Cellular Communication

Towers
(Cellular Communication)

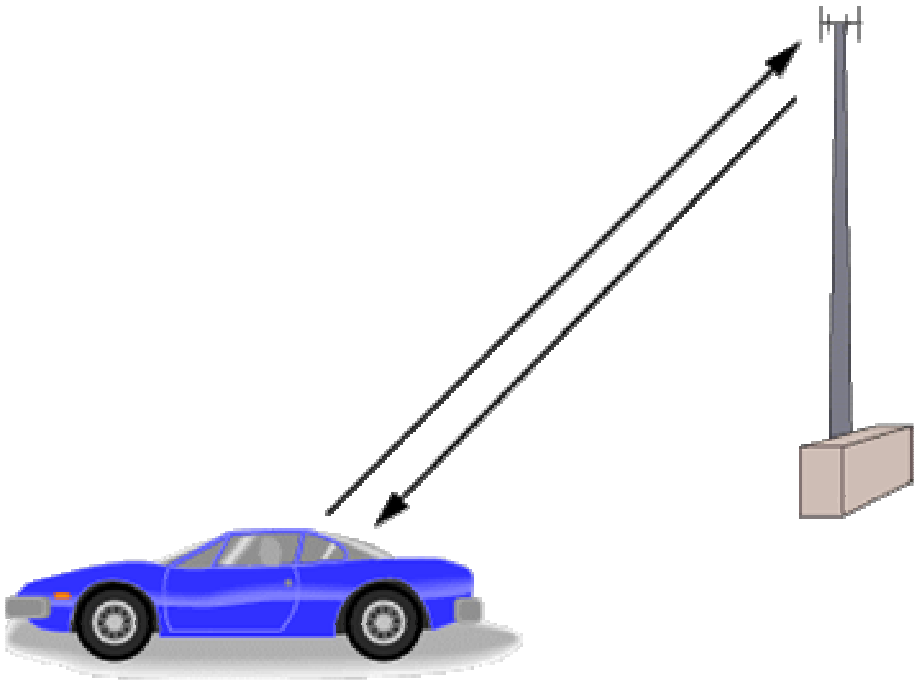

Channels & Frequencies

(Cellular Communication)

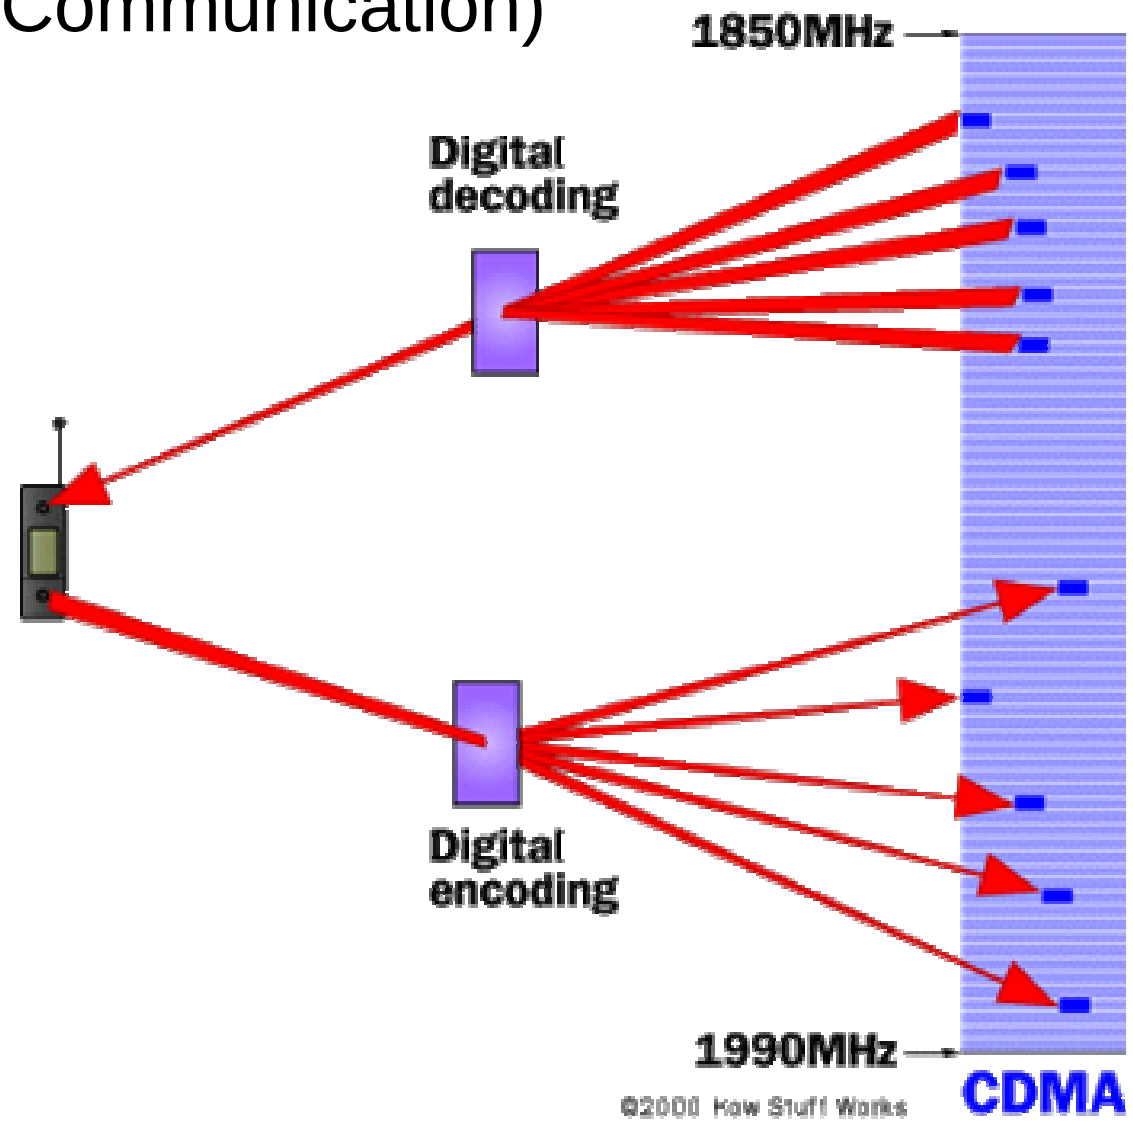

(Animations left out.)

