On Optimal Policies for Energy-Aware Servers

Vincent J. Maccio and Douglas G. Down

Department of Computing and Software
McMaster University
Hamilton, Ontario, Canada

Goal

1. Derive optimal policies, when energy used by the system is a concern
2. Two decisions, when to turn it off, when to turn it on

Previous Work

1. Gandhi et al. [1], [3] determined a set of optimal policies for single server system under the cost function $\mathbb{E}[R]\mathbb{E}[E]$. Solved steady state distribution for a server which instantly turns off when it idles, and begins to turn back on once a job arrives.
2. Various authors [4], [5], [6], [7] use the a cost function of the form $\mathbb{E}[R] + \beta_1\mathbb{E}[E] + \beta_2\mathbb{E}[Sw]$, with great analysis. However, optimal policies remained unknown (even for single server).
3. Various authors [7], [8] examine multi-server systems under certain policies. However, optimal policies remain unknown.

Our Contribution

Give one the tools to derive the optimal policy for a single on/off server system under any cost function of the form:

$$f(\beta, w) = \sum_{i=1}^{M} \beta_i \mathbb{E}[R]^{w_{R,i}} \mathbb{E}[E]^{w_{E,i}} \mathbb{E}[Sw]^{w_{Sw,i}},$$

where all $\beta, w \geq 0$.

Model

1. Single server, infinite queue
2. Arrivals follow Poisson process, rate λ
3. Service times exponentially distributed, rate μ
4. Server can be in one of four (energy) states - *OFF*, *BUSY*, *IDLE*, *SETUP*
5. Corresponding energy values, $E_{Off} = 0$, E_{Setup}, E_{Busy} and E_{Idle}
6. Turn on time, exponentially distributed, rate γ
7. Time to wait before moving to low state, exponentially distributed, rate α (idling times)
8. Time needed to turn off equals zero
9. Server in low state waits for k arrivals before beginning to turn on

