An Analysis of Availability Distributions in Condor

Rich Wolski

Daniel Nurmi

John Brevik

University of California, Santa Barbara (UCSB)

Talk given by: Chandra Krintz (UCSB)

Distributed Computing: How long can we run?

- Modern distributed systems are heterogeneous
  - Large scale SMP
  - Clusters
  - Cycle-harvesting systems
- Resource availability duration is important
  - We are given ‘max walltime’ in supercomputers
  - Resource availability times determined by competing users in cycle-harvesting systems!
- Question: can we predict how much time we are given to compute?

Condor: distributed computing through cycle-harvesting

- Idea: use desktop and cluster resources when they are idle
  - Achieve high-throughput computing
- Processes are scheduled on idle resources until…
  - Process completes
  - Resource is reclaimed
- Submitted ten sensor processes to UW Madison pool
  - Sensor records time between ‘began execution’ and ‘resource was revoked’
  - In total gathered process lifetime intervals from 900 hosts over 26 months

Availability Modeling

- Wish to generate accurate descriptive models
  - Simulation
  - Prediction
- Typically assume exponential distribution
  - Shown in literature to be inaccurate for lifetime modeling
- We choose some other distributions from similar family
  - Weibull, Pareto, Log-normal, Hyper Exponential
- Perform MLE fitting technique on gathered process lifetime intervals
  - Tools for automatic model fitting

