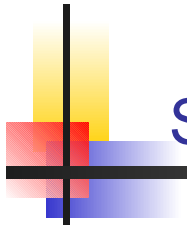

Statistical probabilistic model checking

Håkan L. S. Younes
Carnegie Mellon University
(now at Google Inc.)

Introduction

Model-independent approach to probabilistic model checking

- Relies on simulation and statistical sampling
- Wrong answer possible, but can be bounded (probabilistically)
- Low memory requirements (can handle large/infinite models)
- Trivially parallelizable (distributed sampling gives linear speedup)

Topics covered in this talk:

- Error control
- Hypothesis testing vs. estimation
- Dealing with unbounded properties/infinite trajectories

Probabilistic model checking

Given a model $\mathcal{M}$, a state s, and a property Φ, does Φ hold in s for $\mathcal{M}$?

- Model: stochastic discrete-event system
- Property: probabilistic temporal logic formula

Example: tandem queuing network

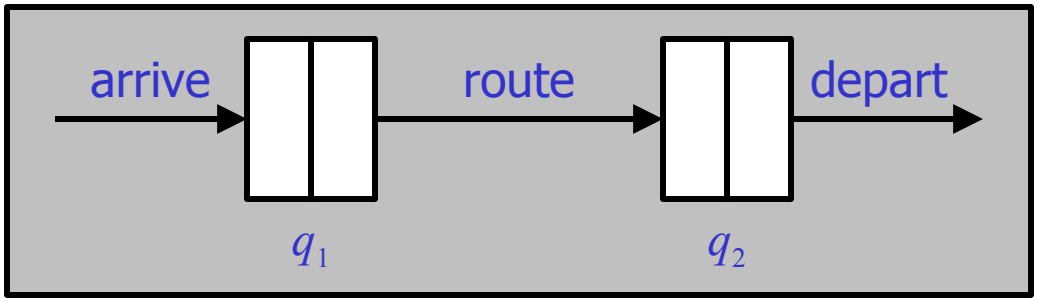

"The probability is at least 0.1 that both queues become full within 5 minutes"

Probabilistic temporal logic (PCTL, CSL)

Standard logic operators: $\neg \Phi$, $\Phi \wedge \Psi$, …

Probabilistic operator: $\mathcal{P}_{\geq \theta}[\varphi]$

- Holds in state s iff probability is at least θ for paths satisfying φ and starting in s

Bounded until: $\Phi \, \mathcal{U}^{\leq T} \Psi$

- Holds over path σ iff Ψ becomes true along σ within time T, and Φ is true until then

Unbounded until: $\Phi \, \mathcal{U} \, \Psi$

- Holds over path σ iff Ψ becomes true eventually along σ, and Φ is true until then

Property examples

"The probability is at least 0.1 that both queues become full within 5 minutes"

- $\mathcal{P}_{\geq 0.1}[\top\ \mathcal{U}^{\leq 5}\ full_1 \wedge full_2\ ]$

"The probability is at most 0.05 that the second queue becomes full before the first queue"

- $\mathcal{P}_{\leq 0.05}[\neg\ full_1\ \mathcal{U}\ full_2\ ]$

The problem (in detail)

Before we propose a solution, we need to fully define the problem

- Possible outcomes of model-checking algorithm
- Ideal vs. realistic error control

Possible outcomes of model-checking algorithm

Given a state s and a formula Φ, a model-checking algorithm $\mathcal{A}$ can:

- Accept Φ as true in s ($s \vdash_{\top} \Phi$)
- Reject Φ as false in s ($s \vdash_{\bot} \Phi$)
- Return an undecided result ($s \vdash_{\mathtt{I}} \Phi$)

An error occurs if:

- $\mathcal{A}$ rejects Φ when Φ is true (false negative)
- $\mathcal{A}$ accepts Φ when Φ is false (false positive)

Note: an undecided result is not an error, but still not desirable

Ideal error control

Bound the probability of false negatives/positives and undecided results under all circumstances

- Bound on false negatives: α $\quad \Pr[s \vdash_{\perp} \Phi \mid s \models \Phi] \leq \alpha$
- Bound on false positives: β $\quad \Pr[s \vdash_{\top} \Phi \mid s \not\models \Phi] \leq \beta$
- Bound on undecided results: γ $\quad \Pr[s \vdash_{?} \Phi] \leq \gamma$

If α, β, and γ are all low, then model-checking algorithm $\mathcal{A}$ produces a correct result with high probability

Unrealistic expectations

Ideal error control for verifying probabilistic formula $\mathcal{P}_{\geq\theta}[\varphi]$ in state s