Markov Chain

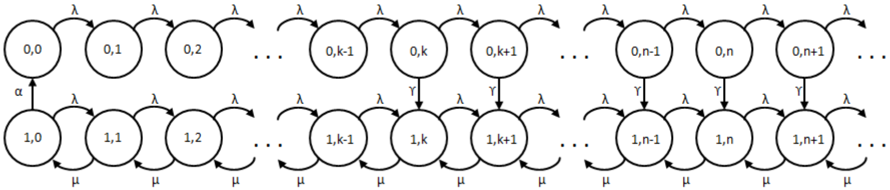

Markov Chain - Off

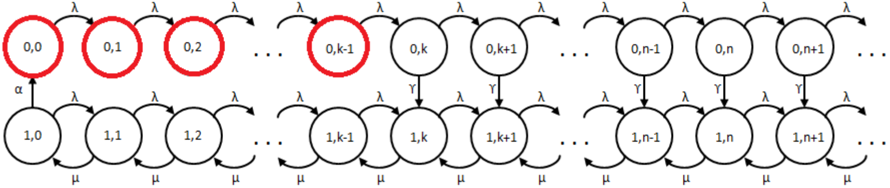

Markov Chain - Setup

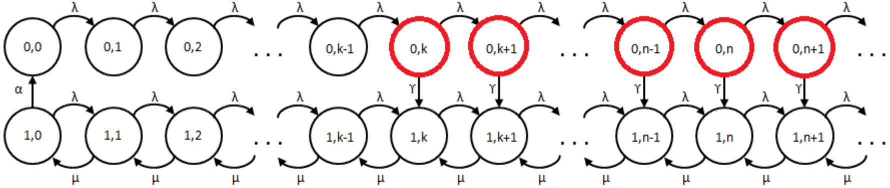

Markov Chain - Busy

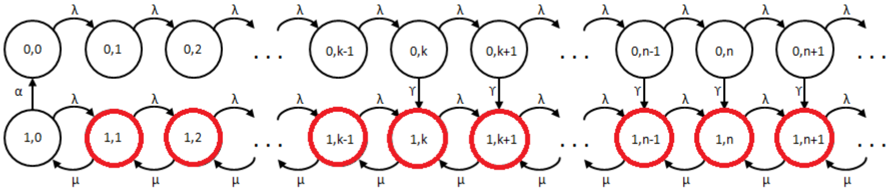

Markov Chain - Idle

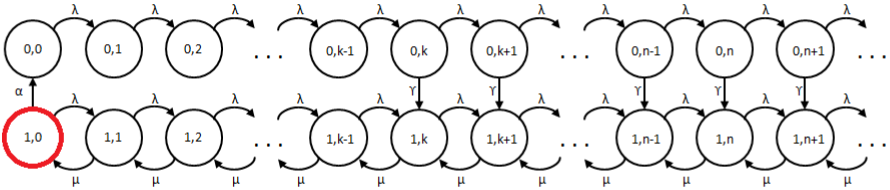

Steady-State Distribution

Using similar methods as in Gandhi et al. [3],

$$
\begin{aligned}
\pi_{0,n} &= \pi_{0,0}, \qquad 0 \leq n < k \\
\pi_{0,n} &= \pi_{0,0}\left(\frac{\lambda}{\lambda+\gamma}\right)^{n-(k-1)}, \qquad k \leq n \\
\pi_{1,n} &= \pi_{0,0}\left(\frac{\lambda}{\alpha}\rho^{n} + \frac{\lambda}{\mu-\lambda}(1-\rho^{n})\right), \qquad 0 \leq n < k \\
\pi_{1,n} &= \pi_{0,0}\left[\left(\frac{\lambda}{\alpha} - \frac{\lambda}{\mu-\lambda}\right)\rho^{n} + \frac{1}{\mu-\lambda-\gamma}\left((\lambda+\gamma)\left(\frac{\lambda}{\lambda+\gamma}\right)^{n-(k-1)} - \frac{\gamma}{1-\rho}\rho^{n-(k-1)}\right)\right], \quad k \leq n \\
\pi_{0,0} &= (1-\rho)\frac{\alpha\gamma}{k\alpha\gamma + \alpha\lambda + \lambda\gamma}
\end{aligned}
$$

Properties of Optimal Policy

The following properties carry over from Gandhi et al. [1]:

1. The decision to start transitioning between lower and higher energy states is made at the moment a job arrives to the system.
2. When the server idles, it is always optimal to immediately turn the server off, or to always keep the server on.
3. If there are jobs in the system and the server is in its higher energy state, the server will never move to its lower energy state.

The Three Metrics

$$\mathbb{E}[Sw] = (1-\rho)\frac{\alpha\lambda\gamma}{k\alpha\gamma + \alpha\lambda + \lambda\gamma}$$

$$\mathbb{E}[E] = \mathbb{E}[E_{M/M/1}] + \frac{(1-\rho)\alpha}{k\alpha\gamma + \alpha\lambda + \lambda\gamma}(\lambda E_{Setup} - (\lambda + k\gamma)E_{Idle})$$

$$\mathbb{E}[R] = \mathbb{E}[R_{M/M/1}] + \frac{1}{\gamma}\frac{\alpha(\lambda + k\gamma)}{k\alpha\gamma + \alpha\lambda + \lambda\gamma} + \frac{1}{2\lambda}\frac{k\alpha\gamma(k-1)}{k\alpha\gamma + \alpha\lambda + \lambda\gamma}$$

	Optimal Values of	
Metric	α	k
$\mathbb{E}[R]$	0	1
$\mathbb{E}[E]$	0 or $\to \infty$	$\to \infty$
$\mathbb{E}[Sw]$	0	$\to \infty$

Weighted Sum Cost Function

$$\mathbb{E}[R] + \beta_1 \mathbb{E}[E] + \beta_2 \mathbb{E}[Sw]$$

Let $r_{idle} = E_{Idle}/E_{Busy}$ and $r_{setup} = E_{Setup}/E_{Busy}$.

