Optimal Sleep Scheduling for a Wireless Sensor Network Node

David Shuman and Mingyan Liu
University of Michigan

40th Annual Asilomar Conference on Signals, Systems, and Computers

October 31, 2006

Introduction

- Wireless sensor networks have recently been utilized in an expanding array of applications

- Energy conservation is a key design issue

- Wide range of solutions proposed
  - Adjust routes and power rates over time
  - Aggregate data to reduce unnecessary traffic
  - Turn nodes off and on periodically (duty-cycling)

- Algorithms utilize different techniques to selectively turn nodes on and off
  - Leverage geographic information provided by GPS (GAF)
  - Distributed algorithms featuring local coordination (Span)
  - Frequent probing of neighboring sensors to actively replace failed nodes without maintaining information about neighbors (PEAS)

Introduction (cont.)

- We also study periodic sleeping, but proceed in a different direction
  - Consider a broad class of sleep scheduling policies, and attempt to identify the optimal
  - Restrict attention to a single node
  - Focus solely on the tradeoffs between energy consumption and packet delay

- Related models
  - Vacation models
    - *A. Federgruen and K.C. So, "Optimality of threshold policies in single-server queueing systems with server vacations," Adv. Appl. Prob., vol. 23, no. 2, pp. 388-405, June 1991*
  - M. Sarkar and R. Cruz (UC San Diego)

Outline

- **Problem Description and Formulation**
- Infinite Horizon Average Expected Cost Problem
- Finite Horizon Expected Cost Problem
- Concluding Remarks

Problem Description
Overview of System Model

Single Node

- Consider a single node in a wireless sensor network
- Modeled as a single-server queue

Two Control Objectives

- Conserve energy through duty-cycling
  - While asleep, the node is unable to transmit packets, but packets continue to arrive at the node
- Minimize packet queuing delay

Key Modeling Assumptions

- Node sleeps for N time slots at a time
  - In place of additional costs or setup time for switching modes
  - Multiple vacations are allowed
- Bernoulli arrival process with success probability p
- Packets arriving in one slot cannot be transmitted until the following slot
- Only one packet transmission per slot, and successful *w.p.1*
- Node has an infinite buffer size

Finite and Infinite Horizon Problem Formulation
Information State, Action Space, and System Dynamics

Information State

- X_t : two-dimensional vector
  - B_t : current queue length
  - S_t : number of slots remaining until node awakes

Action Space

- Two control actions available when node is awake:
  - $U_t = 1$ (“Awake”)
  - $U_t = 0$ (“Sleep”)

System Dynamics

- Controlled Markov Chain model

$$X_{t+1} = f(X_t, U_t, A_t) = \begin{bmatrix} B_{t+1} \\ S_{t+1} \end{bmatrix} = \begin{cases} \begin{bmatrix} B_t + A_t \\ S_t - 1 \end{bmatrix}, & \text{if } S_t > 0 \\ \begin{bmatrix} B_t + A_t \\ N - 1 \end{bmatrix}, & \text{if } S_t = 0 \text{ and } U_t = 0 \\ \begin{bmatrix} [B_t - 1]^+ + A_t \\ 0 \end{bmatrix}, & \text{if } S_t = 0 \text{ and } U_t = 1 \end{cases}$$

Finite and Infinite Horizon Problem Formulation
Cost Structure and Optimization Criteria

Cost Structure

- Constant, positive cost *D* incurred at each time slot the node is awake
- Constant, positive cost *c* incurred at each time slot, by each backlogged packet

Problem (P1)

- Infinite horizon average expected cost problem
- Optimization criterion:

$$J^{\pi} := \limsup_{T \to \infty} \frac{1}{T} \cdot E^{\pi}\left\{ \sum_{t=0}^{T-1} D \cdot U_t + \sum_{t=1}^{T} c \cdot B_t \,\middle|\, \mathsf{F}_0 \right\}$$

Problem (P2)

- Finite horizon expected cost problem
- Optimization criterion:

$$J_T^{\pi} := E^{\pi}\left\{ \sum_{t=0}^{T-1} D \cdot U_t + \sum_{t=1}^{T} c \cdot B_t \,\middle|\, \mathsf{F}_0 \right\}$$

Optimization Space

- In both problems, the minimization is over the space of all randomized and deterministic history-dependent control laws

Outline

- Problem Description and Formulation
- **Infinite Horizon Average Expected Cost Problem**
- Finite Horizon Expected Cost Problem
- Concluding Remarks

Infinite Horizon Average Expected Cost Optimization

Optimal Stationary Policy Exists

- Problem (**P1**) satisfies the (**BOR**) assumptions of Sennott's Theorem 7.5.6, guaranteeing the existence of an optimal stationary Markov policy[1]

When Queue Is Non-Empty

- Optimal policy is to stay awake and serve
  - Eventually, node must serve to avoid infinite average cost
  - Proof via interchange argument utilizes this fact and linear holding cost structure

When Queue is Empty

- Optimal control at boundary state $X = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ is given by the threshold decision rule:

$$\left(\frac{p}{1-p}\right)\cdot\left(\frac{N-1}{2}\right) \underset{Sleep\,(U_t^*=0)}{\overset{Awake\,(U_t^*=1)}{\gtrless}} \frac{D}{c} \qquad (*)$$

[1] See L.I. Sennott, *Stochastic Dynamic Programming and the Control of Queueing Systems*, John Wiley and Sons, 1999.

