# Checkpointing for the RESTART Problem in Markov Networks

Lester Lipsky Derek Doran Swapna Gokhale
(With lots of help from Steve Thompson)

Department of Computer Science & Engineering
University of Connecticut

New Frontiers in Applied Probability at Sandbjerg Estate,
Sønderborg, 1-5 August 2011
Conference in Honour of Søren Asmussen
on the occasion of his 65th Birthday

# Overview 1

# Matrix Exponential (ME) Distributions - I 2

Subsystem with $M$ nodes (phases)

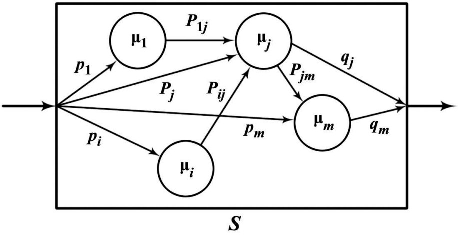

# Matrix Exponential (ME) Distributions - II 3

- Let $\mathbf{P}$ be a transition $M$-Matrix such that $\mathbf{I} - \mathbf{P}$ has an inverse;
- Let $\varepsilon'$ be an $M$ dimensional column-vector of all 1's;
- Let $\mathbf{p}$ be an $M$ row-vector where $(\mathbf{p})_i$ is the probability that the process will start at node $i$, and $\mathbf{p}\varepsilon' = 1$;
- Let each of the $M$ nodes have exponential service time distributions, with rate $\mu_i = (\mathbf{M})_{ii} > 0$ ($\mathbf{M}$ is a diagonal matrix);
- Let $T$ be the time from entry to departure;

# Matrix Exponential (ME) Distributions - III 4

- Define

$$\mathbf{B} = \mathbf{M}(\mathbf{I} - \mathbf{P}) \quad \text{and} \quad \mathbf{V} = \mathbf{B}^{-1};$$

- Then the Probability Distribution (PDF), Reliability, and probability density (pdf) functions for $T$ are

$$F(t) := \mathbb{P}\mathbf{r}[T \leq t] = 1 - \mathbf{p}\exp(-t\mathbf{B})\boldsymbol{\varepsilon}'. \quad \bar{F}(t) = 1 - F(t),$$

$$\text{and} \quad f(t) = \frac{dF}{dt} = \mathbf{p}\exp(-t\mathbf{B})\mathbf{B}\boldsymbol{\varepsilon}'.$$

- Also

$$\mathbb{E}[T^{\ell}] = \ell!\,\mathbf{p}\mathbf{V}^{\ell}\,\boldsymbol{\varepsilon}'.$$

# ME Representation of the Uniform Distribution 5

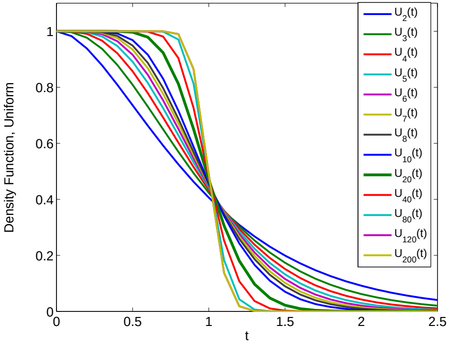

# Truncated Power-tail (TPT) Distributions 6

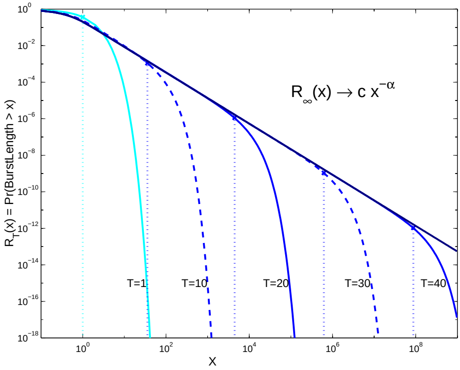

# Recovery Scenarios 7

There have been three general scenarios about recovering after a system crashes during execution.

- preemptive Resume (prs) - RESUME
- preemptive repeat different (prd) - REPLACE
- preemptive repeat identical (pri) - RESTART

RESUME and REPLACE can be analyzed by Markov models. RESTART, however, is more difficult to treat.

