Error Control for Probabilistic Model Checking

Håkan L. S. Younes

Carnegie Mellon University

Contributions

- Framework for expressing correctness guarantees of model-checking algorithms
  - Enables comparison of different algorithms
  - Improves understanding of sampling-based algorithms
- New sampling-based algorithm for probabilistic model checking
  - Better error control through undecided results

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

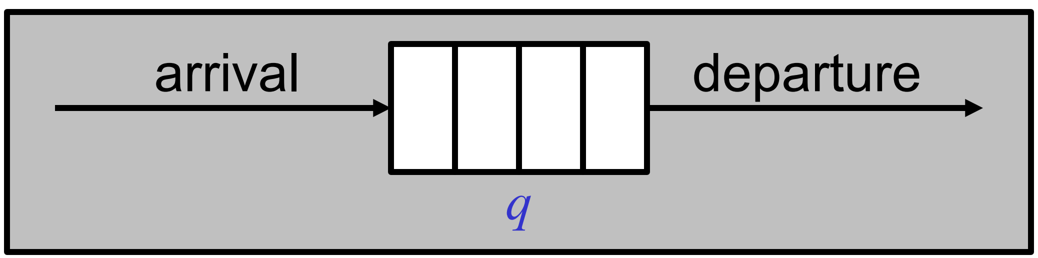

"The probability is at least 0.1 that the queue becomes full within 5 minutes"

Temporal Stochastic Logic (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
- 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

Property Example

- "The probability is at least 0.1 that the queue becomes full within 5 minutes"
  - $\mathcal{P}_{\geq 0.1}[\top \, \mathcal{U}^{\leq 5} \, \mathit{full}\,]$

Possible Results of Model Checking

- 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 $(s \vdash_{\perp} \Phi)$
  - Return an undecided result $(s \vdash_{\mathrm{I}} \Phi)$
- An error occurs if:
  - $\mathcal{A}$ rejects Φ when Φ is true (false negative)
  - $\mathcal{A}$ accepts Φ when Φ is false (false positive)

