Probabilistic Plan Verification through Acceptance Sampling

Håkan L. S. Younes
Carnegie Mellon University

David J. Musliner
Honeywell Laboratories

Introduction

- Probabilistic extension to CIRCA
- Efficient plan verification algorithm
  - Monte Carlo simulation
  - Acceptance sampling
- Guaranteed error bounds

Planning via Model Checking

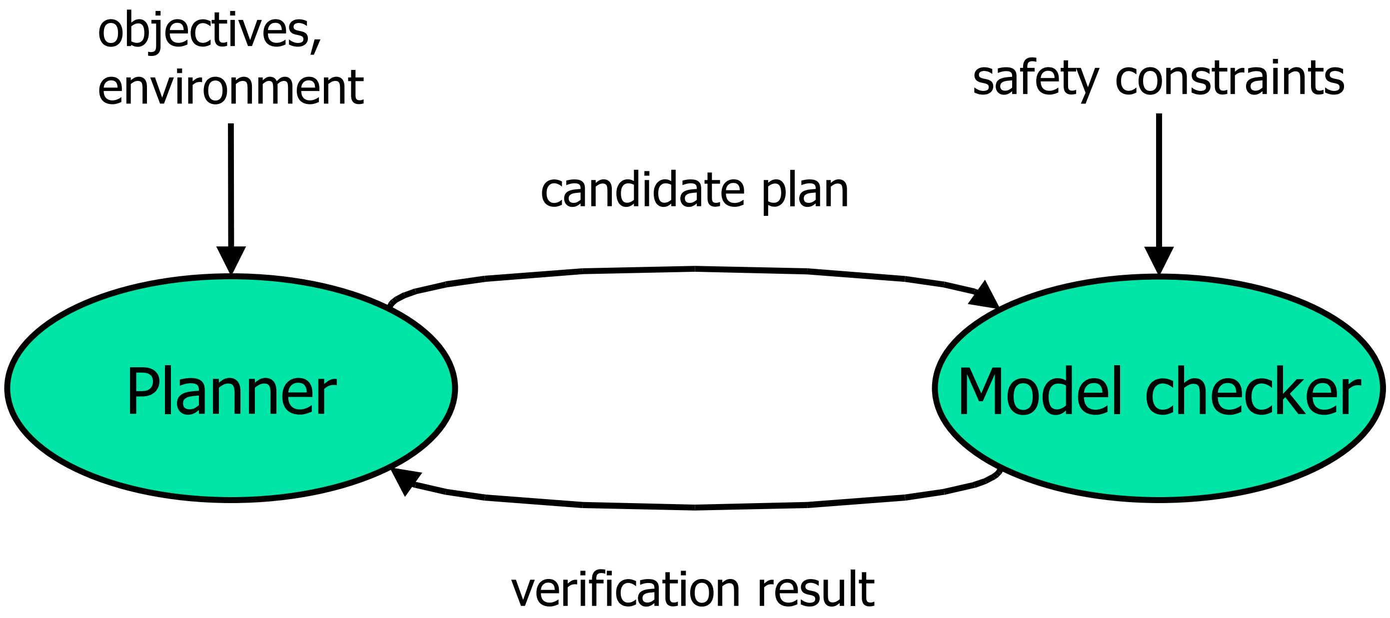

