DIONYSOS
(Dependability, Interoperability and perfOrmaNce analYsiS of networkS)

INRIA Rennes, France

Paris, March 2012

Outline

— 1 — Composition of the team —

— 2 — Overview —

— 3 — First focus on perceptual quality —

— 4 — Second focus on network economics —

— 5 — Third focus on Monte Carlo techniques —

— 6 — Future work —

Members, 2012

Researchers

- G. Rubino, DR, leader
- B. Sericola, DR
- B. Tuffin, CR
- (N. Bouabdallah, CR, on leave from Jan 2011 until 2015)
- C. Viho, Prof., U. of Rennes 1
- A. Ksentini, Ass. Prof., U. of Rennes 1
- Y. Hadjadj, Ass. Prof., U. of Rennes 1
- R. Marie, Prof. Emeritus, U. of Rennes 1
- A. Baire, Eng., U. of Rennes 1

Non permanent members

- 10 PhD students
- 2 temporary engineers
- 1 postdoc

Members, 2012

Researchers

- G. Rubino: QoE, Monte Carlo
- B. Sericola: distrib. sys. perf., fluid queues
- B. Tuffin: Monte Carlo, network economics
- (N. Bouabdallah: wireless, sensors)
- C. Viho: QoE-based networking, interop.
- A. Ksentini: wireless, QoE
- Y. Hadjadj: wireless, QoE
- R. Marie: perf. and depend. eval.
- A. Baire: interop.

Non permanent members

- 10 PhD students
- 2 temporary engineers
- 1 postdoc

Overview of Dionysos's work

Main research directions

- The PSQA project: automatic perceived quality assessment
- Network Economics: rules for better interactions between actors
- Optimizing resource usage in wireless (e.g. sensor) and mobile networks
- Mathematical tools in simulation, fluid models, performability models

Some specificities

- **Perceptual** quality automatic measure
- **Dependability assessment** through probabilistic models
- Rare event analysis, wide spectrum: Monte Carlo and **Quasi-Monte Carlo**, dynamic and **static** models, **four families of methods**
- Qualitative property analyzed: **network components' interoperability**
- Strong collaborative work, including **two associated teams and the MAPI group**

Overview of Dionysos's work

Main research directions

- The PSQA project: automatic perceived quality assessment
- Network Economics: rules for better interactions between actors
- Optimizing resource use in wireless (e.g. sensor) and mobile networks
- Novel mathematical tools in simulation, fluid models, performability models

Some specificities

- **Perceptual** quality automatic measure
- **Dependability assessment** through probabilistic models
- Rare event analysis, wide spectrum: Monte Carlo and **Quasi-Monte Carlo**, dynamic and **static** models, **four families of methods**
- Qualitative property analyzed: **network components' interoperability**
- Strong collaborative work, including **two associated teams and the MAPI group**

The PSQA technique

PSQA in a few lines

- PSQA: Pseudo-Subjective Quality Assessment
- Based on statistical learning
- Learning tool: Random Neural Networks (open queueing network with “signed” customers)
- PSQA maps QoS + application parameters into quality

- **Accurate** (“optimally”)
- **No reference** needed
- **Efficient**: real-time if necessary/useful
- The mapping has **nice mathematical properties** (particular class of rational function) with many possible uses

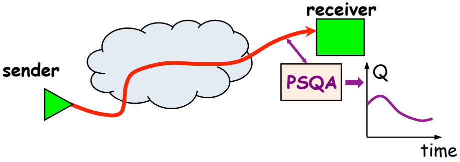

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

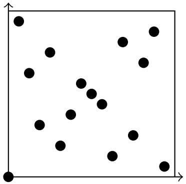

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

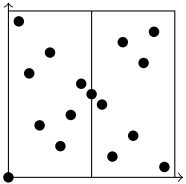

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

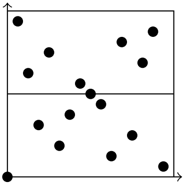

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

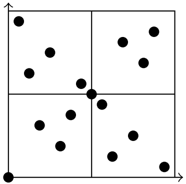

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

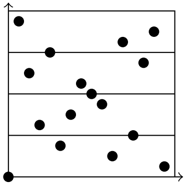

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

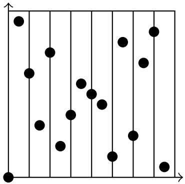

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

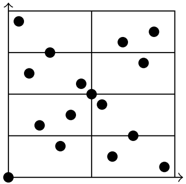

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

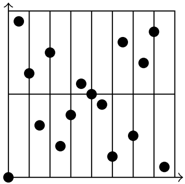

Some details

Procedure

- Choose QoS and application indicators; call them x
- Build many video/audio/voice sequences where x values are carefully sampled
- Perform statistical tests to evaluate quality Q of sequences
- Use statistical learning to build a function $Q = f(x)$
- Then, in operation, measure x, then call function $f()$ and return $f(x) = Q$