The shape of condor resource availability

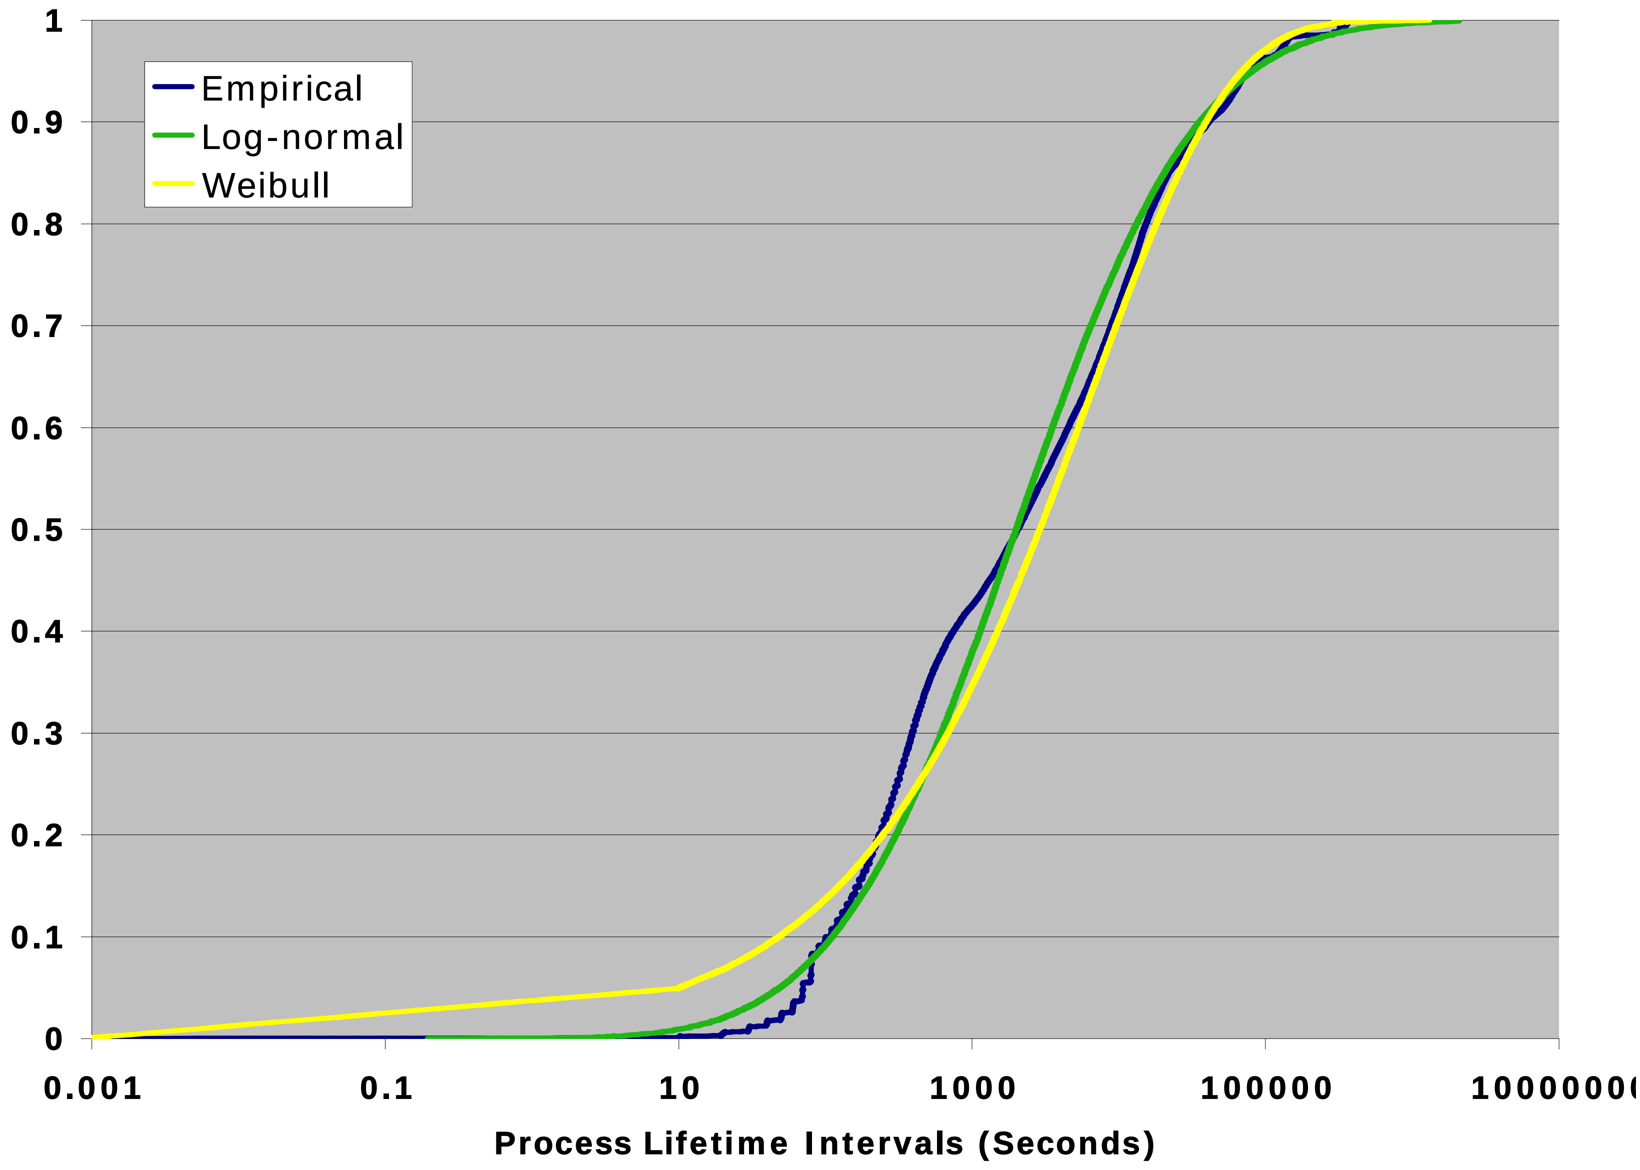

Cluster availability vs. Desktop availability

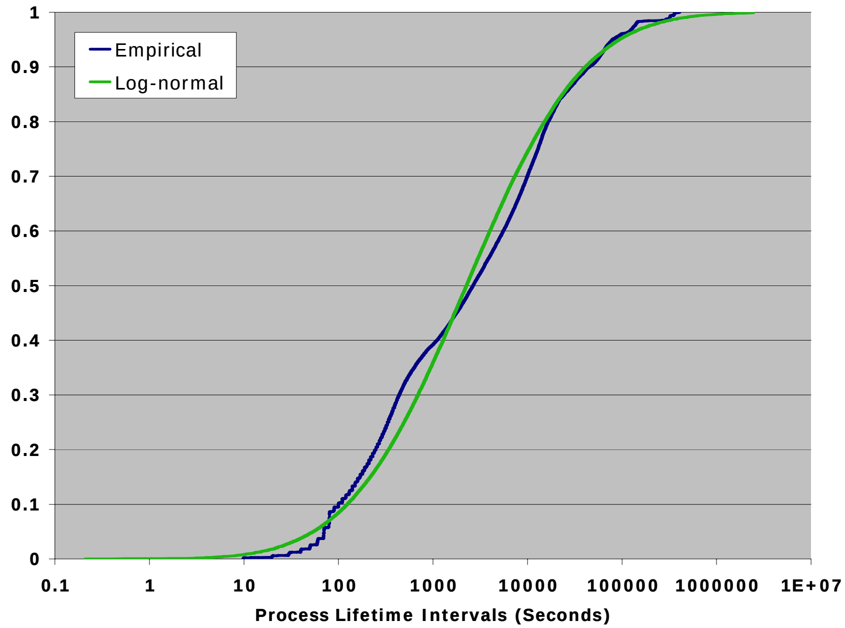

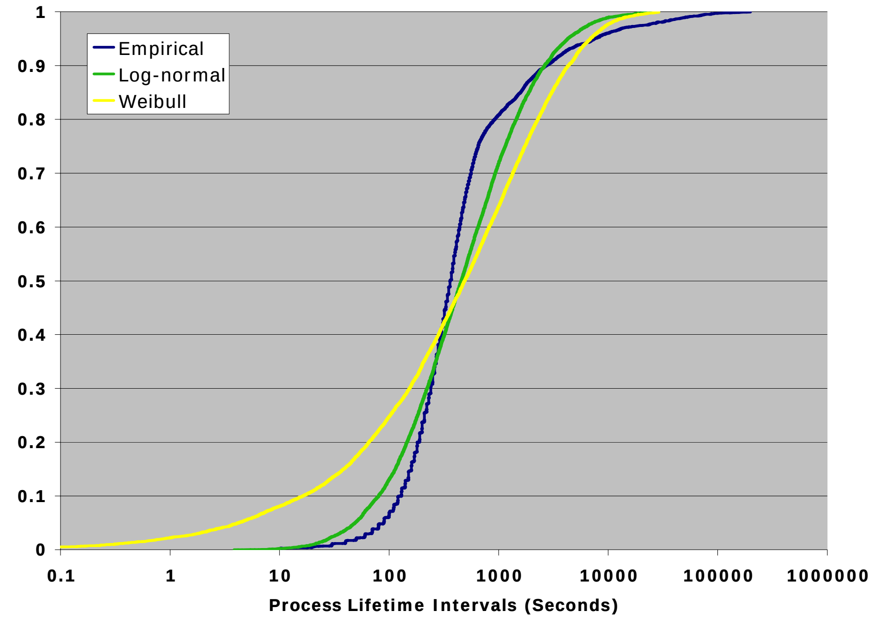

- 'Clustering' results in better fit
- For non-cluster data, neither model fits very well
  - Weibull better for large quantiles
  - Log-normal better for small quantiles

Prediction Results

Data Set	Parametric MLE Method	Bootstrapping	Binomial Method
CSIL	93.3%	96.7%	96.6%
Condor	99.5%	97.1%	95.0%
Long	91.7%	96.0%	97.3%

- Quantile prediction experiment
  - Compare parametric (stat. Model based) method to non-parametric methods
- Non-parametric more accurate, **Binomial Method** (novel) most accurate with smallest data set size

Conclusions

- Log-normal or Weibull distributions are useful for modeling Condor process lifetimes
- Clustering of data into ‘cluster’ and ‘desktop’ sets improves model accuracy
  - Desktop data not very well modeled using parametric techniques
  - Non-parametric better for prediction
- We can accurately predict process lifetime quantiles
  - Optimal checkpoint interval selection
  - Redundancy hints
  - Condor scheduler hints
- Thank You! Questions?