Optimized eeeBond: Energy Efficiency with non-Proportional Router Network Interfaces

Niklas Carlsson

Linköping University, Sweden

LINKÖPING UNIVERSITY

Background and motivation

- Networks provisioned for peak loads
  - Typically very low utilization
  - Opportunity to save energy
- Energy proportionality desirable
  - Energy usage proportional to system utilization
  - However, hardware limitations prevents this in practice
- Two promising approaches
  - EEE: “On/off toggling” of interface
  - eBond: Switch between redundant heterogeneous interfaces

Contributions

1. **Present eeeBond**
   – Hybrid of EEE and eBond

EEE +

Contributions

1. **Present eeeBond**
   - Hybrid of EEE and eBond

EEE +

	"Always on"	"On/off toggling"
Single interface	Naïve/default	EEE
Multi interface	eBond	eeeBond

Contributions

1. **Present eeeBond**
   – Hybrid of EEE and eBond

EEE +

2. **Unified analytic model**
   – Derive closed-form optimized parameter settings for eBond and eeeBond

	"Always on"	"On/off toggling"
Single interface	Naïve/default	EEE
Multi interface	eBond	eeeBond

Contributions

1. **Present eeeBond**
   – Hybrid of EEE and eBond

EEE +

2. **Unified analytic model**
   – Derive closed-form optimized parameter settings for eBond and eeeBond
3. **Performance evaluation**
   – Characterize gains possible with optimized protocols

	"Always on"	"On/off toggling"
Single interface	Naïve/default	EEE
Multi interface	eBond	eeeBond