Ideal Error Control

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

Unrealistic Expectations

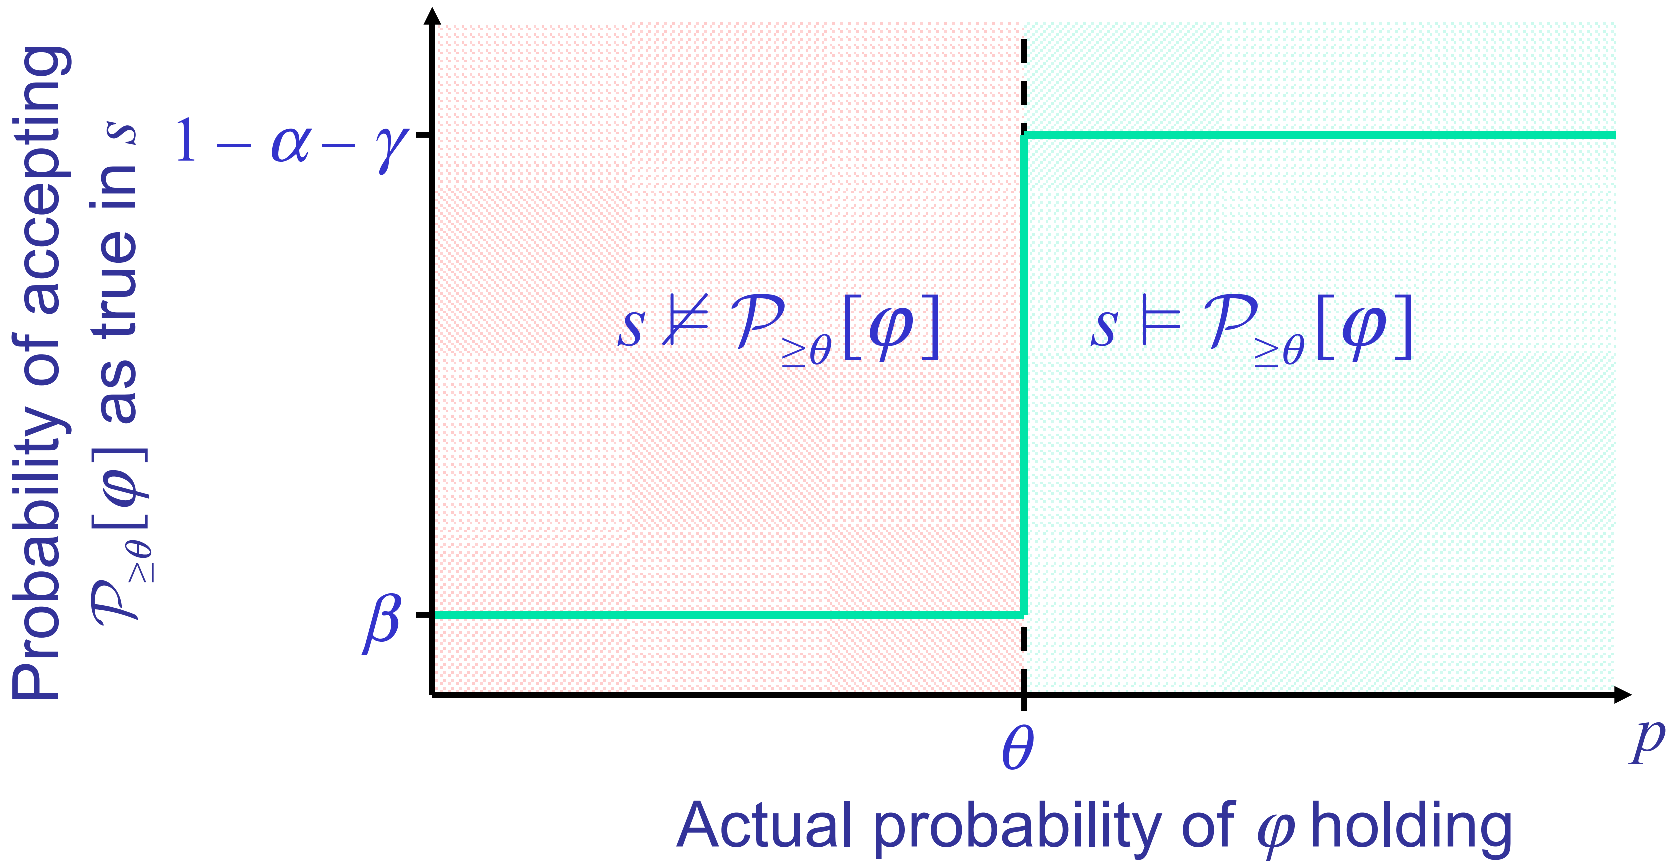

Temporal Stochastic Logic with Indifference Regions (CSL$_\delta$)

- 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 $\theta - \delta$
  - "Too close to call" if probability is within δ distance of θ

