Channel Probing in Communication Systems:
Myopic Policies Are Not Always Optimal

Matt Johnston
Massachusetts Institute of Technology

Joint work with Eytan Modiano and Isaac Keslassy

07/11/13

LIDS

Massachusetts Institute of Technology

Opportunistic Communication

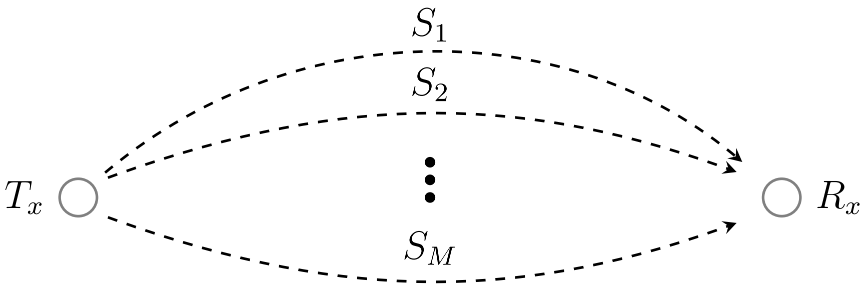

- The quality of wireless channels fluctuates over time

- Objective: Transmit over channels which are in "good" state.
  - "Good" channels yield high throughput
  - Opportunistically selecting channels improves system throughput

- Opportunistic communication requires knowledge of the channel states
  - Transmitter needs to obtain this information (CSI)
  - Obtains CSI via *channel probing*

Control Information

- *General Question:* How much information is necessary to effectively control the network.
  - How often should the transmitter obtain information?
  - What information should the transmitter obtain?

Control Information

- *General Question:* How much information is necessary to effectively control the network.
  - How often should the transmitter obtain information?
  - What information should the transmitter obtain?

- Channel Probing Problem:
  - How often to probe?
    - Last part of the talk
  - What channels to probe?
    - First part of the talk

Previous Work

- Many works looking at channel probing problem
  - See [JMMM ’11], [GMS ‘06], [CP ‘06], [CL ‘07], and references within

Previous Work

- Many works looking at channel probing problem
  - See [JMMM ’11], [GMS ‘06], [CP ‘06], [CL ‘07], and references within

- [Ahmad, Liu, Javidi, Zhao, Krishnamachari; 2009]
  - Channel states vary between ON state and OFF state.
  - Probe one channel in every slot
  - MUST transmit on the probed channel
  - Policy that probes the channel *most likely to be ON* is optimal
  - Myopic Policy: A policy maximizing immediate reward (greedy policy).

Channel States

- Transmitter and Receiver connected through multiple channels

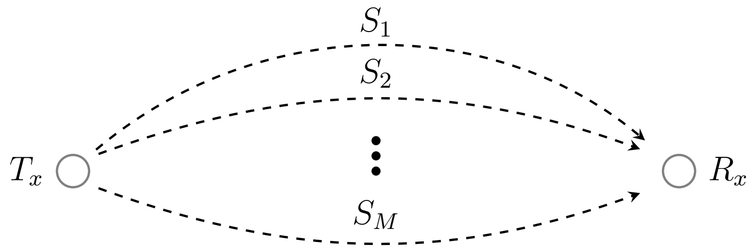

- Channel states are independent of one another

Channel States

- Transmitter and Receiver connected through multiple channels

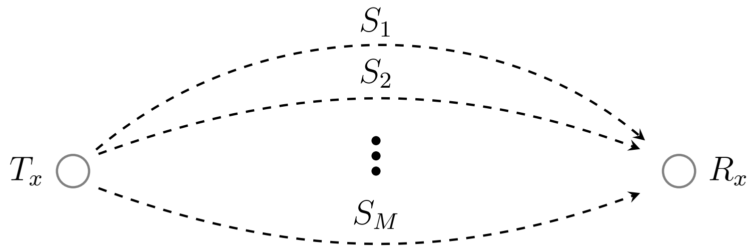

- Channel states are independent of one another

- Assume channel states are ON or OFF:
  - Transmissions across an ON channel are successful
  - Transmissions over an OFF channel are dropped

Channel States

- Transmitter and Receiver connected through multiple channels

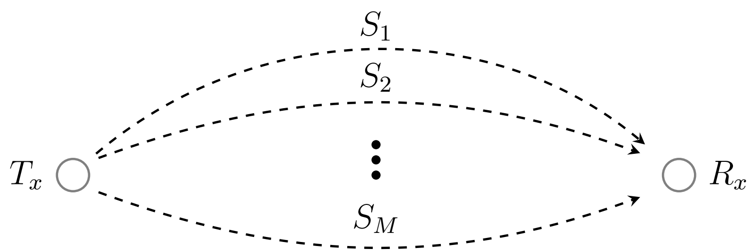

- Channel states are independent of one another

- Assume channel states are ON or OFF:
  - Transmissions across an ON channel are successful
  - Transmissions over an OFF channel are dropped

- Channel States vary over time:
  - Positive channel memory
  - π = the steady state probability of being in the ON state

$S_i(t)$:

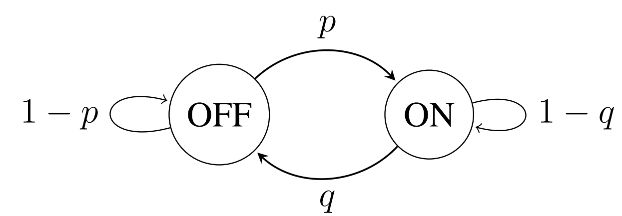

Channel Probing

- Every T slots, the transmitter chooses a channel to probe
  - This is the only way to learn channel state information (CSI)
  - CSI is relevant for multiple time slots