World Model

- States…

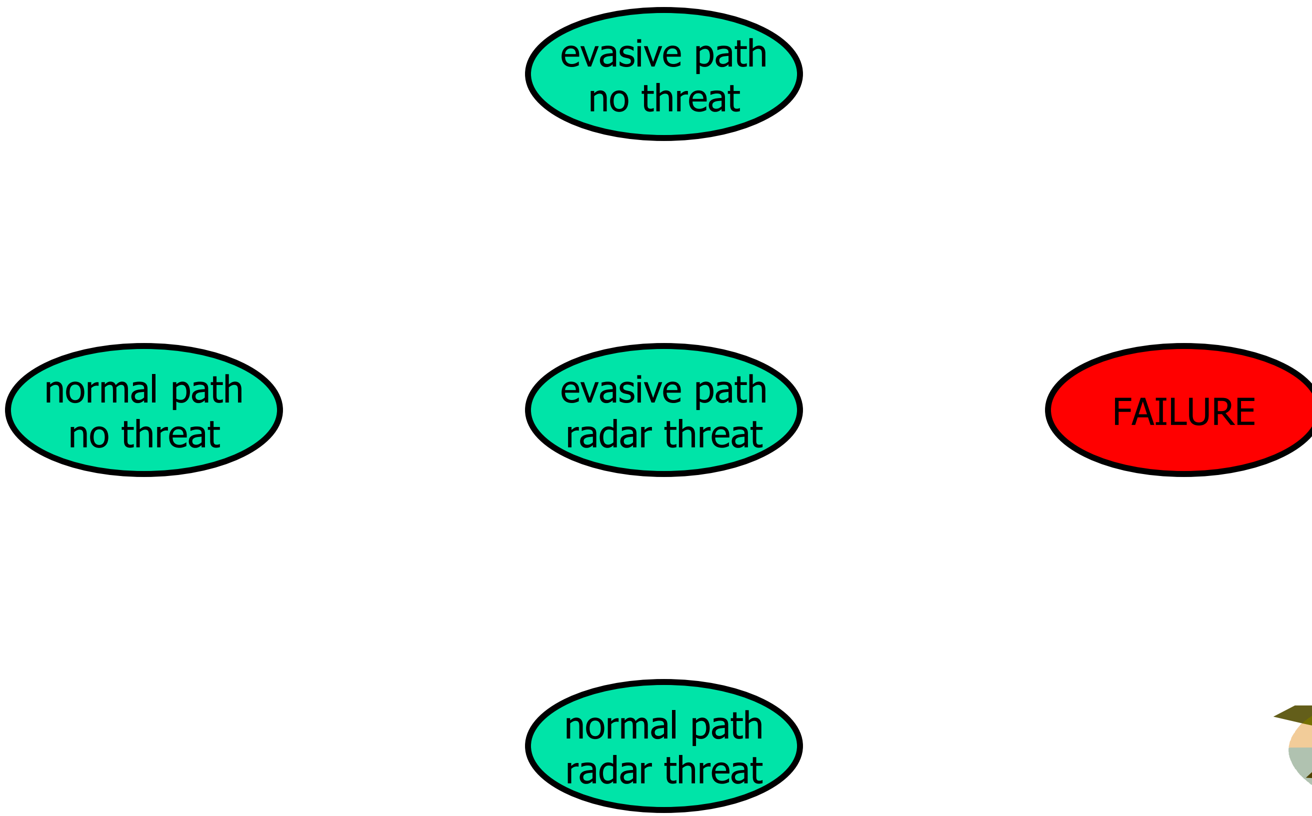

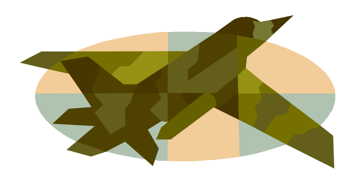

World Model

- States + events = environment

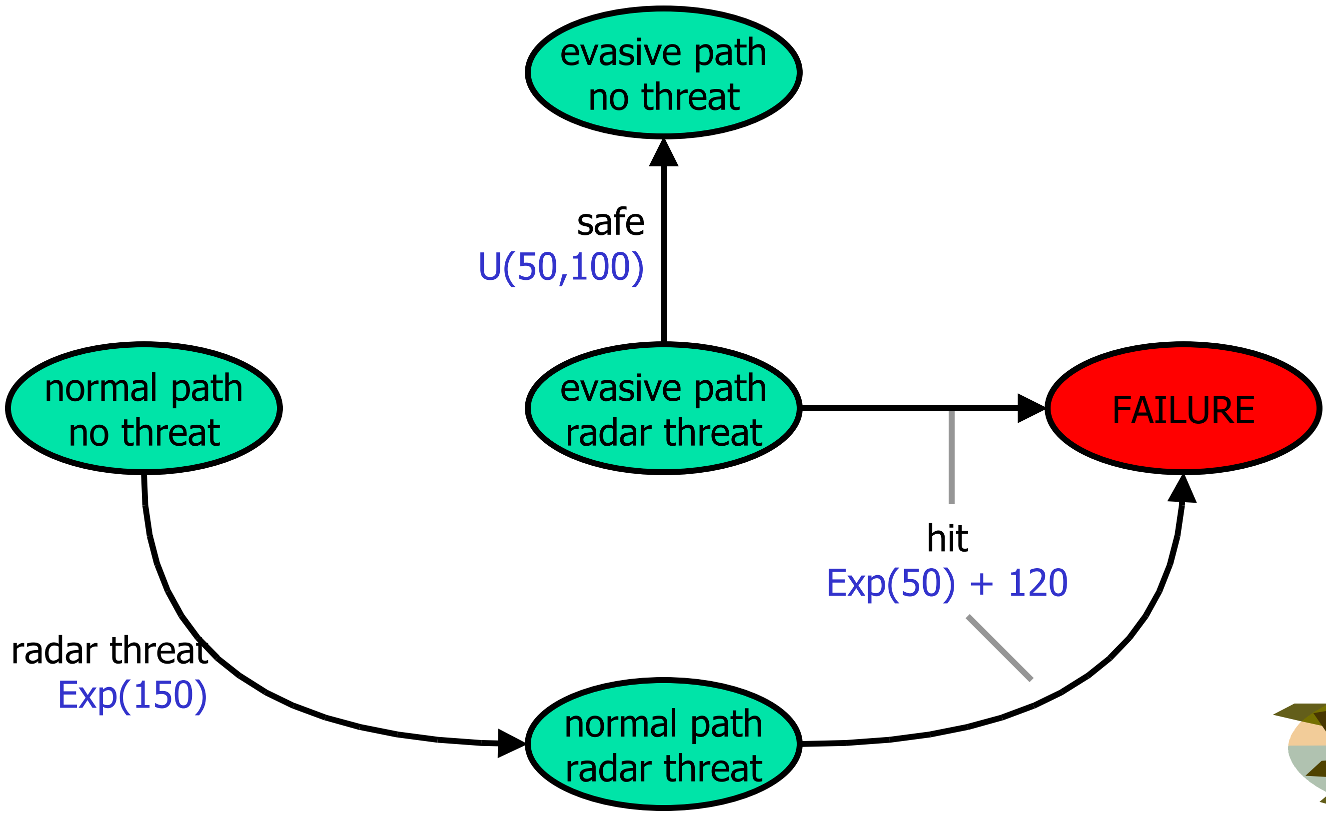

World Model

- A plan maps states to actions

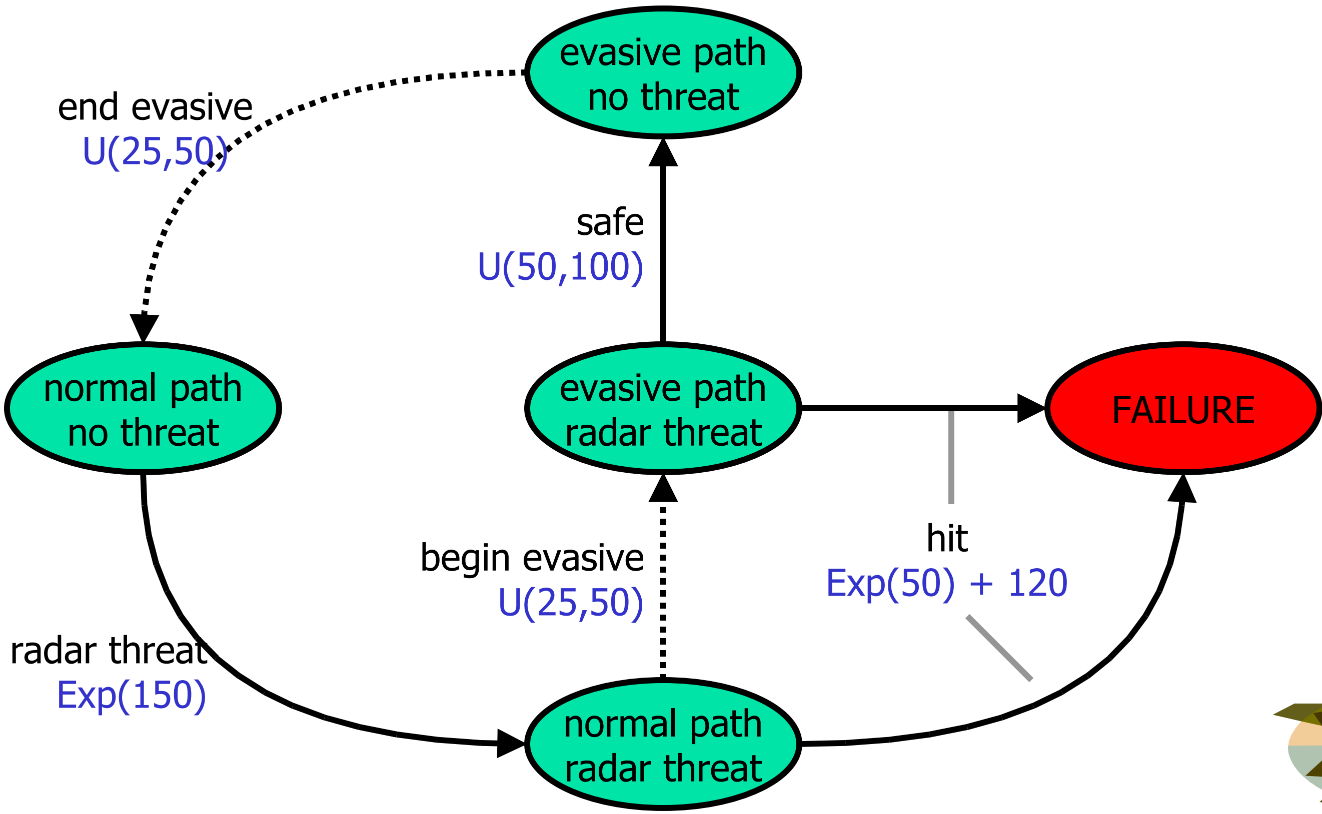

Sample Execution Paths

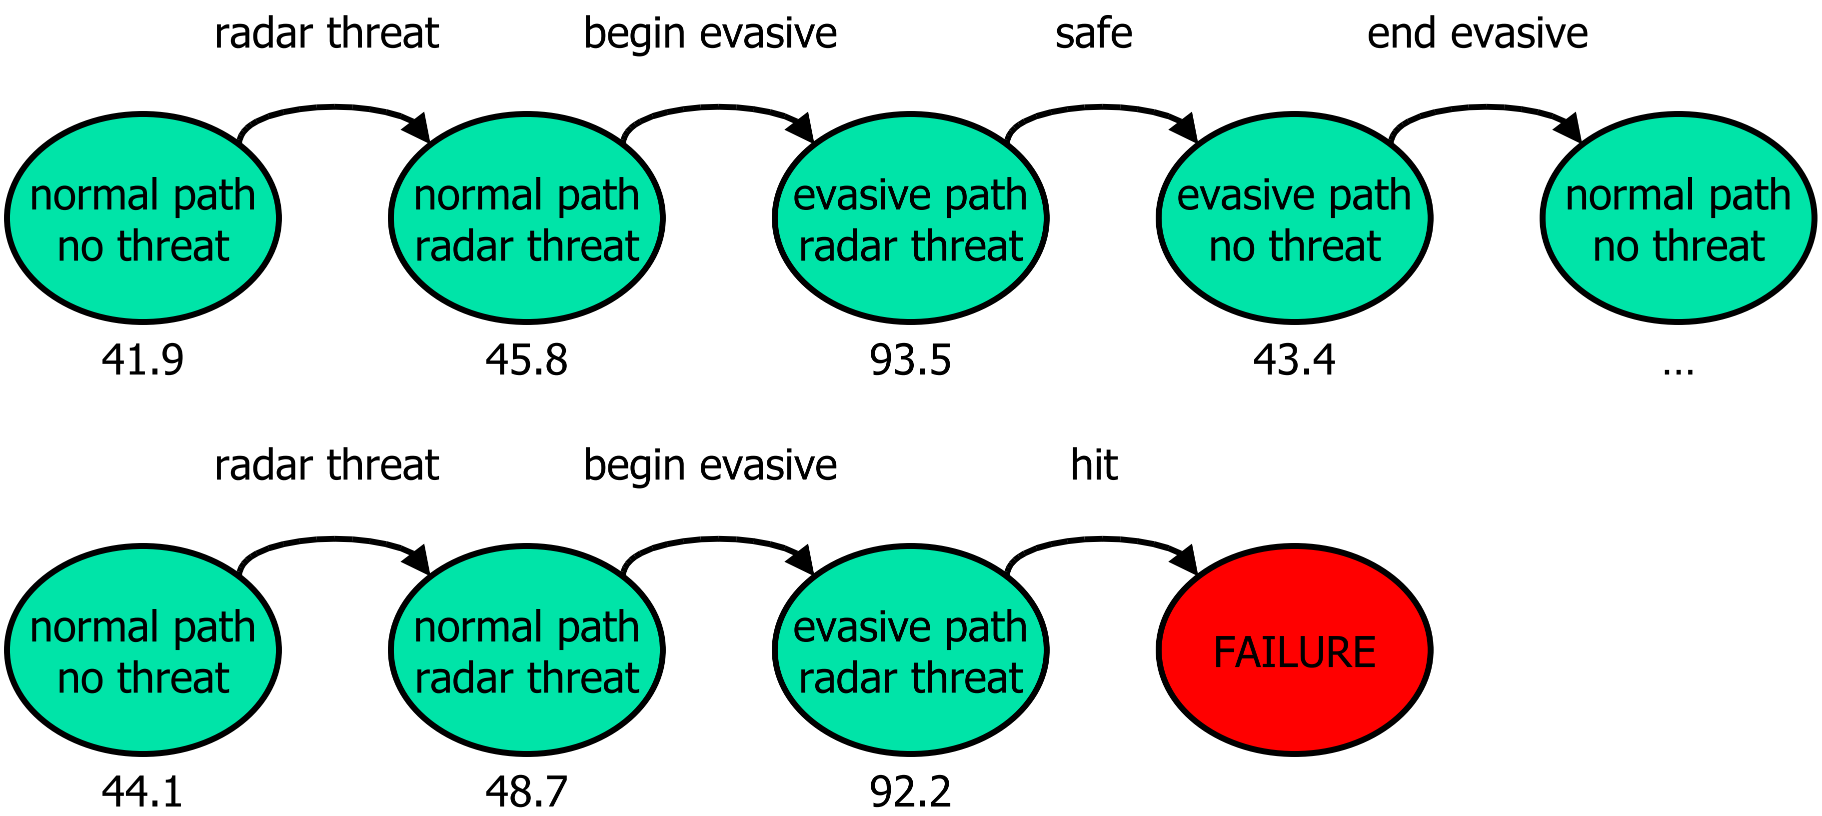

Plan Safety

- Two parameters