RL for Dynamic Channel Allocation in Cellular Telephone Systems

Satinder Singh, Dimitri Bertsekas

Neural Information Processing Systems
1996

The Problem

Area Divided into Cells

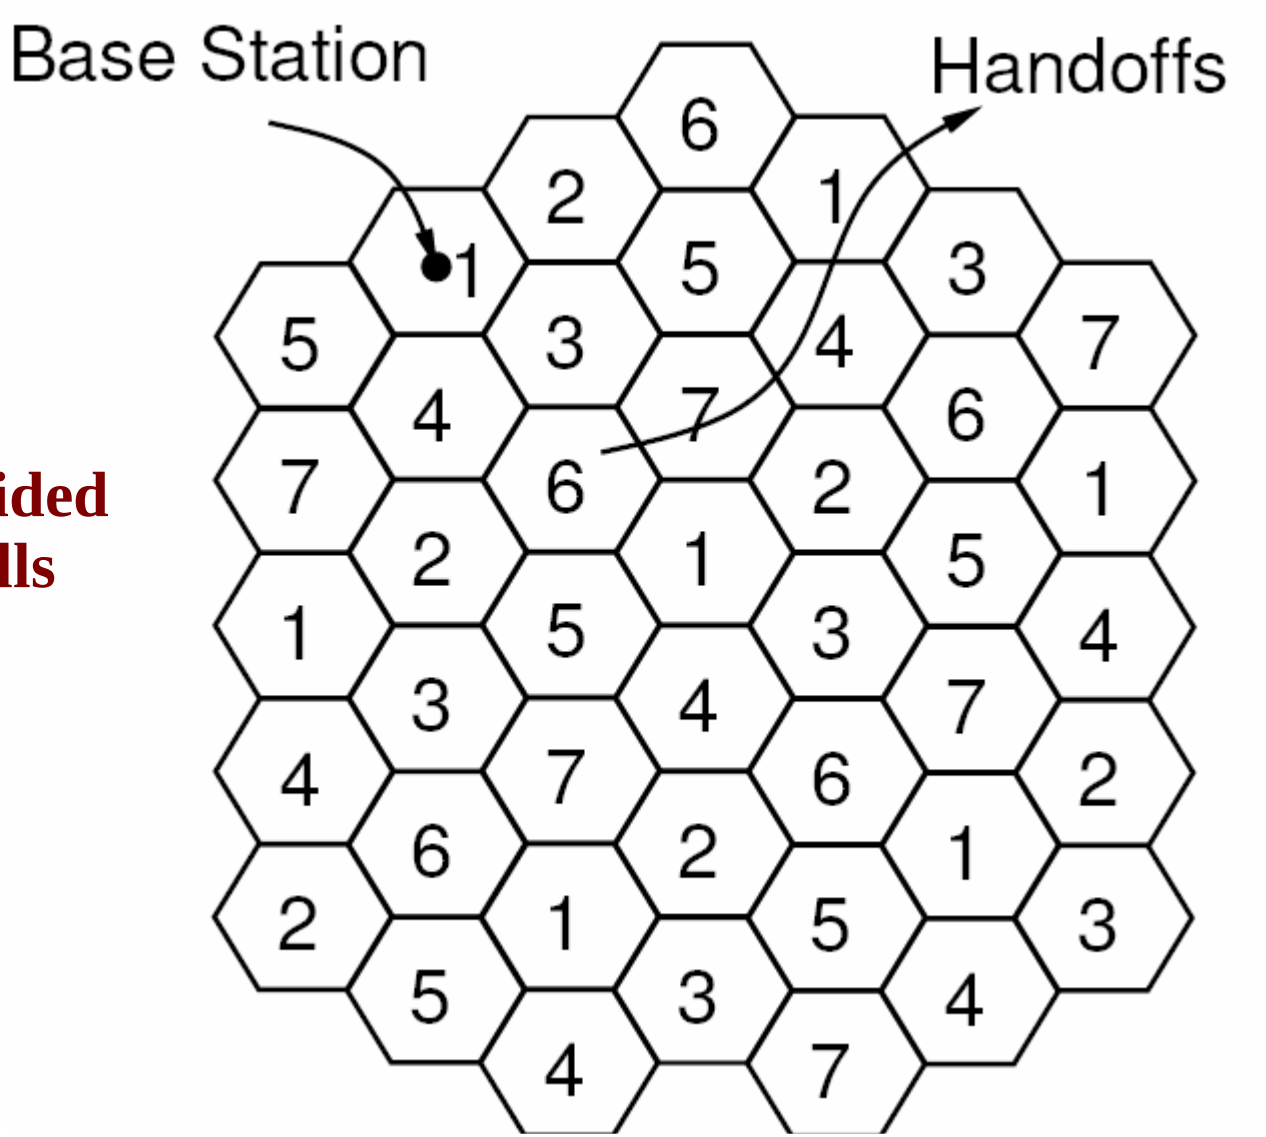

Constraints
(The Problem)

- Channels may be reused in cells far apart
- *Channel Reuse Constraint* is min. distance
- Cell phones may wander from cell to cell

Constraints
(The Problem)

- No channel available:
  - Existing call – dropped call
  - New call – blocked call
  - Blocked calls preferable to dropped calls.