- *Belief* of channel *i:* the probability that channel i is ON given the history of all channel probes.
  - If channel *i* was probed *k* slots ago and was in state *s*.

$$x_i = \mathbf{P}(S_i(t) = \mathrm{ON} | S_i(t-k) = s)$$

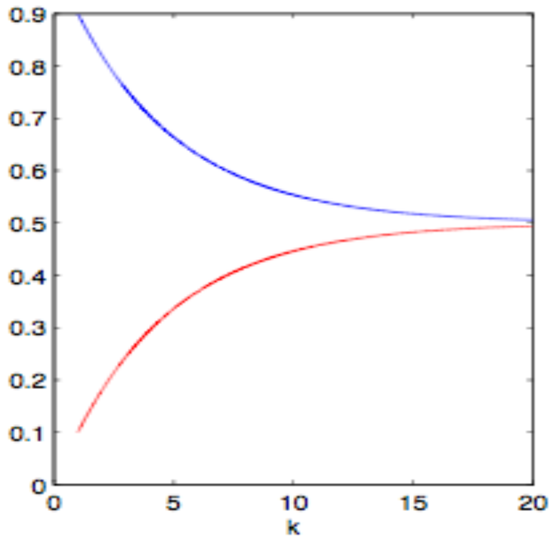

Channel Probing

- Every T slots, the transmitter chooses a channel to probe
  - This is the only way to learn channel state information (CSI)
  - CSI is relevant for multiple time slots

- *Belief* of channel *i:* the probability that channel i is ON given the history of all channel probes.
  - If channel *i* was probed *k* slots ago and was in state *s*,

$$x_i = \mathbf{P}(S_i(t) = \mathrm{ON} | S_i(t-k) = s)$$

- Transmitter transmits over channel with highest belief
  - Expected throughput: $\max_i x_i(t)$
  - Transmitter will transmit over the same channel until new channel probe.
  - Transmitter is not restricted to transmit over the probed channel
    - This restriction was present in previous work [Ahmad et. al, '09]

Infinite Channel System

- We are interested in systems with a large number of channels.

- Infinite channel simplification:
  - When you probe a channel and it’s OFF, it is effectively removed from the system.
  - There always exists a channel that hasn’t been probed for an infinitely long time
    - Belief of such a channel = π

- What is the optimal channel probing policy?
  - Assume fixed probing instances.

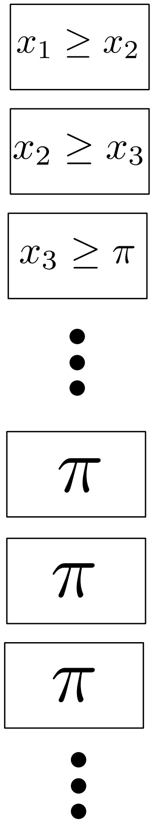

Probe Best Policy

- *Probe best policy [Ahmad et al. '09]:* At each probing instance, probes the channel with the highest belief.

- *Observation:* under the probe best policy, *at most* one channel has belief larger than π.

- Example: (order channels in descending order of belief)
  - Assuming $T = 1$ for illustration, but all intuition holds for $T > 1$.

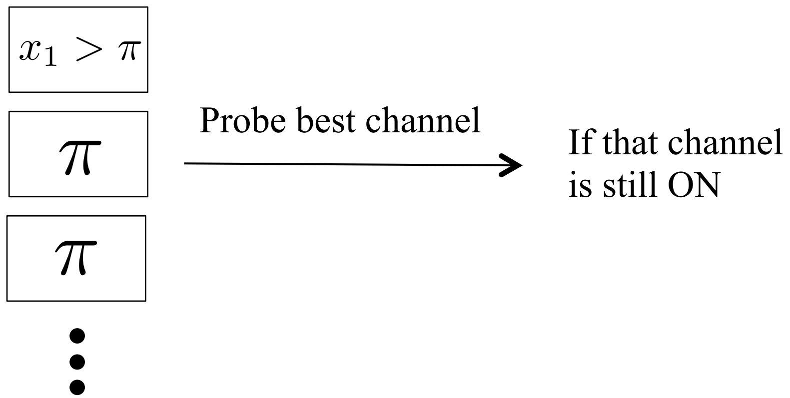

Probe Best Policy

- *Probe best policy [Ahmad et al. '09]:* At each probing instance, probes the channel with the highest belief.

- *Observation:* under the probe best policy, *at most* one channel has belief larger than π.

- Example: (order channels in descending order of belief)
  - Assuming T = 1 for illustration, but all intuition holds for T > 1.

