Cross-Layer Optimization for Wireless Networks with Multi-Receiver Diversity

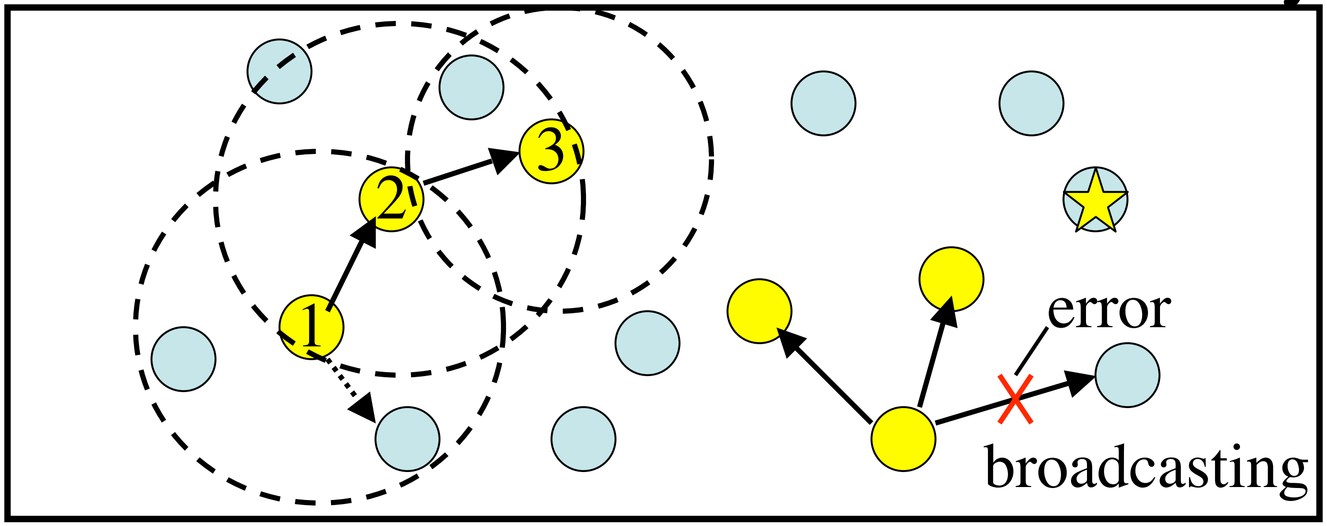

Michael J. Neely
University of Southern California
http://www-rcf.usc.edu/~mjneely/

Comm Theory Workshop, Puerto Rico, May 2006

(Conference paper appears in CISS, March 2006)

*Sponsored by NSF OCE Grant 0520324

Cross-Layer Networking

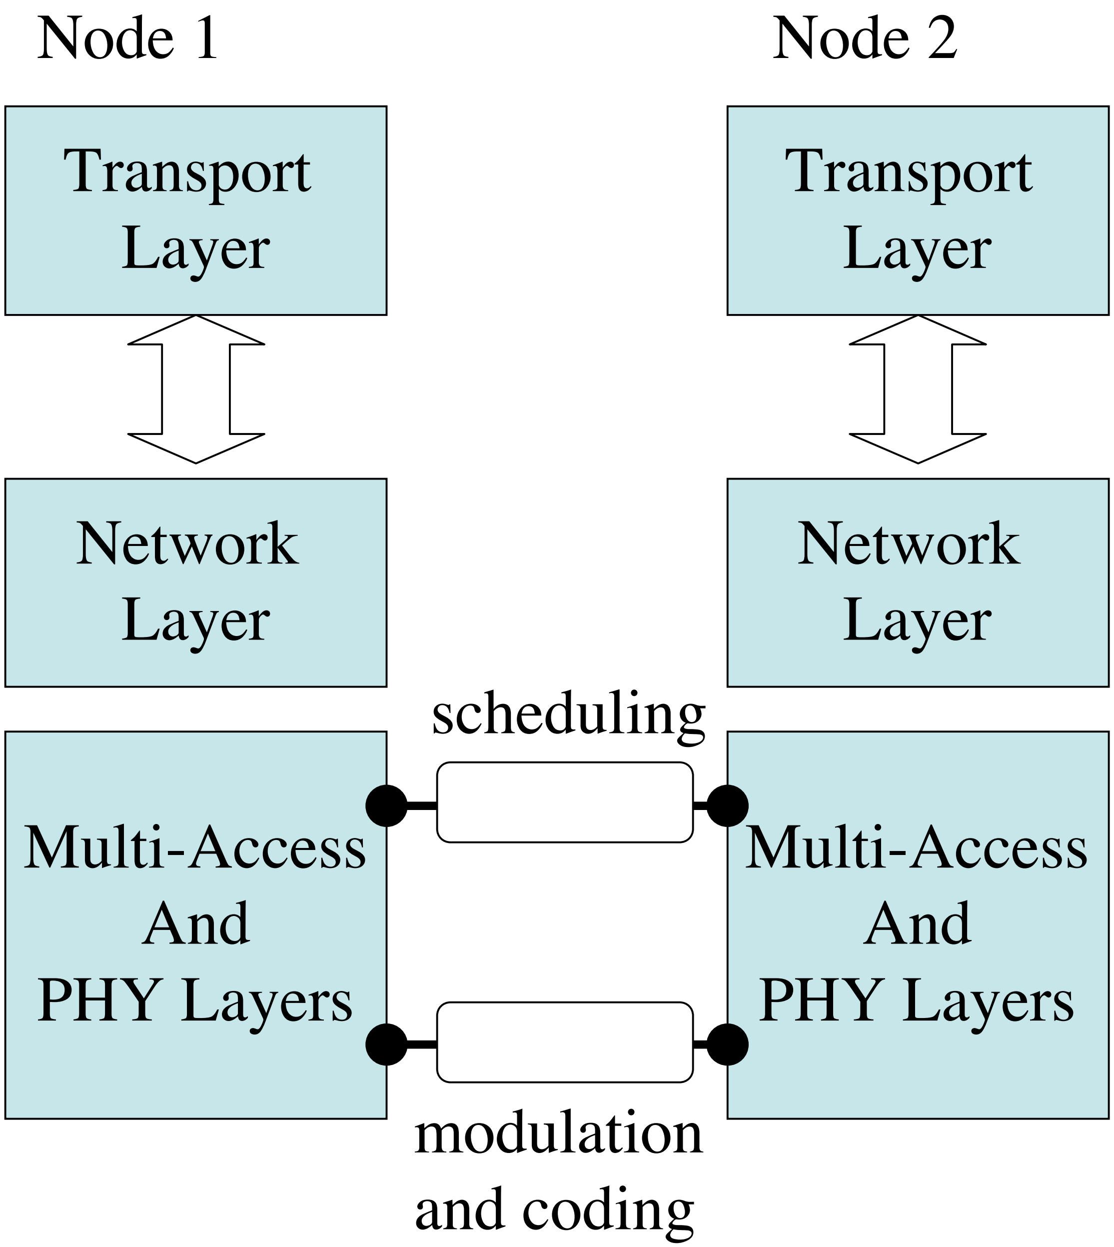

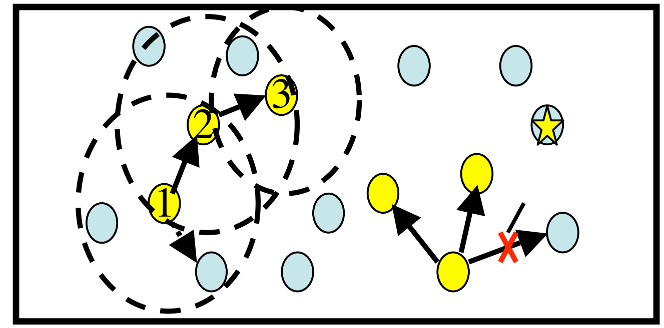

optimization approaches:
[1997-2005]

Kelly, Malloo, Tan
Low
Xiao, Johansson, Boyd
Julian, O'Neill
Chiang
Cruz, Santhanam
Ordonez, Krishnamachari
Lee, Mazumdar, Lin
Shroff
Neely, Modiano

Cross-Layer Networking

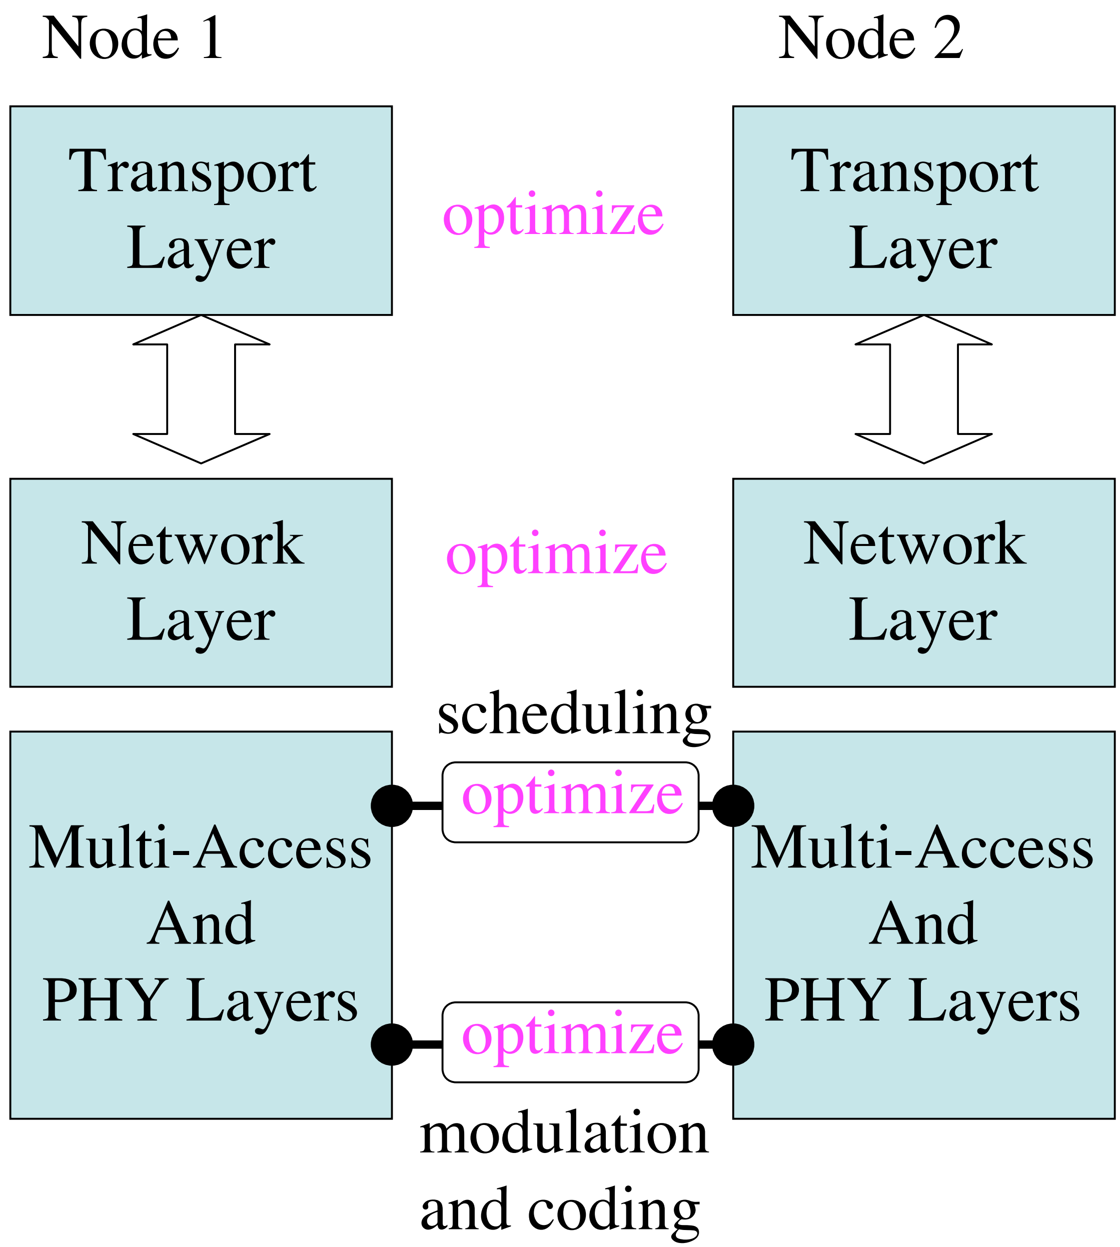

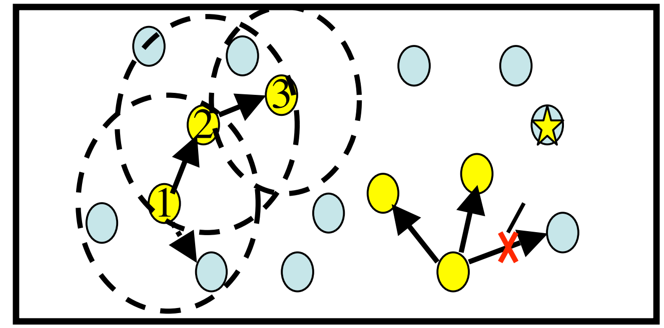

optimization approaches:
[1997-2005]
Kelly, Malloo, Tan
Low
Xiao, Johansson, Boyd
Julian, O'Neill
Chiang
Cruz, Santhanam
Ordonez, Krishnamachari
Lee, Mazumdar, Lin
Shroff
Neely, Modiano

