Cornell University
Operations Research and
Information Engineering

Real-Time Control of Ambulance Services

Shane G. Henderson

Joint work with Matt Maxwell, Huseyin Topaloglu

Thanks to:

NSF CMMI 0758441,
Optima Corporation,
Toronto, Melbourne, Edmonton EMS,
Armann Ingolfsson, Andrew Mason

Pressure on Ambulances

- Traffic congestion
- Increasing call volumes
- Ambulance diversion
- Delays in handovers to emergency departments
  - Can double the time required for a call
- Long term measures needed, but in the meantime…

Redeployment

- AKA system-status management, move up
- Enabled through live status, travel times on road networks, arrival rates in space and time

Outline

- Existing methods
- Approximate DP
- Tuning ADP
- Another service-system application
- Research challenges for simulation folks

SIREN Live Real time View

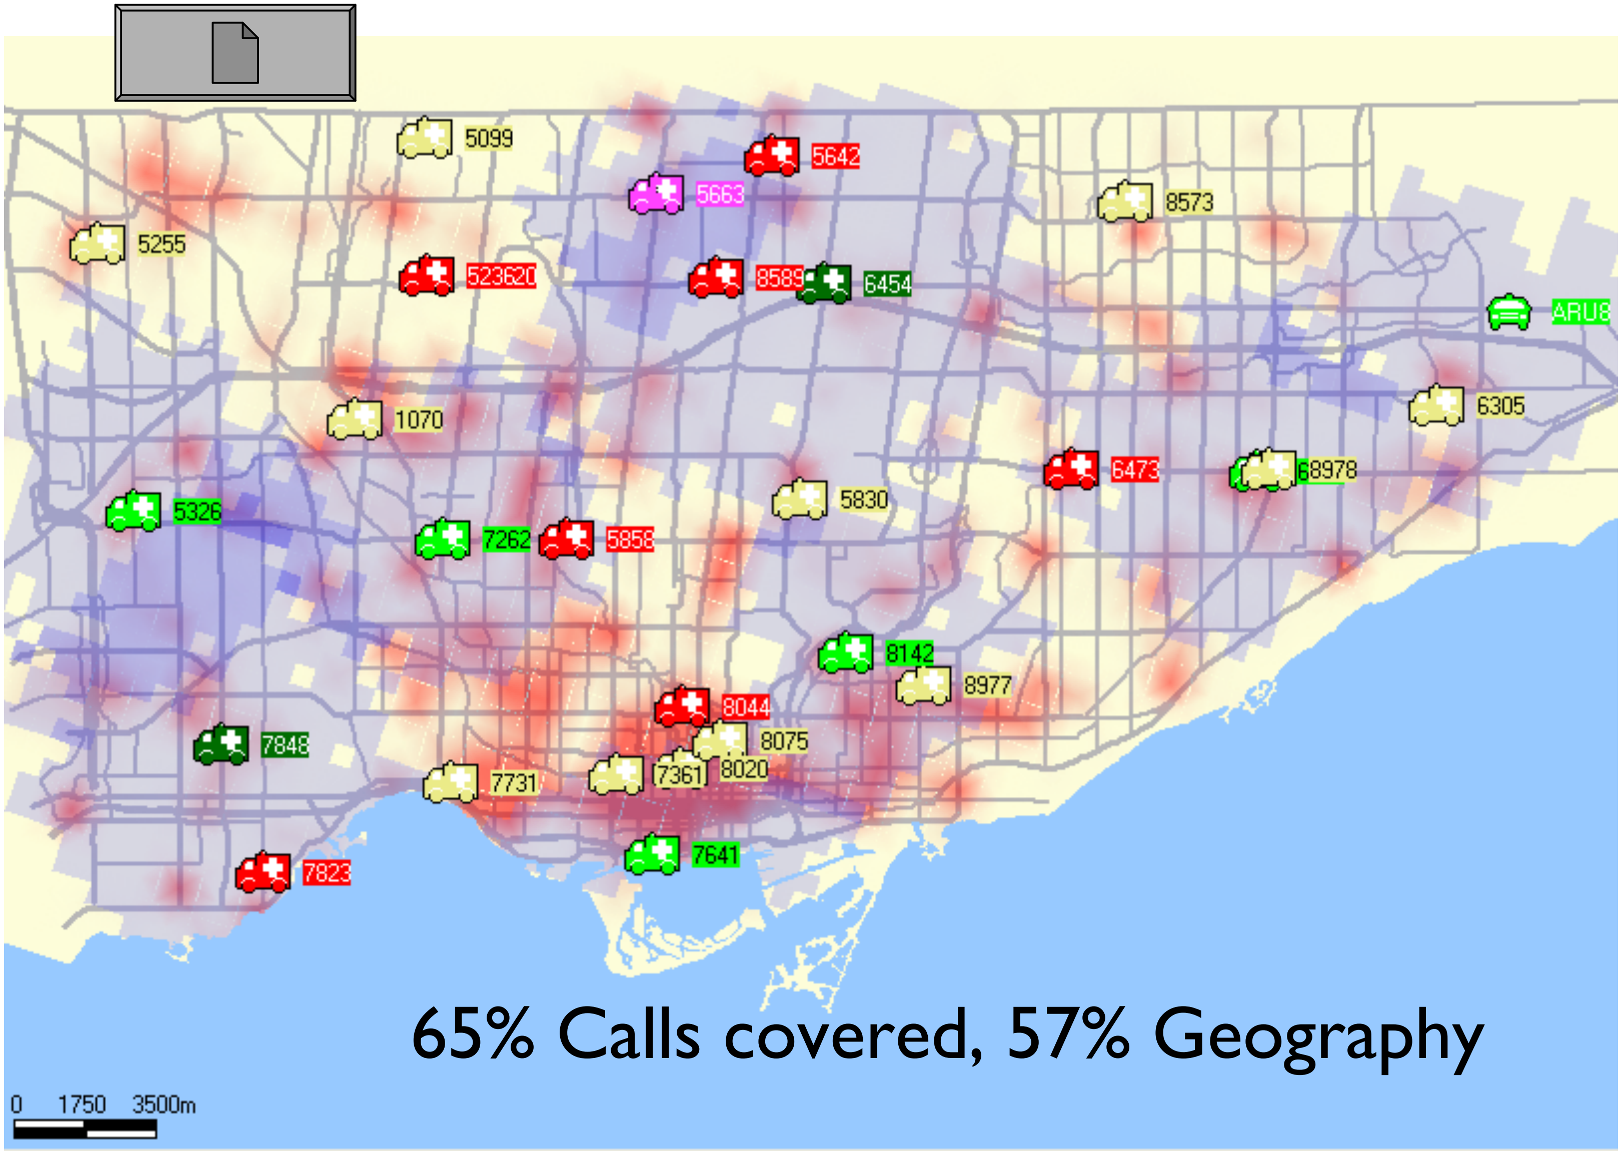

65% Calls covered, 57% Geography

SIREN Live showing Vehicle Moves

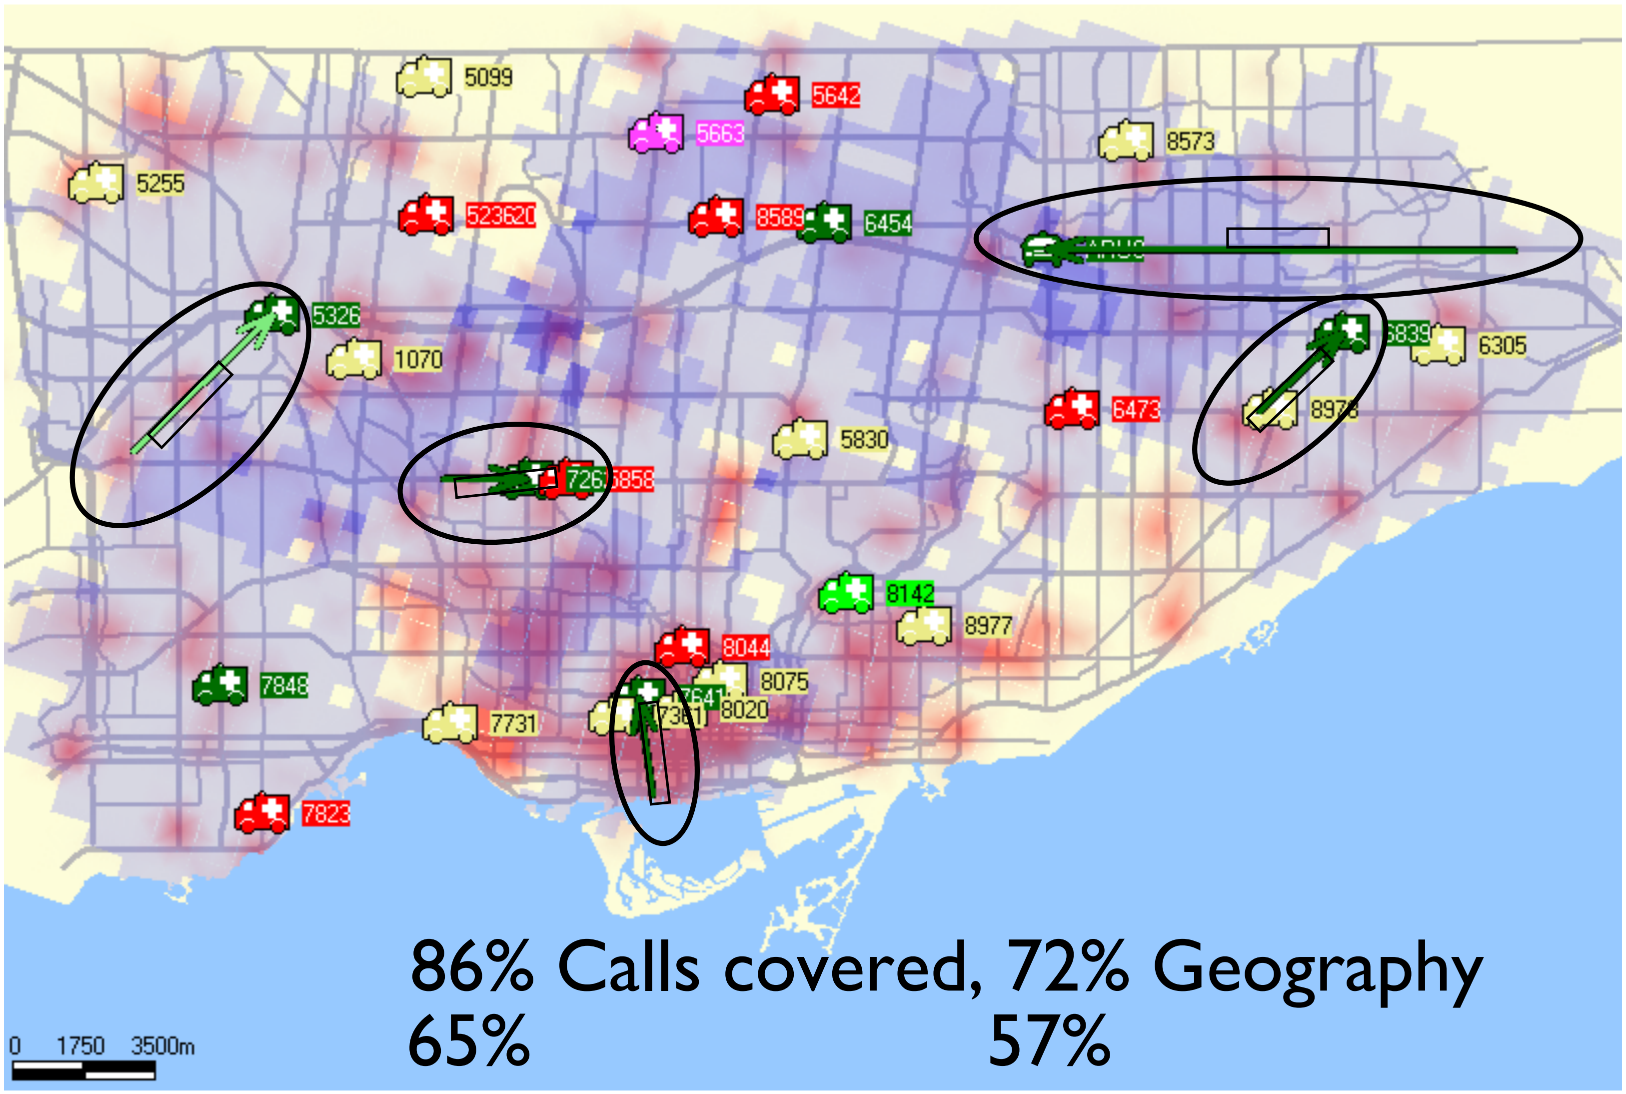

86% Calls covered, 72% Geography
65% 57%

Redeployment: Current Methods

- Lots of “static” locating using IP and sim.
- Spreadsheet tools? Unclear basis
- Solve real-time IP’s, e.g., Montreal, Optima
- Compliance Tables
    - Generate a lookup table saying where *n* free ambulances should be positioned, *n* = 1, 2, …
    - Dispatch to match those locations
- Exact Dynamic Programming
    - Berman et al 1970s
    - Zhang, Mason and Philpott

Approximate Dynamic Programming

- Have a function, V say, that gives the value or quality of a configuration. Use greedy policy wrt V, i.e.,
- When want to redeploy an ambulance, look for configuration that maximizes V
- Keeping in mind that ambulance may not get there before something changes
- So choose action that maximizes
  E(immediate reward + V)