- 49^{49} states – brute force ruled out

Example
(The Problem)

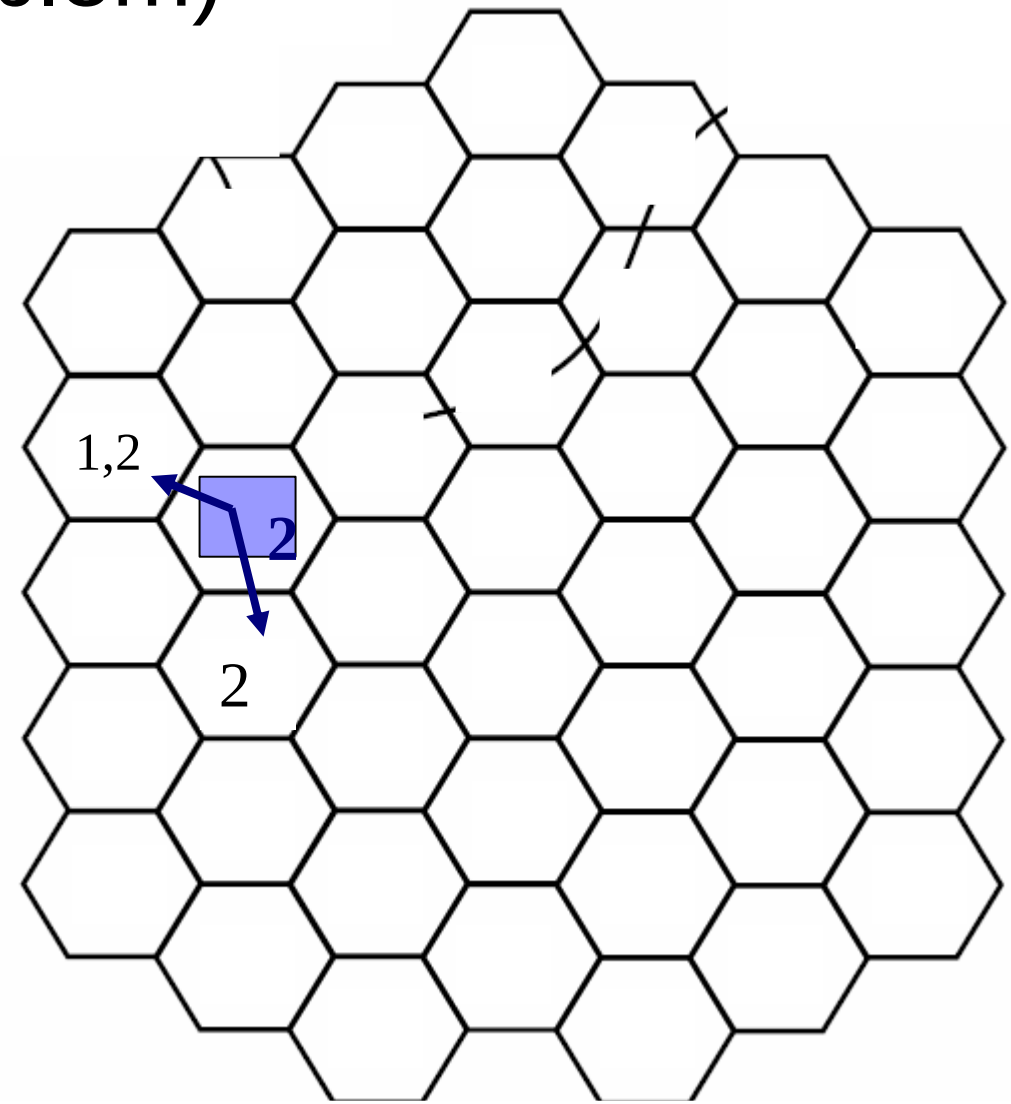

Fixed Assignment Algorithm

- Channels partitioned
- Cells have a fixed subset of channels
- Most popular system
- Busy cell may block calls even when channels available in neighbors.

Dynamic Channel Allocation

- Channels still partitioned between cells.
- Share channels between cells.
- Keeps track of:
  - Occupied and Unoccupied channels.
  - Event type: { arrival, departure, handoff }

Borrowing with Directional Channel Locking

```
partition and assign
on { arrival, handoff } {
    if ( local channel available ) {
        use smallest
    }
    else {
        if ( remote channel available ) {
            borrow largest from neighbor with max free
        }
        else {
            block call
        }
    }
}
on { departure } {
    if ( channel borrowed ) {
        reassign and return
    }
}
```

Borrowing with Directional Channel Locking

- Zhang & Yum (1989)
- Frequent reassignments
- Cascading reassignments / large grids
- Regarded as “best” heuristic algorithm

RL Algorithm

- Unlike Power Grid, rewards are more immediate
- Unlike Robot Problem, system is continual

Decision Function
(RL Algorithm)

Infinite Horizon

$$J = E\left\{\int_0^{\infty} e^{-\beta t} c(t) dt\right\}$$

Ongoing calls at t

Discounting factor

Temporal Difference Learning

"TD learning is a combination of Monte Carlo ideas and dynamic programming ...TD methods can learn directly from raw experience without a model of the environment's dynamics."