Cross-Layer Networking

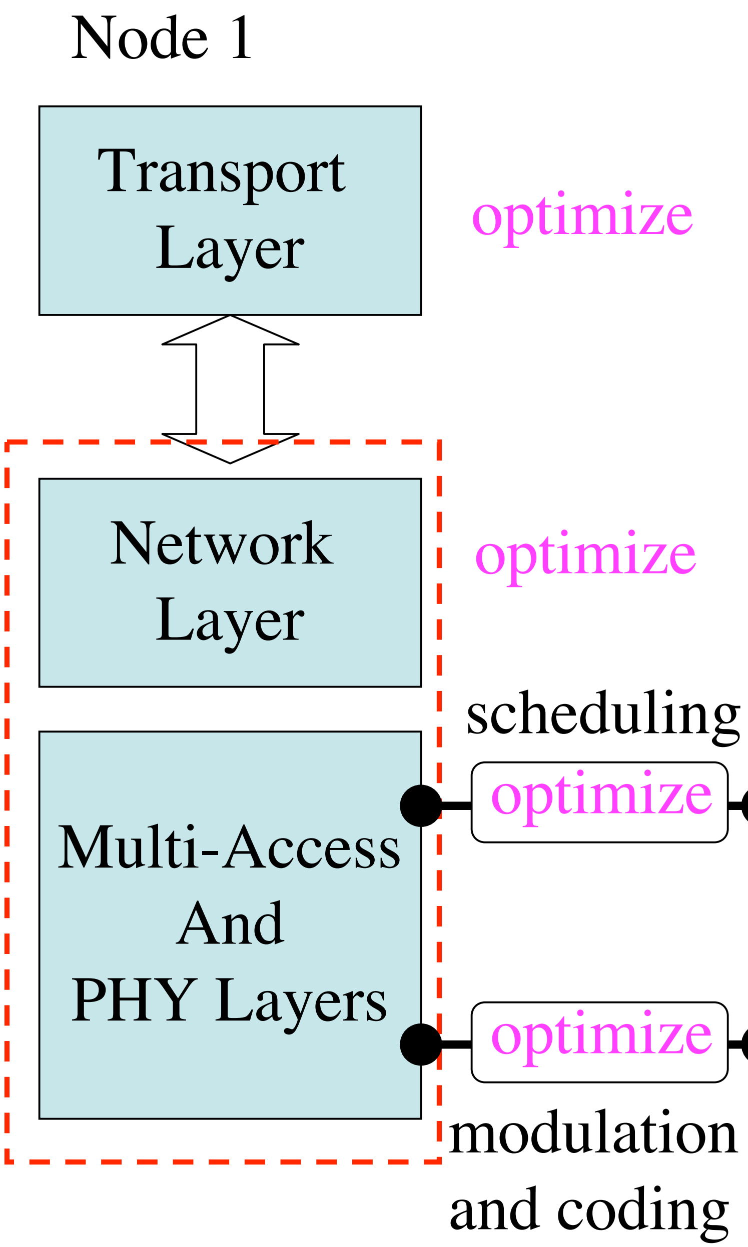

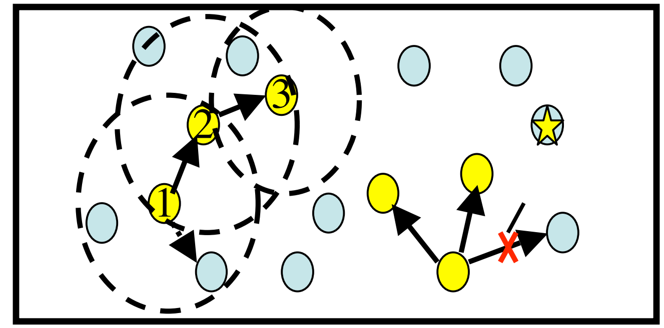

Networking,
Multiple Access,
PHY Layers

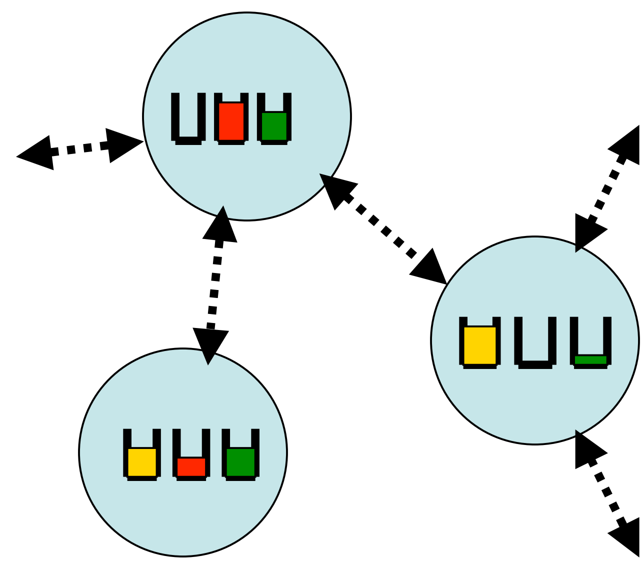

Cross-Layer Networking

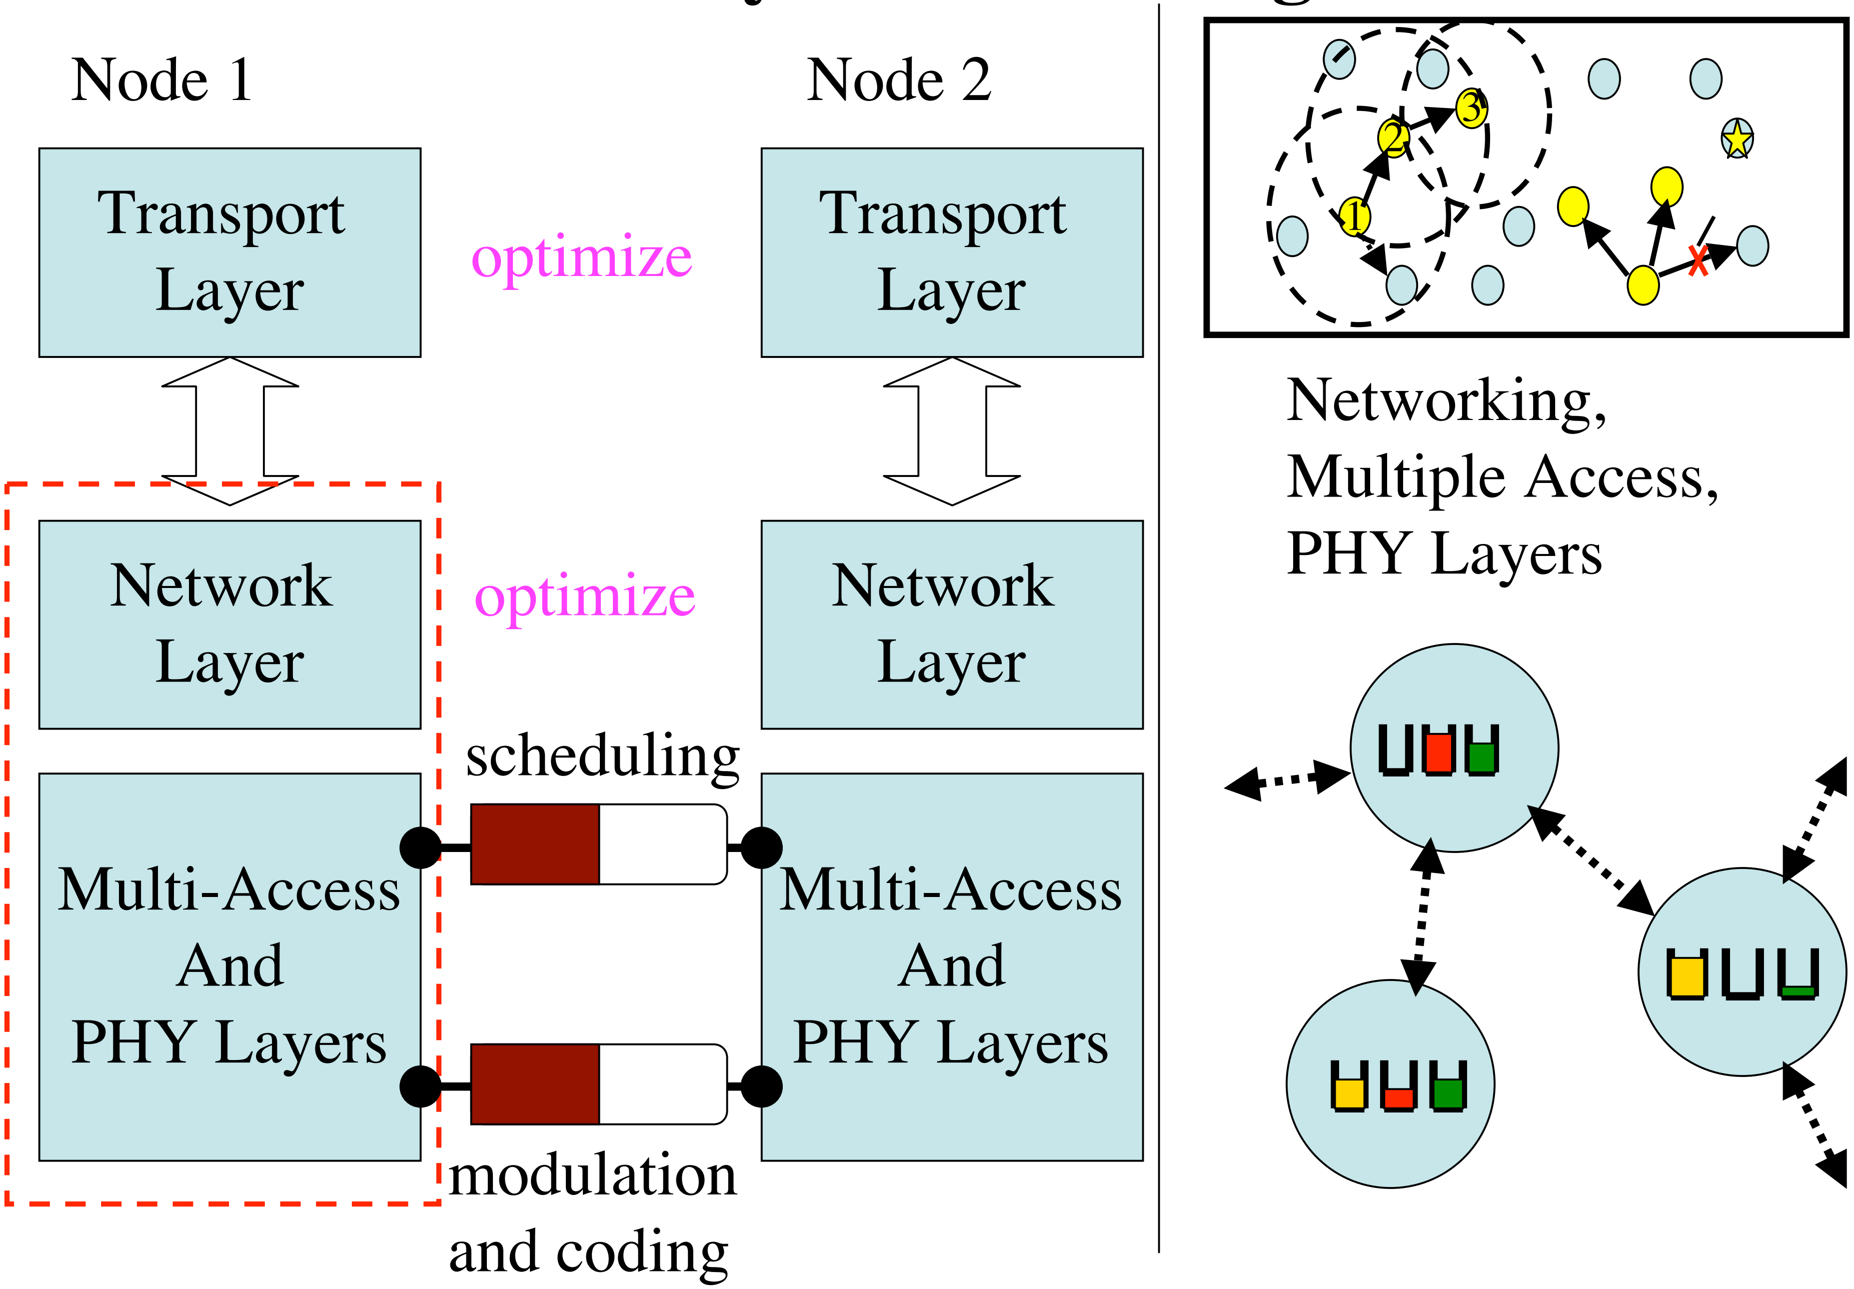

Cross-Layer Networking

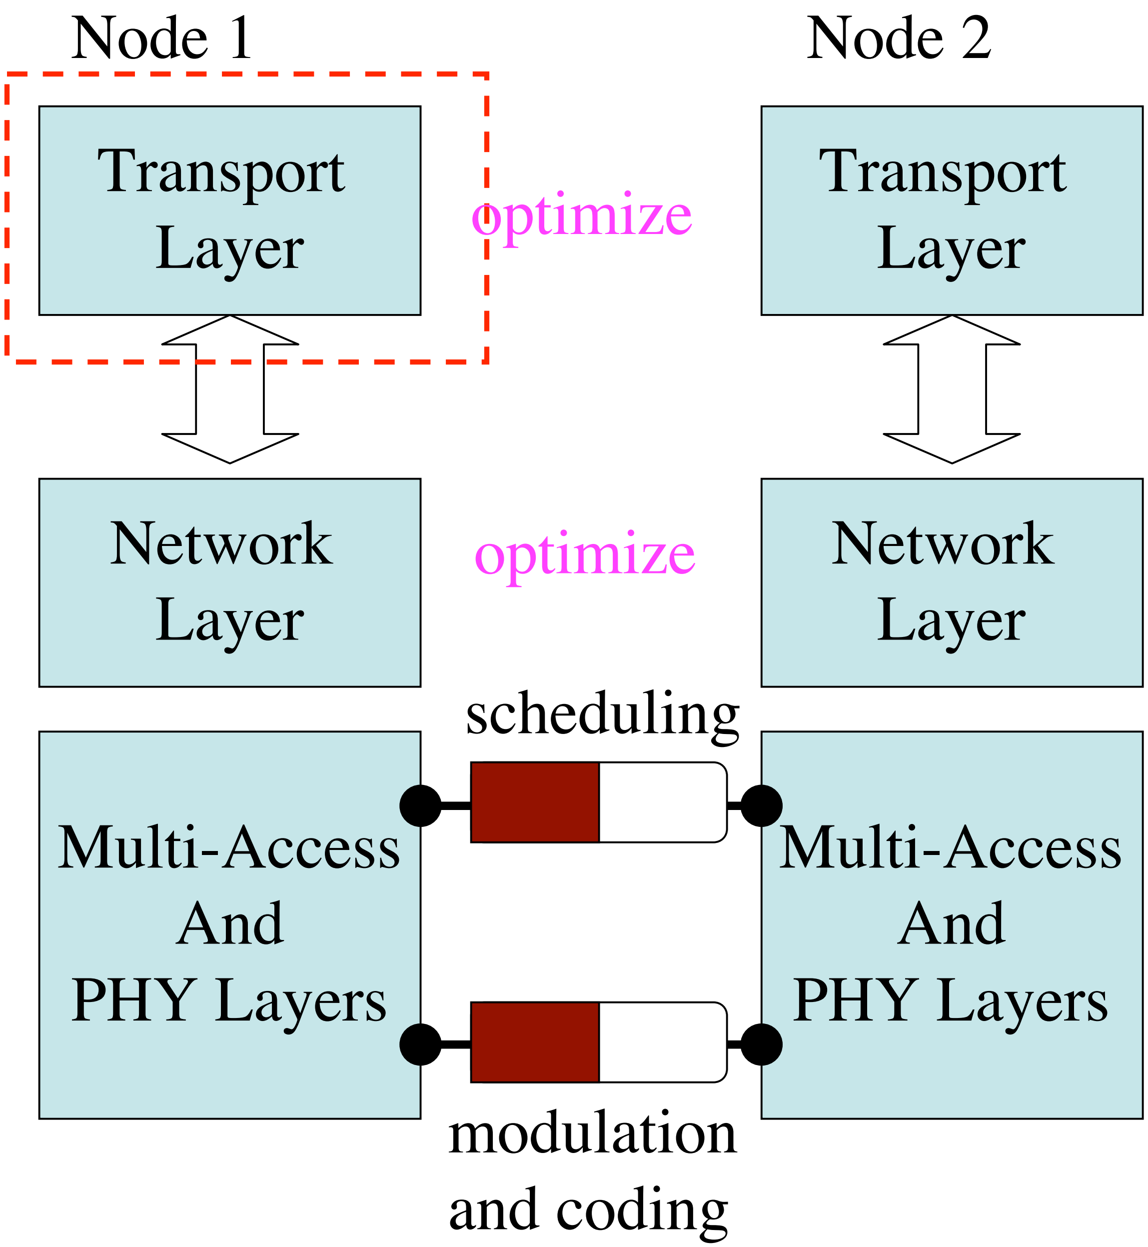

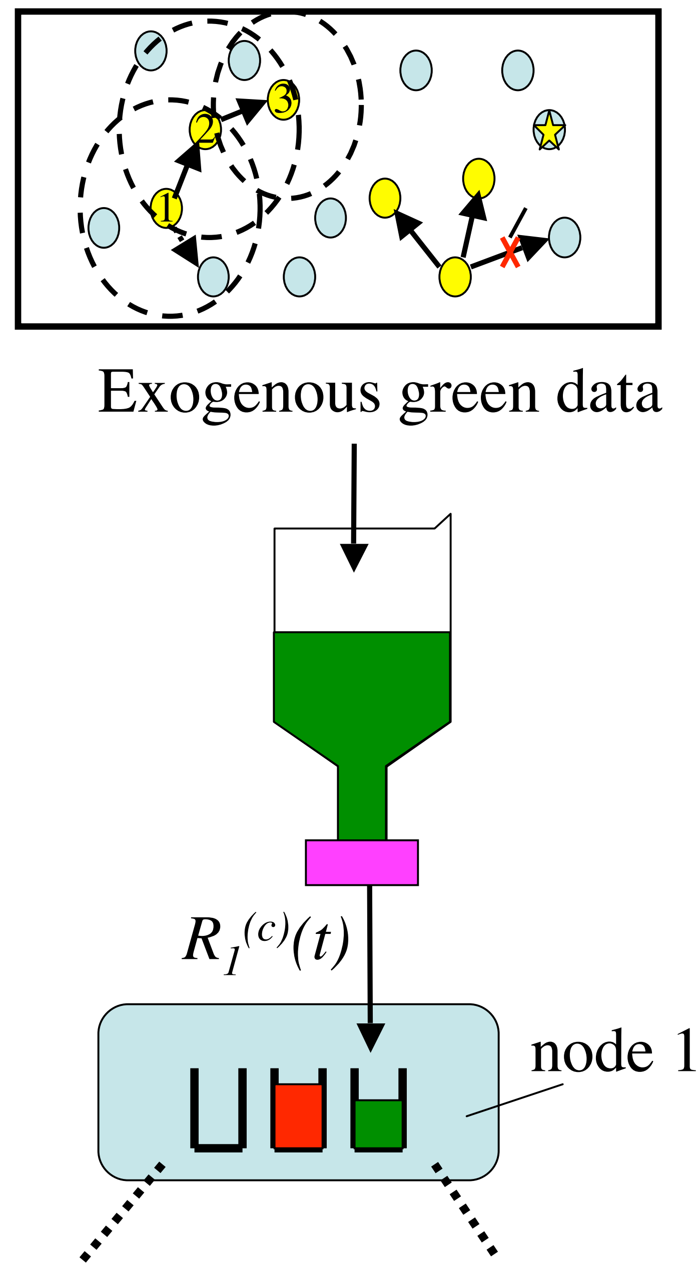

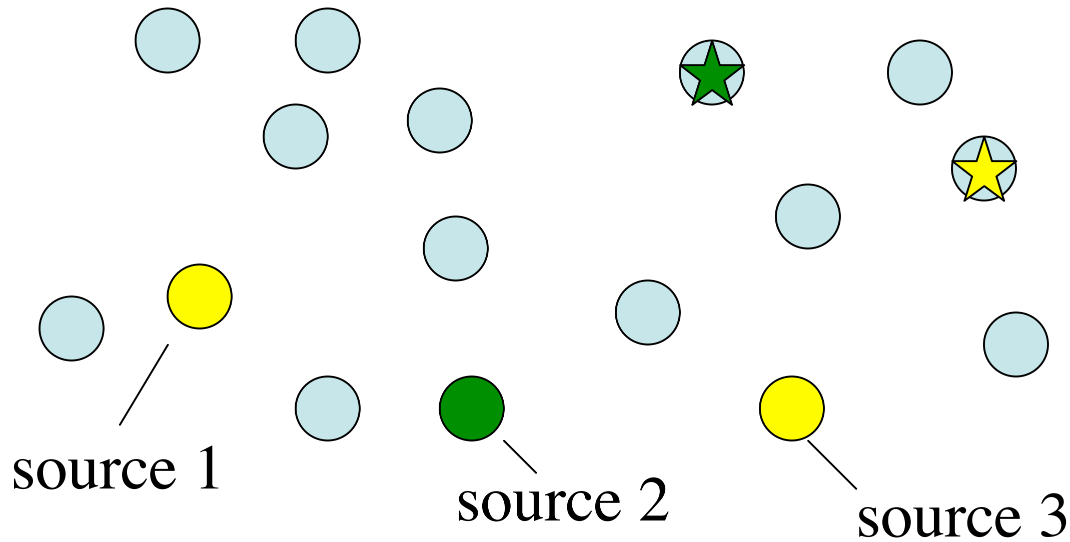

-Multi-Node Wireless Network (possibly mobile)
-Operates in slotted time (t = 0, 1, 2, …)
-Broadcast Advantage, Channel Errors
-Time Varying Transmission Success
Probabilities $q_{ab}(t)$

Example: Suppose Source 1 transmits…

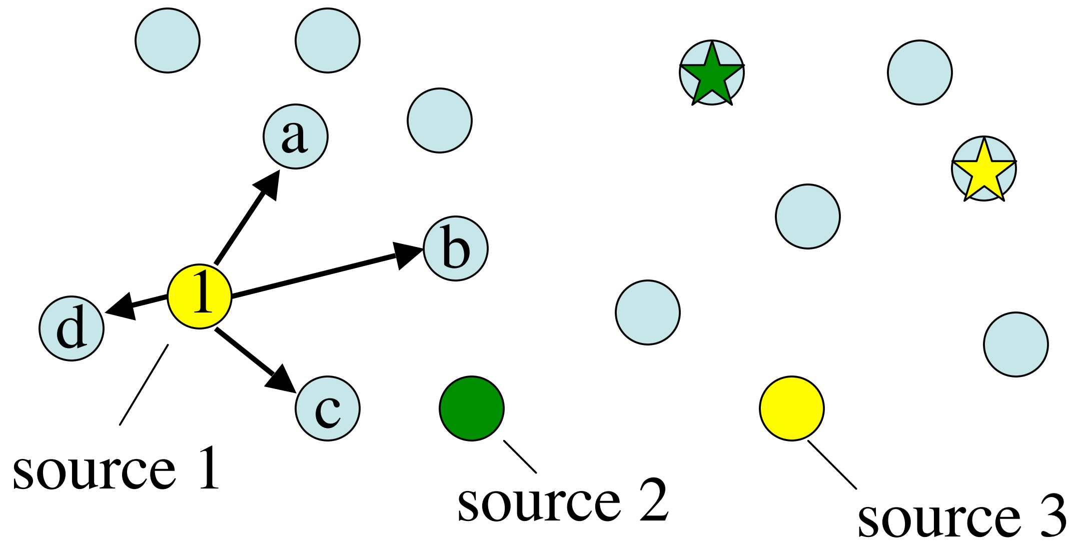

-Multi-Node Wireless Network (possibly mobile)
-Operates in slotted time (t = 0, 1, 2, …)
-Broadcast Advantage, Channel Errors
-Time Varying Transmission Success
Probabilities $q_{ab}(t)$

Example: Suppose Source 1 transmits…

There are 4 possible receivers…

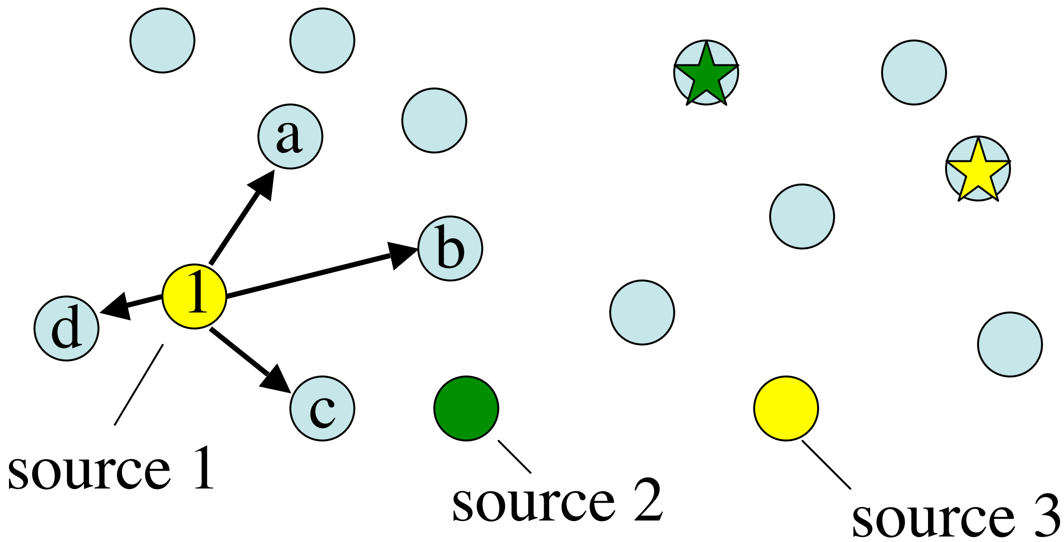

-Multi-Node Wireless Network (possibly mobile)
-Operates in slotted time ($t = 0, 1, 2, \ldots$)
-Broadcast Advantage, Channel Errors
-Time Varying Transmission Success
Probabilities $q_{ab}(t)$

