Mean-Payoff Optimization in Continuous-Time Markov Chains with Parametric Alarms

Christel Baier Clemens Dubslaff Ľuboš Korenčiak
Antonín Kučera Vojtěch Řehák

IFIP WG 2.2 2017 Meeting, Bordeaux

Dynamic power management of a disk drive

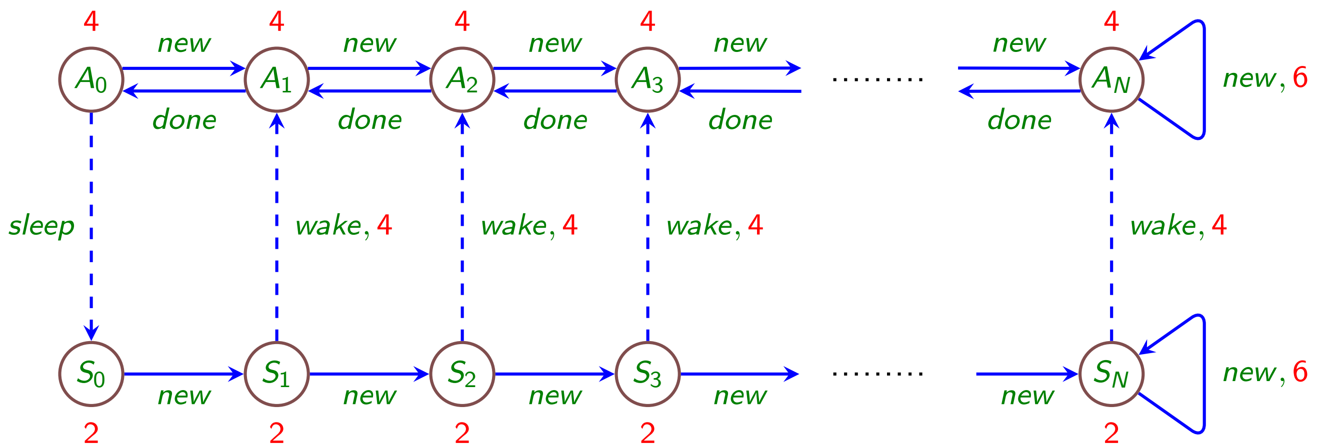

- What timeout values achieve the minimal long-run average power consumption?

CTMC with parametric alarms (1)

"Ordinary" CTMC

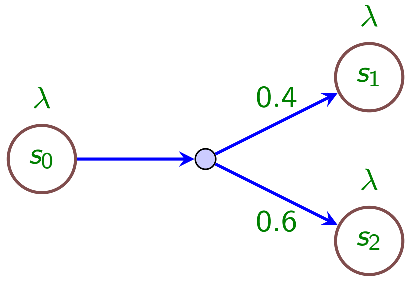

CTMC with alarms $\{a_1, \ldots, a_n\}$

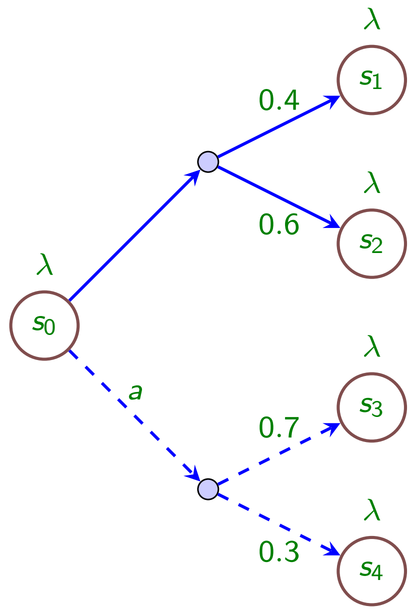

CTMC with parametric alarms (2)

- In CTMC with parametric alarms $\{a_1, \ldots, a_n\}$, the distributions associated to $\{a_1, \ldots, a_n\}$ are not fixed but parameterized by a single parameter.
- Restrictions:
  - At most one alarm is active in each state.
  - Each alarm is set in precisely one state.
- After fixing the parameters, we obtain a fully stochastic CTMC with alarms.
- Can we compute parameter values achieving ε-optimal mean-payoff?

Computing ε-optimal parameter values

1. Given a CTMC with parametric alarms and $\varepsilon > 0$, we compute a discretization constant $\kappa > 0$ such that ε-optimal parameter values are among the (finitely many) κ-discretized values.
2. We construct a semi-Markov decision process $\mathcal{M}$ where the actions correspond to discretized parameter values. Thus, the original problem reduces to computing an optimal strategy for $\mathcal{M}$.
3. The set of states of $\mathcal{M}$ is small but the number of actions is very large. We employ a symbolic technique which avoids the explicit construction of these actions.