eBond overview

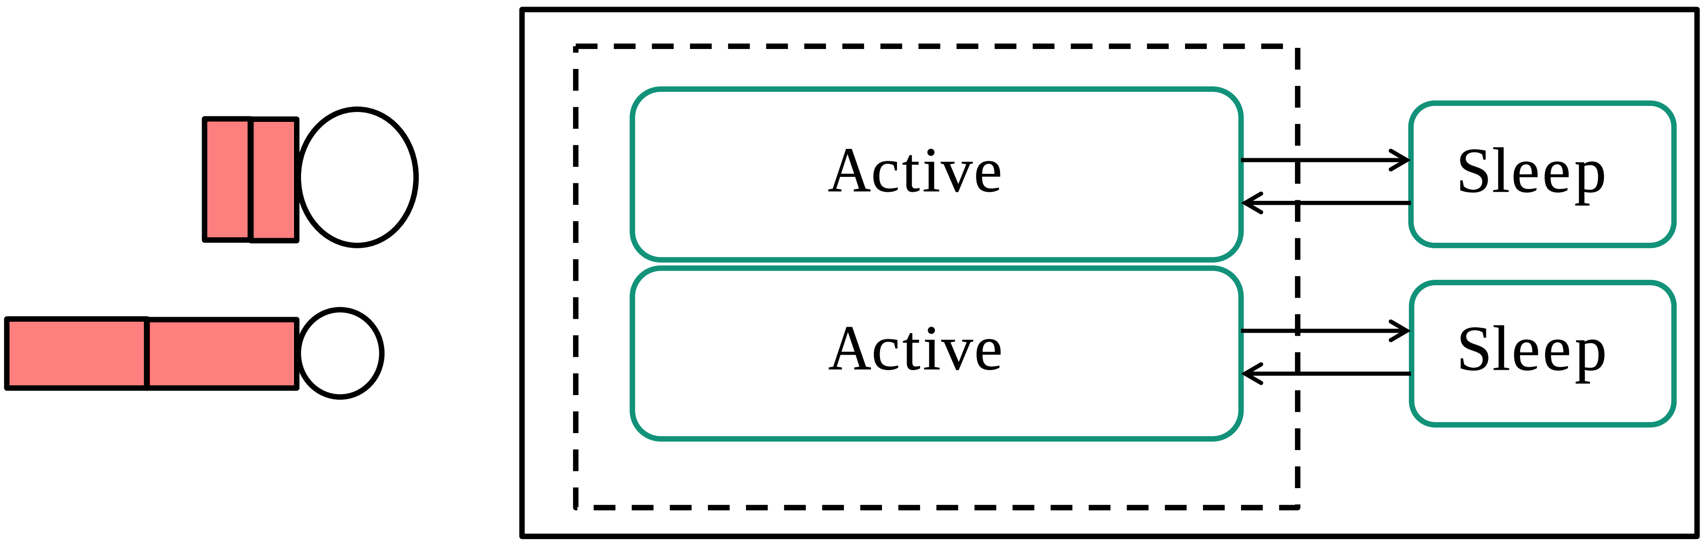

- Redundant heterogeneous interfaces
  - Toggle between interfaces based on load

eBond overview

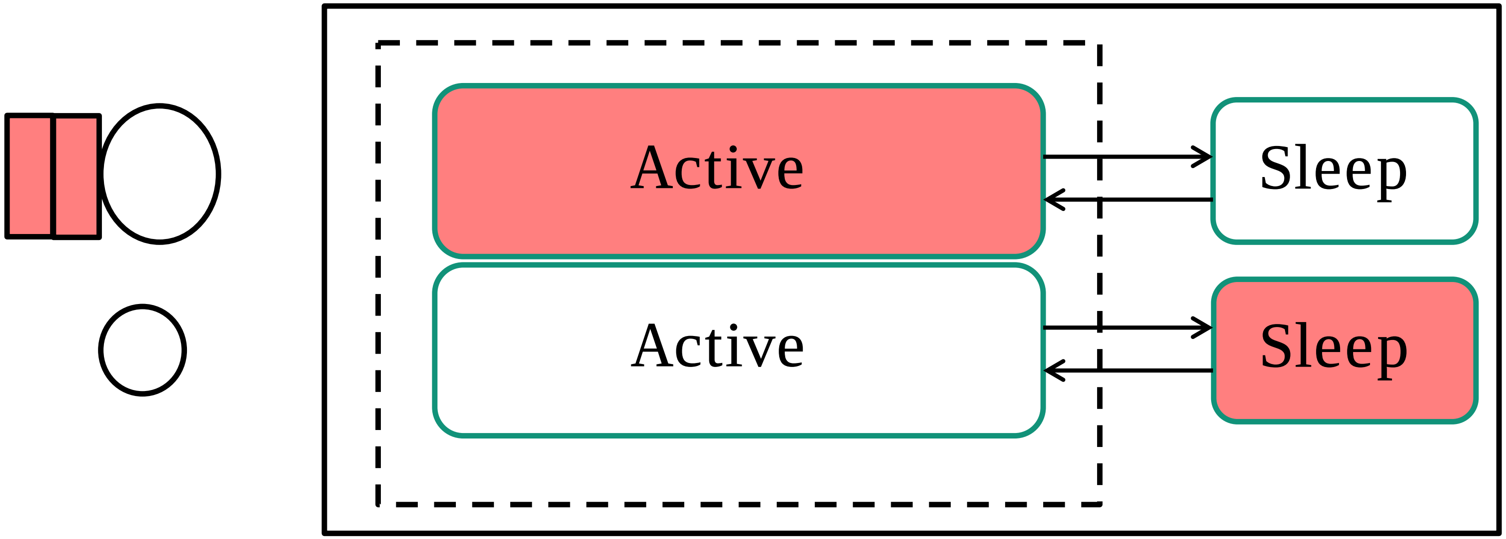

- Redundant heterogeneous interfaces
  - Toggle between interfaces based on load

eBond overview

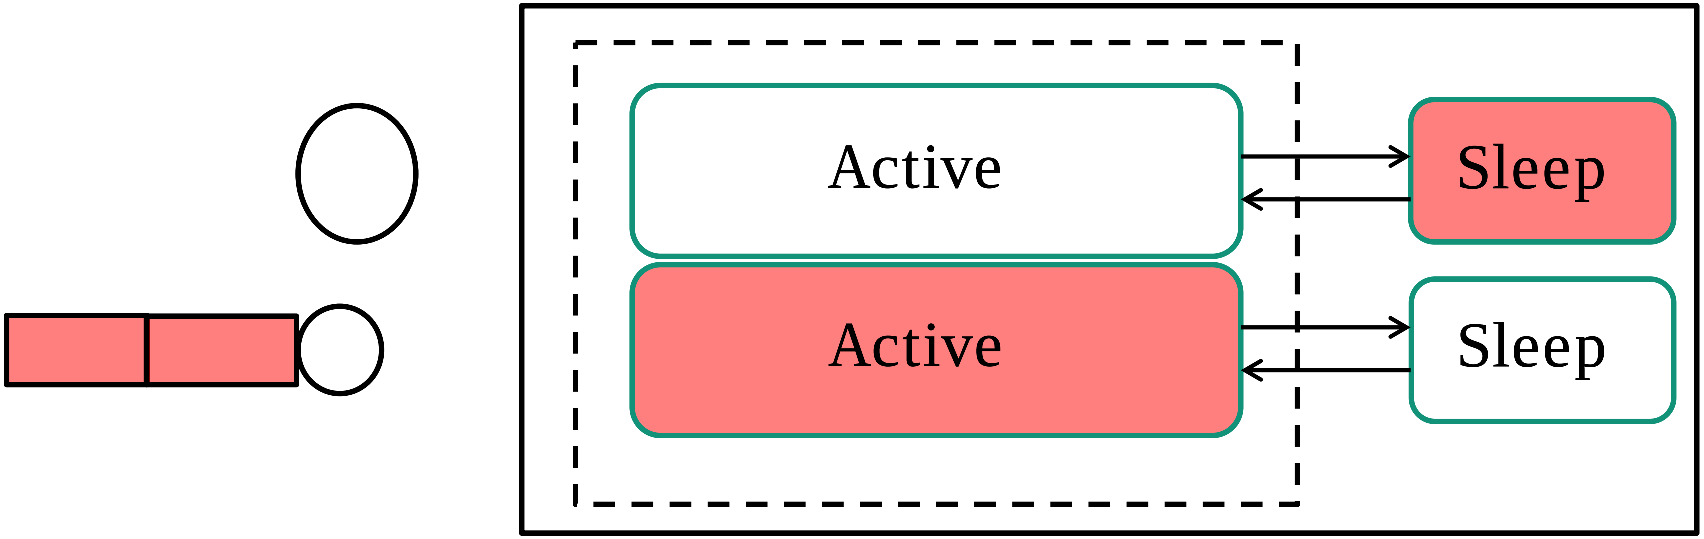

- Redundant heterogeneous interfaces
  - Toggle between interfaces based on load

eBond overview

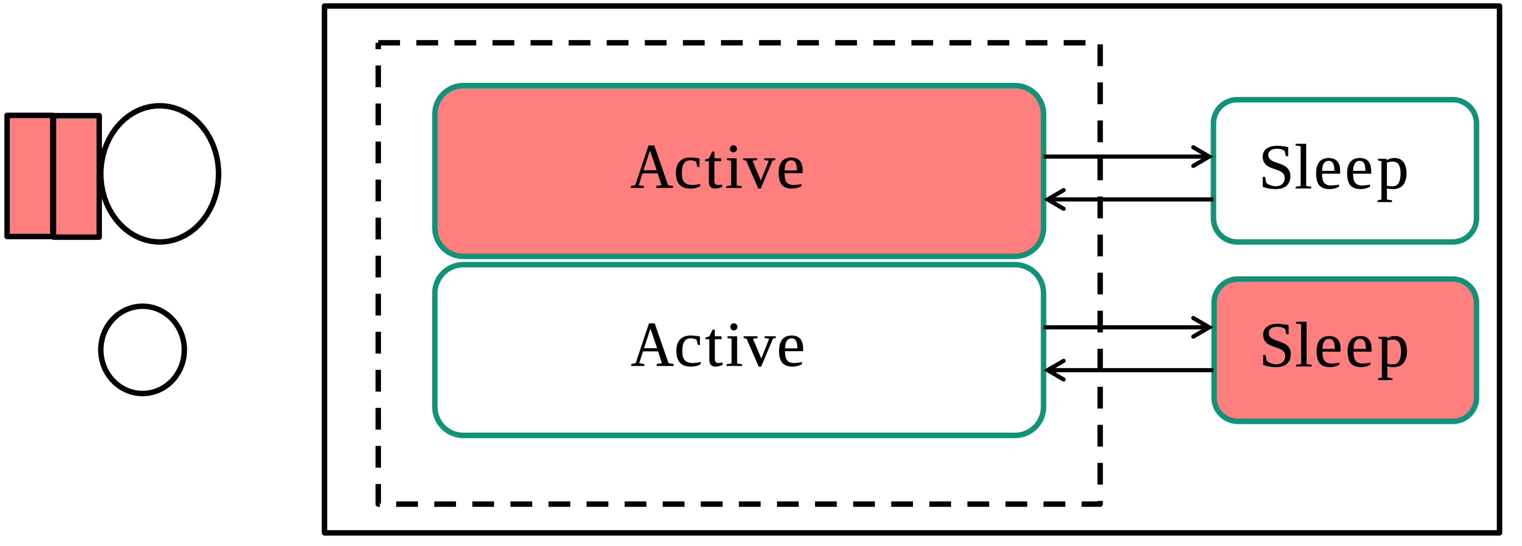

- Redundant heterogeneous interfaces
  - Toggle between interfaces based on load

EEE overview

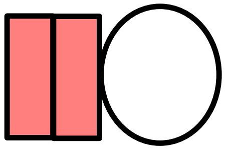

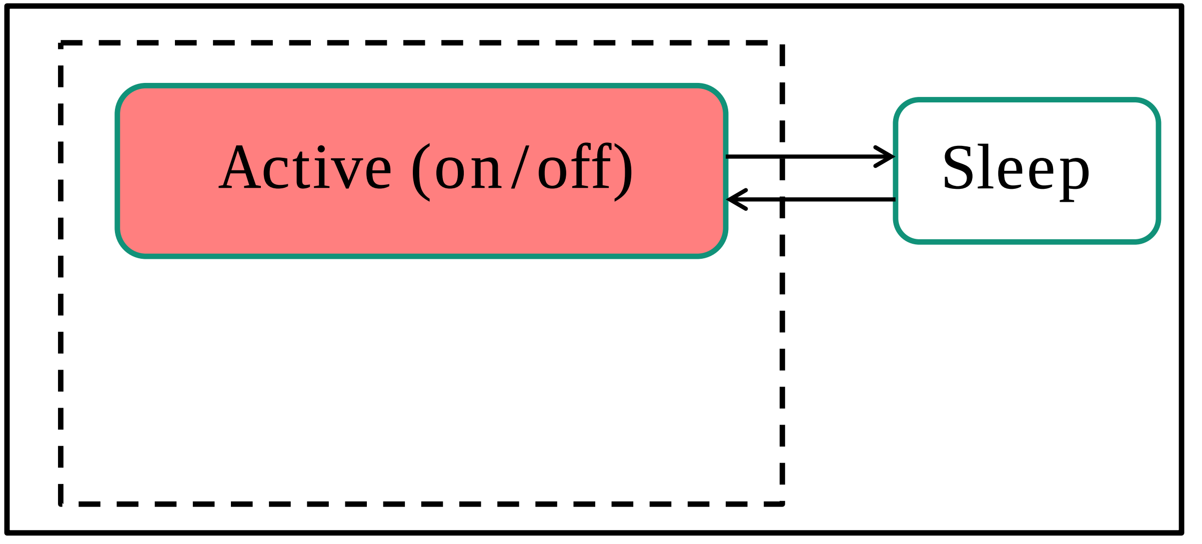

- Single interface
  - "On/off toggling" between active and (light) sleep state

EEE overview

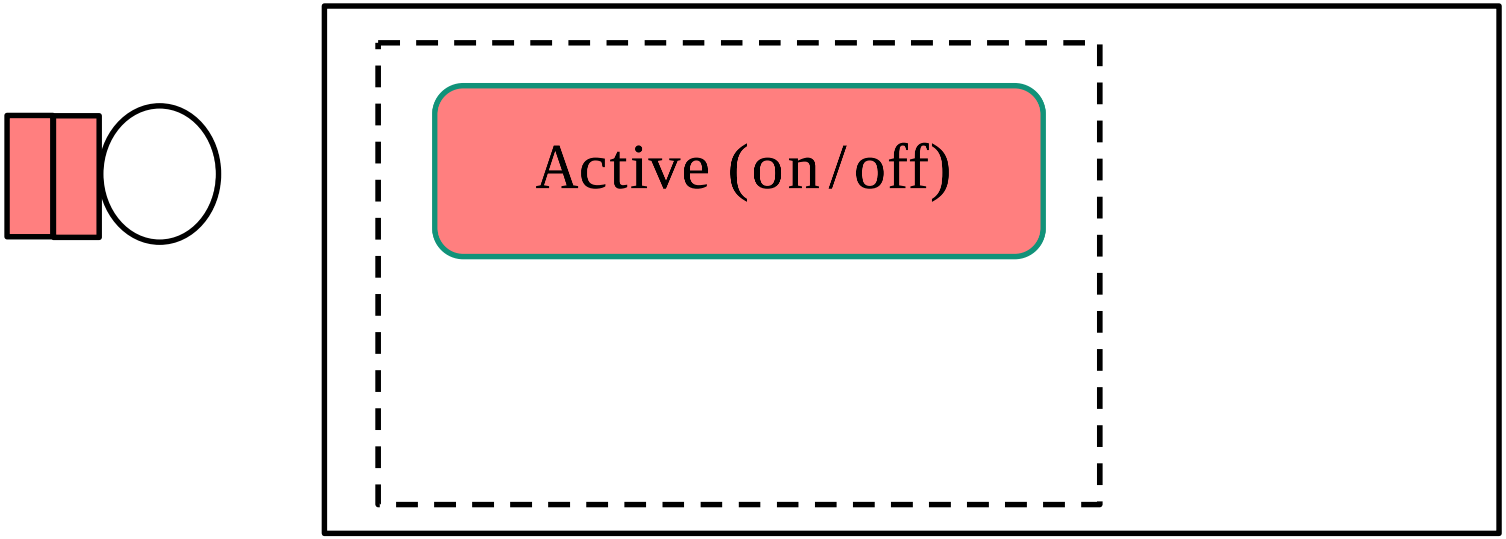

- Single interface
  - "On/off toggling" between active and (light) sleep state

EEE overview

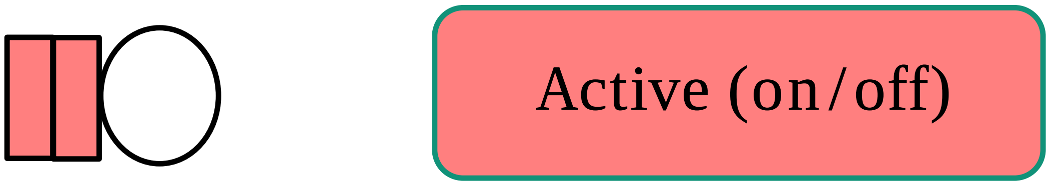

- Single interface
  - “On/off toggling” between active and (light) sleep state

EEE overview

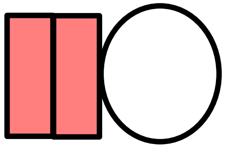

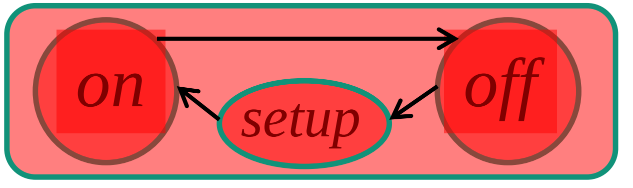

- Single interface
  - “On/off toggling” between active and (light) sleep state

EEE overview

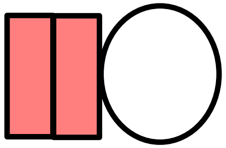

- Single interface
  - "On/off toggling" between active and (light) sleep state

EEE overview

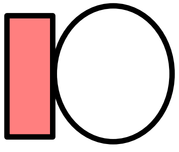

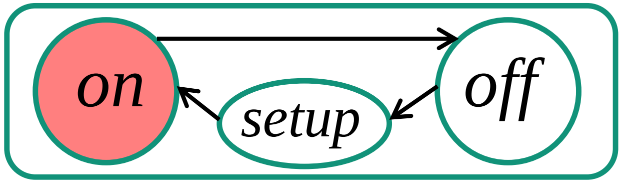

- Single interface
  - "On/off toggling" between active and (light) sleep state