Outline

- Problem Description and Formulation
- Infinite Horizon Average Expected Cost Problem
- **Finite Horizon Expected Cost Problem**
- Concluding Remarks

Finite Horizon Expected Cost Optimization
Goal: Identify Optimal Markov Policy at Each State and Time Slot Pair

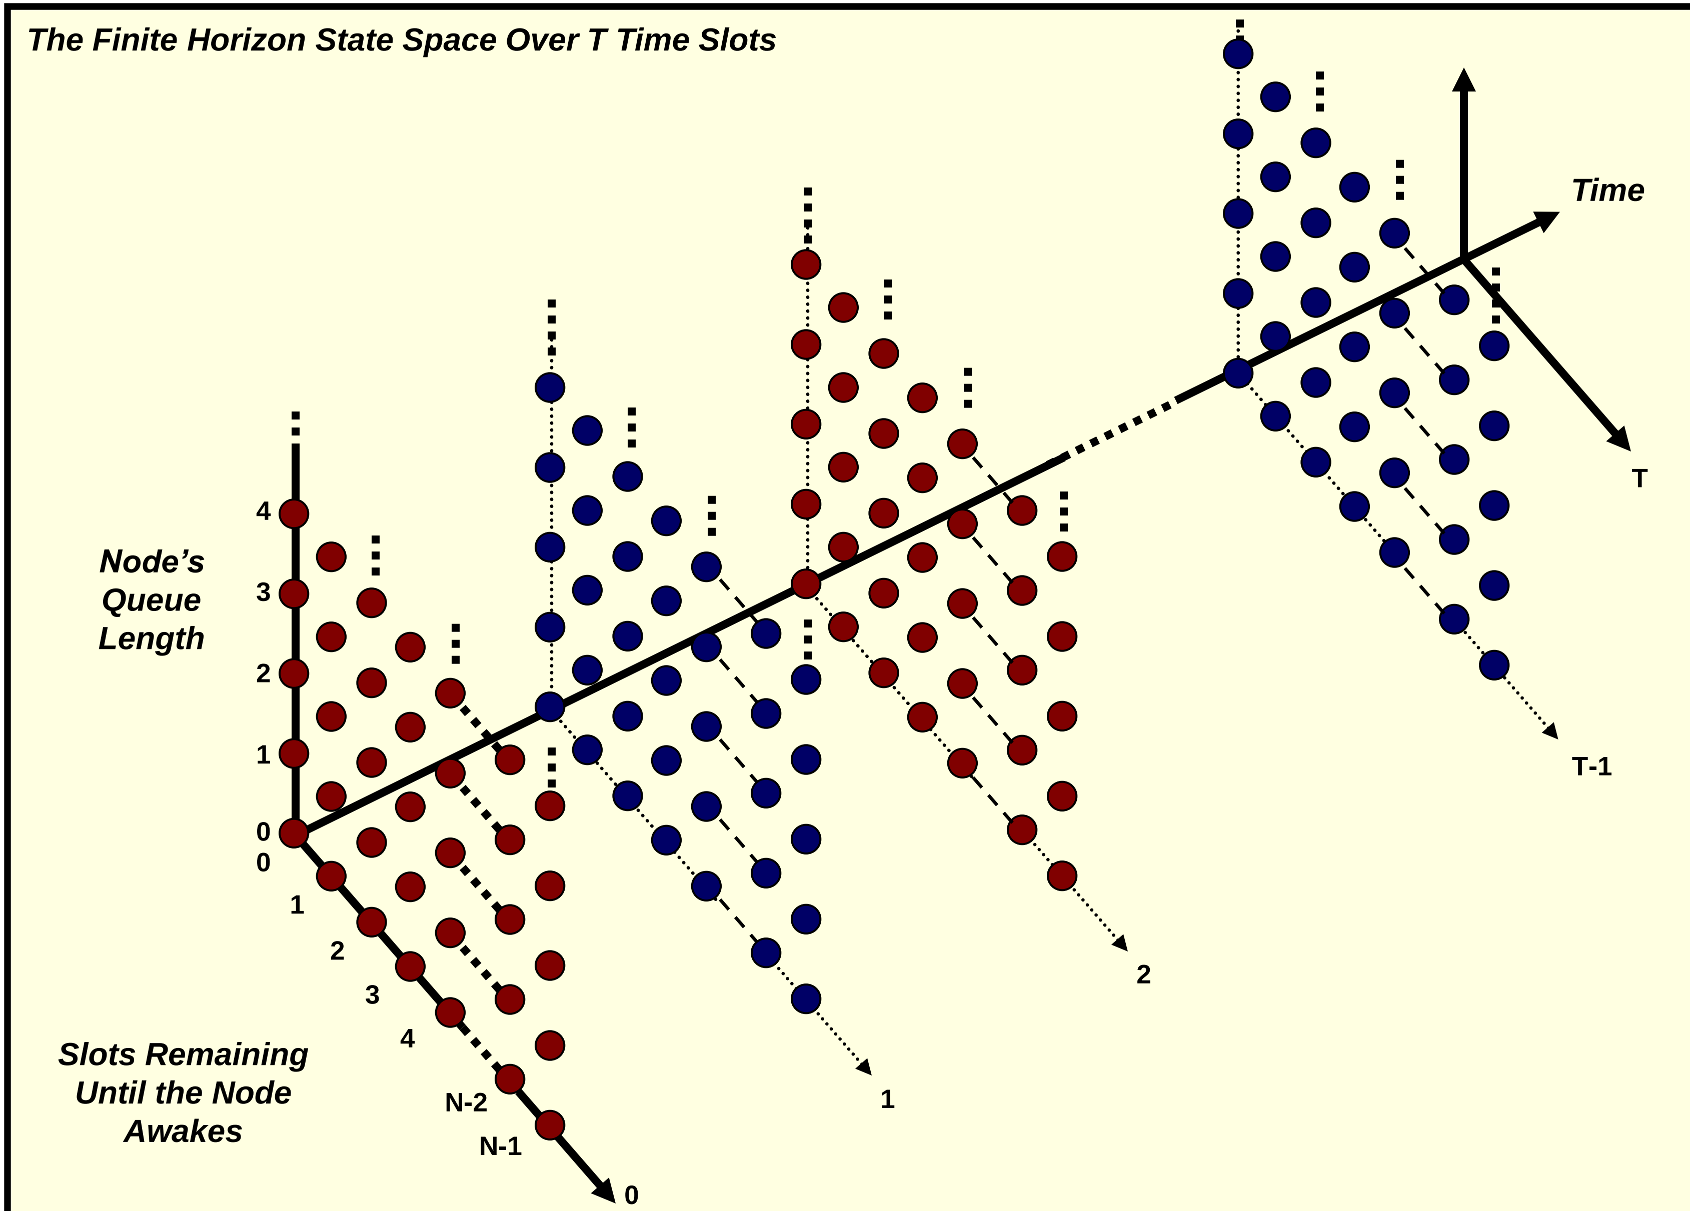

Finite Horizon Expected Cost Optimization
Goal: Identify Optimal Markov Policy at Each State and Time Slot Pair