# The Performance of Systems Under RESTART - I 8

- Let $T$ be the time for a job to complete without failures, .
- Let $F(t)$, $f(t)$ and $\bar{F}(t) = 1 - F(t)$ be the *PDF, pdf*, and *reliability functions* for T.
- Assume that the failure distribution is exponential with failure rate $\beta$. Then for $T = t$, let $X(t, \beta)$ be the completion time with failures, under RESTART, with PDF $H(x|t)$. Then its Laplace transform was shown to be

$$H^*(s|t) = \frac{(s+\beta)e^{-(s+\beta)t}}{s+\beta e^{-(s+\beta)t}}.$$

- Since this is the *moment generating function* of $H(x|t)$, we have in general

$$\mathbb{E}[X(t, \beta)^\ell] = (-1)^\ell \left[\frac{d^\ell H^*(s|t)}{ds^\ell}\right]_{s=0}.$$

UCONN

## The Performance of Systems Under RESTART - II 9

- Since $T = t$ throughout a RESTART process, it follows that

$$\mathbb{E}[X(\beta)^\ell] = \int_0^\infty \mathbb{E}[X(t,\, \beta)^\ell]\, f(t)\, dt.$$

- In particular, for $\ell = 1$ we have

$$\mathbb{E}[X(t,\, \beta)] = \frac{e^{\beta t} - 1}{\beta} \quad \text{and}$$

$$\mathbb{E}[X(\beta)] = \int_0^\infty \frac{e^{\beta t} - 1}{\beta}\, f(t)\, dt$$

.

# The Performance of Systems Under RESTART - III 10

Define:

$$\lambda_s := \sup \left\{ \lambda \mid \int_0^\infty \exp(\lambda t)\, f(t)\, dt < \infty \right\}.$$

Also define

$$\alpha := \sup \left\{ \ell \mid \int_0^\infty x^\ell\, h(x)\, dx < \infty \right\}$$

where $h(x)$ is the pdf for $X(\beta)$ (total completion time under *RESTART*). Then $X(\beta)$ is *power-tailed* (PT) with index $\alpha$ if $0 < \alpha < \infty$.

# The Performance of Systems Under RESTART - IV 11

From these definitions we have the following.

- if $T$ has infinite support, $X(\beta)$ is *sub-exponential*.
- $f(t)$ has an *exponential tail* with parameter $\lambda_s$ if $0 < \lambda_s < \infty$. If $\lambda_s = 0$ then $f(t)$ is *sub-exponential*.
- if $T$ has an exponential tail with parameter $\lambda_s$, then $X(\beta)$ will be PT with index

$$\alpha = \lambda_s/\beta.$$

Thus as $\beta$ becomes bigger, $\alpha$ becomes smaller, and the system behavior becomes more unstable.

UCONN

# Markov Models of Software (MMS model) 12

- Software systems (among others) are highly modular, where the system control is passed among independent components.
- The passing of control between the $M$ components (nodes) maps to an $M$ dimensional Markov matrix, $\mathbf{P}$.
- Assume that:
  - the service time at each node is exponentially distributed with rate $\mu_i := [\mathbf{M}]_{ii} > 0$;
  - there is a path to exit the system from each node;

Then, as previously described, the distribution for the total execution time $T$ is ME distributed (actually, *PHase*).

# The MMS Model Under RESTART 13

For ME distributions, $\lambda_s := \mathrm{Min}[|\lambda_i|]$, where $\{\lambda_i \,|\, 1 \leq i \leq M\}$ is the set of eigenvalues of **B** whose eigenvectors are not orthogonal to **p** or $\boldsymbol{\varepsilon}'$.

- If the MMS model is subject to exponential failures, and must RESTART, $X(\beta)$ will be PT distributed with $\alpha = \lambda_s/\beta$
- The first two moments of $X(\beta)$ are given by:

$$\mathbb{E}[X(\beta)] = \mathbf{p}\left[\mathbf{V}(\mathbf{I} - \beta\mathbf{V})^{-1}\right]\boldsymbol{\varepsilon}' \quad (\beta < \lambda_s)$$

$$\mathbb{E}[X(\beta)^2] = 2\mathbf{p}\left[\mathbf{V}^2(\mathbf{I} - 2\beta\mathbf{V})^{-2}(\mathbf{I} - \beta\mathbf{V})^{-1}\right]\boldsymbol{\varepsilon}' \quad (\beta < \lambda_s/2)$$

even though $X(\beta > 0)$ is *not* ME.