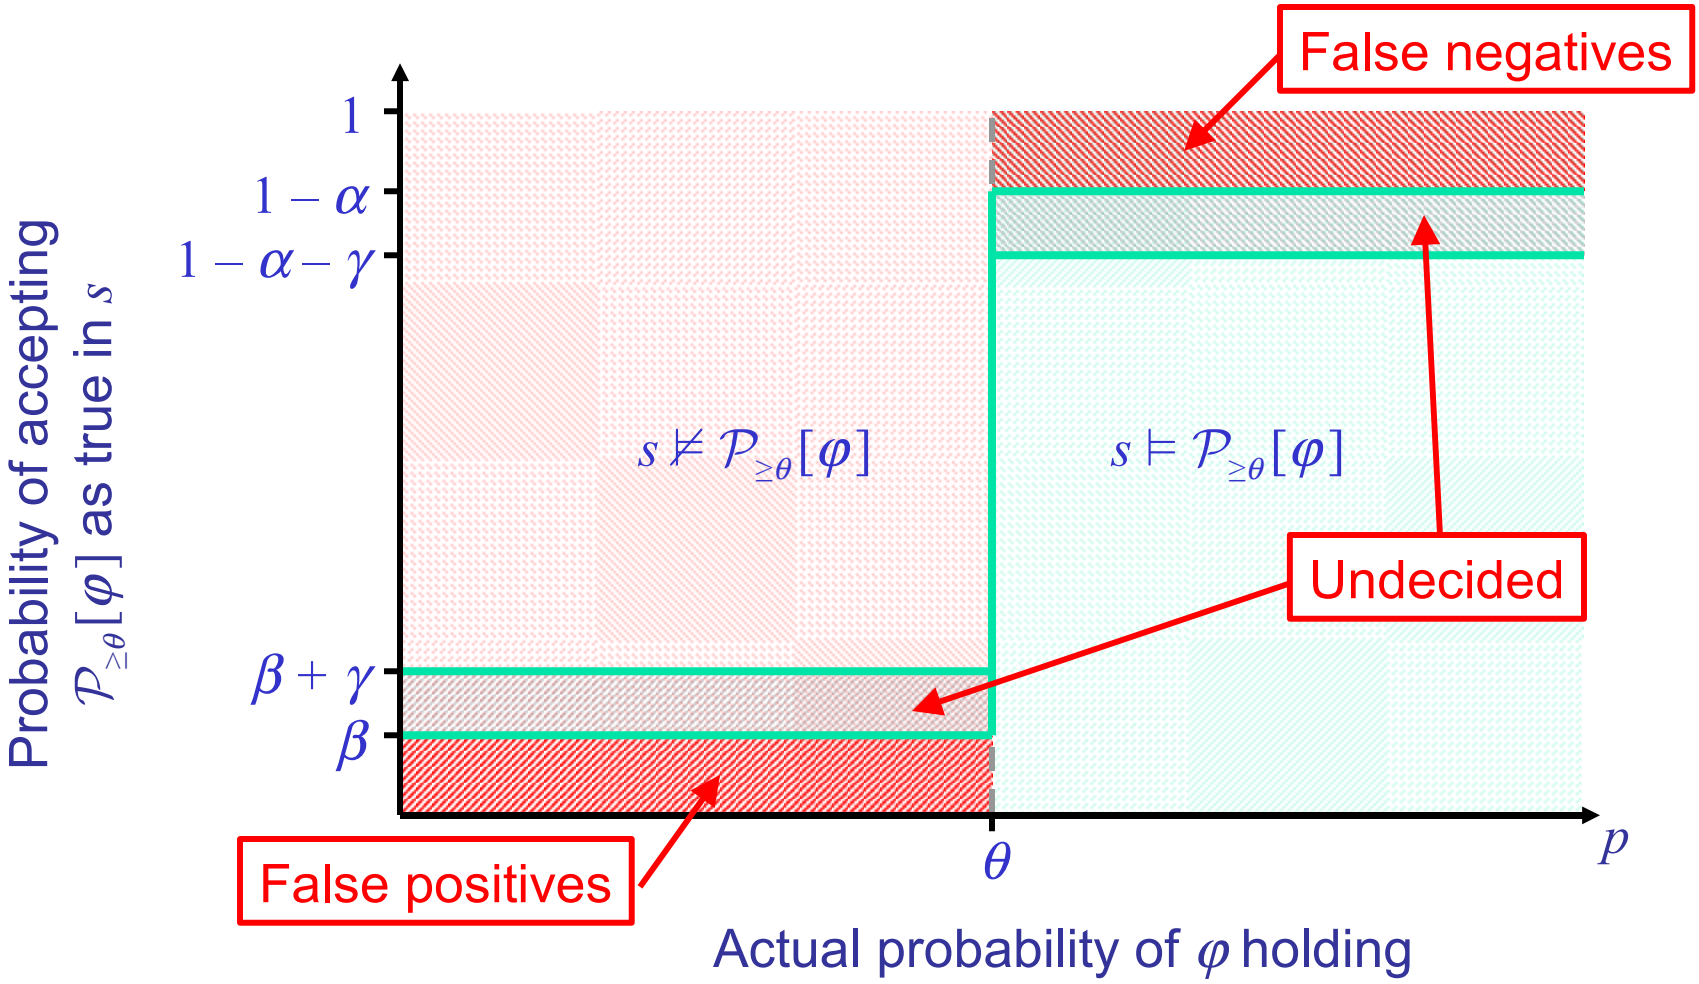

Relaxing the problem

Indifference region of width 2δ centered around probability thresholds

Probabilistic operator: $\mathcal{P}_{\geq\theta}[\varphi]$

- Holds in state s if probability is at least $\theta+\delta$ for paths satisfying φ and starting in s
- Does not hold if probability is at most $\theta-\delta$ for paths satisfying φ and starting in s
- "Too close to call" if probability is within δ distance of θ (indifference)

Essentially three-valued logic, but we care only about true and false

Error control for relaxed problem

Option 1: bound the probability of false positives/negatives outside of the indifference region; no undecided results

- Bound on false negatives: α $\quad \Pr[s \vdash_{\perp} \Phi \mid s \models_{\top}^{\delta} \Phi] \leq \alpha$
- Bound on false positives: β $\quad \Pr[s \vdash_{\top} \Phi \mid s \models_{\perp}^{\delta} \Phi] \leq \beta$
- No undecided results: $\gamma = 0$ $\quad \Pr[s \vdash_{?} \Phi] = 0$

Option 2: bound the probability of undecided results outside of the indifference region; low error probability under all circumstances

- Bound on false negatives: α $\quad \Pr[s \vdash_{\perp} \Phi \mid s \models \Phi] \leq \alpha$
- Bound on false positives: β $\quad \Pr[s \vdash_{\top} \Phi \mid s \not\models \Phi] \leq \beta$
- Bound on undecided results: γ $\quad \Pr[s \vdash_{?} \Phi \mid (s \models_{\top}^{\delta} \Phi) \vee (s \models_{\perp}^{\delta} \Phi)] \leq \gamma$

Realistic error control—no undecided results

Error control for verifying probabilistic formula $\mathcal{P}_{\geq\theta}[\varphi]$ in state s

