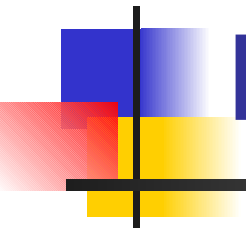

Numerical vs. Statistical Probabilistic Model Checking

Håkan L. S. Younes
Carnegie Mellon University

Marta Kwiatkowska Gethin Norman David Parker
University of Birmingham

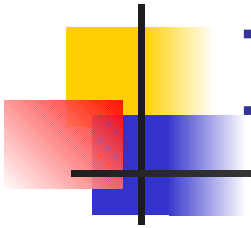

Introduction

- Model checking of stochastic systems
  - Continuous-time Markov chains
  - Continuous Stochastic Logic (CSL)
    - Probabilistic time-bounded properties
- Comparison of two techniques
  - Numerical computation of probabilities
  - Statistical hypothesis testing

Example: Tandem Queuing Network

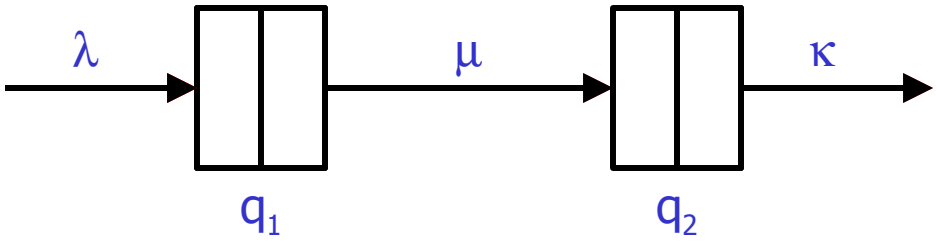

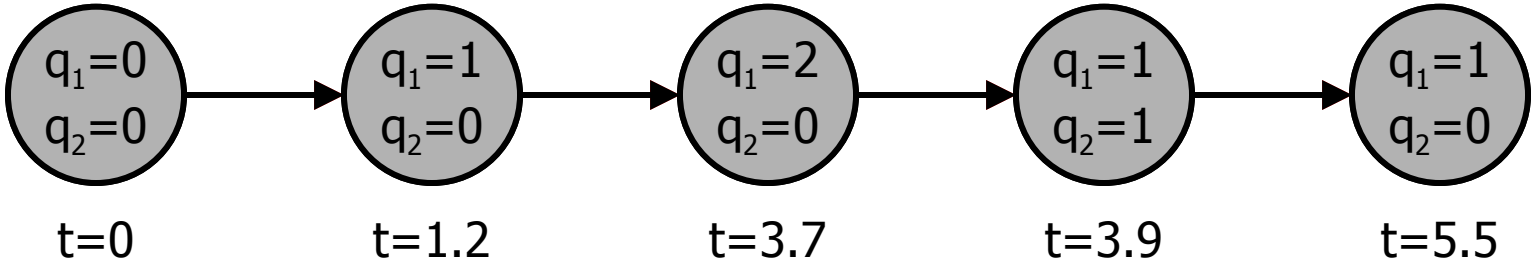

With both queues empty, is the probability less than 0.5 that both queues become full within 5 seconds?

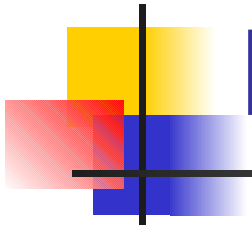

Probabilistic Model Checking

- Given a model M, a state s, and a property φ, does φ hold in s for M?
  - Model: continuous-time Markov Chain
  - Property: Continuous Stochastic Logic (CSL) formula

Continuous-Time Markov Chain

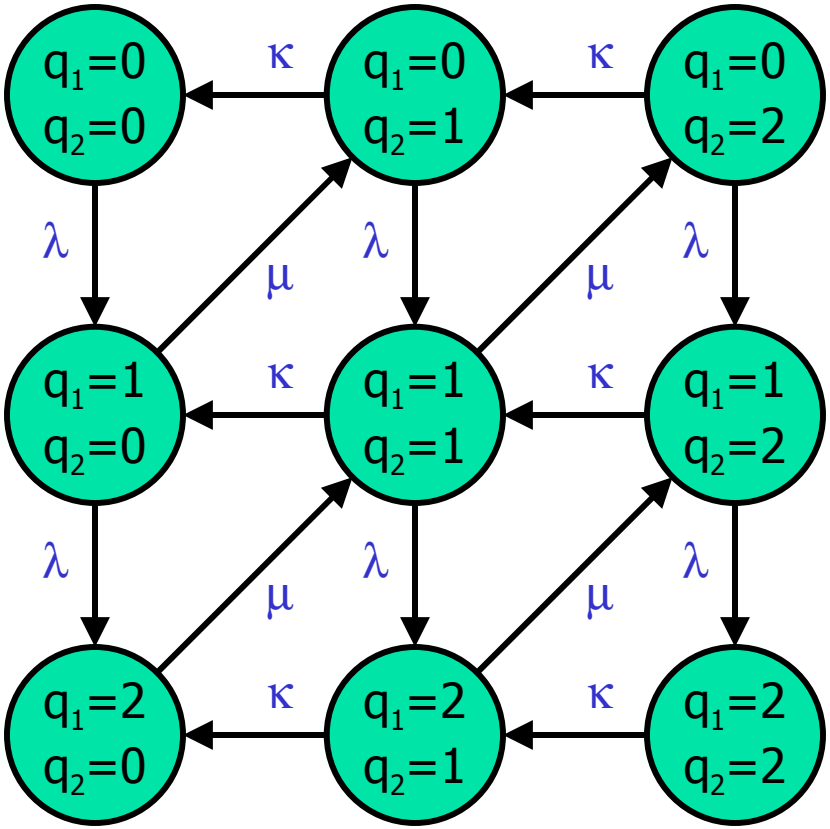

Continuous-Time Markov Chain

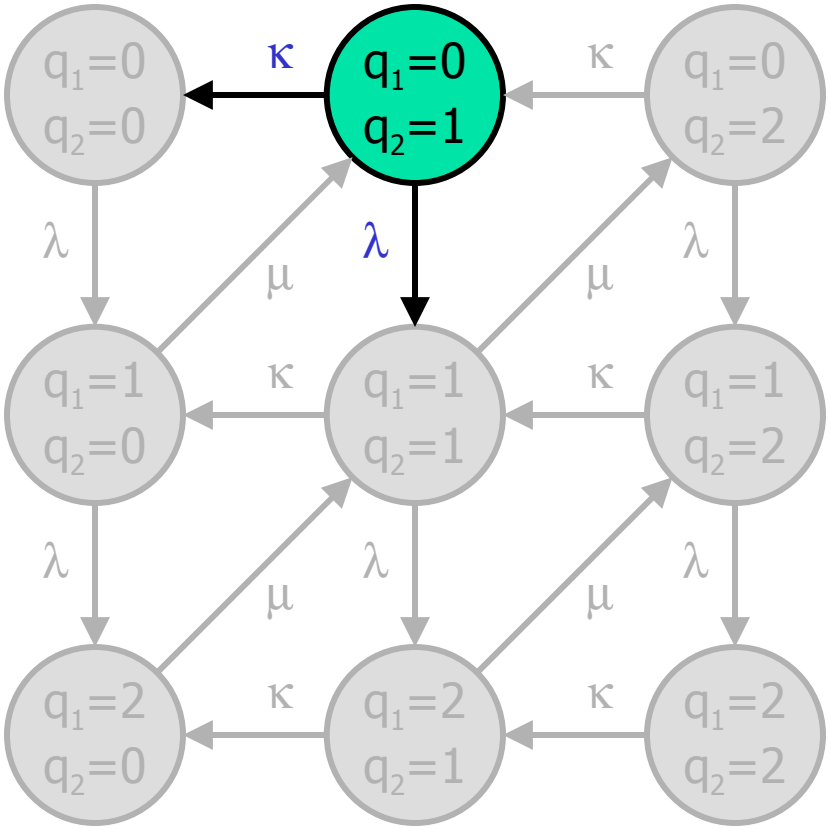

$$\Pr(T \le t) = 1 - e^{-(\lambda+\kappa)t}$$

$$\Pr(q'_1=1, q'_2=1) = \lambda/(\lambda+\kappa)$$

$$\Pr(q'_1=0, q'_2=0) = \kappa/(\lambda+\kappa)$$

Continuous Stochastic Logic (CSL)

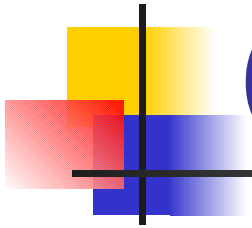

- State formulas
  - Truth value is determined in a single state
- Path formulas
  - Truth value is determined over a path

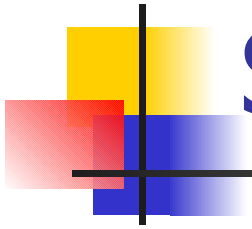

State Formulas

- Standard logic operators: $\neg\varphi$, $\varphi_1 \wedge \varphi_2$, ...
- Probabilistic operator: $Pr_{\geq\theta}(\rho)$
  - Holds in state s iff probability is at least θ that ρ holds over paths starting in s

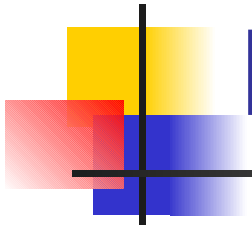

Path Formulas

- Until: $\varphi_1 \; U^{\leq T} \; \varphi_2$
  - Holds over path σ iff φ_2 becomes true in some state along σ before time T, and φ_1 is true in all prior states