# Markov Chains with Two Absorbing States - I 14

- Consider an $(M+2)$-dimensional Markov matrix $\bar{\mathbf{P}}$ with two absorbing states, $a$ and $b$. That is,

$$\bar{\mathbf{P}}\bar{\varepsilon'} = \bar{\varepsilon'} \quad \text{and} \quad (\bar{\mathbf{P}})_{aa} = (\bar{\mathbf{P}})_{bb} = 1$$

- Deleting the rows and columns of $a$ and $b$ gives $\mathbf{P}$.
- Then,

$$[\mathbf{Z}]_{ij} := [(\mathbf{I} - \mathbf{P})^{-1}]_{ij}$$

is the expected number of visits to $j$ before absorption, given that the chain started at $i$.

UCONN

# Markov Chains with Two Absorbing States - II 15

- Now define the $M$-dimensional column vectors

$$(\mathbf{q}'_{\mathbf{a}})_i := \bar{P}_{ia} \quad \text{and} \quad (\mathbf{q}'_{\mathbf{b}})_i := \bar{P}_{ib}, \quad \text{where} \quad i \neq a, \; b.$$

  These are the probability vectors of being absorbed by $a$ and $b$, respectively.

- It follows that the $i^{th}$ components of

$$\boldsymbol{\varepsilon}'_{\mathbf{a}} := \mathbf{Z}\,\mathbf{q}'_{\mathbf{a}} \quad \text{and} \quad \boldsymbol{\varepsilon}'_{\mathbf{b}} := \mathbf{Z}\,\mathbf{q}'_{\mathbf{b}}$$

  are the probabilities that the process will end at $a$ or $b$, respectively, given that the process started at $i$.
  Note that $\boldsymbol{\varepsilon}'_{\mathbf{a}} + \boldsymbol{\varepsilon}'_{\mathbf{b}} = \boldsymbol{\varepsilon}'$.

# Markov Chains with Two Absorbing States - III 16

- Let $\mathbf{p}_\text{o}$ be the entrance vector. Then

$$p_a = \mathbf{p}_\text{o}\boldsymbol{\varepsilon}'_\mathbf{a} \quad \text{and} \quad p_b = \mathbf{p}_\text{o}\boldsymbol{\varepsilon}'_\mathbf{b}, \quad \text{where} \quad p_a + p_b = 1$$

  are the probabilities that the process will be absorbed by $a$ or $b$.
- It is well known that $[\mathbf{p}_\text{o}\exp(-\mathbf{B}t)]_i$ is the probability that absorption has not occured by time $t$, and the system is in state $i$. This all leads to the following:

UCONN

# Markov Chains with Two Absorbing States - IV



- **Theorem:** Let $\mathbf{q}'_\mathbf{u}$, $\boldsymbol{\varepsilon}'_\mathbf{u}$, $\mathbf{p_o}$, $\mathbf{B}$ and $\mathbf{V}$, where $u \in \{a, b\}$, be defined as above. Then $T_u$ has distribution

$$\bar{F}_u(t) := \mathbb{P}\mathbf{r}[T_u > t] = \mathbf{p}_o \exp(-\mathbf{B}t)\boldsymbol{\varepsilon}'_\mathbf{u}/p_u, \quad u = a,\ b.$$

The moments of these distributions come from above:

$$\mathbb{E}[T_u^\ell] = \ell!\, \mathbf{p_o}\, [\mathbf{V}^\ell] \boldsymbol{\varepsilon}'_\mathbf{u}/p_u$$

We then say that $\bar{F}_u(t)$ is generated by the triplet $\langle\, \mathbf{p}_o,\ \mathbf{B}, \boldsymbol{\varepsilon}'_\mathbf{u} \,\rangle$.