For sampling, we mix Monte Carlo and Quasi-Monte Carlo, that is, we use pseudo-random numbers and low-discrepancy numbers:

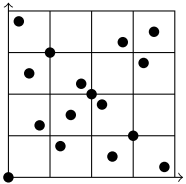

A research prototype

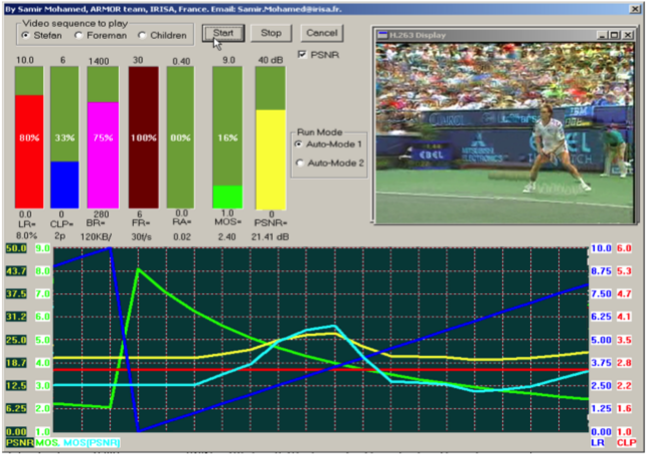

PSQA in the period

- Developed for different cases, in different projects:
  - DVB-H (*QoSmobile*)
  - SVC - DVB-T2 (*SVC4QOE*)
  - P2P video (*P2Pim@ges*)
  - Video streaming with HTTP (*ViPEER*)
- Monitoring system under development: NQAS (Network Quality Assessment System)
- Use of PSQA for
  - choosing the access network (1 PhD)
  - designing multi-path video streaming system (1 PhD)
- Complete PSQA-based design of a P2P TV network (1 PhD)
  - Validated on PlanetLab
  - Implemented by an operator abroad (Uruguay)
  - Multi-source approach
  - Control variates: frame-type-dependent load and redundancy level (among paths)
  - Almost no-signaling, by means of a non-standard use of pseudo-random number generators
  - First version of NQAS used
  - Related modeling work for dimensioning

NQAS: Network Quality Assessment System

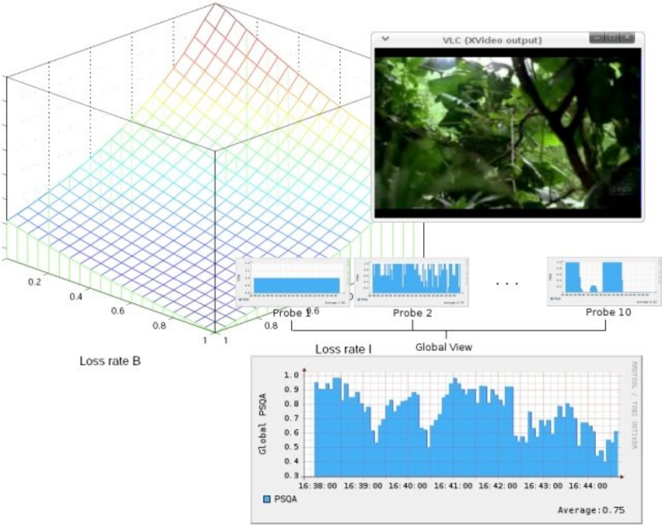

[1]

[1]Snapshot from our twice awarded prototype

Usage examples

Video quality analysis (P2P)

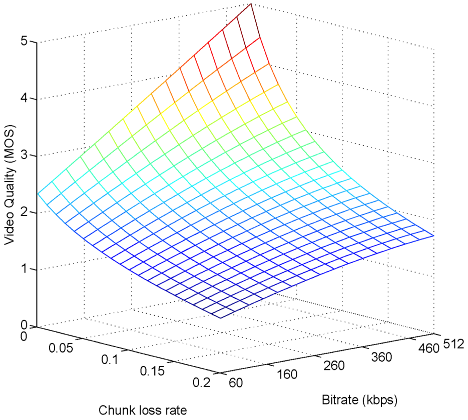

Analytical developments in VoIP applications

- Consider a bottleneck modeled as a $M/M/1/N$ with load $\rho \neq 1$.
- Classical approach: choose N and ρ such that $p_L < \varepsilon$, using

$$p_L = \frac{1-\rho}{1-\rho^{N+1}}\rho^N$$

- PSQA-based approach: choose N and ρ such that *the perceived quality* Q satisfies $Q \geq Q_0$, using

$$Q = \frac{13.3 + 2.01\rho + 6.74\rho^N - 20.0\rho^{N+1} - 2.01\rho^{N+2}}{16.6 + 3.99\rho + 93.3\rho^N - 110\rho^{N+1} - 3.99\rho^{N+2}}$$

From pricing to network economics

- Situations with more than one player, in different contexts: (i) ISPs inside a single domain; (ii) WiFi against WiMax; (iii) sharing the radio spectrum.
- In (i) and (ii), price of anarchy = 1
- In (iii), different interests on licensing or not the bandwidth

Sponsored links of search engines (SEs): auctions used.