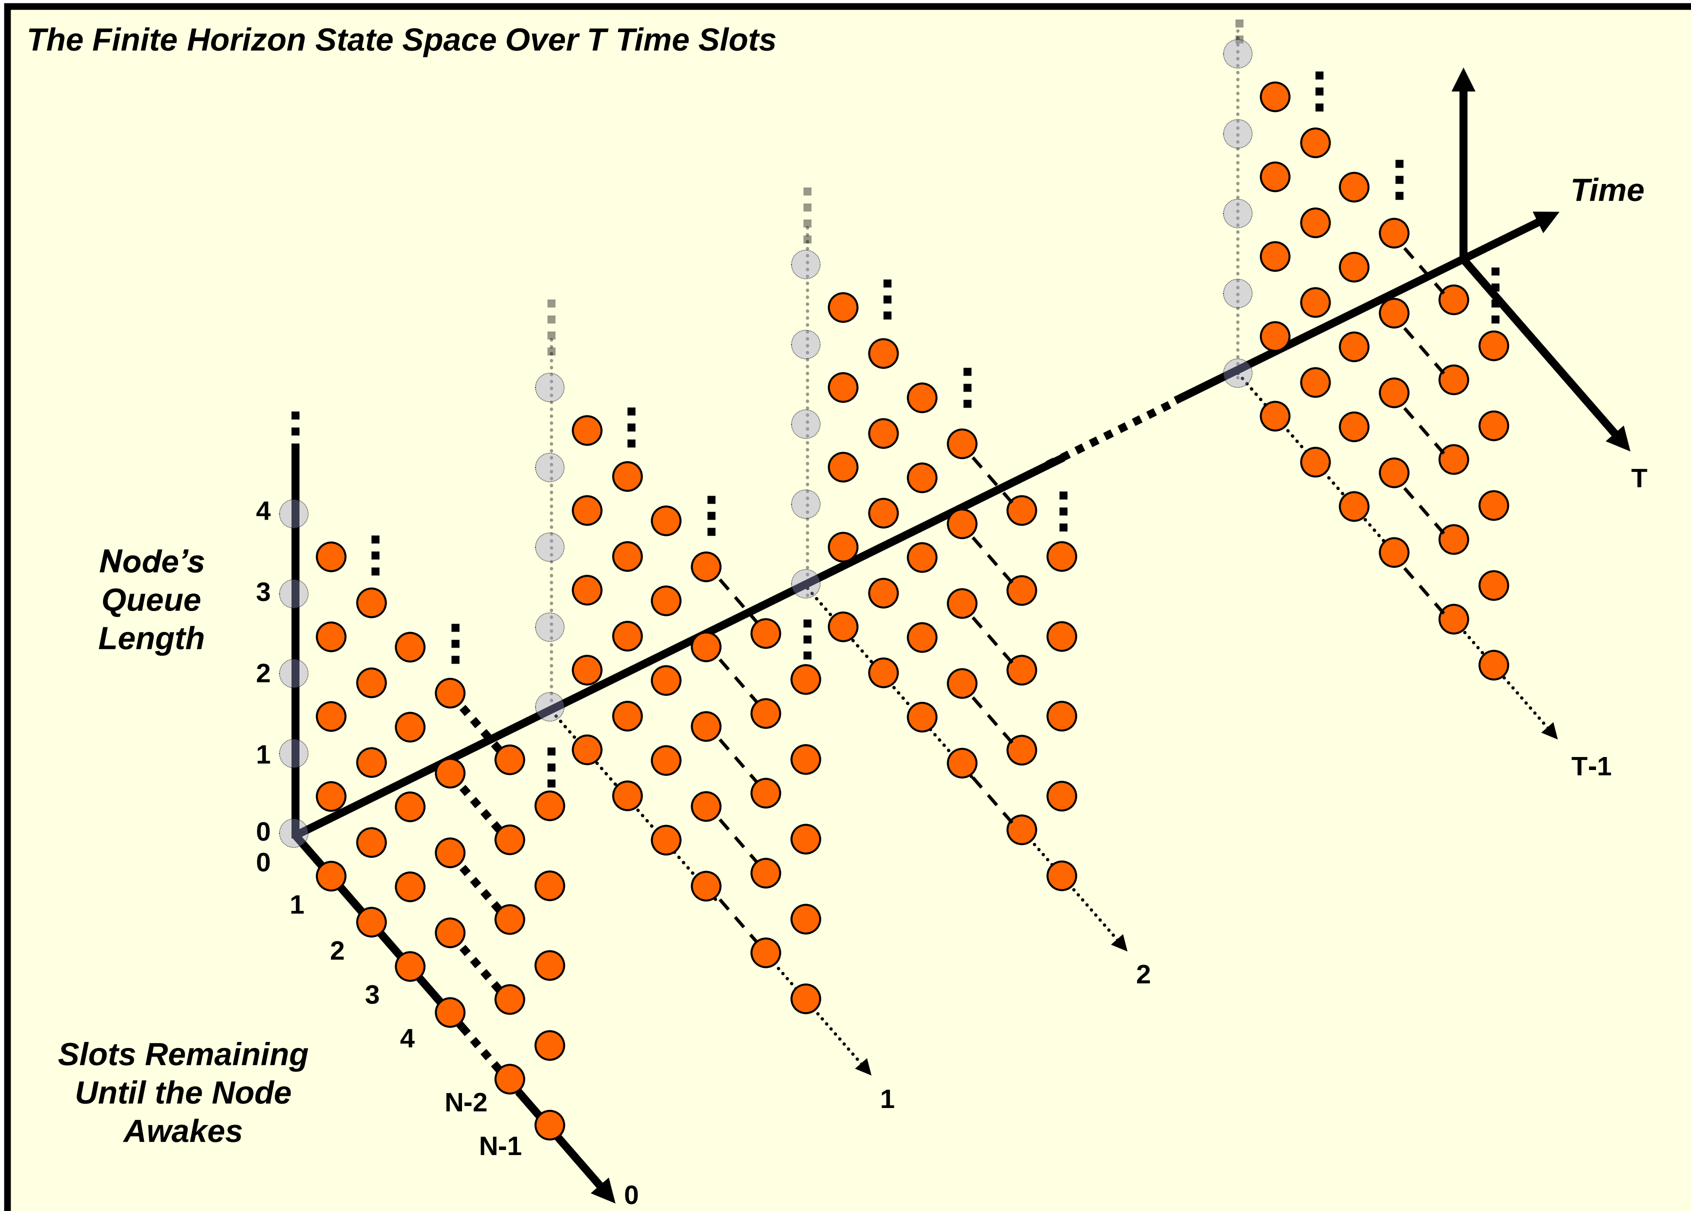

Finite Horizon Expected Cost Optimization
Optimal Policy at the End of the Time Horizon and When Queue is Non-Empty

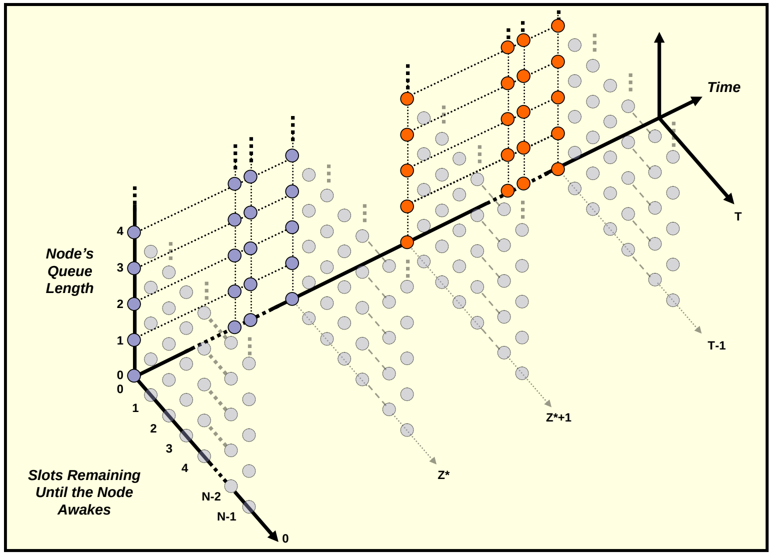

Node Awake at the End of the Time Horizon

- When $T - \frac{D}{c} \le t < T$, the optimal control is to sleep
- Basic idea is that marginal benefit of serving is at most $c \cdot \left\lfloor \frac{D}{c} \right\rfloor \le D$, the marginal cost of serving
- Proof by backwards induction
- For notation purposes, we define $z^* := \left\lfloor T - \frac{D}{c} \right\rfloor$