Example: Suppose Source 1 transmits…

Each with different success probs…

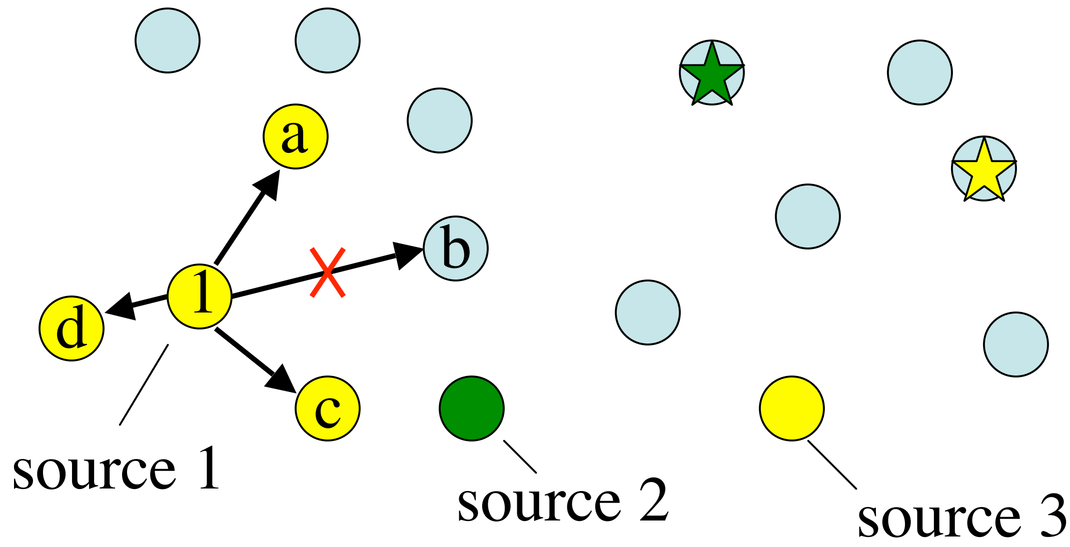

-Multi-Node Wireless Network (possibly mobile)
-Operates in slotted time (t = 0, 1, 2, …)
-Broadcast Advantage, Channel Errors
-Time Varying Transmission Success
 Probabilities $q_{ab}(t)$

Example: Suppose Source 1 transmits…

Only 3 successfully receive…

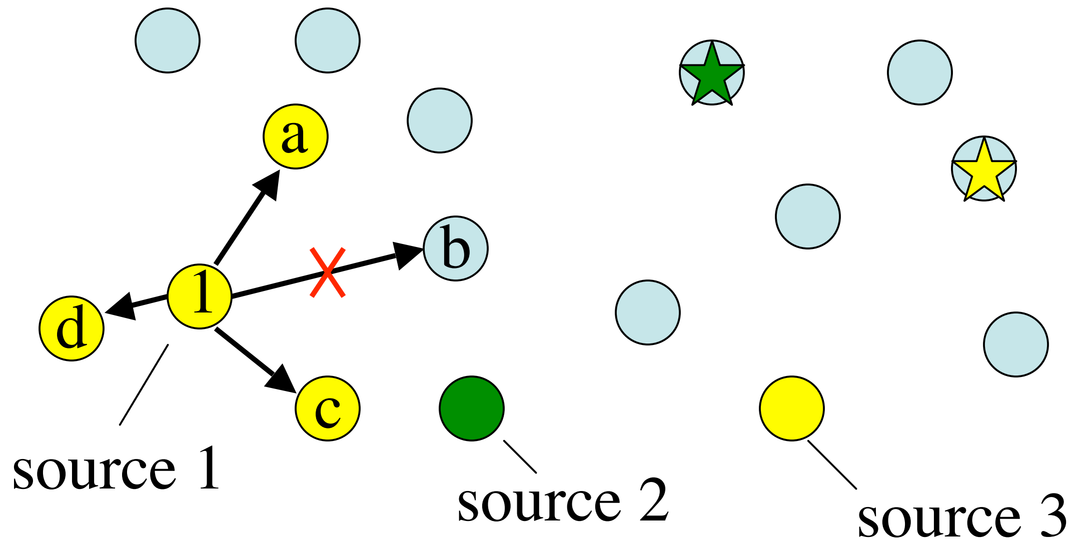

-Multi-Node Wireless Network (possibly mobile)
-Operates in slotted time (t = 0, 1, 2, …)
-Broadcast Advantage, Channel Errors
-Time Varying Transmission Success
Probabilities $q_{ab}(t)$

__Multi-Receiver Diversity!__

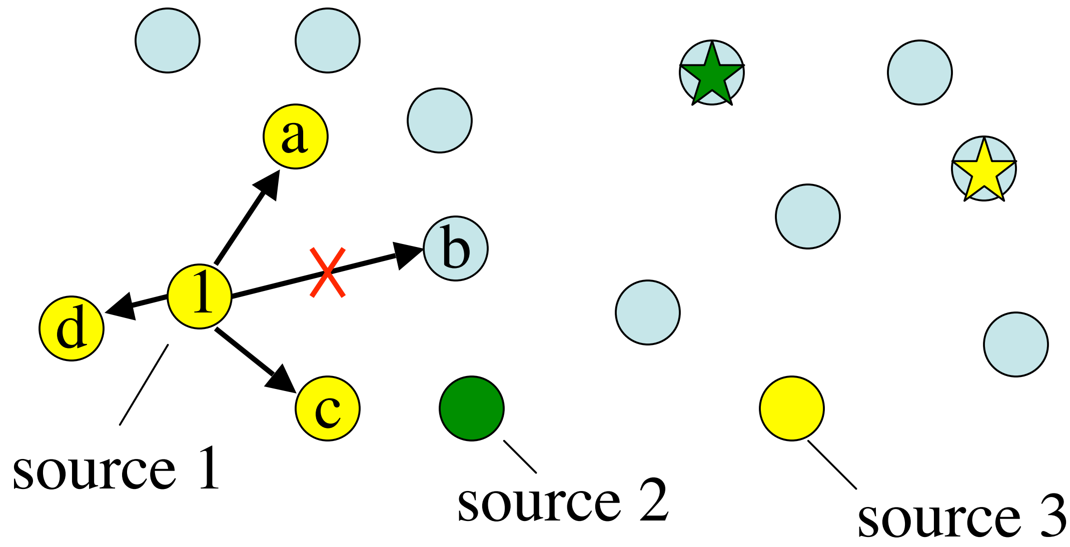

Fundamental Questions:

1) How to Fully Utilize Multi-Receiver Diversity?

2) How to Maximize Throughput? Minimize Av. Power?

3) How to choose which node takes charge of the packet?

4) Should we allow redundant forwarding of different copies of the same packet?

5) How to schedule multiple traffic streams?

GeRaF:

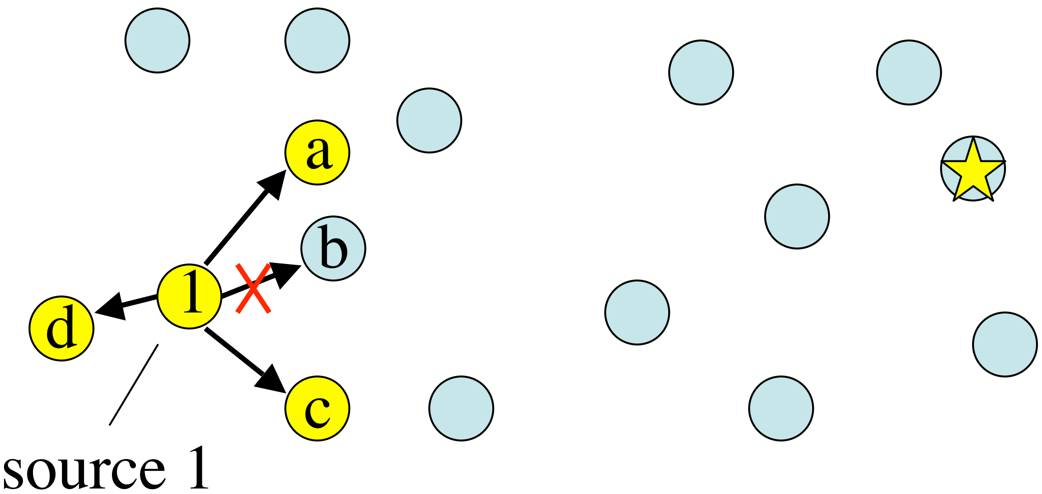

A Hot Topic Area:

Zorzi and Rao: “Geographic Random Forwarding” (GeRaF) [IEEE Trans. on Mobile Computing, 2003].

Biswas and Morris: “Extremely Opportunistic Routing” (EXOR) [Proc. of Sigcomm, 2005].

Baccelli, et. al. [IEEE Trans. Information Theory 2006]

GeRaF:

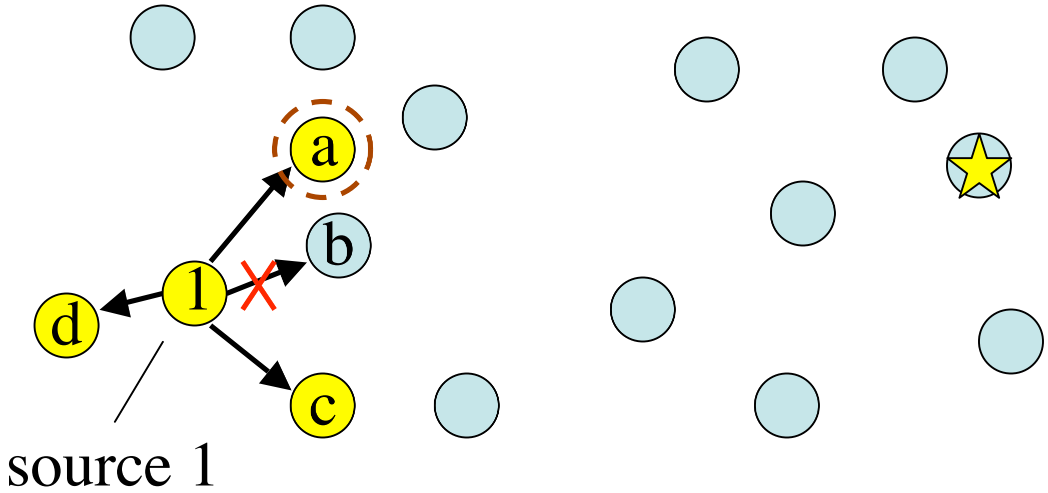

A Hot Topic Area:

- Zorzi and Rao: "Geographic Random Forwarding" (GeRaF) [IEEE Trans. on Mobile Computing, 2003].

  "*Closest-to-Destination*" Heuristic

Biswas and Morris: "Extremely Opportunistic Routing" (EXOR) [Proc. of Sigcomm, 2005].

Baccelli, et. al. [IEEE Trans. Information Theory 2006]

GeRaF:

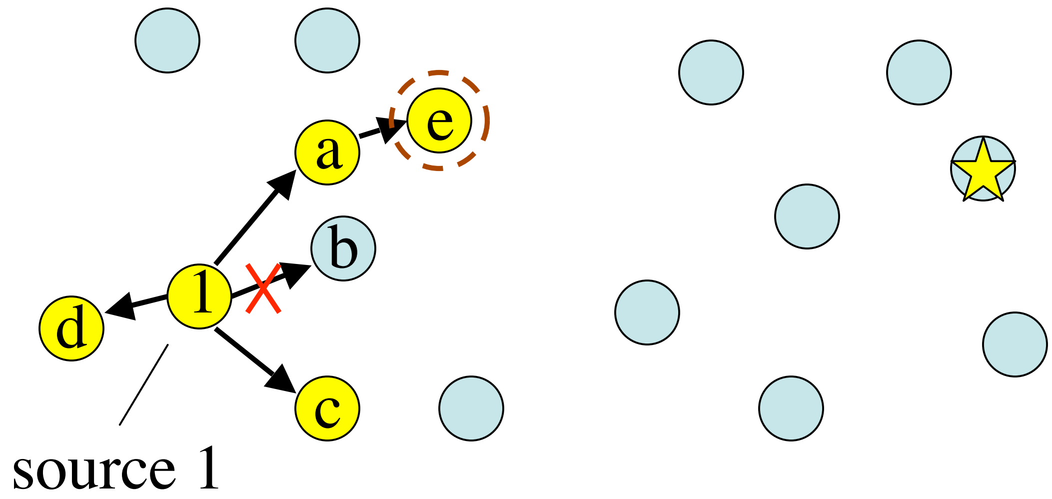

A Hot Topic Area:

- Zorzi and Rao: "Geographic Random Forwarding" (GeRaF) [IEEE Trans. on Mobile Computing, 2003].

  "*Closest-to-Destination*" Heuristic

Biswas and Morris: "Extremely Opportunistic Routing" (EXOR) [Proc. of Sigcomm, 2005].

Baccelli, et. al. [IEEE Trans. Information Theory 2006]

GeRaF:

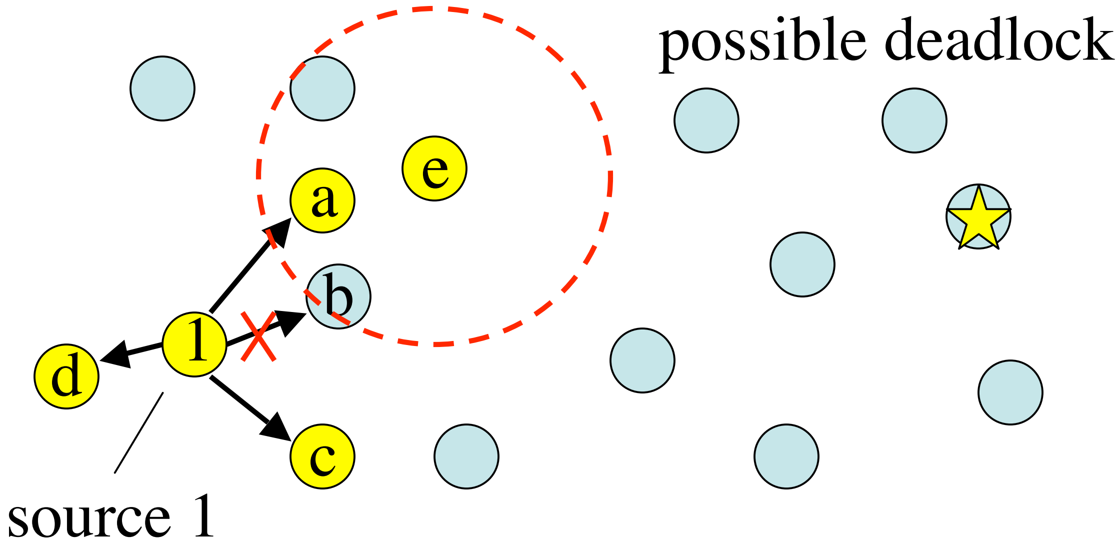

A Hot Topic Area:

- Zorzi and Rao: "Geographic Random Forwarding" (GeRaF) [IEEE Trans. on Mobile Computing, 2003].

  "*Closest-to-Destination*" Heuristic

Biswas and Morris: "Extremely Opportunistic Routing" (EXOR) [Proc. of Sigcomm, 2005].

Baccelli, et. al. [IEEE Trans. Information Theory 2006]

EXOR:

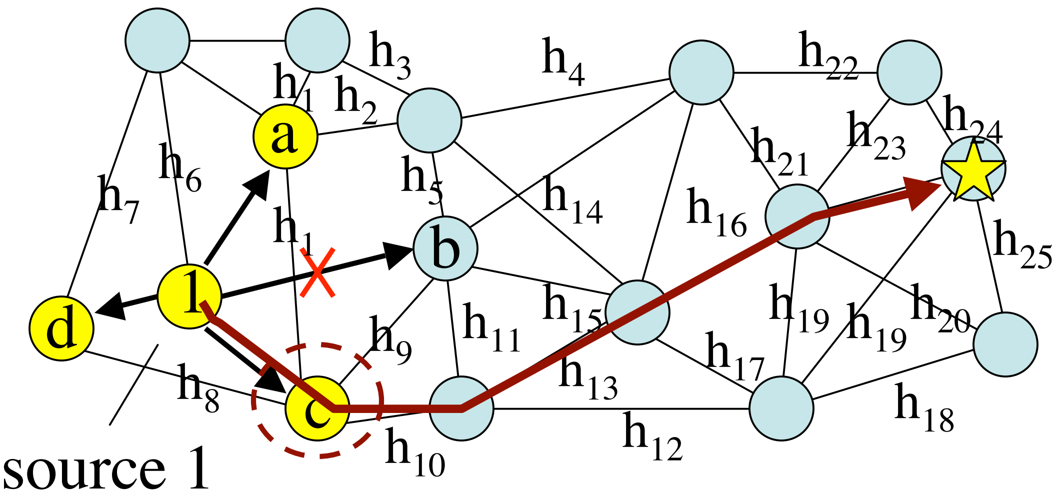

A Hot Topic Area:

Zorzi and Rao: "Geographic Random Forwarding" (GeRaF) [IEEE Trans. on Mobile Computing, 2003].

"*Closest-to-Destination*" Heuristic

- Biswas and Morris: "Extremely Opportunistic Routing" (EXOR) [Proc. of Sigcomm, 2005].

"*Fewest Expected Hops to Destination*" Heuristic
(using a traditional shortest path based on error probs)

How to achieve thruput and energy optimal routing?

A Big Challenge: Complexity!

Example: Suppose a node transmits a packet, and there are k potential receivers…

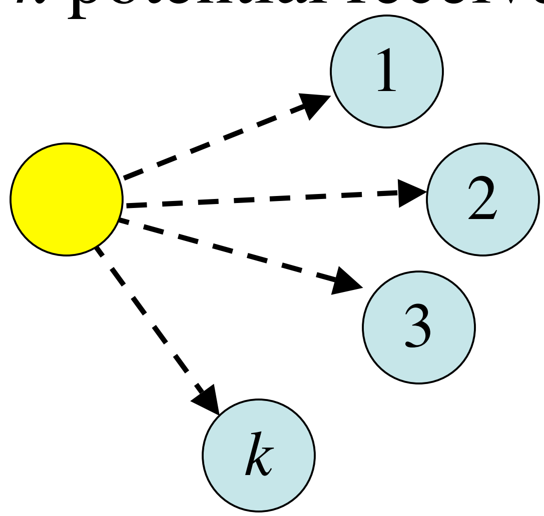

Then there are 2^k possible outcomes. An optimal algorithm must specify a contingency plan for each possible outcome.

How to achieve thruput and energy optimal routing?

A Big Challenge: Complexity!
Example: Suppose a node transmits a packet, and there are k potential receivers…

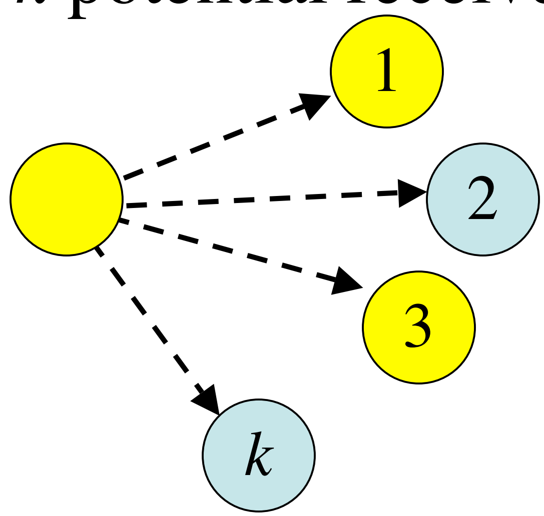

Then there are 2^k possible outcomes. An optimal algorithm must specify a contingency plan for each possible outcome.

How to achieve thruput and energy optimal routing?

A Big Challenge: Complexity!

Example: Suppose a node transmits a packet, and there are k potential receivers…

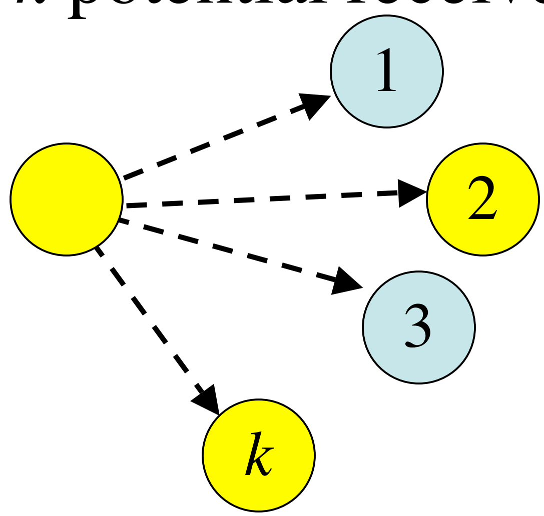

Then there are 2^k possible outcomes. An optimal algorithm must specify a contingency plan for each possible outcome.

How to achieve thruput and energy optimal routing?