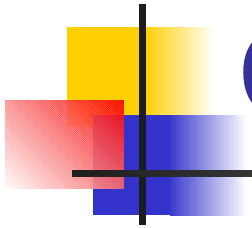

CSL Example

- With both queues empty, is the probability less than 0.5 that both queues become full within 5 seconds?
  - State: $q_1=0 \wedge q_2=0$
  - Property: $Pr_{<0.5}(\text{true } U^{\leq 5}\ q_1=2 \wedge q_2=2)$
    - equivalent to $\neg Pr_{\geq 0.5}(\text{true } U^{\leq 5}\ q_1=2 \wedge q_2=2)$

Numerical vs. Statistical Probabilistic Model Checking

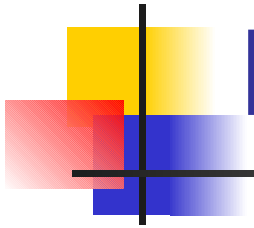

- Numerical Method
  - Highly accurate results
  - Expensive for systems with many states
- Statistical Method
  - Low memory requirements
  - Adapts to difficulty of problem (sequential)
  - Expensive if high accuracy is required

Numerical Solution Method

- Verify $Pr_{\geq\theta}(\varphi_1\ U^{\leq T}\ \varphi_2)$ using transient analysis [Baier et al. 2000]
  - Make states satisfying $\neg\varphi_1 \vee \varphi_2$ absorbing
  - Compute probability p of being in a state satisfying φ_2 at time T in modified model
  - $Pr_{\geq\theta}(\varphi_1\ U^{\leq T}\ \varphi_2)$ holds iff $p \geq \theta$

Probability Computation

- Uniformization [Jensen 1953]
  - Transform model into discrete time Markov chain with transition matrix **P**
  - Compute p for all states as follows:

$$\sum_{k=0}^{\infty} \frac{e^{-q \cdot T} (q \cdot T)^k}{k!} \left(\mathbf{P}^k \cdot \vec{\varphi}_2 \right)$$

  - Truncated summation from L_ε to R_ε with truncation error ε [Fox & Glynn 1988]

Model Transformation: Tandem Queuing Network

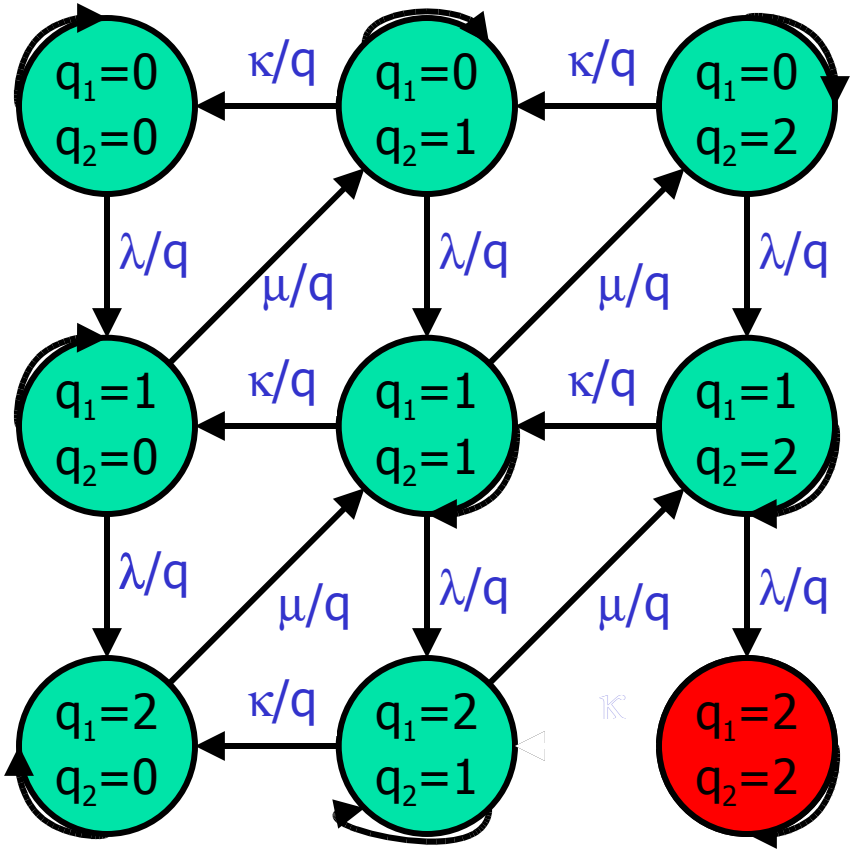

Uniformization constant q

$Pr_{<0.5}(true\ U^{\leq 5}\ q1=2 \wedge q2=2)$

Role of Truncation Error

- We know that $p \geq \tilde{p}$ and $p \leq \tilde{p} + \varepsilon$
  - If $\tilde{p} \geq \theta$ then $p \geq \theta$
  - If $\tilde{p} + \varepsilon \leq \theta$ then $p \leq \theta$
  - Otherwise, can't tell if $Pr_{\geq\theta}(\varphi_1\ U^{\leq T}\ \varphi_2)$ holds
- Good news: $\varepsilon = 10^{-10}$ possible without noticeable performance degradation

Complexity of Numerical Solution Method

- $O(q \cdot T)$ matrix vector multiplications
  - Rates, time bound, and number of states
- All states for same cost
  - In practice, memory and time savings for single state

$$\sum_{k=L_\varepsilon}^{R_\varepsilon} \frac{e^{-q \cdot T} (q \cdot T)^k}{k!} \left(\mathbf{P}^k \cdot \vec{\varphi}_2 \right) \qquad \mathbf{P}^k \cdot \vec{\varphi}_2 = \mathbf{P} \cdot \left(\mathbf{P}^{k-1} \cdot \vec{\varphi}_2 \right)$$

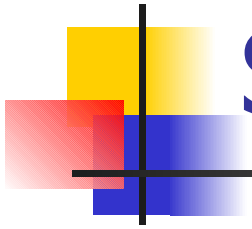

Speedup Techniques

- Steady-state detection [Malhotra et al. 1994]
  - If $\mathbf{P}^k \approx \mathbf{P}^{k-1}$ then stop after k iterations
  - Can lead to significant savings
- Sequential stopping rule
  - Stop if $\tilde{p} \geq \theta$ after k iterations
  - At most $O(\sqrt{q \cdot T})$ fewer iterations

Statistical Solution Method [Younes & Simmons 2002]

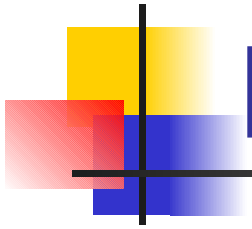

- Use discrete event simulation to generate sample paths
- Use sequential acceptance sampling to verify probabilistic properties
  - Hypothesis: $Pr_{\geq\theta}(\rho)$

Not estimation!

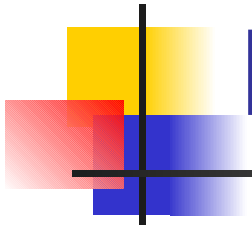

Error Bounds

- Probability of false negative: $\leq \alpha$
  - We say that φ is false when it is true
- Probability of false positive: $\leq \beta$
  - We say that φ is true when it is false

