SDC

STORAGE DEVELOPER CONFERENCE

SNIA ■ SANTA CLARA, 2016

Modeling Disk Failures with Absorbing Markov Chains

Rachel Traylor, Ph.D.
EMC

What is a Markov Chain?

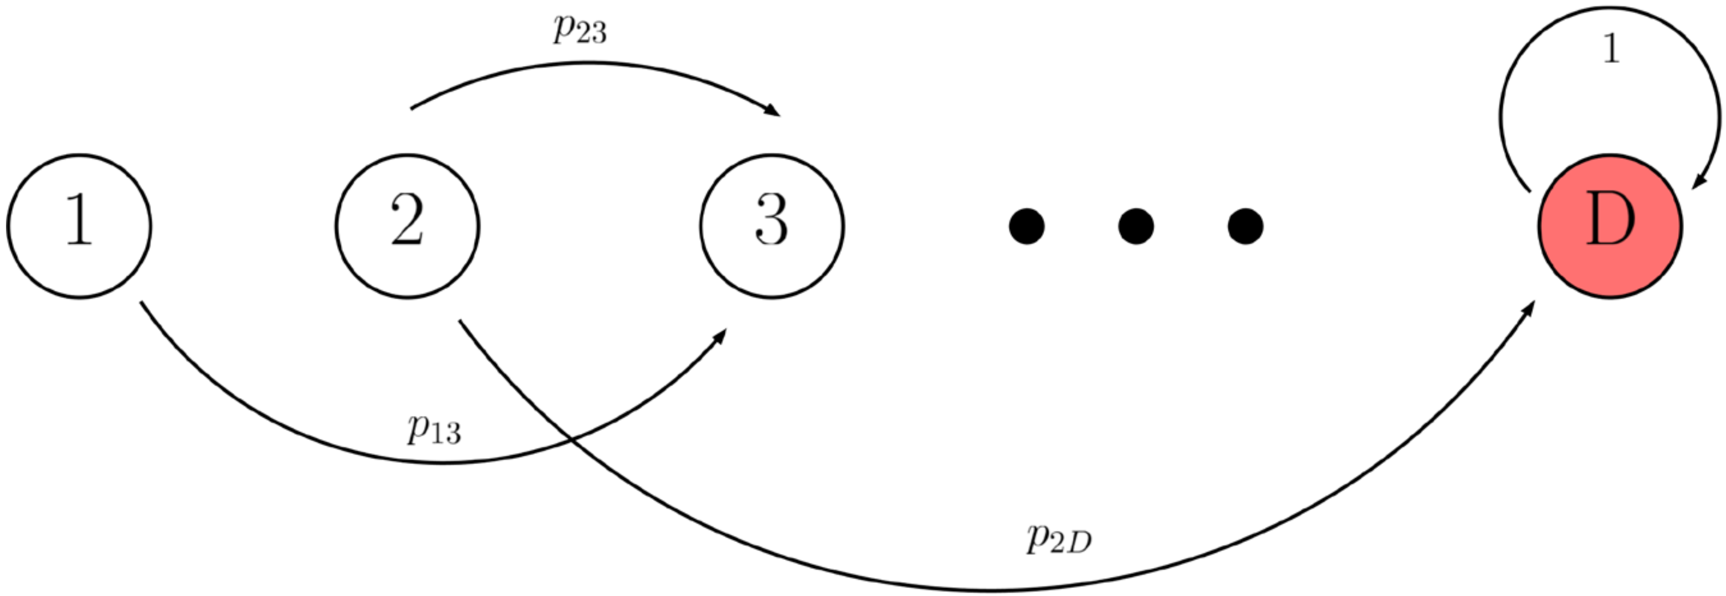

Terminology

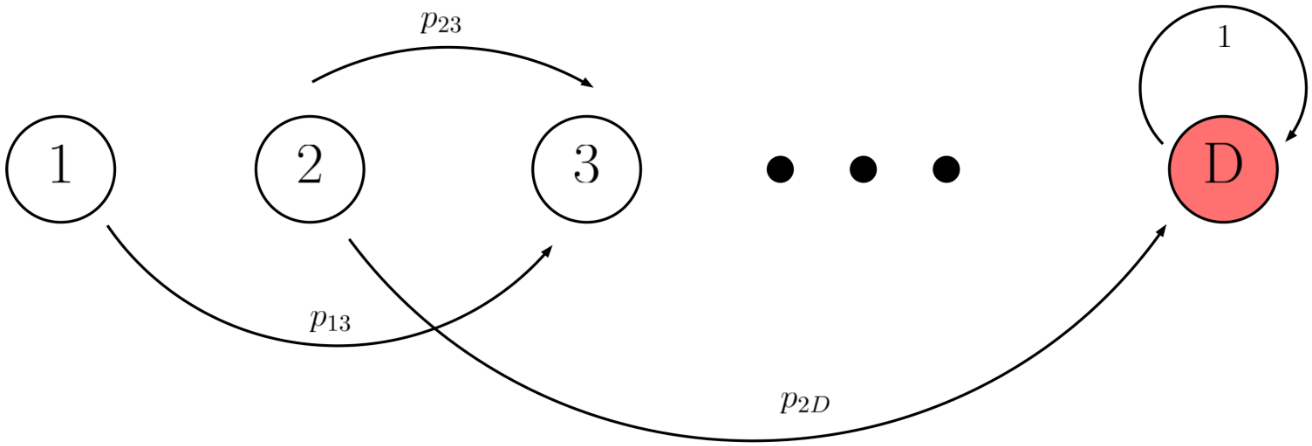

State Space: the set of possible states the process can reach

Transition Probability (p_{ij}): the probability the process transitions from state *i* to state *j* in one step

Absorbing State: a state that, once entered, can never be left.

Transition Matrix

$$\begin{bmatrix} p_{11} & p_{12} & p_{13} & \cdots & p_{1D} \\ 0 & p_{22} & p_{23} & \cdots & p_{2D} \\ 0 & 0 & p_{33} & \cdots & p_{3D} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 1 \end{bmatrix}$$

- For failure modeling, error counts always increase, so for $j < i$, $p_{ij} = 0$
- Note that the diagram below does not depict all possible transitions given in the matrix