Book:

$$V(s_t) \leftarrow V(s_t) + \alpha\left[r_{t+1} + \gamma V(s_{t+1}) - V(s_t)\right].$$

Paper:

$$J_{new}(\tilde{x}) \quad = \quad (1-\alpha)J_{old}(\tilde{x}) + \alpha\left(c(x,a,\Delta t) + \gamma(\Delta t)J_{old}(\tilde{y})\right)$$

RL Algorithm

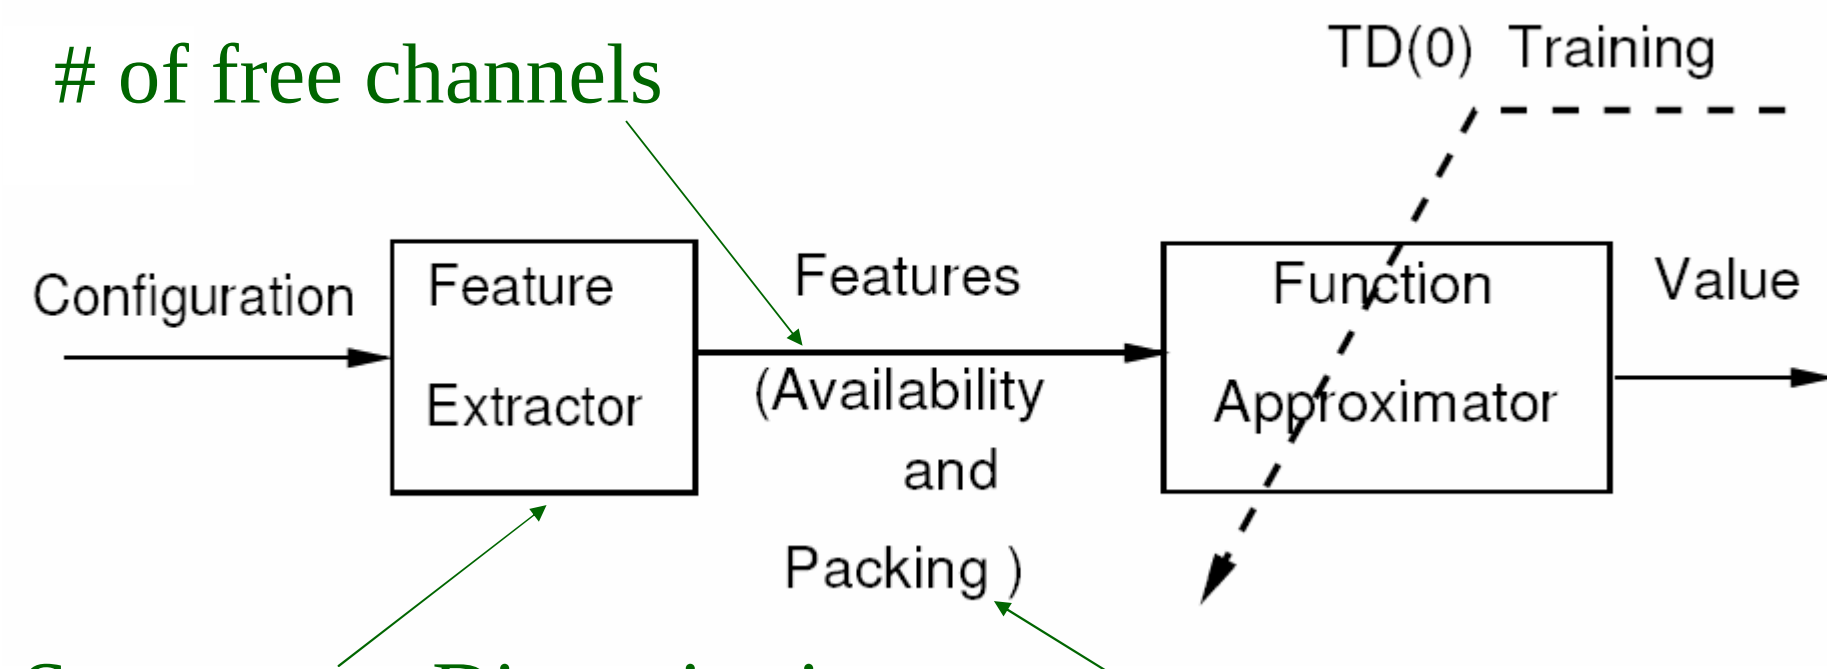

RL Algorithm

```
partition and assign
on { arrival, handoff } {
    if ( local channel available ) {
        use smallest
    }
    else {
        if ( remote channel available ) {
            borrow from neighbor with maximum expected reward
        }
        else {
            block call
        }
    }
}
on { departure } {
    if ( channel borrowed ) {
        reassign and return with minimum expected penalty
    }
}
```