Performance of Test

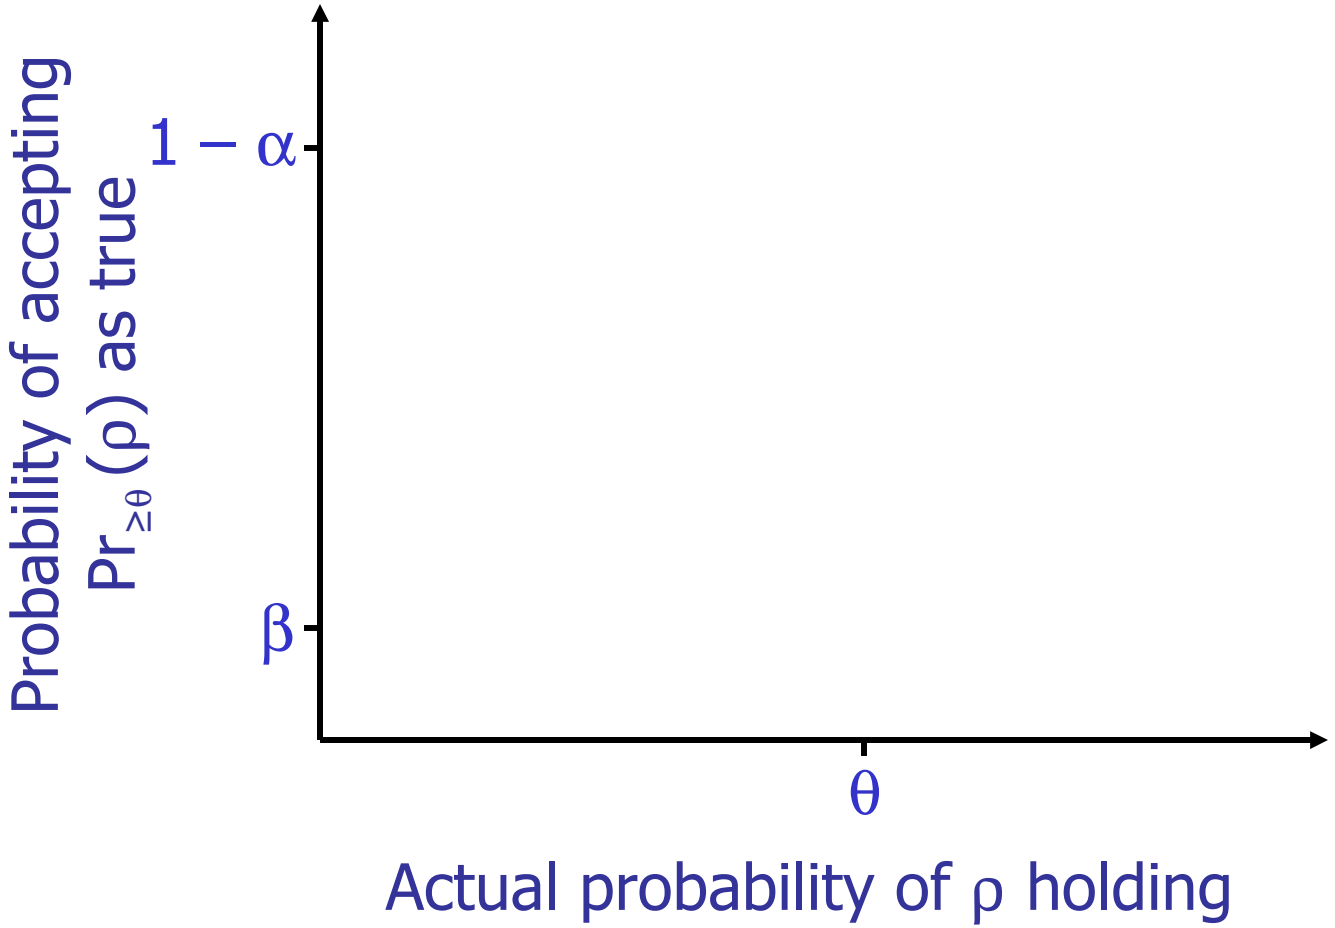

Ideal Performance

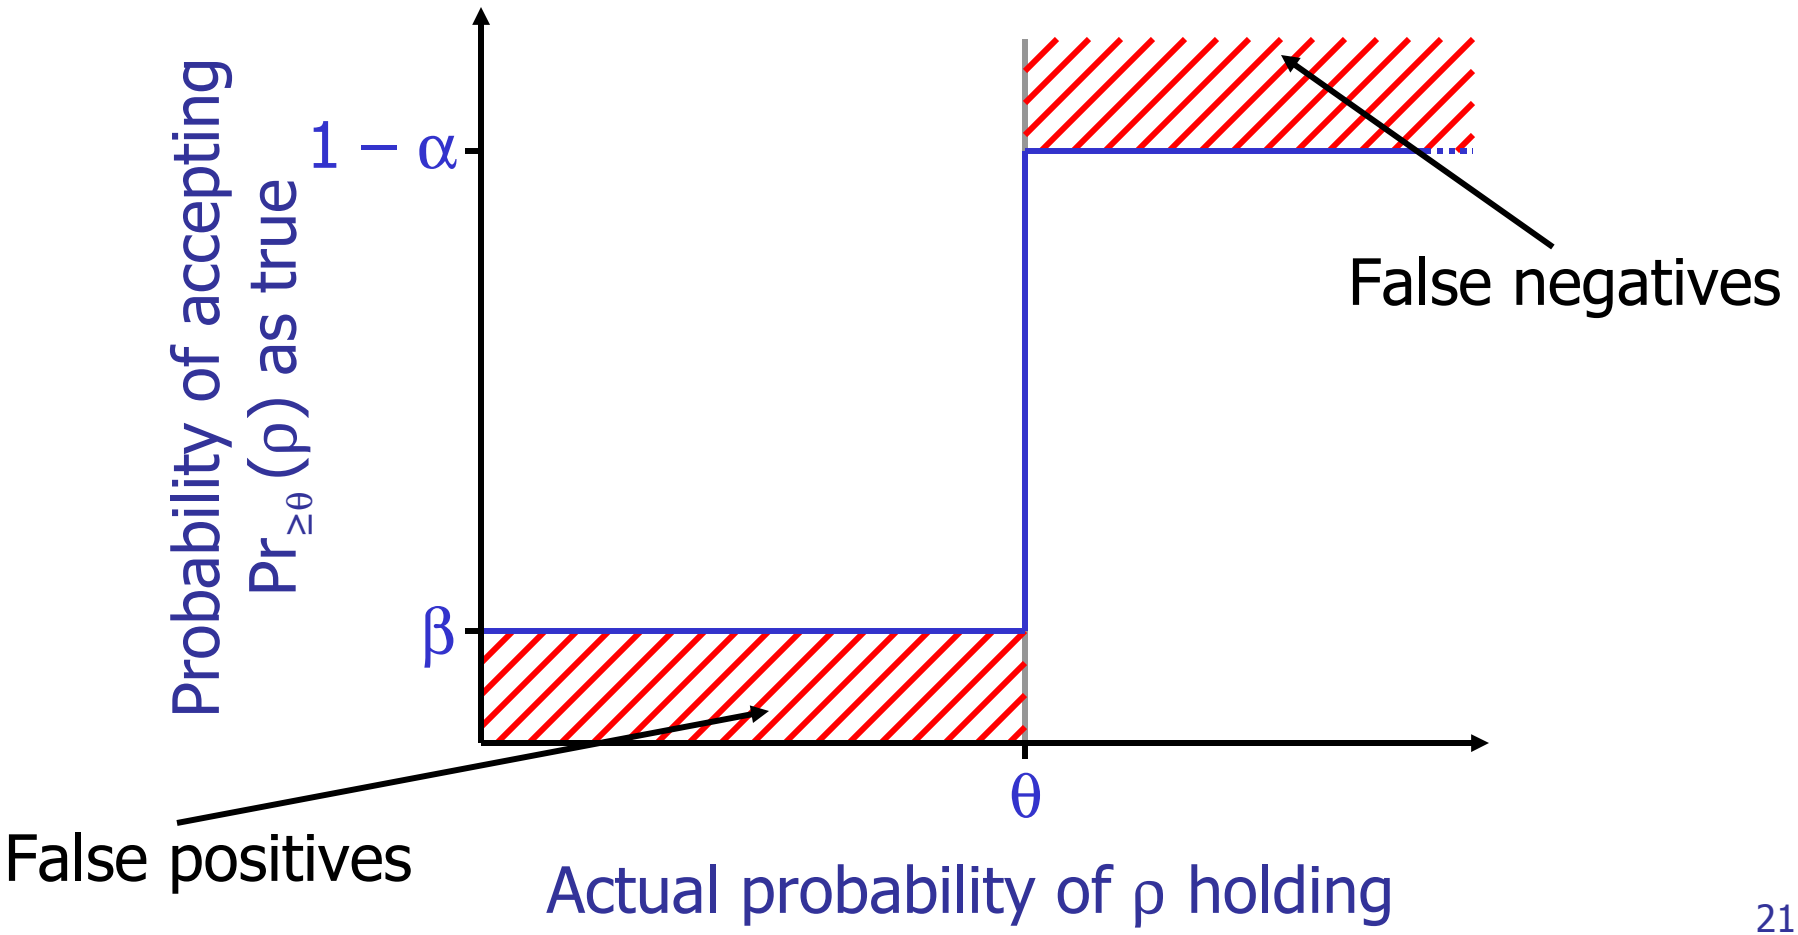

Actual Performance

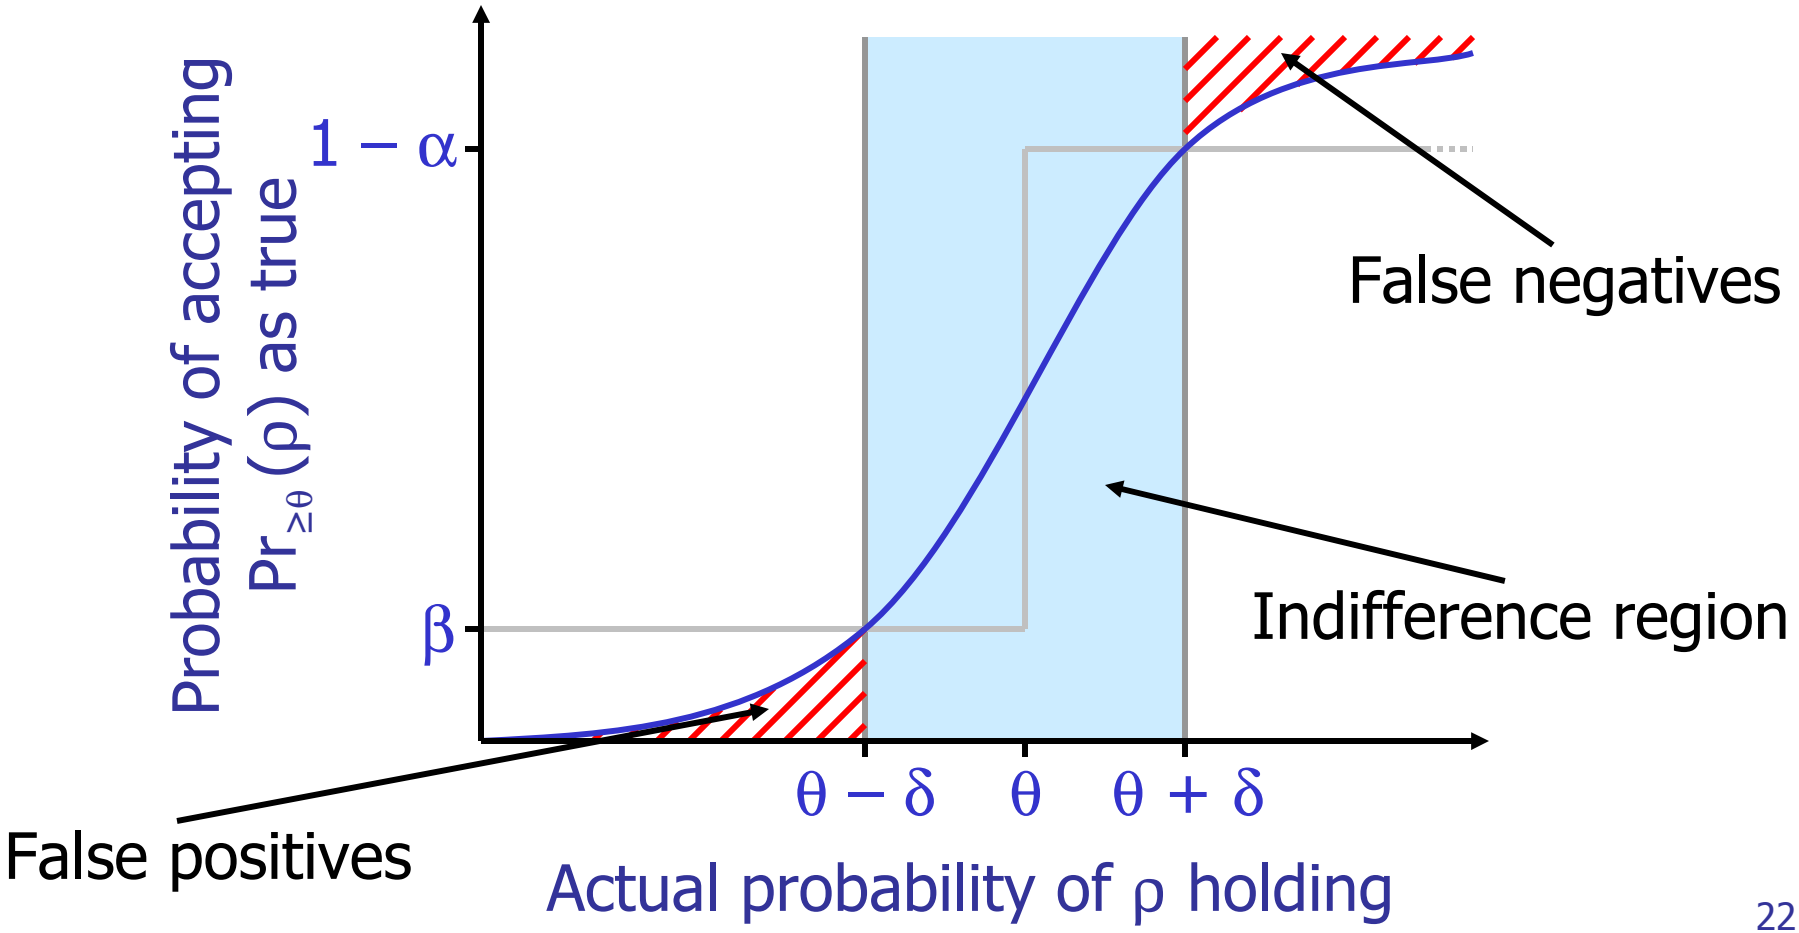

Sequential Acceptance Sampling [Wald 1945]