A Big Challenge: Complexity!
Example: Suppose a node transmits a packet, and there are k potential receivers…

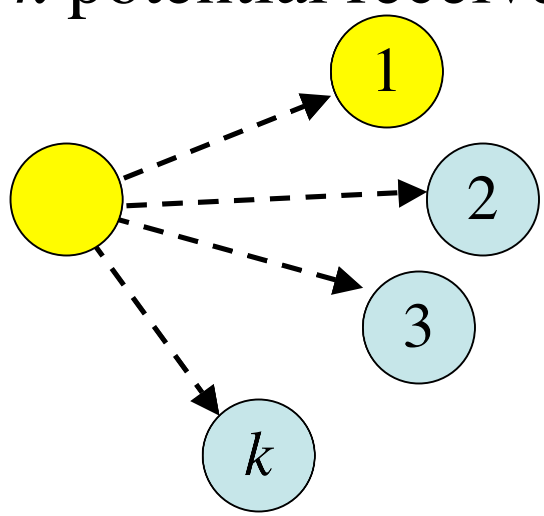

Then there are 2^k possible outcomes. An optimal algorithm must specify a contingency plan for each possible outcome.

How to achieve thruput and energy optimal routing?

A Big Challenge: Complexity!
Example: Suppose a node transmits a packet, and there are k potential receivers…

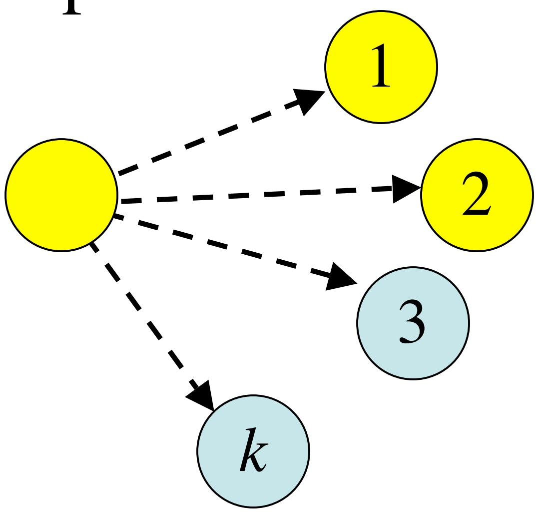

Then there are 2^k possible outcomes. An optimal algorithm must specify a contingency plan for each possible outcome.

How to achieve thruput and energy optimal routing?

<u>A Big Challenge</u>: Complexity!

<u>Example</u>: Suppose a node transmits a packet, and there are k potential receivers…

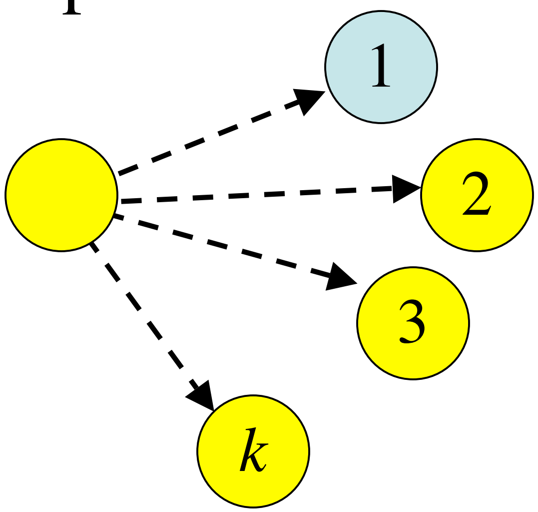

Then there are 2^k possible outcomes. An optimal algorithm must specify a <u>contingency plan</u> for each possible outcome.

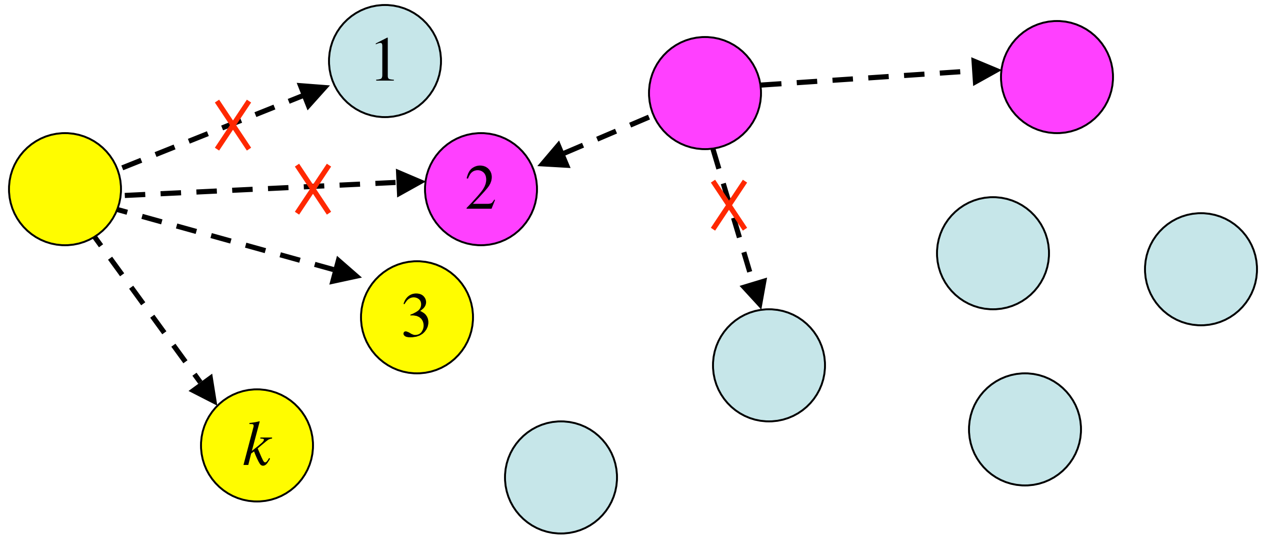

<u>Further Challenges</u>:

1) How to Handle Multiple Simultaneous Transmissions?
2) How to Handle Multiple Traffic Sessions?
3) How to Handle Mobility and/or Time Varying Channel Success Probabilities?

Our Main Results: (Algorithm **DIVBAR**)

1. Show that redundant packet forwarding is not necessary for optimal routing.

2. Achieve Thruput and Energy Optimality via a simple ***Backpressure Index*** between neighboring nodes.

3. **DIVBAR**: "Diversity Backpressure Routing." Distributed alg. Uses local link success probability info.

4. Admits a ***Channel Blind Transmission Mode*** (channel probs. not needed) in special case of single commodity networks and when power optimization is neglected.

The Seminal Paper on Backpressure Routing for Multi-Hop Queueing Networks:

L. Tassiulas, A. Ephremides [IEEE Trans. Aut. Contr. 1992]

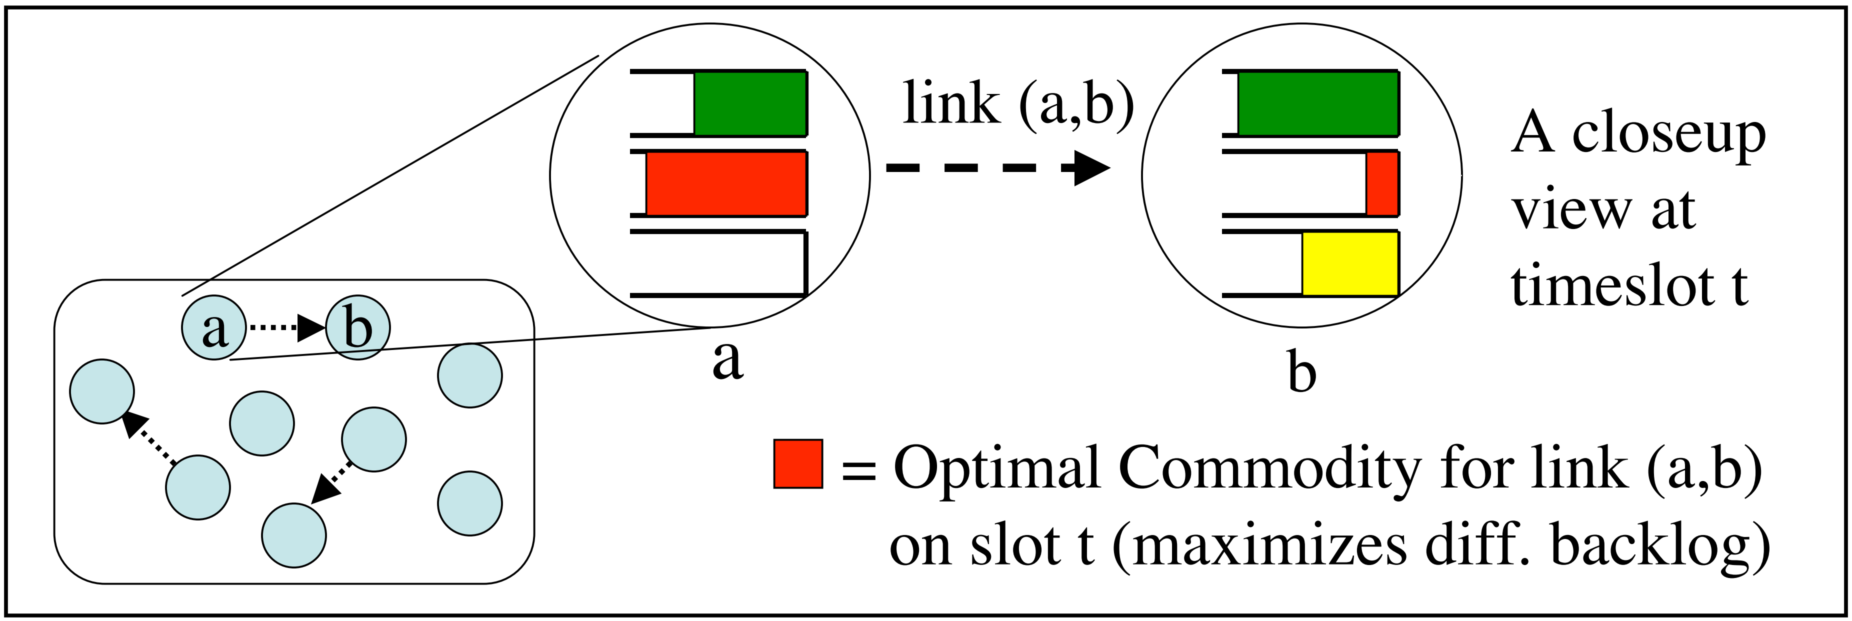

Fundamental Results of Tassiulas-Ephremides [92]:

a. Dynamic Routing via Differential Backlog
b. Max Weight Matchings
c. Stability Analysis via Lyapunov Drift

A brief history of **Lyapunov Drift** for Queueing Systems:

Lyapunov Stability:

Tassiulas, Ephremides [91, 92, 93]
P. R. Kumar, S. Meyn [95]
McKeown, Anantharam, Walrand [96, 99]
Kahale, P. E. Wright [97]
Andrews, Kumaran, Ramanan, Stolyar, Whiting [2001]
Leonardi, Mellia, Neri, Marsan [2001]
Neely, Modiano, Rohrs [2002, 2003, 2005]

Lyapunov Stability with Stochastic Performance Optimization:

Neely, Modiano [2003, 2005] (Fairness, Energy)
Georgiadis, Neely, Tassiulas [NOW Publishers, F&T, 2006]

Alternate Approaches to Stoch. Performance Optimization:

Eryilmaz, Srikant [2005] (Fluid Model Transformations)
Stolyar [2005] (Fluid Model Transformations)
Lee, Mazumdar, Shroff [2005] (Stochastic Gradients)

Problem Formulation:

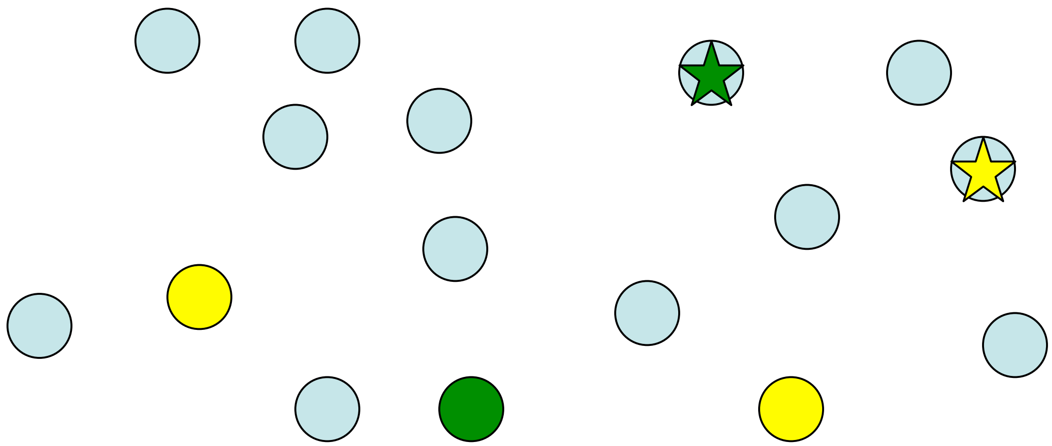

1. Slotted Time ($t = \{0, 1, 2, \ldots\}$)
2. Can transmit 1 packet (power P_{tran}) or else idle.
3. Traffic: $A_i^c(t)$ i.i.d. over slots, rates $E[A_i^c(t)] = \lambda_i^c$
4. Topology state process $S(t)$:
   - Transmission opportunities: $\chi_i(t) = \hat{\chi}_i(S(t)) \in \{0, 1\}$

     (Pre-specified MAC: $\chi_i(t) = 1$ $\Rightarrow$ node i can transmit 1 packet)
   - Channel Probabilities: $q_{i,\Omega}(t) = \hat{q}_{i,\Omega}(S(t))$

     (Ω = A particular subset of receivers)

Problem Formulation:

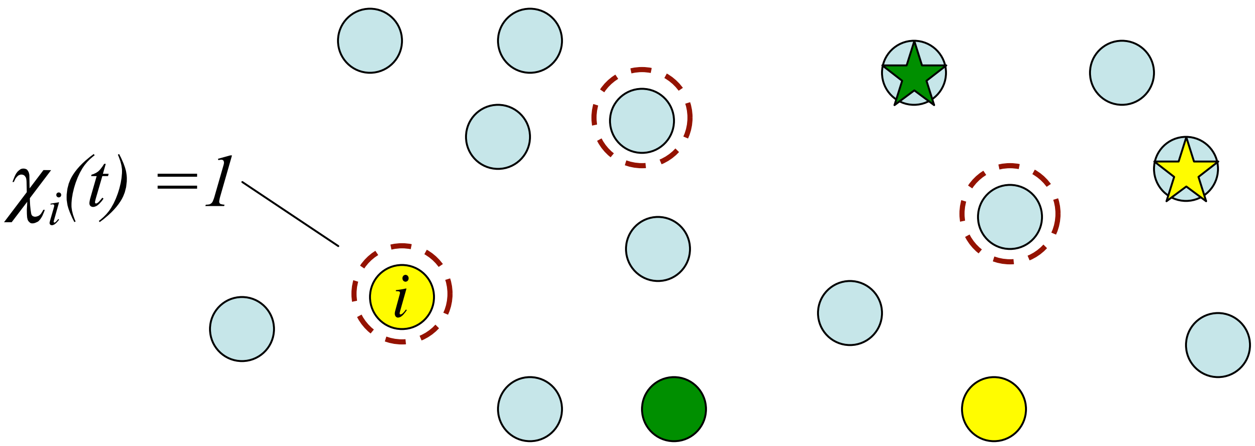

1. Slotted Time ($t = \{0, 1, 2, \ldots\}$)
2. Can transmit 1 packet (power P_{tran}) or else idle.
3. Traffic: $A_i^c(t)$ i.i.d. over slots, rates $E[A_i^c(t)] = \lambda_i^c$
4. Topology state process $S(t)$:
   - Transmission opportunities: $\chi_i(t) = \hat{\chi}_i(S(t)) \in \{0, 1\}$

     (Pre-specified MAC: $\chi_i(t) = 1$ $\Rightarrow$ node i can transmit 1 packet)

   - Channel Probabilities: $q_{i,\Omega}(t) = \hat{q}_{i,\Omega}(S(t))$

     (Ω = A particular subset of receivers)

Problem Formulation:

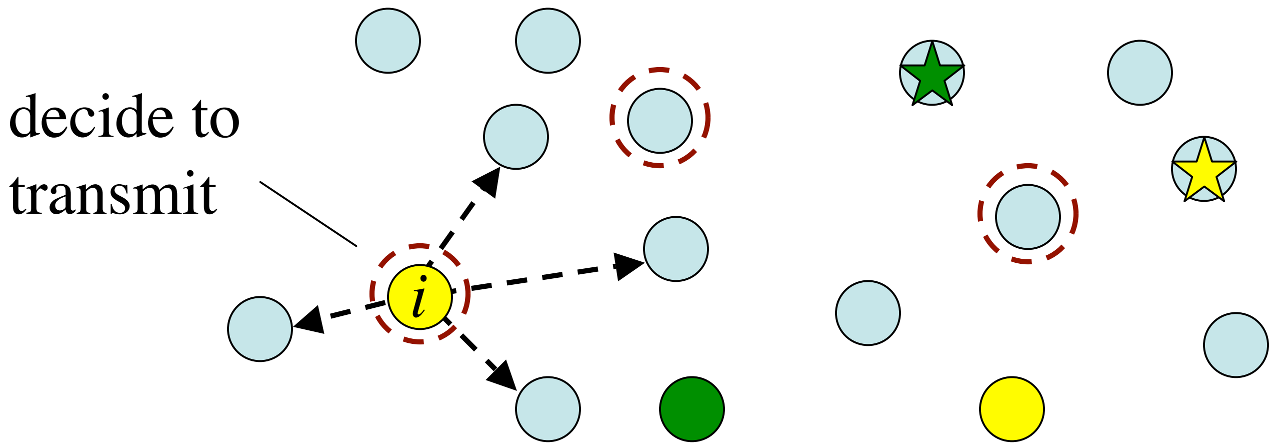

1. Slotted Time ($t = \{0, 1, 2, \ldots\}$)
2. Can transmit 1 packet (power P_{tran}) or else idle.
3. Traffic: $A_i^c(t)$ i.i.d. over slots, rates $E[A_i^c(t)] = \lambda_i^c$
4. Topology state process $S(t)$:
   - Transmission opportunities: $\chi_i(t) = \hat{\chi}_i(S(t)) \in \{0, 1\}$

     (Pre-specified MAC: $\chi_i(t) = 1$ $\Rightarrow$ node i can transmit 1 packet)
   - Channel Probabilities: $q_{i,\Omega}(t) = \hat{q}_{i,\Omega}(S(t))$

     (Ω = A particular subset of receivers)

Problem Formulation:

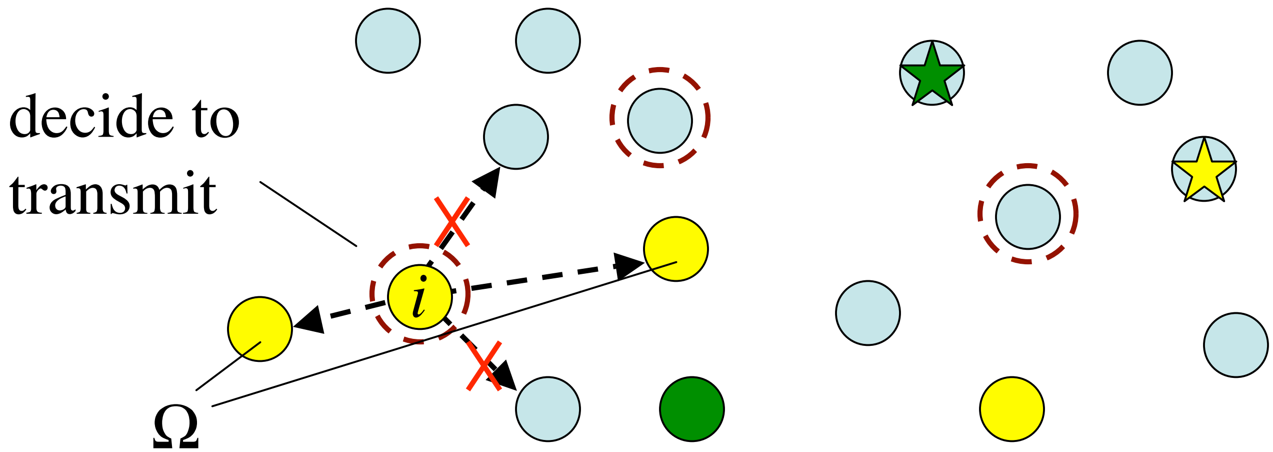

1. Slotted Time ($t = \{0, 1, 2, \ldots\}$)
2. Can transmit 1 packet (power P_{tran}) or else idle.
3. Traffic: $A_i^c(t)$ i.i.d. over slots, rates $E[A_i^c(t)] = \lambda_i^c$
4. Topology state process $S(t)$:
   - Transmission opportunities: $\chi_i(t) = \hat{\chi}_i(S(t)) \in \{0, 1\}$

