Bounding the Mean Cumulated Reward up to Absorption

A. P. Couto and G. Rubino

COPPE & INRIA/IRISA

MAM 2006

Outline

Motivations

- paper area: evaluating performance, dependability and performability measures using Markov models
- models considered: no closed-form solutions, large state spaces
- problem: no exact (numerically speaking) solution available
- goal: obtain approximations of standard metrics *with error bounds*, avoiding the large state spaces
- starting point: in usual dependability contexts, we have *stiffness*

Main ideas

- We compute lower and upper bound of the Expected Cumulated Reward up to Absorption (ECRA),
- for a system modeled by an absorbing Markov chain with states weighted by rewards.
- The approach merges the idea of state space reduction using stochastic complementation with a path-based procedure.
- The idea of fast and slow transitions and states is used.

Markov model

- state space: $\{1, \cdots, M\} \cup \{M+1\}$, with $M \gg 1$
- CTMC X where $1, \cdots, M$ are transient and $M+1$ is absorbing; $\Omega = \{1, \cdots, M\}$
- infinitesimal generator of X and its initial distribution are decomposed with respect to $(\Omega, \{M+1\})$:

$$\begin{pmatrix} Q & q \\ 0 & 0 \end{pmatrix}, \quad q = -Q\,\mathbb{1}; \qquad \begin{pmatrix} \alpha \\ 0 \end{pmatrix}$$

- from any transient state i there is a path to state $M+1$

Adding rewards

- r_i: reward rate at state i; $r_{M+1} = 0$.
- The Cumulated Reward until Absorption is the random variable

$$R_\infty^X = \int_0^\infty r_{X(t)}\, dt.$$

- Target: $\mathrm{E}(R_\infty^X)$, = ECRA.

Working in discrete time

- The canonical embedded DTMC is decomposed on $(\Omega, \{M+1\})$:

$$\begin{pmatrix} P & p \\ 0 & 1 \end{pmatrix}, \quad P_{ij} = Q_{ij}/(-Q_{ii}), \quad p = (I - P)\,1$$

- Let $c = (\cdots r_i/q_i \cdots)^{\mathsf{T}}$.
- Let $\mathrm{E}(R_\infty^X \mid X(0) = i) = \rho_i$, and $\rho = (\cdots \rho_i \cdots)^{\mathsf{T}}$.
- Then, $\rho = (I - P)^{-1} c$ (and ECRA $= < \alpha\, , \, \rho >$).
- From now, goal = bounding ECRA through bounding ρ.

Key concept: fast and slow states

- transition (i,j) is fast iff $Q_{i,j} \geq \theta > 0$
- transition (i,j) is slow iff $0 < Q_{i,j} < \theta$
- state i is fast iff for some j, transition (i,j) is fast
- state $i \neq M+1$ [1] is slow iff it is not fast
- $S = \{i : \; i \text{ is slow}\} \ni X(0)$; $F = \{i : \; i \text{ is fast}\}$
- Assumption: there is no circuit only composed of fast transitions.
- We can decompose matrices such as Q or P, and vectors such as α, or ρ, on S, F. For instance,

$$P = \begin{pmatrix} P_S & P_{SF} \\ P_{FS} & P_F \end{pmatrix}, \quad \rho = \begin{pmatrix} \rho^S \\ \rho^F \end{pmatrix}.$$

[1] State $M+1$ excluded just to simplify the presentation.

Going to discrete time through uniformization

- Y: DTMC obtained by uniformizing X w.r.t. η

$$\text{t.p.m. (on } \Omega, \{M+1\}\text{):} \quad \begin{pmatrix} U & u \\ 0 & 1 \end{pmatrix}$$

$$U = I + Q/\eta, \; u = (I - U)\mathbb{1} = q/\eta$$

- Also, $U = \begin{pmatrix} U_S & U_{SF} \\ U_{FS} & U_F \end{pmatrix}$.
- Define $R_\infty^Y = \sum_{k=0}^{\infty} r_{Y(k)}$ and $\text{ECRA}^Y = \text{E}(R_\infty^Y)$.

Going to discrete time through uniformization (cont.)

- We have $\text{ECRA}^Y = < \alpha\,,\, (I - U)^{-1} c >$
- and $\text{ECRA}^X = \dfrac{\text{ECRA}^Y}{\eta}$

So, from this point, our target is ECRA^Y.

- To simplify, let us rename $\rho_i = \text{E}(R_\infty^Y \mid Y(0) = i)$, $i \in S$, and ρ the corresponding vector.
- We will bound this conditional expected cumulated reward vector ρ (so, working on the DTMC Y).

Going to discrete time through uniformization (cont.)