- Hypothesis: $Pr_{\geq\theta}(\rho)$

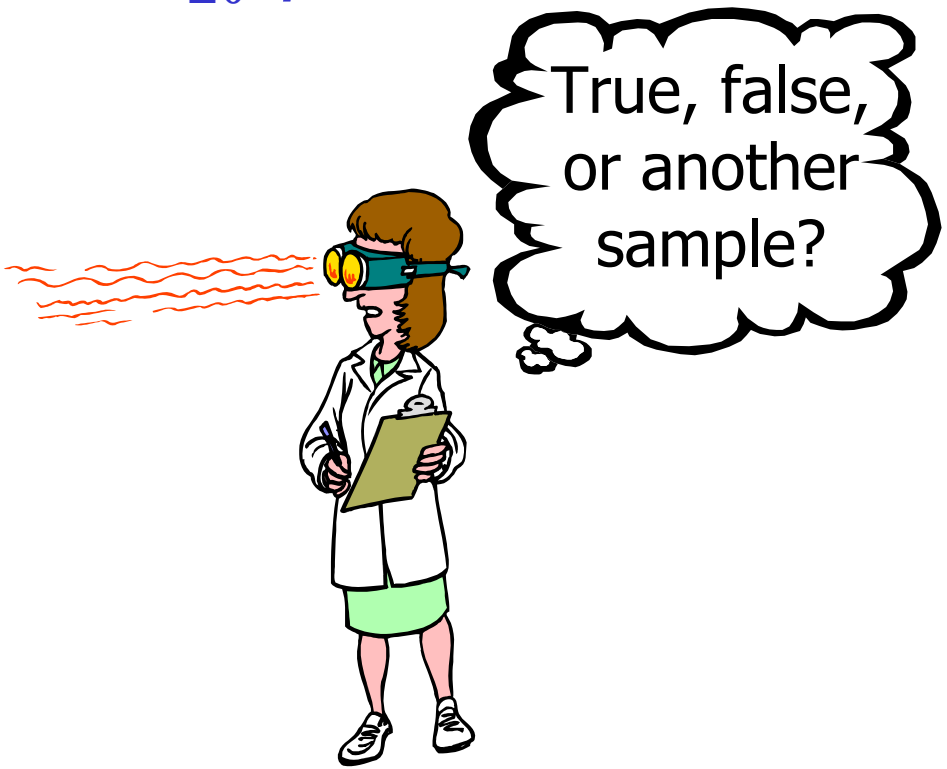

Graphical Representation of Sequential Test

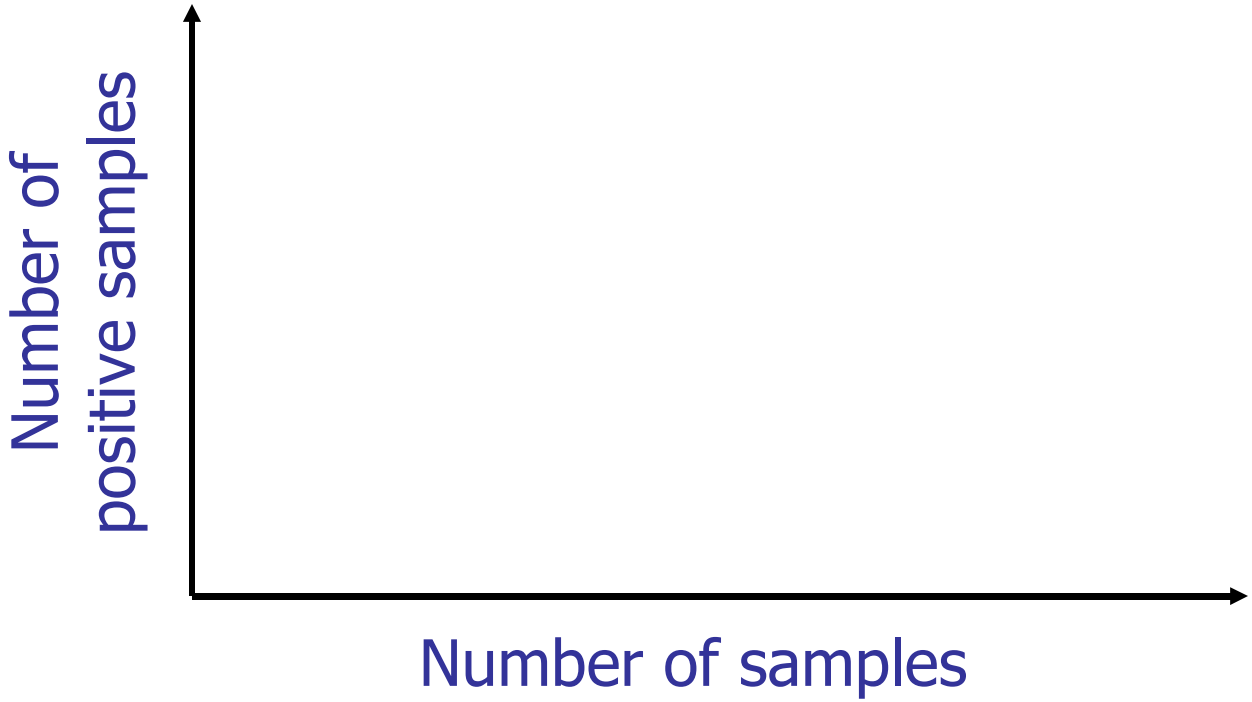

Graphical Representation of Sequential Test

- We can find an acceptance line and a rejection line given θ, δ, α, and β

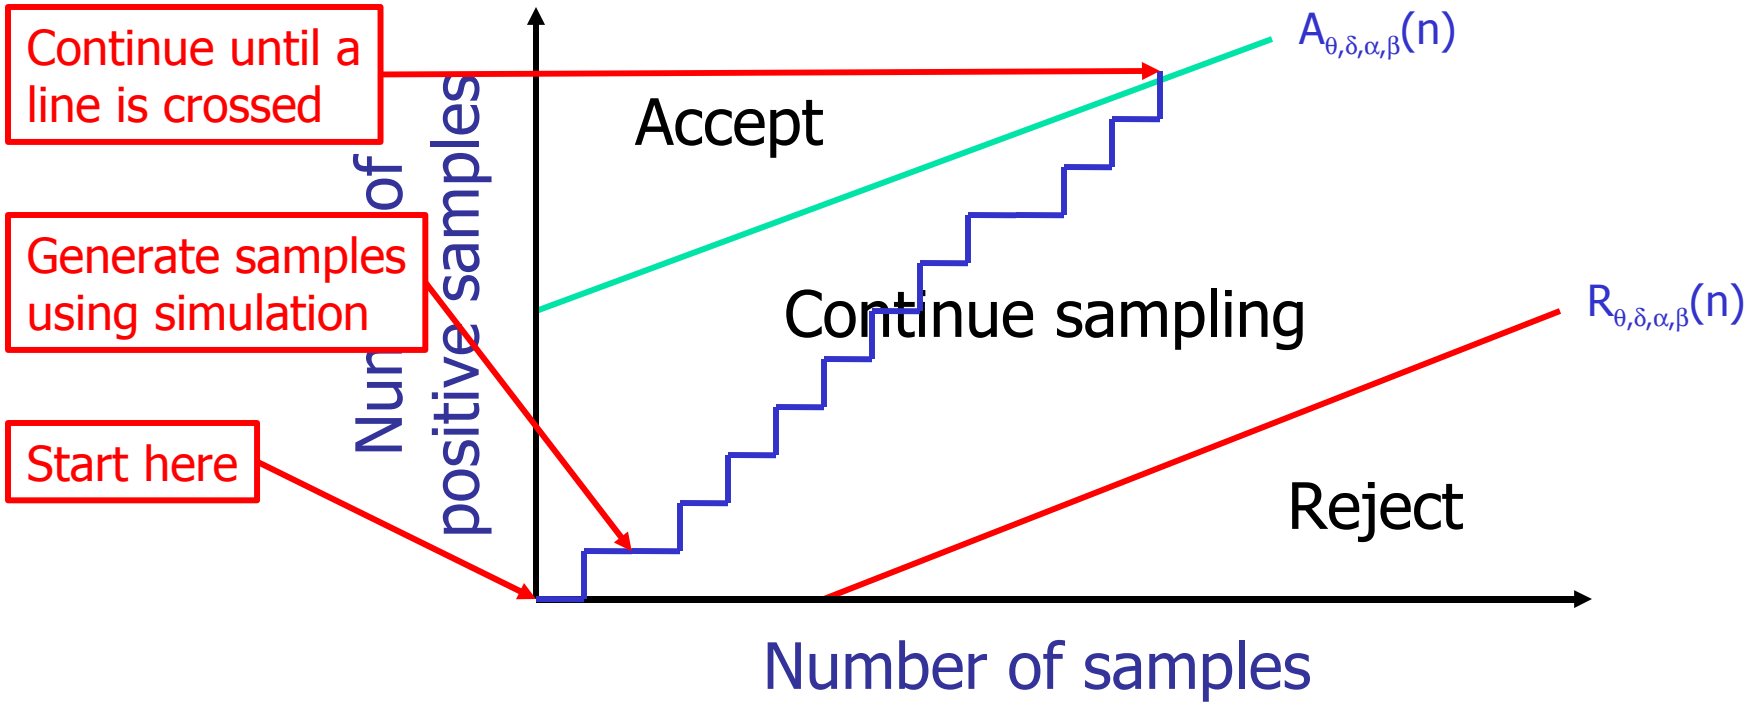

Verifying Probabilistic Properties

- Verify $Pr_{\geq\theta}(\rho)$ with error bounds α and β
  - Generate sample paths using simulation
  - Verify ρ over each sample path
    - If ρ is true, then we have a positive sample
    - If ρ is false, then we have a negative sample