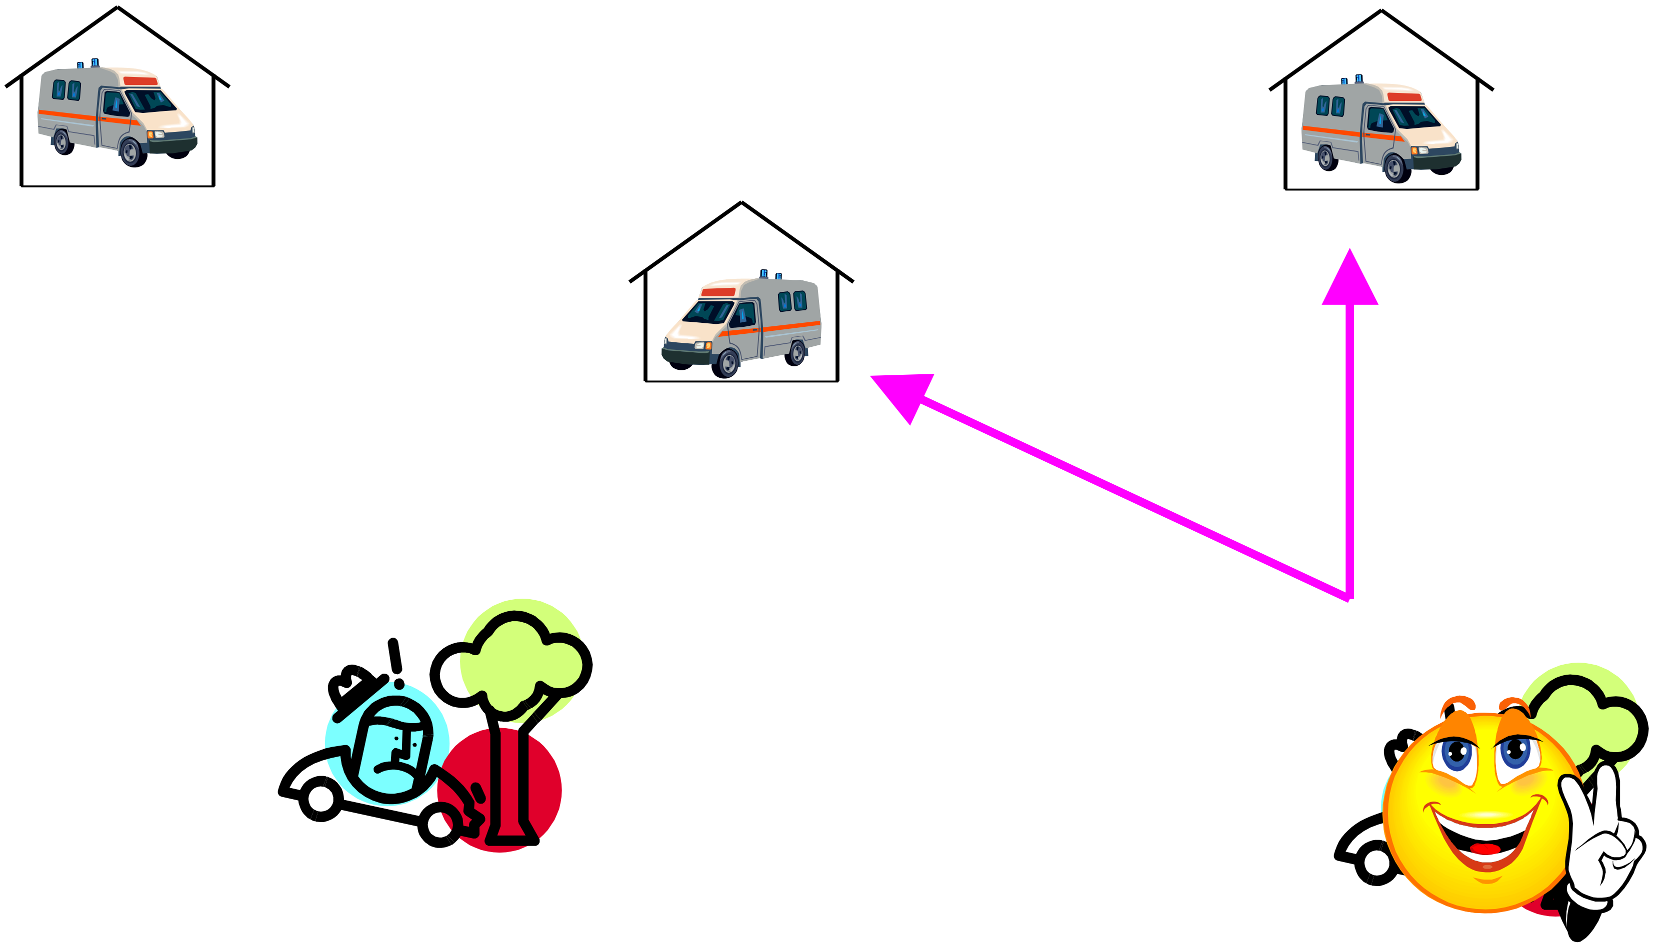

Assume send ambulance to 1

Do 10 times:

Simulate immediate future, and look up V for final ambulance positions, status

Compute average of V values

= 50.1 say

Assume send ambulance to 2

Do 10 times:

Simulate immediate future, and look up V for final ambulance positions, status

Compute average of V values

= 63.2

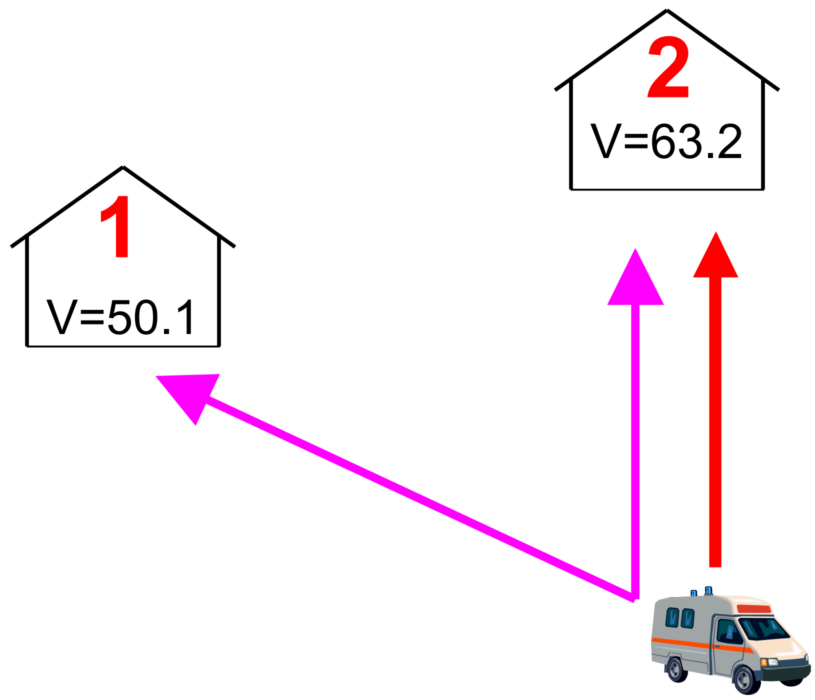

Where Does V Come From?

- Can't store V values for every possible state, so need to approximate V
- We use $V = r_1 V_1 + \ldots + r_n V_n$
- $V_1, V_2, \ldots, V_n$ are fixed basis functions that we choose
- Choose r_i's in initial training stage

Basis Functions (V_i's)

- For each base, rate of calls arriving to surrounding area, that will likely (Erlang loss) be missed, assuming vehicles reach current destinations

Choosing Coefficients

Training was approximate value iteration (TD learning, other tricks came later)

1. Choose some *r_i's* … gives a V
2. Simulate performance of V
3. V was supposed to match observed performance (principle in DP)
4. Perform a regression (r_i's) to try to get V to match observed performance

"Convergence"

Surprise?

A powerful function is given by the sum of the basis functions, and…

regression doesn't find it (nor does LP)

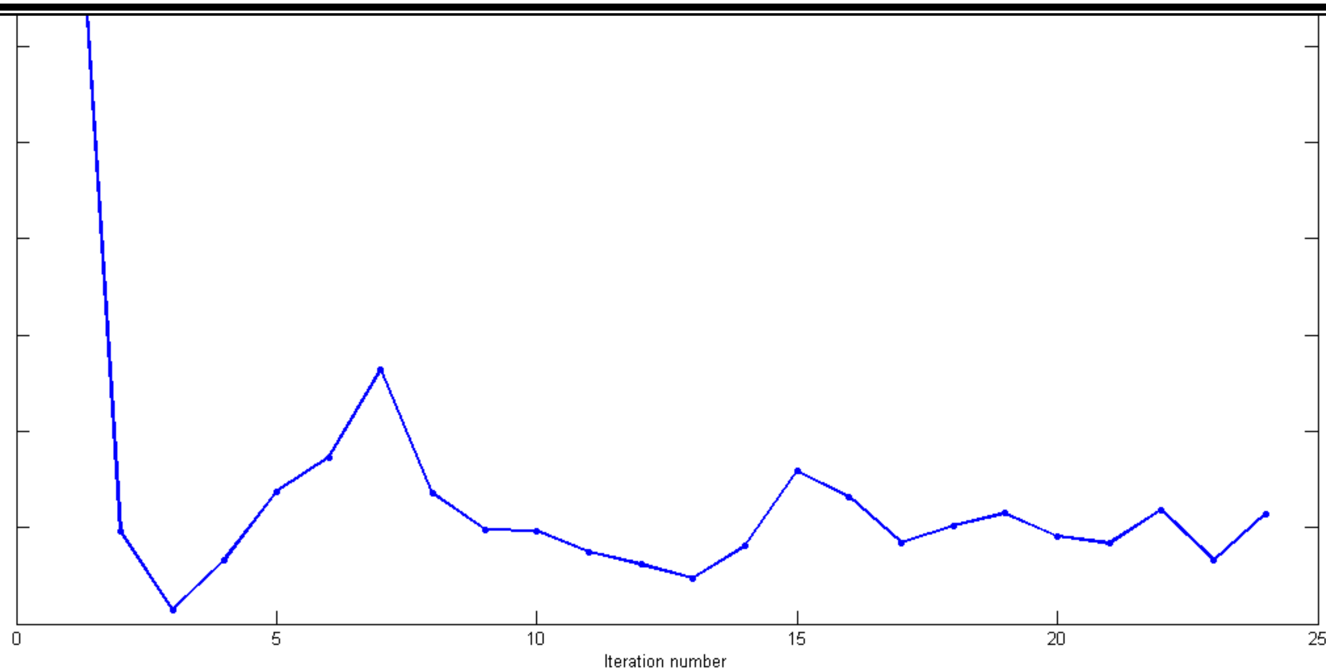

Direct Search

- So we tried a simulation optimization to try to find a good value function
- Nelder-Mead algorithm searching 11-dimensional space (Edmonton) for coefficients
- No attempt to be fancy!
- Each function evaluation takes 40 - 60 minutes
- Would take about a year, so …