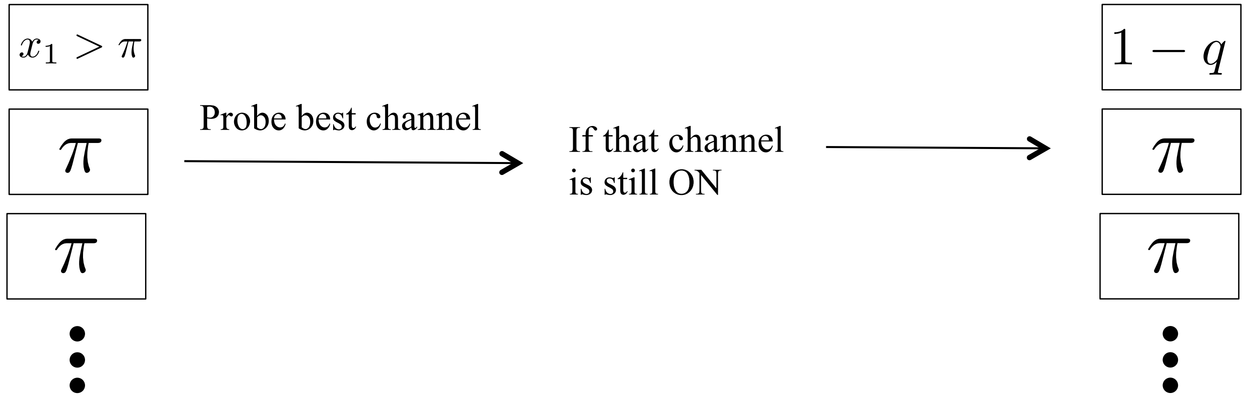

Probe Best Policy

- *Probe best policy [Ahmad et al. '09]:* At each probing instance, probes the channel with the highest belief.

- *Observation:* under the probe best policy, *at most* one channel has belief larger than π.

- Example: (order channels in descending order of belief)
  - Assuming T = 1 for illustraiton, but all intuition holds for T > 1.

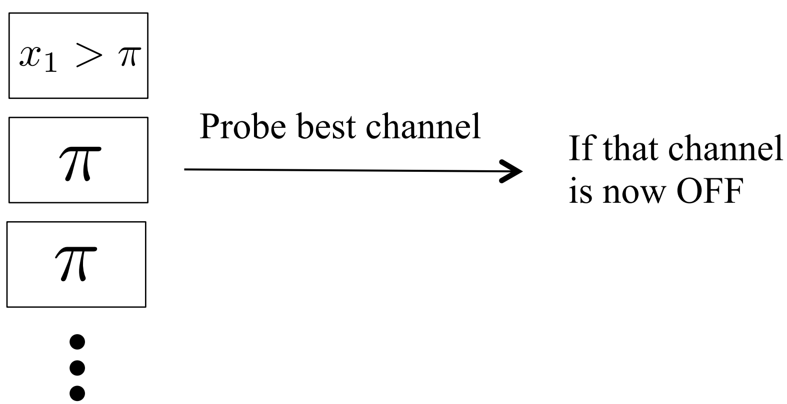

Probe Best Policy

- *Probe best policy [Ahmad et al. '09]:* At each probing instance, probes the channel with the highest belief.

- *Observation:* under the probe best policy, *at most* one channel has belief larger than π.

- Example: (order channels in descending order of belief)
  - Assuming T = 1 for illustraiton, but all intuition holds for T > 1.

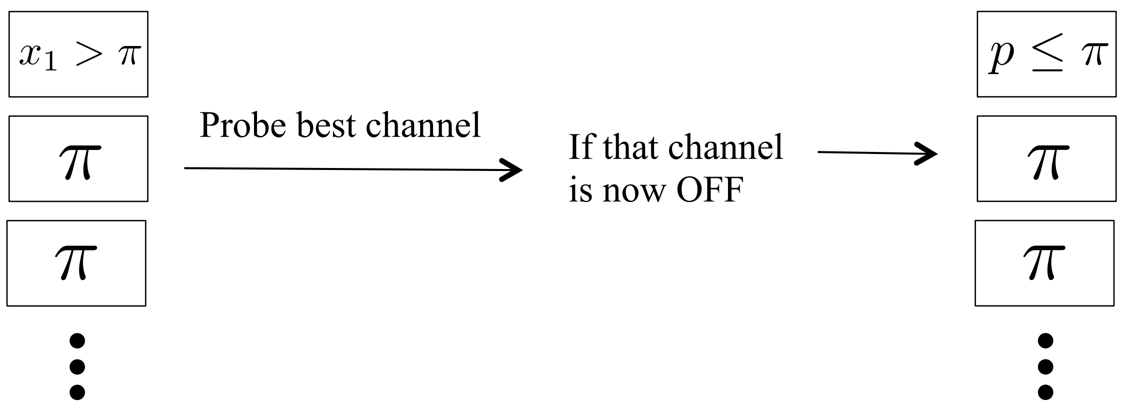

Probe Best Policy

- *Probe best policy [Ahmad et al. '09]:* At each probing instance, probes the channel with the highest belief.
- *Observation:* under the probe best policy, *at most* one channel has belief larger than π.
- Example: (order channels in descending order of belief)
  - Assuming T = 1 for illustraiton, but all intuition holds for T > 1.