  - Use sequential acceptance sampling to test the hypothesis $Pr_{\geq\theta}(\rho)$

Complexity of Statistical Solution Method

- Number of samples
  - Complex dependency on θ, δ, α, and β
- Length of sample paths
  - Expected length at most $q \cdot T$
  - Shorter paths if $\neg\varphi_1 \vee \varphi_2$ is satisfied early
- No direct dependence on size of state space

Case Study: Tandem Queuing Network

- $M/Cox_2/1$ queue sequentially composed with $M/M/1$ queue
- Each queue has capacity n
- State space of size $O(n^2)$

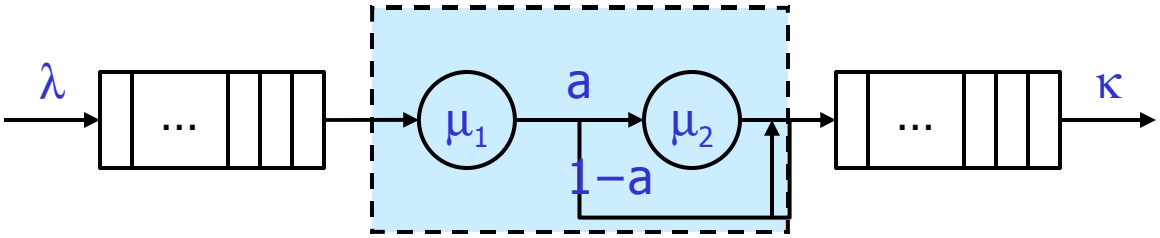

Tandem Queuing Network (results)

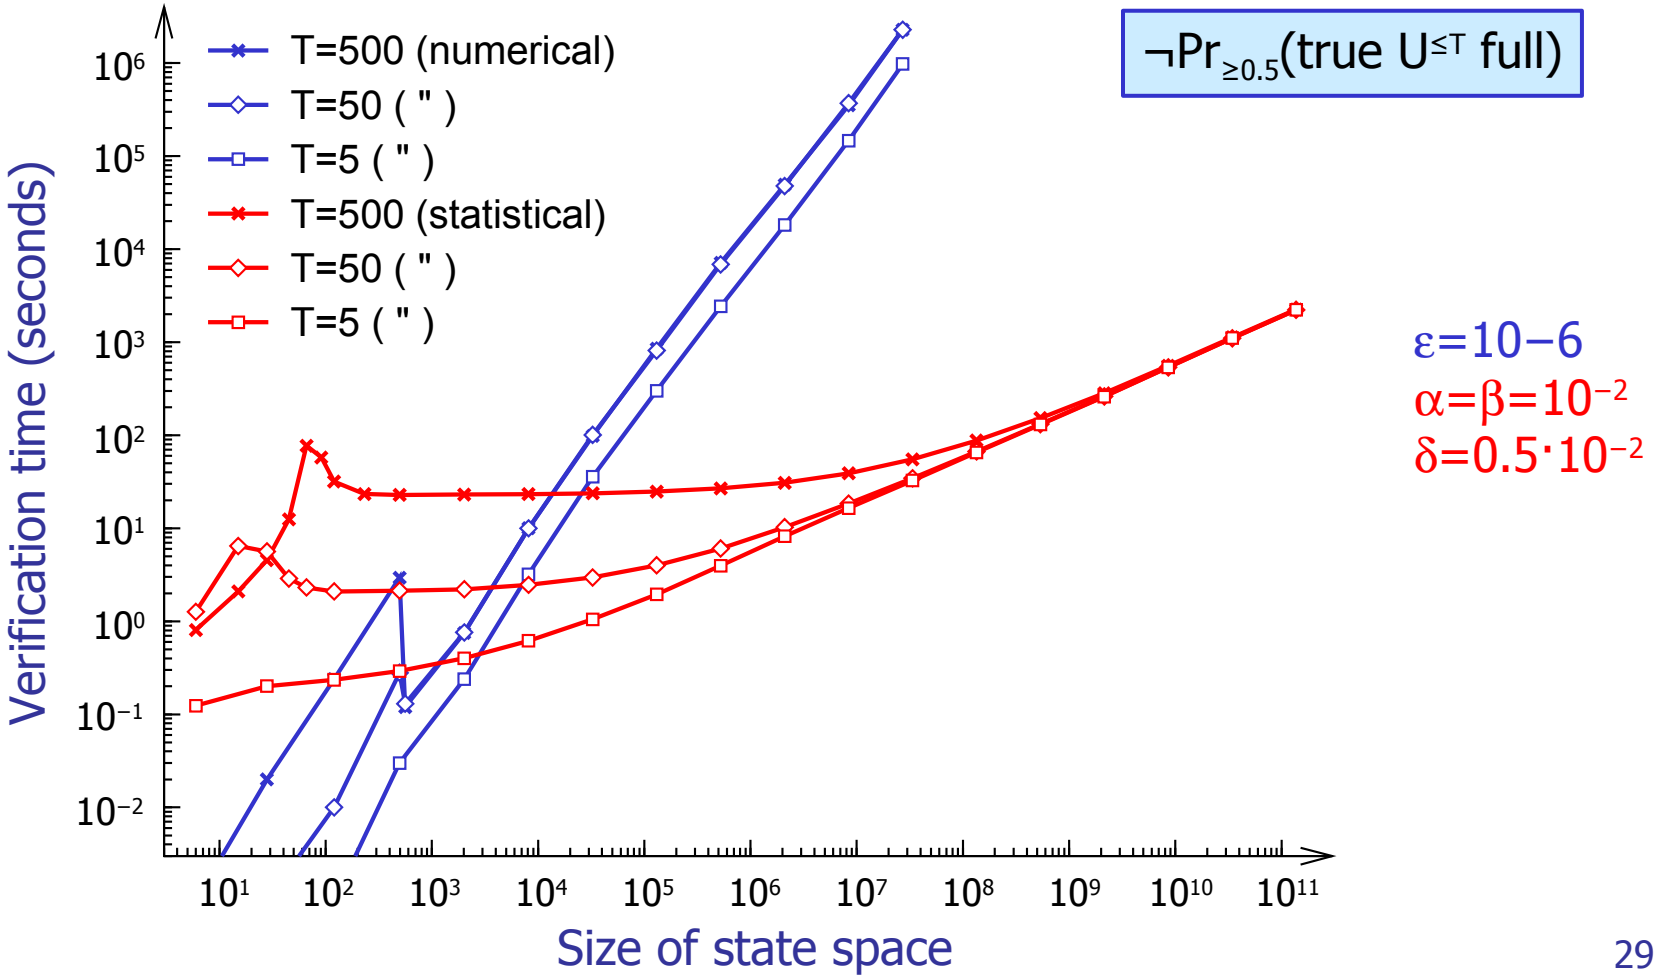

Tandem Queuing Network (results)