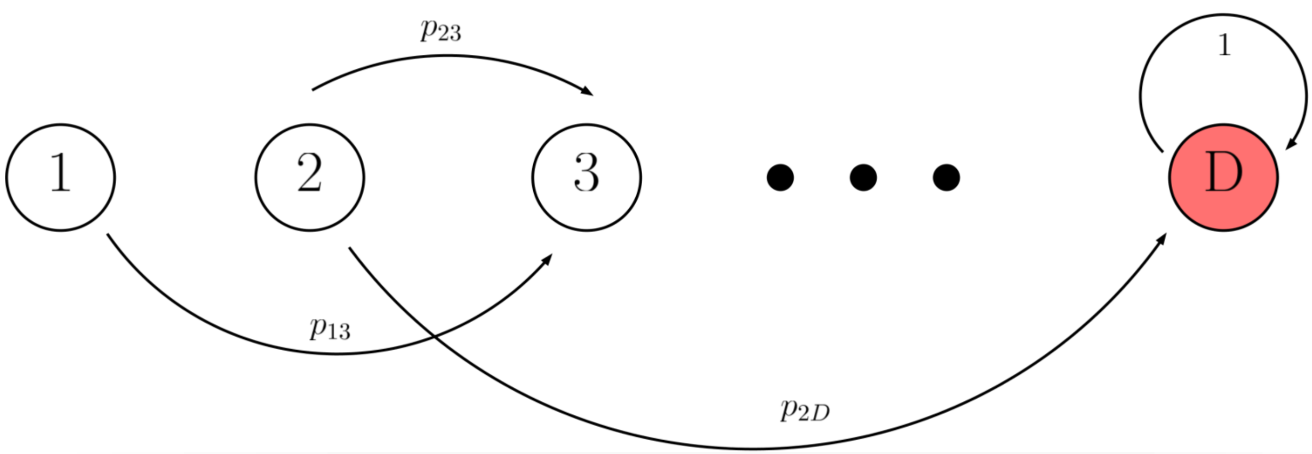

Modeling Medium Error Evolution

- **Medium error:** physical issue on the drive was encountered during access attempt
  - most commonly head error
- Medium Errors have been shown to be a good predictor of HDD failure.
- They are fairly rare in the field—approximately 1% of disks experience medium errors
- Conditioning on the initial NMEC, we see a decrease in survival rates fairly quickly

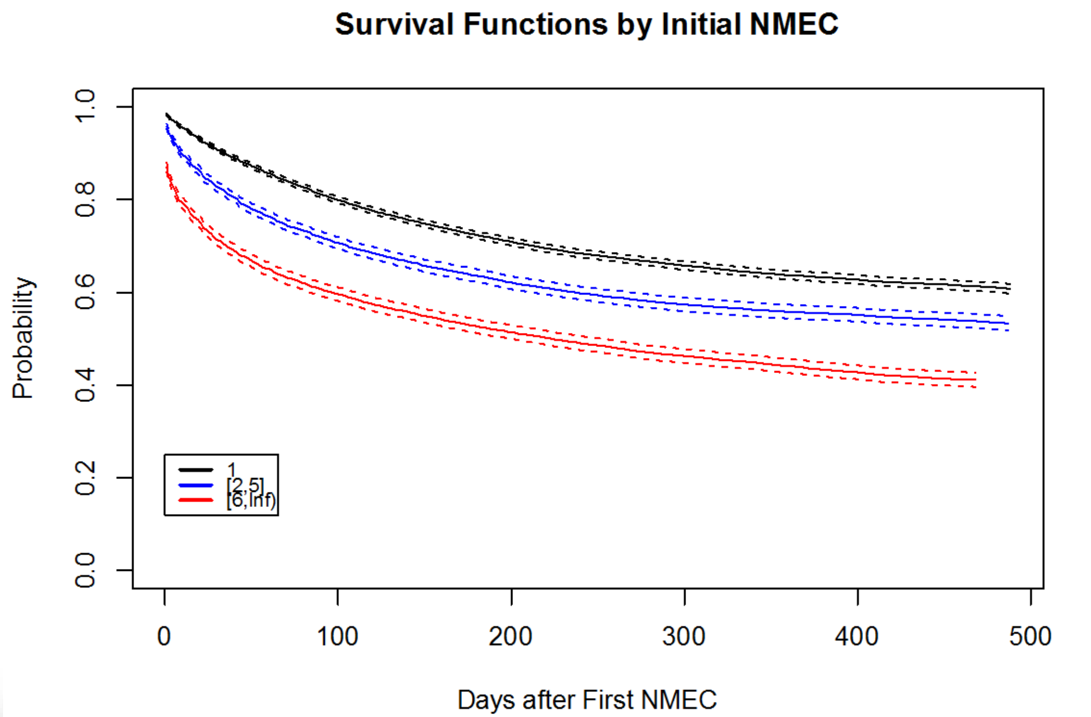

Methodology-Conditional Markov Chain

- Create transition matrix conditioning on the existence of medium errors
- Let τ be a transition. Then