1. Choose $\alpha = \infty$ if

$$\beta_1(1-\rho)\lambda\gamma(k\gamma+\lambda) \leq \lambda(\lambda+k\gamma)+\gamma k(k-1)+\beta_1(1-\rho)\lambda^2\gamma r_{setup}+\beta_2(1-\rho)\lambda^2\gamma^2$$

   Otherwise, choose $\alpha = 0$.

2. Solve

$$0 = \frac{\gamma^2}{2\lambda}k^2 + \gamma k - (\lambda + (1-\rho)\lambda\gamma(\beta_1 r_{setup} + \beta_2\gamma))$$

   If a valid solution, check $\lceil k^* \rceil$ and $\lfloor k^* \rfloor$.

Generalizing Distributional Assumptions

1. Allowing all distributions to be arbitrary, other than the arrival process being Poisson, $\mathbb{E}[E]$ and $\mathbb{E}[Sw]$ are the same as for the setting with all distributions exponential.
2. $\mathbb{E}[E]$ and $\mathbb{E}[Sw]$ are dependent only on the means of the underlying distributions.
3. If $r_{idle} < \frac{\lambda}{k\gamma+\lambda} r_{setup}$ it is always optimal to leave the server on.

Generalizing Distributional Assumptions (cont)

If the arrival stream remains Poisson, and the idling time of the system remains exponentially distributed, can derive the mean response time using M/G/1-type arguments (embedded chain).

$$\mathbb{E}[N] = \mathbb{E}[N_{M/G/1}] + \frac{\alpha}{\alpha+\lambda}\left[\rho\frac{(k-1)\gamma+\lambda}{\gamma} + \frac{1}{2}\Gamma\right] + \alpha\frac{(k-1)\gamma+\lambda}{k\alpha\gamma+\alpha\lambda+\lambda\gamma}\left[\frac{1}{2} - \rho - \rho\frac{\alpha}{\alpha+\lambda}\left(\frac{(k-1)\gamma+\lambda}{\gamma}\right) - \frac{1}{2}\frac{\alpha}{\alpha+\lambda}\Gamma\right]$$

where,

$$\Gamma = (k-1)^2 + (2k-1)\frac{\lambda}{\gamma} + \lambda^2\sigma_\Gamma^2$$

Constrained Optimization

Minimize the expected energy subject to $\mathbb{E}[R] \leq T$, where $\frac{1}{\mu-\lambda} \leq T \leq \frac{1}{\mu-\lambda} + \frac{1}{\gamma}$. Fix $k = 1$.

$$\alpha = \frac{\lambda\gamma^2}{\lambda+\gamma} \frac{T - \mathbb{E}[R_{M/M/1}]}{1 - \gamma(T - \mathbb{E}[R_{M/M/1}])}$$

Random Routing

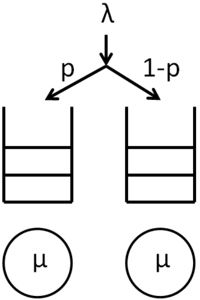

- Five parameters to choose: α_1, α_2, k_1, k_2, and p.
- Minimize $\mathbb{E}[N] + \beta\mathbb{E}[E]$

With $\mu = 2$, $\lambda = 1.9$, $\gamma = 0.2$, $\beta = 22$

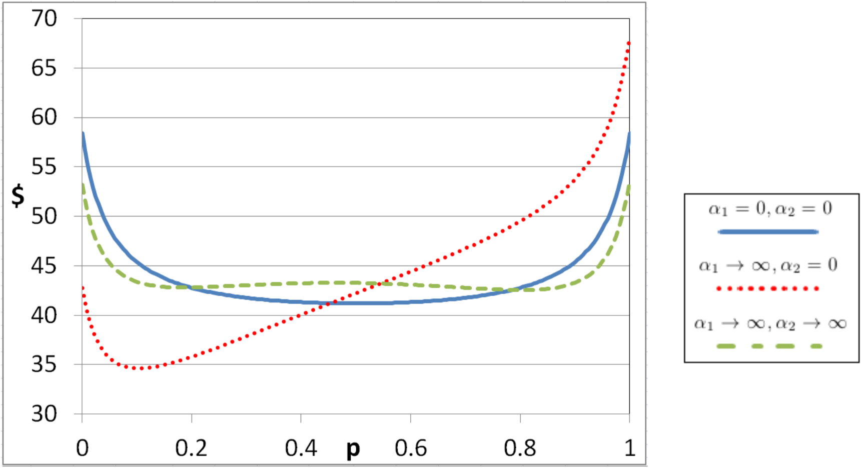

Here $k_1 = 5$ for the optimal configuration and $p = 0.105$.