EEE overview

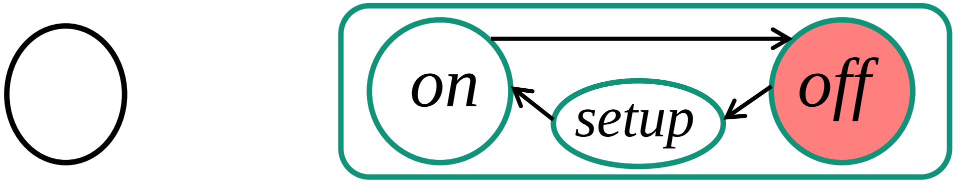

- Single interface
  - "On/off toggling" between active and (light) sleep state

EEE overview

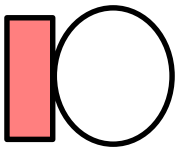

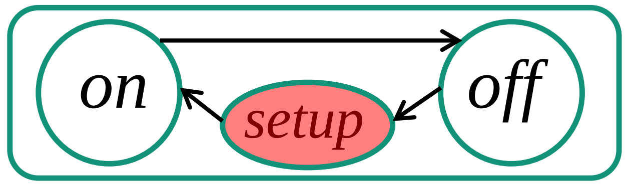

- Single interface
  - “On/off toggling” between active and (light) sleep state

EEE overview

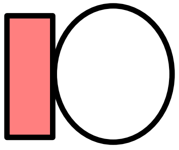

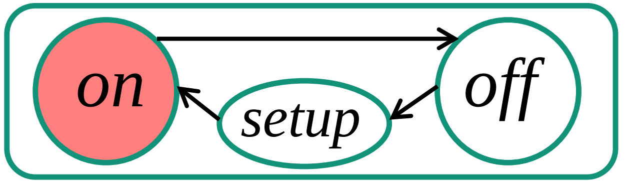

- Single interface
  - “On/off toggling” between active and (light) sleep state

EEE overview

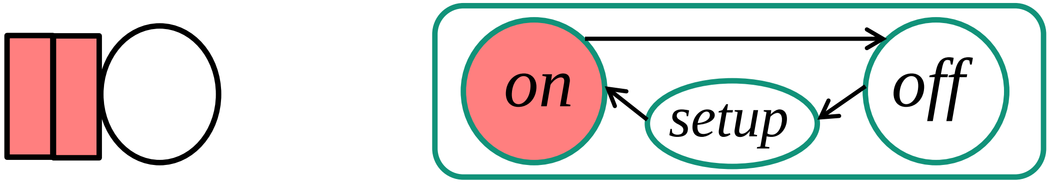

- Single interface
  - "On/off toggling" between active and (light) sleep state

Router model and power states

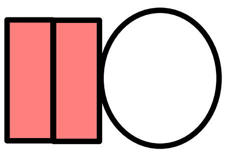

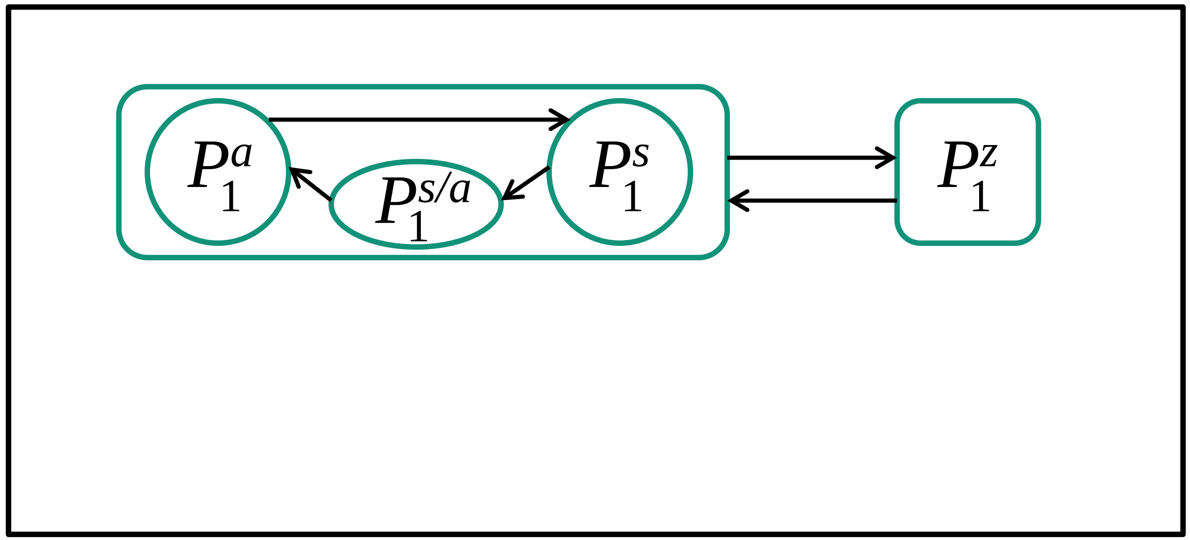

- Single interface
  - “On/off toggling” between active and (light) sleep state

Router model and power states

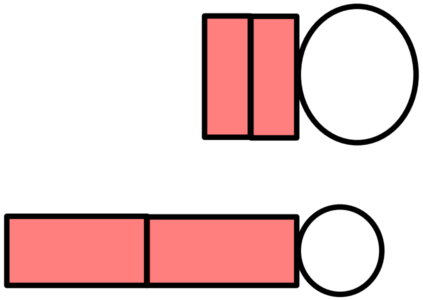

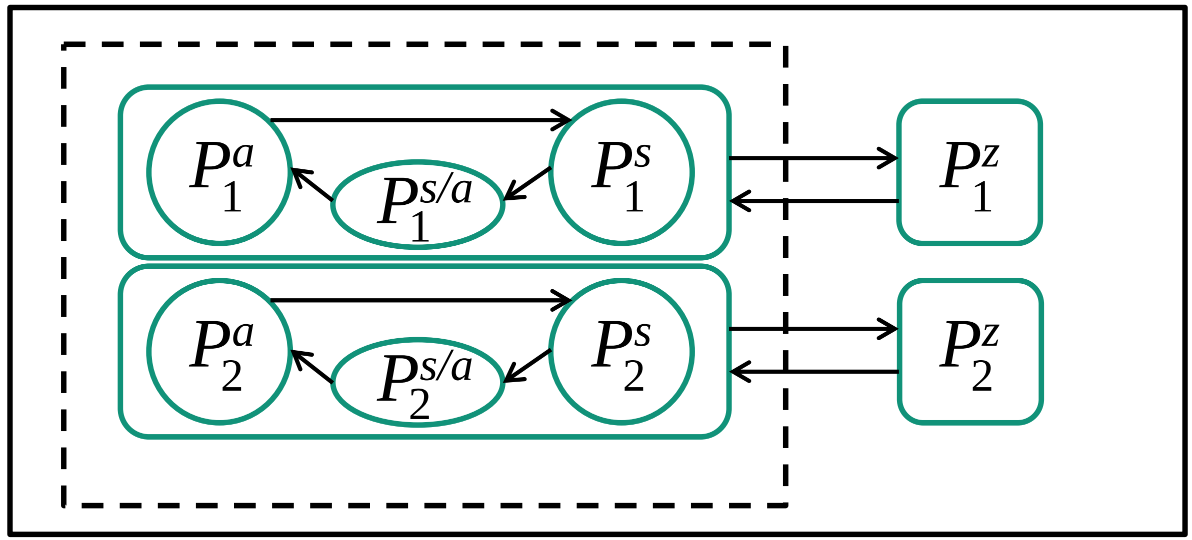

Router model and power states

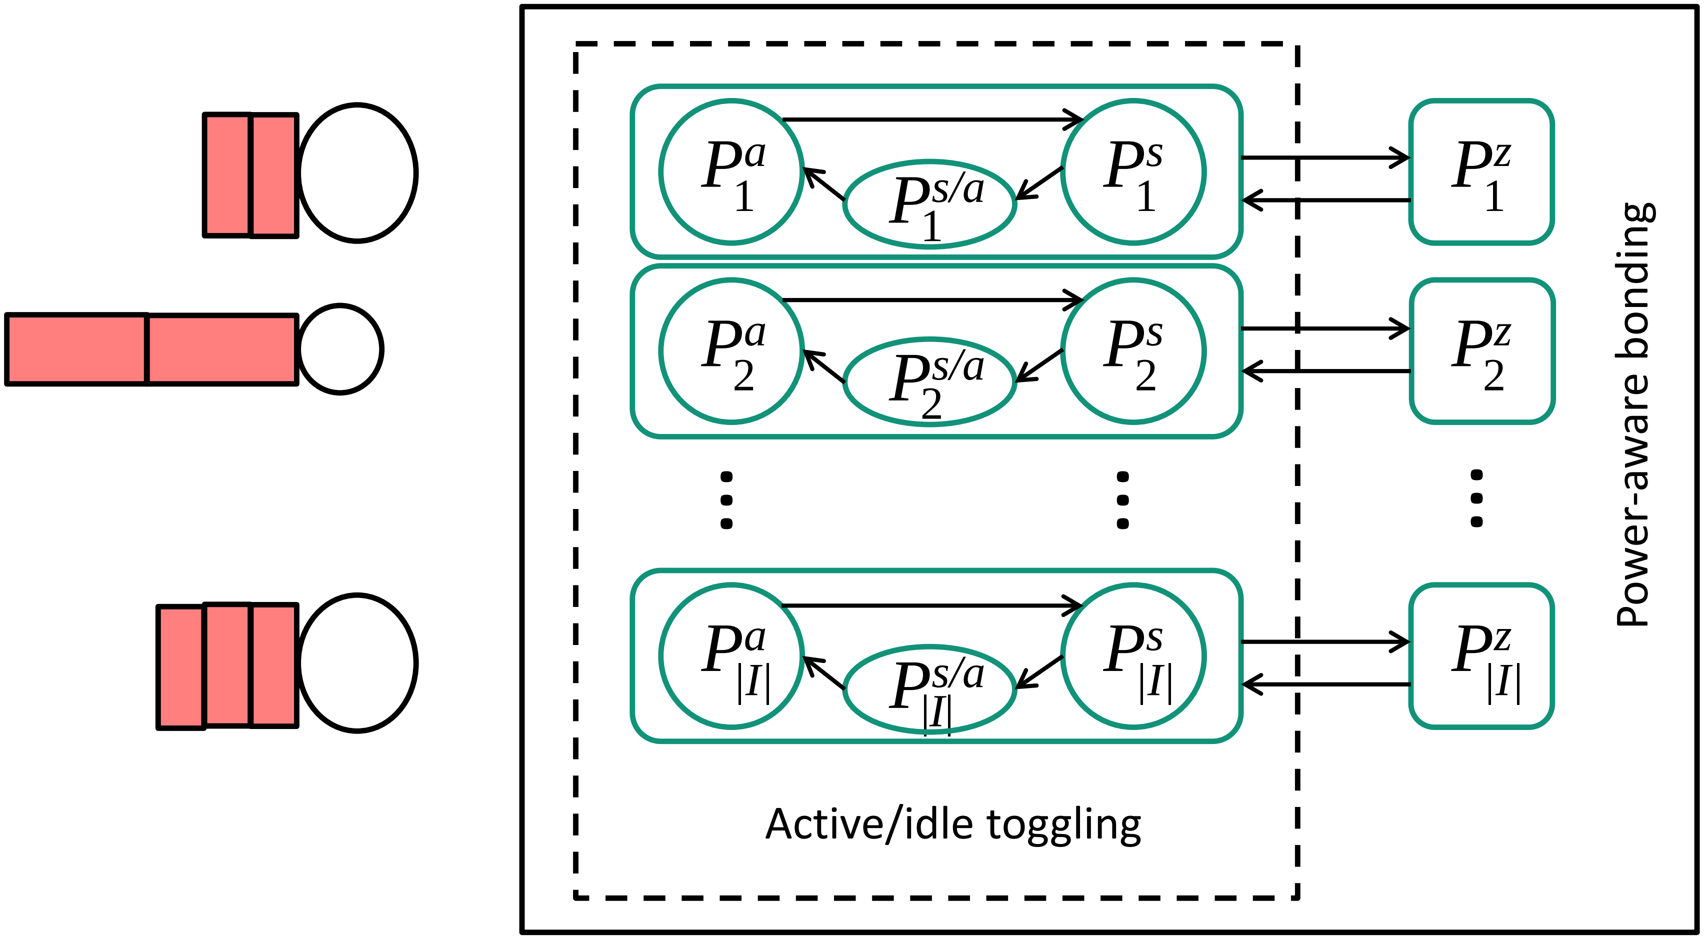

eeeBond overview

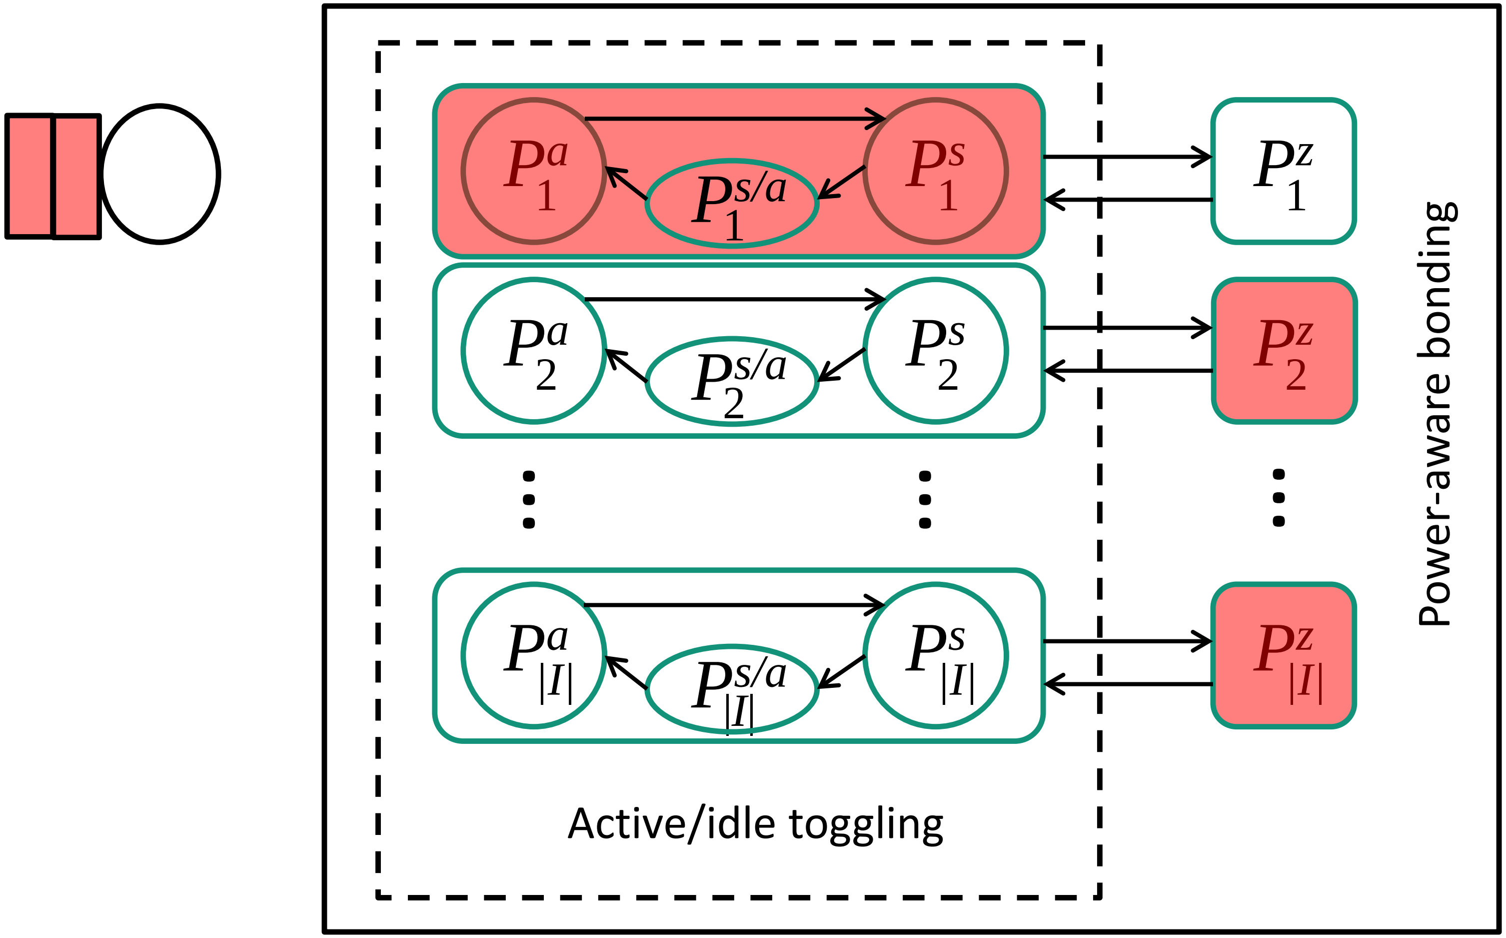

eeeBond overview

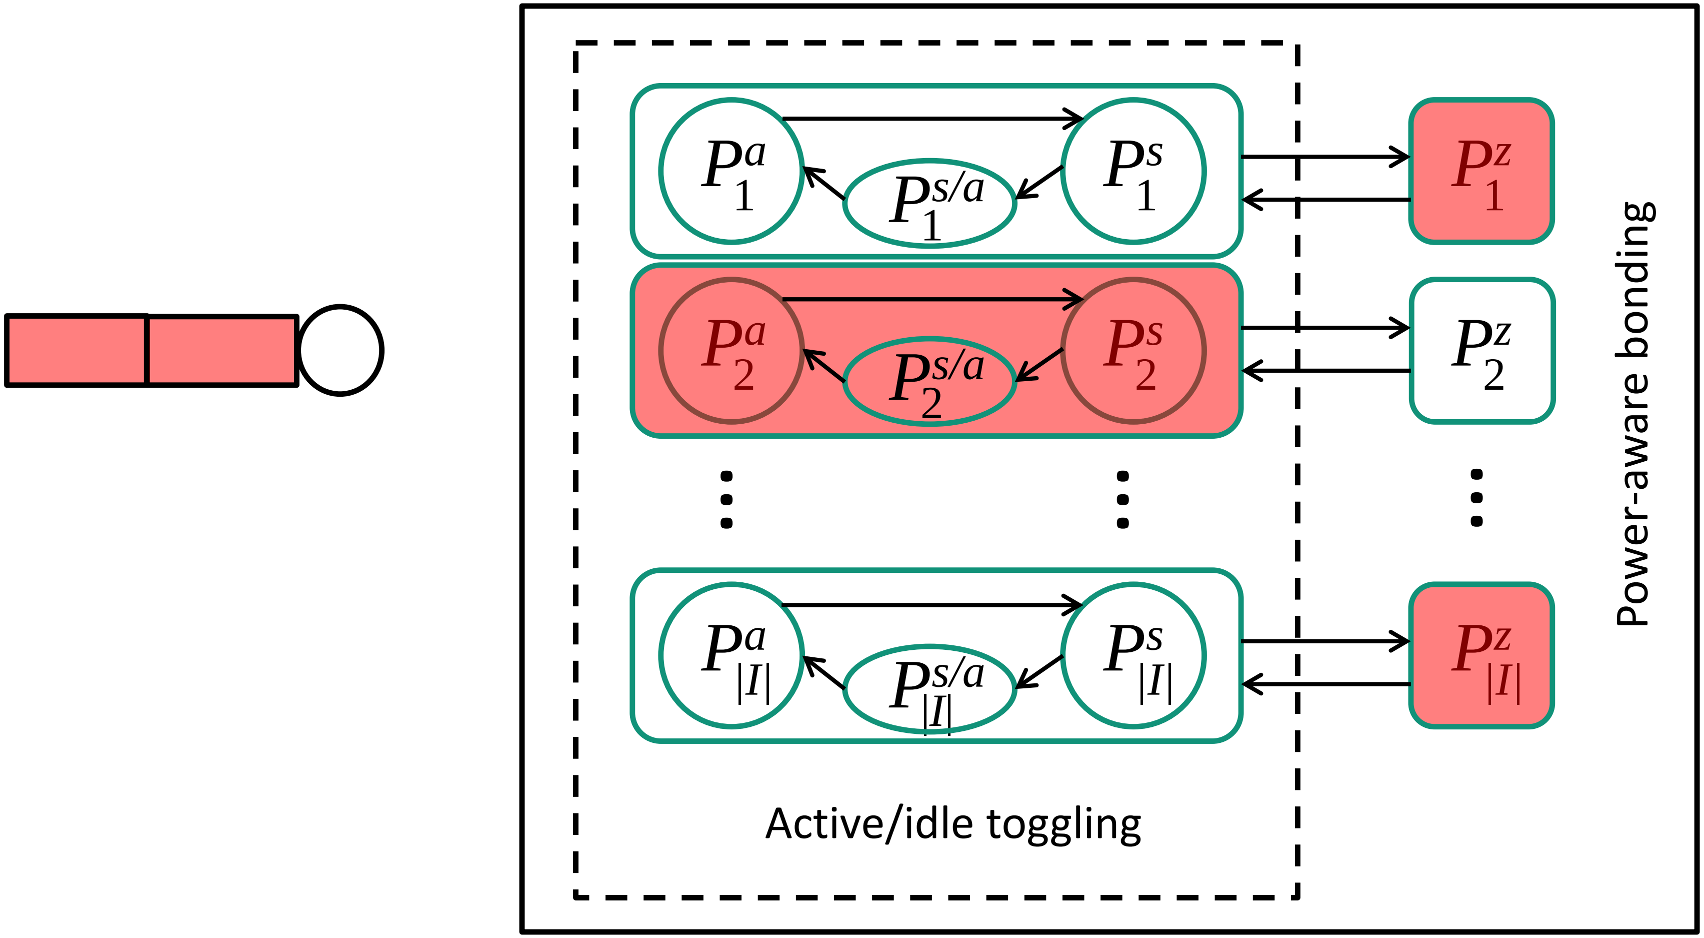

eeeBond overview

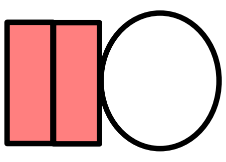

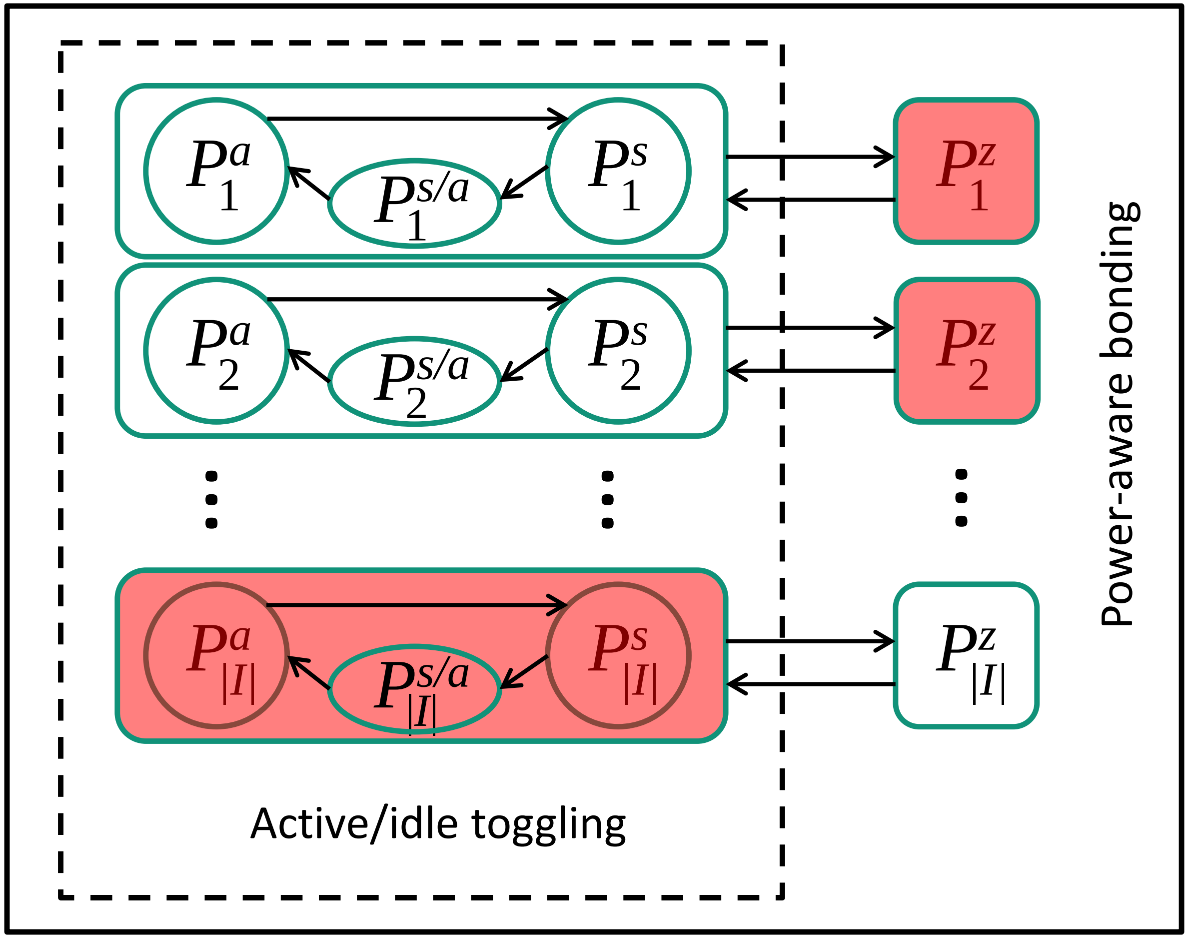

eeeBond overview

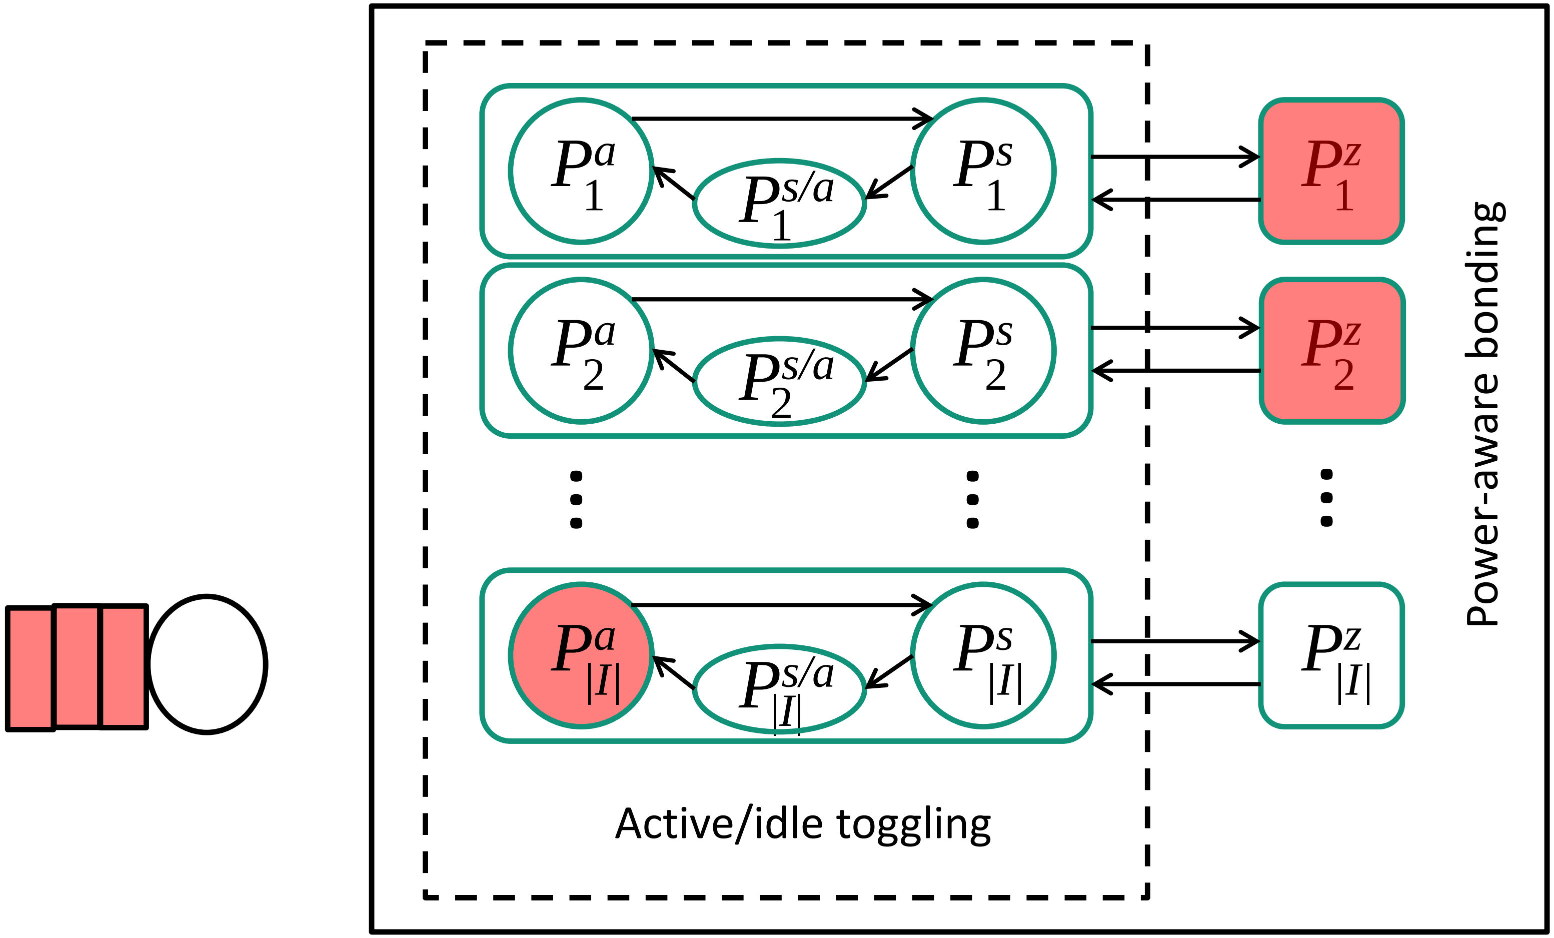

eeeBond overview

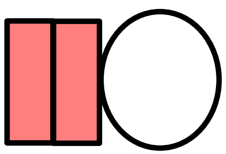

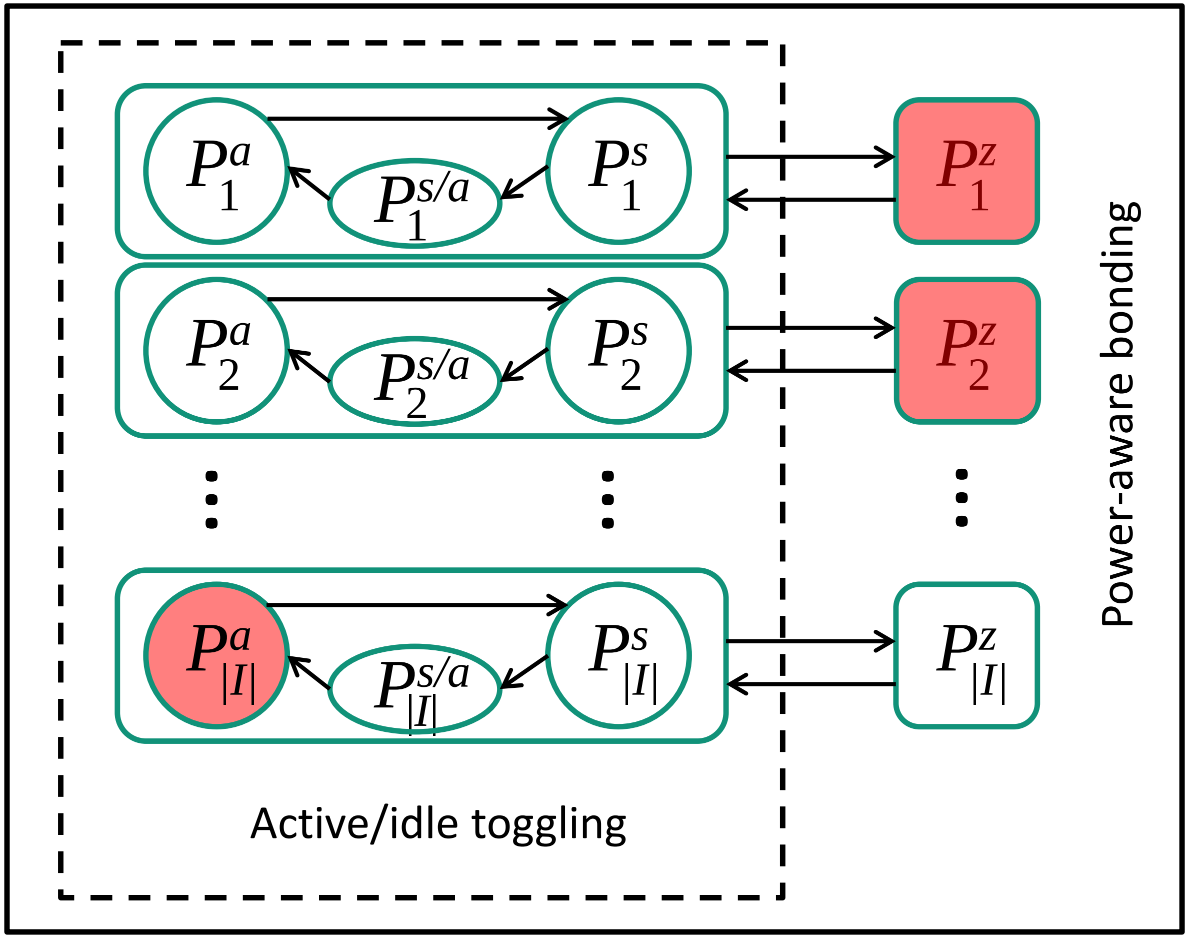

eeeBond overview

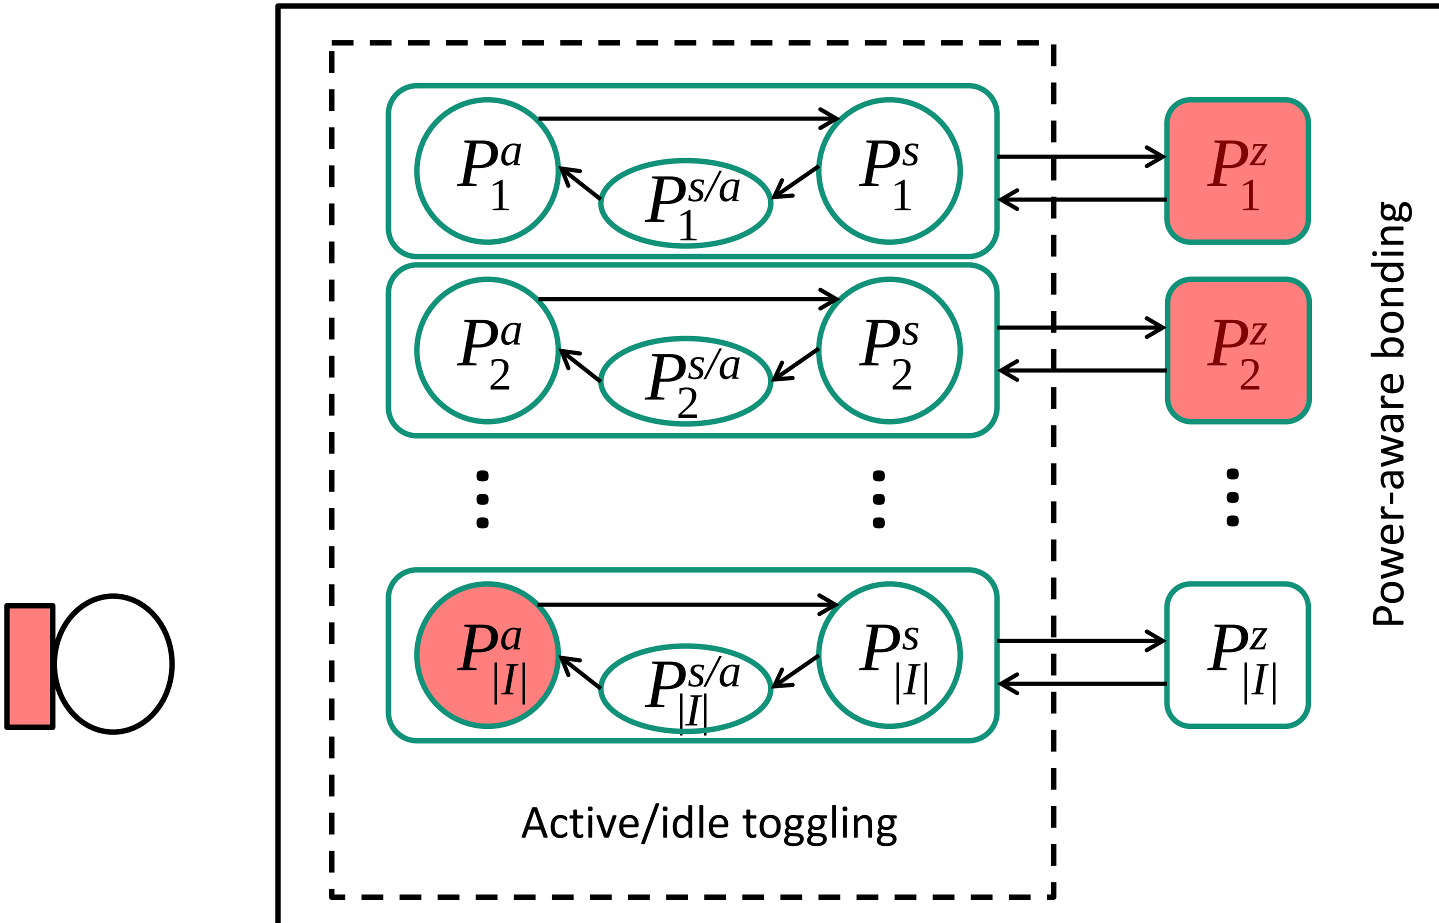

eeeBond overview

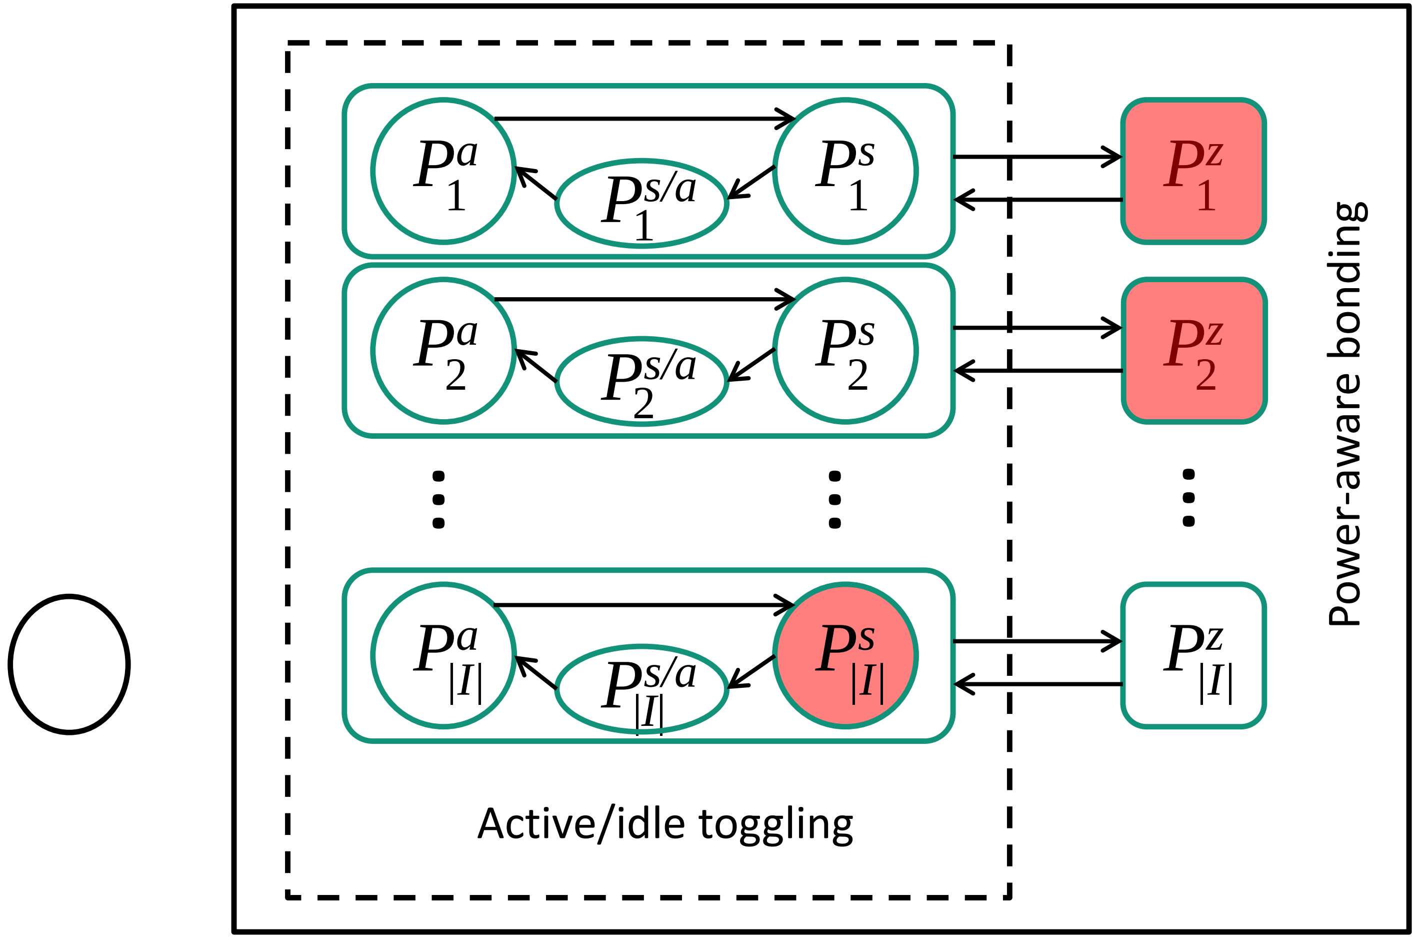