  - Failure probability threshold: θ
  - Maximum execution time: t_{max}
- A plan is safe if the probability of reaching a failure state within t_{max} time units is at most θ

Safety Over Sample Execution Paths

- Given $t_{max} = 200$:

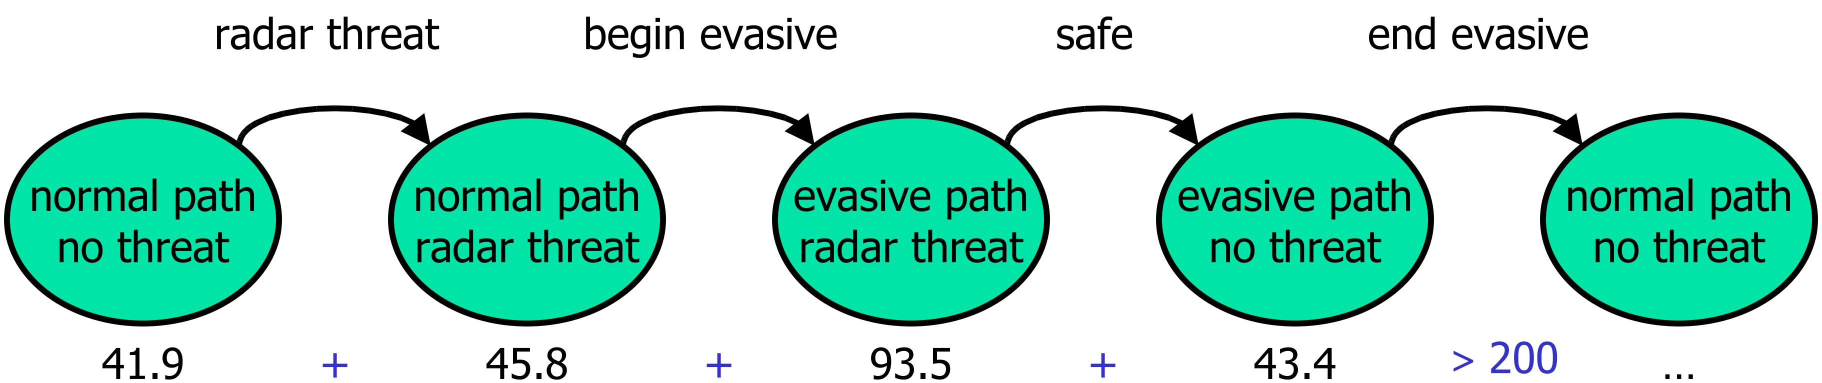

Safe!

Safety Over Sample Execution Paths

- Given $t_{max} = 200$:

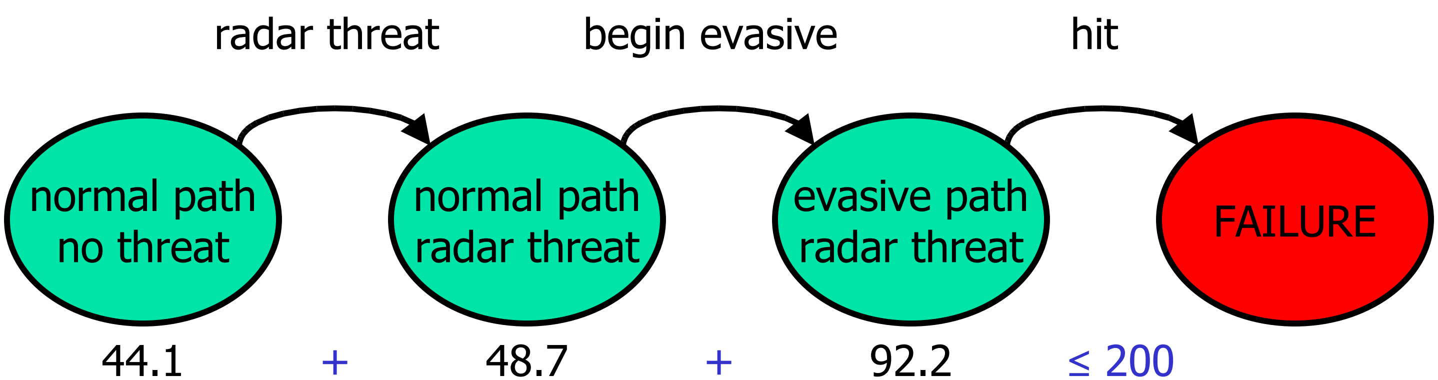

Not safe!