Fixed-Rate Traffic Results

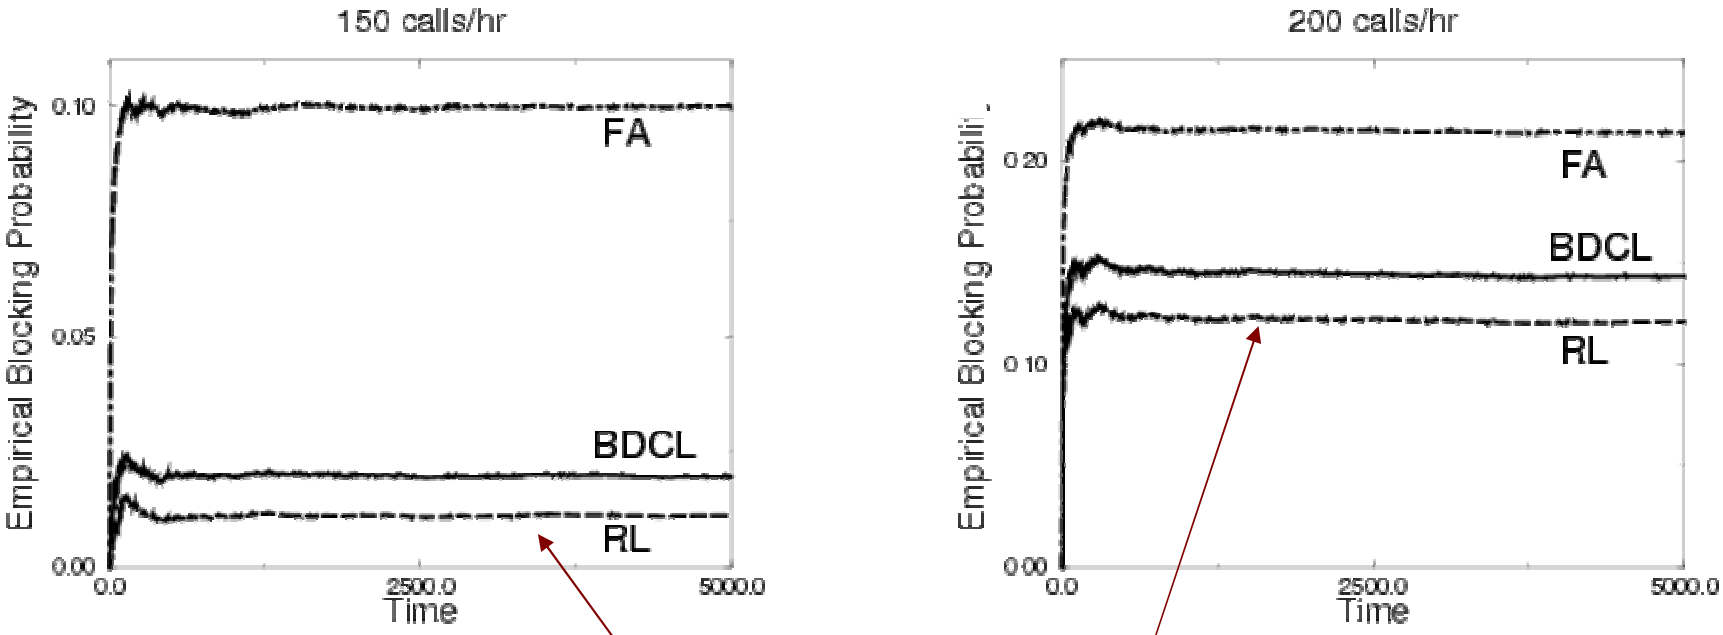

- Quick learning due to TD(0) function.
- Learning curve “disadvantage” ruled out.

Fixed-Rate Traffic Results

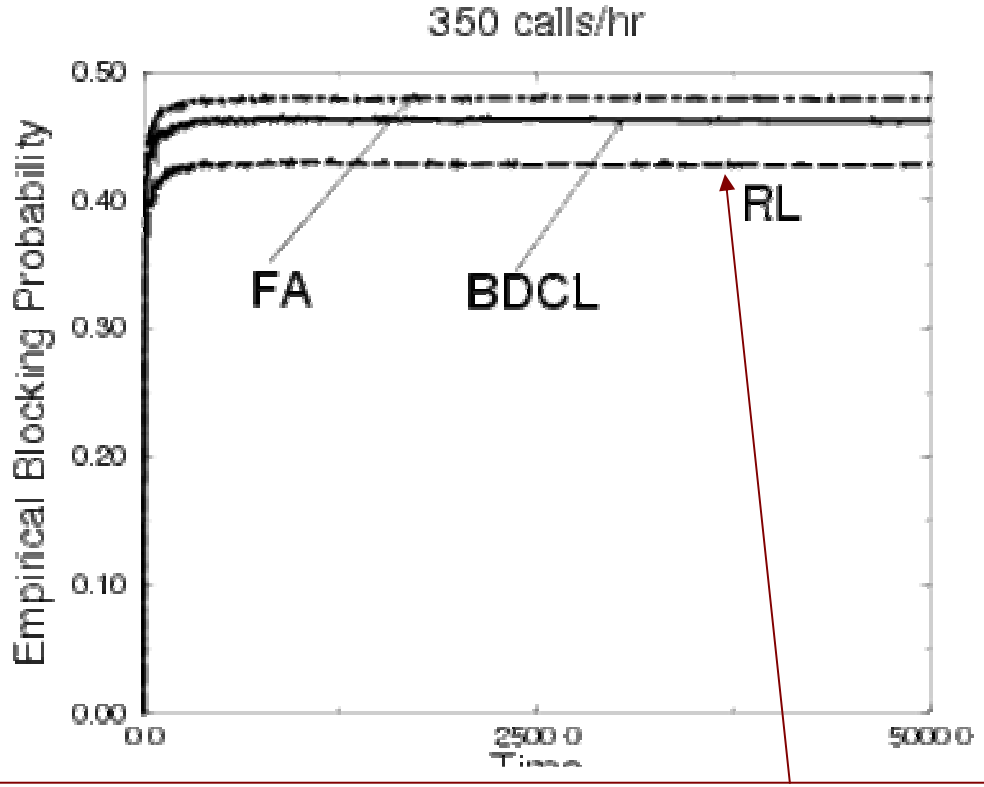

- FA optimal as # calls → ∞
- No short term fluctuations to exploit.