eeeBond overview

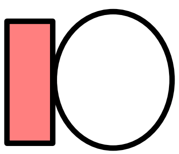

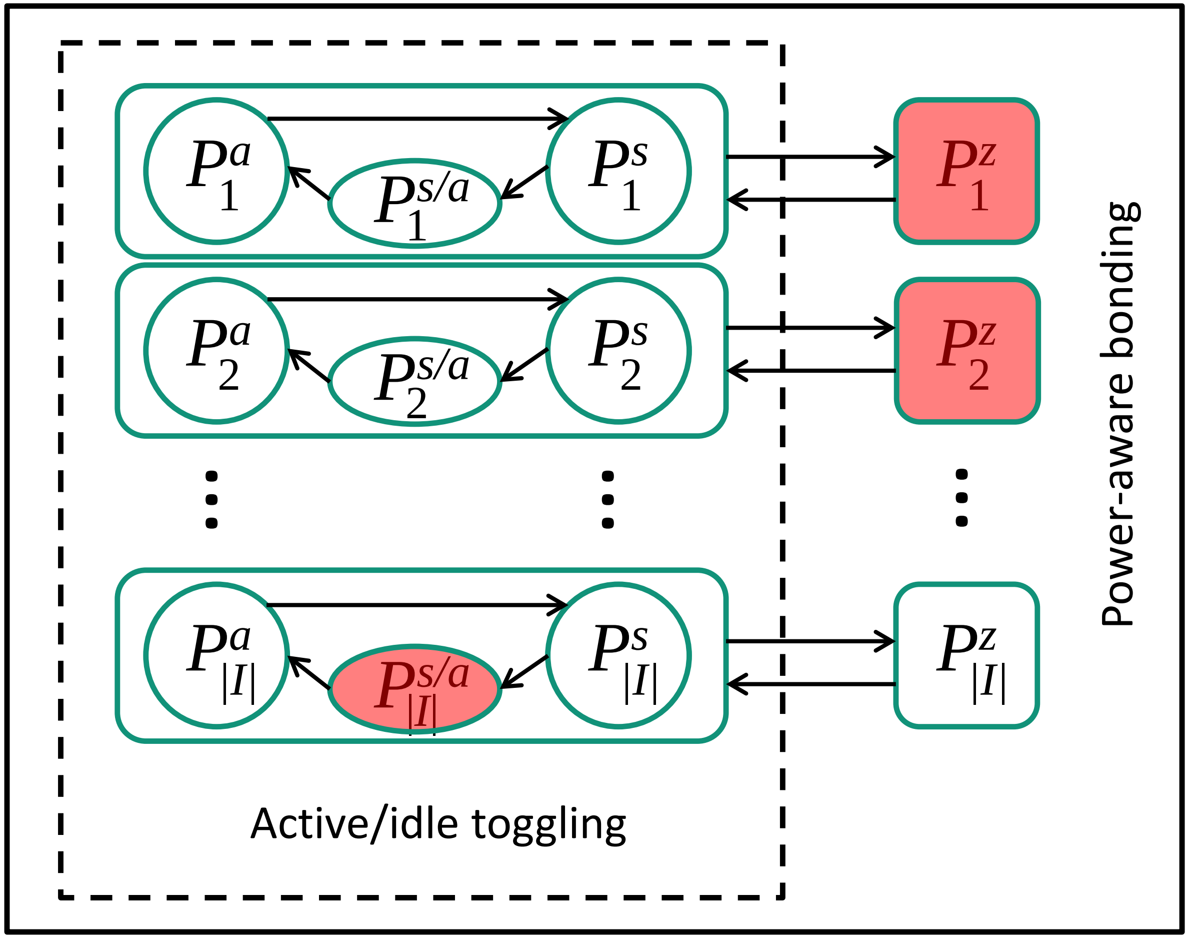

eeeBond overview

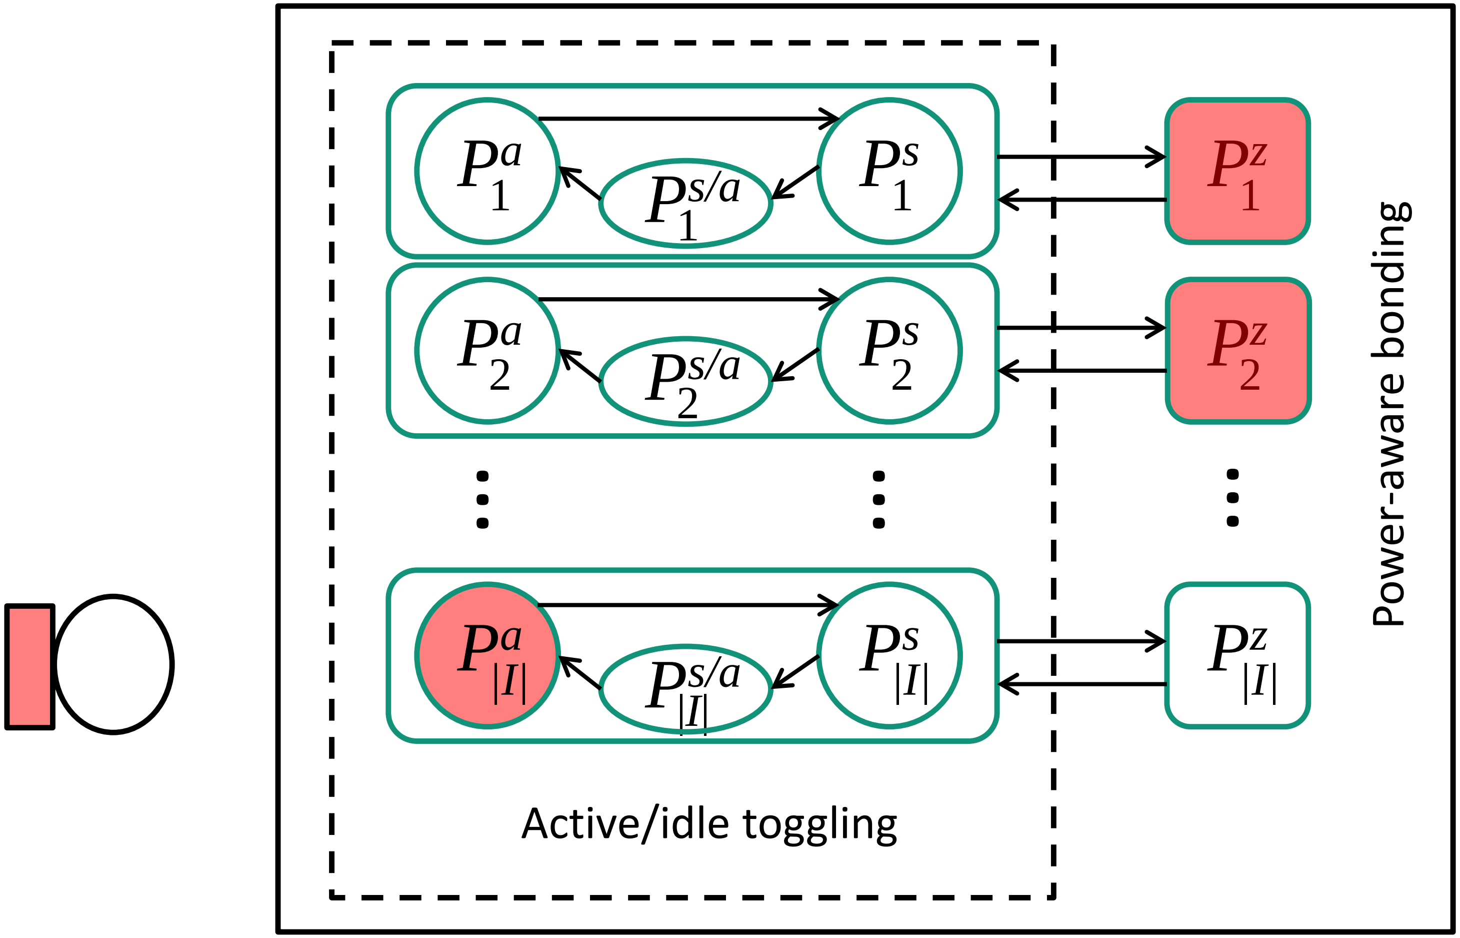

eeeBond overview

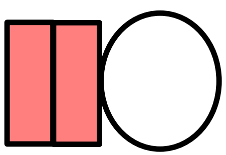

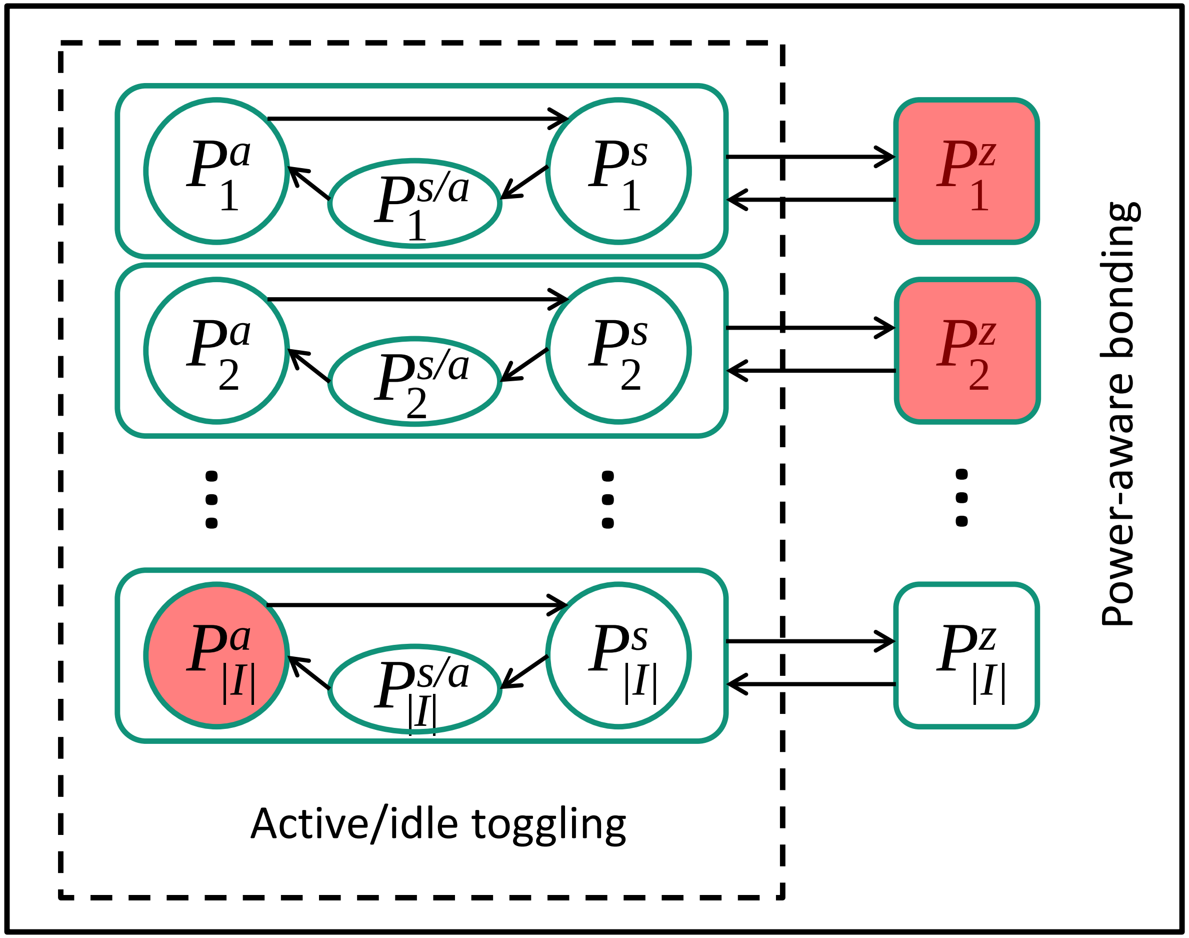

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

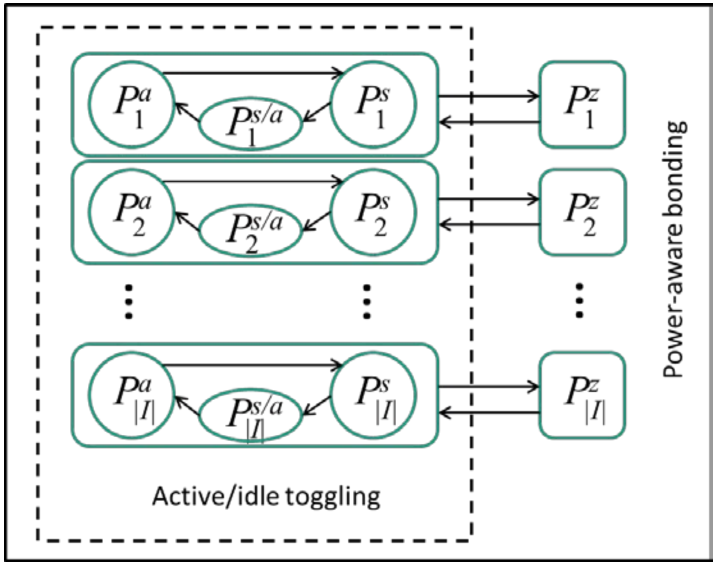

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Sum over all interfaces

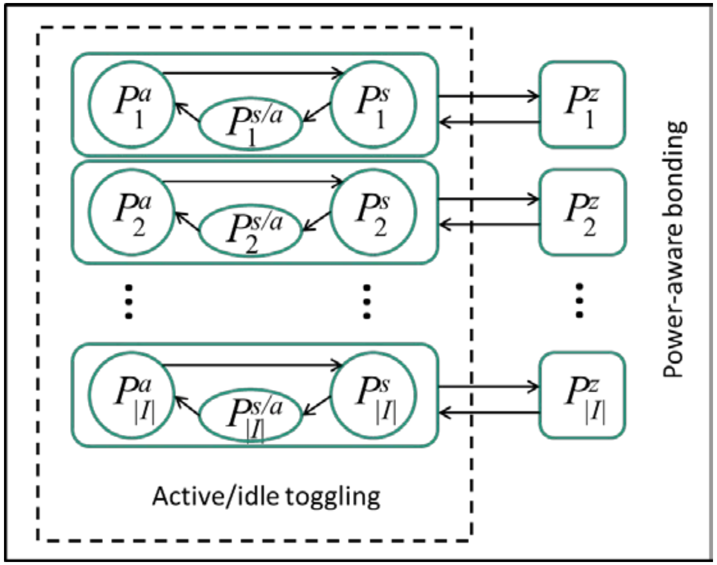

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Active high-power

Short-term sleep

Short setup period

Long-term sleep

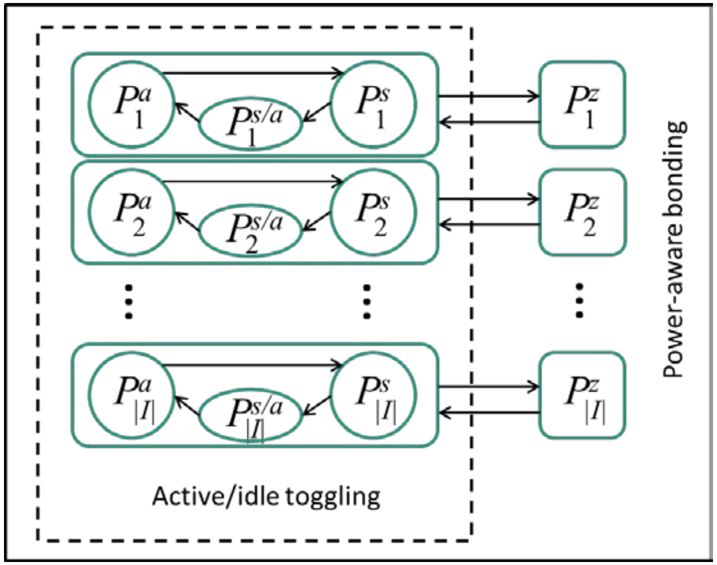

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Active high-power

Short-term sleep

Short setup period

Long-term sleep

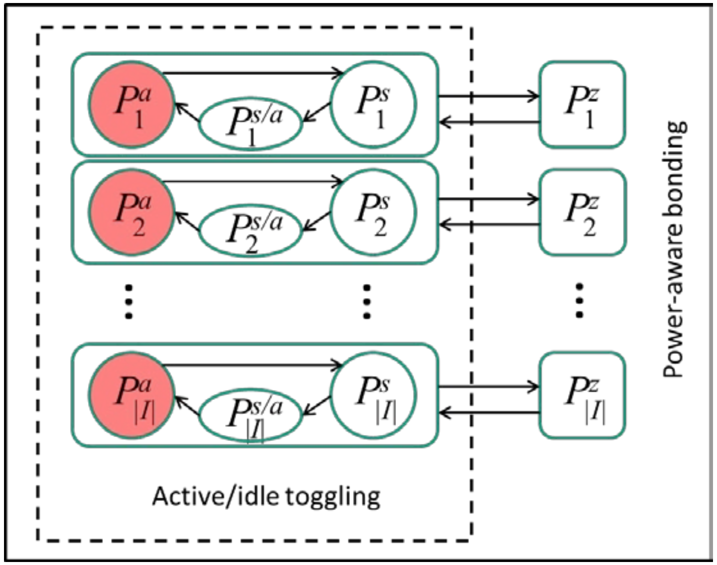

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Active high-power

Short-term sleep

Short setup period

Long-term sleep

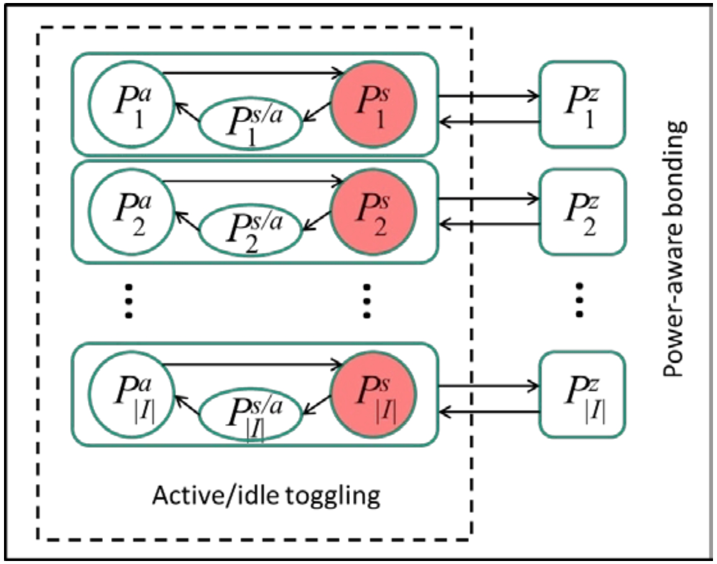

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Active high-power

Short-term sleep

Short setup period

Long-term sleep

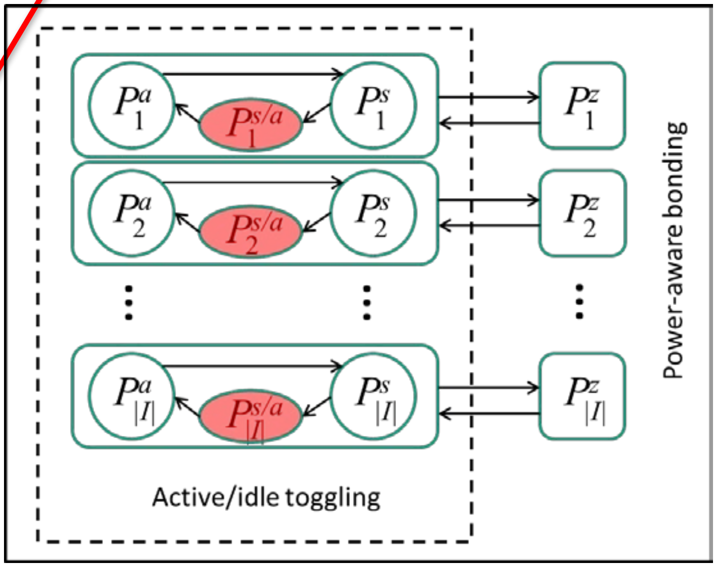

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

Active high-power

Short-term sleep

Short setup period

Long-term sleep

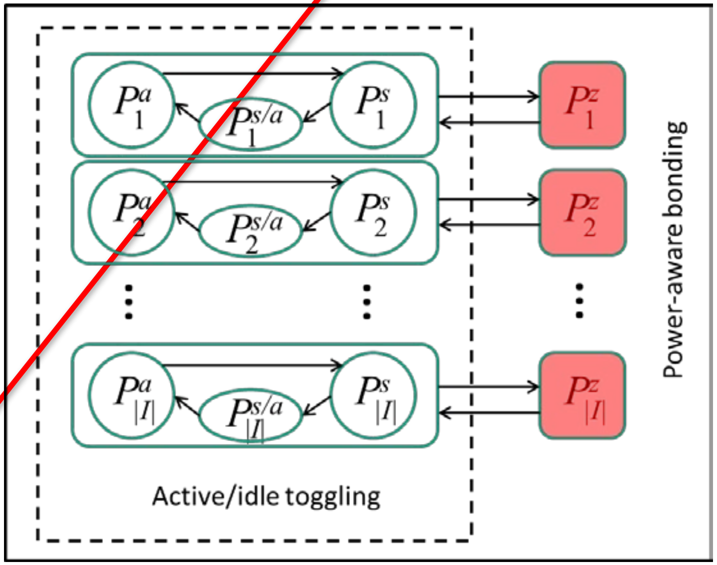

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a}$ $(0 \leq c_i \leq 1)$ $g_i = \frac{P_i^{s/a}}{P_i^a}$ $g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