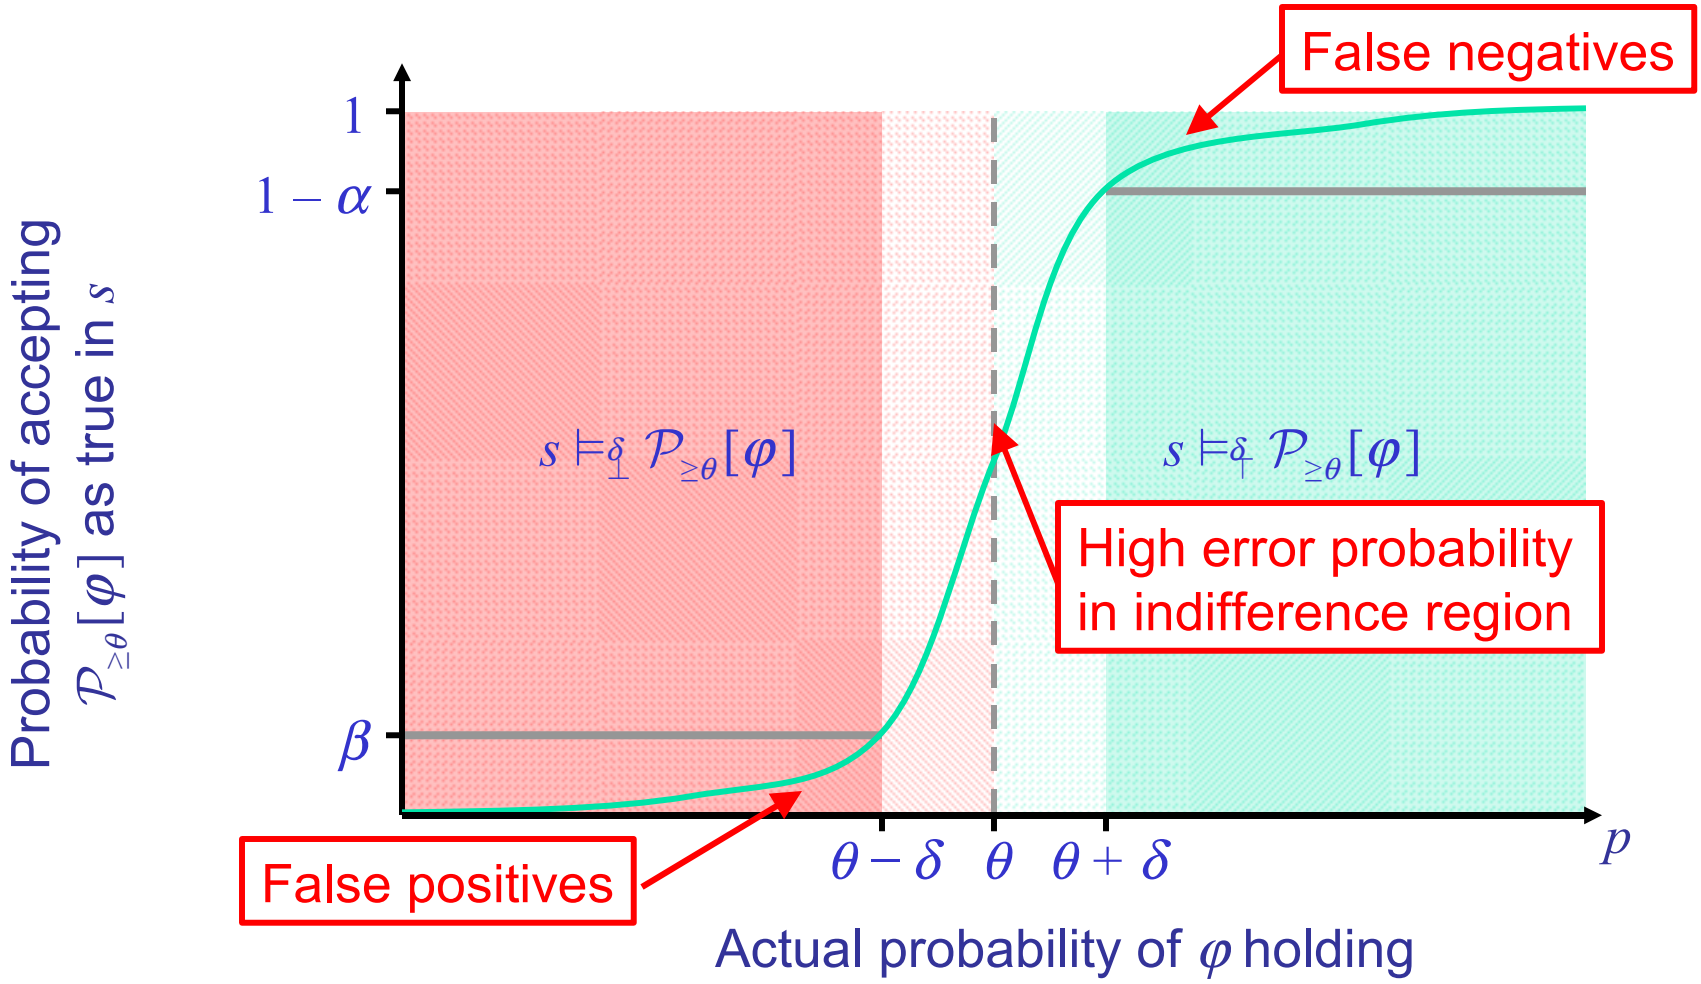

Realistic error control—with undecided results

Error control for verifying probabilistic formula $\mathcal{P}_{\geq\theta}[\varphi]$ in state s

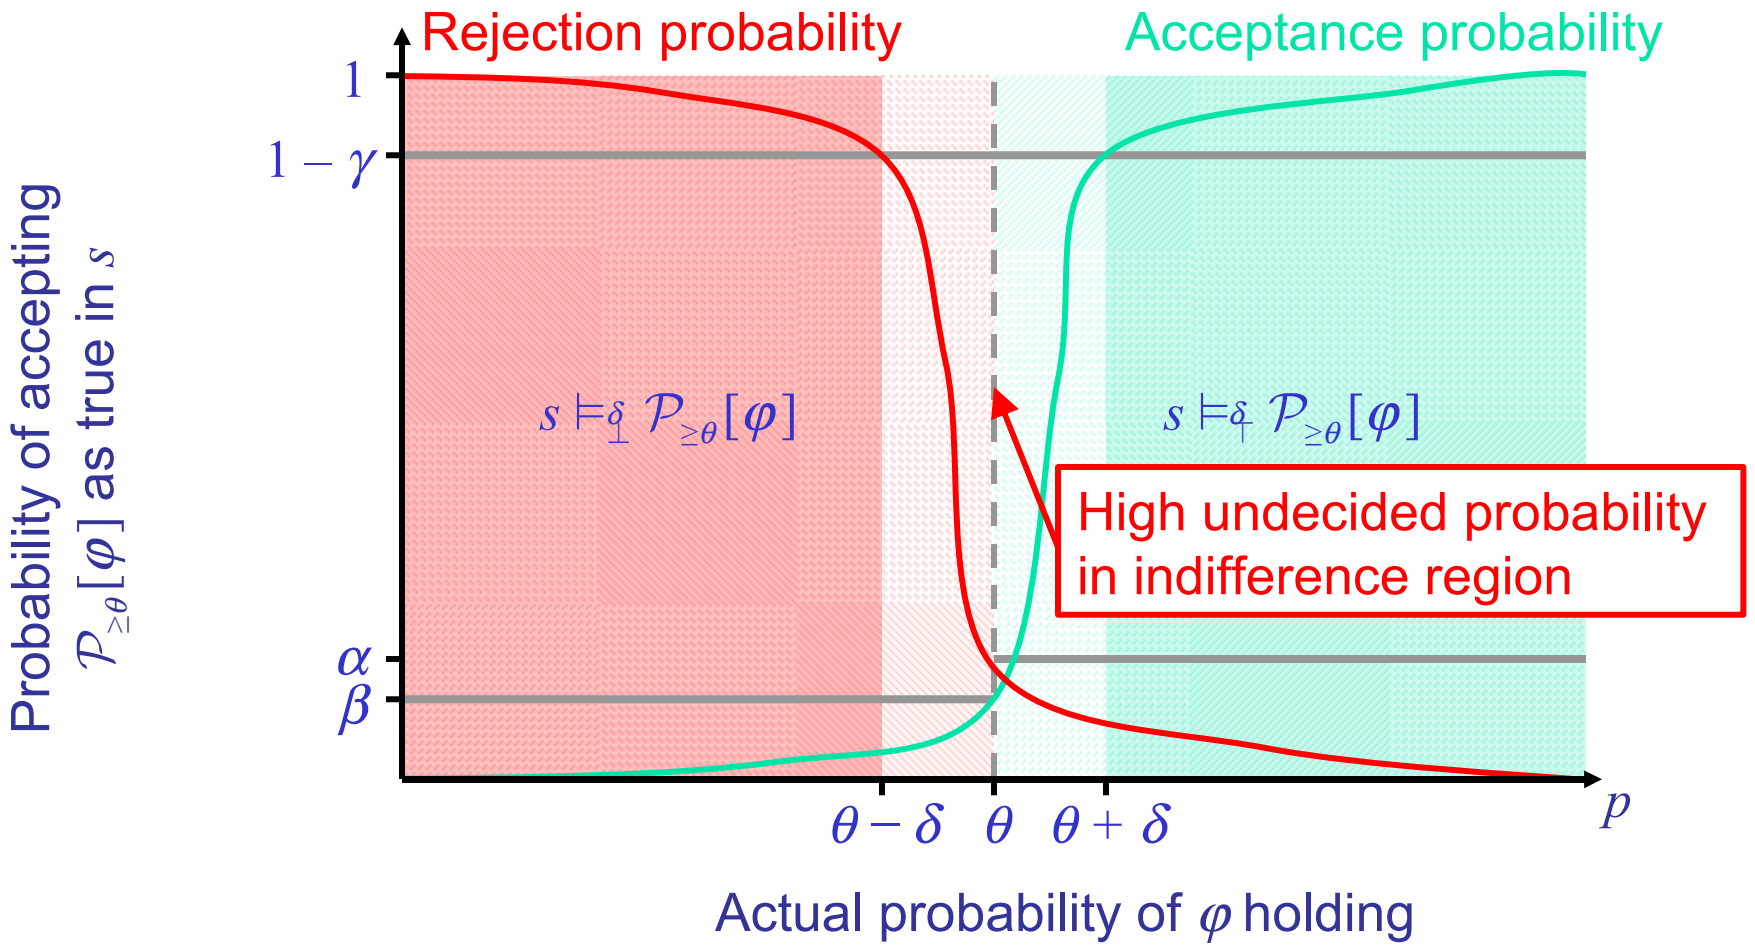

The solution

- Statistical sampling (hypothesis testing vs. estimation)
- Undecided results
- Avoiding infinite sample trajectories in simulation for unbounded until