Variable Traffic Results

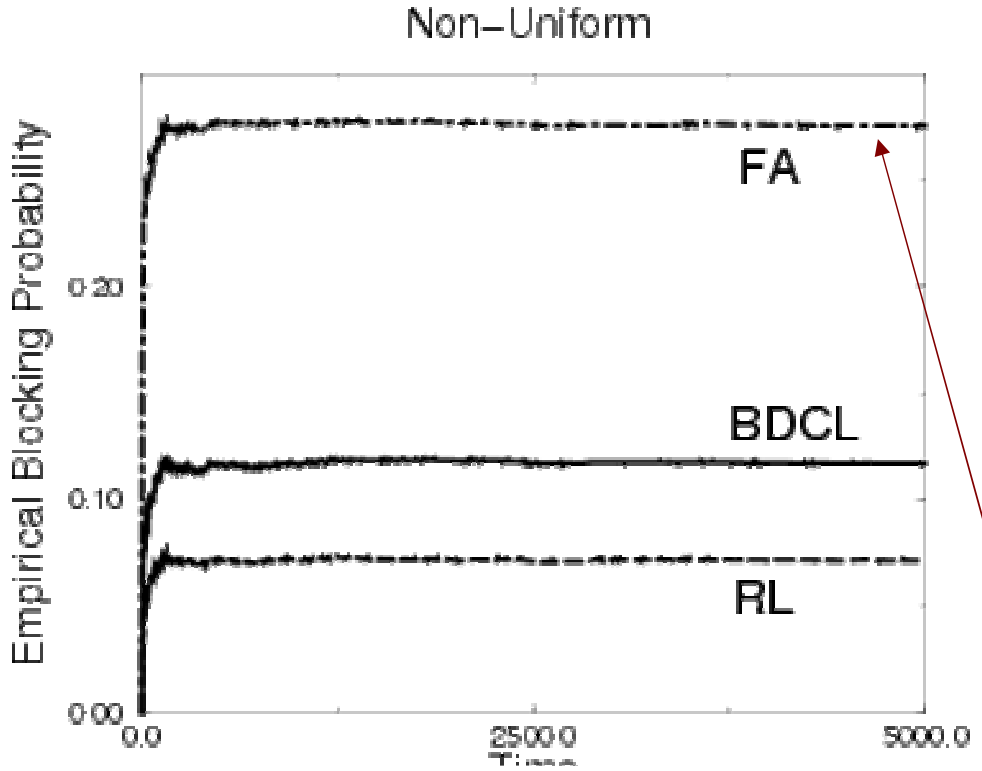

- FA, BDCL ignore traffic patterns
- Learning allows pattern recognition

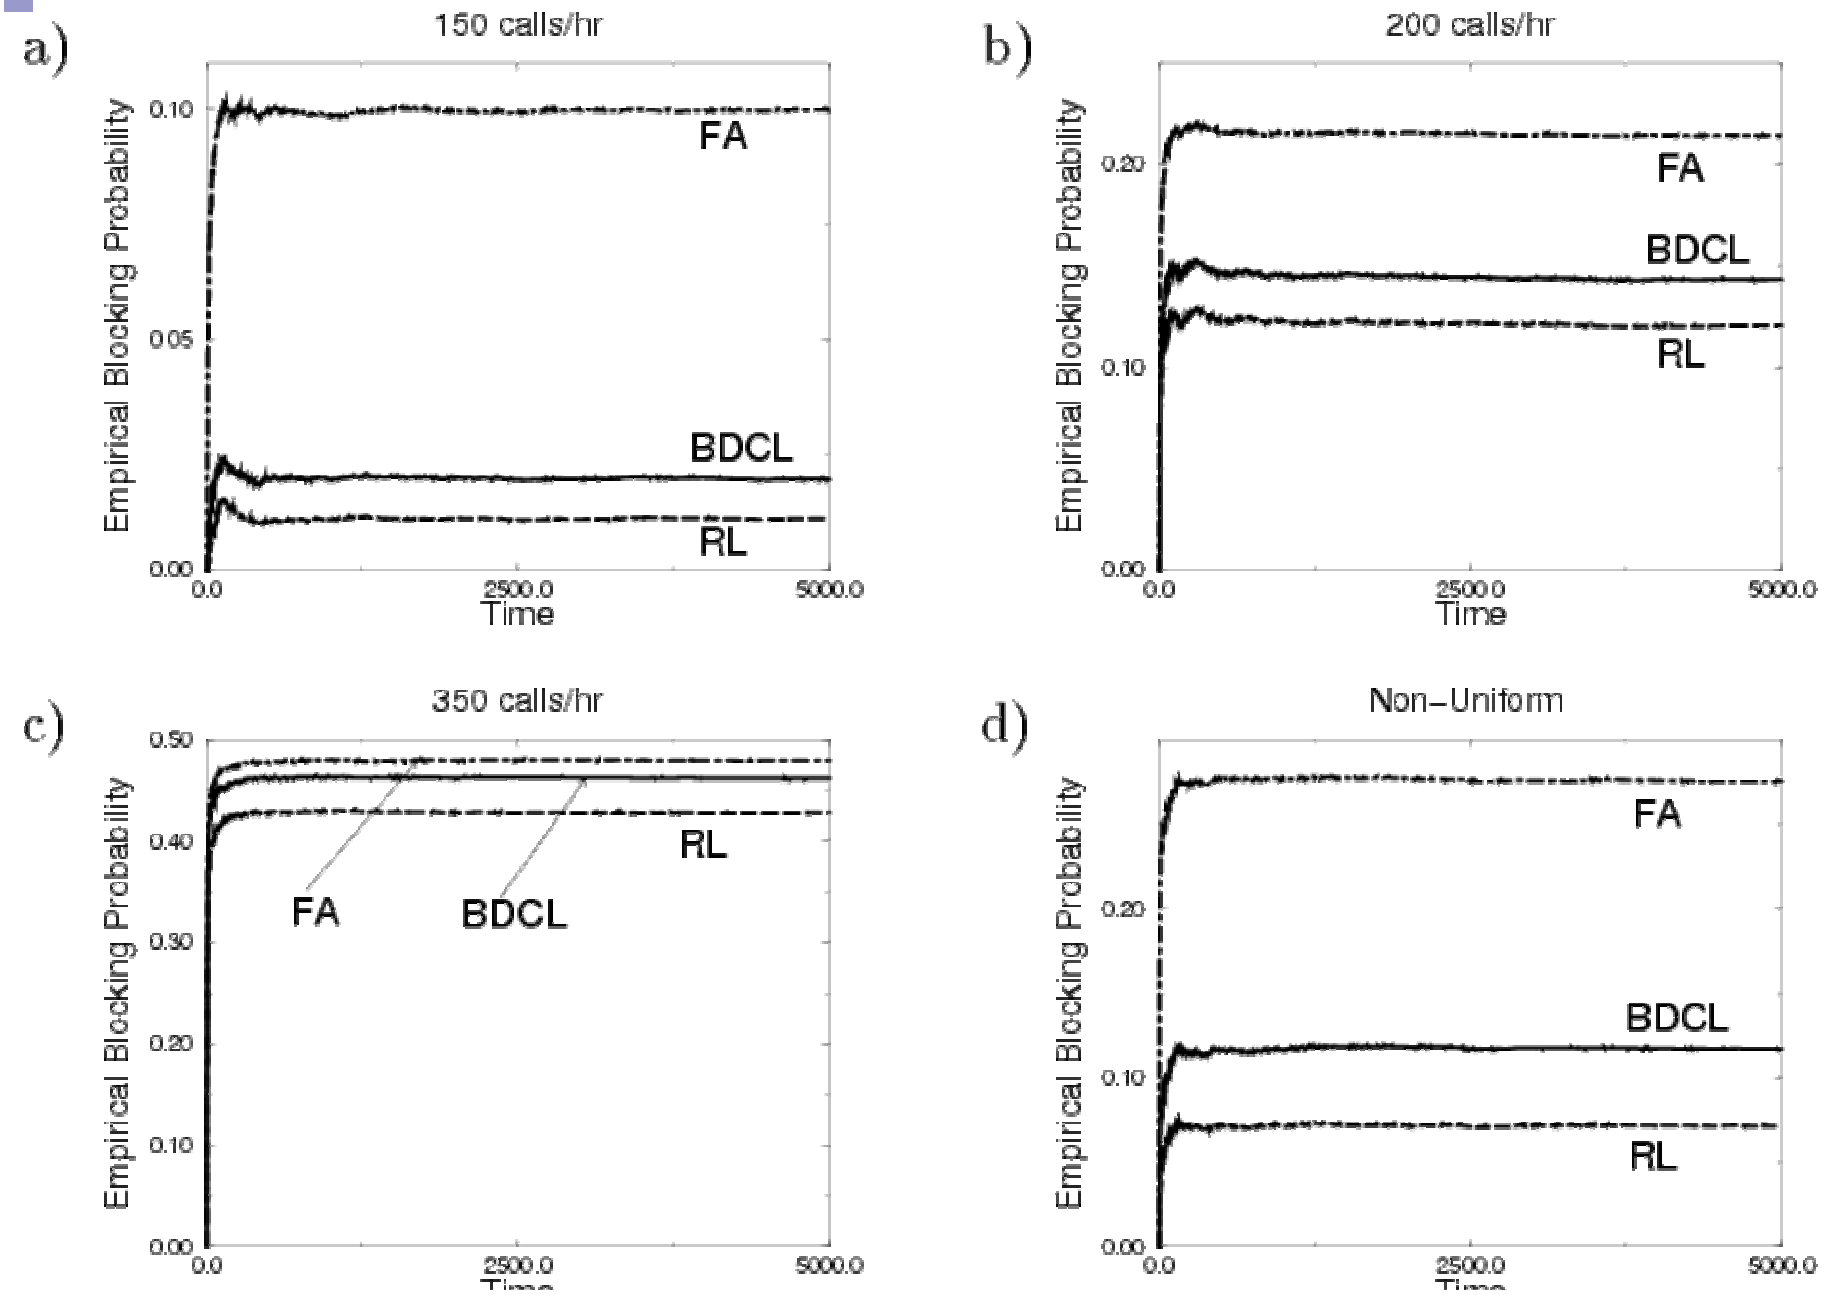
a)
150 calls/hr
Empirical Blocking Probability
0.10
0.05
0.00
FA
BDCL
RL
0.0
2500.0
5000.0
Time
b)
200 calls/hr
Empirical Blocking Probability
0.20
0.10
0.00
FA
BDCL
RL
0.0
2500.0
5000.0
Time
c)
350 calls/hr
Empirical Blocking Probability
0.50
0.40
0.30
0.20
0.10
0.00
RL
FA
BDCL
0.0
2500.0
5000.0
d)
Non–Uniform
Empirical Blocking Probability
0.20
0.10
0.00
FA
BDCL
RL
0.0
2500.0
5000.0