     (Pre-specified MAC: $\chi_i(t) = 1$ ➡ node i can transmit 1 packet)
   - Channel Probabilities: $q_{i,\Omega}(t) = \hat{q}_{i,\Omega}(S(t))$

     (Ω = A particular subset of receivers)

Anatomy of a Single Timeslot:

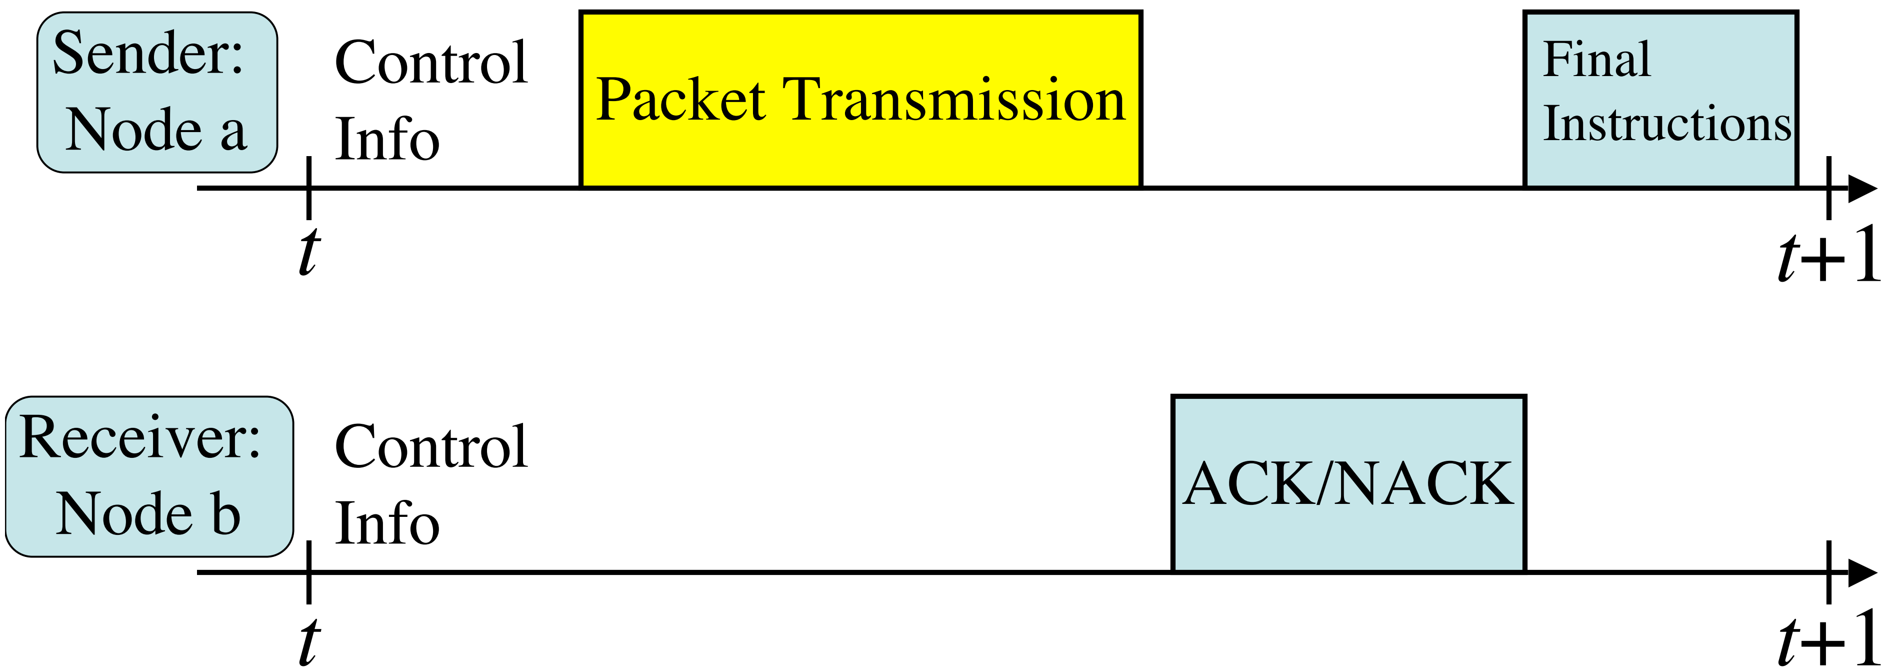

Idealistic Assumptions:

-No errors on control channels.

-After a packet transmission, the "handshake" enables the transmitter to know the successful recipients.

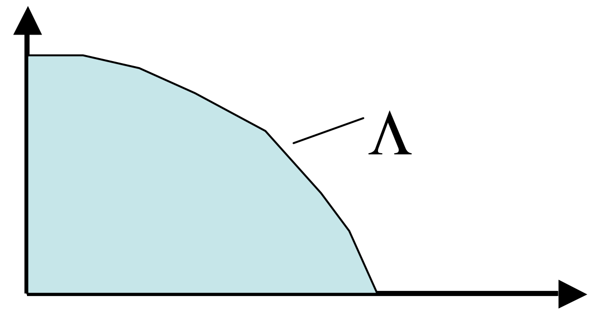

Definition: The **network layer capacity region Λ** is the set of all rate matrices (λ_i^c) that can be stably supported, considering all possible routing/scheduling algorithms that conform to the network model (possibly forwarding multiple copies of the same packet).

Lemma: The capacity region (and minimum avg. energy) can be achieved without redundant packet forwarding.

Note: Our network model does not include:
-Signal enhancement via cooperative communication
-Network coding

(Network capacity can be increased by extending the valid control actions to include such options).

<u>Theorem 1</u>: (Network Capacity and Minimum Avg. Energy)

(a) Network Capacity Region Λ is given by all (λ_i^c) such that:

There exist variables:

$$\boxed{\{f_{nk}^{(c)}\},\ \{\alpha_n^{(c)}(s)\},\ \{\theta_{nk}^{(c)}(\Omega_n)\},\ (\text{for all } n, k, c,\ s \in \mathcal{S},\ \Omega_n)}$$

Such that:

$$f_{ab}^{(c)} \geq 0\ ,\ f_{cb}^{(c)} = 0\ ,\ f_{aa}^{(c)} = 0$$

$$\sum_a f_{an}^{(c)} + \lambda_n^{(c)} \leq \sum_b f_{nb}^{(c)} \quad \text{for all } n \neq c$$

$$\sum_c f_{nk}^{(c)} \leq \sum_c \sum_{s \in \mathcal{S}} \pi_s \alpha_n^{(c)}(s) \left[\sum_{\Omega_n \in \mathcal{H}_n} \hat{q}_{n,\Omega_n}(s) \theta_{nk}^{(c)}(\Omega_n) \right]$$

$$\theta_{nk}^{(c)}(\Omega_n) = 0 \text{ if } k \notin \{\Omega_n \cup \{n\}\}\ \ ,\ \ \sum_{k=1}^{N} \theta_{nk}^{(c)}(\Omega_n) = 1$$

$$\sum_{c=1}^{N} \alpha_n^{(c)}(s) \leq 1\ ,\ \ \alpha_n^{(c)}(s) = 0 \text{ if } \hat{\chi}_n(s) = 0$$

Theorem 1 part (b): The Minimum Avg. Energy is given by the solution to:

Minimize: $$\sum_{s \in \mathcal{S}} \pi_s \left[\sum_{n=1}^{N} \sum_{c=1}^{N} \alpha_n^{(c)}(s) P_{tran} \right]$$

Subject to: The constraints of part (a)

Note: Just writing down the optimal solution takes an ***Exponential Number of Parameters***!

Parameters Used:

$$\{f_{nk}^{(c)}\}, \{\alpha_n^{(c)}(s)\}, \{\theta_{nk}^{(c)}(\Omega_n)\}$$

for all $n, k, c \in \{1, \ldots, N\}$
for all topology states $s \in \mathcal{S}$
for all subsets Ω_n (subsets of $\{1, 2, \ldots, N\} - \{n\}$)

A Simple Backpressure Solution (in terms of a control parameter V):
Algorithm **DIVBAR** “Diversity Backpressure Routing”

Let $\chi_n(t) = 1$

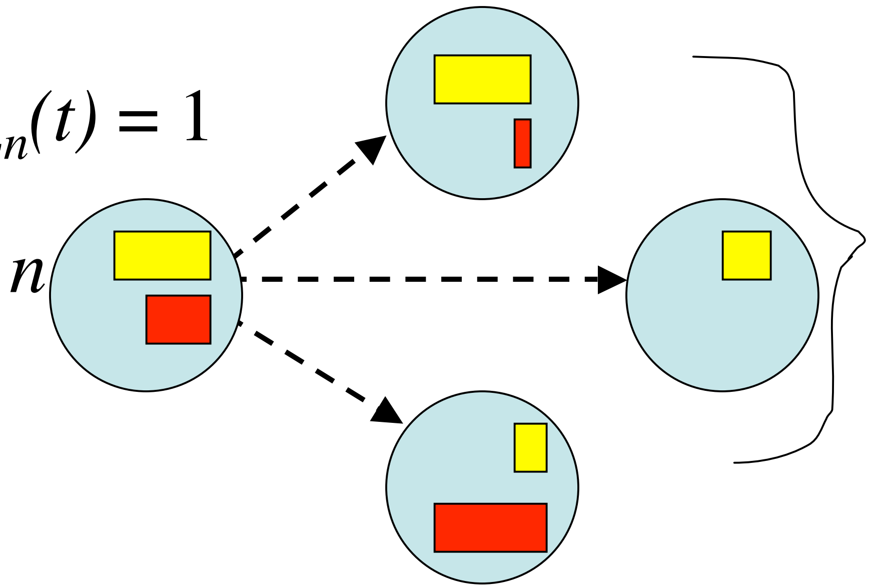

$K_n(t)$ = Set of potential receivers at time t.

1. For each $k \in K_n(t)$, compute $W_{nk}^{(c)}(t)$:

$$\boxed{W_{nk}^{(c)}(t) = \max[U_n^{(c)}(t) - U_k^{(c)}(t), 0]}$$

(Differential Backlog)

($U_k^{(c)}(t)$=# commodity c packets in node n at slot t)

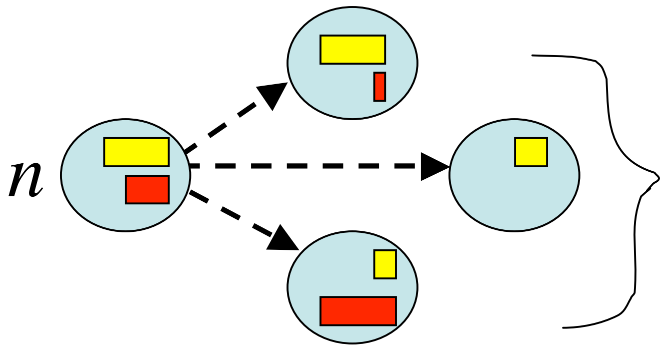

$K_n(t)$ = Set of potential receivers at time t.

---

2. Node n ***rank orders*** its $W_{nk}^{(c)}(t)$ values for all $k \in K_n(t)$:

$$W_{nk(n,c,t,1)}^{(c)}(t) > W_{nk(n,c,t,2)}^{(c)}(t) > W_{nk(n,c,t,3)}^{(c)}(t) > \ldots$$

(where $k(n,c,t,b)$ = bth largest weight in rank ordering)

3. Define $\phi_{nk}^{(c)}(t)$ = *Probability that a packet transmitted by node n (at slot t) is correctly received at node k, but not received by any other nodes with rank order higher than k.*
(for $k \in K_n(t)$)

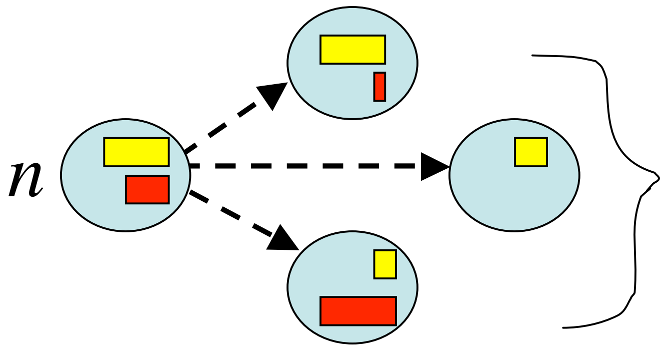

$K_n(t)$ = Set of potential receivers at time t.

---

4. Define the ***optimal commodity*** $c^*_n(t)$ as the maximizer of:

$$\sum_{b=1}^{|\mathcal{K}_n(t)|} W^{(c)}_{n,k(n,c,t,b)}(t)\phi^{(c)}_{n,k(n,c,t,b)}(t)$$

Define $W_n^*(t)$ as the above maximum weighted sum.

5. If $W_n^*(t) > V\, P_{tran}$ then transmit a packet of commodity $c^*_n(t)$. Else, remain idle.

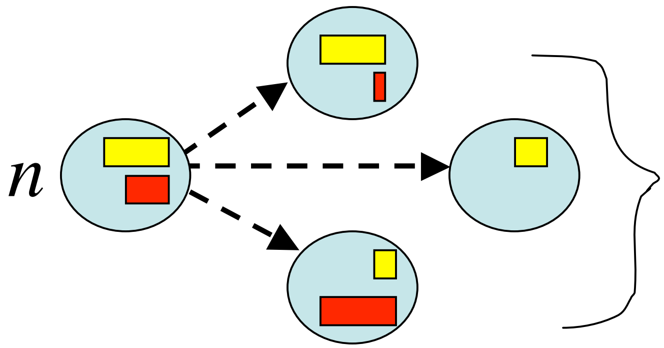

$K_n(t)$ = Set of potential receivers at time t.

Final step of DIVBAR:

<u>If we transmit</u>: After receiving ACK/NACK feedback about successful reception, node n sends a final instruction that ***transfers responsibility of the packet to the receiver with largest differential backlog*** $W_{nk}^{(c^*)}(t)$. If no successful receivers have positive differential backlog, node n retains responsibility for the packet.

<u>Theorem 2</u> (DIVBAR Performance): If arrivals i.i.d. and topology state S(t) i.i.d. over timeslots, and if input rates are strictly interior to capacity region Λ, then implementing DIVBAR for any control parameter V>0 yields:

$$\limsup_{t\rightarrow\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\sum_{n,c}\mathbb{E}\left\{U_n^{(c)}(t)\right\} \leq \frac{N(B+VP_{tran})}{\epsilon_{max}}$$

$$\limsup_{t\rightarrow\infty}\frac{1}{t}\sum_{\tau=0}^{t-1}\sum_{n}\mathbb{E}\left\{P_n(\tau)\right\} \leq P^*_{min} + NB/V$$

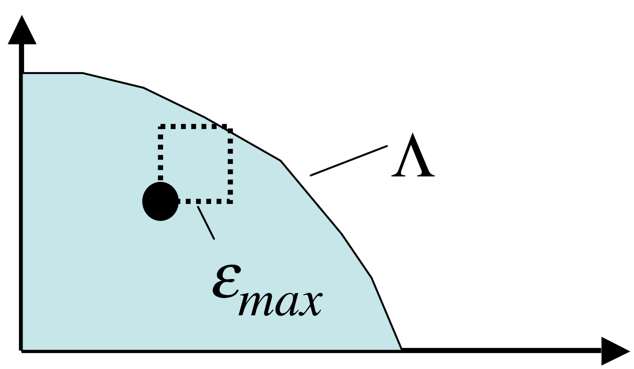

(B = system constant)

Important Special Case…

Channel Blind Transmission:

-One commodity (multiple sources, single sink)

-Neglect Average Power Optimization (set V=0)

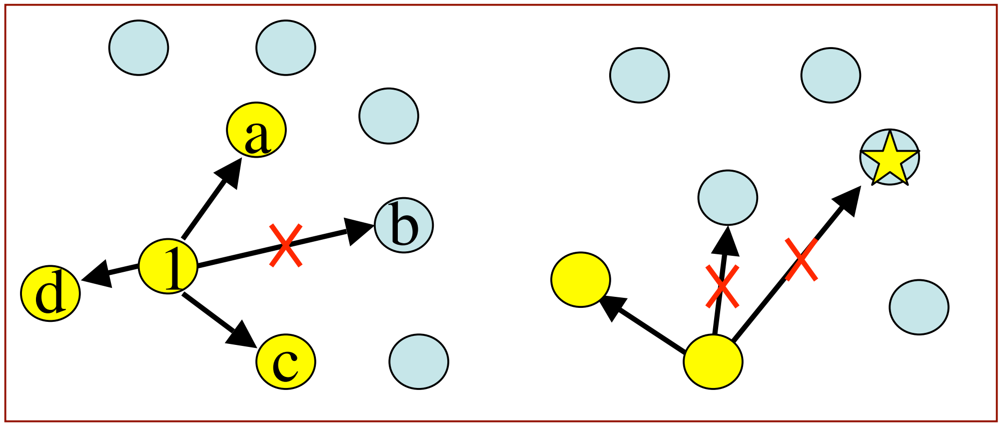

Skip steps 1-5: Just transmit whenever $\chi_n(t)=1$, and transfer responsibility to receiver that maximizes differential backlog. *Achieves throughput optimality without requiring knowledge of (potentially time varying) channel probabilities!*

Extensions:
-Variable Rate and Power Control
-Optimizing the MAC layer

$\boldsymbol{\mu}(t)=(\mu_1(t), \mu_2(t), \ldots, \mu_N(t))$ (# packets transmitted)
$\boldsymbol{P}(t) = (P_1(t), P_2(t), \ldots, P_N(t))$ (Power allocation vector)

$I(t) = (\boldsymbol{\mu}(t); \boldsymbol{P}(t))$ = Collective Control Action

$$q_{n,\, W_n}(t) = \hat{q}_{n,\, W_n}(I(t), S(t))$$

Jointly choose $I(t)$, $c_n^*(t)$ to maximize:

$$\sum_n \left[\left(\sum_{b=1}^{|\mathcal{K}_n(t)|} W^{(c_n^*)}_{n,k(n,c_n^*,t,b)}(t)\hat{\phi}_{n,k(n,c_n^*,t,b)}(I(t), S(t)) \right) - VP_n(t) \right]$$

DIVBAR can easily be integrated with other cross-layer performance objectives using ***<u>stochastic Lyapunov optimization,</u>*** using techniques of *Virtual Power Queues*, *Auxiliary Variables*, *Flow State Queues* developed in:

Flow Control, Fairness, Energy:

[Neely, Modiano 2003, 2005] (fairness, stochastic utility opt.)