Some results:

- interest in adding randomness to the assignment rule
- comparisons between revenue-based and bid-based rules
- competition between SEs

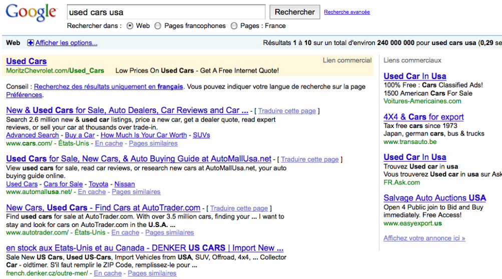

- Inter-domain: transporting traffic from/to other domains
  - problem: find appropriate incentive mechanisms
  - result: no perfect mechanism possible
    - individual rationality
    - incentive compatibility
    - budget balance
    - efficiency
  - new problem: the design of "as good as possible" mechanisms
- Networks economics, utility functions and PSQA
  - most models rely on appropriate utility functions (sometimes defined by the properties satisfied)
  - PSQA allows to work with the "ultimate target", user's perception of the quality provided
  - problem: build economic solutions based on quality, not on mathematical utility functions

Rare event analysis

- Application domain: dependability properties
- Models: **static (network reliability)** and dynamic (e.g. Markov)
- In the latter case, the analyzed metric can be transient or considered for the system in equilibrium
- We work on four different approaches:
  - Importance Sampling, and in particular, **Zero-Variance techniques**
  - Splitting
  - **recursive factoring-based procedures**
  - **combined Monte Carlo + Quasi-Monte Carlo methods**
- We also develop a research axis on **theoretical properties of estimators when the target is rare**

Rare event analysis

- Application domain: dependability properties
- Models: **static (network reliability)** and dynamic (e.g. Markov)
- In the latter case, the analyzed metric can be transient or considered for the system in equilibrium
- We work on four different approaches:
  - Importance Sampling, and in particular, **Zero-Variance techniques**
  - Splitting
  - **recursive factoring-based procedures**
  - **combined Monte Carlo + Quasi-Monte Carlo methods**
- We also develop a research axis on **theoretical properties of estimators when the target is rare**

Zoom: Zero-Variance approaches

- Importance Sampling:

$$\mathbb{E}[X] = \int x \, \mathrm{d}\mathbb{P}(x) = \int x \frac{\mathrm{d}\mathbb{P}(x)}{\mathrm{d}\tilde{\mathbb{P}}(x)} \, \mathrm{d}\tilde{\mathbb{P}}(x) = \tilde{\mathbb{E}}\left[X \frac{\mathrm{d}\mathbb{P}(x)}{\mathrm{d}\tilde{\mathbb{P}}(x)}\right].$$

- Consider a discrete-time Markov chain $Y = \{Y_j, j \geq 0\}$, with state space $\mathcal{Y}$ (discrete, possibly infinite and high-dimensional), evolving up to a stopping time $\tau < \infty$, and driven by transition matrix P and initial probability $\pi_0(y_0) = \mathbb{P}[Y_0 = y_0]$.
- Each transition from y to z brings a reward $c(y, z)$. We want to compute $\mathbb{E}[X]$ with

$$X = \sum_{i=1}^{\tau} c(Y_{i-1}, Y_i).$$

- Let $\gamma(x) = \mathbb{E}_x[X]$.

- If we consider the change of measure that makes use of the new transition matrix $\tilde{P}$ defined by

$$\tilde{P}_{y,z} = \frac{P_{y,z}\big[c(y,z) + \gamma(z)\big]}{\sum_w P_{y,w}\big[c(y,w) + \gamma(w)\big]},$$

  it can be shown that the resulting estimator has zero variance.
- Requires to know $\gamma(x)$! ... but can be approached by *some* $\hat{\gamma}(x)$ (or can even be learned).
- Current applications on dependability modeling and analysis:
  - MTTF estimation for Markovian models. Much more stable (and accurate) results than with existing heuristics, if we replace $\hat{\gamma}$ by probabilities (or adjusted probabilities) of direct paths to failure.
  - In the case of network reliability estimation: used by sampling links one after the other; there, $\hat{\gamma}$ is taken as the probability of specific mincuts.