(safe if $t_{max} < 185$)

Verifying Plan Safety

- Symbolic Methods
  - Pro: Exact solution
  - Con: Works only for restricted class of models
- Sampling
  - Pro: Works for any model that can be simulated
  - Con: Uncertainty in correctness of solution

Our Approach

- Use simulation to generate sample execution paths
- Use sequential acceptance sampling to verify plan safety

Error Bounds

- Probability of false negative: $\leq\alpha$
  - We say that a plan is not safe when it is
- Probability of false positive: $\leq\beta$
  - We say that a plan is safe when it is not

Acceptance Sampling

- Test hypothesis $Pr_{\leq\theta}(X)$
- In our case
  - θ is the failure probability threshold
  - X is the proposition that a failure state is reached within the time limit

Sequential Acceptance Sampling

- Test hypothesis $Pr_{\leq\theta}(X)$

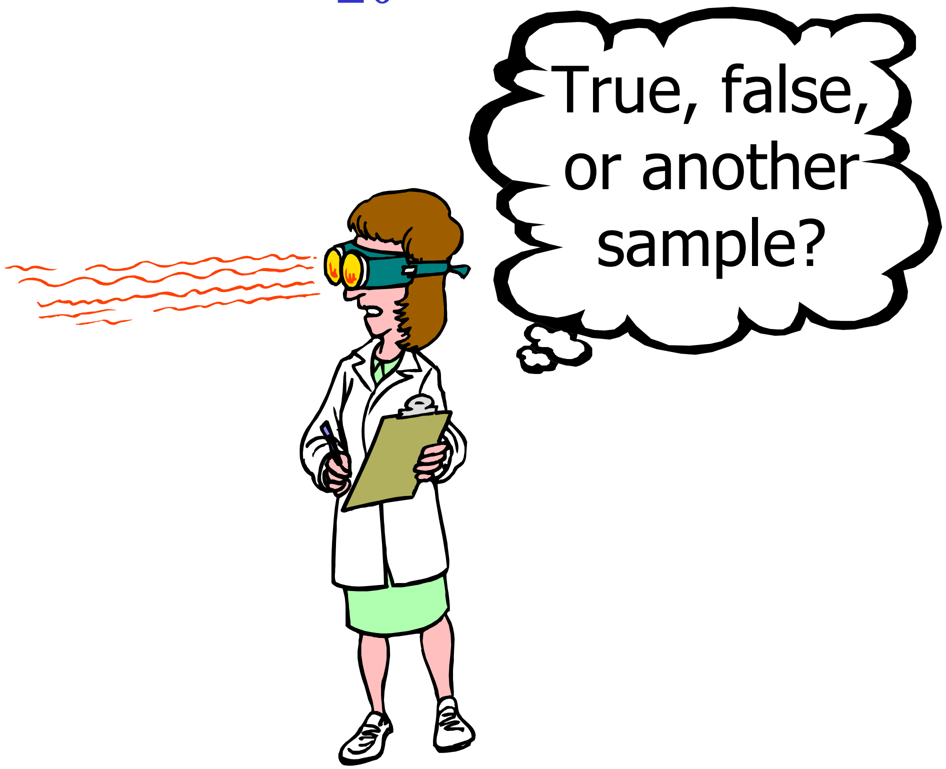