[Neely Infocom 2005] (energy optimal control)

[Georgiadis, Neely, Tassiulas NOW 2006]

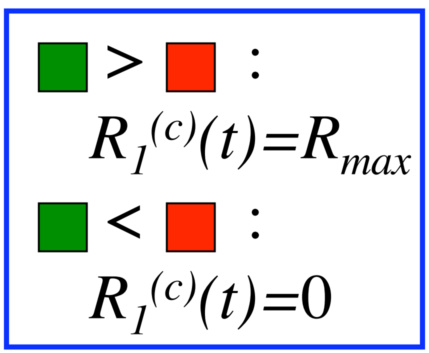

max:

$$Vg_1^{(c)}(\gamma_1^{(c)}) - \gamma_1^{(c)} Z_1^{(c)}(t)$$

$$0 \leq \gamma_1^{(c)} \leq R_{max}$$

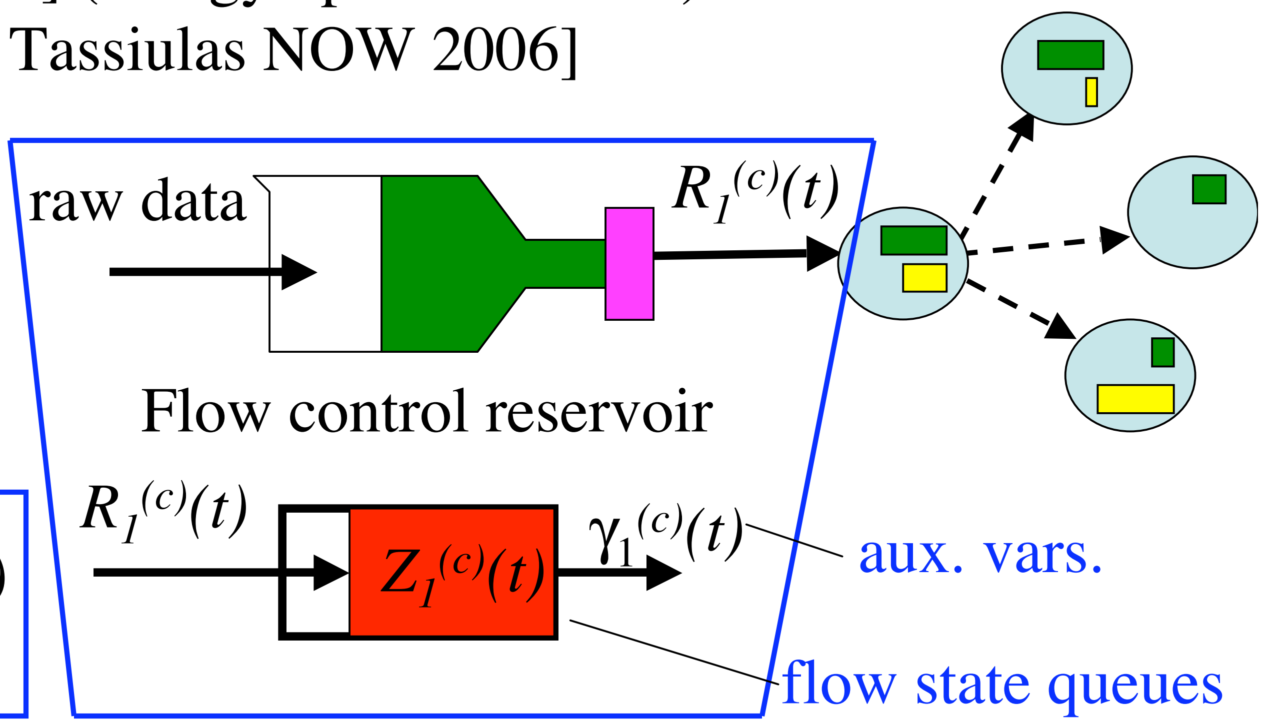

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

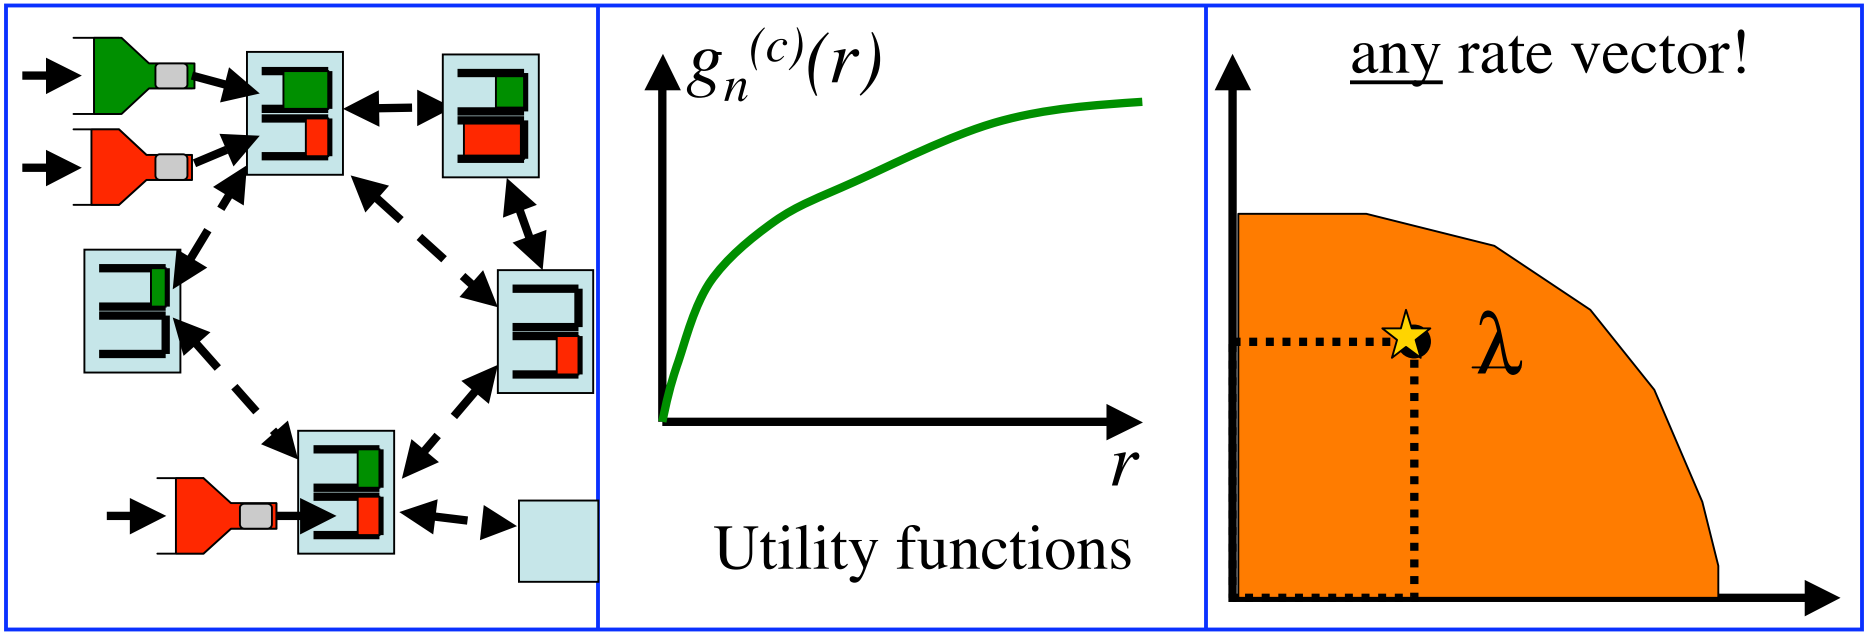

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

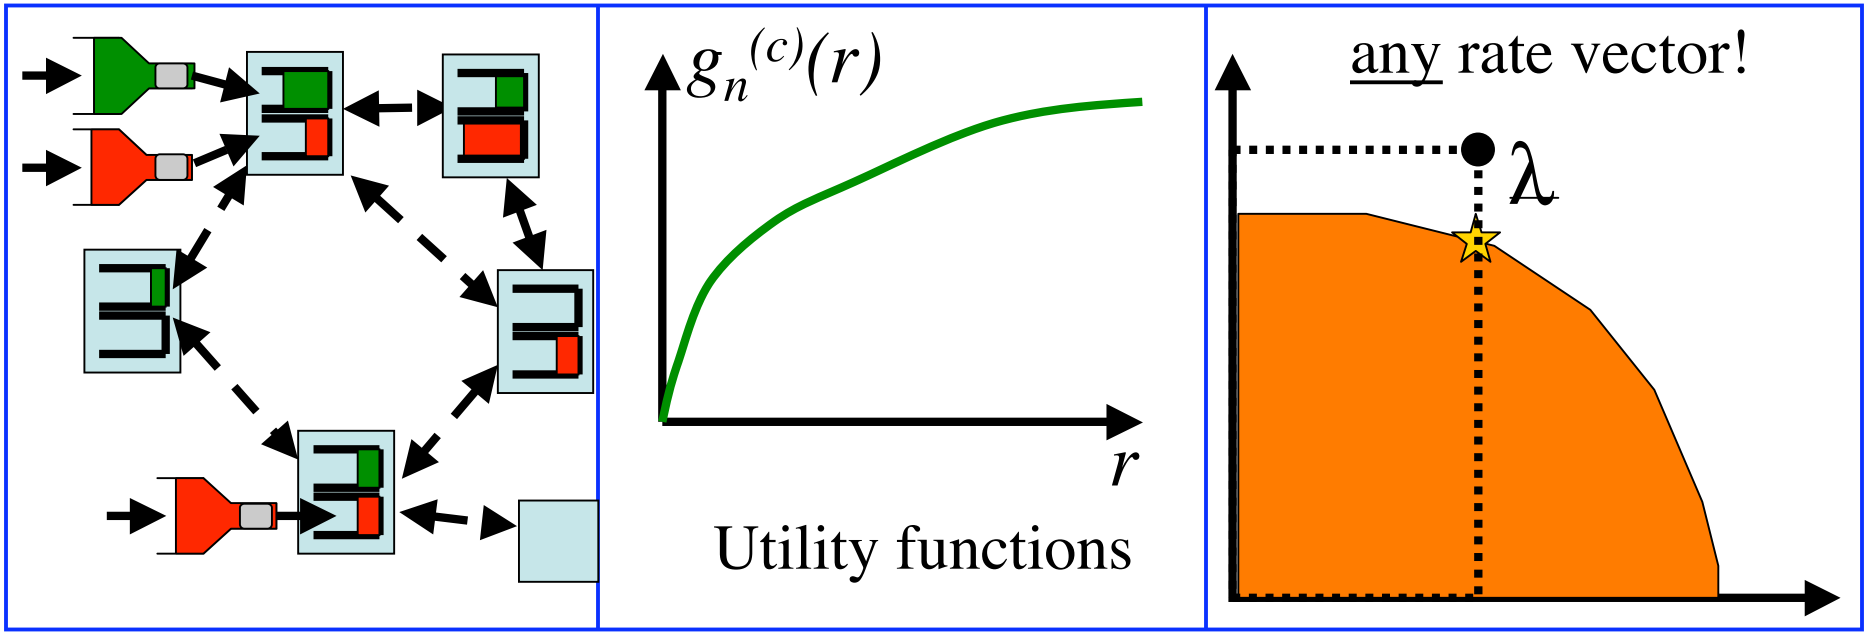

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

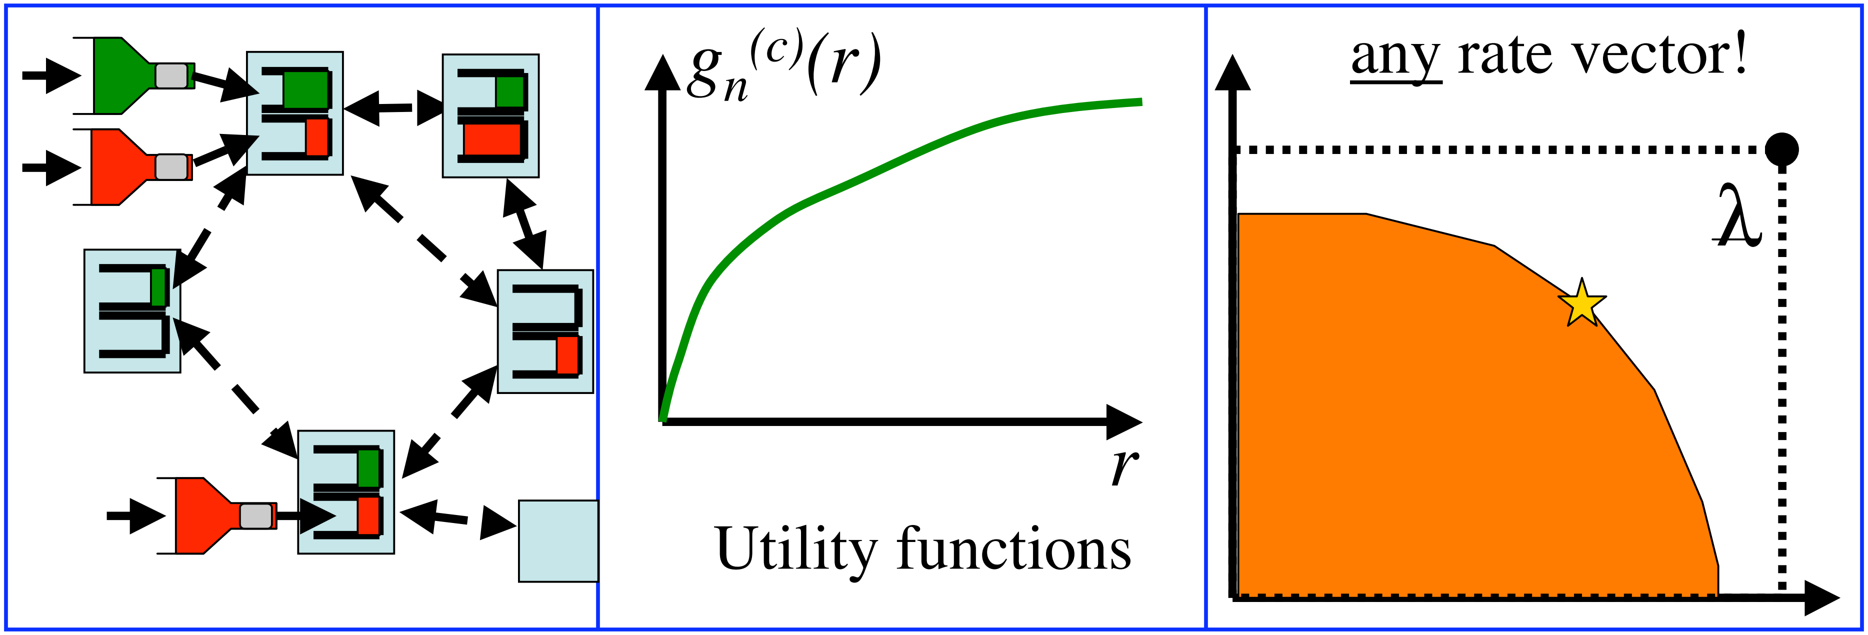

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

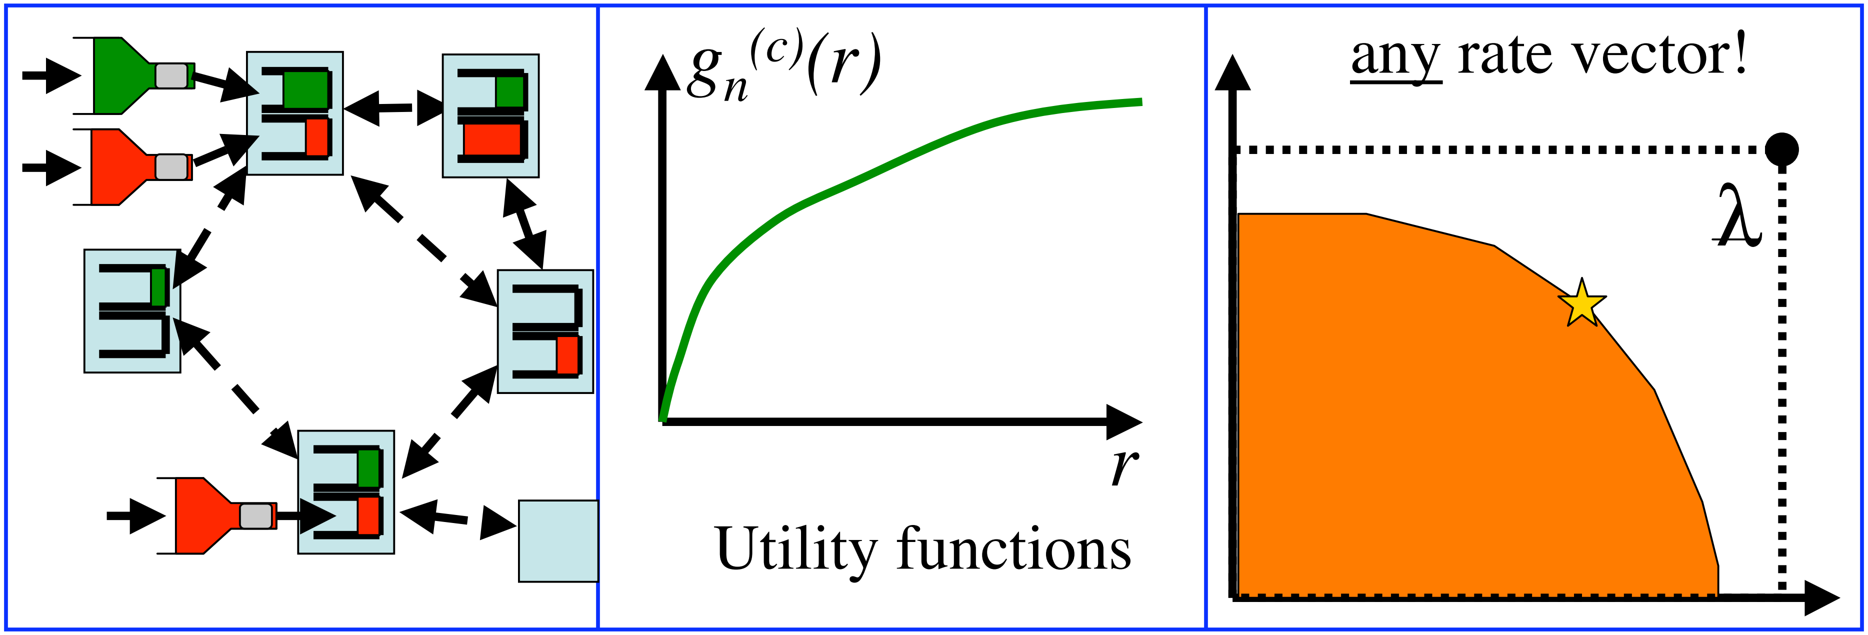

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

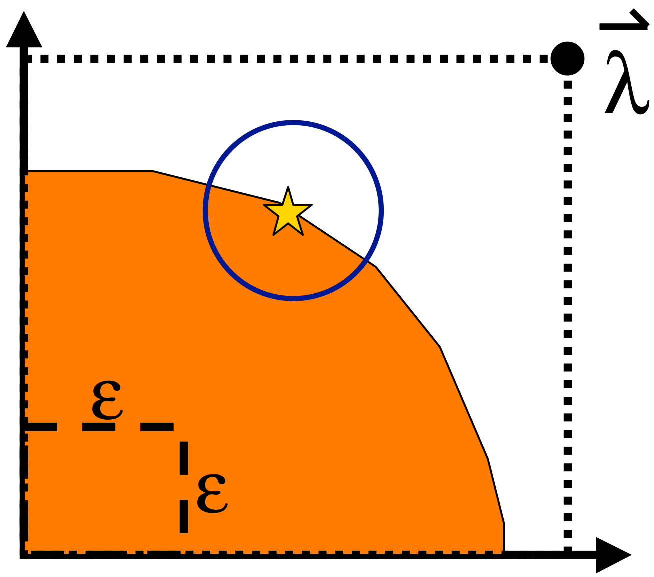

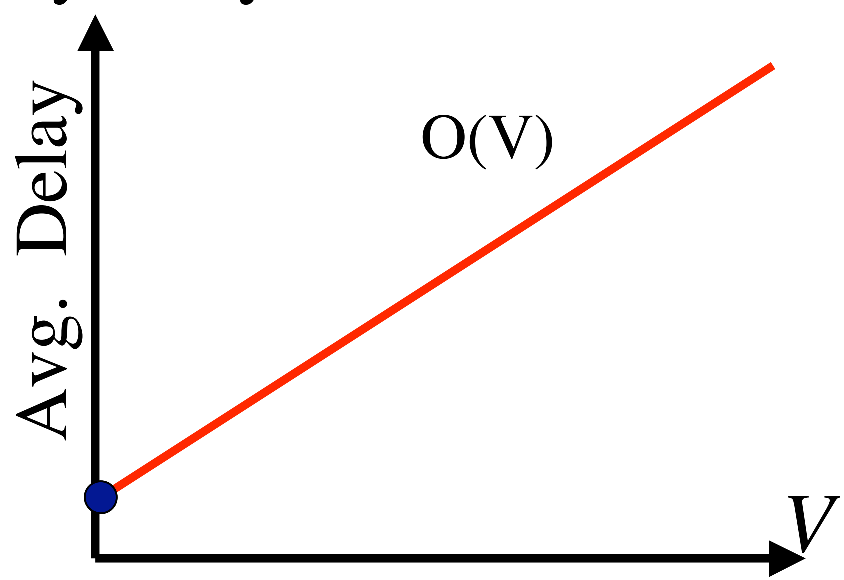

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

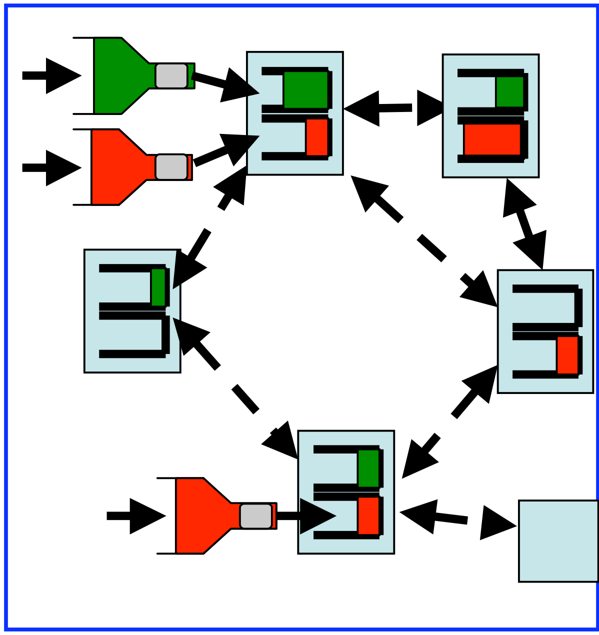

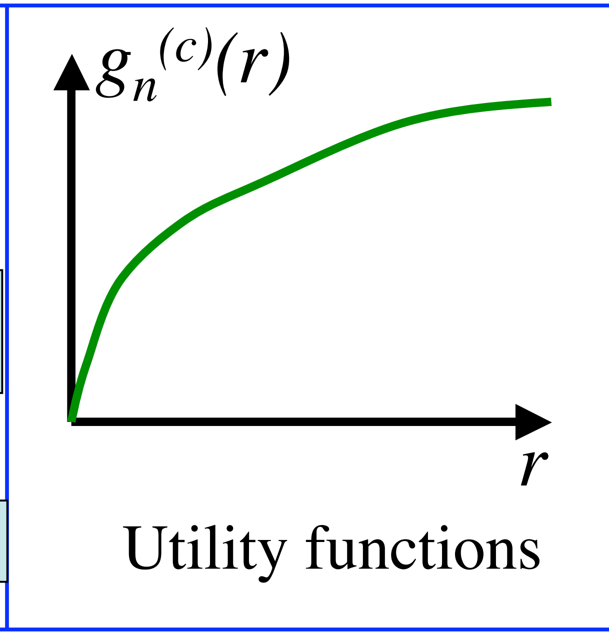

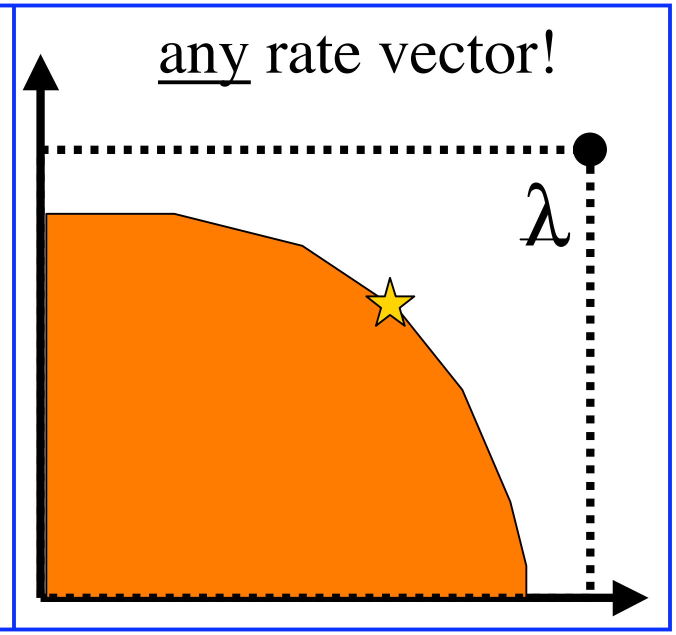