General model

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} \ (0 \le c_i \le 1)$ $\quad g_i = \frac{P_i^{s/a}}{P_i^a} \ g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

General model

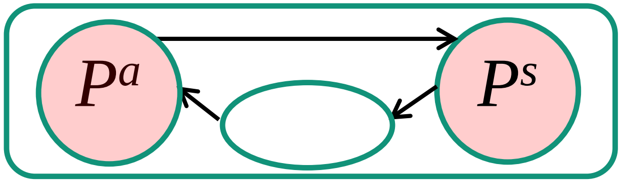

Sleep savings ratio

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} (0 \leq c_i \leq 1)$ $g_i = \frac{P_i^{s/a}}{P_i^a}$ $g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

General model

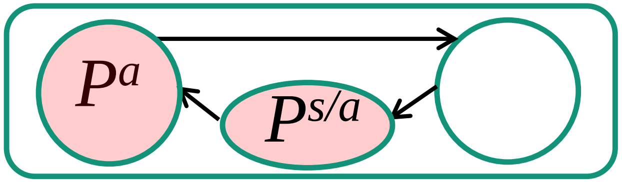

Setup power

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} \ (0 \leq c_i \leq 1)$ $\quad g_i = \frac{P_i^{s/a}}{P_i^a}$ $\quad g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

General model

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} \; (0 \leq c_i \leq 1) \quad g_i = \frac{P_i^{s/a}}{P_i^a} \; g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

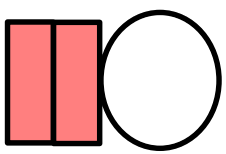

P^a

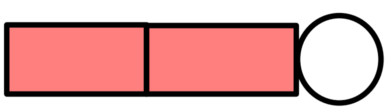

P^a

Power scaling between interfaces

General model

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} \; (0 \leq c_i \leq 1)$ $\quad g_i = \frac{P_i^{s/a}}{P_i^a} \; g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$

General model

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a}\ (0 \leq c_i \leq 1)$ $\quad g_i = \frac{P_i^{s/a}}{P_i^a}\ g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a(\frac{\mu_i}{\mu_0})^x$
- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

General model

- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

General model

Setup time + service time

If arriving to empty system

- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

General model

Time stuck behind earlier packet + service time

If arriving to "busy" system

- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

General model

- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

General model

- Energy usage

$$P_{\mathcal{I}} = \sum_{i \in \mathcal{I}} \left[q_i^a P_i^a + q_i^s P_i^s + q_i^{s/a} P_i^{s/a} + q_i^z P_i^z \right].$$

- Hardware comparison
  - Per-interface modes: $c_i = \frac{P_i^s}{P_i^a} \; (0 \leq c_i \leq 1)$ $\quad g_i = \frac{P_i^{s/a}}{P_i^a} \quad g_i \approx 1$
  - Interface differences: $P_i^a = f(\mu_i)$, where $f(\mu_i) = P_0^a (\frac{\mu_i}{\mu_0})^x$
- Per-interface delay
  - Setup time:

$$w_k = \begin{cases} \Delta_i + \frac{l_k}{\mu_i}, & \text{if } t_k > t_{k-1} + w_{k-1} \\ w_{k-1} + t_{k-1} - t_k + \frac{l_k}{\mu_i}, & \text{otherwise.} \end{cases}$$

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

M/G/1 without vacations

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

M/G/1 without vacations

Setup delay penalty

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a} \right]$$

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho) E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Average power usage, conditioned on "active"

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a} \right]$$

Active high-power

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1 - \rho_i}{1 + \lambda E[\Delta_i]} P_i^s + \frac{\lambda(1 - \rho) E[\Delta_i]}{1 + \lambda E[\Delta_i]} P_i^{s/a} \right]$$

Short-term sleep

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Short setup period

Steady-state model

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1 - q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Steady-state model

- M/G/1(E,SU) model
- Waiting times

$$\overline{W_i} = E[w_k|i] = \frac{\lambda E[S_i^2]}{2(1-\rho_i)} + \frac{2E[\Delta_i] + \lambda E[\Delta_i^2]}{2(1+\lambda E[\Delta_i])},$$

- Energy usage of active interfaces

$$\frac{P_i - q_i^z P_i^z}{1-q_i^z} = \left[\rho_i P_i^a + \frac{1-\rho_i}{1+\lambda E[\Delta_i]} P_i^s + \frac{\lambda(1-\rho)E[\Delta_i]}{1+\lambda E[\Delta_i]} P_i^{s/a}\right]$$

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an average target waiting time W^* and an estimated packet inter arrival rate λ, the optimal eBond policy always picks the interface with the lowest service rates μ_i that can support a packet arrival rate*

$$\lambda \leq \lambda_i^* = \frac{2(W^* - E[S_i])}{E[S_i^2] + 2E[S_i](W^* - E[S_i])}, \qquad (10)$$

where $E[S_i] = \frac{E[l_k]}{\mu_i}$ *and* $E[S_i^2] = \frac{E[l_k^2]}{\mu_i^2}$.

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an average target waiting time* W^* *an* … *eB* … *vi* …

$$E[S_i^2] + 2E[S_i](W^* - E[S_i])$$

where $E[S_i] = \frac{E[l_k]}{\mu_i}$ *and* $E[S_i^2] = \frac{E[l_k^2]}{\mu_i^2}$.

LEMMA 1. *The expected waiting time* W_i *is a monotonically non-decreasing function of the arrival rate* λ *for the region in which* $W_i \geq \Delta_i$*, and for any* $\lambda \leq \lambda^*$ *for which* $W_i^* = W_i(\lambda^*) \geq \Delta_i$*, the waiting time* $W_i(\lambda) \leq W_i(\lambda^*)$.

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an average target waiting time* W^* *an… eB… vi…*

LEMMA 1. *The expected waiting time* W_i *is a monoton-ic… re…* $W_i = W_i(\lambda^*) \leq \Delta_i$*, the waiting time* $W_i(\lambda) \leq W_i(\lambda^*)$*.*

LEMMA 2. *The expected waiting time* W_i *is a monotonically non-increasing function of the service rate* μ_i*.*

$$\ldots E[S_i^2] + 2E[S_i](W^* - E[S_i])$$

where $E[S_i] = \frac{E[l_k]}{\mu_i}$ *and* $E[S_i^2] = \frac{E[l_k^2]}{\mu_i^2}$.

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an average target waiting time* W^* *an … eB … vi …* *where* $E[S$ …

LEMMA 1. *The expected waiting time* W_i *is a monoton- ic … re …* W_i …

LEMMA 2. *The expected waiting time* W_i *is a monotoni- c …*

LEMMA 3. *Given a target delay* $W^* \geq \Delta_i$*, unless there does not exist any interface with higher service rate, the optimal policy never picks a low-power interface with service rate* μ_i *when the packet arrival rate* λ *exceeds an upper bound*

$$\lambda_i^u = \frac{-a_1 + \sqrt{a_1^2 - 4a_2a_0}}{2a_2}, \quad (14)$$

where $a_2 = \Delta_i E[S_i](2W^* - \Delta_i) + \Delta_i E[S_i^2]$*,* $a_1 = E[S_i^2] + 2E[S_i](W^* - \Delta_i) + \Delta_i(\Delta_i - 2W^*)$*, and* $a_0 = 2(\Delta - W^*)$*.*

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an* … *an* … *eB* … *vi* … *where* $E[S_i$ …

LEMMA 1. *The exp* … *ic* … *re* … W_i …

LEMMA 2. *The* … *co* … $=$ …

LEMMA 3. … *does not exist* … *timal policy n* … *rate* μ_i *when th* …

where $a_2 = \Delta_i E[S_i](2W^* - \Delta_i) + \Delta_i E[S_i^2]$, $a_1 = E[S_i^2] + 2E[S_i](W^* - \Delta_i) + \Delta_i(\Delta_i - 2W^*)$, *and* $a_0 = 2(\Delta - W^*)$.

LEMMA 4. *The expected power usage* P_i *is a monotonically non-decreasing function of the service rate* μ *whenever the relative energy scaling parameter* x *satisfies the condition:*

$$x \geq x^* = \frac{\Delta_i + E[S_i]}{G + \lambda H}, \quad (15)$$

where $G = c + (1+c)\Delta_i + (1-c)E[S_i]$, $H = \Delta_i(\Delta_i + (1 - c)E[S_i])$, *and* $c = \frac{P_i^s}{P_i^a}$. *Otherwise, the expected power usage* P_i *is a monotonically non-increasing function of* μ.

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

LEMMA 5. *Unless there does not exist another interface with higher service rate, the optimal policy never picks a low-rate, low-power interface with service rate* μ_i *(over an interface with higher service rate) when the packet arrival rate* λ *is less than a lower bound*

$$\lambda_i^l = \frac{1}{H}(\Delta_i + E[S_i] - xG), \quad (18)$$

where G *and* H *are defined as in Lemma 4.*

where $a_2 = \Delta_i E[S_i](2W^* - \Delta_i) + \Delta_i E[S_i^2]$, $a_1 = E[S_i^2] + 2E[S_i](W^* - \Delta_i) + \Delta_i(\Delta_i - 2W^*)$, *and* $a_0 = 2(\Delta - W^*)$.

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

THEOREM 1. *Given an* …

LEMMA 1. *The exp* …

LEMMA 2. *The* …

LEMMA 3. … *does not exist* … *timal policy n* … *rate μ_i when th* …

where $a_2 = \Delta_i E[S_i](2W^* - \Delta_i) + \Delta_i E[S_i^2]$, $a_1 = E[S_i^2] + 2E[S_i](W^* - \Delta_i) + \Delta_i(\Delta_i - 2W^*)$, *and* $a_0 = 2(\Delta - W^*)$.

LEMMA 4. *The expected power usage P_i is a monotoni-* …

LEMMA 5. *Unless there does not exist another interface* …

THEOREM 2. *Given a target waiting time $W^* \geq \max_i \Delta_i$ and arrival rate λ, the optimal eeeBond policy picks the lowest powered interface that satisfy both (i) $\lambda_i^l \leq \lambda$, and (ii) $\lambda \leq \lambda_i^u$, where λ_i^l and λ_i^u are given by equations (18) and (14), respectively. In the case no interface satisfies both constraints, the optimal policy picks the highest capacity interface.*

Protocol optimization

- M/G/1(E,SU) model
  - Optimized eBond
  - Optimized eeeBond
- High-level summary
  - Theorems/lemmas specifying interface selection

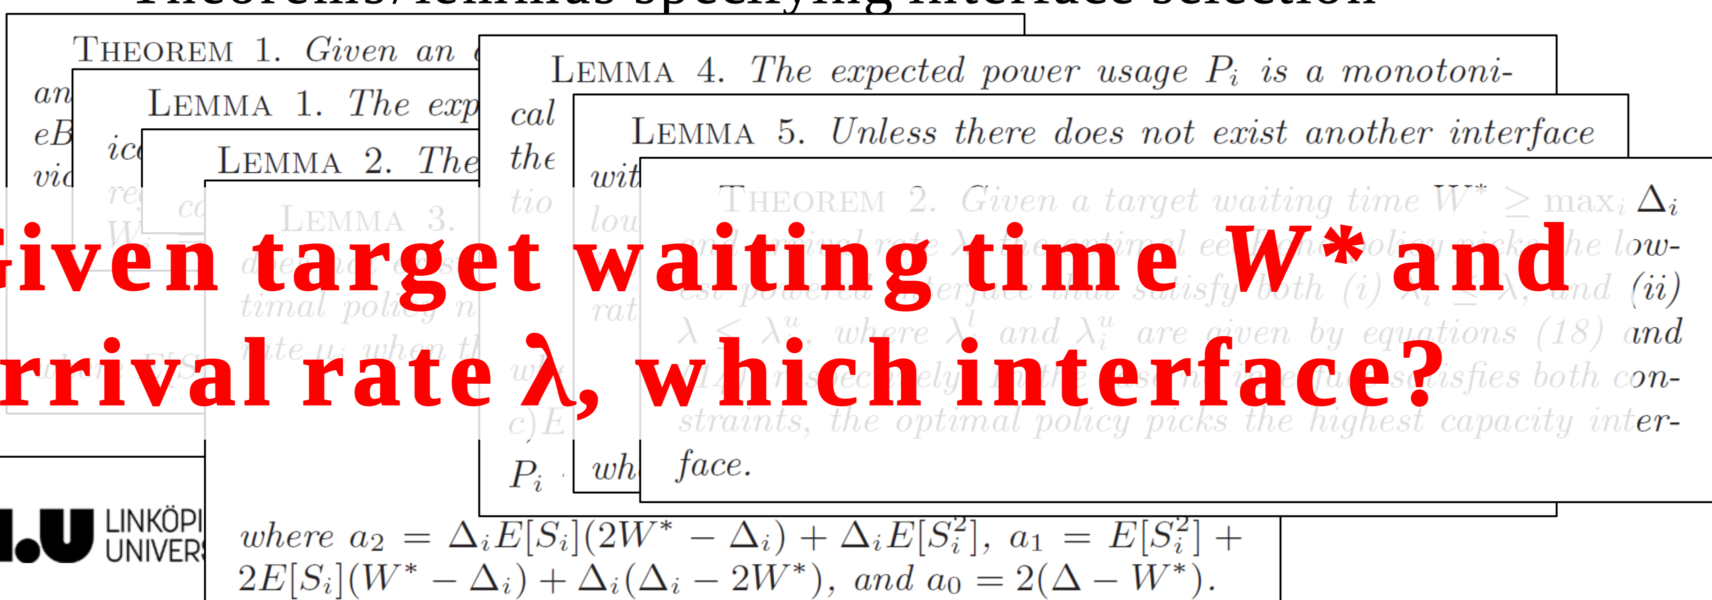

Given target waiting time W^* and arrival rate λ, which interface?