Post-Decision State ADP

- Use post-decision state formulation (Powell and van Roy 2004)
- i.e., take limit of micro-simulations as their length goes to 0
- Don’t do micro-simulations, just compute V for post-decision state
- Now sim-opt is feasible
- One short week and a half later…

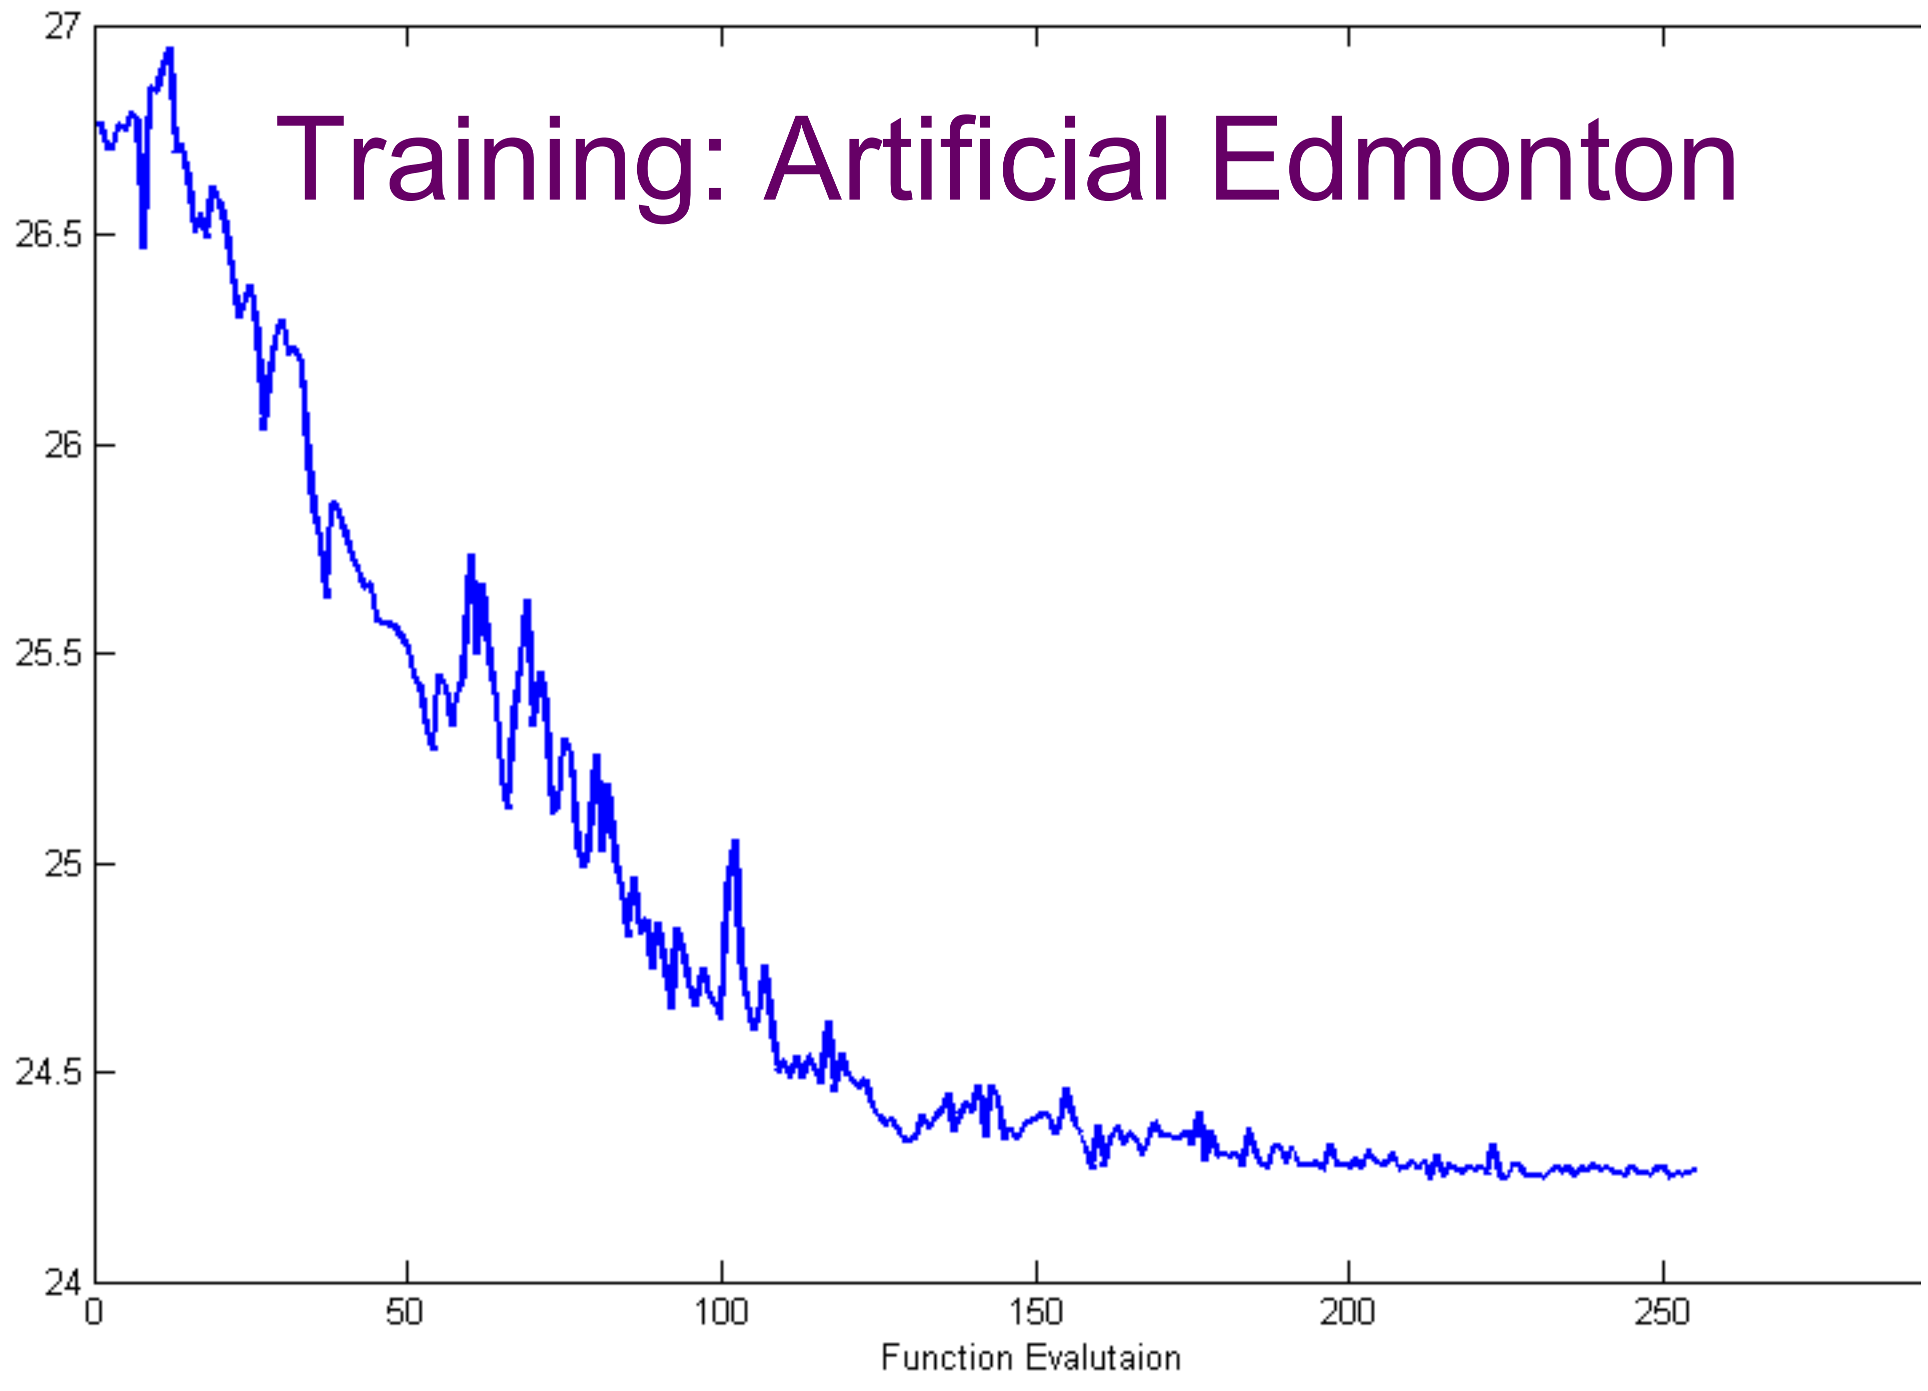
Training: Artificial Edmonton
27
26.5
26
25.5
25
24.5
24
0
50
100
150
200
250
Function Evalutaion