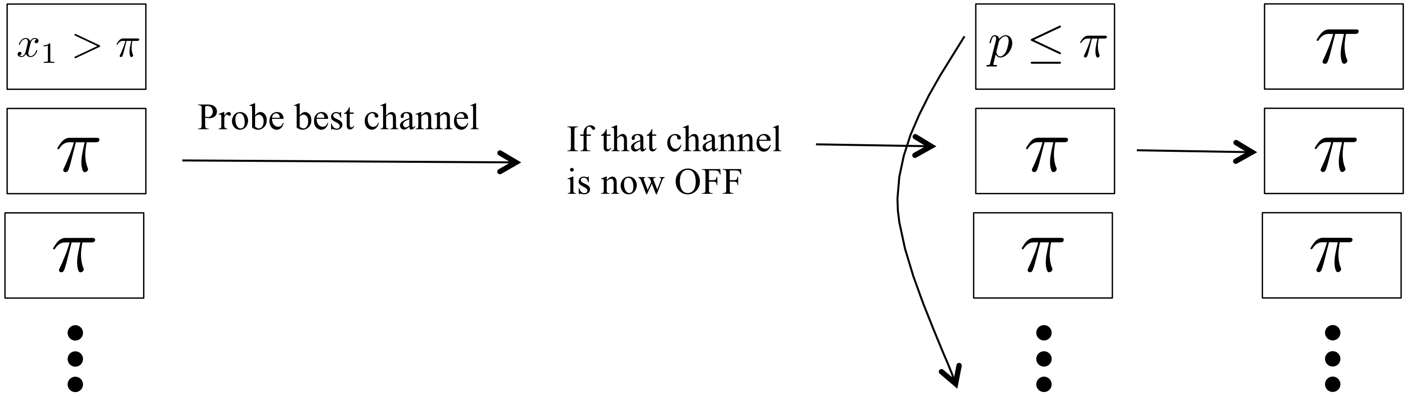

Probe Best Policy

- As long as the transmitter probes OFF channels, the state stays the same.

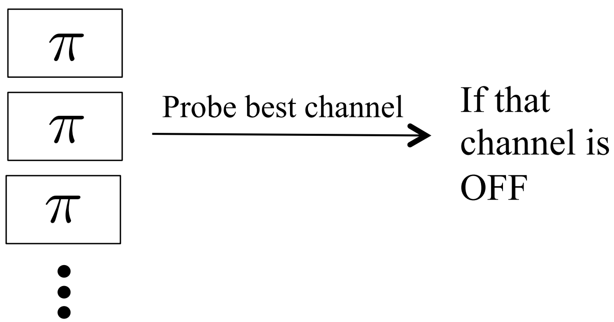

Probe Best Policy

- As long as the transmitter probes OFF channels, the state stays the same.

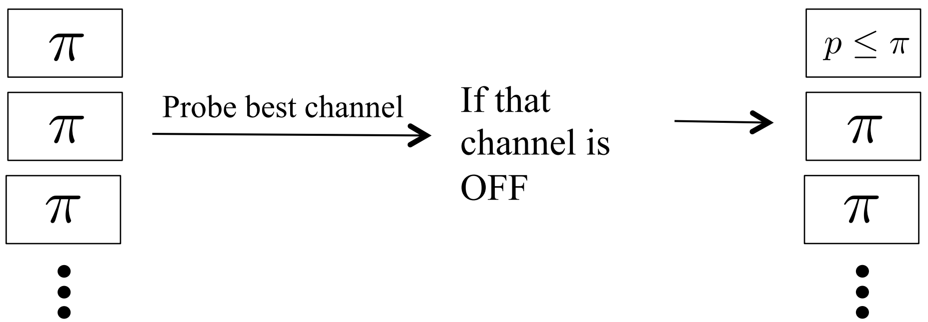

Probe Best Policy

- As long as the transmitter probes OFF channels, the state stays the same.

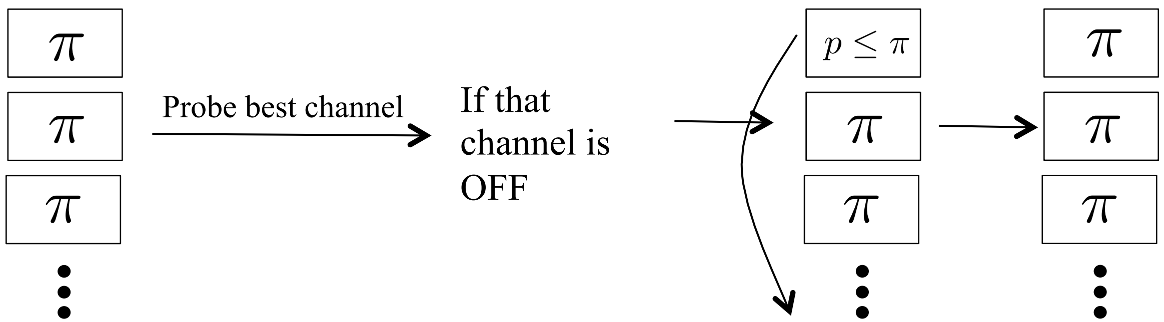

- This repeats until an ON channel is found.

- Renewal Channel Process: renewal occurs upon OFF channel probe

- A renewal occurs when the ordered belief vector is $\mathbf{x} = (\pi, \pi, \ldots)$
  - If an ON channel is found, that channel is probed until found OFF → renewal

- Use renewal-reward theory to compute average throughput.

Probe Best Policy Discussion

- Advantages to Probe Best Policy
  - Probing the channel with the highest belief maximizes the immediate probability of finding an ON channel
  - Maximizes Immediate Throughput (greedy).

- Disadvantages to Probe Best Policy
  - When an OFF channel is found, the transmitter has no knowledge of which channel to probe to find an ON channel.
  - Until an ON channel is found, the transmitter sends packets over a channel with belief $\pi \rightarrow$ Low expected throughput.