Performance of Test

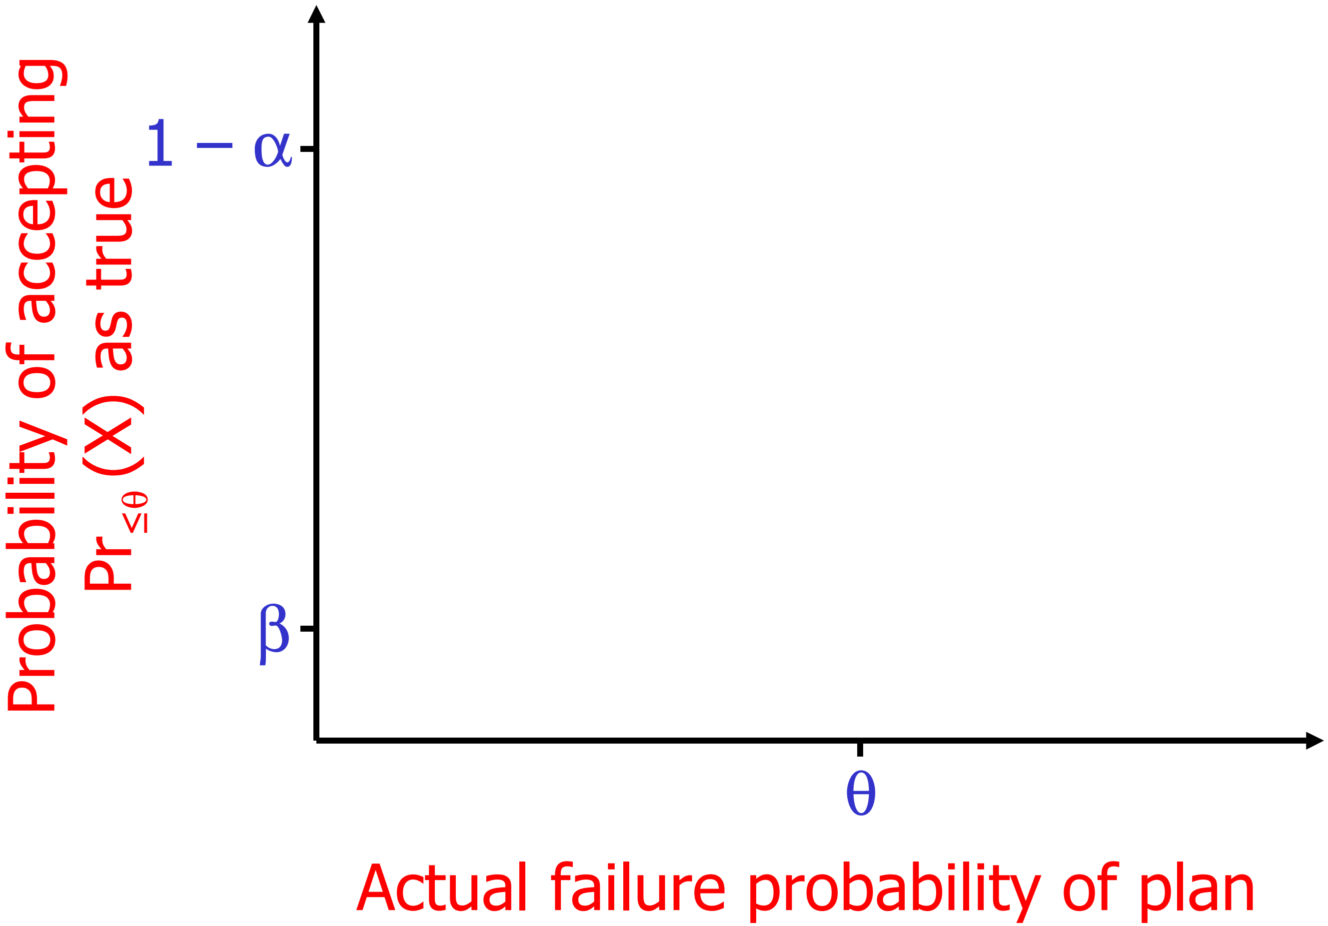

Ideal Performance

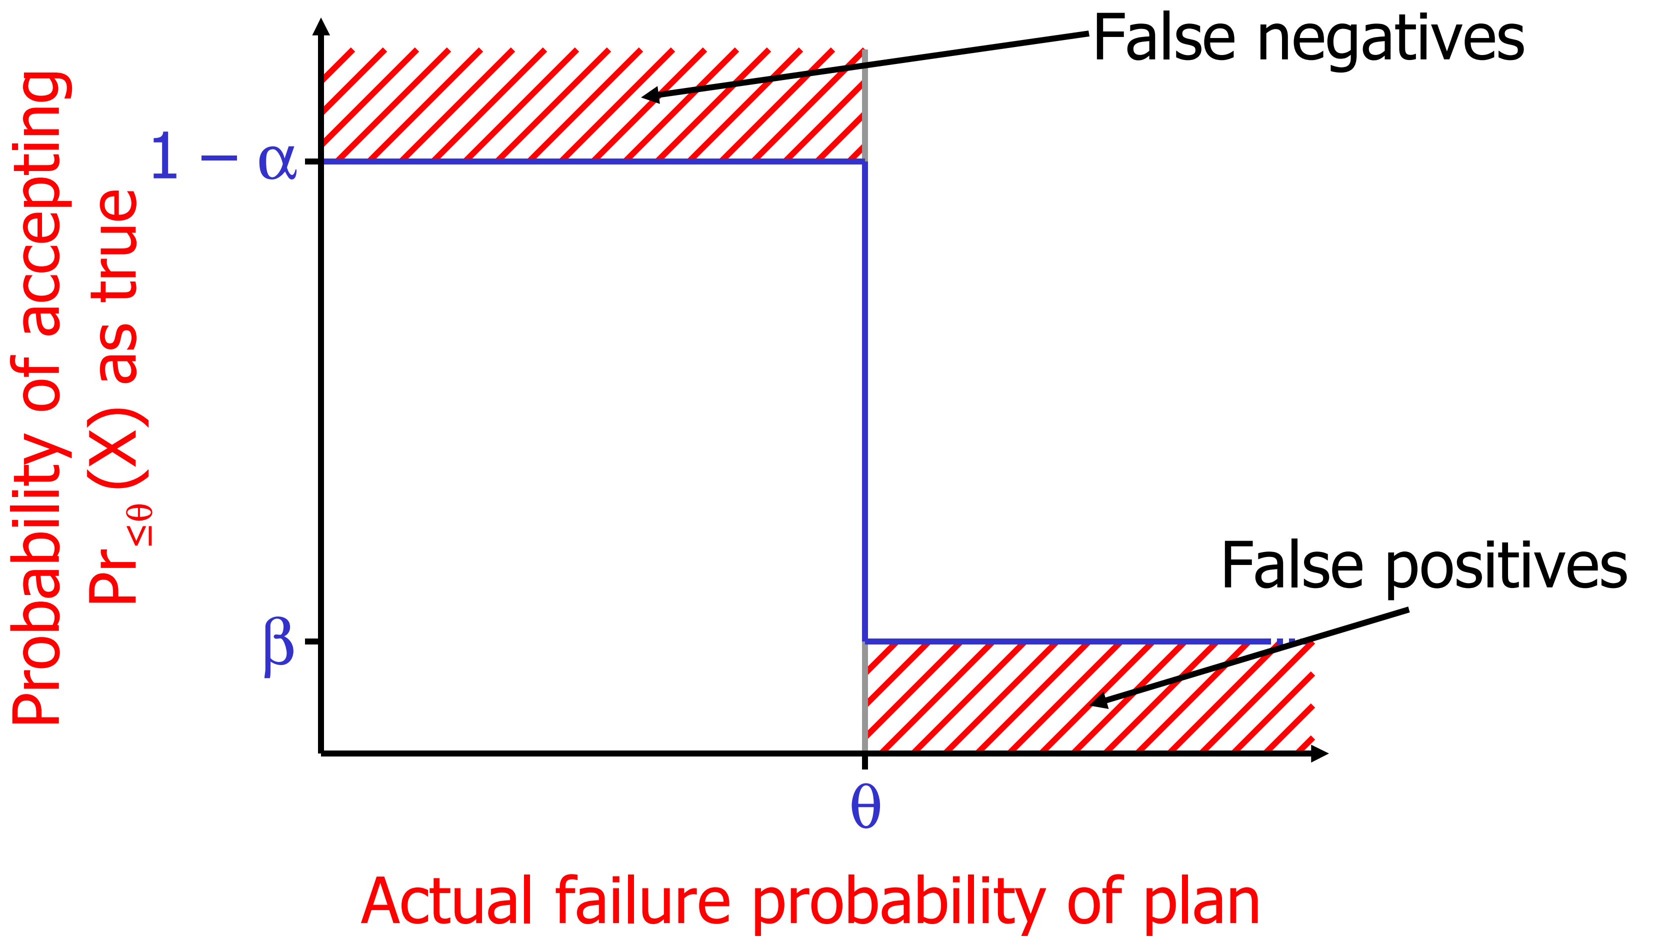

Actual Performance

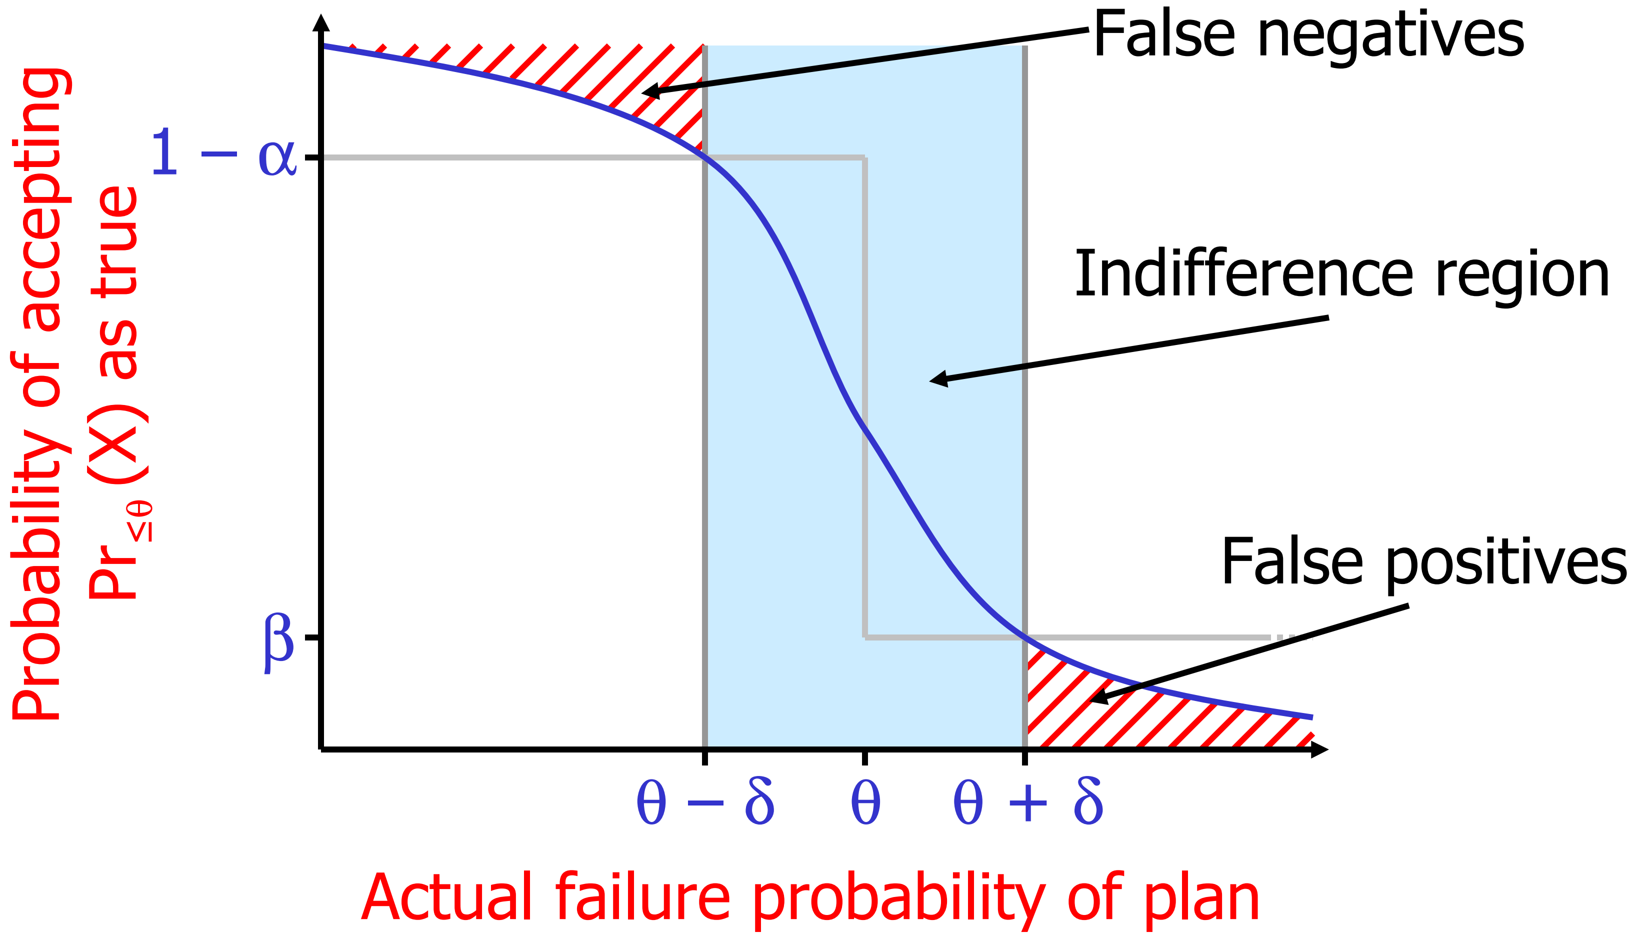

Graphical Representation of Sequential Test

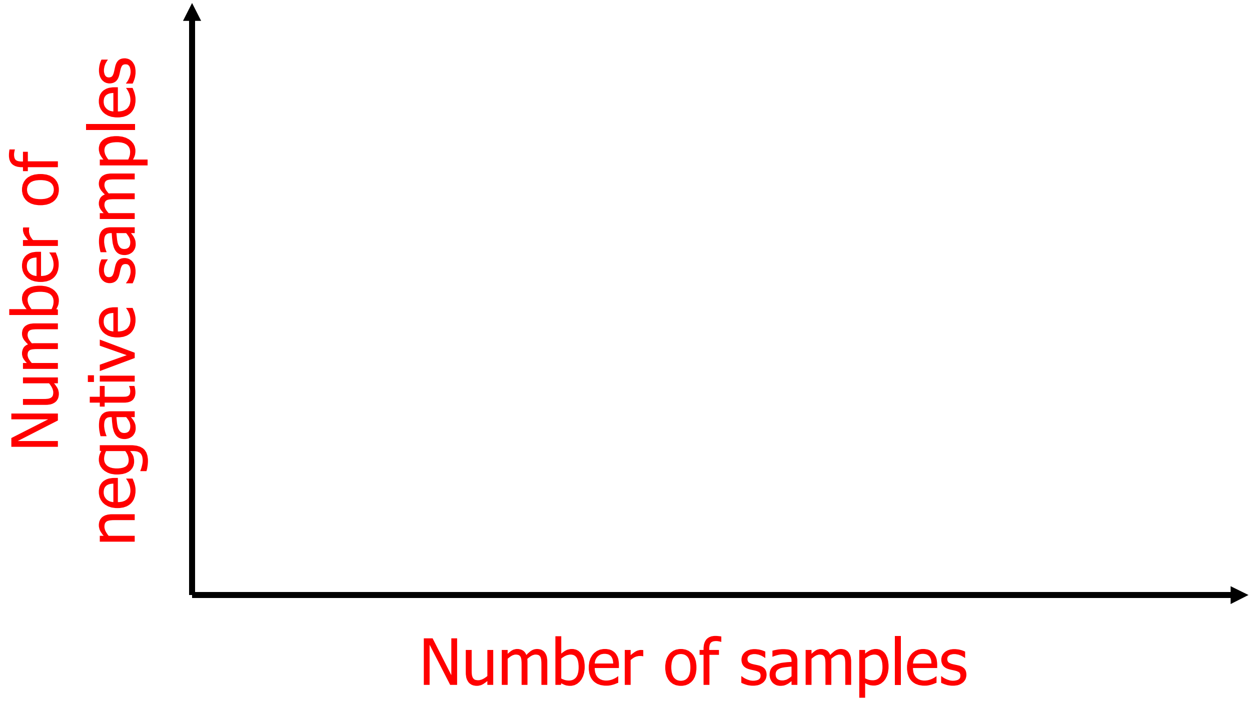

Graphical Representation of Sequential Test

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

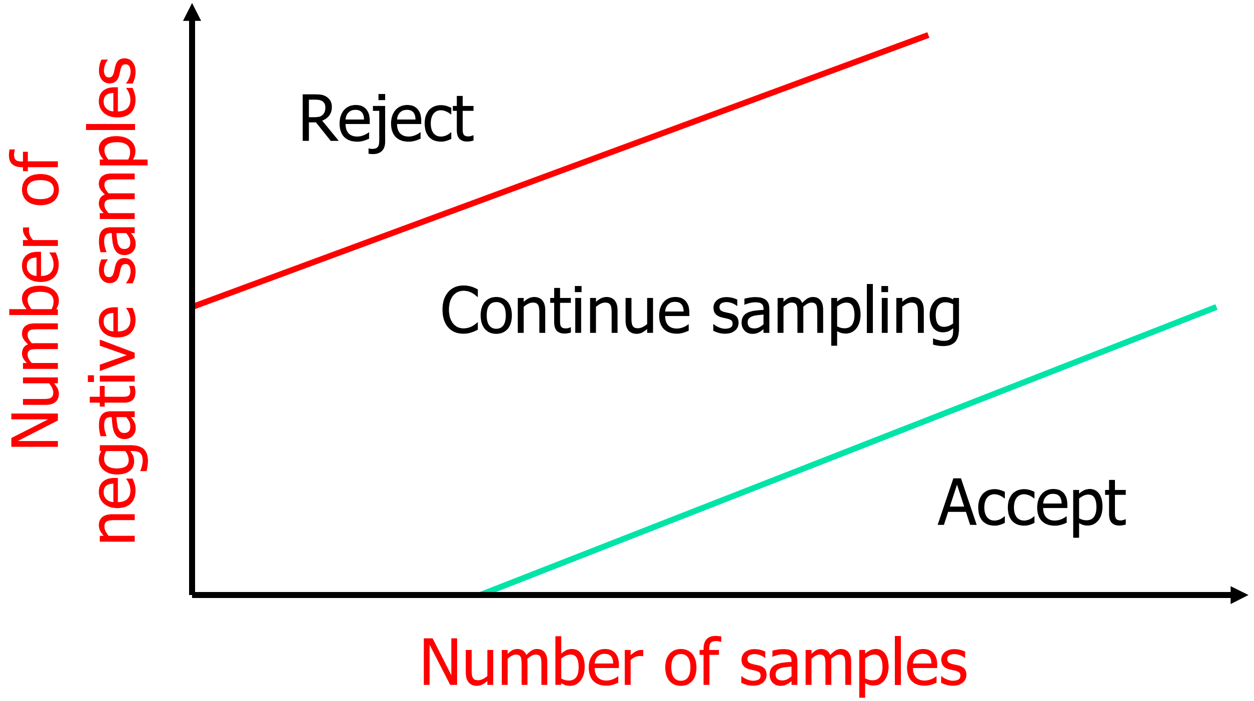