Error Control for Current Solution Methods

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

Probabilistic Model Checking with Indifference Regions

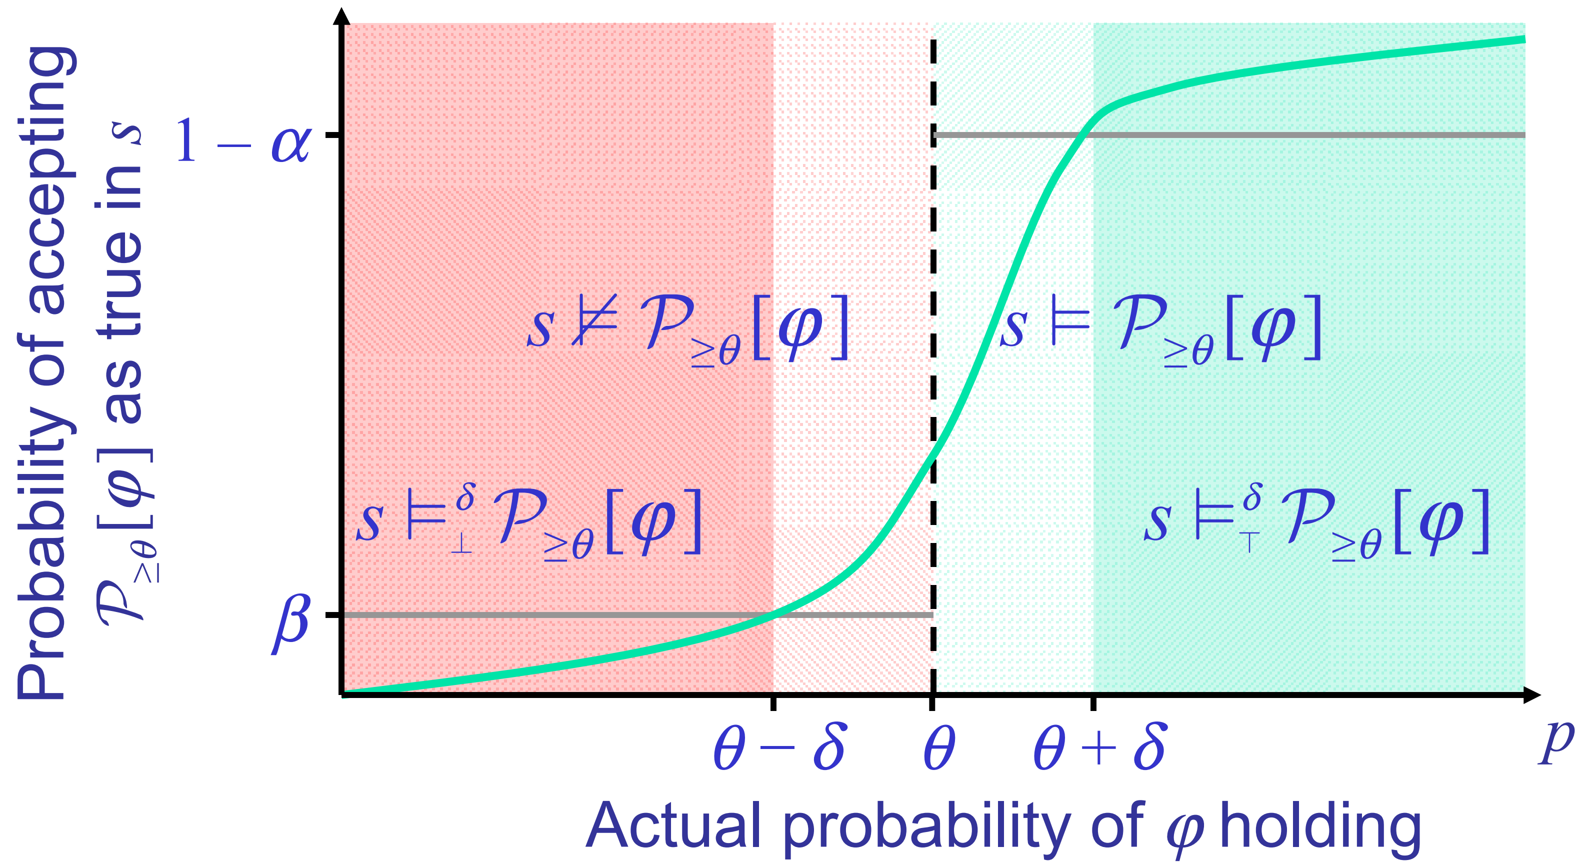

Hypothesis Testing
Younes & Simmons (CAV'02)

- Single sampling plan: $\langle n, c \rangle$
  - Generate n sample execution paths
  - Accept $\mathcal{P}_{\geq\theta}[\varphi]$ iff more than c paths satisfy φ
  - Probability of accepting $\mathcal{P}_{\geq\theta}[\varphi]$ as true:

$$1 - F(c; n, p) = 1 - \sum_{i=0}^{c} \binom{n}{i} p^i (1-p)^{n-i}$$

- Sequential acceptance sampling

Statistical Estimation
Hérault et al. (VMCAI'04)