- (Note that $\mathbb{E}[T^\ell] = p_a\mathbb{E}[T_a^\ell] + p_b\mathbb{E}[T_b^\ell] = \ell!\, \mathbf{p_o}\, [\mathbf{V}^\ell]\boldsymbol{\varepsilon}'$.)

UCONN

# Applying Checkpointing to the MMS 18

- Checkpointing can easily be applied to the model to combat the PT service times under RESTART.
- After execution of a selected node $m$, a system checkpoint operation can be applied, saving the system state.
- Ideally, the designer will apply checkpointing for each state, and select the one that yields the best performance.
- To analyze this system we need the conditional distributions for the time to absorption at each of two absorbing states.

UCONN

# Checkpointing in Markov Systems (MMSC model) 19

- For the original MMSC model, select node $m$ as the one that is followed by a system checkpoint. Then,

$$q_m = [\mathbf{q}']_m := [(\mathbf{I} - \mathbf{P})\boldsymbol{\varepsilon}']_m$$

  is the probability that execution will end after finishing at $m$.
- Add one row and one column to $\mathbf{P}$ at index $M + 1$, representing the system checkpoint state, to produce the matrix $\mathbf{P_c}$.

# The MMSC Model - I



- $\mathbf{P_c}$ has the following properties: for $i \neq m,\ M+1$ and $j \neq M+1$,

$$[\mathbf{P_c}]_{ij} = \mathbf{P}_{ij}, \quad [\mathbf{P_c}]_{i,M+1} = 0,$$

$$[\mathbf{P_c}]_{mi} = 0, \quad [\mathbf{P_c}]_{m,M+1} = 1 - q_m,$$

$$[\mathbf{P_c}]_{M+1,k} = 0, \ \forall \ k.$$

- This defines a Markov chain with two absorbing states, $e$ (for end) and $c$ (for checkpoint).
- To use the established theorem we need $\mathbf{q_e'}$ and $\mathbf{q_c'}$.
- $\mathbf{q_e'}$ is the same exit vector as that for the original model, with additional component $[\mathbf{q_e'}]_{M+1} = 0$, so

$$[\mathbf{q_e'}]_i = [(\mathbf{I} - \mathbf{P})\boldsymbol{\varepsilon}']_i, \quad \text{but} \quad [\mathbf{q_e'}]_{(M+1)} = 0$$

# The MMSC Model - II



- $\mathbf{q}'_\mathbf{c}$ is given as: $[\mathbf{q}'_\mathbf{c}]_i = 0$, for $i \leq M$ and $[\mathbf{q}'_\mathbf{c}]_{M+1} = 1$.
- We define the $(M+1)$-matrix $\mathbf{Z_c} = (\mathbf{I} - \mathbf{P_c})^{-1}$ to get

$$\boldsymbol{\epsilon}'_{\boldsymbol{e}} = \mathbf{Z_c}\,\mathbf{q}'_\mathbf{e} \quad \text{and} \quad \boldsymbol{\epsilon}'_{\boldsymbol{c}} = \mathbf{Z_c}\,\mathbf{q}'_\mathbf{c}$$

- The probability of finishing the process without checkpointing is:

$$p_{oe} := \mathbf{p_o}\boldsymbol{\epsilon}'_{\boldsymbol{e}}$$

- We can also get the probability of reaching the checkpoint before finishing:

$$p_{oc} := \mathbf{p_o}\boldsymbol{\epsilon}'_{\boldsymbol{c}}$$

# The MMSC Model - III



- Now we apply the theorem to get the conditional distributions for the time to finish given no checkpoint ($T_{oe}$) and the time to reach and execute the checkpoint ($T_{oc}$).
- Define the diagonal matrix

$$[\mathbf{M_c}]_{ii} = [\mathbf{M}]_{ii} \quad \text{and} \quad [\mathbf{M_c}]_{M+1,M+1} = \mu_c,$$

where $t_c = 1/\mu_c$ is the mean time to process a checkpoint.