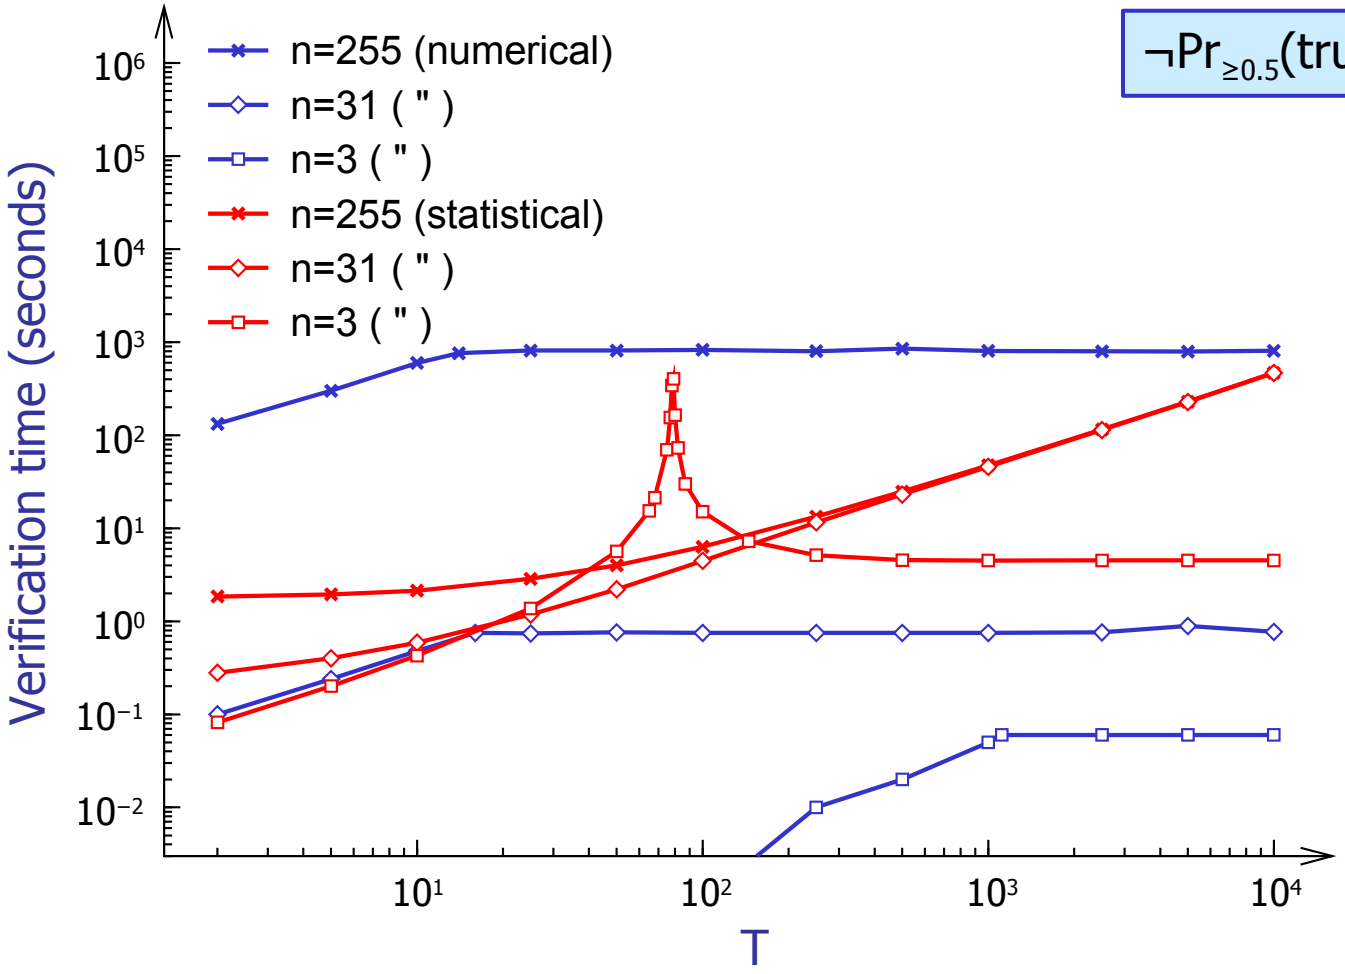

$\neg Pr_{\geq 0.5}(\text{true } U^{\leq T} \text{ full})$

$\varepsilon = 10-6$
$\alpha = \beta = 10^{-2}$
$\delta = 0.5 \cdot 10^{-2}$

Case Study: Symmetric Polling System

- Single server, n polling stations
- Stations are attended in cyclic order
- Each station can hold one message
- State space of size $O(n \cdot 2^n)$

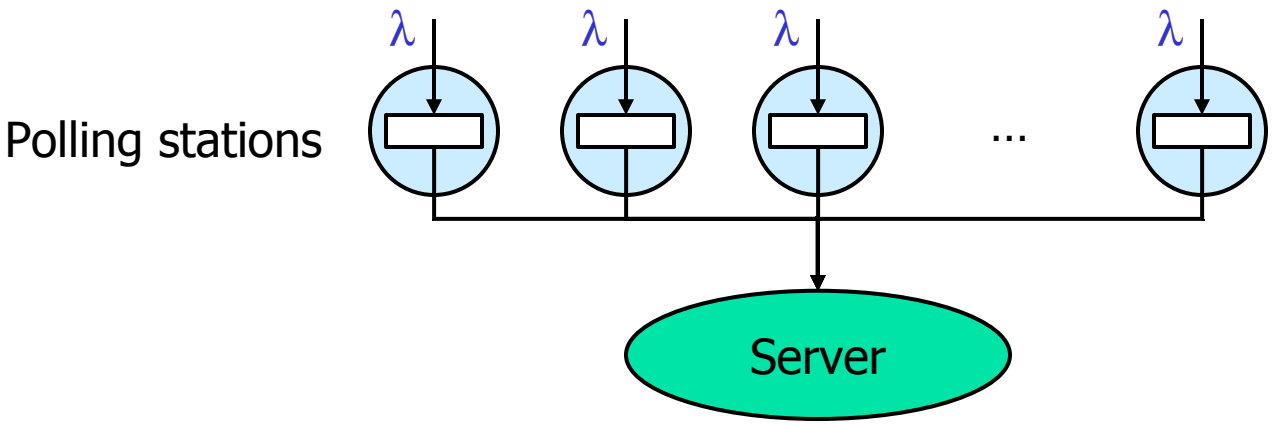

Symmetric Polling System (results)