Variable Traffic Results

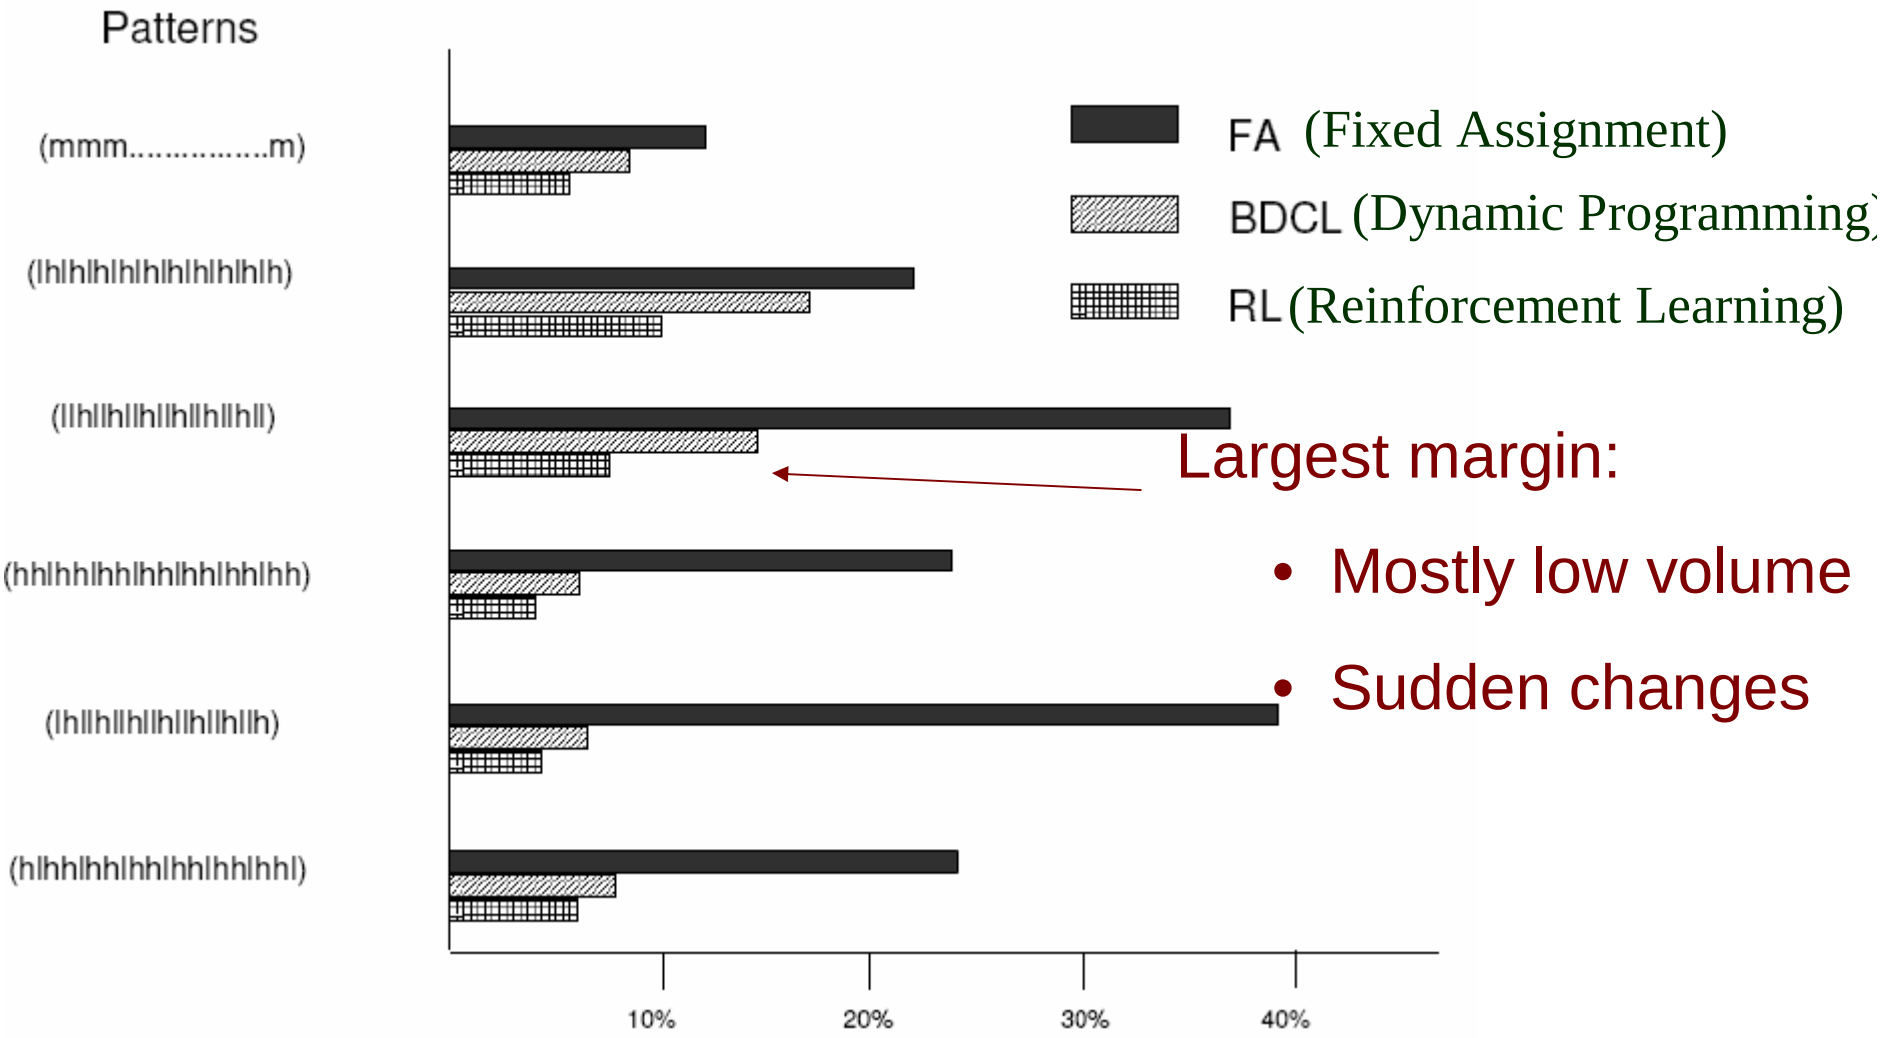

Largest margin:

- Mostly low volume
- Sudden changes

L = low, M = medium, H = high

Decentralized Training

(Future Work)

- Policies and values are on a per-cell basis.
- Small networks pose a problem, decisions impact other cells. (Competition.)
- Large networks more feasible, cells with large distance have little relation.

Related Works

- Admission Control

  S. Singh, T.X. Brown, H. Tong in NIPS 1998

  http://www.eecs.umich.edu/~baveja/Papers/NIPS98Hui.ps.gz

- Predicting Temporal Differences (H0)

  R. Sutton in Machine Learning, 3:9-44