Finite Horizon Expected Cost Optimization
Optimal Policy at the End of the Time Horizon and When Queue is Non-Empty

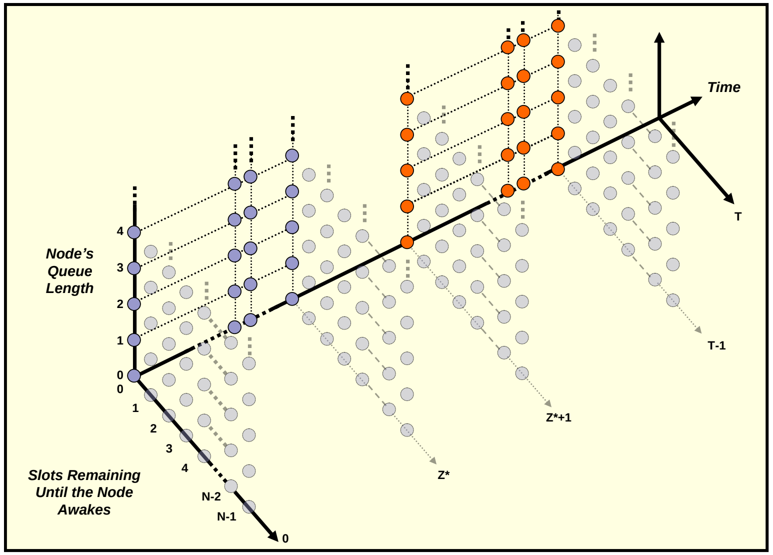

Node Awake at the End of the Time Horizon

- When $T-\frac{D}{c}\leq t<T$, the optimal control is to sleep
- Basic idea is that marginal benefit of serving is at most $c\cdot\left\lfloor\frac{D}{c}\right\rfloor\leq D$, the marginal cost of serving
- Proof by backwards induction
- For notation purposes, we define $z^{*}:=\left\lfloor T-\frac{D}{c}\right\rfloor$

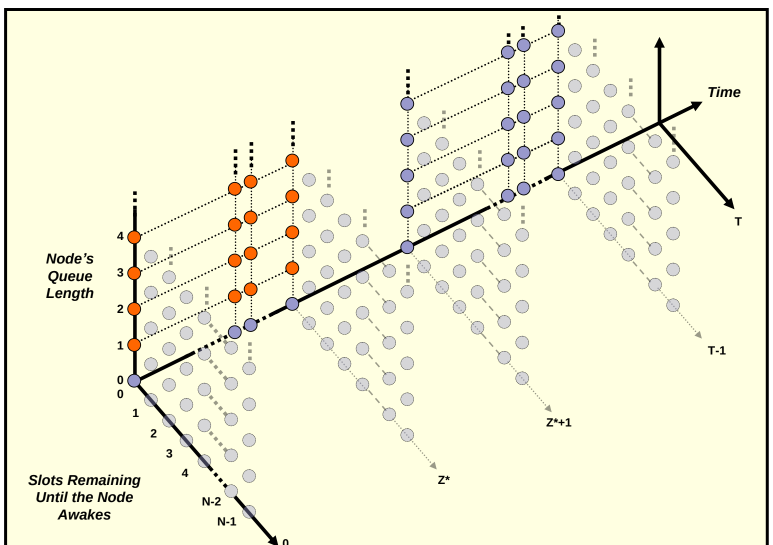

Node Awake Before End and Queue Non-Empty

- Optimal policy is to stay awake and serve
- Proof follows from similar interchange argument as the infinite horizon problem

Finite Horizon Expected Cost Optimization
Optimal Policy at the Boundary State, Before the End of the Time Horizon

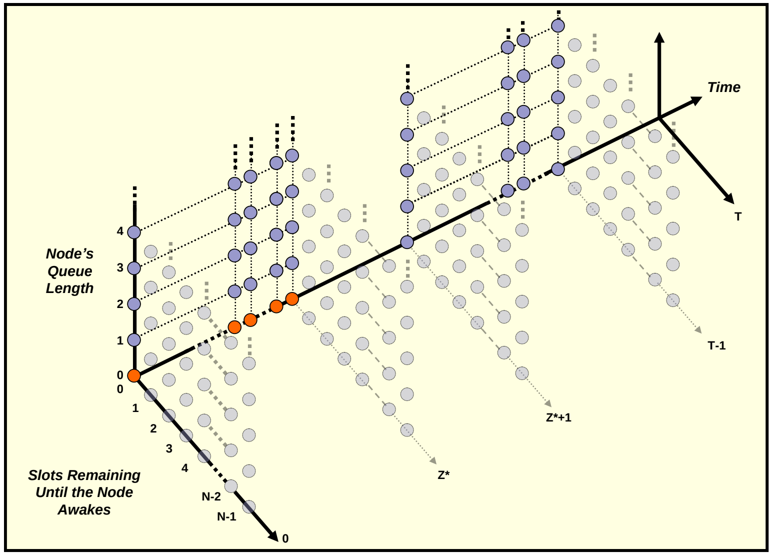

Finite Horizon Expected Cost Optimization
Optimal Policy at the Boundary State, Before the End of the Time Horizon