- The conditional distributions are then:

$$\mathbf{B_c} := \mathbf{M_c}(\mathbf{I} - \mathbf{P_c})$$

$$\bar{F}_{\rm ou}(t) := \mathbb{P}\mathbf{r}[T_{\rm ou} > t] = \mathbf{p_o} \exp(-t\mathbf{B_c})\, \boldsymbol{\epsilon}'_{\boldsymbol{u}} \,/\, p_{\rm ou}$$

$$u \in \{e,\, c\}$$

# The MMSC Model - IV



- If the system execution takes the path described by *oe*, the process ends. But if the path leads to $m$, the system checkpoints after it's execution.
- We must define a *restart vector* $\mathbf{p_c}$ as an entrance vector into the system corresponding to where the execution of the system begins again after checkpointing.
- $\mathbf{p_c}$ is composed of the transition probabilities out of state $m$:

$$\mathbf{p_c} := [\mathbf{P}_{m1}, \mathbf{P}_{m2}, \dots, \mathbf{P}_{mM}, 0]/(1 - q_m)$$

# The MMSC Model - V



- So the probability of the system finishing after checkpointing without returning to $m$ is

$$p_{ce} := \mathbf{p_c} \boldsymbol{\epsilon}'_{\boldsymbol{e}}$$

- The probability of the system returning to $m$ after already checkpointing (to save a more recent state of the system) is

$$p_{cc} := \mathbf{p_c} \boldsymbol{\epsilon}'_{\boldsymbol{c}}$$

- The time distribution for these two events are (for $u \in \{c, e\}$):

$$\bar{F}_{cu}(t) := \mathbb{P}\mathbf{r}[T_{cu} > t] = \mathbf{p_c} \exp(-t\mathbf{B_c}) \boldsymbol{\epsilon}'_{\boldsymbol{u}} \, / \, p_{cu}.$$

# The MMSC Model - VI



- What has been described can be thought of as an embedded Markov chain with four nodes whose service time distributions are given by each of the $\bar{F}_{ab}$.
- The transition matrix for this process is:

$$\hat{\boldsymbol{P}}_{\boldsymbol{c}} := \begin{array}{c|cccc|} & \mathrm{oe} & \mathrm{oc} & ce & cc \\ \hline \mathrm{oe} & 0 & 0 & 0 & 0 \\ \mathrm{oc} & 0 & 0 & p_{ce} & p_{cc} \\ ce & 0 & 0 & 0 & 0 \\ cc & 0 & 0 & p_{ce} & p_{cc} \end{array} \quad \text{with} \quad \hat{\boldsymbol{p}}_{\boldsymbol{c}} := [p_{\mathrm{oe}},\ p_{\mathrm{oc}},\ 0,\ 0]$$

- expected number of visits to $C$:
  $\mathbb{E}[N_c] = p_{oc} + p_{oc}\, p_{cc} \,/\, p_{oe} = p_{oc} \,/\, p_{oe}$

# Diagrams of the Markov Chain 26

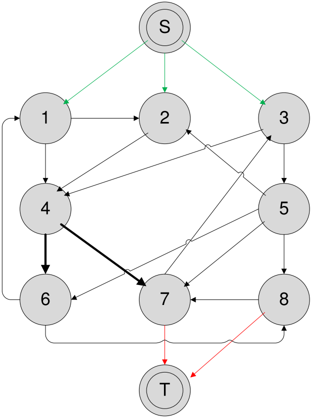


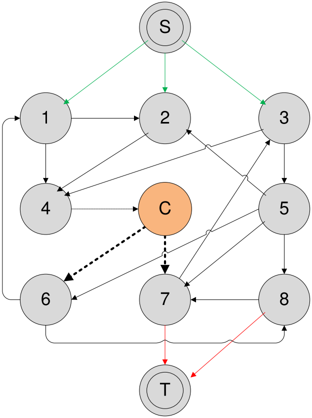

# Applying RESTART to the MMSC 27

- $\hat{\boldsymbol{P}}_{\boldsymbol{c}}$, together with the ME service time distributions of each node is an ME representation (but only for $\beta = 0$).
- If there is a failure, the system only has to redo whatever work had been accomplished within the node that had failed.
- Thus we can get $\mathbb{E}[X_u(\beta)]$ and $\mathbb{E}[X_u^2(\beta)]$ for $u \in \{\mathrm{oe}, \mathrm{o}c, cc, ce\}$.
- With the first two moments of the distribution for each node, we can get $\mathbb{E}[X_c^\ell(\beta)]$ $(\ell = 1, 2)$ as follows:

# Mean and Variance of $X_c(\beta)$



- Define the 4-matrices

$$[\hat{\boldsymbol{T}}_{\boldsymbol{c}}]_{uu} := \mathbb{E}[X_u(\beta)], \quad \hat{\boldsymbol{V}}_{\boldsymbol{c}} := [\hat{\boldsymbol{I}} - \hat{\boldsymbol{P}}_{\boldsymbol{c}}]^{-1}\, \hat{\boldsymbol{T}}_{\boldsymbol{c}}, \qquad \text{and}$$

$$[\hat{\boldsymbol{\Gamma}}]_{uu} := C_u^2 - 1,$$

where $C_u^2 = \sigma_u^2(\beta) / \left(\mathbb{E}[X_u(\beta)]\right)^2$ is the squared coefficient of variation of $X_u(\beta)$.

- Then

$$\mathbb{E}[X_c(\beta)] = \hat{\boldsymbol{p}}_{\boldsymbol{c}}\, \hat{\boldsymbol{V}}_{\boldsymbol{c}}\, \hat{\boldsymbol{\epsilon}}'$$

and

$$\sigma_c^2(\beta) = \sigma_{exp}^2 + \hat{\boldsymbol{p}}_{\boldsymbol{c}}\, \hat{\boldsymbol{V}}_{\boldsymbol{c}}\, \hat{\boldsymbol{T}}_{\boldsymbol{c}}\, \hat{\boldsymbol{\Gamma}}\, \hat{\boldsymbol{\epsilon}}'$$

where $\sigma_{exp}^2 = 2(\hat{\boldsymbol{p}}_{\boldsymbol{c}}\, \hat{\boldsymbol{V}}_{\boldsymbol{c}}{}^2\, \hat{\boldsymbol{\epsilon}}') - (\hat{\boldsymbol{p}}_{\boldsymbol{c}}\, \hat{\boldsymbol{V}}_{\boldsymbol{c}}\, \hat{\boldsymbol{\epsilon}}')^2$ is the variance of the similar exponential network.

UCONN

# Asymptotic Properties of $T$, $X(\beta)$ and $X_c(\beta)$ - I



- The exponential tail for $T$ is determined by $\lambda_s = \lambda_{min}$, where $\lambda_{min}$ is the smallest eigenvalue of **B**.
  - If **P** is a *feed-forward* matrix, then the eigenvalues of **B** are the service rates, $\mu_i$, of the nodes (assuming $P_{ii} = 0$), so $\lambda_s = \mathrm{Min}\{\mu_\mathrm{i}\}$.
  - If there are some *feed-back* loops, then $\lambda_s$ may be smaller. In any case, $\lambda_s \leq \mathrm{Min}\{\mu_\mathrm{i}\}$.

UCONN

## Asymptotic Properties of $T$, $X(\beta)$ and $X_c(\beta)$ - II



- The PT index for $X(\beta)$ is $\alpha = \lambda_s/\beta$
- Let $\lambda_{us}$ ($u \in \{oc, oe, cc, ce\}$) be the exponential parameter for $F_u(t)$. Then $\lambda_{cs} := \mathrm{Min}\{\lambda_{us}\}$ determines $\alpha_c = \lambda_{cs}/\beta$
- If **P** is feed-forward, then the index for $X_c(\beta)$ is the same as for $X(\beta)$ (although $\mathbb{E}[X_c(\beta)] < \mathbb{E}[X(\beta)]$)
- If **P** has loops, and the checkpoint is inserted within a loop then $\alpha_c$ can be much larger.
- Even if $\hat{\boldsymbol{P}}$ has feedback ($p_{cc} > 0$), $\alpha_c$ does NOT change.

