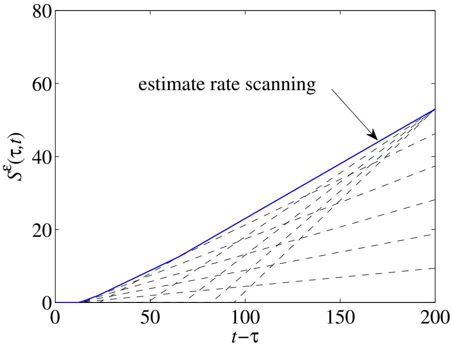

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

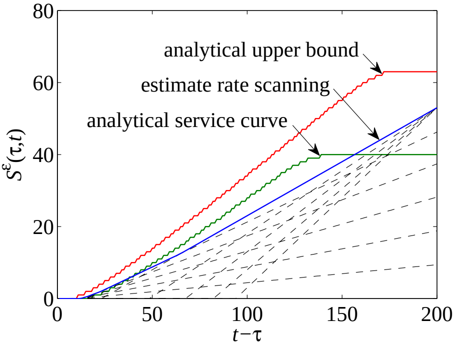

Rate Scanning [LFV '10, LFL '14]

- Example for the random sleep scheduler, with $p = 0.1$ and $q = 0.5$.
- For every rate $r \in \{0.05, 0.1, \ldots, 0.5\}$ we get 10^5 backlog samples.
- $\xi = 10^{-4}$ so that $\epsilon = \sum_{r \in \mathcal{R}} \xi = 10^{-3}$

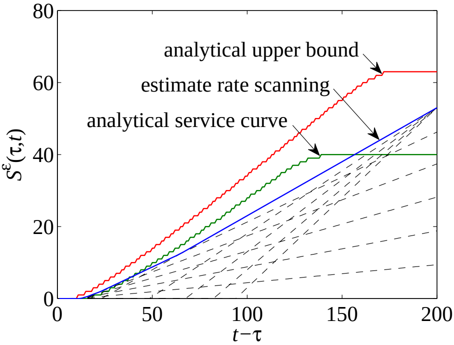

The service curve cannot recover the non-convex part of the analytical results.

Burst Response

- Uses canonical probes for system identification, i.e.,

$$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$

Burst Response

- Uses canonical probes for system identification, i.e.,

$$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$

- $D(t) = \inf_{\tau \in [0,t]} \{A(\tau) + S(\tau, t)\}$

Burst Response

- Uses canonical probes for system identification, i.e.,
$$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$
- $D(t) = \inf_{\tau \in [0,t]} \{\delta(\tau) + S(\tau, t)\}$

Burst Response

- Uses canonical probes for system identification, i.e.,
  $$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$
- $D(t) = \inf_{\tau \in [0,t]} \{\delta(\tau) + S(\tau, t)\} = S(0, t)$

Burst Response

- Uses canonical probes for system identification, i.e.,
$$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$
- $D(t) = \inf_{\tau \in [0,t]} \{\delta(\tau) + S(\tau, t)\} = S(0, t)$
- For additive service processes: $S(\tau, t) = S(0, t) - S(0, \tau)$

Burst Response

- Uses canonical probes for system identification, i.e.,
  $$A(\tau) = \delta(\tau) = \begin{cases} 0 & \text{for } \tau = 0, \\ \infty & \text{for } \tau > 0. \end{cases}$$
- $D(t) = \inf_{\tau \in [0,t]}\{\delta(\tau) + S(\tau, t)\} = S(0, t)$
- For additive service processes: $S(\tau, t) = S(0, t) - S(0, \tau)$
- Repeat measurements to get the set of all feasible sample Ω
- Remove the worst-cases and compute from the remaining set Ψ the non-stationary service curve

$$S_{br}^{\varepsilon}(\tau, t) = \inf_{\psi \in \Psi}\{S_{\psi}(\tau, t)\} \tag{3}$$

Service curve estimates

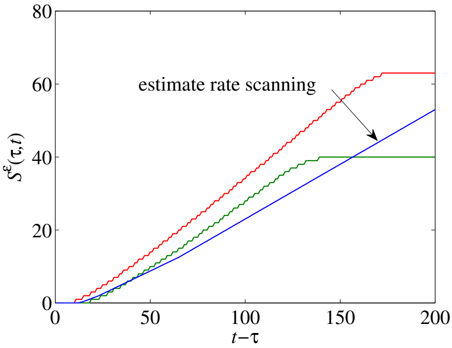

Service curve estimates

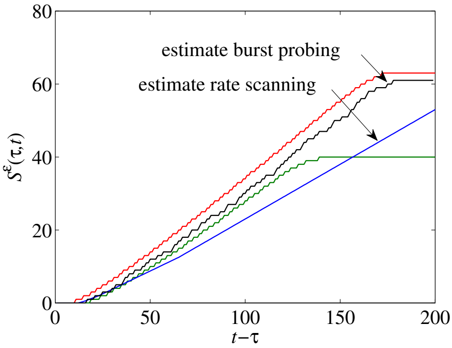

- Burst probes cause non-linear behavior of certain systems.
- Preempt other traffic, resulting in a too optimistic service estimate.