- Estimate p using sample of size n: $\tilde{p} = \frac{1}{n}\sum_{i=1}^{n} x_i$
- Choosing n:

$$n = \left\lceil \frac{1}{2\delta^2} \log \frac{2}{\alpha} \right\rceil \quad \Rightarrow \quad \Pr\left[\left| \tilde{p} - p \right| < \delta \right] \geq 1 - \alpha$$

- Acceptance condition for $\mathcal{P}_{\geq\theta}[\varphi]$: $\tilde{p} \geq \theta$

$\boxed{\text{Same as single sampling plan } \langle n, \lfloor n\theta + 1 \rfloor \rangle\text{!}}$

Statistical Estimation vs. Hypothesis Testing

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

Numerical Transient Analysis
Baier et al. (CAV'00)

- Estimate p with truncation error ε:

$$\tilde{p} \leq p \leq \tilde{p} + \varepsilon$$

- Acceptance condition for $\mathcal{P}_{\geq\theta}[\varphi]$: $\tilde{p} + \frac{\varepsilon}{2} \geq \theta$
  - $\Pr[s \vdash_{\perp} \Phi \mid s \models_{\top}^{\delta} \Phi] = 0$
  - $\Pr[s \vdash_{\top} \Phi \mid s \models_{\perp}^{\delta} \Phi] = 0$