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

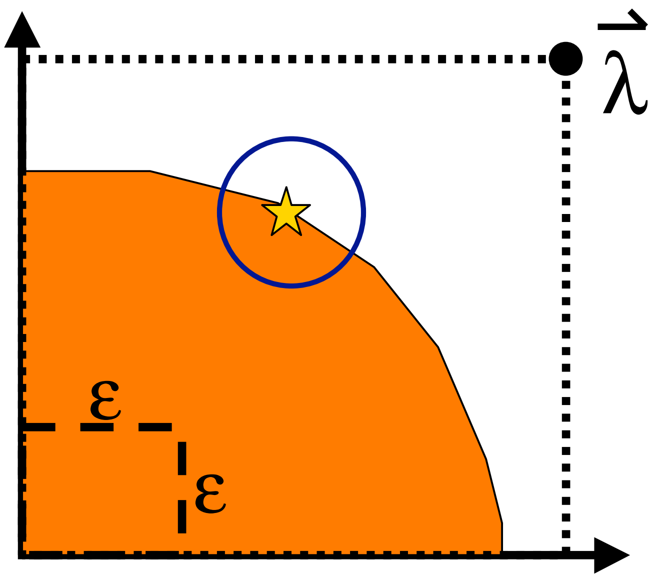

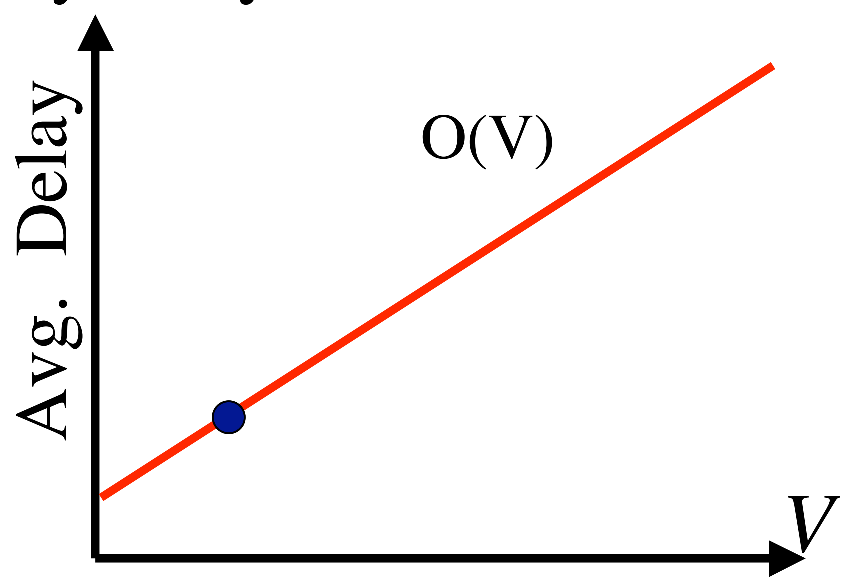

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

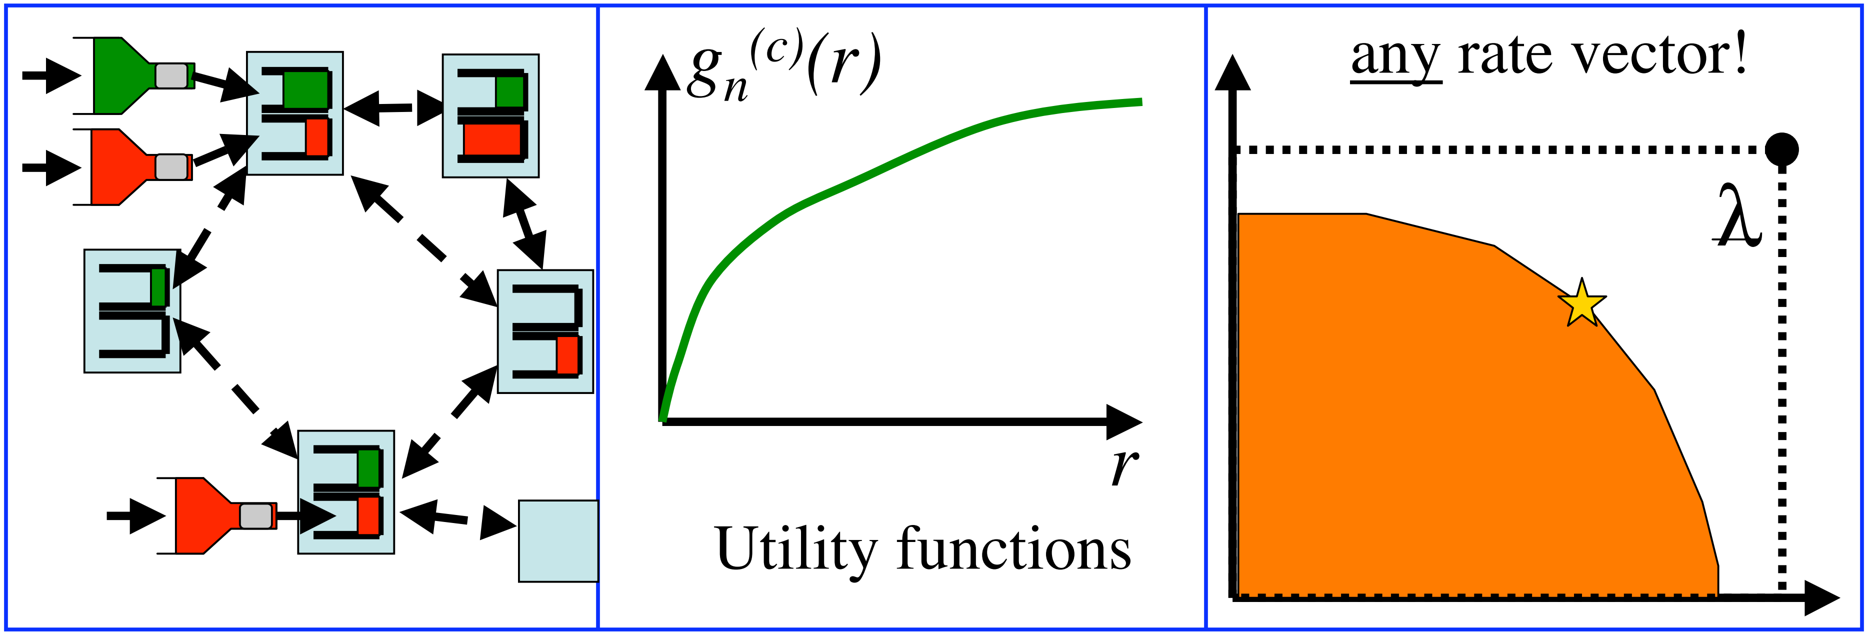

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

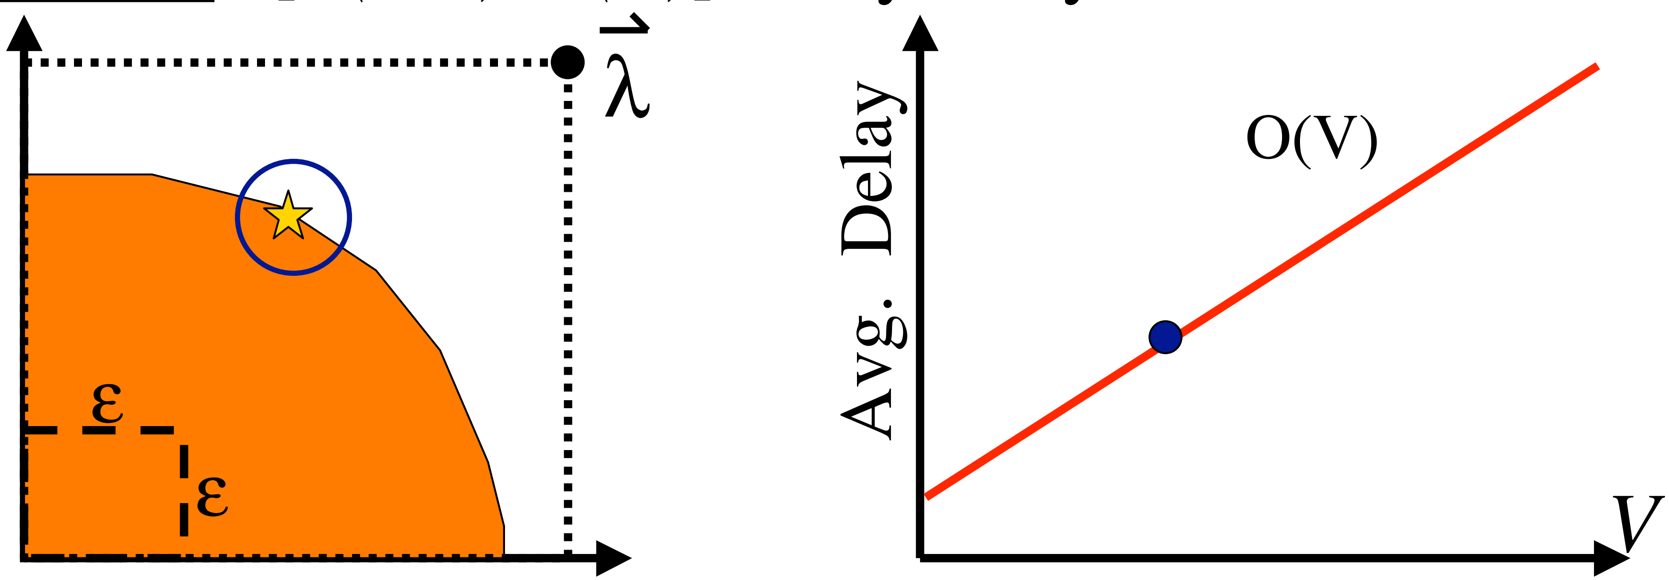

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

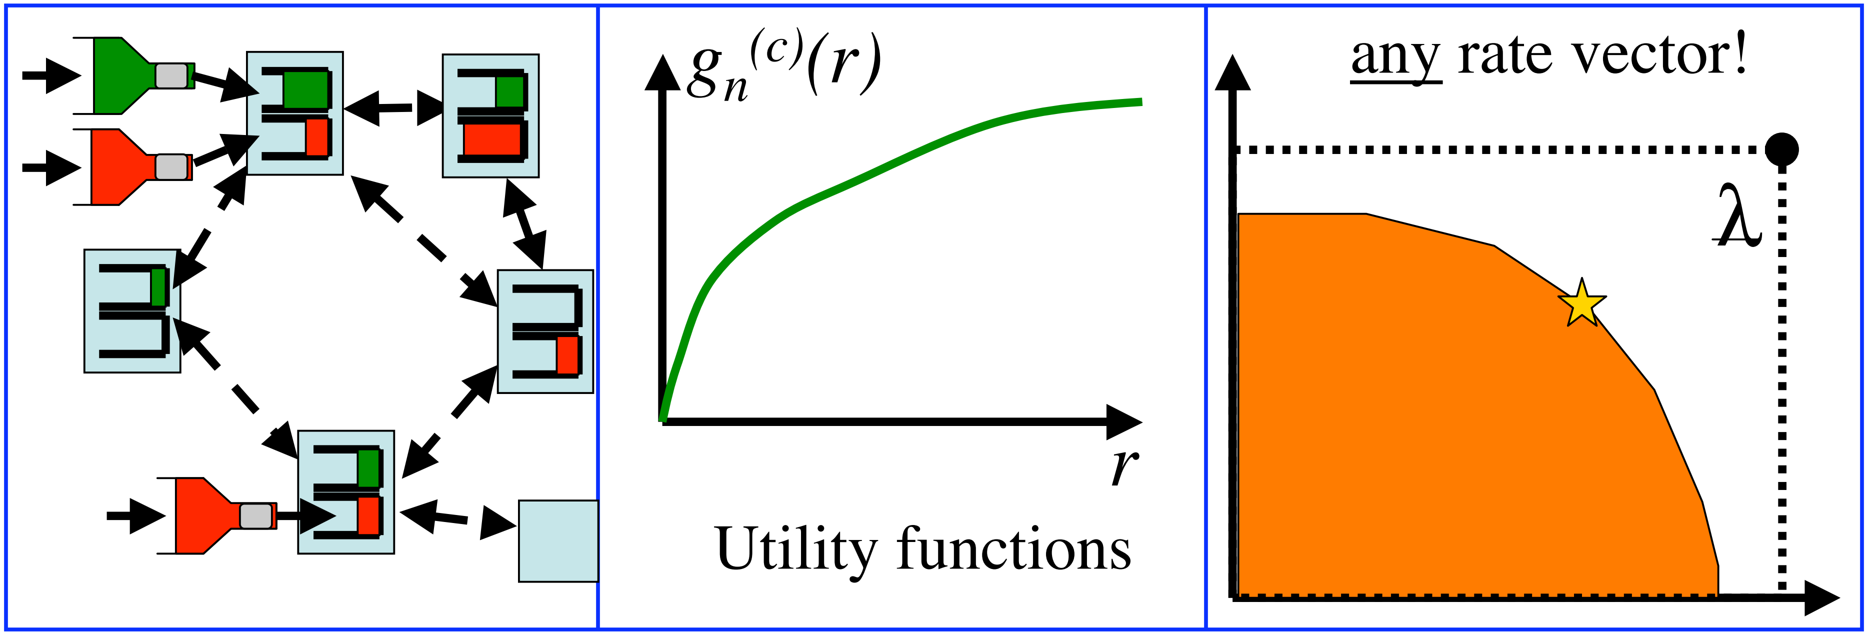

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

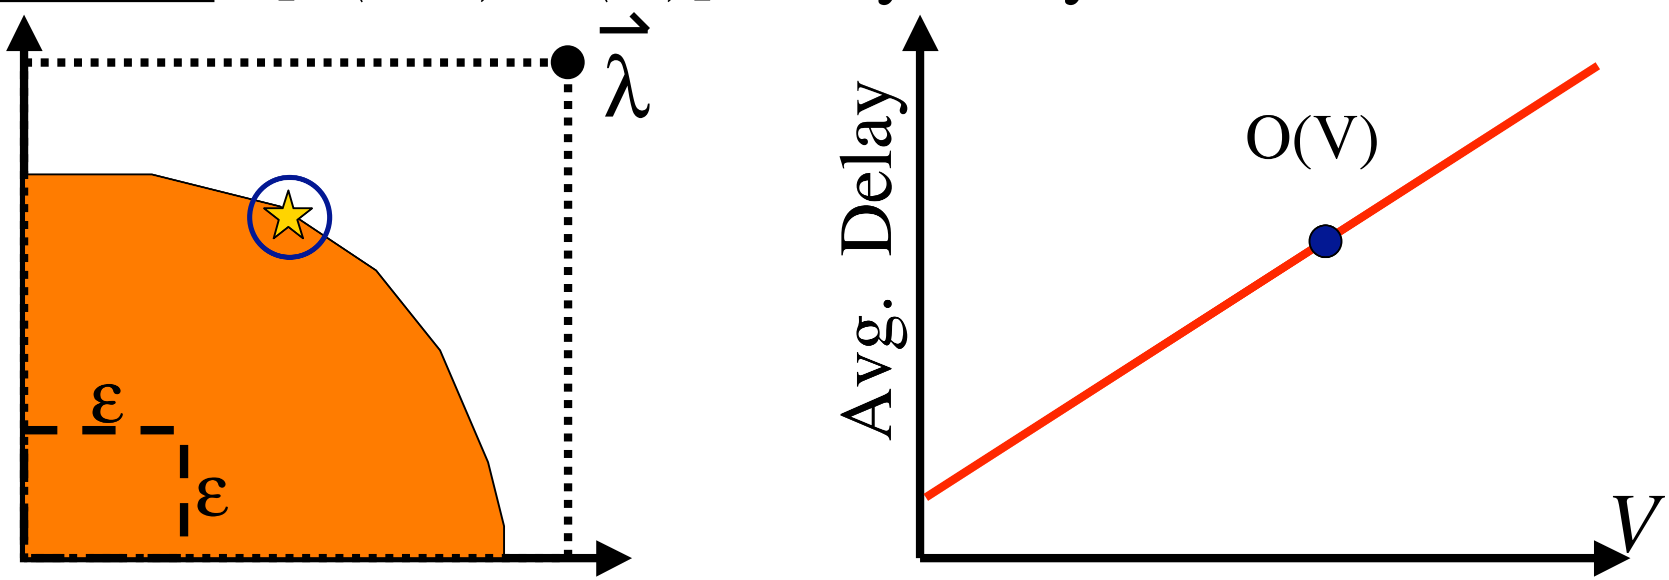

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

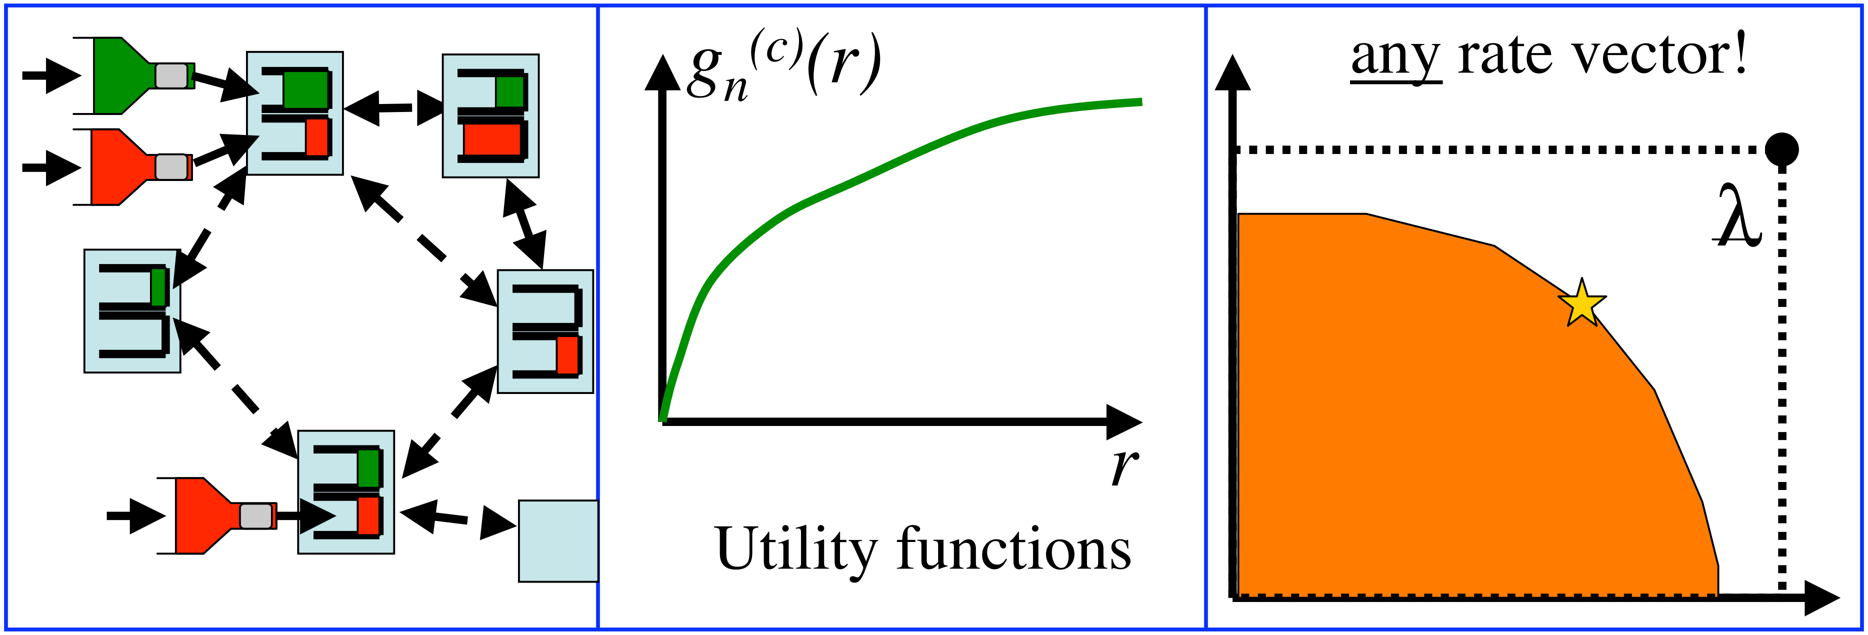

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

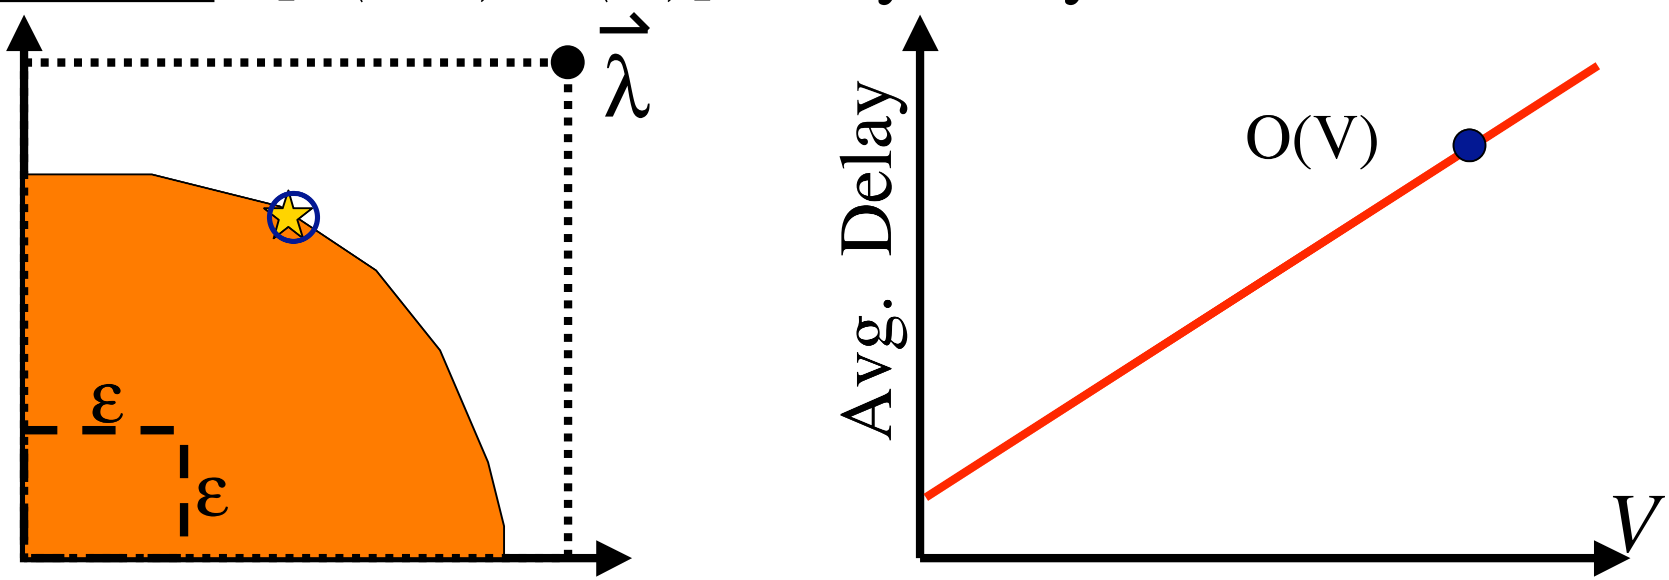

Flow Control + DIVBAR (similar to Neely, Modiano 03, 05)

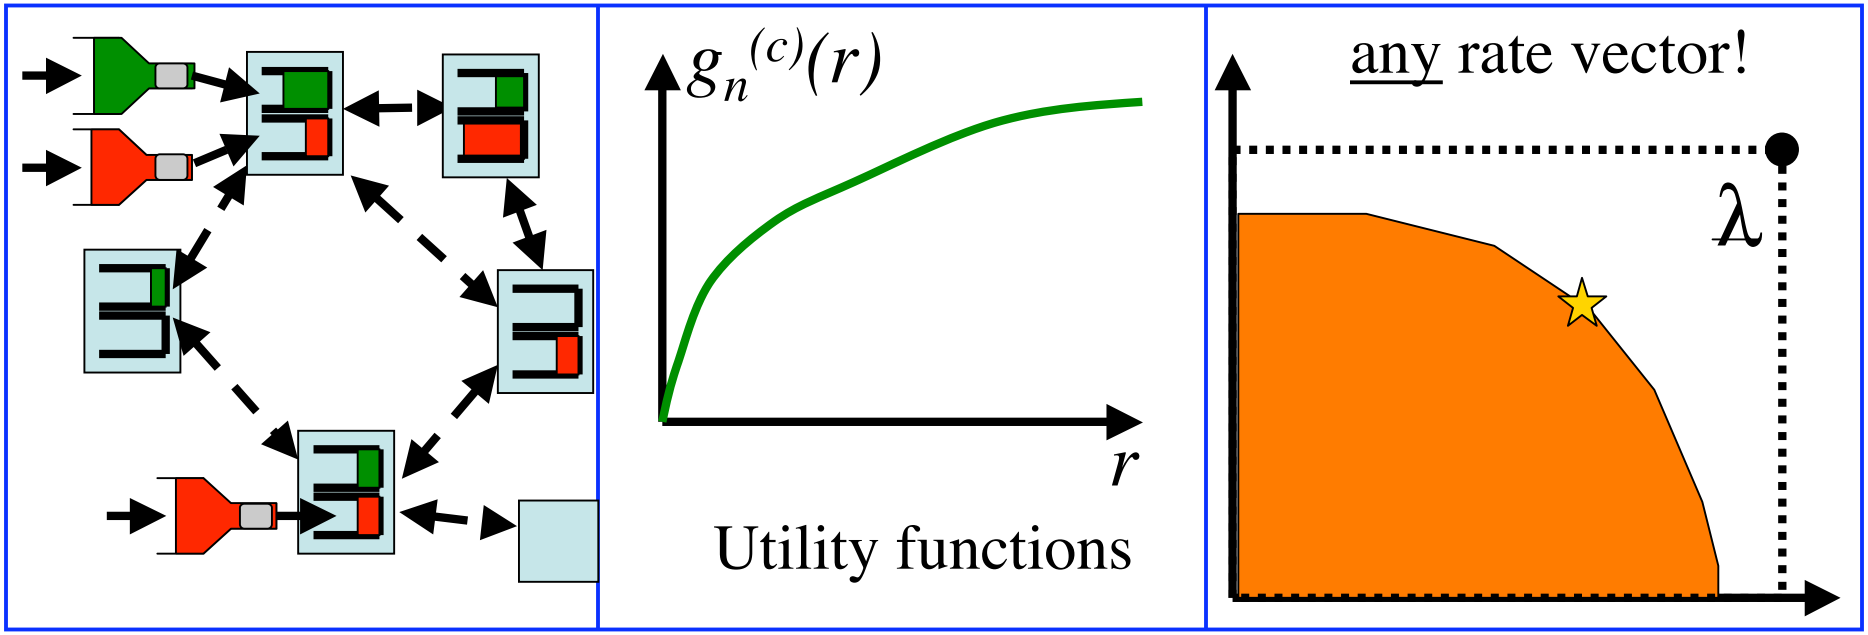

Achieves: [$O(1/V)$, $O(V)$] utility-delay tradeoff!

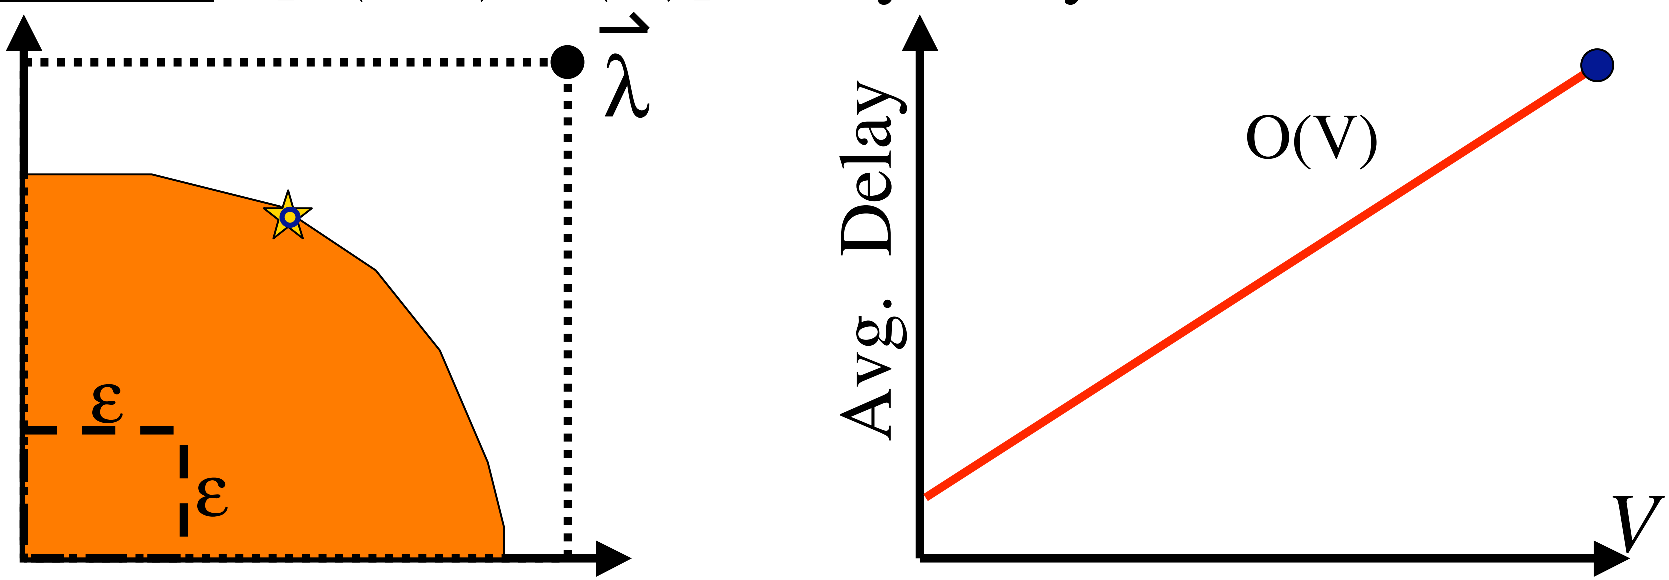

DIVBAR also works for:

-Non-i.i.d. arrivals and channel states