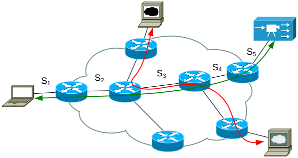

For additive and univariate service S^i $(i = 1, 2, \ldots n)$:
$S^{net}(\tau, t) = S^1 \otimes S^2 \otimes \cdots \otimes S^n(\tau, t)$

Lemma (Super-additivity of $\otimes$)
Given two bivariate functions $f(s,t)$ and $g(s,t)$ for $t \geq s \geq 0$ where $f(t,t), g(t,t) = 0$ for all $t \geq 0$. Define $h(s,t) = f \otimes g(s,t)$.

If f and g are super-additive, then h is super-additive.

Super-additivity of $\otimes$

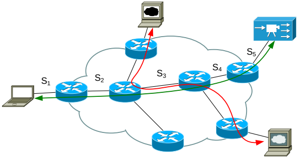

For additive and univariate service S^i $(i = 1, 2, \dots n)$:
$S^{net}(\tau, t) = S^1 \otimes S^2 \otimes \dots \otimes S^n(\tau, t)$

Lemma (Super-additivity of $\otimes$)
Given two bivariate functions $f(s,t)$ and $g(s,t)$ for $t \geq s \geq 0$ where $f(t,t), g(t,t) = 0$ for all $t \geq 0$. Define $h(s,t) = f \otimes g(s,t)$.

If f and g are super-additive, then h is super-additive.

$$\Rightarrow S^{net}(\tau, t) \leq S^{net}(0, t) - S^{net}(0, \tau)$$

Note, that it includes additive processes, as well!

Minimal Probing

- We seek to find the minimal probe that satisfies for a fixed t $D(t) = \inf_{\tau \in [0,t]}\{A_{mp}(\tau) + S(\tau, t)\} = S(0, t)$,
- i.e., the minimal probe that allows estimating the service from observations of the departures.
- The minimal probe is $A_{mp}(\tau) = S(0, t) - S(\tau, t)$.
- For any other larger or smaller probe it leads to a lower service.
- We do not know $S(\tau, t)$ in advance

1. Compute the burst response estimate , i.e., $S^{\varepsilon}_{br}(\tau, t)$
2. Use $S^{\varepsilon}_{br}(\tau, t)$ to compute the minimal probe, i.e.,

$$\tilde{A}_{mp}(\tau) = S^{\varepsilon}_{br}(0, t) - S^{\varepsilon}_{br}(\tau, t) \tag{4}$$

and repeat the measurements to get the service for the minimal probe, $S^{\varepsilon}_{mp}(\tau, t)$.

Accuracy of Minimal Probing

For $\widetilde{A}_{mp}(\tau) = \mathcal{S}^{\varepsilon}_{br}(0,t) - \mathcal{S}^{\varepsilon}_{br}(\tau,t)$ we conclude that $B^{\varepsilon}(t)$ observed by minimal probing is a measure of accuracy that separates the conservative estimate of minimal probing from the possibly too optimistic estimate of burst probing, i.e.,

$$S^{\varepsilon}_{mp}(\tau,t) = S^{\varepsilon}_{br}(\tau,t) - B^{\varepsilon}(t) \qquad (5)$$

LTE Service Curve Estimates

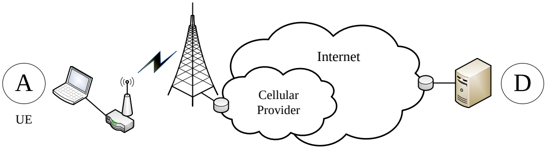

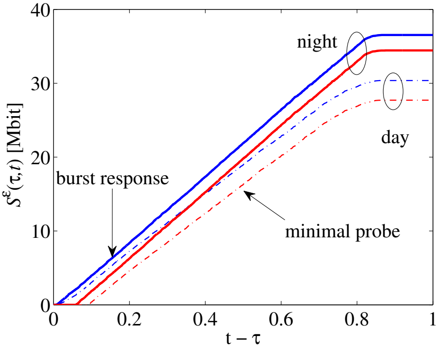

Conclusion

- Analysis of non-stationary service curves
- Evaluated the effect on transient phases (also in comparison to stationary service curves)
- Devised a novel two-phase method to obtain an accurate service curve estimate
- Simulation results confirmed the fidelity of the approach
- Measurements in LTE show that the method is applicable in practice.