$$\mathbf{P} = \begin{bmatrix} 0 & .7 & 0 & 0.30 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1.00 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.75 & .25 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & .4 & .6 & 0 \\ 0 & .3 & 0 & 0 & 0 & .3 & .1 & .3 \\ .8 & 0 & 0 & 0 & 0 & 0 & 0 & .2 \\ 0 & 0 & .75 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & .1 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{q}' = \begin{bmatrix} .00 \\ .00 \\ .00 \\ .00 \\ .00 \\ .00 \\ .25 \\ .90 \end{bmatrix}$$

$[\mathbf{q}' = (\mathbf{I} - \mathbf{P})\boldsymbol{\varepsilon}']$, with entrance vector

$$\mathbf{p} = [0.60,\ 0.20,\ 0.20,\ 0,\ 0,\ 0,\ 0,\ 0], \quad \text{and}$$

$$\mathbf{M} = \mathrm{Diag}[1.2,\ 2.3,\ 3.4,\ 0.8,\ 2.0,\ 2.4,\ 6.5,\ \mu_c]$$

.

# Diagrams of the Markov Chain With Node Service rates 32

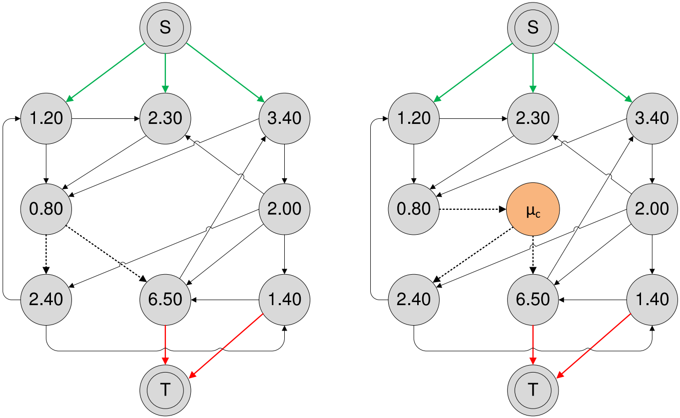

# Reliability Functions, $\bar{F}$, $\bar{F}_u$ $(u \in \{oc, oe, cc, ce\})$ 33

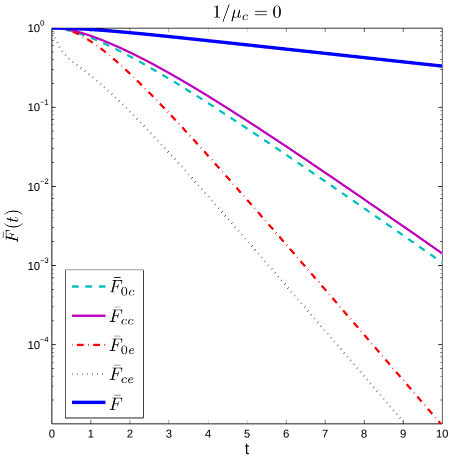

# Asymptotic Tail Parameter



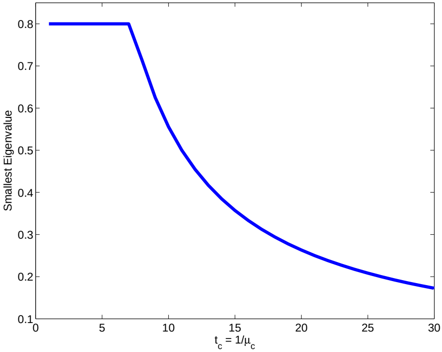


$$\lambda_{sc} = \mathrm{Min}\{\mu_4,\ \mu_c\}$$