Graphical Representation of Sequential Test

- Accept hypothesis

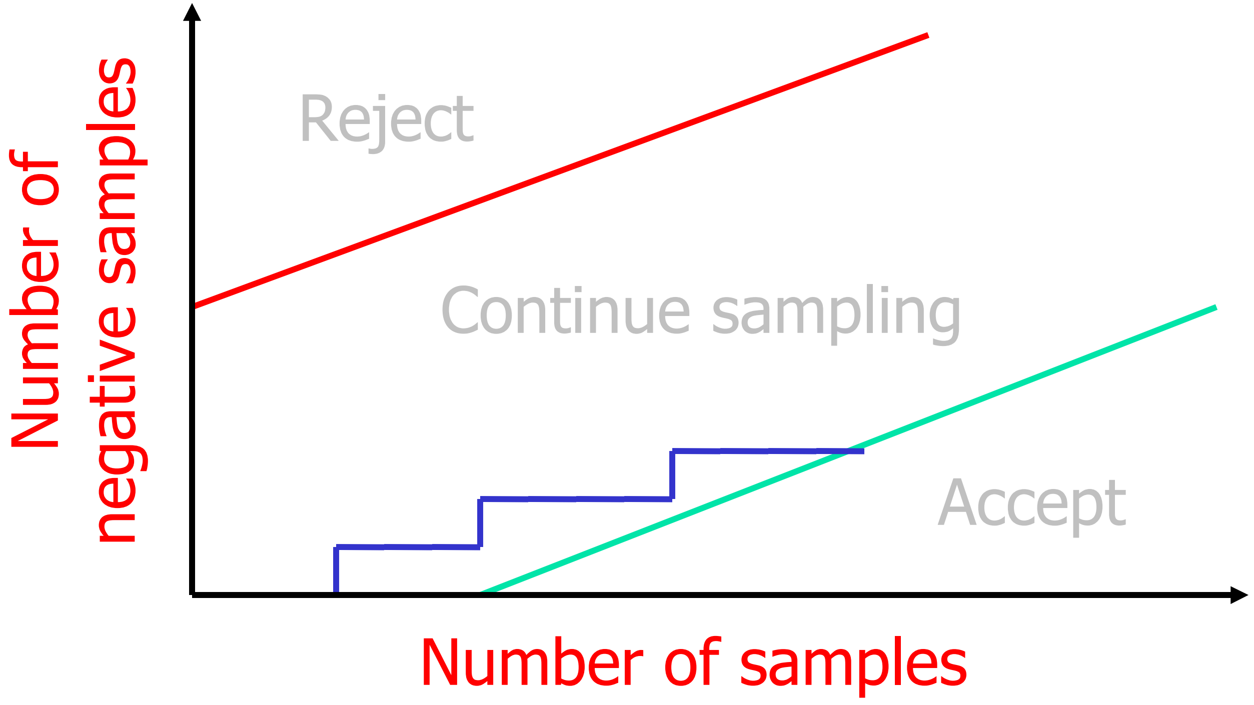

Graphical Representation of Sequential Test

- Reject hypothesis

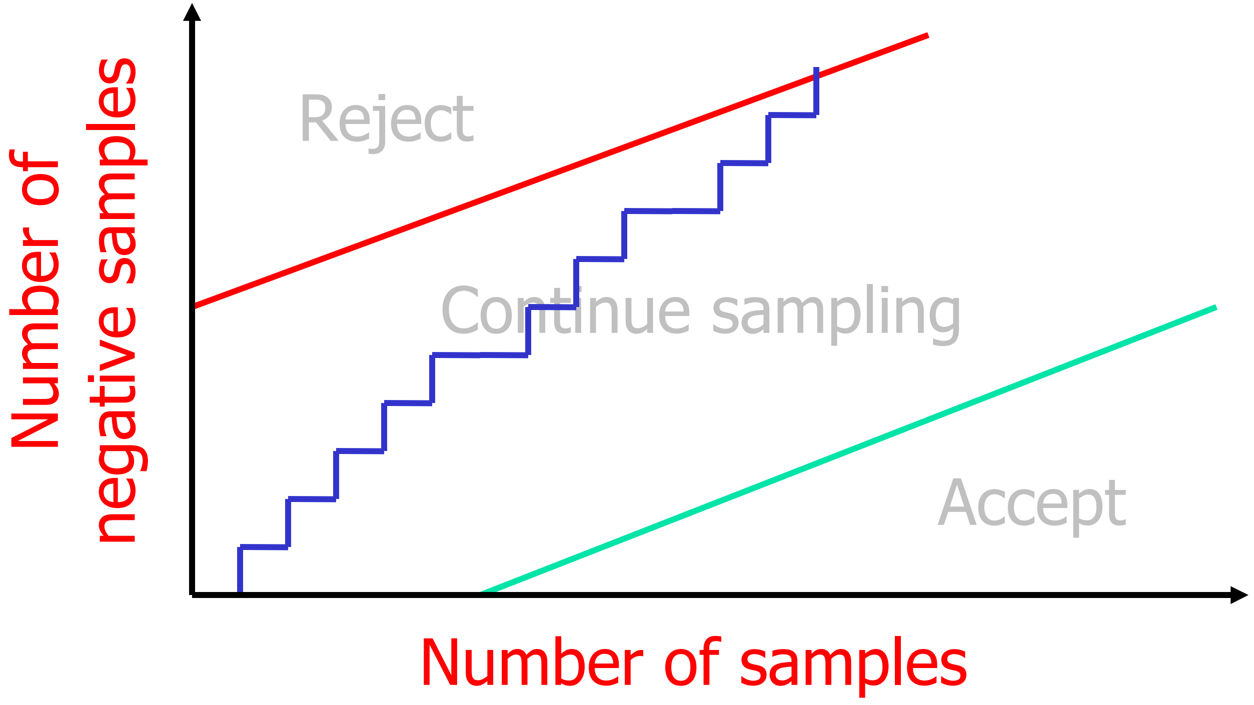

Example

- Verify plan with $\theta=0.05$, $\delta=0.01$, $\alpha=\beta=0.05$, $t_{max}=200$

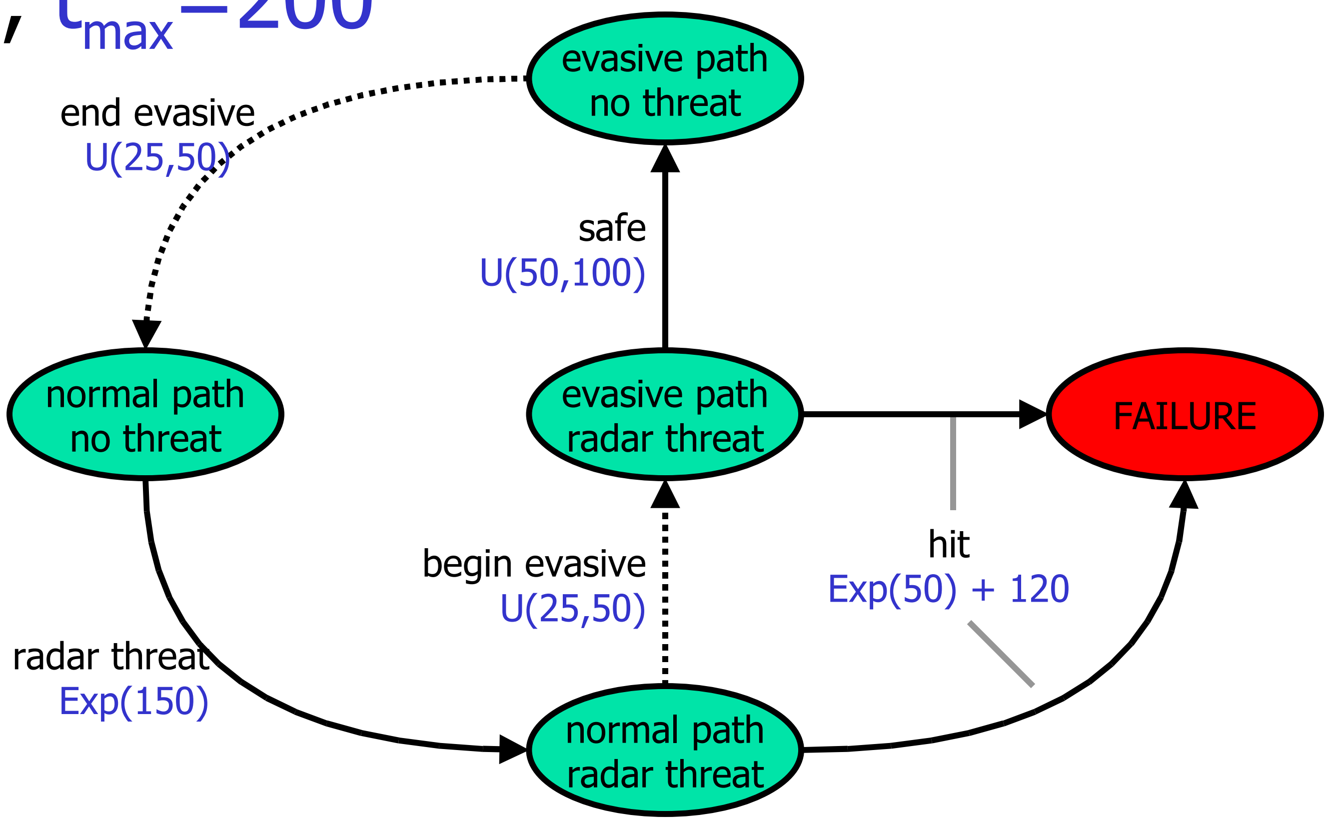

Example

- Verify plan with $\theta=0.05$, $\delta=0.01$, $\alpha=\beta=0.05$, $t_{max}=200$