- Write

$$R_{\infty}^{Y,S} = \sum_{k=0}^{\infty} r_{Y(k)} 1_{Y(k)\in S}, \quad R_{\infty}^{Y,F} = \sum_{k=0}^{\infty} r_{Y(k)} 1_{Y(k)\in F},$$

then $\rho_i^S = \mathrm{E}(R_{\infty}^{Y,S} \mid Y(0) = i)$, $\rho_i^F = \mathrm{E}(R_{\infty}^{Y,F} \mid Y(0) = i)$.

- We will start from

$$\rho = \rho^S + \rho^F.$$

- Intuitively, in case of a highly dependable system, $\rho \approx \rho^S$.

Decomposing the state space

- Let $\widetilde{Y}$ be the DTMC obtained by stochastically complementing the subset S of slow states (plus state $M+1$) in Y.
- Process $\widetilde{Y}$ is called the reduction of Y with respect to $S \cup \{M+1\}$. It can be also defined by "sampling" from Y when Y visits $S \cup \{M+1\}$.
- The t.p.m. of $\widetilde{Y}$ is

$$\begin{pmatrix} \widetilde{U} & \widetilde{u} \\ 0 & 1 \end{pmatrix}, \; \widetilde{u} = (I - \widetilde{U})1.$$

where $\widetilde{U} = U_S + U_{SF}(I - U_F)^{-1} U_{FS}$

Decomposing the state space (cont.)

- Define $R_\infty^{\widetilde{Y}} = \sum_{k=0}^{\infty} r_{\widetilde{Y}(k)}$ and $\mathrm{ECRA}^{\widetilde{Y}} = \mathrm{E}(R_\infty^{\widetilde{Y}})$
- For $i \in S$, define $\widetilde{\rho} = \mathrm{E}(R_\infty^{\widetilde{Y}} \mid \widetilde{Y}(0) = i)$, and denote by $\widetilde{\rho}$ the corresponding vector.
- Observe that $\widetilde{\rho} = \rho^S$.

Lower bounding procedure

- Fix some $N \geq 2$. For all $i, j \in S$, we compute the $|S| \times |S|$ matrix U'_N where $(U'_N)_{ij}$ is the probability for Y to go from i to j by remaining in F in between, in no more than N steps.
- For this, we use a BFS (Breadth First Search) with depth N, to build all paths of the form

$$(i, f_1, f_2, \cdots, f_{N-1}, j),$$

where $f_1, f_2, \cdots, f_{N-1} \in F$.
- Clearly, $U'_N \leq \widetilde{U}$.
- If we compute now $\rho' = (I - U'_N)^{-1} c^S$, we have

$$\rho'_i \leq \widetilde{\rho}_i = \rho_i^S \leq \rho_i^S + \rho_i^F = \rho_i.$$

- Our lower bound is thus ρ' obtained by solving the linear system with size $|S|$

$$(I - U'_N)\rho' = c^S.$$

Remark: if by chance the model allows to compute $\widetilde{U}$, then our lower bound of ρ will be simply $\widetilde{\rho}$.

Upper bounding procedure

- Writing $\rho_i = \rho_i^S + \rho_i^F$,
  we will upper-bound both terms in this representation.

Upper-bound of $\rho^S = \widetilde{\rho}$

- We had $U'_N \le \widetilde{U}$; write $U''_N = \widetilde{U} - U'_N \ge 0$.

$$U''_N = U_{SF}(U_F^{N-1} + U_F^N + \cdots)U_{FS} = U_{SF}\,U_F^{N-1}\,V\,U_{FS},$$

where $V = (I - U_F)^{-1}$.
V_{ij} is the mean # of visits made by Y to state j, starting at i, before leaving F $(i, j \in F)$.

A bounding Lemma

Assuming there are no circuits in F only composed of fast transitions, the mean sojourn time of Y in F can be upper bounded by

$$\sigma = \frac{\eta}{\mu} \frac{1 - (1 - \psi)^D}{\psi (1 - \psi)^D},$$

where

$$\lambda = \max_{i \in F} \sum_{j:\,(i,j)\ \text{slow}} Q_{ij}, \ < \ \mu = \min_{i \in F} \sum_{j:\,(i,j)\ \text{fast}} Q_{ij},$$

$\psi = \lambda / \mu < 1$ and D is the maximal # of fast transitions starting in F, before leaving F.

Upper-bound of $\rho^S = \widetilde{\rho}$ (cont.)