Zoom: theoretical analysis of estimators

- Use a *rarity* parameter ε
- Defined and studied properties of estimators $f(X)$ of some target γ, functions of sample X:
  - Bounded Relative Error, BRE (as $\varepsilon \downarrow 0$)
  - Logarithmic Efficiency, or Asymptotic Optimality (as $\varepsilon \downarrow 0$)
  - **Work-Normalized versions** of preceding concepts
  - **Bounded Normal Approximation** (as $\varepsilon \downarrow 0$)
  - **Vanishing Variance** (as $\varepsilon \downarrow 0$)
  - **Coverage properties**
  - More complex situations: e.g., **BRE doesn't hold but "it is numerically observed"**

See all Part I of

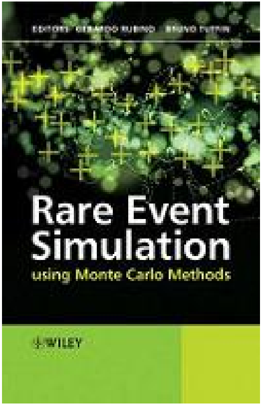

Next 4 years: main points

- perceptual quality
  - PSQA "2.0"
    - "enriched" RNN models
    - exploiting the math properties of function Q: ∇Q, Q^{-1}
    - avoiding human panels in the learning process using signal-based tools
    - adding message meaning (for voice)
  - perceptual quality $\longrightarrow$ QoE
    - assessing perceptual quality for Web service/applications in general
    - control issues (as for video/audio quality)
  - industrial transfer
    - with *Perceptiva Labs*
    - first project around NQAS
- Monte Carlo techniques
  - rare event analysis methods
    - mixed methods (e.g. around Zero-Variance approaches)
    - compare/combine with bounding techniques
    - the dependent case

- network economics
    - when all the actors are represented
    - neutrality issues (among others): which trade-off is possible?
    - specific challenge: replace utility functions with PSQA functions

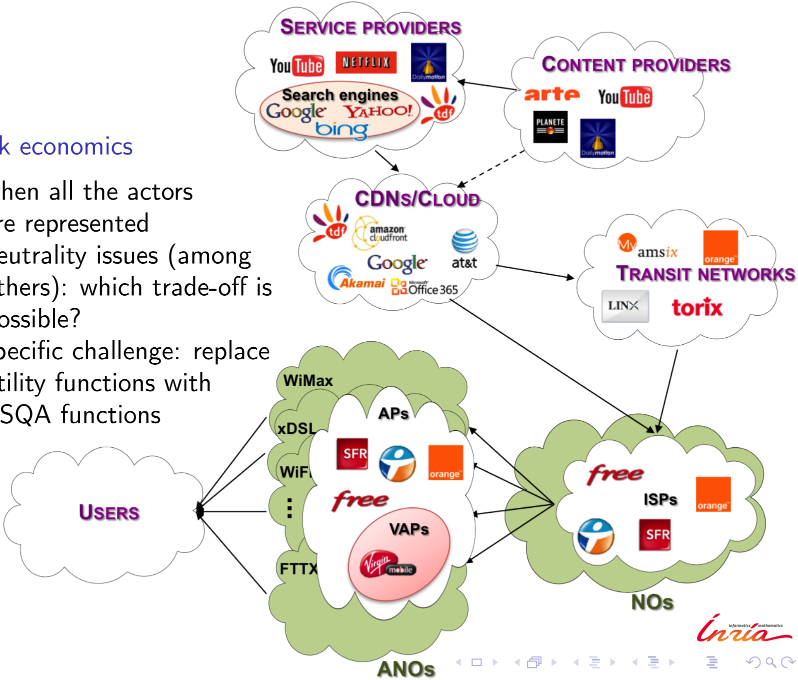

- inside MAPI group
  - cluster-based large scale distributed systems
  - one cluster model $\rightsquigarrow$ large models
  - n clusters competing for ressources $\rightsquigarrow$?
    - branching processes techniques...?
    - the case of $n \longrightarrow \infty$?
- on the MAPI group itself
  - first "scientific service" at INRIA?
- bounding techniques
  - in probabilistic model checking
    - new transient performability metrics
    - bounding techniques to deal with the state explosion problem
  - in rare event analysis
    - linear algebra-based approaches
    - target: steady-state measures and MTTF evaluation)
    - comparison/mixing with Monte Carlo

- Content Centric Networks (CCN)
  - what?
    - overlay solution (Gnutella-like)
    - in the intra-domain
    - for adaptive video streaming (DASH)
    - **using QoE indicators**
  - why?
    - current technological evolution, new usages, digital switchover
    - our previous work on P2P and on CDNs, for video delivery
    - keeping the current network architecture
    - **exploiting our results on QoE** (e.g., adding indicators to the chunks' names)
  - challenges:
    - replication and caching strategies
    - congestion avoidance/control (in particular, optimizing signaling overhead)
    - **QoE-aware mechanisms usage**
    - cache selection algorithms
    - **economical models**