$$p_{ij} = \frac{|\{\tau : ME = j\}|}{|\{\tau : ME\_prev = i\}|}$$

Binning

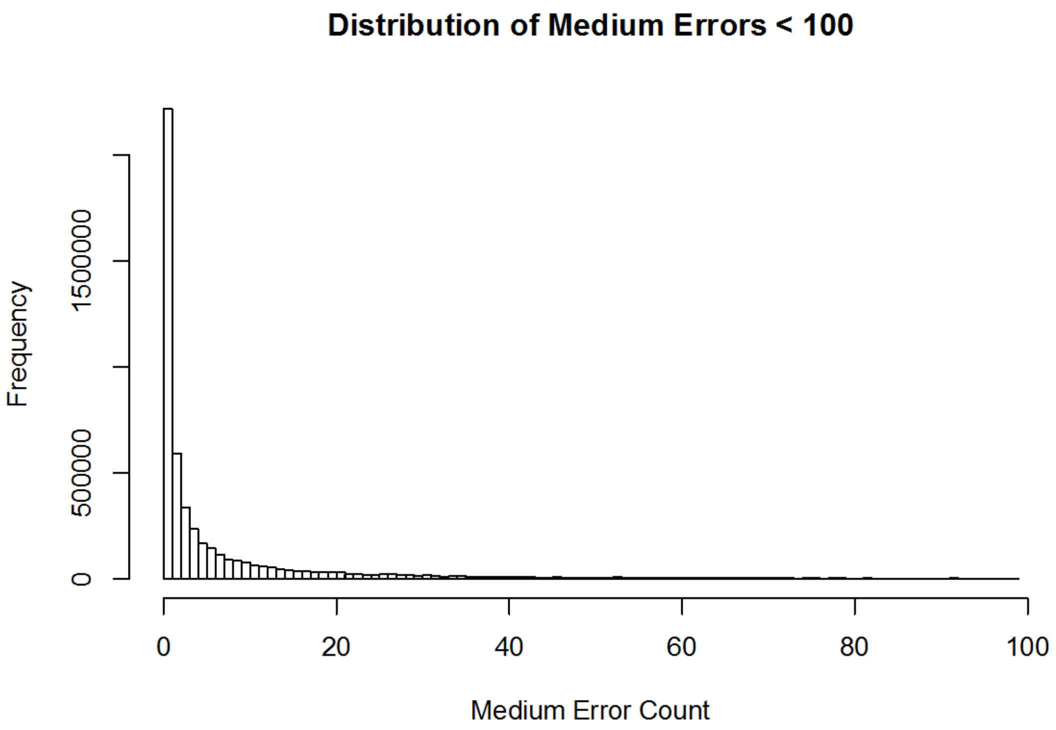

- A Markov chain remains such even when condensing states
- Empirical bins were used to condense the state space to produce a more manageable transition matrix

Min	Q1	Med	Q3	Max
1	1	4	20	144250

Beyond the Transition Probabilities

- The transition matrix simply gives the probability of moving from State *i* to State *j*
- The transition matrix be used to answer many deeper questions about the evolution of the state space

N-step transition probabilities

- Of greater interest is the probability of transitioning from to State *j* in *N* steps, given that the process began in State *i*
- The Markov process is memoryless
  - The information contained in the process's entire past is the same as the information contained in the previous step

$$P(X_n = x_n \mid X_{n-1} = x_{n-1}, \ldots, X_0 = x_0) = P(X_n = x_n \mid X_{n-1} = x_{n-1})$$

- So, we may assume the day of observation is the "start" of a new process.