t = z*

- The optimal control at $X_{z^*} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ is to sleep

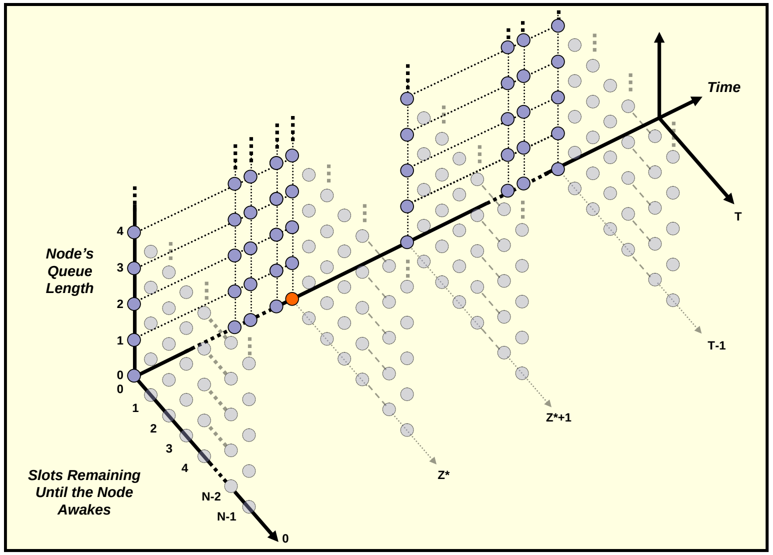

Finite Horizon Expected Cost Optimization
Optimal Policy at the Boundary State, Before the End of the Time Horizon

$t = z^*$

- The optimal control at $X_{z^*} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ is to sleep

$z^*-N < t < z^*$

- If $z^*-N < t < z^*$ and $X_t = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, the optimal control at slot t to minimize J_t^{π} is given by the threshold decision rule:

$$c \cdot \sum_{j=1}^{z^*-t} \left\{ p^j \left(T - t - j\right) \right\} - D \cdot \sum_{j=0}^{z^*-t} p^j \underset{Sleep\,(U_t^*=0)}{\overset{Awake\,(U_t^*=1)}{\gtrless}} 0$$

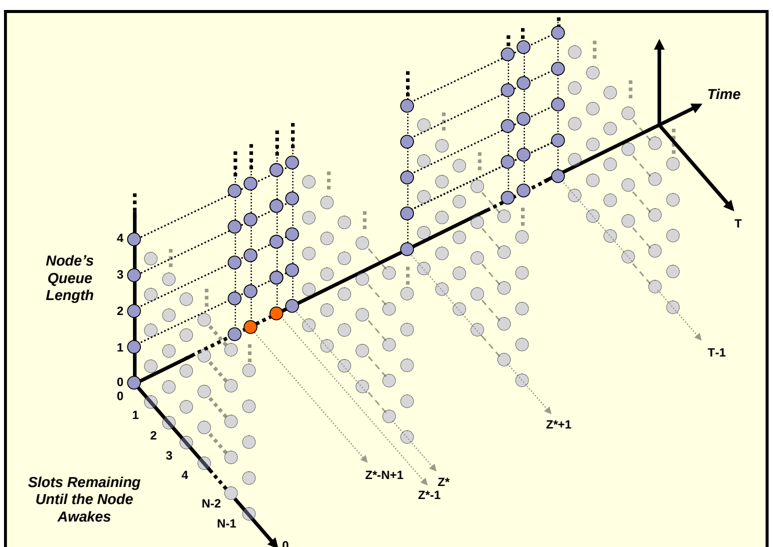

Finite Horizon Expected Cost Optimization
Optimal Policy at the Boundary State, Before the End of the Time Horizon

t = z*

- The optimal control at $X_{z^*} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ is to sleep

z*-N < t < z*

- If *z*-N* < *t* < *z** and $X_t = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, the optimal control at slot *t* to minimize J_t^{π} is given by the threshold decision rule:

$$c \cdot \sum_{j=1}^{z^*-t} \left\{ p^j \left(T - t - j \right) \right\} - D \cdot \sum_{j=0}^{z^*-t} p^j \; \underset{Sleep\,(U_t^*=0)}{\overset{Awake\,(U_t^*=1)}{\gtrless}} \; 0$$

Implication

- The optimal control when the node is awake and the queue is empty is non-increasing over time, from *z*-N+1* until the end of the time horizon