Immediate Reward

- Assume we sort the channels by belief (high to low)
  - $(x_1, x_2, x_3, \ldots)$

$$E[\text{Reward} \mid \text{Probe Ch. i}] = \Pr(\text{Ch. i is ON})E[\text{Reward} \mid \text{Ch. i is ON}] + \Pr(\text{Ch. i is OFF})E[\text{Reward} \mid \text{Ch. i is OFF}]$$

- Probe the best channel:

$$E[\text{Reward}] = x_1 \cdot 1 + (1 - x_1) \cdot x_2 = x_1 + x_2 - x_1 x_2$$

Immediate Reward

- Assume we sort the channels by belief (high to low)
  - $(x_1, x_2, x_3, \ldots)$

$$E[\text{Reward} \mid \text{Probe Ch. i}] = \Pr(\text{Ch. i is ON}) E[\text{Reward} \mid \text{Ch. i is ON}] + \Pr(\text{Ch. i is OFF}) E[\text{Reward} \mid \text{Ch. i is OFF}]$$

- Probe the best channel:

$$E[\text{Reward}] = x_1 \cdot 1 + (1 - x_1) \cdot x_2 = x_1 + x_2 - x_1 x_2$$

- Probe the 2nd best channel:

$$E[\text{Reward}] = x_2 \cdot 1 + (1 - x_2) \cdot x_1 = x_1 + x_2 - x_1 x_2$$

Immediate Reward

- Assume we sort the channels by belief (high to low)
  - $(x_1, x_2, x_3, \ldots)$

$$E[\text{Reward} \mid \text{Probe Ch. i}] = \Pr(\text{Ch. i is ON}) E[\text{Reward} \mid \text{Ch. i is ON}] + \Pr(\text{Ch. i is OFF}) E[\text{Reward} \mid \text{Ch. i is OFF}]$$

- Probe the best channel:

$$E[\text{Reward}] = x_1 \cdot 1 + (1 - x_1) \cdot x_2 = x_1 + x_2 - x_1 x_2$$

- Probe the 2nd best channel:

$$E[\text{Reward}] = x_2 \cdot 1 + (1 - x_2) \cdot x_1 = x_1 + x_2 - x_1 x_2$$

- Probe the kth best channel:

$$E[\text{Reward}] = x_k \cdot 1 + (1 - x_k) \cdot x_1 = x_1 + x_k - x_1 x_k$$

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- Intuition:
  - Under the *probe second best policy*, there can be *two* channels with belief greater than π.

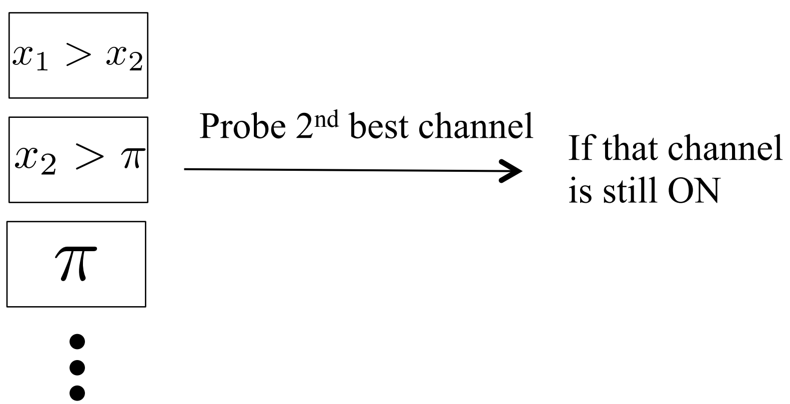

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- Intuition:
  - Under the *probe second best policy*, there can be *two* channels with belief greater than π.

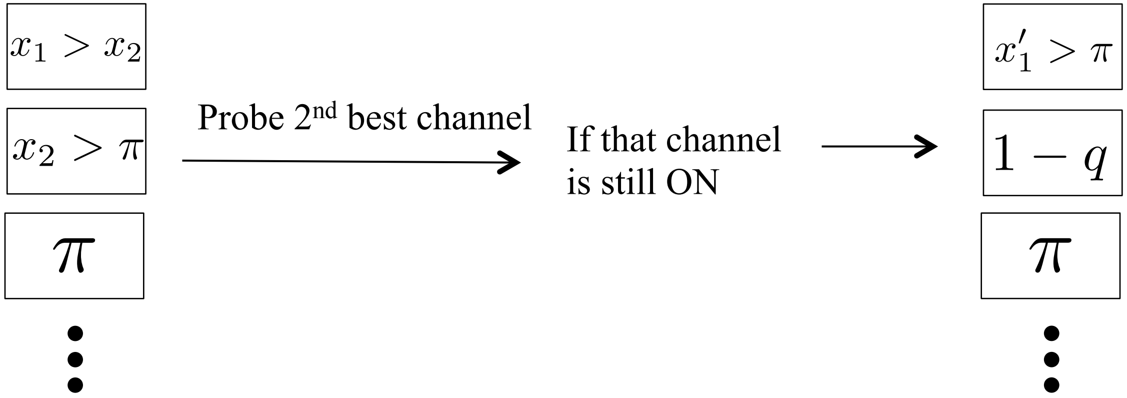

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- Intuition:
  - Under the *probe second best policy*, there can be *two* channels with belief greater than π.