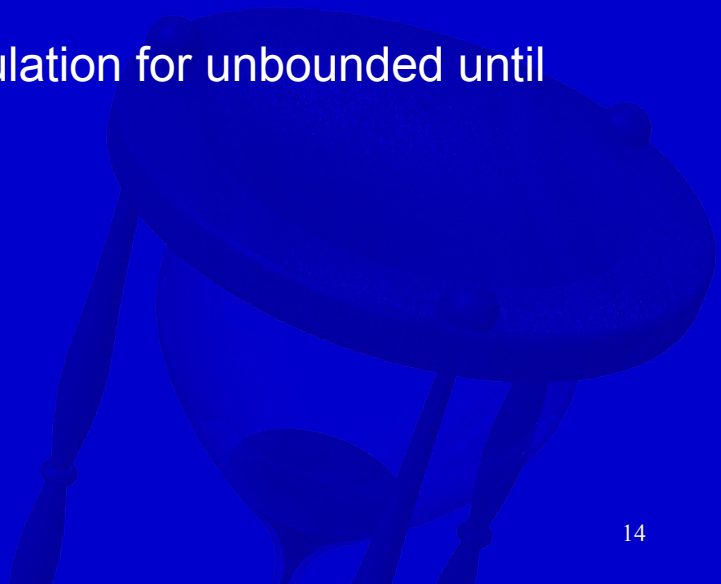

Verifying probabilistic properties—no undecided results

Younes & Simmons (CAV'02, Information and Computation'06)

Use acceptance sampling to verify $\mathcal{P}_{\geq\theta}[\varphi]$ in state s

- Test hypothesis $H_0: p \geq \theta + \delta$ against hypothesis $H_1: p \leq \theta - \delta$
- Observation: verify φ over sample trajectories generated using simulation

Acceptance sampling with fixed sample size

Single sampling plan: $\langle n, c \rangle$

- Generate n sample trajectories
- Accept H_0: $p \geq \theta + \delta$ iff more than c paths satisfy φ
- Pick n and c such that:
  - Probability of accepting H_1 when H_0 holds is at most α
  - Probability of accepting H_0 when H_1 holds is at most β

Sequential single sampling plan:

- Accept H_0 after $m < n$ observations if more than c observations are positive
- Accept H_1 after $m < n$ observations if at most k observations are positive and $k + (n - m) \leq c$

Graphical representation of sequential single sampling plan

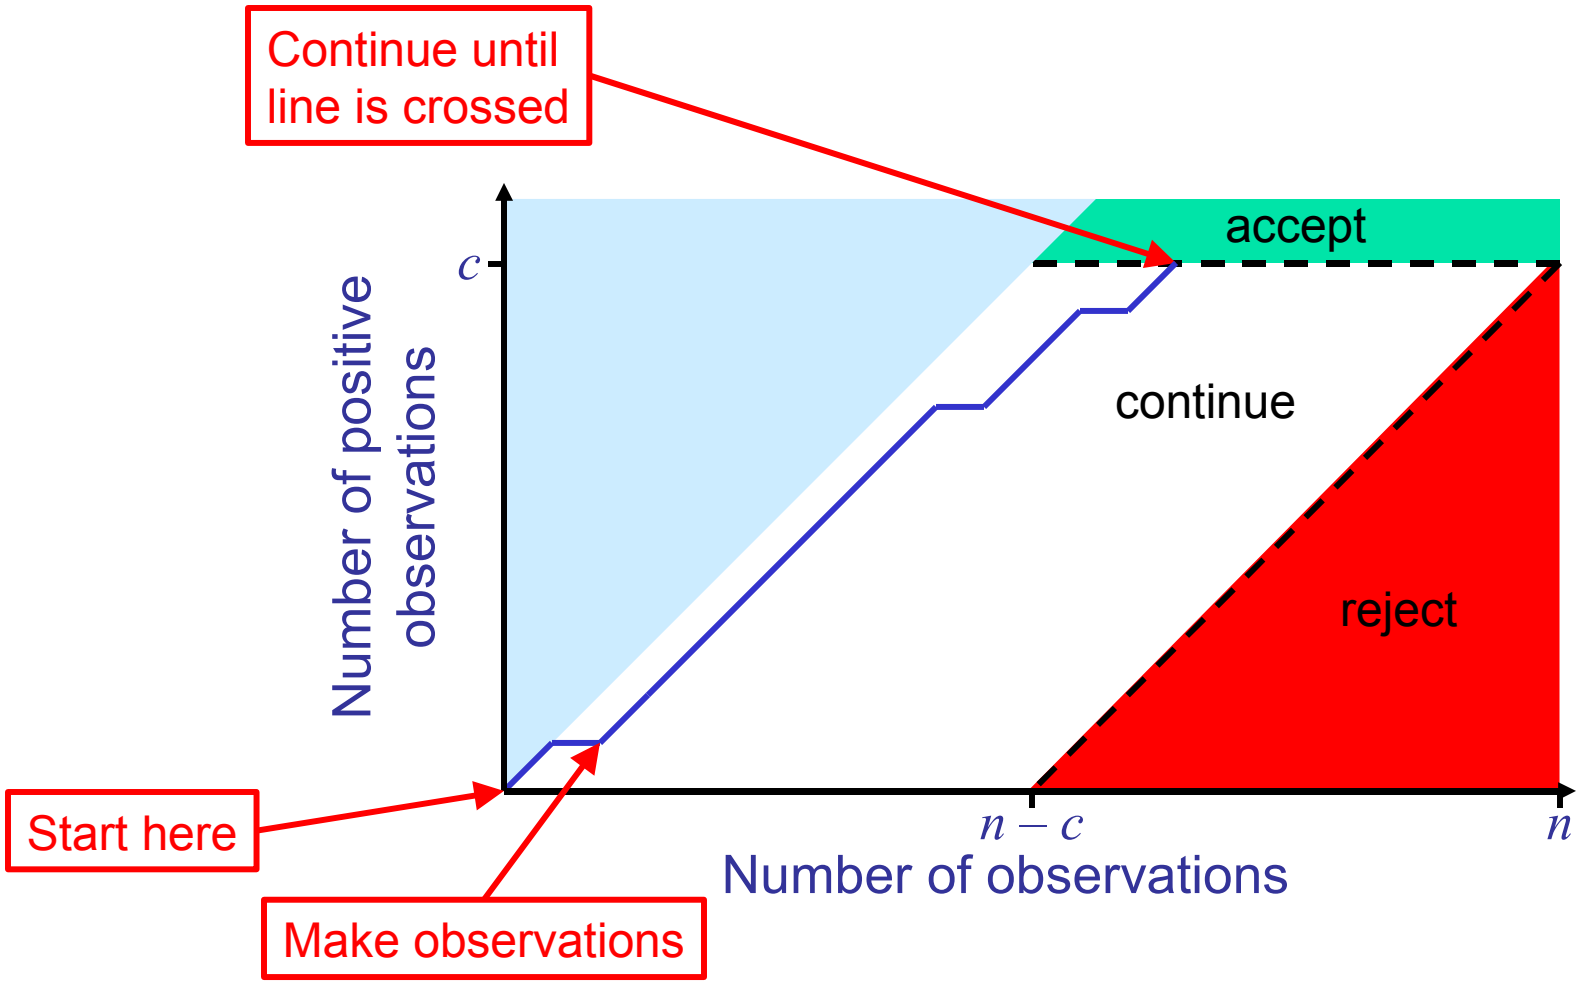

Sequential probability ratio test (SPRT)