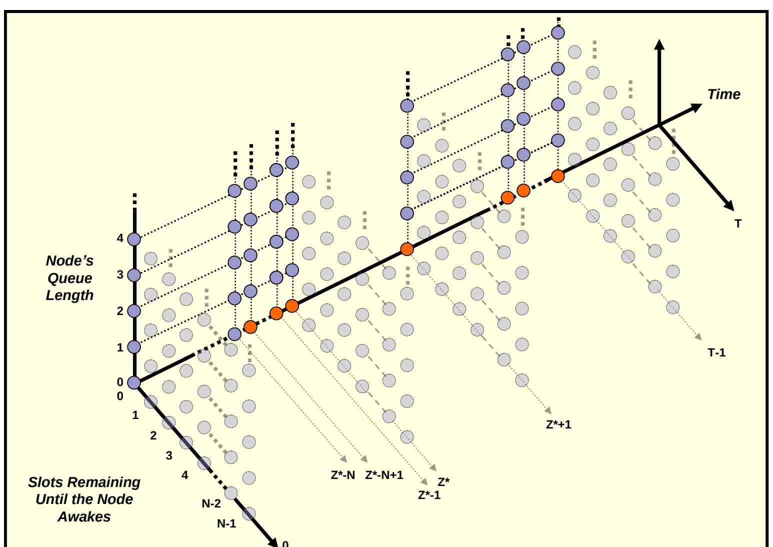

Finite Horizon Expected Cost Optimization
Optimal Policy at the Boundary State, Before the End of the Time Horizon

t = z*

- The optimal control at $X_{z^*} = \begin{bmatrix}0\\0\end{bmatrix}$ is to sleep

z*-N < t < z*

- If *z*-N* < *t* < *z** and $X_t = \begin{bmatrix}0\\0\end{bmatrix}$, the optimal control at slot *t* to minimize J_t^{π} is given by the threshold decision rule:

$$c\cdot\sum_{j=1}^{z^*-t}\left\{p^j\left(T-t-j\right)\right\}-D\cdot\sum_{j=0}^{z^*-t}p^j \underset{Sleep\,(U_t^*=0)}{\overset{Awake\,(U_t^*=1)}{\gtrless}} 0$$

Implication

- The optimal control when the node is awake and the queue is empty is non-increasing over time, from *z*-N+1* until the end of the time horizon

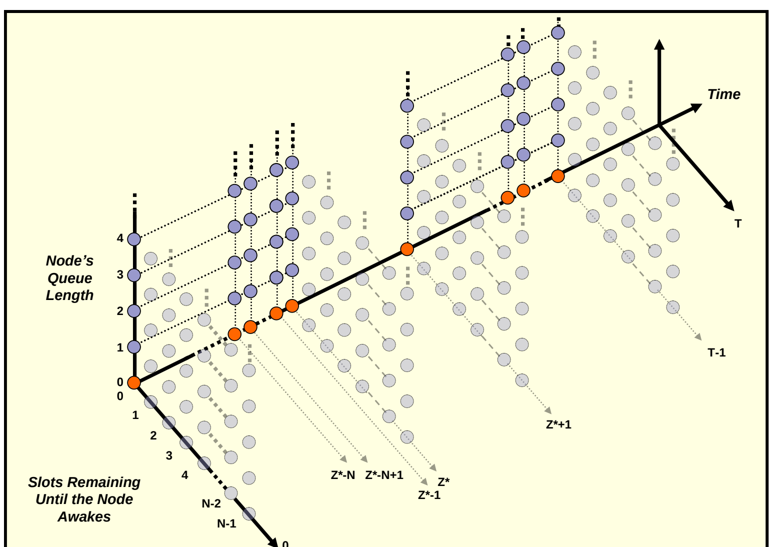

<u>Question</u>:

Is the optimal policy at the boundary state necessarily monotonic over the entire time horizon?

Finite Horizon Expected Cost Optimization
The Optimal Policy at the Boundary State Is Not Necessarily Monotonic in Time

Answer

- No, as the following counterexample demonstrates

Optimal Control at $X_t = [0,0]^T$ When $T = 15$, $N = 3$, $c = 10$, $D = 21$, and $p = 2/3$