$$\delta = \frac{\varepsilon}{2}$$

Alternative Error Control

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

Probabilistic Model Checking with Undecided Results

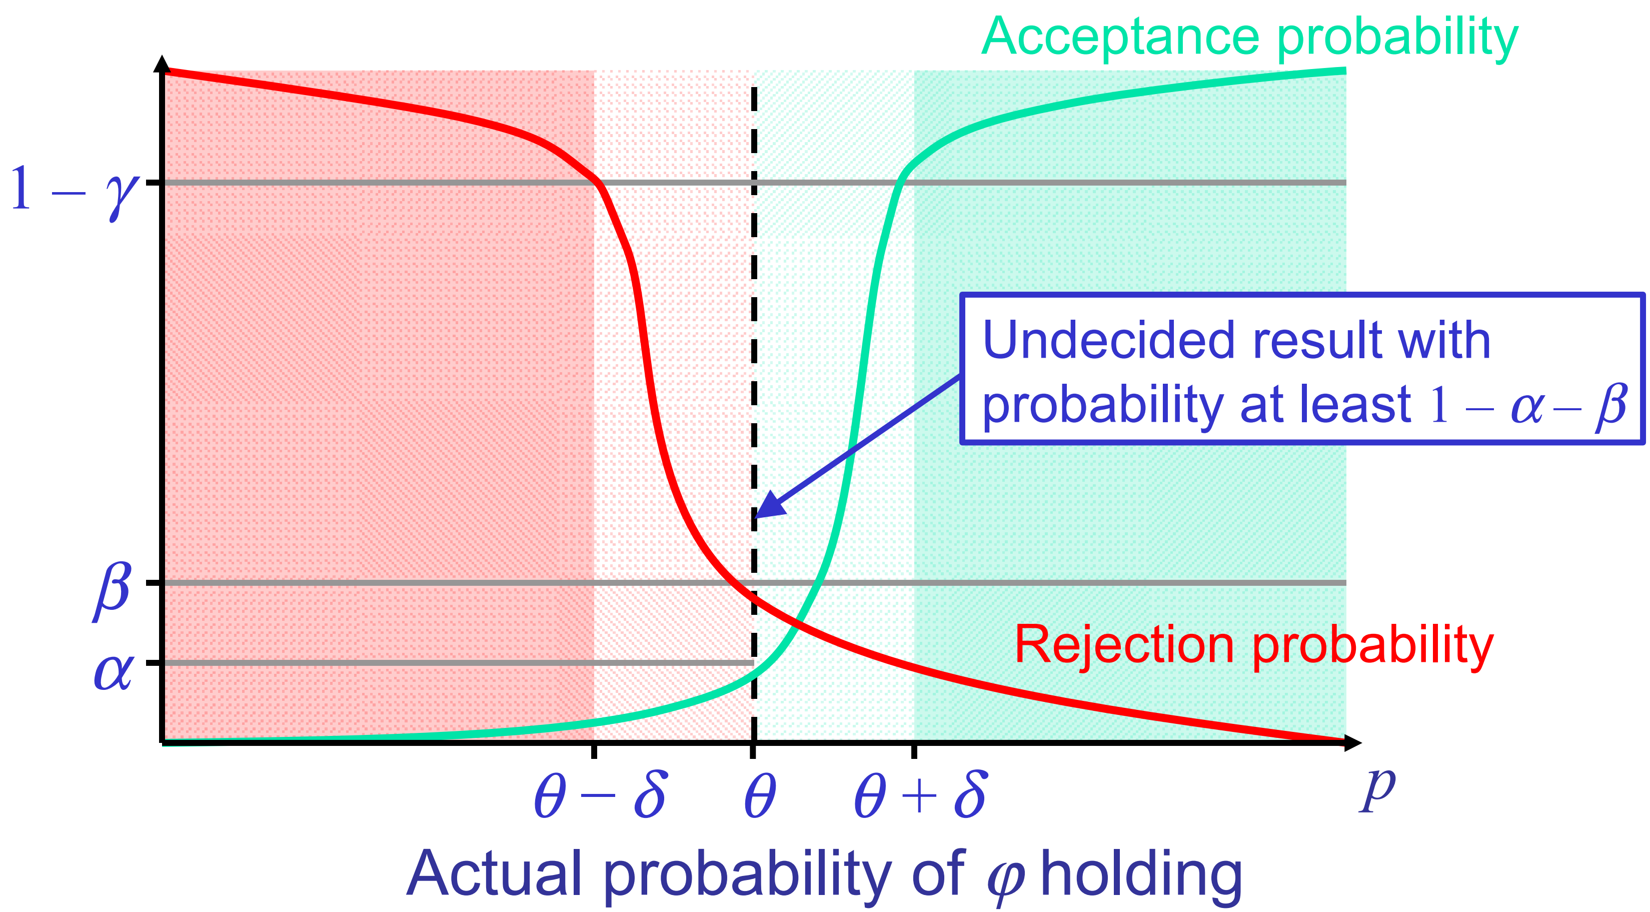

Statistical Solution Method

- 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

Empirical Evaluation (Symmetric Polling System)