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

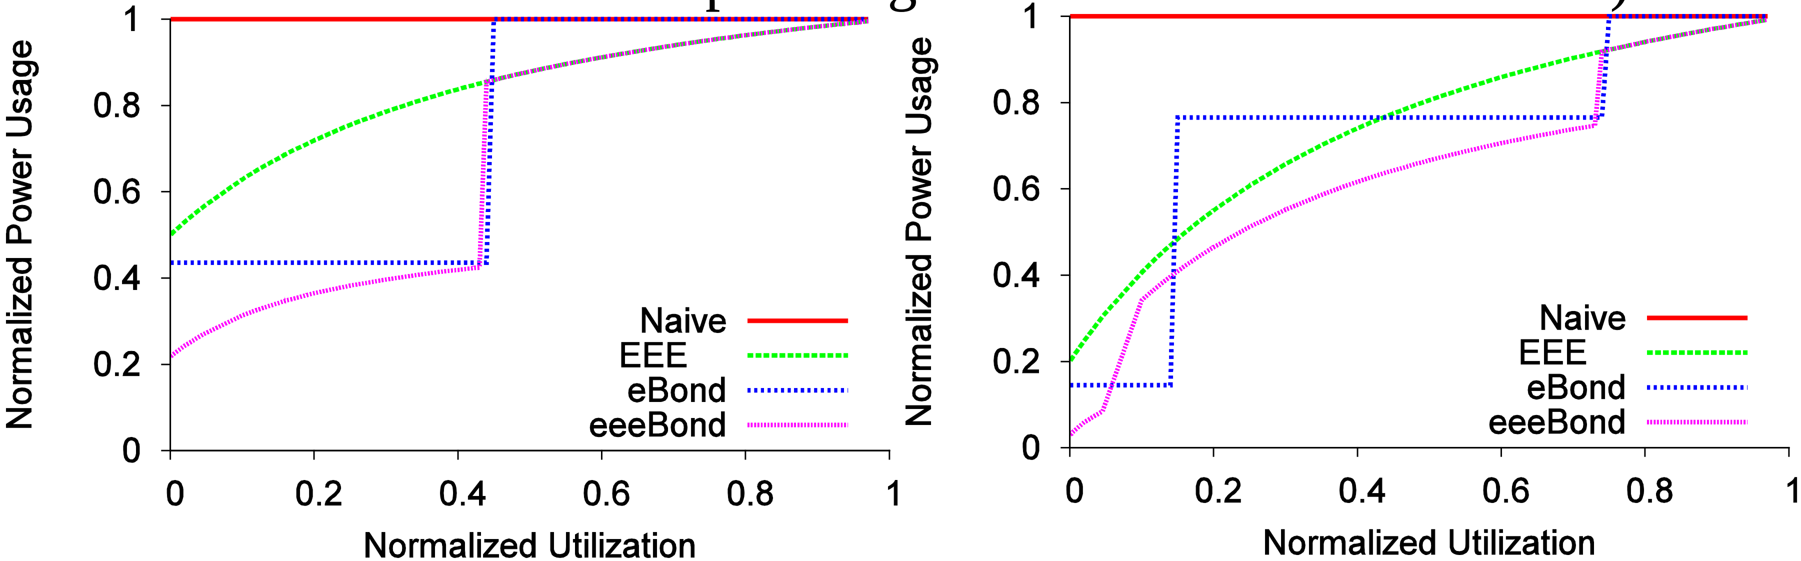

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

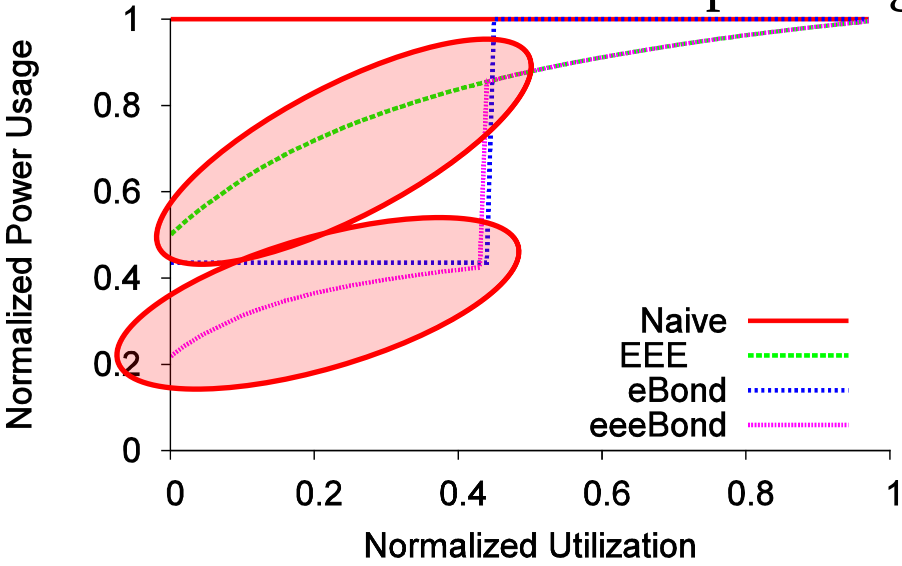

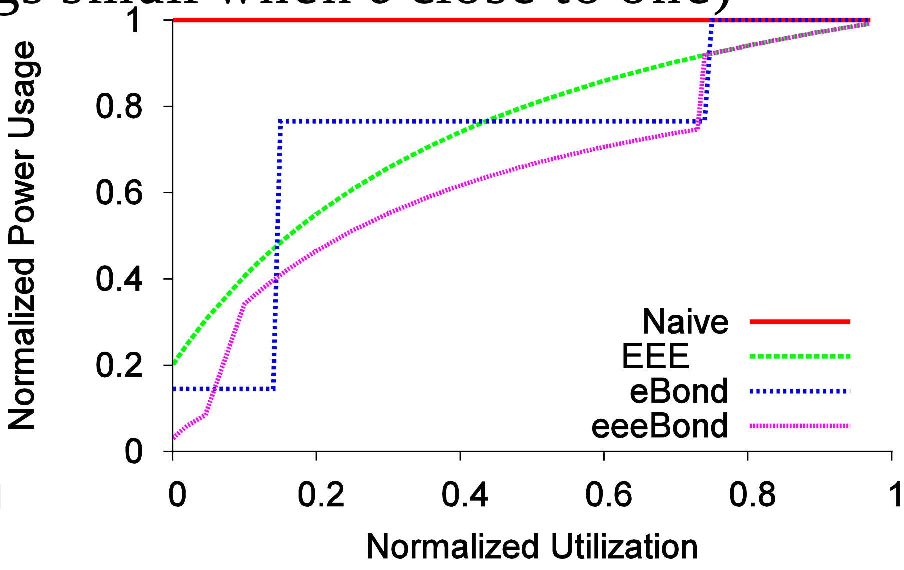

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

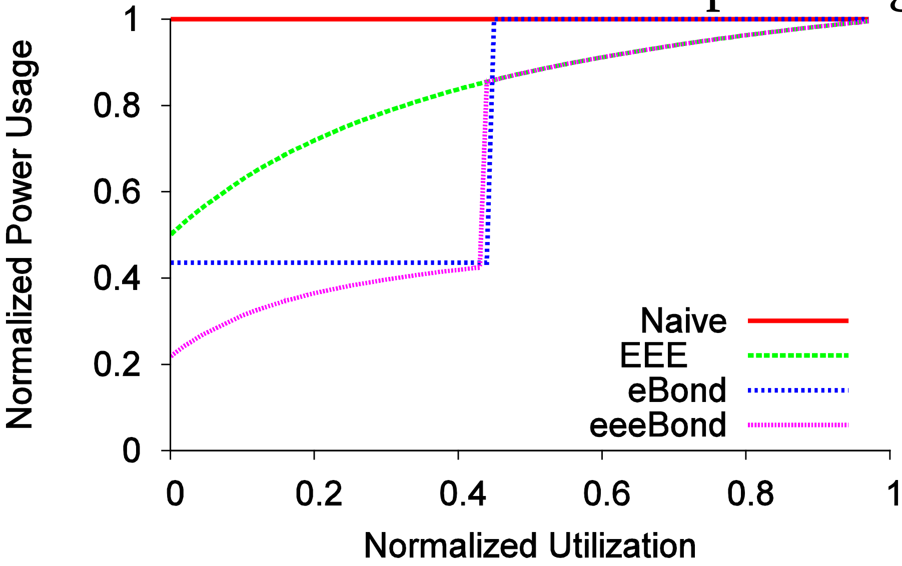

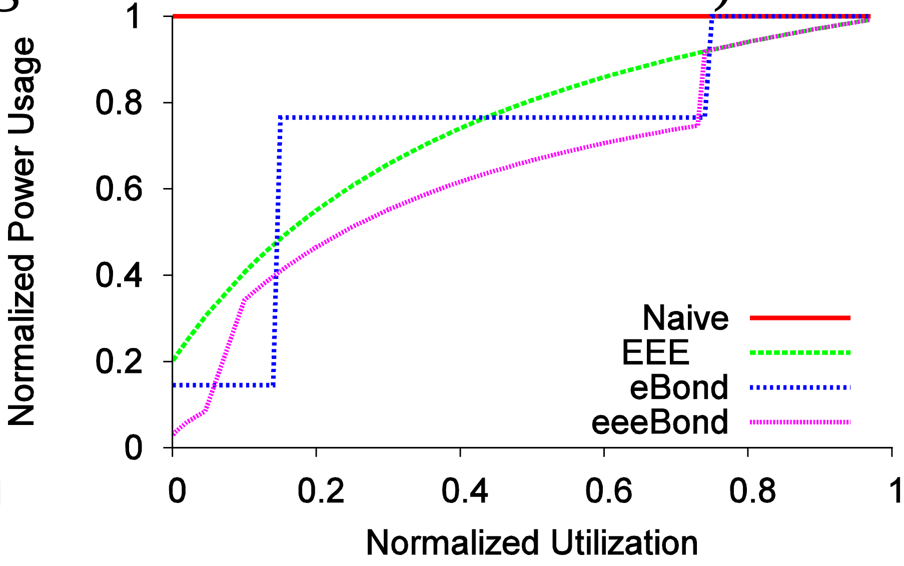

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

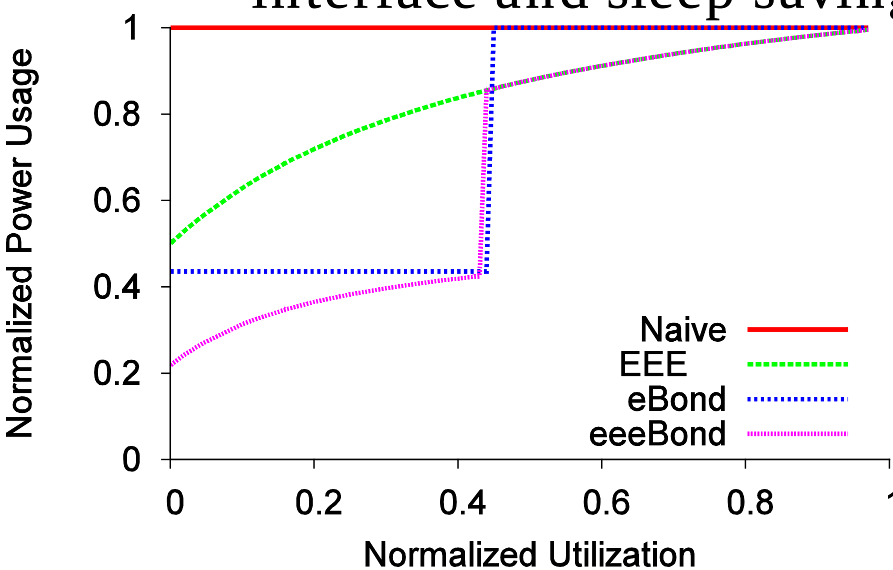

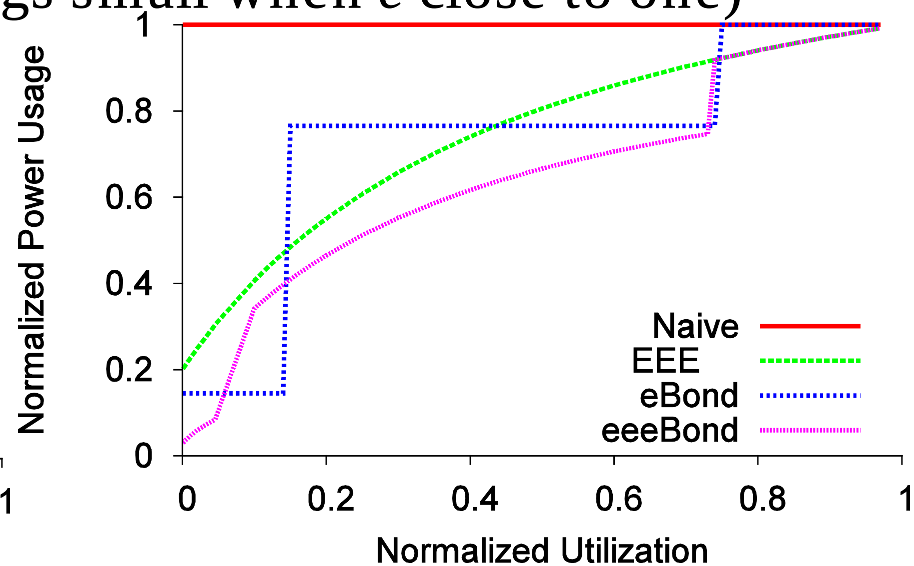

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

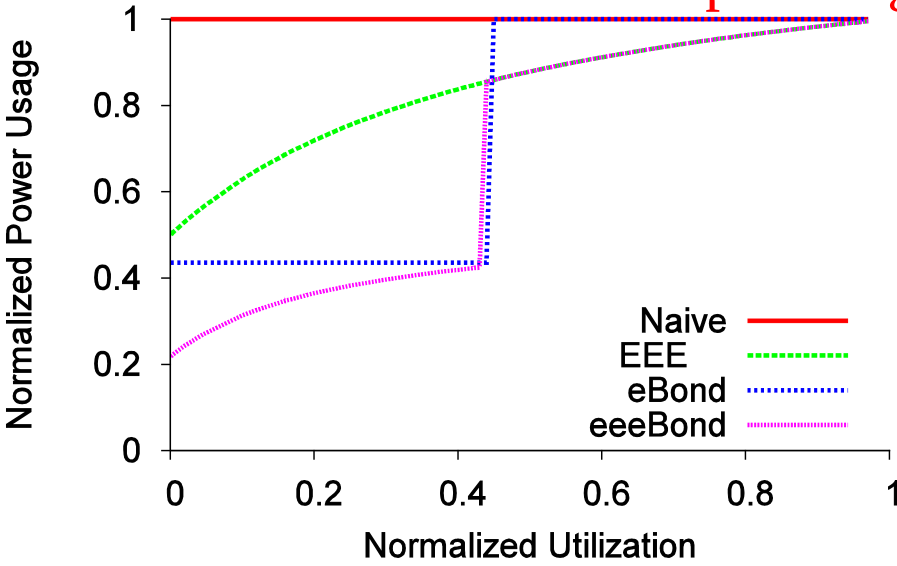

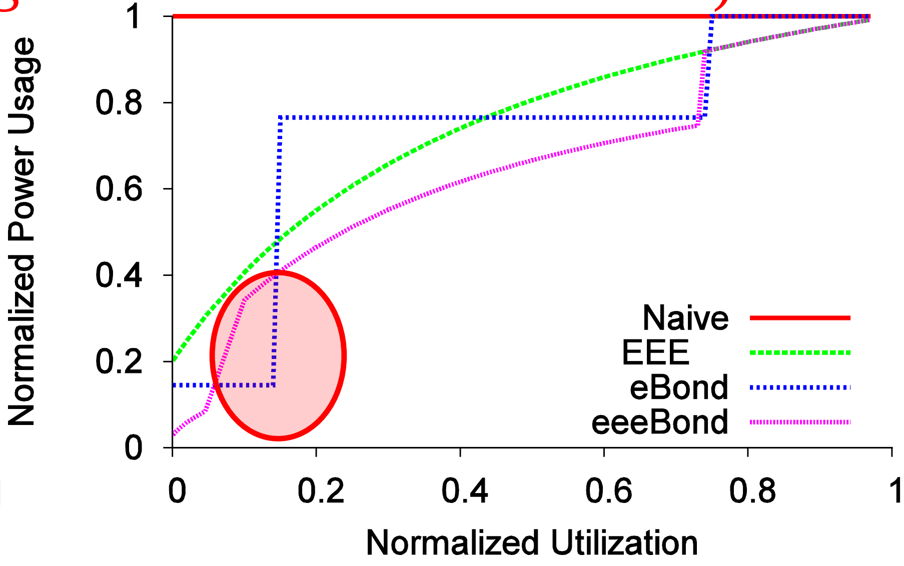

Example scenarios

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

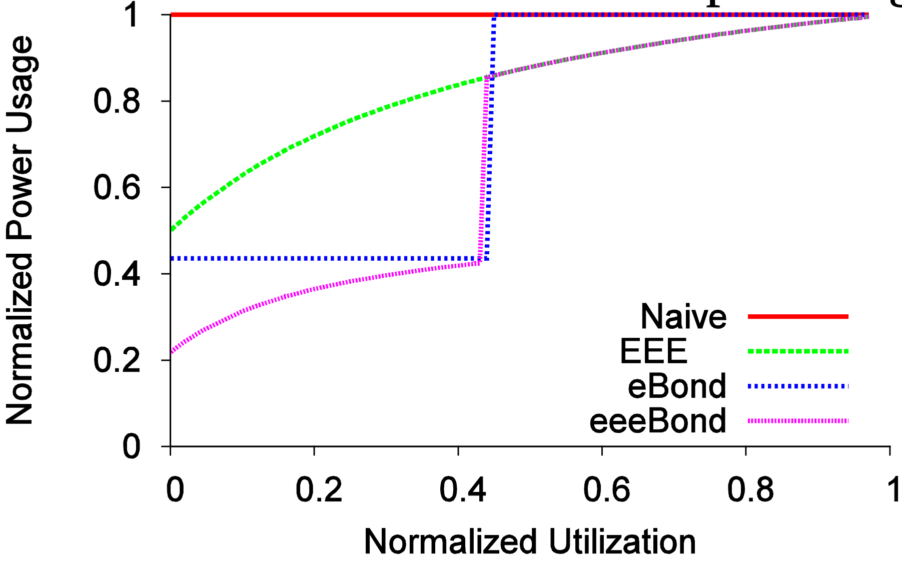

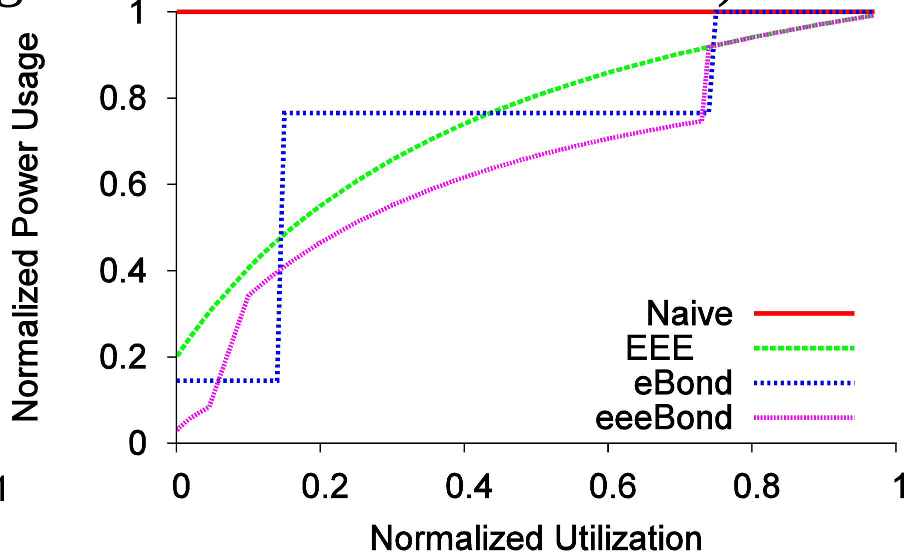

Time-of-day comparison

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

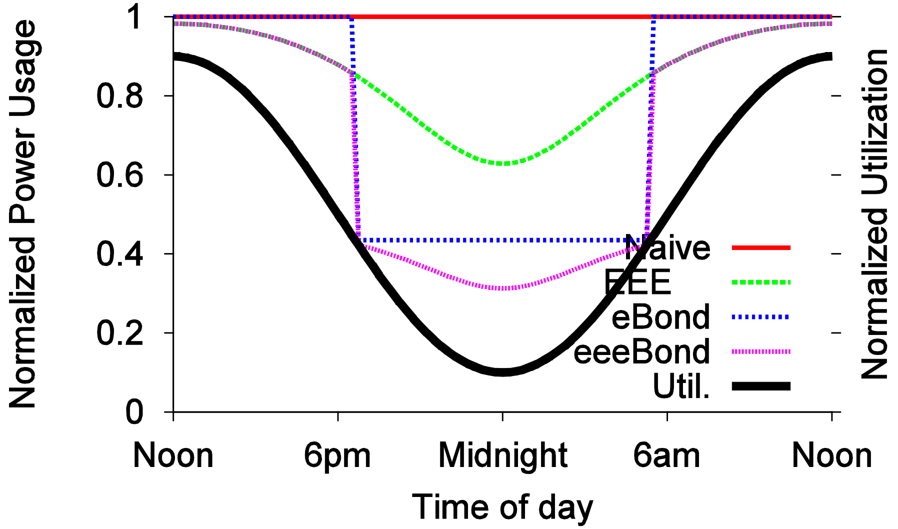

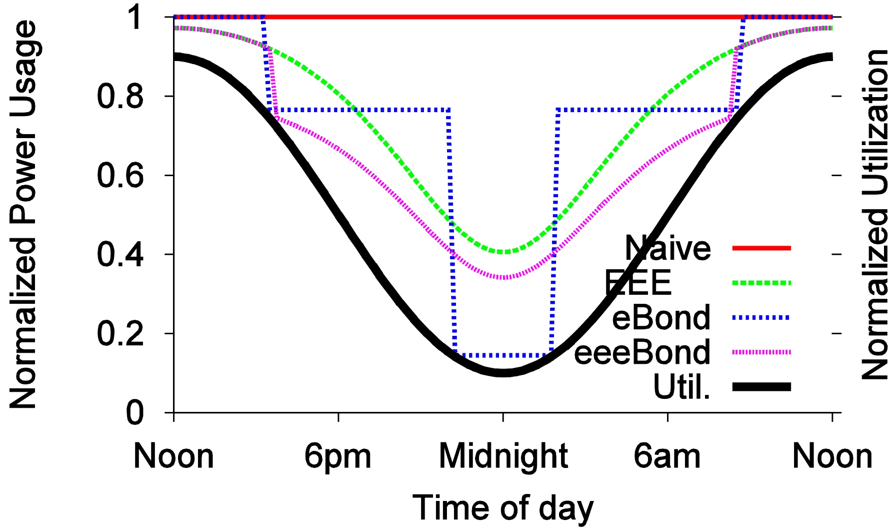

Time-of-day comparison

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

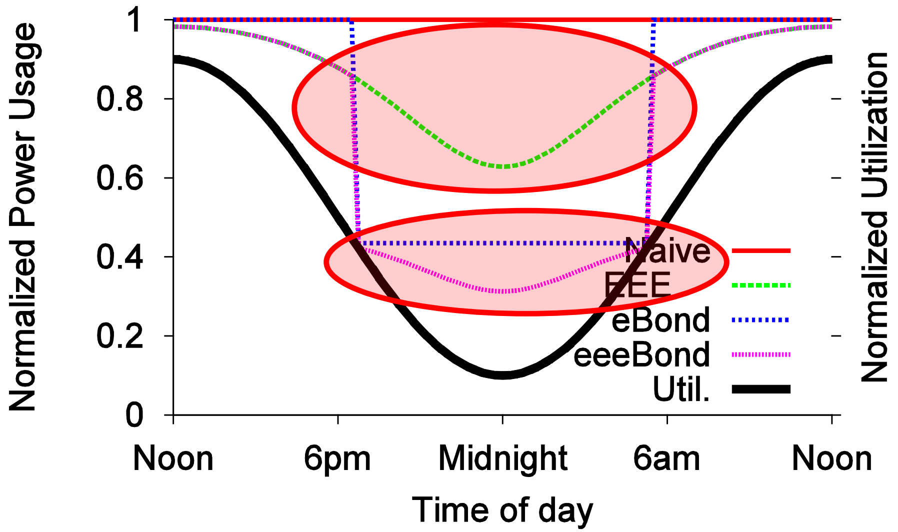

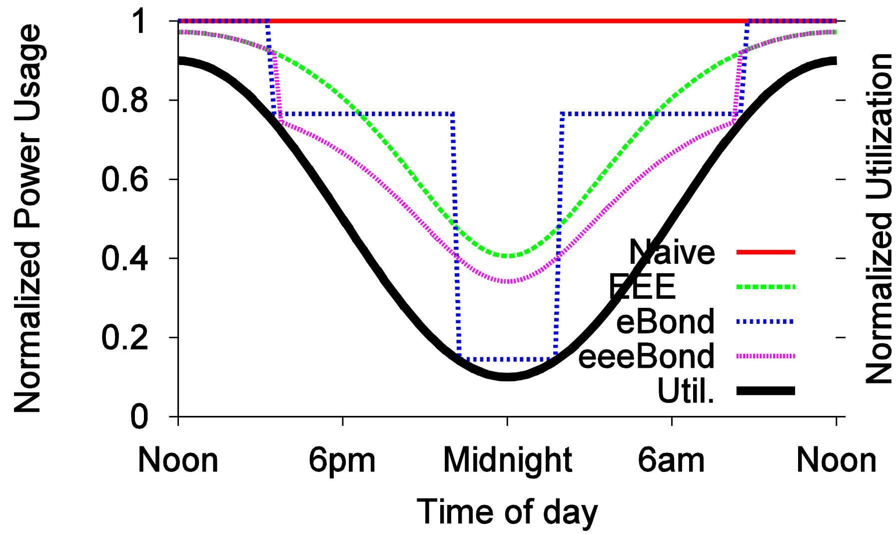

Time-of-day comparison

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when c close to one)