Wald (Annals of Mathematical Statistics'45)

More efficient than sequential single sampling plan

- After m observations, k positive, compute ratio: $f = \frac{(\theta-\delta)^k(1-\theta+\delta)^{m-k}}{(\theta+\delta)^k(1-\theta-\delta)^{m-k}}$
- Accept $H_0: p \geq \theta + \delta$ if $f \leq \beta/(1-\alpha)$
- Accept $H_1: p \leq \theta - \delta$ if $f \geq (1-\beta)/\alpha$

No fixed upper bound on sample size, but much smaller on average

Graphical representation of SPRT

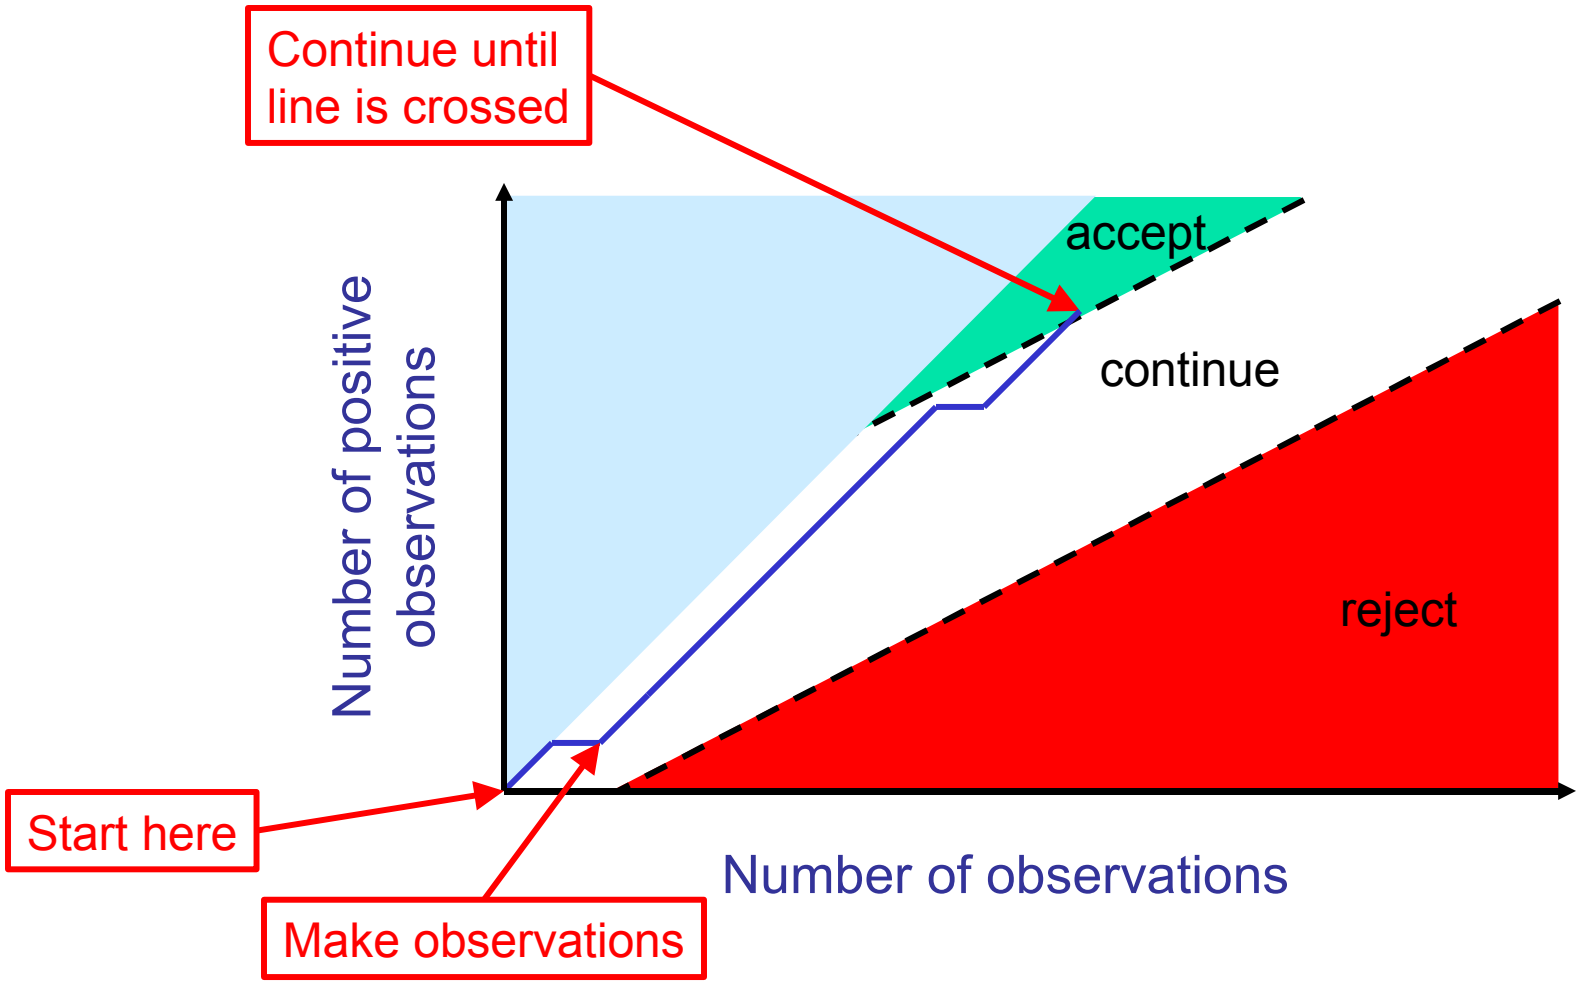

Statistical estimation

Hérault et al. (VMCAI'04)

Estimate p with confidence interval of width 2δ

- Accept H_0: $p \geq \theta + \delta$ iff center of confidence interval is at least θ
- Choosing sample size: $n = \left\lceil \frac{1}{2\delta^2} \log \frac{2}{\alpha} \right\rceil \quad \Rightarrow \quad \Pr\left[|\tilde{p} - p| < \delta\right] \geq 1 - \alpha$

Same as single sampling plan $\langle n, \lfloor n\theta + 1 \rfloor \rangle$; never more efficient!