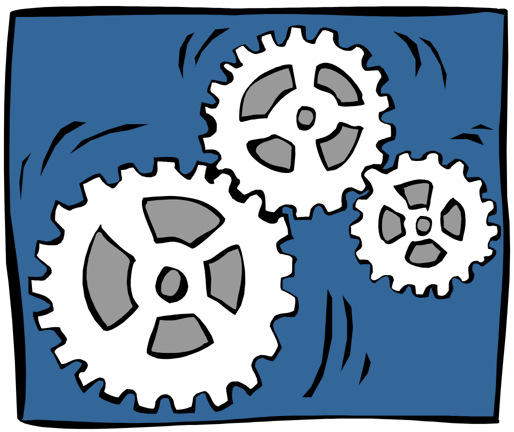

Simulator

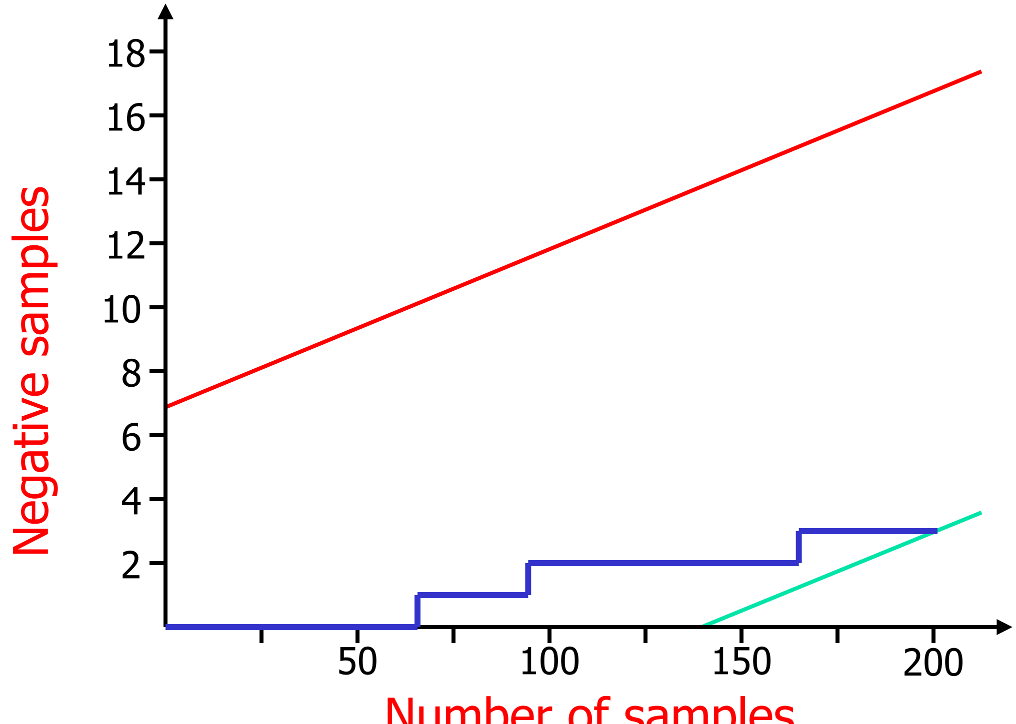

Performance

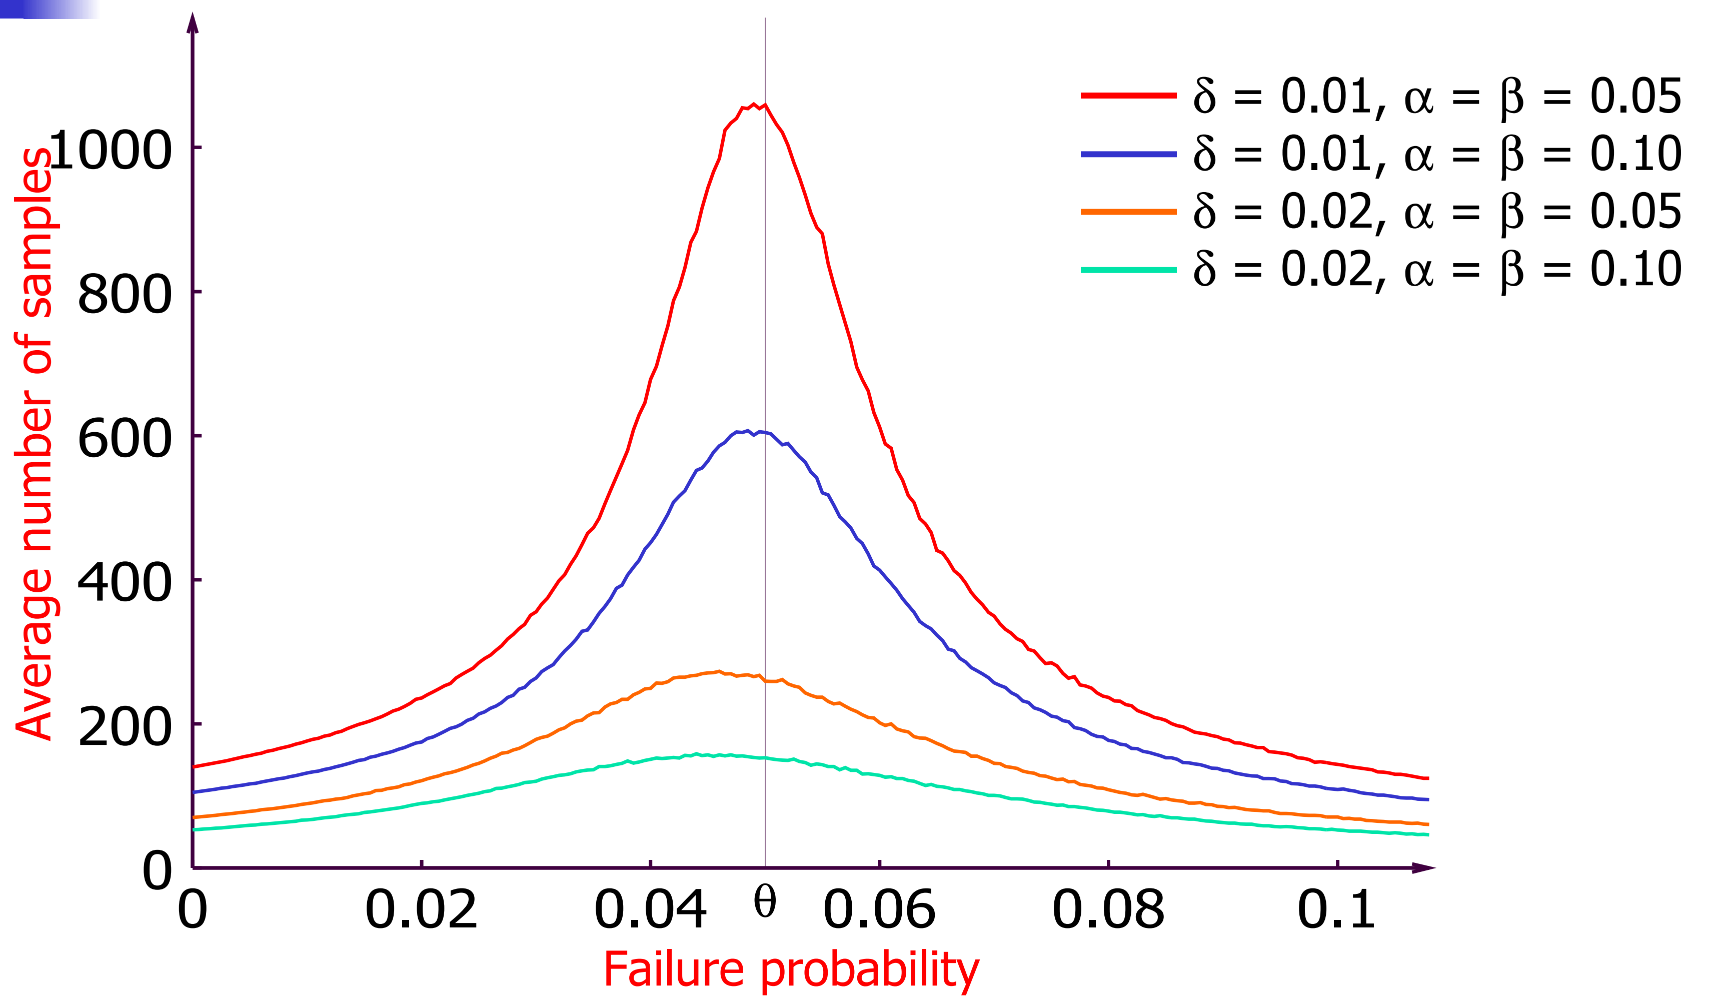

Summary

- Probabilistic extension to CIRCA
  - Allows for plans with non-zero failure probability
- Efficient plan verification algorithm based on acceptance sampling
- Guaranteed error bounds
- Easy to trade efficiency for accuracy

Future Work

- Sensitivity analysis
- Using verification result to guide plan generation
- “Generalized semi-Markov Decision Processes”