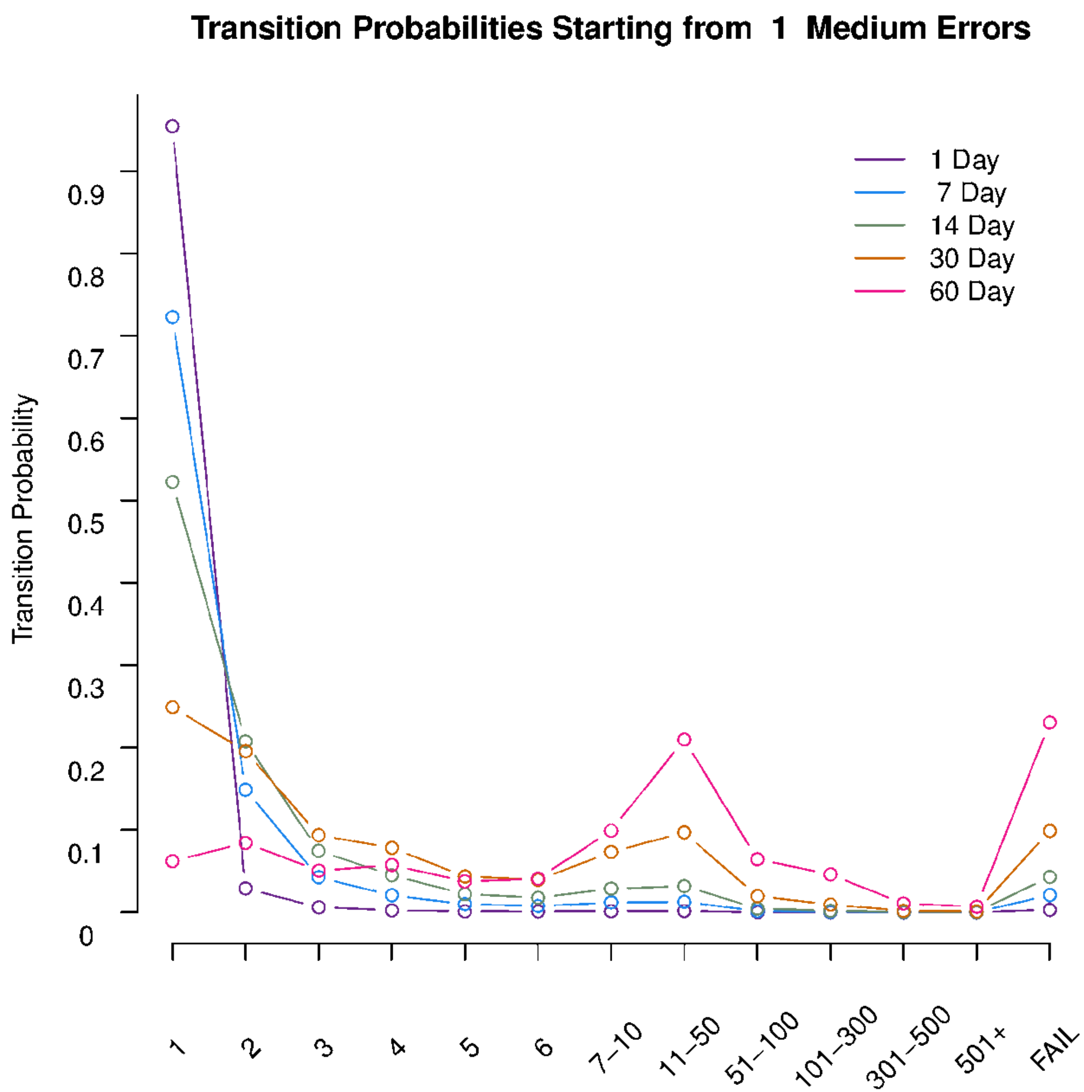

The N-step transition probabilities are given by the entries of the matrix

$$P^N$$

Other uses for the transition matrix

What is the probability that a disk with 3 medium errors today will have more than 10 medium errors in 7 days?

- The timestep is 7 days – $\mathbf{P}^7$
- The current state is 3 medium errors

$$P(ME_7 > 10 \mid ME_c = 3) = p_{3,8}^{(7)} + p_{3,9}^{(7)} + p_{3,10}^{(7)} + p_{3,11}^{(7)} + p_{3,12}^{(7)} + p_{3,D}^{(7)} = 0.074$$