θ	α	β	n_{est}	n_{opt}	$n_{\text{est}} / n_{\text{opt}}$
0.5	10^{-2}	10^{-2}	26,492	13,527	1.96
0.5	10^{-8}	10^{-2}	95,570	39,379	2.43
0.5	10^{-8}	10^{-8}	95,570	78,725	1.21
0.9	10^{-2}	10^{-2}	26,492	4,861	5.45
0.9	10^{-8}	10^{-2}	95,570	13,982	6.84
0.9	10^{-8}	10^{-8}	95,570	28,280	3.38

Acceptance sampling with undecided results

Younes (VMCAI'06)

Simultaneous acceptance sampling plans

- $H_0^{\perp}: p \geq \theta$ against $H_1^{\perp}: p \leq \theta - \delta$
- $H_0^{\top}: p \geq \theta + \delta$ against $H_1^{\top}: p \leq \theta$

Combining the results

- Accept $\mathcal{P}_{\geq\theta}[\varphi]$ if $H_0^{\top}$ and $H_0^{\perp}$ are accepted
- Reject $\mathcal{P}_{\geq\theta}[\varphi]$ if $H_1^{\top}$ and $H_1^{\perp}$ are accepted
- Undecided result otherwise

Graphical representation of SPRT with undecided results

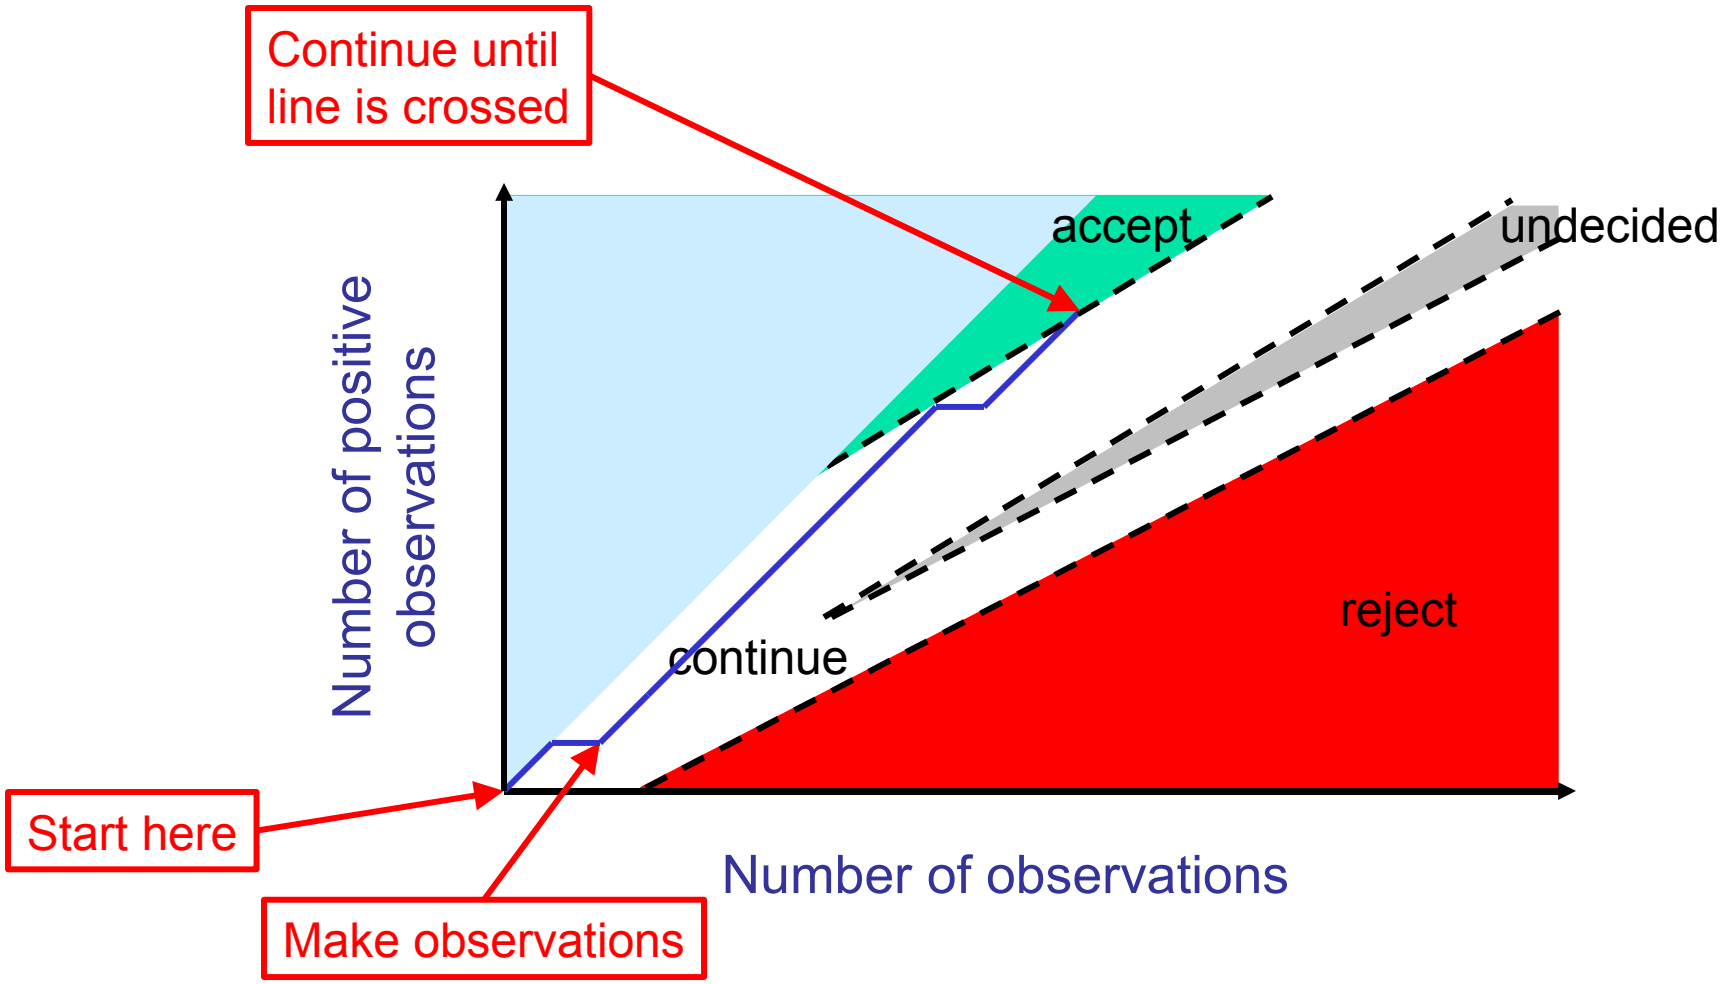

Unbounded until—avoiding infinite sample trajectories

Younes (unpublished manuscript)

Premature termination with probability p_t after each state transition

- Ensures finite sample trajectories
- Change value of positive sample trajectory ω from 1 to $(1 - p_t)^{-|\omega|}$
- Inspired by Monte Carlo method for matrix inversion by Forsythe & Leibler (1950)

Observations no longer 0 or 1: previous methods do not apply

- Use sequential estimation by Chow & Robbins (1965)
- Lower p_t means fewer samples, by longer trajectories

Note: Sen et al. (CAV'05) tried to handle unbounded until with termination probability, but flawed because observations are still 0 or 1

Empirical evaluation

Numerical vs. statistical (tandem queuing network)

Younes et al. (TACAS'04)