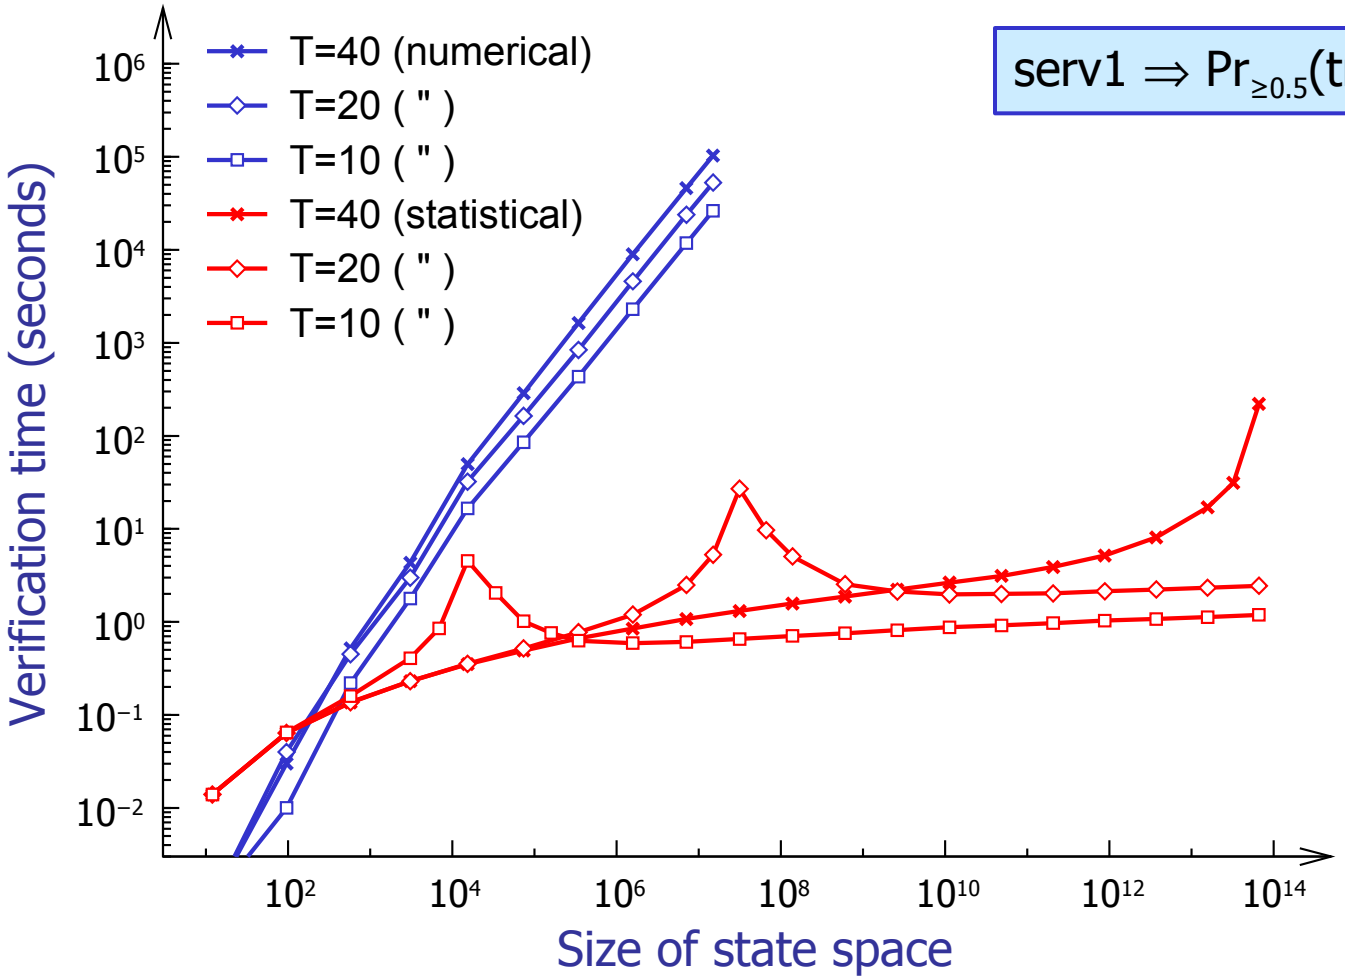

$serv1 \Rightarrow Pr_{\geq 0.5}(\text{true } U^{\leq T} \text{ poll1})$

$\varepsilon = 10-6$
$\alpha = \beta = 10^{-2}$
$\delta = 0.5 \cdot 10^{-2}$

Symmetric Polling System (results)

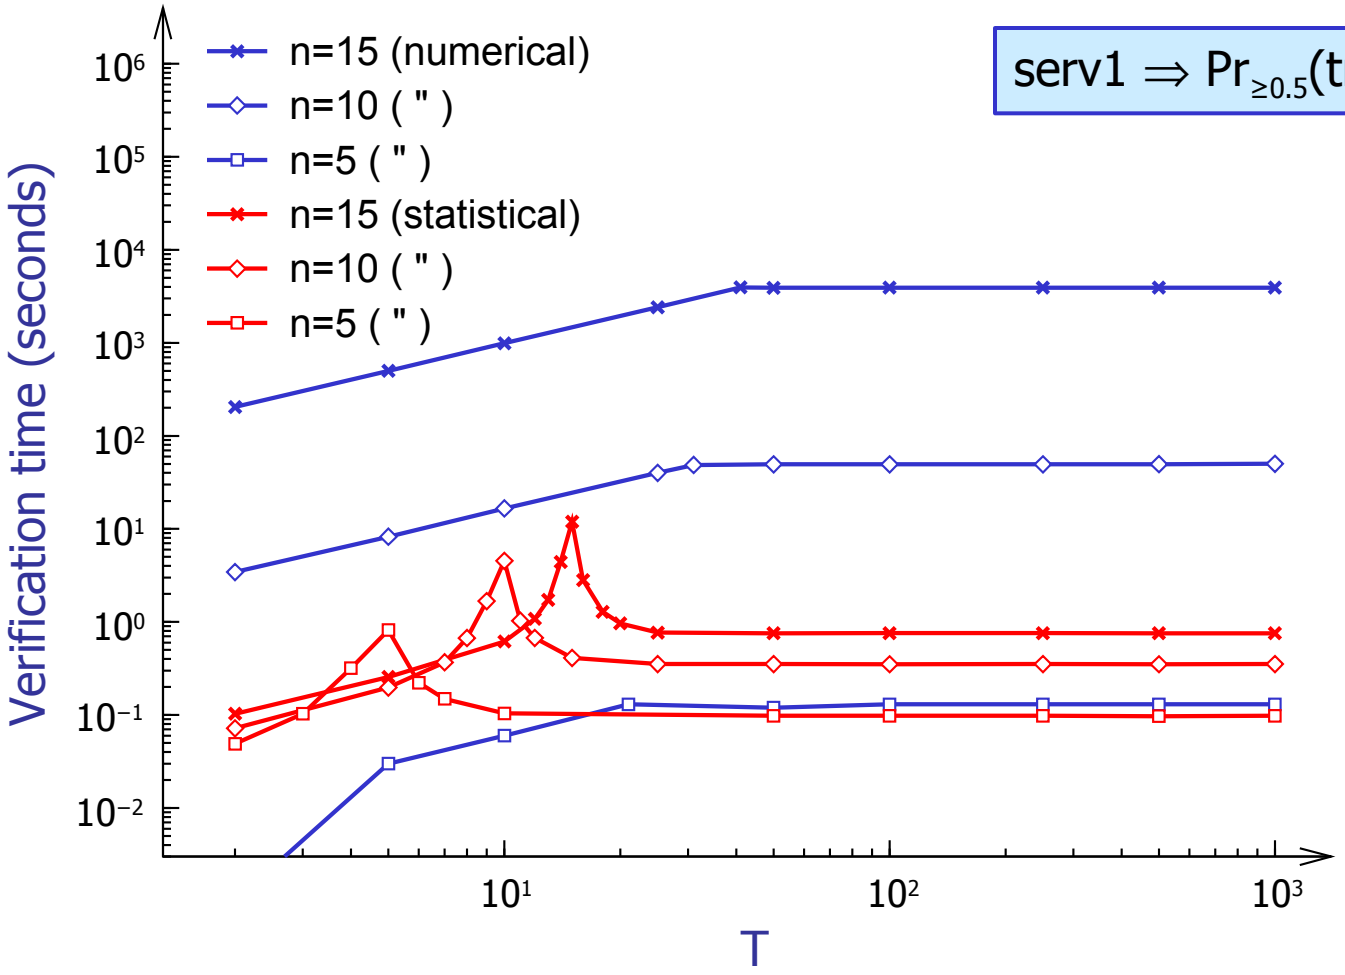

$serv1 \Rightarrow Pr_{\geq 0.5}(\text{true } U^{\leq T} \text{ poll1})$

$\varepsilon = 10-6$

$\alpha = \beta = 10^{-2}$

$\delta = 0.5 \cdot 10^{-2}$

Symmetric Polling System (results)