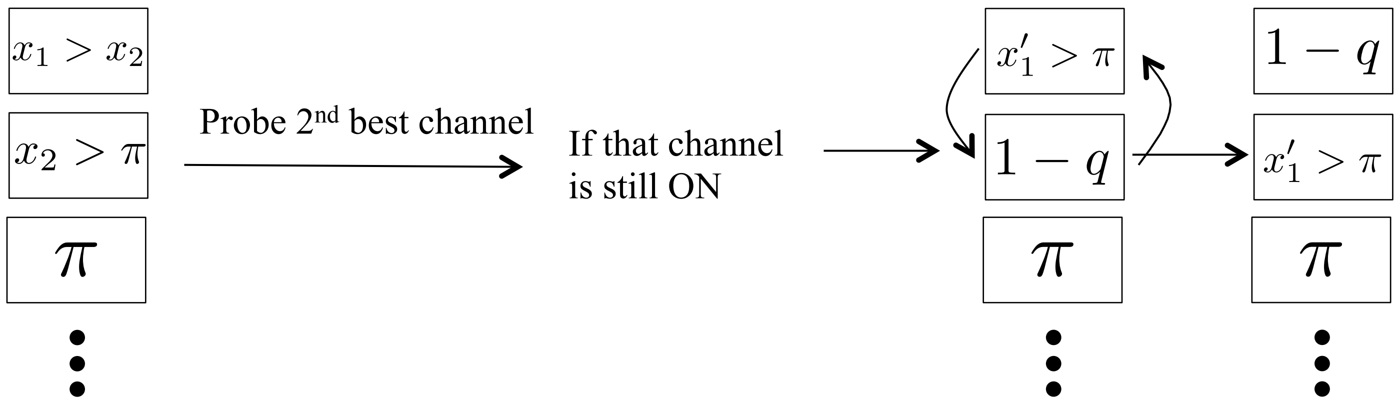

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

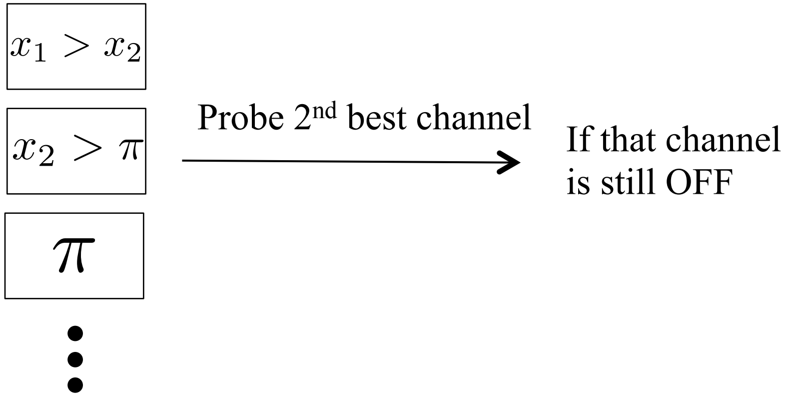

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

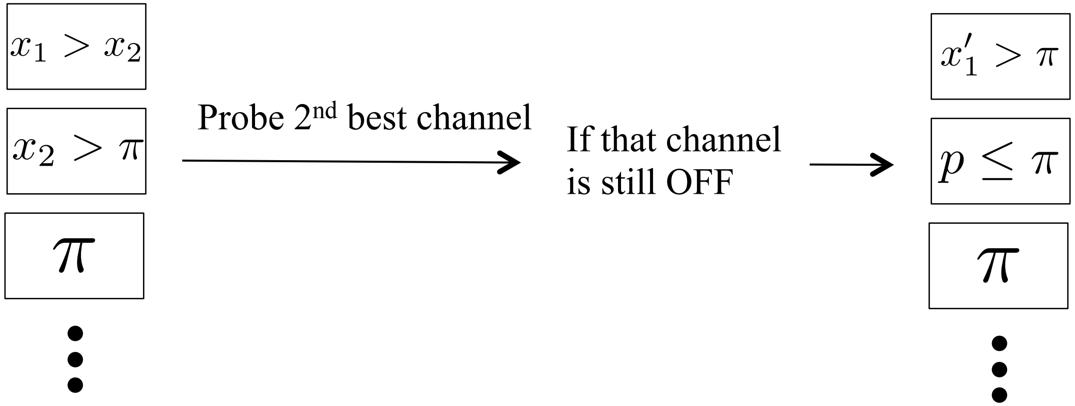

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

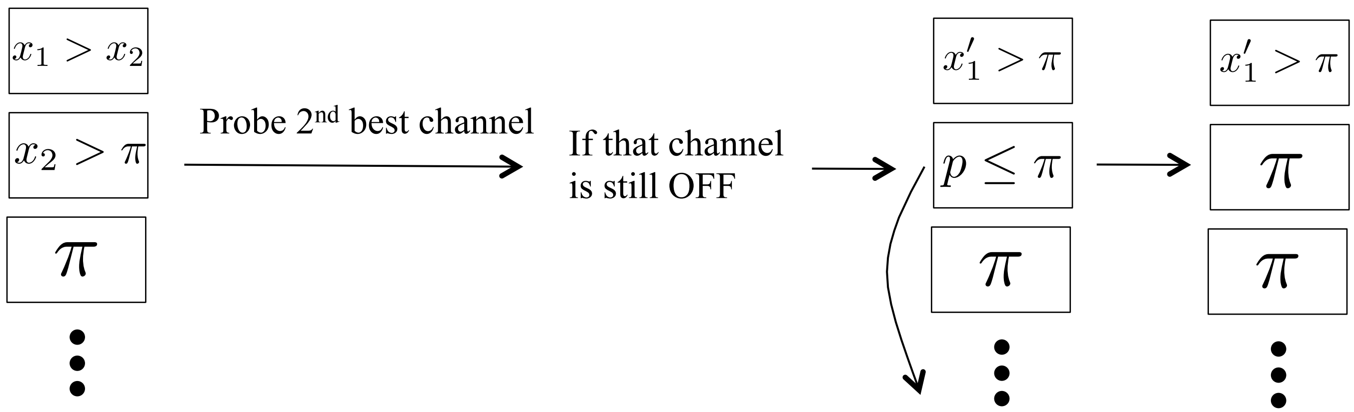

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

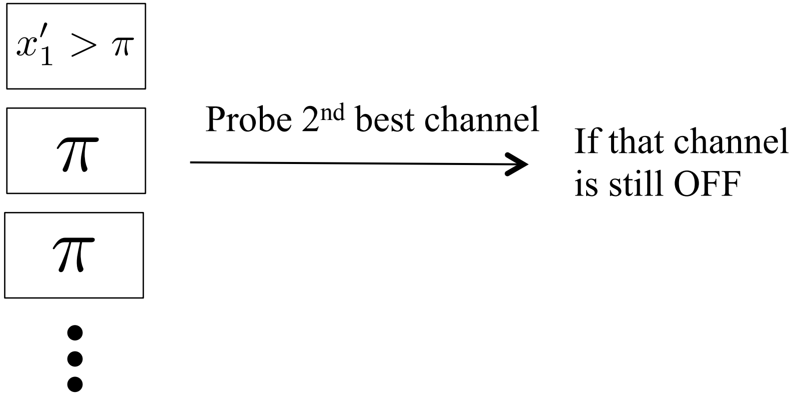

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