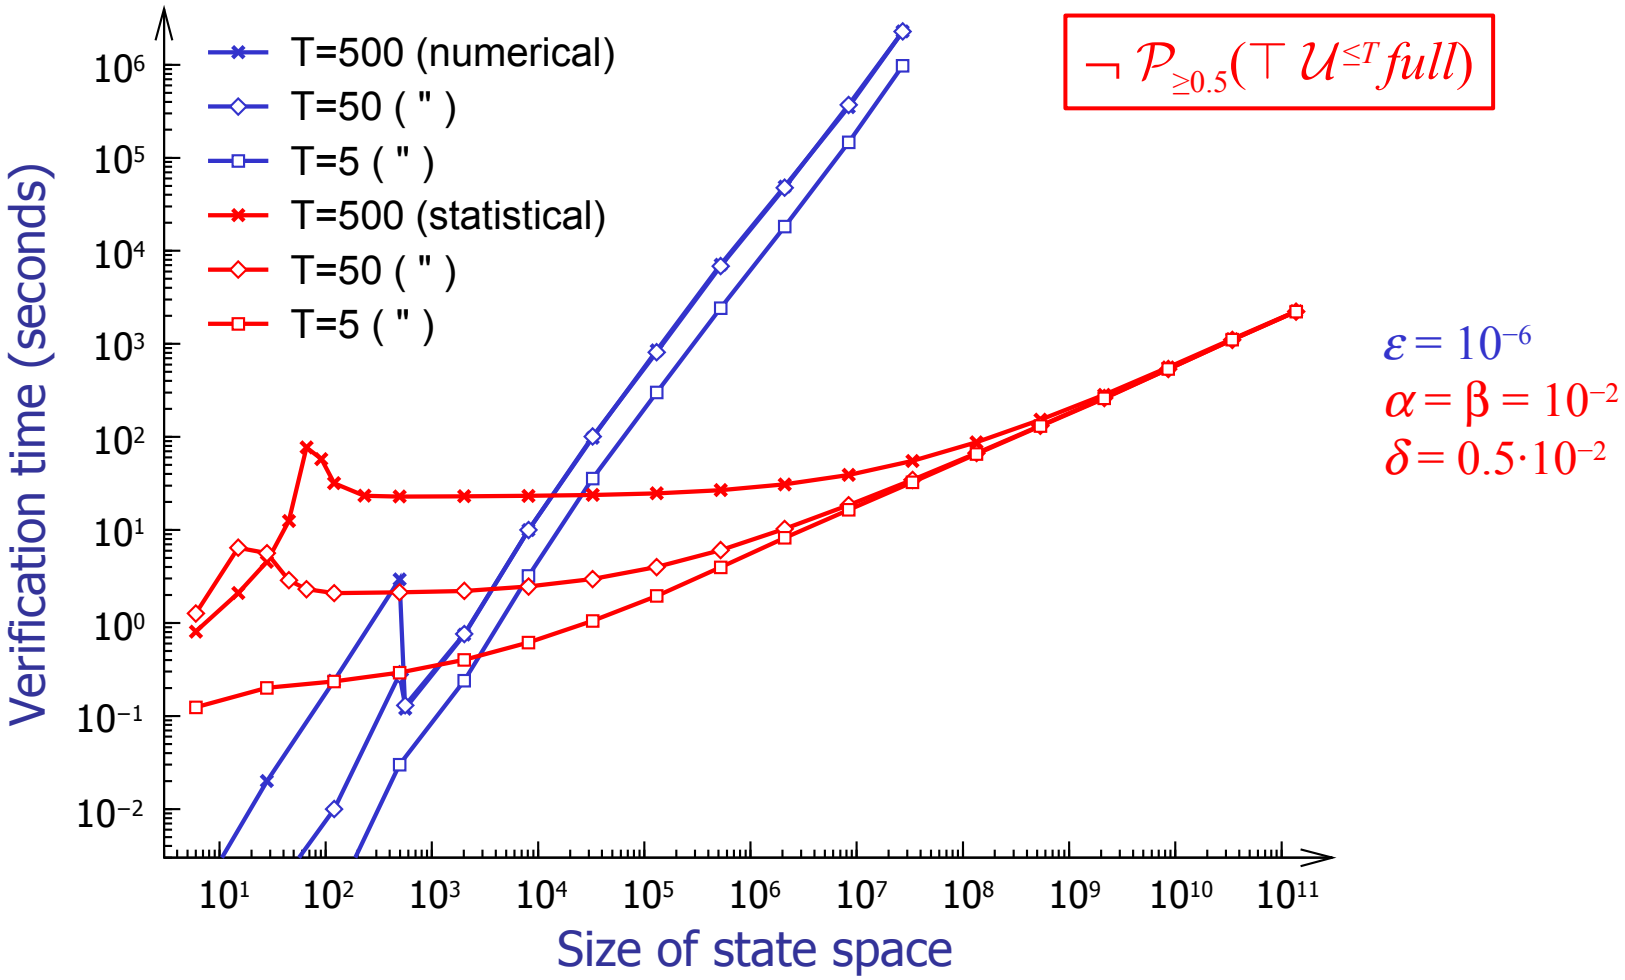

Numerical vs. statistical (symmetric polling system)

Younes et al. (TACAS'04)