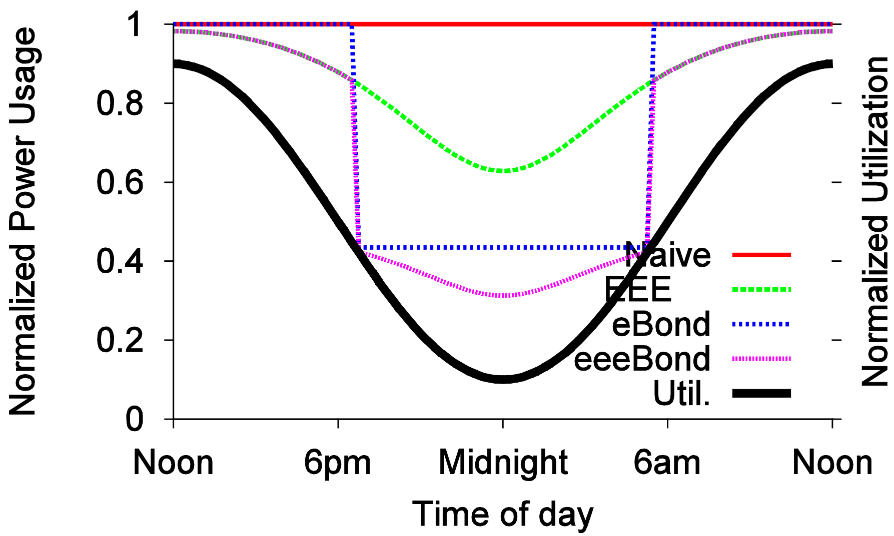

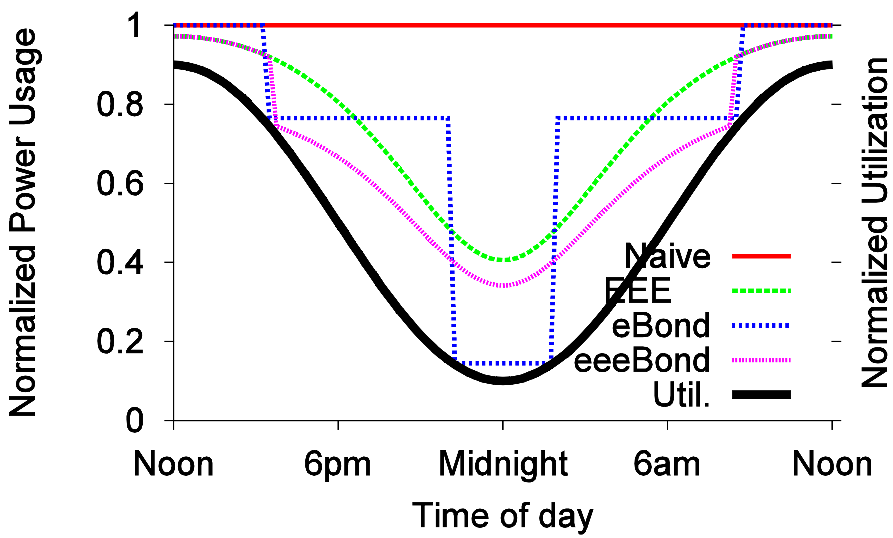

Time-of-day comparison

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

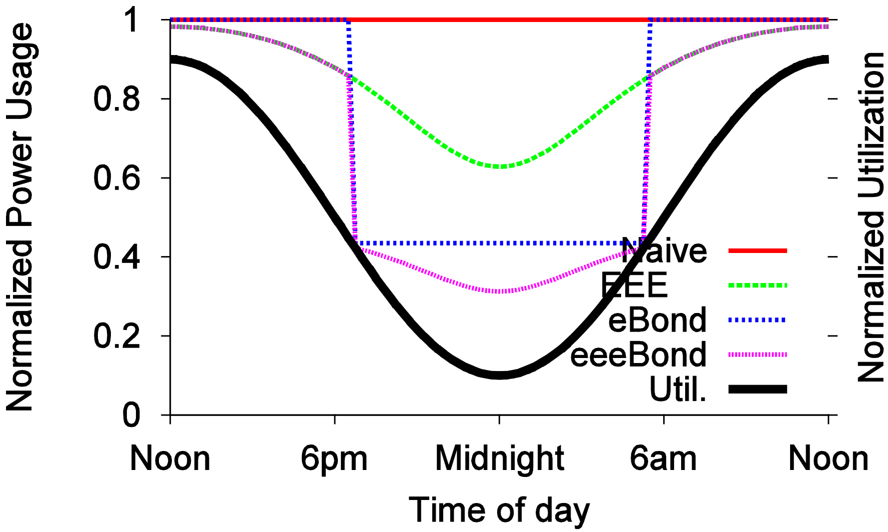

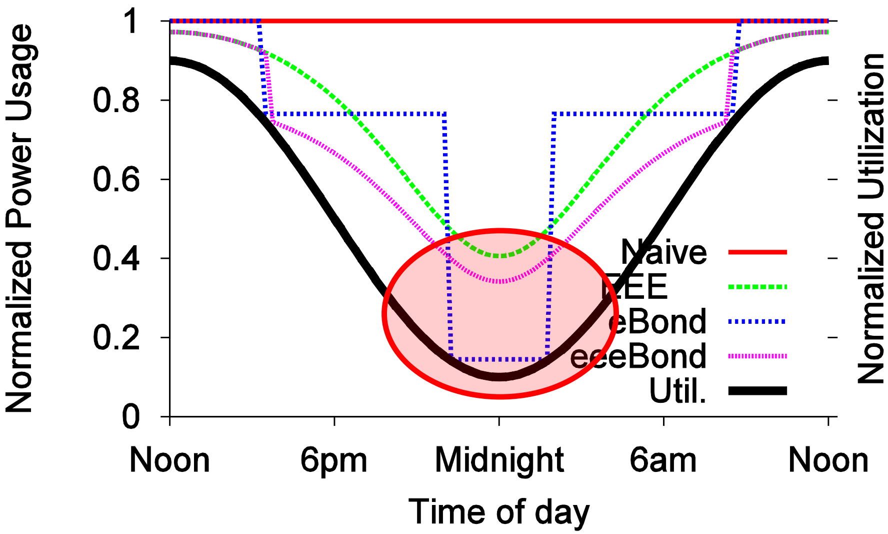

Time-of-day comparison

- EEE and eeeBond adopt power usage for individual interface(s)
- eeeBond often the winner, but cases where eBond even better
  - increase in delay prevent eeeBond using lower-power interface and sleep savings small when *c* close to one)

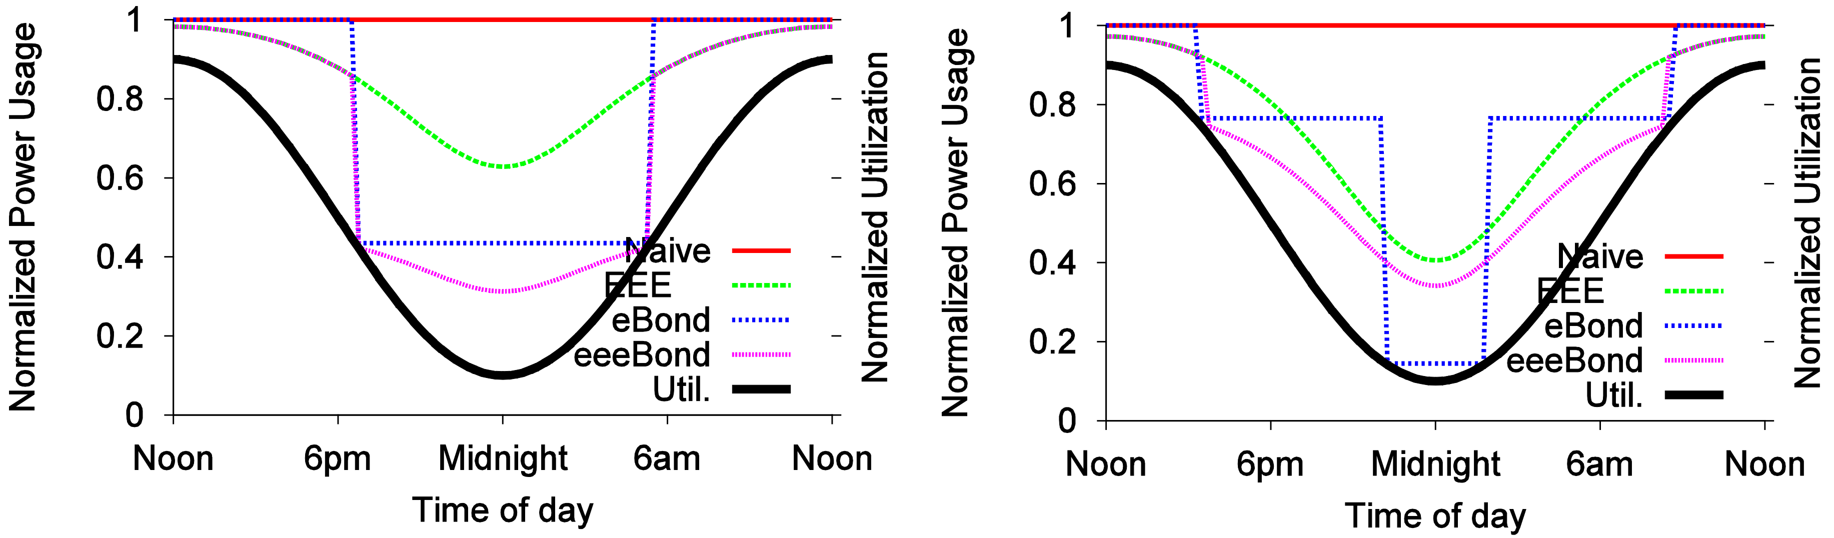

Normalized power usage with diurnal model

- eBond outperforms EEE (sometimes even eeeBond)
- When significant sleep savings (i.e., smaller *c*), eeeBond is the best protocol

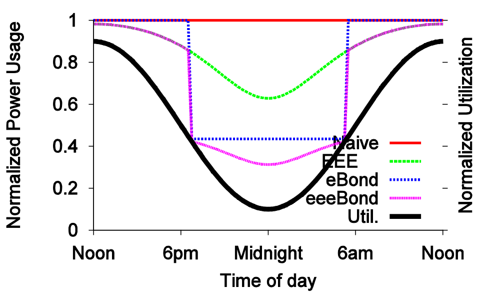

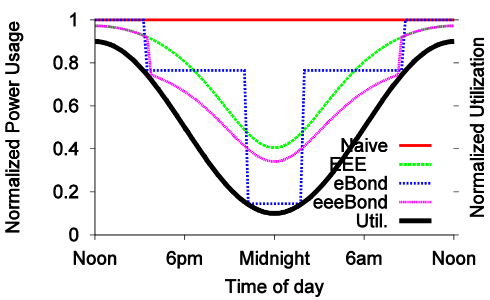

Scenario			Current $c = 0.8$			Future $c = 0.2$		
	$\overline{u}$	x	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	0.5	1.2	0.95	0.76	**0.73**	0.76	0.76	**0.65**
	0.25	1.2	0.90	**0.49**	0.49	0.58	0.49	**0.38**
	0.125	1.2	0.87	0.44	**0.39**	0.44	0.44	**0.25**
	0.25	0.8	0.90	0.62	**0.60**	0.58	0.62	**0.47**
	0.25	2	0.90	**0.32**	0.34	0.58	0.32	**0.27**
Three	0.5	1.2	0.95	**0.74**	0.81	0.76	0.74	**0.68**
	0.25	1.2	0.90	**0.57**	0.59	0.58	0.57	**0.45**
	0.125	1.2	0.87	**0.43**	0.47	0.44	0.43	**0.30**
	0.25	0.8	0.90	**0.66**	0.67	0.58	0.66	**0.51**
	0.25	2	0.90	**0.45**	0.48	0.58	0.45	**0.36**

Normalized power usage with diurnal model

- eBond outperforms EEE (sometimes even eeeBond)
- When significant sleep savings (i.e., smaller *c*), eeeBond is the best protocol

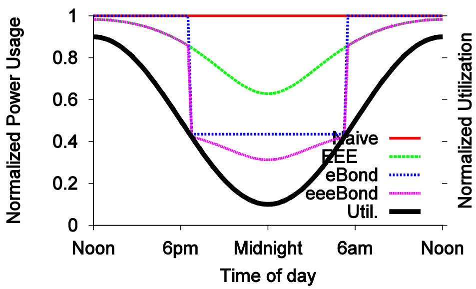

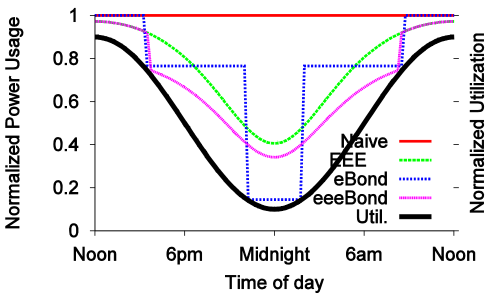

Scenario			Current $c = 0.8$			Future $c = 0.2$		
	$\overline{u}$	x	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	0.5	1.2	0.95	0.76	**0.73**	0.76	0.76	**0.65**
	0.25	1.2	0.90	**0.49**	0.49	0.58	0.49	**0.38**
	0.125	1.2	0.87	0.44	**0.39**	0.44	0.44	**0.25**
	0.25	0.8	0.90	0.62	**0.60**	0.58	0.62	**0.47**
	0.25	2	0.90	**0.32**	0.34	0.58	0.32	**0.27**
Three	0.5	1.2	0.95	**0.74**	0.81	0.76	0.74	**0.68**
	0.25	1.2	0.90	**0.57**	0.59	0.58	0.57	**0.45**
	0.125	1.2	0.87	**0.43**	0.47	0.44	0.43	**0.30**
	0.25	0.8	0.90	**0.66**	0.67	0.58	0.66	**0.51**
	0.25	2	0.90	**0.45**	0.48	0.58	0.45	**0.36**

Normalized power usage with diurnal model

- eBond outperforms EEE (sometimes even eeeBond)
- When significant sleep savings (i.e., smaller *c*), eeeBond is the best protocol

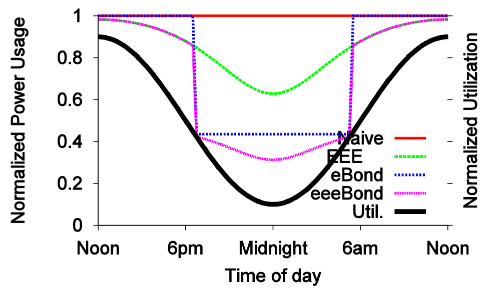

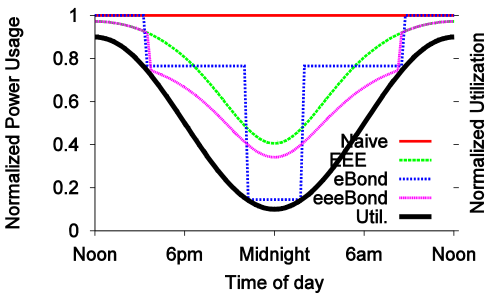

Scenario			Current $c = 0.8$			Future $c = 0.2$		
	$\overline{u}$	x	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	0.5	1.2	0.95	0.76	**0.73**	0.76	0.76	**0.65**
	0.25	1.2	0.90	**0.49**	0.49	0.58	0.49	**0.38**
	0.125	1.2	0.87	0.44	**0.39**	0.44	0.44	**0.25**
	0.25	0.8	0.90	0.62	**0.60**	0.58	0.62	**0.47**
	0.25	2	0.90	**0.32**	0.34	0.58	0.32	**0.27**
Three	0.5	1.2	0.95	**0.74**	0.81	0.76	0.74	**0.68**
	0.25	1.2	0.90	**0.57**	0.59	0.58	0.57	**0.45**
	0.125	1.2	0.87	**0.43**	0.47	0.44	0.43	**0.30**
	0.25	0.8	0.90	**0.66**	0.67	0.58	0.66	**0.51**
	0.25	2	0.90	**0.45**	0.48	0.58	0.45	**0.36**

Normalized power usage with diurnal model

- eBond outperforms EEE (sometimes even eeeBond)
- When significant sleep savings (i.e., smaller *c*), eeeBond is the best protocol

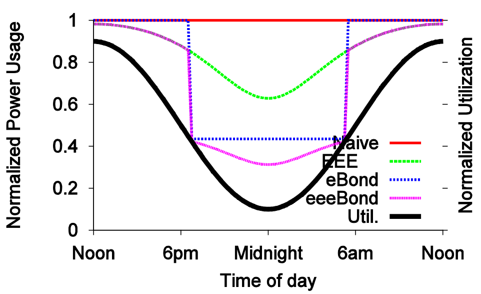

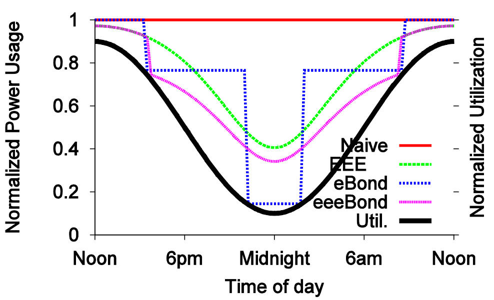

Scenario			Current $c = 0.8$			Future $c = 0.2$		
	$\overline{u}$	x	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	0.5	1.2	0.95	0.76	**0.73**	0.76	0.76	**0.65**
	0.25	1.2	0.90	**0.49**	0.49	0.58	0.49	**0.38**
	0.125	1.2	0.87	0.44	**0.39**	0.44	0.44	**0.25**
	0.25	0.8	0.90	0.62	**0.60**	0.58	0.62	**0.47**
	0.25	2	0.90	**0.32**	0.34	0.58	0.32	**0.27**
Three	0.5	1.2	0.95	**0.74**	0.81	0.76	0.74	**0.68**
	0.25	1.2	0.90	**0.57**	0.59	0.58	0.57	**0.45**
	0.125	1.2	0.87	**0.43**	0.47	0.44	0.43	**0.30**
	0.25	0.8	0.90	**0.66**	0.67	0.58	0.66	**0.51**
	0.25	2	0.90	**0.45**	0.48	0.58	0.45	**0.36**

Normalized power usage with diurnal model

- eBond outperforms EEE (sometimes even eeeBond)
- When significant sleep savings (i.e., smaller *c*), eeeBond is the best protocol