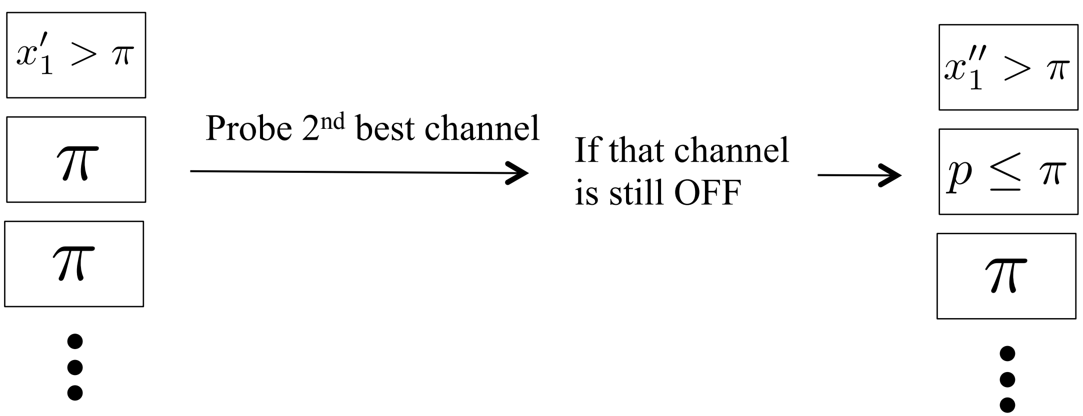

Probe Second Best Policy

- *Probe Second Best Policy*: The policy that probes the channel with the second highest belief.
- When a channel is OFF, the transmitter has a channel with high belief to transmit over until a new ON channel is found.

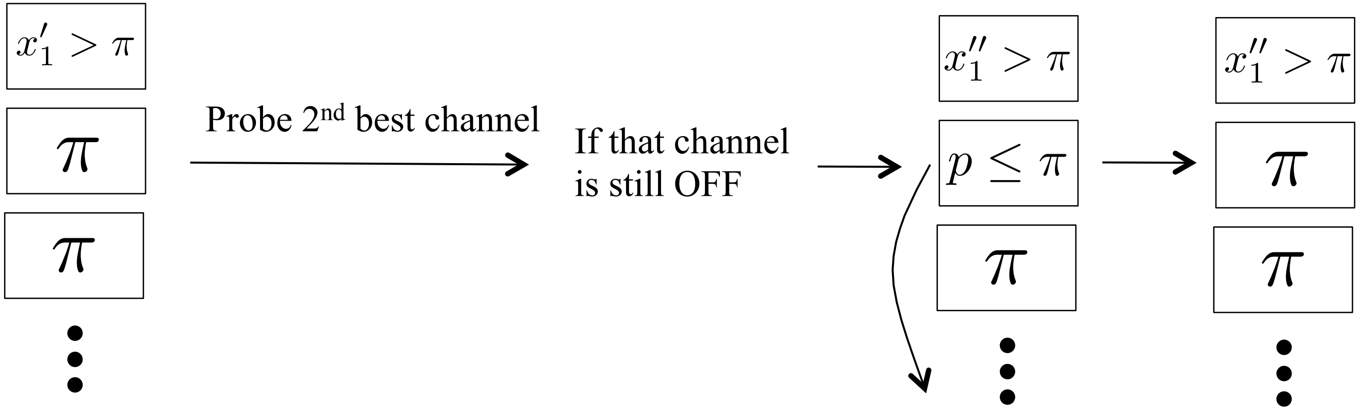

Renewal Theory analysis

- A renewal occurs upon two consecutive ON channel probes
- Eventually, two consecutive probes will be ON:

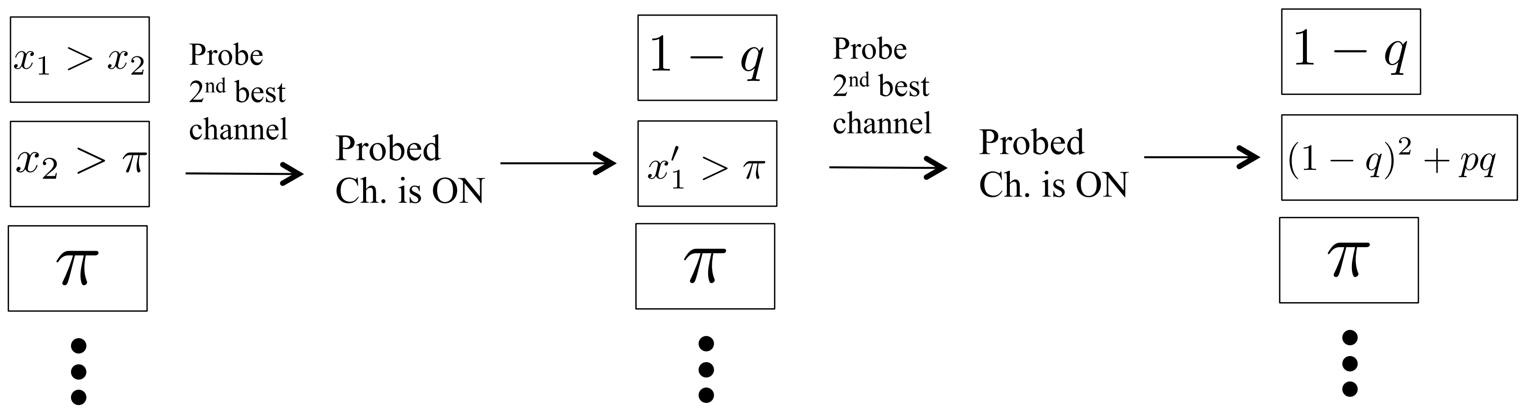

- From this state, the time to arrive to this state again is i.i.d.

- We can use renewal-reward theory to calculate average throughput.
  - Function of p, q, and T.

Policy Comparison

- **Theorem**: For fixed probing times T, the *probe second best* policy has a **higher expected throughput** than the probe best policy.
- Numerically, for p=q=0.05:

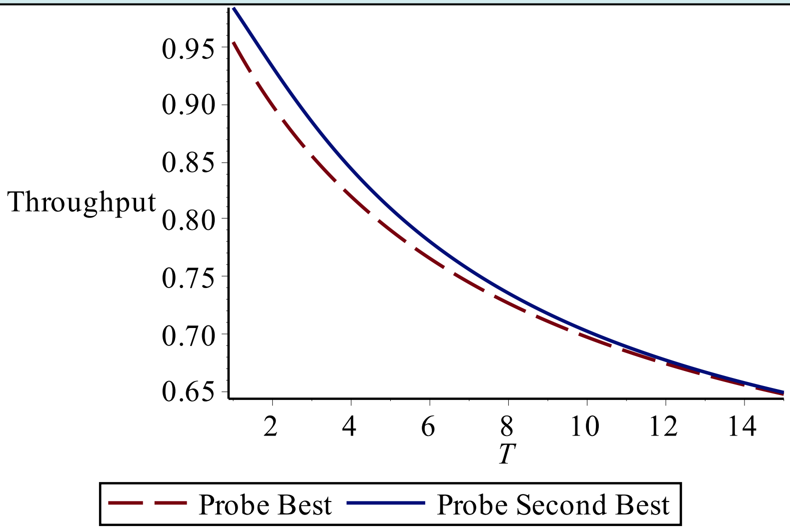

- Note that in the case where the transmitter must send over the probed channel, the probe second best policy has smaller immediate reward
  - Probe best policy shown to be optimal in this case *[Ahmad et al. '09]*.

Other Policies

- *Round robin policy*: Probes the channel for which the transmitter has the least knowledge.
  - In an infinite channel system, the belief of the probed channel is always π, while the transmitter will send over the channel that was last found to be in an ON state.

Other Policies

- *Round robin policy*: Probes the channel for which the transmitter has the least knowledge.
  - In an infinite channel system, the belief of the probed channel is always π, while the transmitter will send over the channel that was last found to be in an ON state.
- **Theorem:** The round robin policy has the *same expected throughput* as the probe best policy.
  - This policy maximizes the amount of knowledge the transmitter has about all channels
  - Probe best policy is greedy, and has very little knowledge of the rest of the channels (other than the best).
  - However, both policies perform the same in terms of expected throughput.

Dynamic Probing Times

- Suppose there is a fixed cost associated with probing
- Question: How often should the transmitter probe a channel?

Dynamic Probing Times

- Suppose there is a fixed cost associated with probing
- Question: How often should the transmitter probe a channel?
- Results:
  1) For a fixed probing interval, can compute optimal probing interval for probing policies discussed previously.

Dynamic Probing Times

- Suppose there is a fixed cost associated with probing
- Question: How often should the transmitter probe a channel?
- Results:
  1) For a fixed probing interval, can compute optimal probing interval for probing policies discussed previously.

  2) When interval length can vary from probe to probe, for probe best and round robin policies:
     - If probed channel is OFF, immediately probe again
     - If probed channel is ON, wait a predetermined interval before probing again
  - Optimal probing interval under probe second best policy is unknown

Conclusion

- Considered channel probing policies, where a transmitter probes a channel, and then chooses which channel to transmit over.
- Using renewal theory, computed average throughput for the probe best policy, probe second best policy, and round robin policy.
- Probe second best policy outperforms the probe-best policy, which was previously shown to be optimal for a slightly different model.

Looking Forward

- What about an optimal policy?
  - **Conjecture**: The *probe second best* policy is the optimal probing policy for fixed probing intervals T.
  - Simulation results / numerical results supporting claim.
  - Proof of Optimality is still under investigation.

Looking Forward

- What about an optimal policy?
  - **Conjecture**: The *probe second best* policy is the optimal probing policy for fixed probing intervals T.
  - Simulation results / numerical results supporting claim.
  - Proof of Optimality is still under investigation.

- Fundamental Limits
  - This talk: focused on optimal channel probing strategies
  - What is the theoretical minimum amount of information exchange required?
  - How do channel probing policies perform in comparison with this fundamental limit?