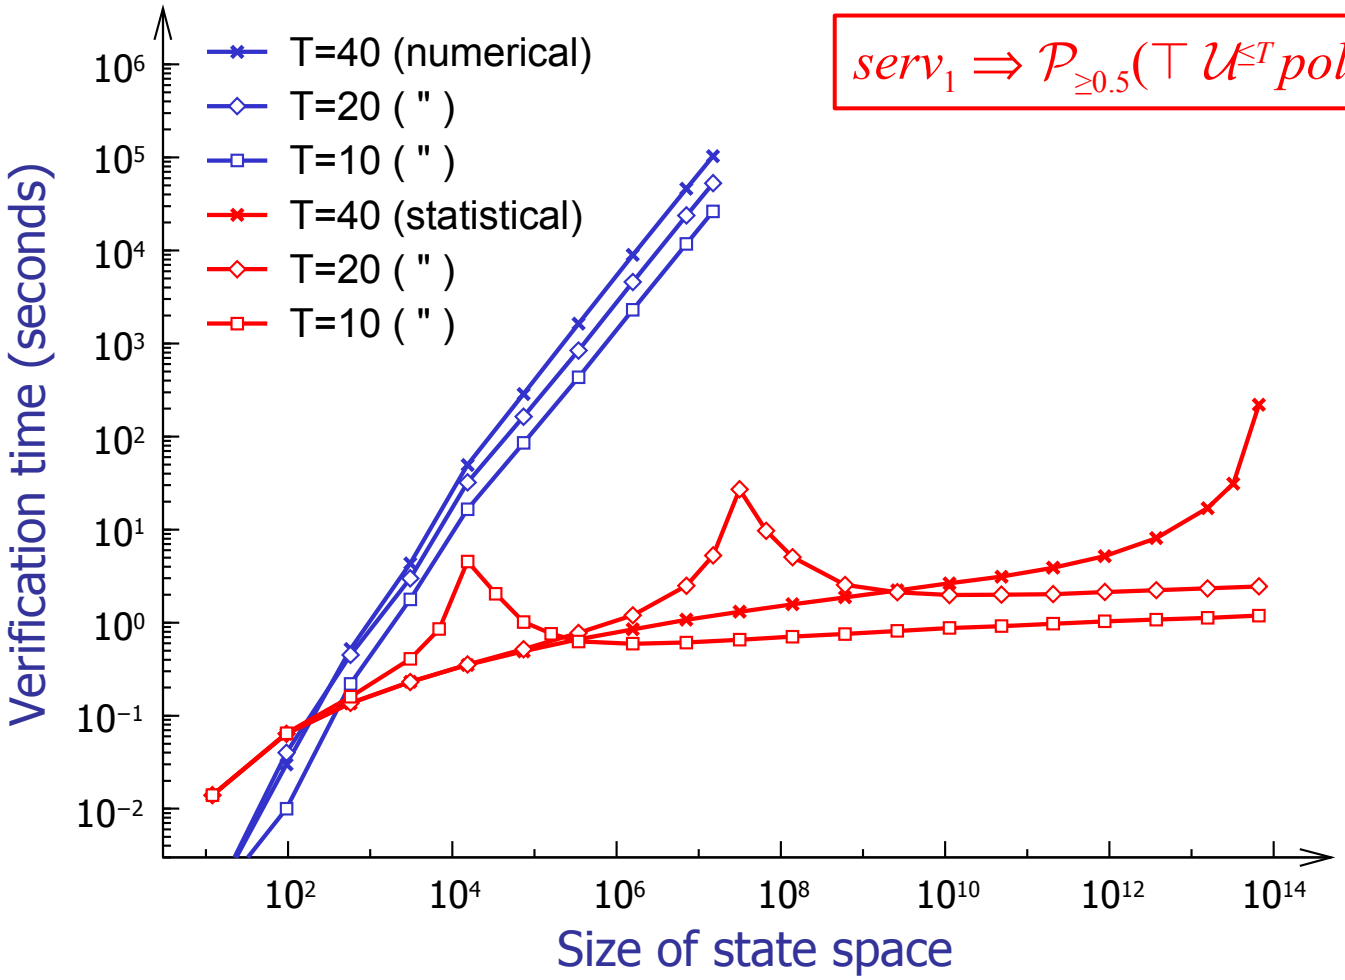

$$serv_1 \Rightarrow \mathcal{P}_{\geq 0.5}(\top\ \mathcal{U}^{\leq T}\ poll_1)$$

$\varepsilon = 10^{-6}$

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

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

Undecided results (symmetric polling system)

Younes (VMCAI'06)

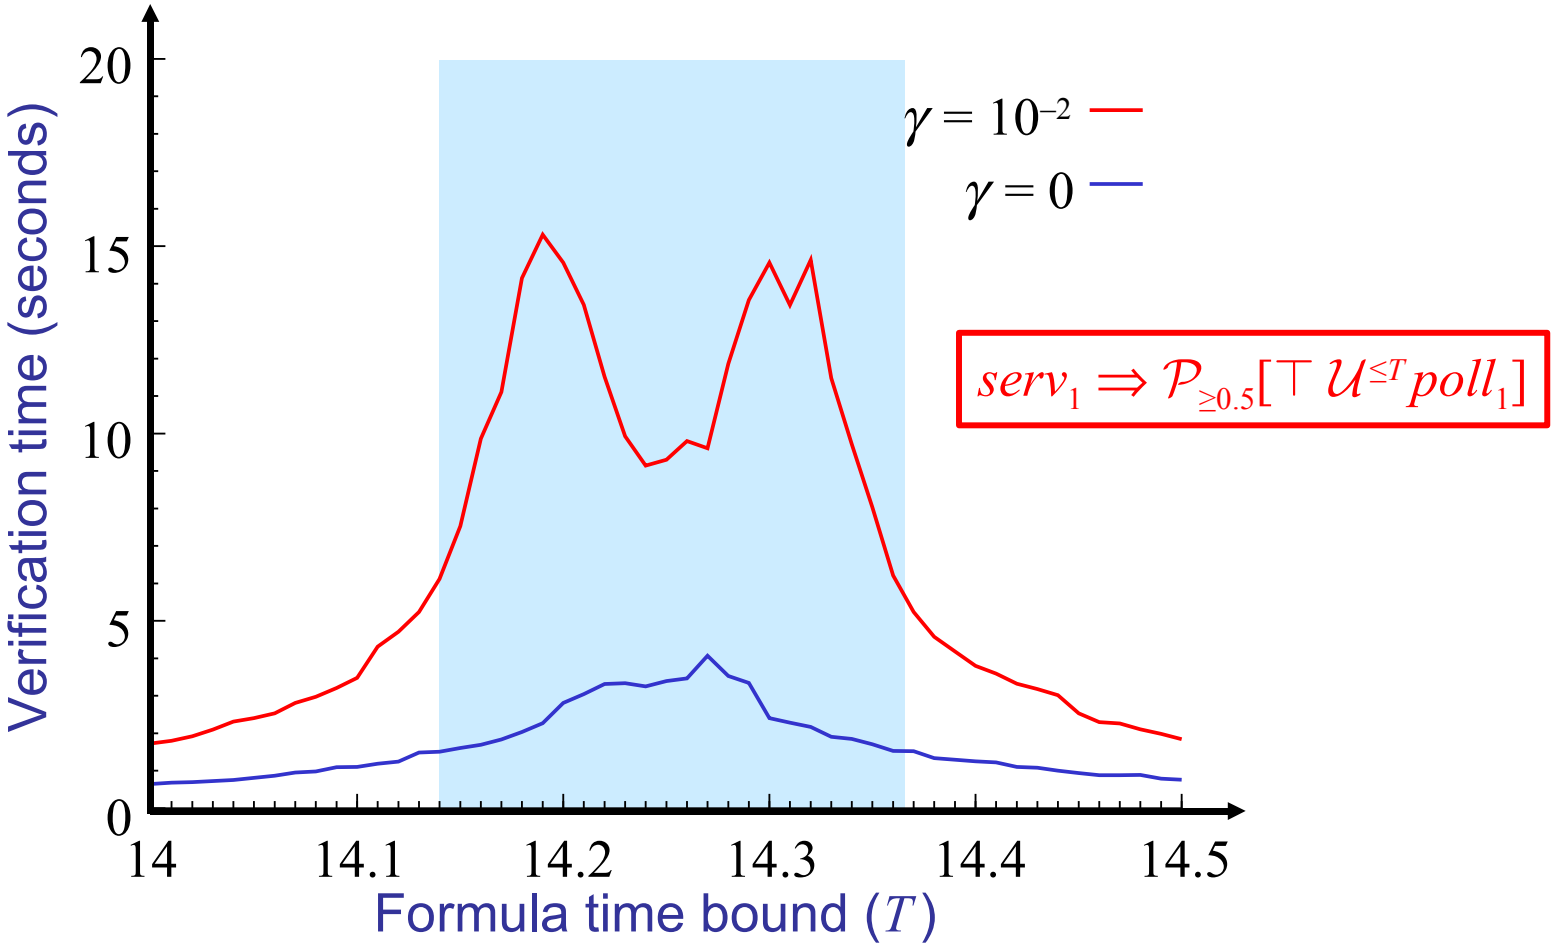

Undecided results (symmetric polling system)

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

result	14.10	14.15	14.20	14.25	14.30	14.35	14.40
accept	0	3	9	50	88	97	100
reject	100	97	91	50	12	3	0
accept	0	0	0	0	32	99	100
reject	100	99	42	1	0	0	0
undecided	0	1	58	99	68	1	0

Thank you!

Questions?