-“Enhanced DIVBAR” (EDR) (improve delay via shortest path metric)

-Distributed MAC via Random Access

(similar to analysis in Neely 2003, JSAC 2005)

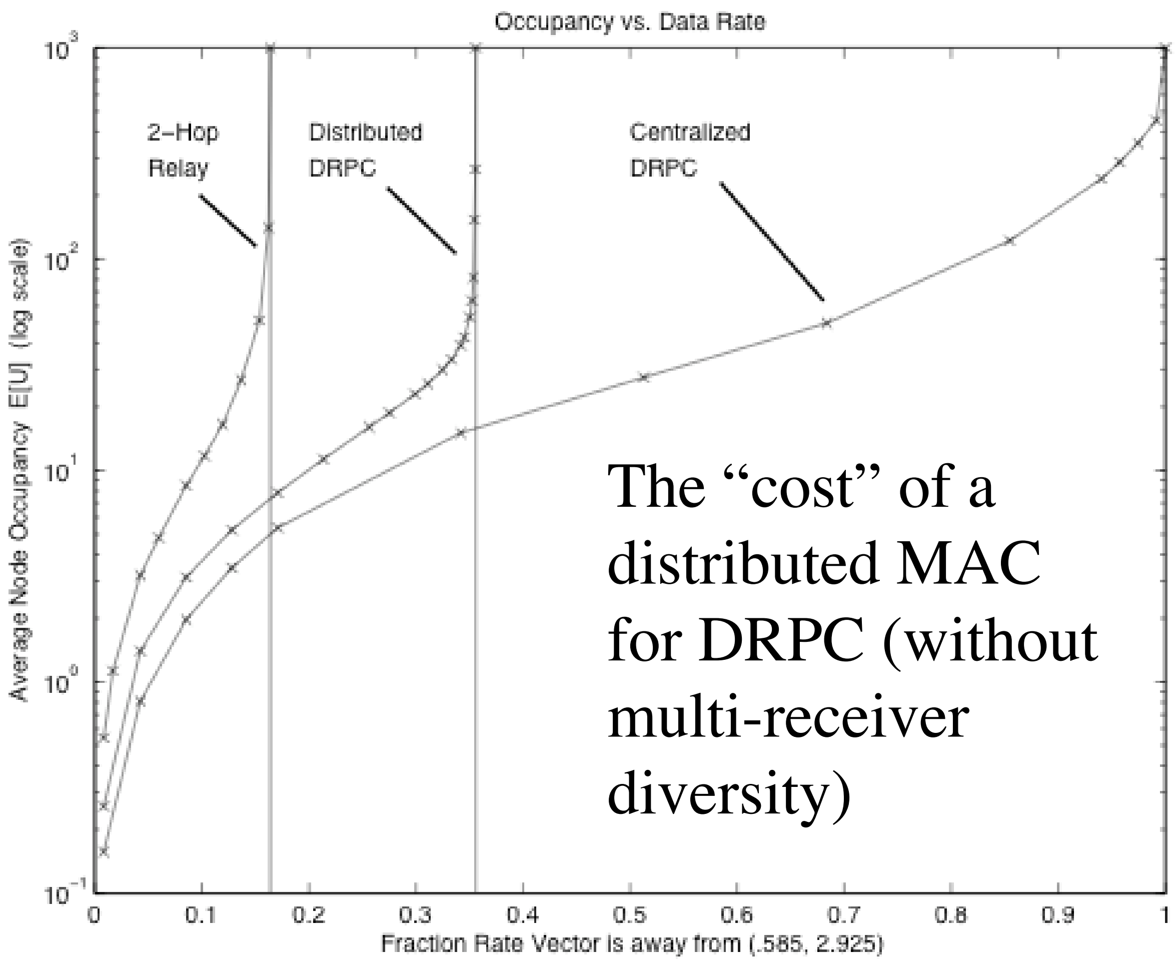

The “cost” of a distributed MAC for DRPC (without multi-receiver diversity)

(DRPC Alg. Of JSAC 2005)

Conclusions:

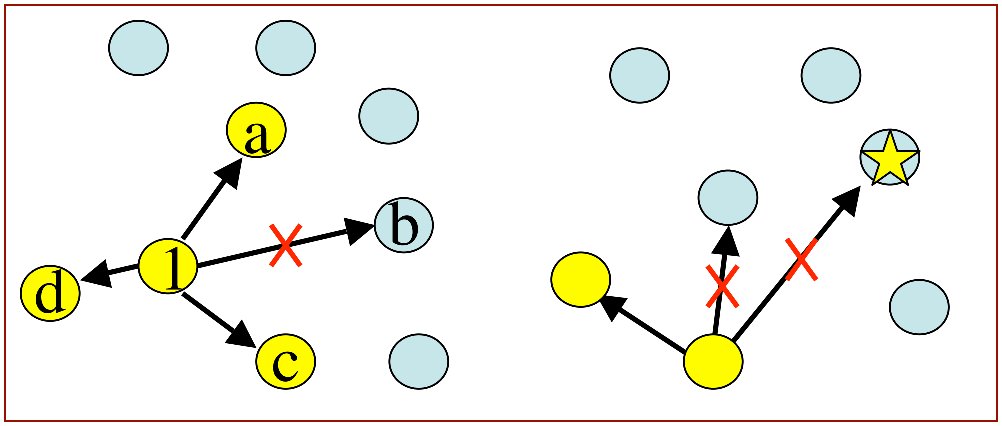

1. DIVBAR takes advantage of Multi-Receiver Diversity.

2. Achieves thruput and energy optimality via a simple backpressure index control law.

3. Channel Blind Transmission Mode: when V=0 and there is only one commodity, DIVBAR achieves thruput optimality without knowledge of channel error probabilities.

4. Flexible algorithm that can be used with other cross layer control techniques and objectives.

…Super-Fast Tradeoffs:

Optimal Energy-Delay Tradeoffs (Square Root Law)
-Berry, Gallager IEEE Trans. on Information Theory 2002
-Neely Infocom 2006

Optimal Utility-Delay Tradeoffs (Logarithm Law)
-Neely Infocom 2006, JSAC 2006