With $\mu = 2$, $\lambda = 1$, $\gamma = 0.2$, $\beta = 15$

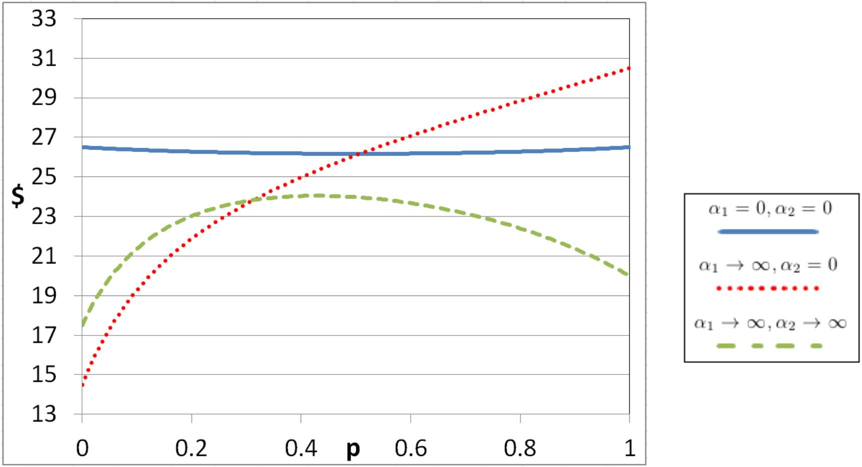

Conclusion

1. Complete analysis of single server systems.
2. In random routing scenario, heterogeneous servers may be desirable.

Future Work

1. Add more options and configurations to the single server model (sleep states, speed scaling etc.)
2. Look at the multi-server problem (what properties does the optimal policy have in this setting?)

[1] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. Kozuch. "Optimality Analysis of Energy-Performance Trade-off for Server Farm Management." *Performance Evaluation*, vol. 67, 2010, pp. 1155-1171.

[2] A. Gandhi and M. Harchol-Balter. "M/M/k with Exponential Setup." *CMU Technical Report CMU-CS-09-166*, 2010.

[3] A. Gandhi, M. Harchol-Balter, and I. Adan. "Server Farms with Setup Costs." *Performance Evaluation*, vol. 67, no. 11, 2010, pp. 1123-1138.

[4] A. Penttinen, E. Hyytia, and S. Aalto. "Energy-Aware Dispatching in Parallel Queues with On-Off Energy Consumption." *IEEE International Performance Computing and Communications Conference*, 2011.

[5] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wand, and N. Gautam. "Managing Server Energy and Operational Costs in Hosting Centers." *ACM SIGMETRICS*, 2005, pp. 303-314.

[6] A. Wierman, L. L. H. Andrew, and M. Lin. "Speed Scaling: An Algorithmic Perspective." *Chapter in Handbook on Energy-Aware and Green Computing*, CRC Press, 2012, pp. 385-406.

[7] J. Slegers, N Thomas, and I. Mitrani. "Dynamic Server Allocation for Power and Performance." *Performance Evaluation: Metrics, Models and Benchmarks Lecture Notes in Computer Science*, vol. 5119, 2008, pp. 247-261.

[8] A. Gandhi, S. Doroudi, M. Harchol-Balter and A. Scheller-Wolf. "Exact Analysis of the M/M/k/setup Class of Markov Chains via Recursive Renewal Reward." *ACM SIGMETRICS*, 2013.

Questions?

Thank you.