- Now, we can bound U''_N as follows:
$$U''_N \leq B = \sigma U_{SF} U_F^{N-1} \mathbb{1} U_{FS}.$$
(1 is here a $|F| \times |F|$ matrix composed of ones.)
- We have $\lim_{N \to \infty} B = 0$.

$$\begin{aligned} \text{Finally,} \quad \rho^S = \widetilde{\rho} &= \left(\sum_{k=0}^{\infty} \widetilde{U}^k \right) c^S \\ &= \left(\sum_{k=0}^{\infty} (U'_N + U''_N)^k \right) c^S \\ &\leq \left(\sum_{k=0}^{\infty} (U'_N + B)^k \right) c^S = \varphi. \end{aligned}$$

See that φ is computed again solving a new linear system with size $|S|$ (the system $(I - U'_N - B)\varphi = c^S$).

Upper-bound of ρ^F

- For upper bounding ρ^F, consider ν_i, the mean # of visits that Y makes to subset F when starting from $i \in S$; we have

$$\rho_i^F \le \nu_i \sigma r^*, \quad r^* = \max\{r_k,\ k \in F\}.$$

- If $W = (I - \widetilde{U})^{-1}$, then $\nu_i = \sum_{j \in S} W_{ij} U_{jF}$
- From the preceding results,

$$\rho_i^F \le \sigma r^* \sum_{j \in S} W_{ij} U_{jF}.$$

Upper-bound of ρ^F (cont.)

- If u_F is the vector indexed on S whose jth entry is U_{jF}, then $\nu = W\, u_F$ and
$$\rho^F \le \sigma r^* W\, u_F.$$
- Now, $Wu_F = \sum_{n\ge 0} \widetilde{U}^n u_F \le \sum_{n\ge 0} (U'_N + B)^n u_F = \tau$.
- So, our upper bound of ρ^F is $\sigma r^* \tau$ where τ is the solution to $(I - U'_N - B)\tau = u_F$ (for N large enough so that $\|U'_N + B\| < 1$).

Remark: if $\widetilde{U}$ is available, then ν is computed from $(I - \widetilde{U})\nu = u_F$ and the upper bound is $\sigma r^* \nu$.

A system with identical components

- The system is composed of m identical and independent components.
- Each component evolves among 3 states: operational, down and being repaired, definitely down.
- If Φ = failure rate, c = covering factor and γ = repair rate, we used $c\Phi = 10^{-6}$, $(1-c)\Phi = 10^{-7}$ and $\gamma = 1$.
- The threshold θ was $\theta = \gamma$.
- The symmetries of the system allowed us to work on a reduced state space (strong lumpability).
- We used $m = 10, 20, 30, 40, 50$ leading to $66, 231, 496, 861, 1326$ states resp.
- Here, $|S| = m$.

A system with identical components: some results

m	Lower Bound	Exact	Upper Bound
10	2.9285×10^7	2.9285×10^7	2.9285×10^7
20	3.5974×10^7	3.5975×10^7	3.5975×10^7
30	3.9932×10^7	3.9933×10^7	3.9933×10^7
40	4.2765×10^7	4.2765×10^7	4.2767×10^7
50	4.4959×10^7	4.4959×10^7	4.4961×10^7

Same system, but without using strong lumpability

- Two cases considered:
  1. $m = 11$ components $\rightsquigarrow$ 177.147 states; $|S| = 2.048$
  2. $m = 12$ components $\rightsquigarrow$ 531.441 states; $|S| = 4.096$
- depth in the BFS: $N = 6$

Numerical values:

m	Lower Bound	Exact	Upper Bound
11	3.0116×10^7	3.0177×10^7	3.0206×10^7
12	3.0925×10^7	3.1026×10^7	3.1031×10^7

Same system, but now the components belong to 3 classes ("CPU", "memory", "disks")

- Associated with each class there is a cost when the component is down.
- We bound here the ECRA.
- We don't use the strong lumpability property.
- Example: with 4 CPU, 5 memory units and 3 disks, we get 531 441 states; $|S| = 12$.
- Depth in the BFS: $N = 6 \rightsquigarrow$ relative error $\approx 10^{-4}$.

As before, but we replace exponentials by equivalent Coxian distributions

- The exact value is still available,
- and the state space can be huge.
- We don't use the strong lumpability property.
- An example: in a configuration we get ≈ 430 M states and $|S| \approx 1.7$ M states.
- Depth in the BFS: $N = 7 \rightsquigarrow$ a relative error $\approx 10^{-3}$.

Some concluding remarks

- Some refinements under study.
- In cases where $|S| \ll 1$, we are exploring the use of the stochastic complement with respect to a subset $S' \supset S$ plus the absorbing state (some results also in the paper [2]).
- Other linear algebra ideas are also under study, based on previous work done on stationary measures (on irreducible models).
  This can be of interest in particular in cases where $|S|$ can still be large.

[2] Presented at the "A. A. Markov Anniversary Meeting" (MAM), Charleston, USA, June 12–14 2006.