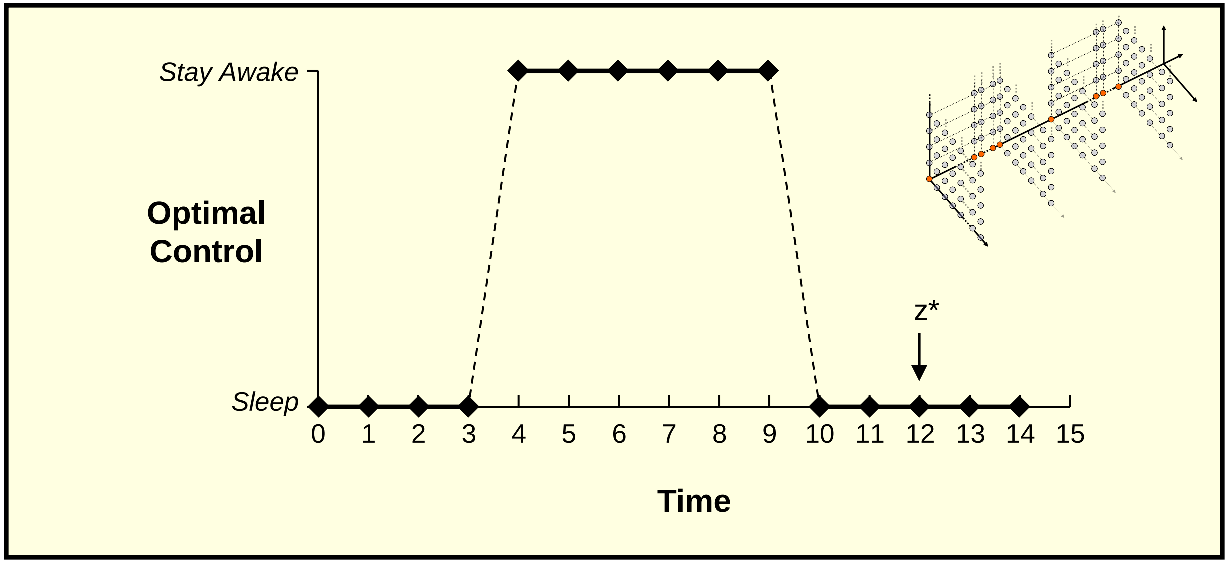

More Questions

- Can we find sufficient conditions to guarantee the optimal policy at the boundary state is non-increasing over the entire time horizon
- What behavior is possible in the optimal control at the boundary state when such conditions are not met?

Finite Horizon Expected Cost Optimization Conjectures

Conjecture 1

- If the parameters of Problem (**P2**) satisfy the following condition:

$$\left(\frac{p}{1-p}\right)\cdot\left(\frac{N-1}{2}\right)\geq\frac{D}{c} \qquad (SC),$$

the optimal policy when the node is awake and the queue is empty is non-increasing in time for the entire time horizon

Conjecture 2

- At most one jump

Implication

- Only three possible structural forms of the optimal policy at the boundary:

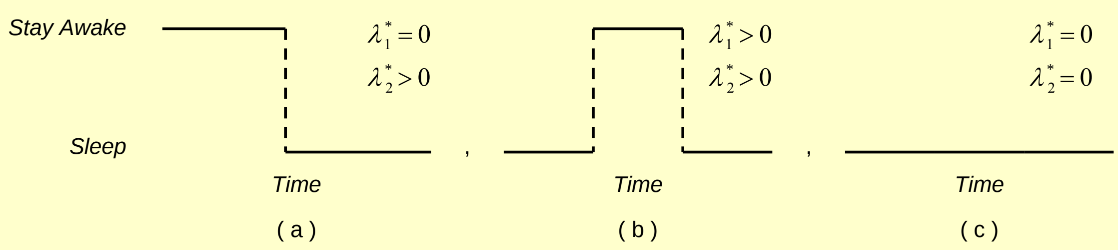

Finite Horizon Expected Cost Optimization
Observations on Numerical Results

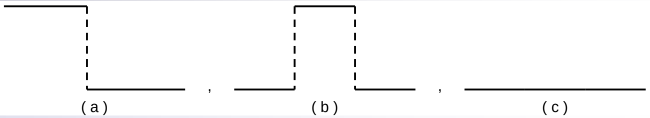

Observation 1

- If the time horizon is sufficiently long, then the optimal control is of the form (a) if (**SC**) holds, but of the form (b) or (c) if (**SC**) does not hold
  - Sufficient condition (SC) is identical to ($*$) from the infinite horizon problem

Observation 2

- The three possible structural forms lie on a spectrum in a sense

Why (b)?

- Underlying tradeoff at the boundary state is between extra backlog costs from sleeping, and energy costs incurred during unutilized slots

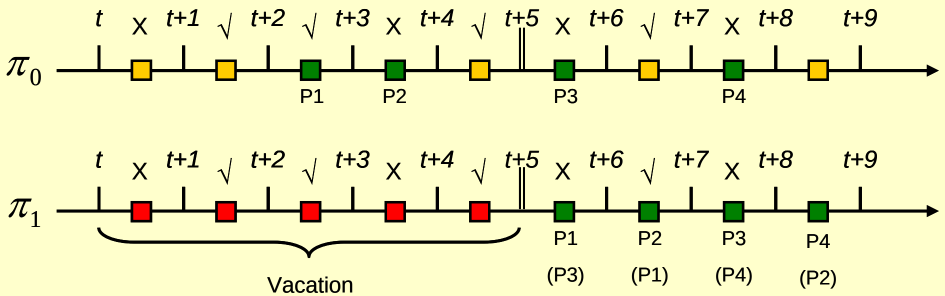

Summary and Future Work

- Infinite horizon average expected cost problem
  - Demonstrated existence of optimal stationary Markov policy
  - Completely characterized optimal control
- Finite horizon expected cost problem
  - Characterized optimal control away from the boundary
  - Posed two conjectures concerning structure of optimal control at boundary
- Possible extensions
  - Formulate as constrained optimization problem instead of assigning energy costs
  - Extend to multiple nodes