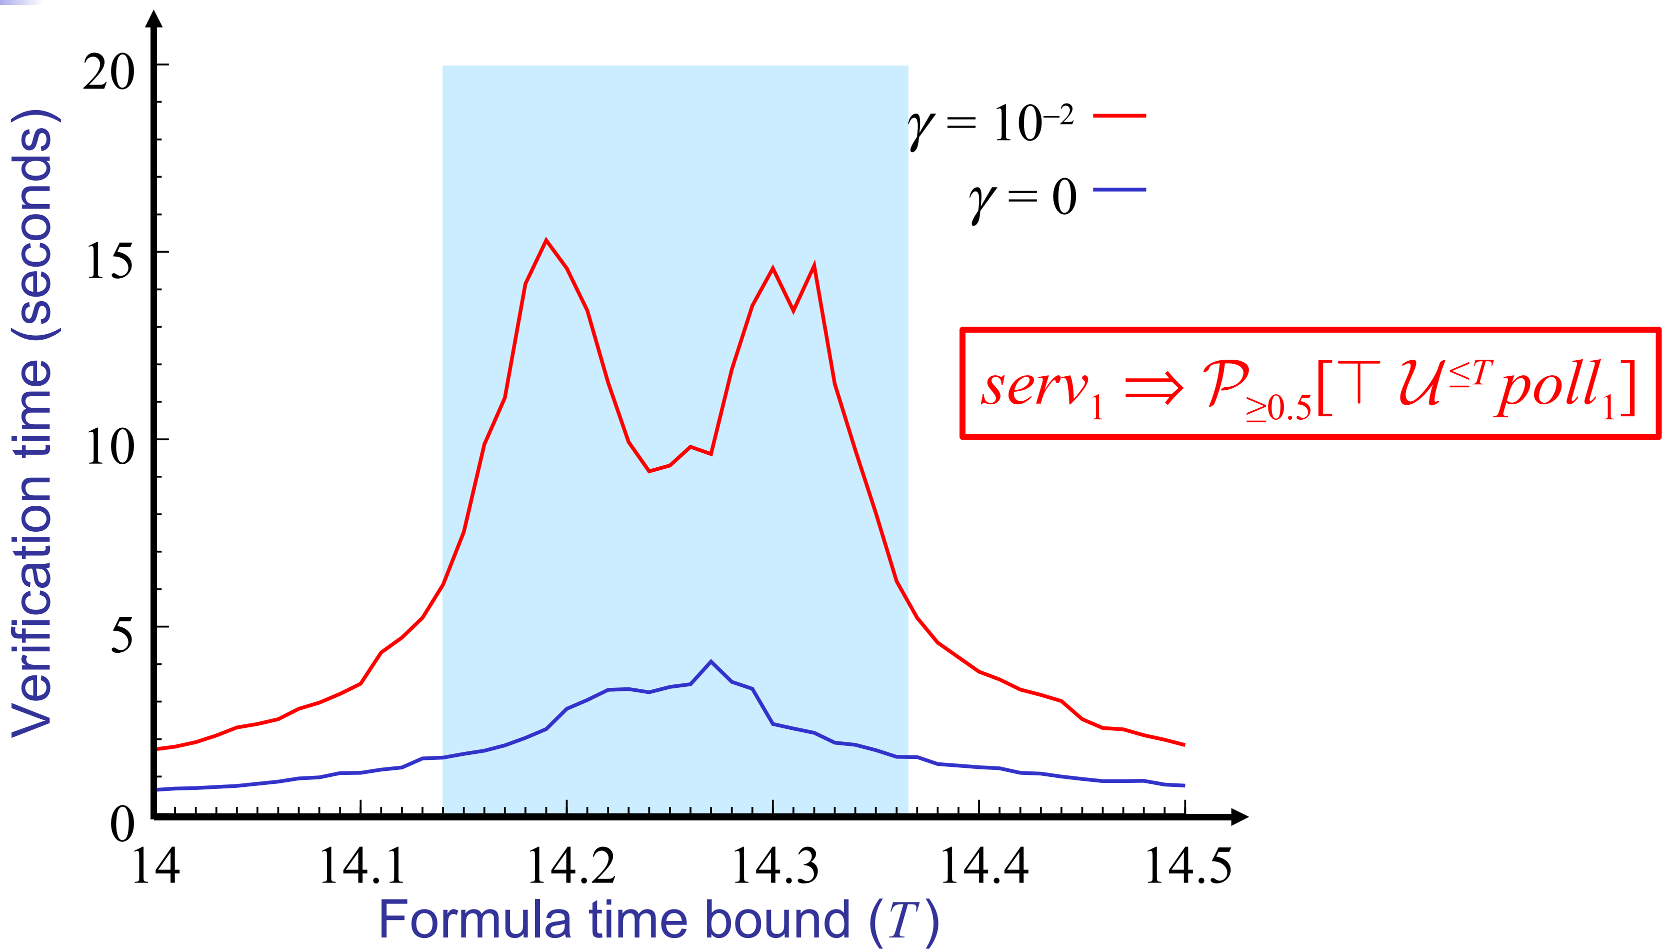

Empirical Evaluation (Symmetric Polling System)

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

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

Summary

- Statistical estimation is never more efficient than hypothesis testing
- Statistical methods are randomized algorithms for CSL_δ model checking
- Numerical methods are exact algorithms for CSL_δ model checking
- New statistical solution method with finer error control (γ parameter)