Expected Time to Absorption

$$E[T] = (I - Q)^{-1}\mathbf{1}$$

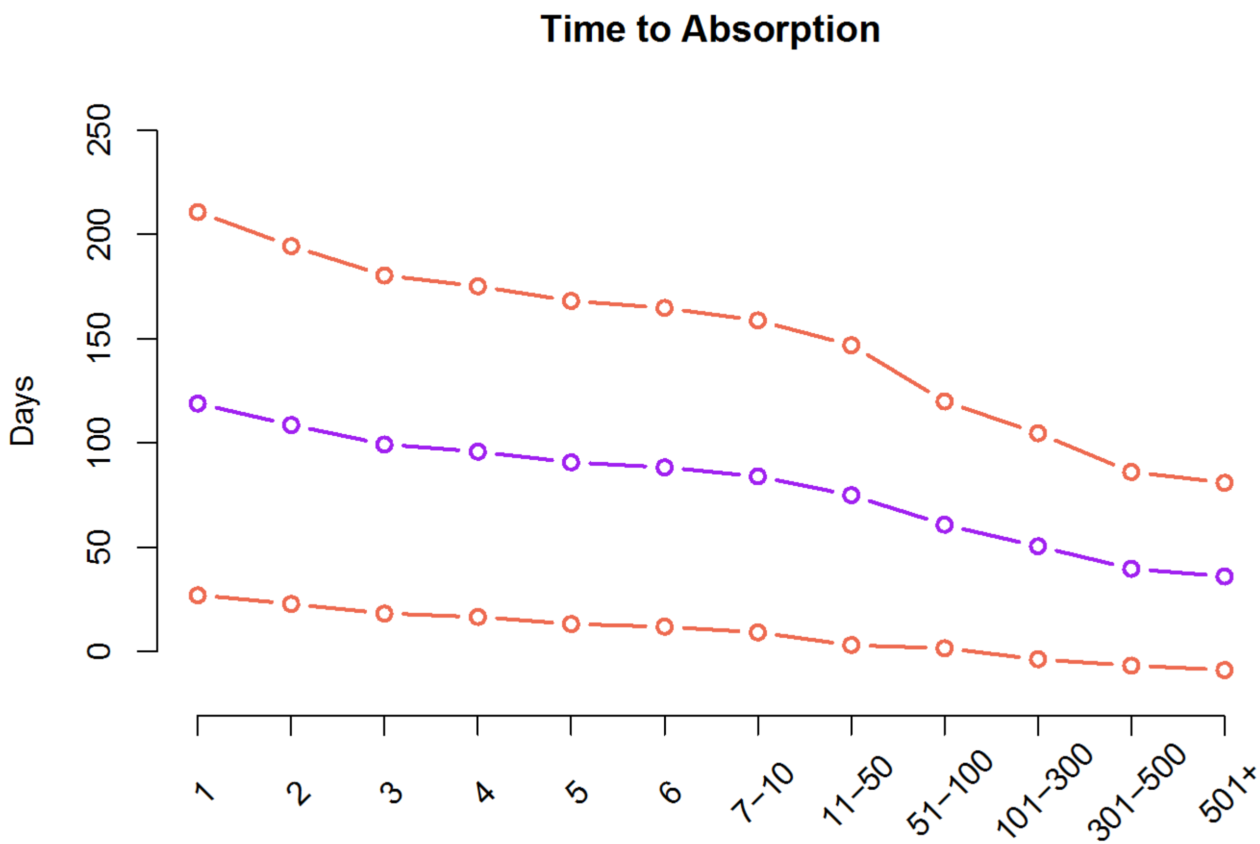

Applications: RAID Group Risk algorithm

- The results of Markov chain modeling were used to develop an algorithm to predict the risk of data loss in a RAID group