Related Work

[BF '15] Becker, Fidler. : A Non-stationary Service Curve Model for Performance Analysis of Transient Phases.
in *Proc. of ITC 27*, Sep 2015

[Chang '00] C.-S. Chang. : Performance Guarantees in Communication Networks. *Springer-Verlag*, 2000

[LFL '14] Lübben, Fidler, Liebeherr. : Stochastic Bandwidth Estimation in Networks With Random Service.
IEEE/ACM *Trans. Netw.*, vol. 22, no.2, pp. 484-497, Apr 2014

[LFV '10] Liebeherr, Fidler, Valaee. : A System Theoretic Approach to Bandwidth Estimation.
IEEE/ACM *Trans. Netw.*, vol. 18, no.4, pp. 1040-1053, Aug 2010

Minimal probing

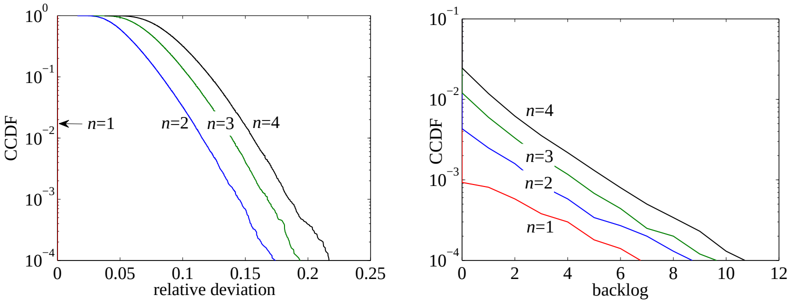

(a) Deviation from additivity (b) Accuracy of minimal probing

Figure: Network of n systems with random sleep scheduling in series. (a) The network service process deviates from additivity. (b) Minimal probing achieves small backlogs, corresponding to a high accuracy of the estimate.

Consider an service process $S(t)$. The process is stationary, if

$$P[S(\tau, t) \leq x] = P[S(\tau + \delta, t + \delta) \leq x], \quad (6)$$

for any $\tau, t, \delta \geq 0$, i.e. the probability to see a certain amount of service in an interval does not depend on the time instance at which the interval starts but only on the duration of the interval.

Let $S(\tau, t)$ be a bivariate random service process. Then,

i. any function $S^{\varepsilon}(t)$ that satisfies

$$P[S(\tau, t) \geq S^{\varepsilon}(t-\tau),\ \forall \tau \in [0, t]] \geq 1-\varepsilon, \tag{7}$$

is an **ε-effective service curve**

ii. any function $S^{\varepsilon}(\tau, t)$ that satisfies

$$P[S(\tau, t) \geq S^{\varepsilon}(\tau, t),\ \forall \tau \in [0, t]] \geq 1-\varepsilon, \tag{8}$$

is a **non-stationary service curve**

for all $t \geq 0$, where $\varepsilon \in (0, 1]$ is the underflow probability.

Sample path selection

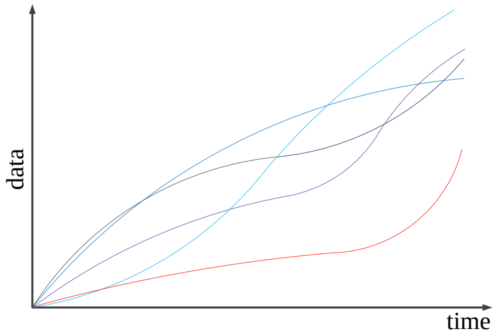

Sample path selection

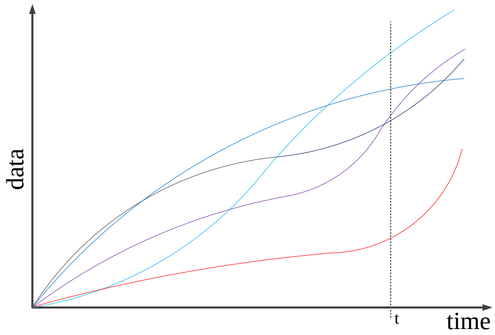

Sample path selection

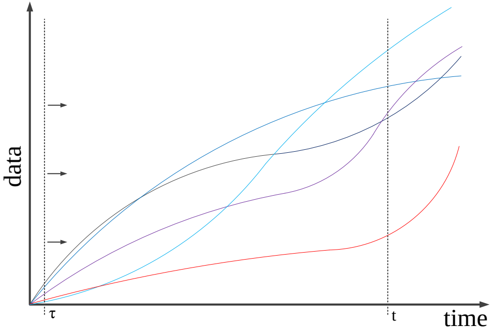

Sample path selection

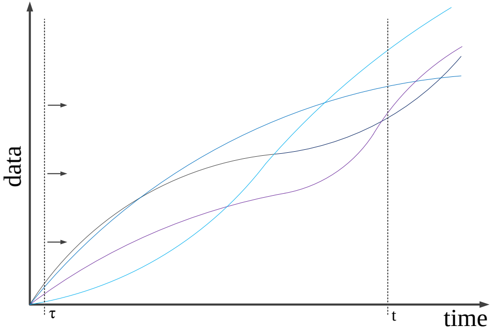

$$S_{br}^{\varepsilon}(\tau, t) = \inf_{\psi \in \Psi_t} \{S_{\psi}(\tau, t)\}$$