Computing ε-optimal parameter values (2)

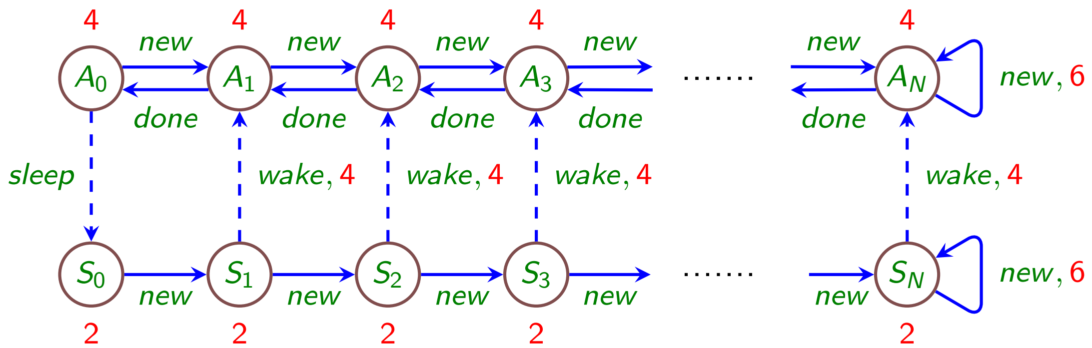

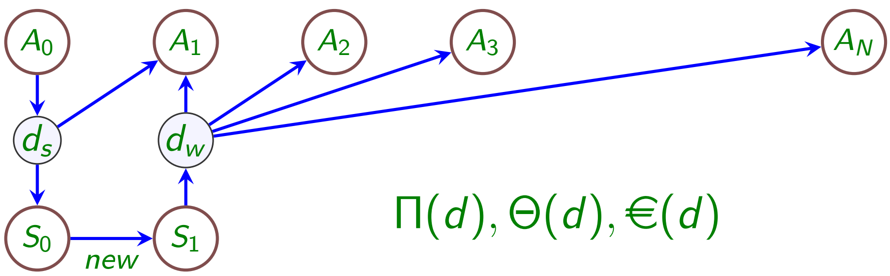

$\Pi(d), \Theta(d), \text{€}(d)$

Computing ε-optimal parameter values (3)

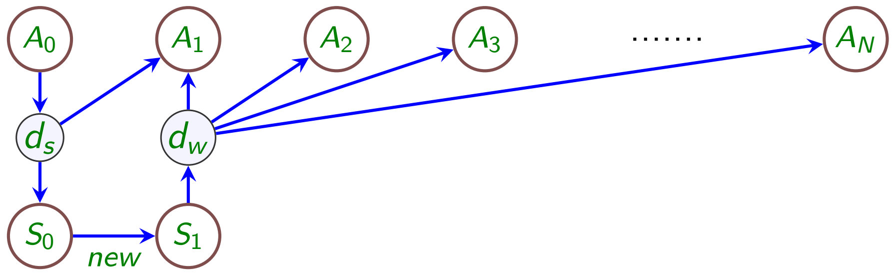

- An optimal for $\mathcal{M}$ can be computed by strategy iteration.
- Each action d is ranked by a function $F(d)$ depending on $\Pi(d)$, $\Theta(d)$, and €(d). The goal is to find an action with minimal $F(d)$.
- We express $F(d)$ analytically, compute its derivative, and consider only a small set of actions close to the local minima of $F(d)$.
- Applicable to alarms with Dirac (fixed-delay), uniform, and Weibull distributions.

Experiments (disk drive example)

- We considered $N \in \{2, 4, 6, 8\}$, $\varepsilon \in \{0.1, 0.01, 0.001, 0.0005\}$.
- The upper and lower bounds for the timeouts were 0.1 and 10 time units, respectively.
- The required discretization step ranges from 10^{-25} to 10^{-19}.

Experiments (disk drive example), cont.

N	ε	creating time [s]	solving time [s]	poly degree
2	0.1	0.15	0.24	46
	0.01	0.15	0.25	46
	0.001	0.16	0.28	53
	0.0005	0.16	0.33	53
4	0.1	0.14	0.25	46
	0.01	0.16	0.25	46
	0.001	0.16	0.28	53
	0.0005	0.16	0.33	53
6	0.1	0.16	0.35	46
	0.01	0.16	0.35	46
	0.001	0.17	0.40	53
	0.0005	0.18	0.43	53
8	0.1	0.19	0.35	46
	0.01	0.19	0.35	46
	0.001	0.20	0.43	53
	0.0005	0.22	0.44	53

Limitations, future work

- Each alarm has to be set in precisely one state. Hence, we cannot model systems of concurrently running components. POMPD techniques might help?
- Other objectives?
- Multi-criteria parameter optimizations.