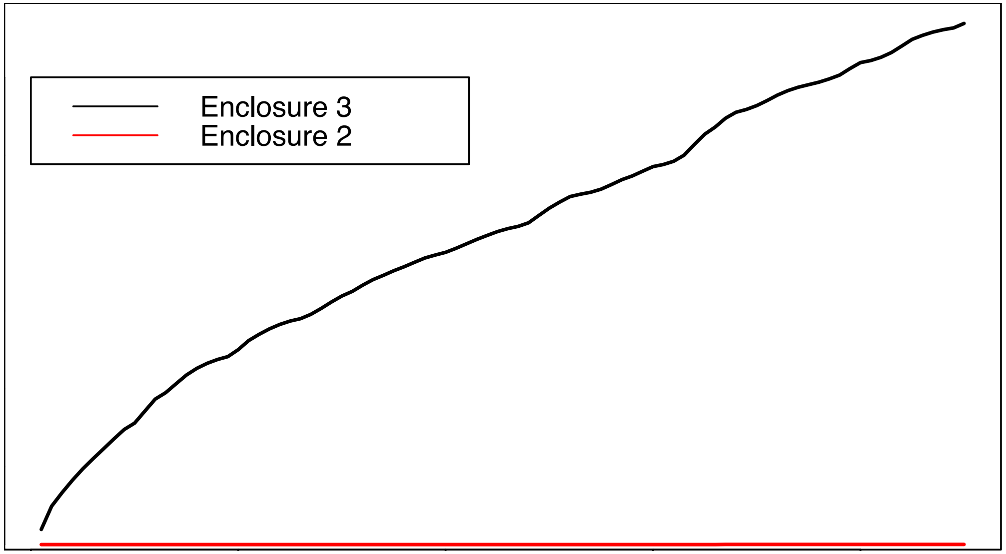

Other Applications

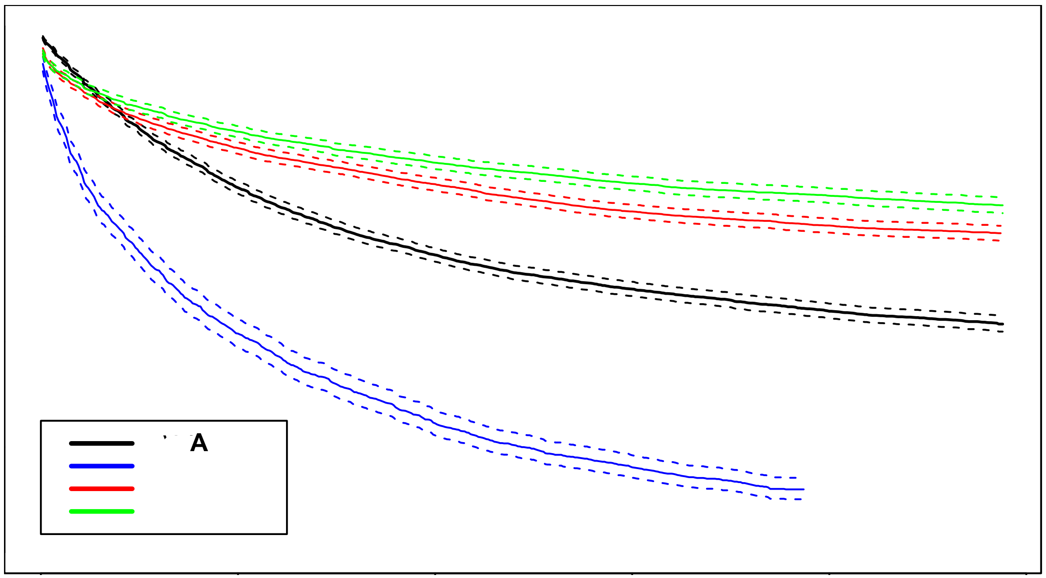

- State Space used for profiling
- Multidimensional/continuous Markov chain

Conclusions

- Probabilistic foundation
- Computationally cheap, requires little updating
- Requires less data to maintain/evaluate
- Elegant and interpretable
- Compact form still allows for a wide variety of questions to be answered

SDC 16

Questions?

Rachel Traylor, Ph.D.
Senior Research Scientist
Office of the CTO
Core Technology Division
EMC²