# The Checkpointing Effect - I ( $\mathbb{E}[X_c(\beta)]$ ) 35

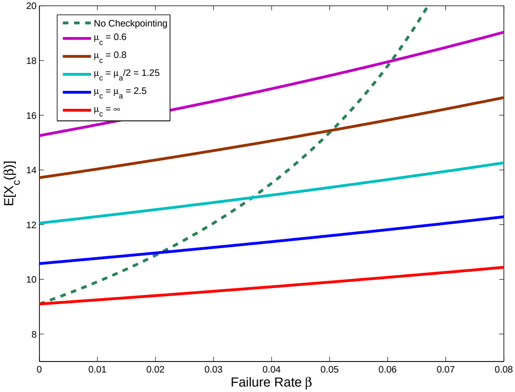

# The Checkpointing Effect - II ($\mathbb{E}[X_c(\beta)]\,(1-\beta/\lambda_{sc})$) 36

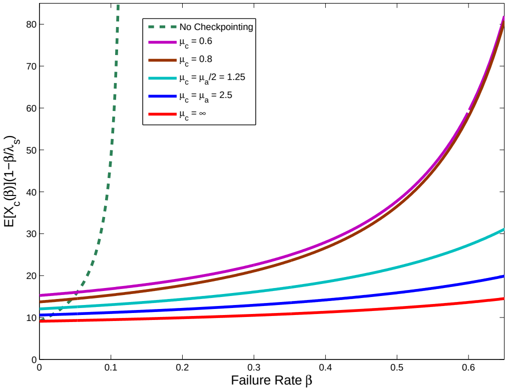

# Squared Coefficient of Variaton ($C_v^2 := \sigma^2 / \mathbb{E}[X_c(\beta)]^2$) 37

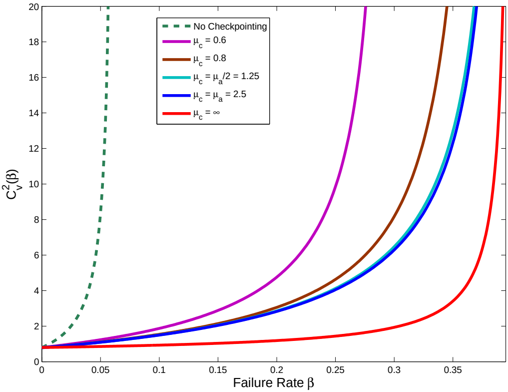

# Blowup of Squared Coefficient of Variaton ($C_v^2 := \sigma^2 / \mathbb{E}[X_c(\beta)]^2$) 38

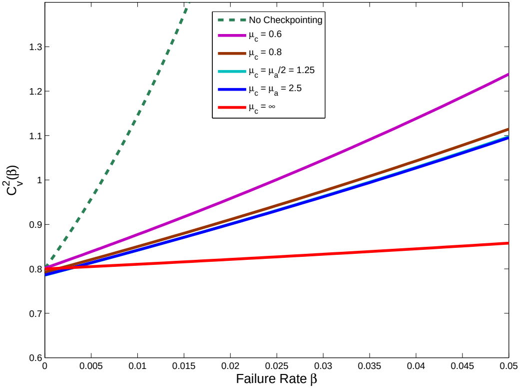


UCONN

# Some Unresolved Questions 39

- How large must $x$ be before the asymptotic formula is a "good" approximation to $\bar{H}(x)$?
- How robust is the method if the nodes have non-exponential service times?
- What is to be done if the failure distribution is not exponential?

# Simulation of $X(\beta)$ and Asymptotic Formulas for Exponential, Hyperexponential, and Erlangian Functions 40

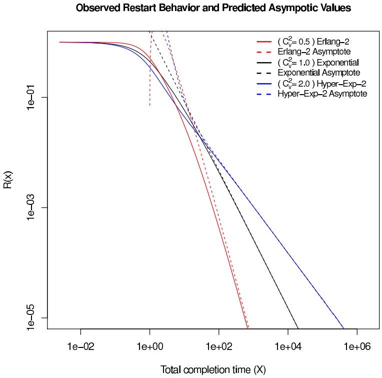

# Relative Difference Between Simulation and Analytic Asymptotic formula [Abs(Sim − Asymp)/Asymp]



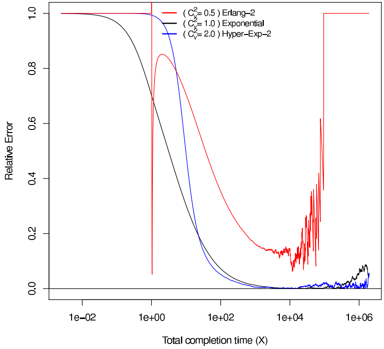

## Conclusion



- We can compute the moments of $\bar{H}(x)$;
- We can get the asymptotic index, $\alpha_c$;
- We can't get $\bar{H}(x)$.

'