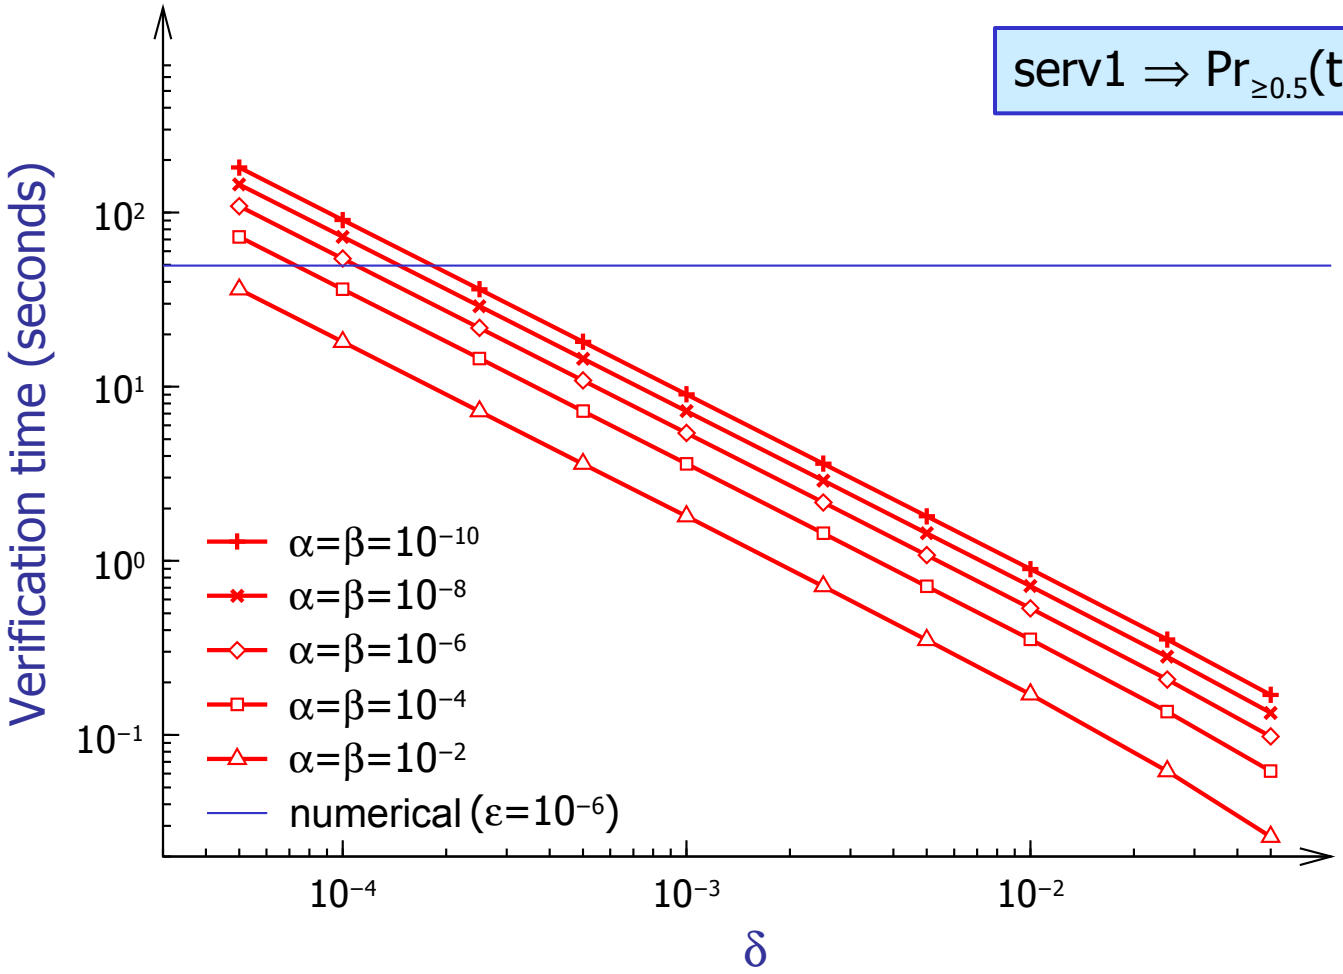

Nested Queries

- $Pr_{\geq 0.9}(Pr_{\geq 0.5}(\text{true } U^{\leq 5} \text{ comm.}) U^{\leq 20} \text{ gold})$
- Statistical method: hypothesis testing problem in each state along a path!

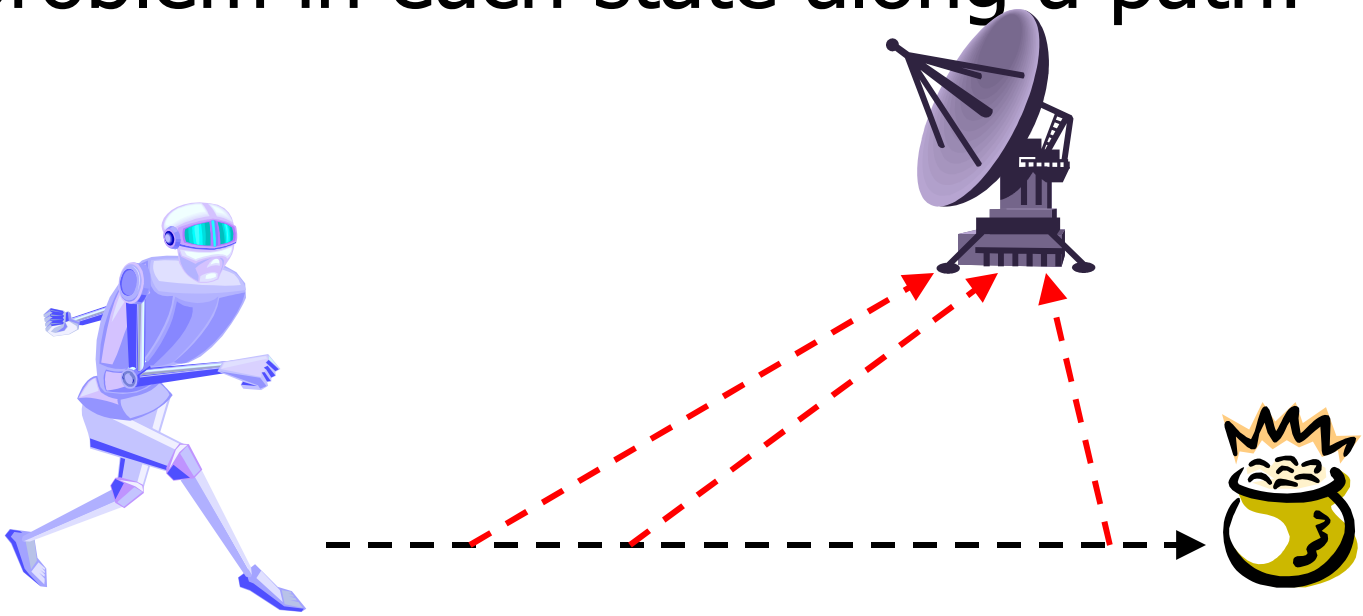

Nested Queries: Combining the Methods

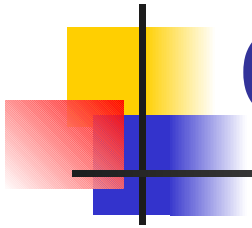

- Verify inner probabilistic statement for all states using numerical method
- Verify outer probabilistic statement using statistical method

Nested Queries (results)

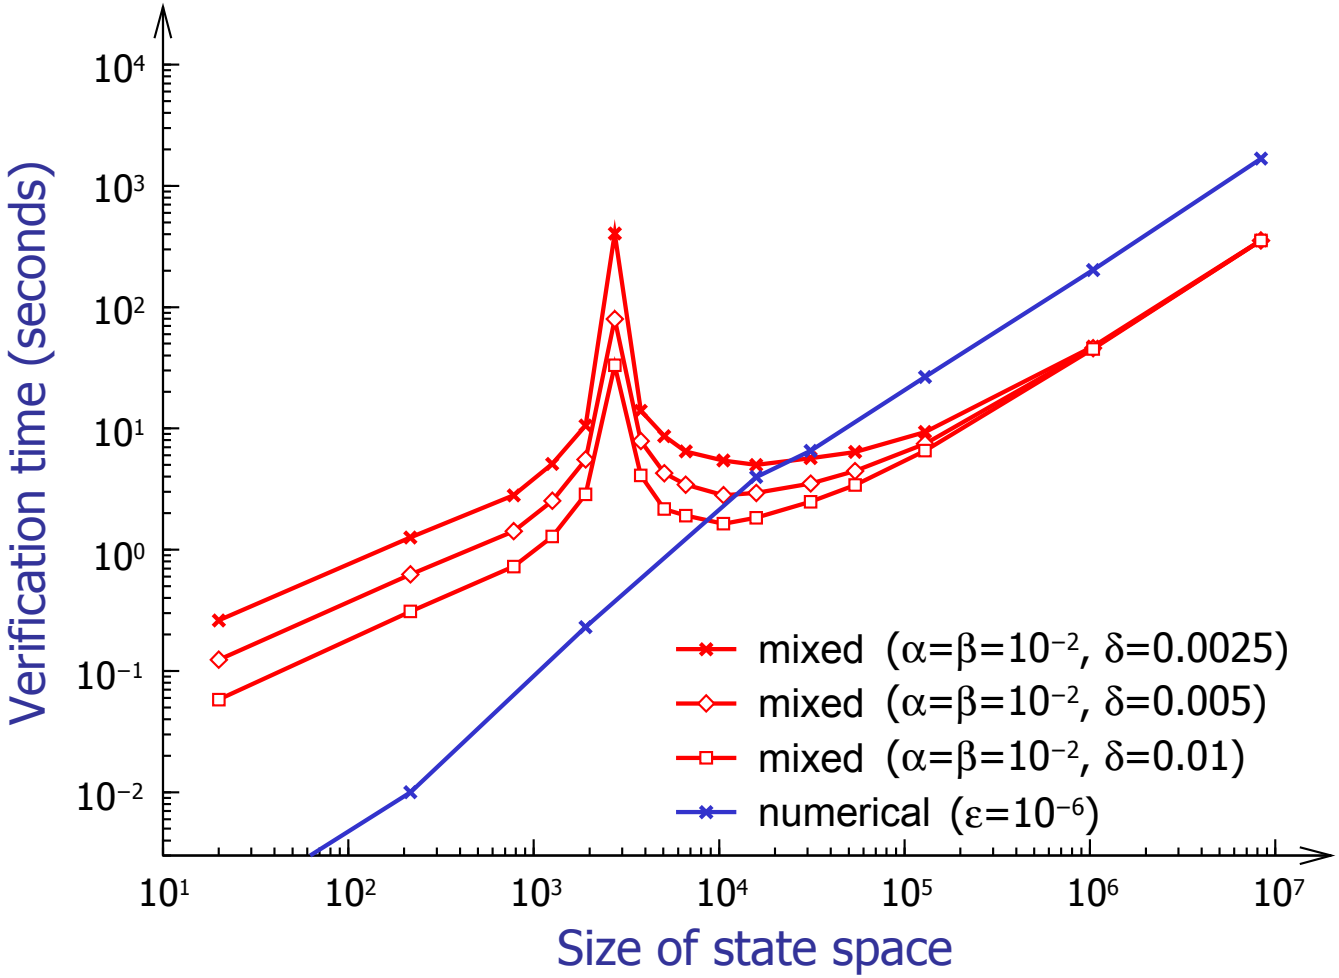

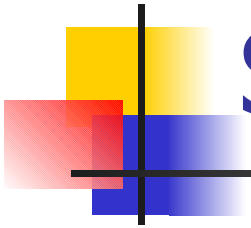

Summary

- Benefits of numerical method
  - All states at the price of one
  - Steady-state detection
  - High accuracy
- Benefits of statistical method
  - Easy to trade accuracy for speed
  - Scales well with size of state space
  - Parallelizable
  - Model independent

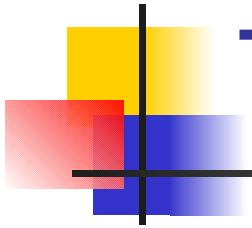

Tools

- PRISM
  - http://www.cs.bham.ac.uk/~dxp/prism/
- Ymer
  - http://www.cs.cmu.edu/~lorens/ymer.html