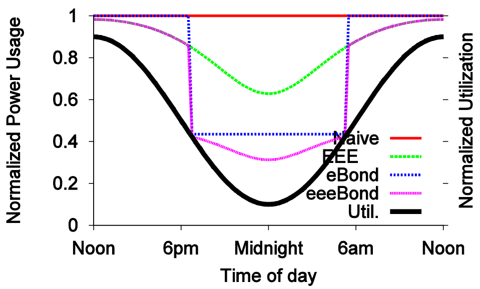

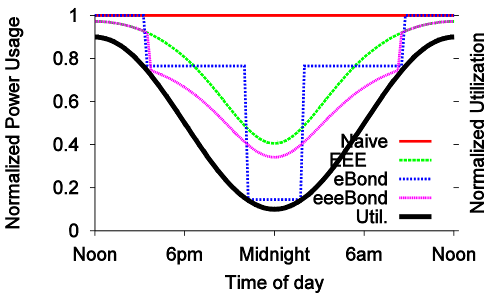

Scenario			Current $c = 0.8$			Future $c = 0.2$		
	$\overline{u}$	x	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	0.5	1.2	0.95	0.76	**0.73**	0.76	0.76	**0.65**
	0.25	1.2	0.90	**0.49**	0.49	0.58	0.49	**0.38**
	0.125	1.2	0.87	0.44	**0.39**	0.44	0.44	**0.25**
	0.25	0.8	0.90	0.62	**0.60**	0.58	0.62	**0.47**
	0.25	2	0.90	**0.32**	0.34	0.58	0.32	**0.27**
Three	0.5	1.2	0.95	**0.74**	0.81	0.76	0.74	**0.68**
	0.25	1.2	0.90	**0.57**	0.59	0.58	0.57	**0.45**
	0.125	1.2	0.87	**0.43**	0.47	0.44	0.43	**0.30**
	0.25	0.8	0.90	**0.66**	0.67	0.58	0.66	**0.51**
	0.25	2	0.90	**0.45**	0.48	0.58	0.45	**0.36**

Trace-driven simulations

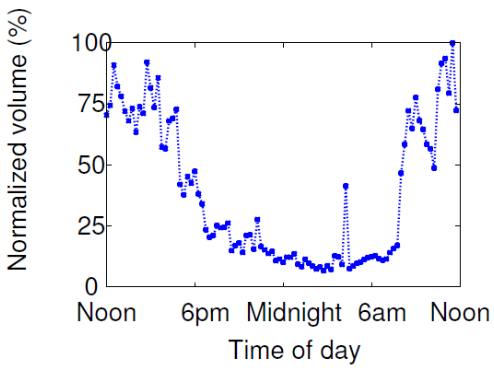

(a) Edge, incoming

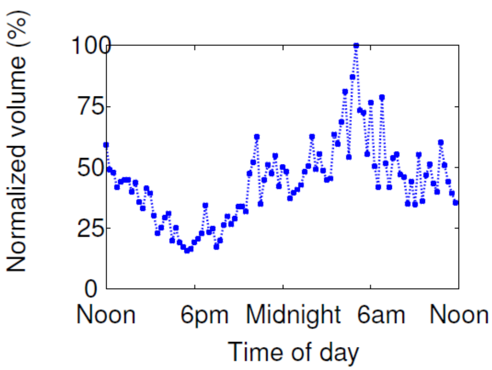

(b) Core, dir-A

- eBond (and eeeBond) always better than EEE
- eeeBond has significant benefits over others when greater sleep benefits (i.e., smaller *c*)

Scenario		Current $c = 0.8$			Future $c = 0.2$		
	Trace	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	Edge, in	0.81	**0.16**	0.43	0.22	0.16	**0.13**
	Edge, out	0.81	**0.38**	0.74	0.22	0.38	**0.20**
	Core, dirA	0.81	0.15	**0.12**	0.21	0.15	**0.04**
	Core, dirB	0.82	0.15	**0.12**	0.24	0.15	**0.04**

Trace-driven simulations

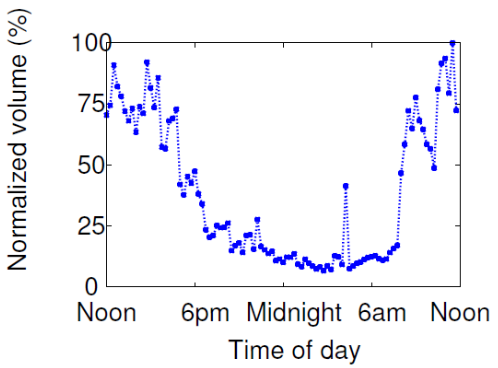

(a) Edge, incoming

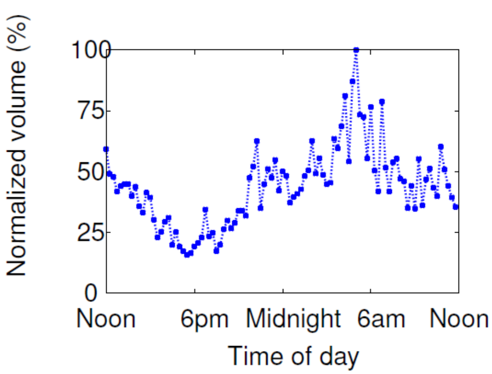

(b) Core, dir-A

- eBond (and eeeBond) always better than EEE
- eeeBond has significant benefits over others when greater sleep benefits (i.e., smaller *c*)

Scenario		Current $c = 0.8$			Future $c = 0.2$		
	Trace	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	Edge, in	0.81	**0.16**	0.43	0.22	0.16	**0.13**
	Edge, out	0.81	**0.38**	0.74	0.22	0.38	**0.20**
	Core, dirA	0.81	0.15	**0.12**	0.21	0.15	**0.04**
	Core, dirB	0.82	0.15	**0.12**	0.24	0.15	**0.04**

Trace-driven simulations

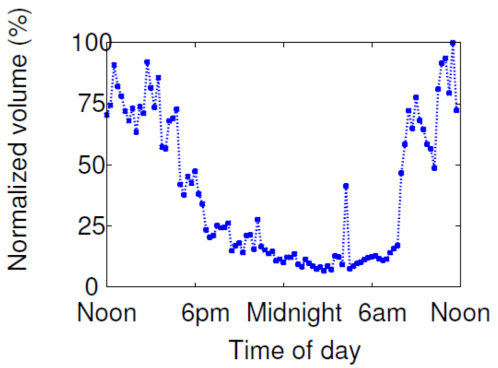

(a) Edge, incoming

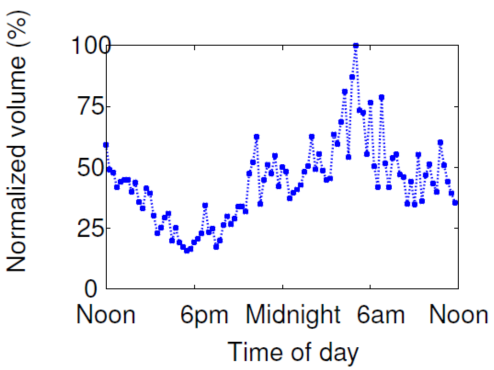

(b) Core, dir-A

- eBond (and eeeBond) always better than EEE
- eeeBond has significant benefits over others when greater sleep benefits (i.e., smaller *c*)

Scenario		Current $c = 0.8$			Future $c = 0.2$		
	Trace	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	Edge, in	0.81	**0.16**	0.43	0.22	0.16	**0.13**
	Edge, out	0.81	**0.38**	0.74	0.22	0.38	**0.20**
	Core, dirA	0.81	0.15	**0.12**	0.21	0.15	**0.04**
	Core, dirB	0.82	0.15	**0.12**	0.24	0.15	**0.04**

Trace-driven simulations

- eBond (and eeeBond) always better than EEE
- eeeBond has significant benefits over others when greater sleep benefits (i.e., smaller *c*)

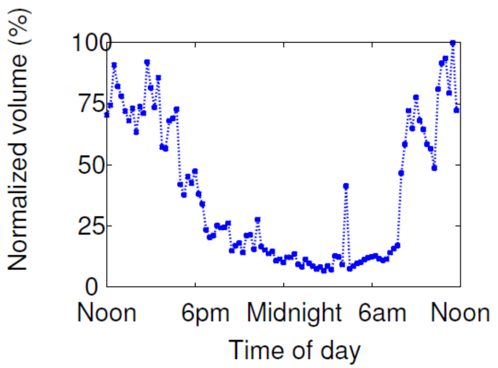

(a) Edge, incoming

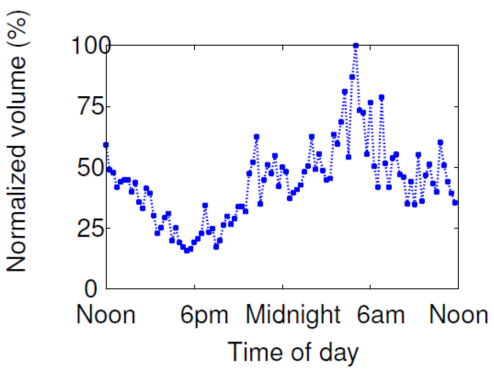

(b) Core, dir-A

Scenario		Current $c = 0.8$			Future $c = 0.2$		
	Trace	EEE	eBond	e^3B	EEE	eBond	e^3B
Two	Edge, in	0.81	**0.16**	0.43	0.22	0.16	**0.13**
	Edge, out	0.81	**0.38**	0.74	0.22	0.38	**0.20**
	Core, dirA	0.81	0.15	**0.12**	0.21	0.15	**0.04**
	Core, dirB	0.82	0.15	**0.12**	0.24	0.15	**0.04**

Model validation

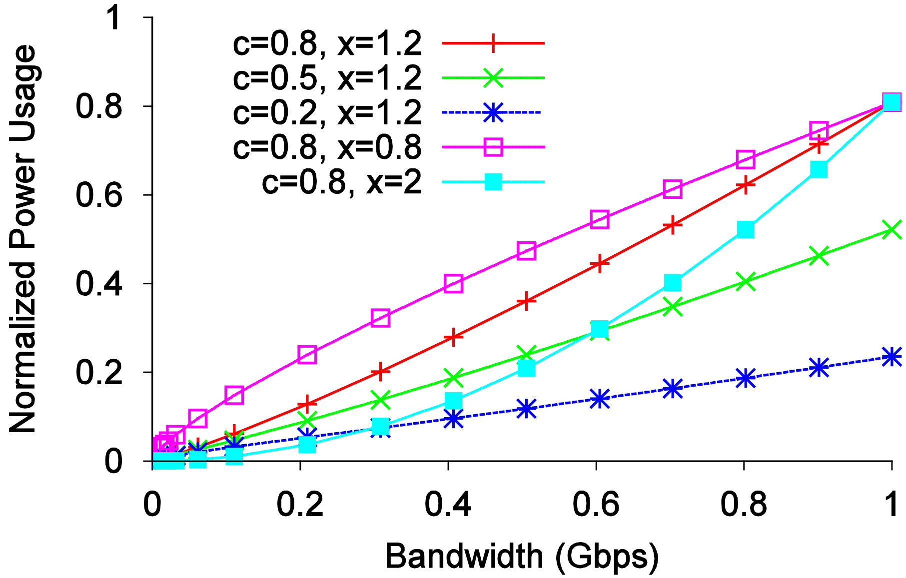

(a) Power model

- Good match

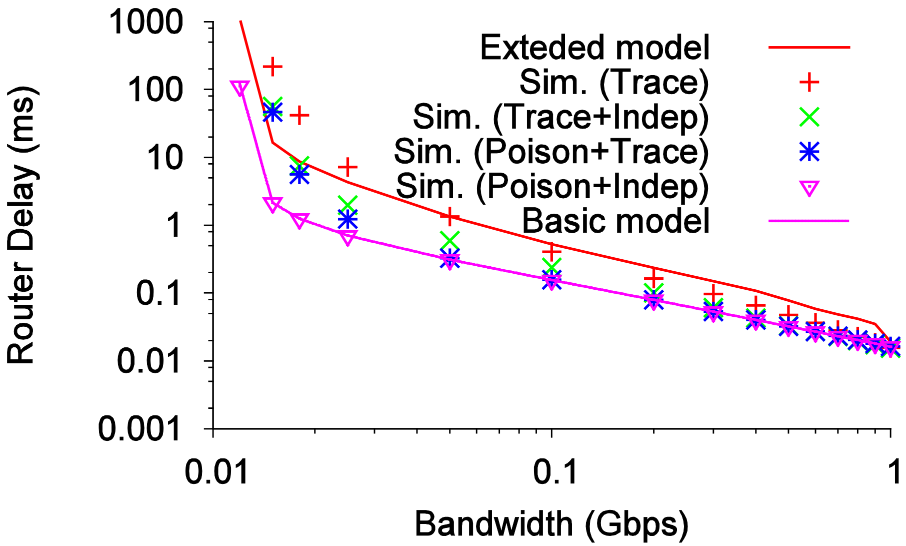

(b) Delay model

- Impact of assumptions

- Extended model improves

Model validation

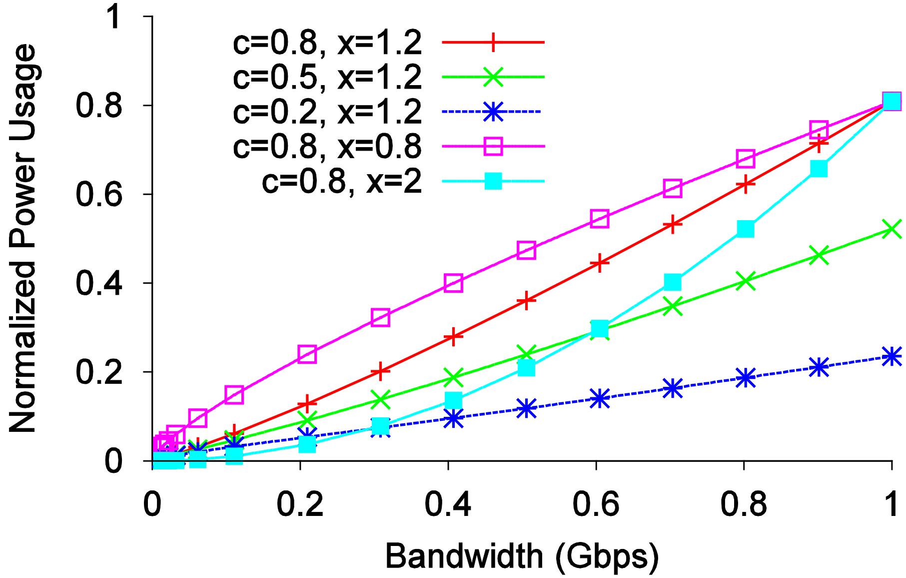

(a) Power model

- Good match

Model validation

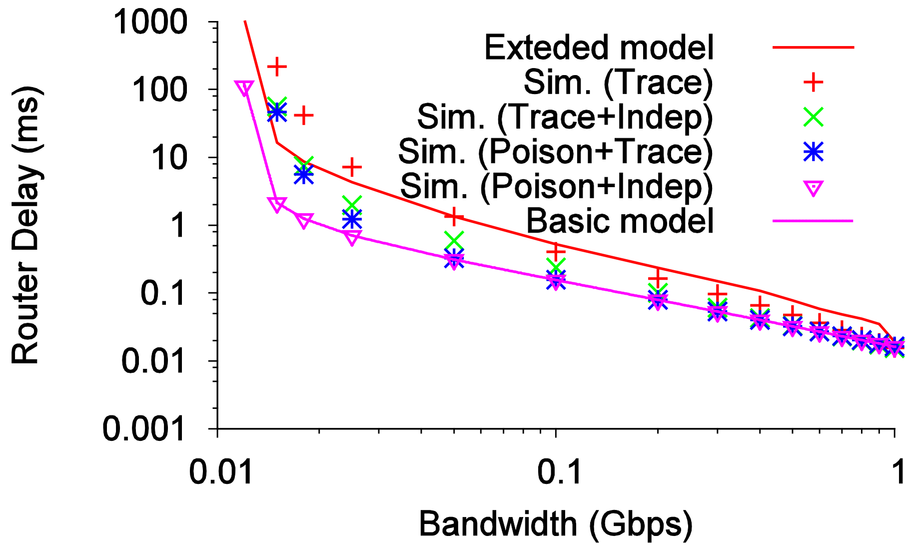

(b) Delay model

- Impact of assumptions
- Extended model improves

Model validation

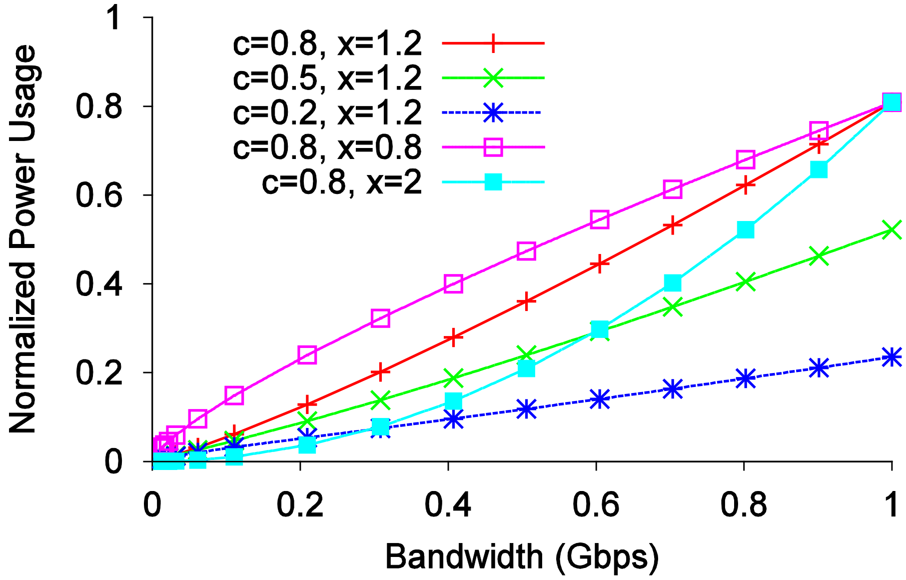

(a) Power model

- Good match

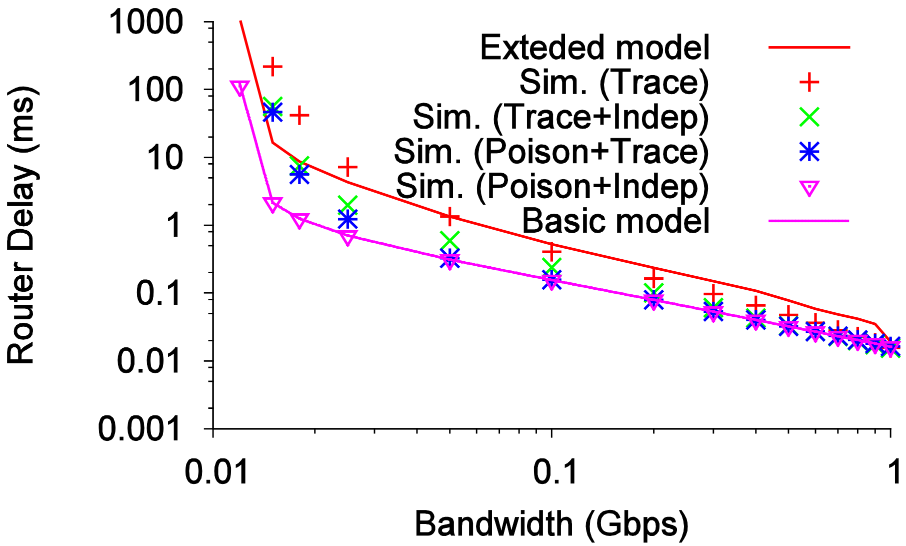

(b) Delay model

- Impact of assumptions

- Extended model improves

Conclusions

- Presented eeeBond
  - Hybrid protocol that combines benefits of EEE and eBond
- Presented a generalized protocol evaluation framework
- Performed protocol optimization of eBond and eeeBond
  - Closed-form expressions of delay and energy tradeoffs
- Characterized energy savings with the different protocols
  - Significant benefits of eBond and eeeBond over EEE (when $x>1$), even when EEE itself becomes more energy proportional
- Future work will refine our extended model and optimize additional aspects of eeeBond

Optimized eeeBond: Energy Efficiency with non-Proportional Router Network Interfaces

$EEE\ +$  $=$ 

Niklas Carlsson (niklas.carlsson@liu.se)

LINKÖPING UNIVERSITY