Training: Artificial Melbourne

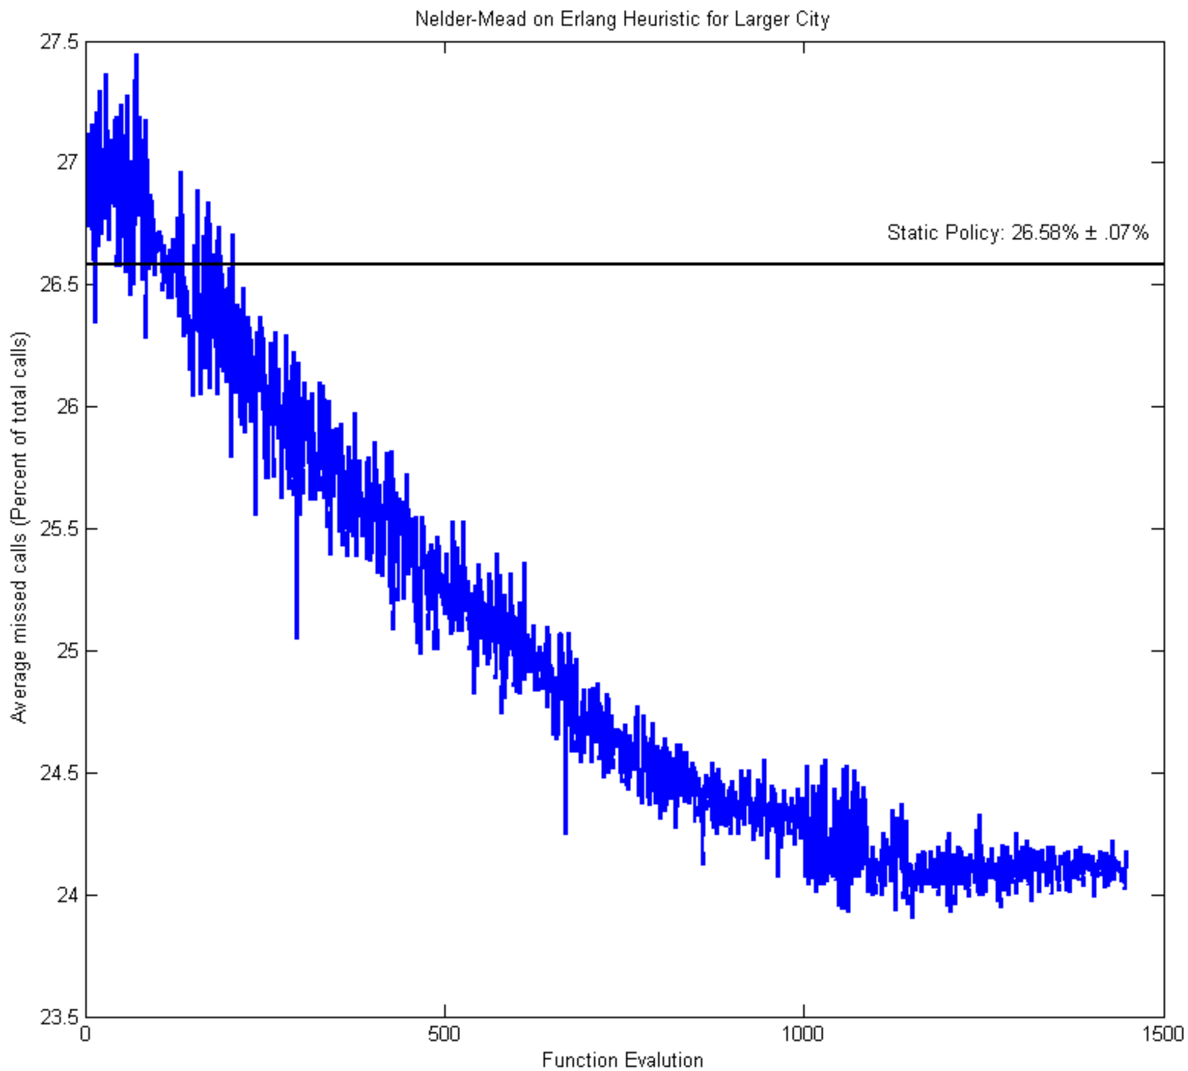

Missed Calls: Artificial Edmonton

- Reasonable Static Policy: (32.3 ± 0.1)%
- Best ADP policy using regression-based search (26.5 ± 0.2)%
- ADP using sim opt: (24.4 ± 0.2)%
- This is just redeploying newly free ambulances. No wake ups!

ADP Folks Know About This…

- ADP folks are aware that regression isn't always effective
- Average Tetris scores:
  - 20K using regression (Desai, Farias, Moallemi 2010)
  - 350K using cross-entropy based simulation optimization (Szita, Lorincz 2006)

What Goes Wrong?

- Regression tries to fit value function globally, but local changes are the key to good performance
- Regression matches value function to observations, but we care instead about performance of greedy policy induced by approximation
- So perform slow simulation optimization

Simulation in ADP

	Using Micro Simulations	*Post Decision State*
Training	Optimization over (simulation model + micro sims)	Optimization over simulation model
Real Time	Track system state Micro sims + V for decisions	Track system state V for decisions

Outline

- Existing methods
- Approximate DP
- Tuning ADP
  - Practically significant improvements
  - Real-time calculations are fast (< 1 sec)
  - Tuning SLOW: OK for application
  - Use ~~regression~~ sim. opt. for tuning
- Another service-system application
- Research challenges

One View of a Call Centre

Our View of a Call Centre

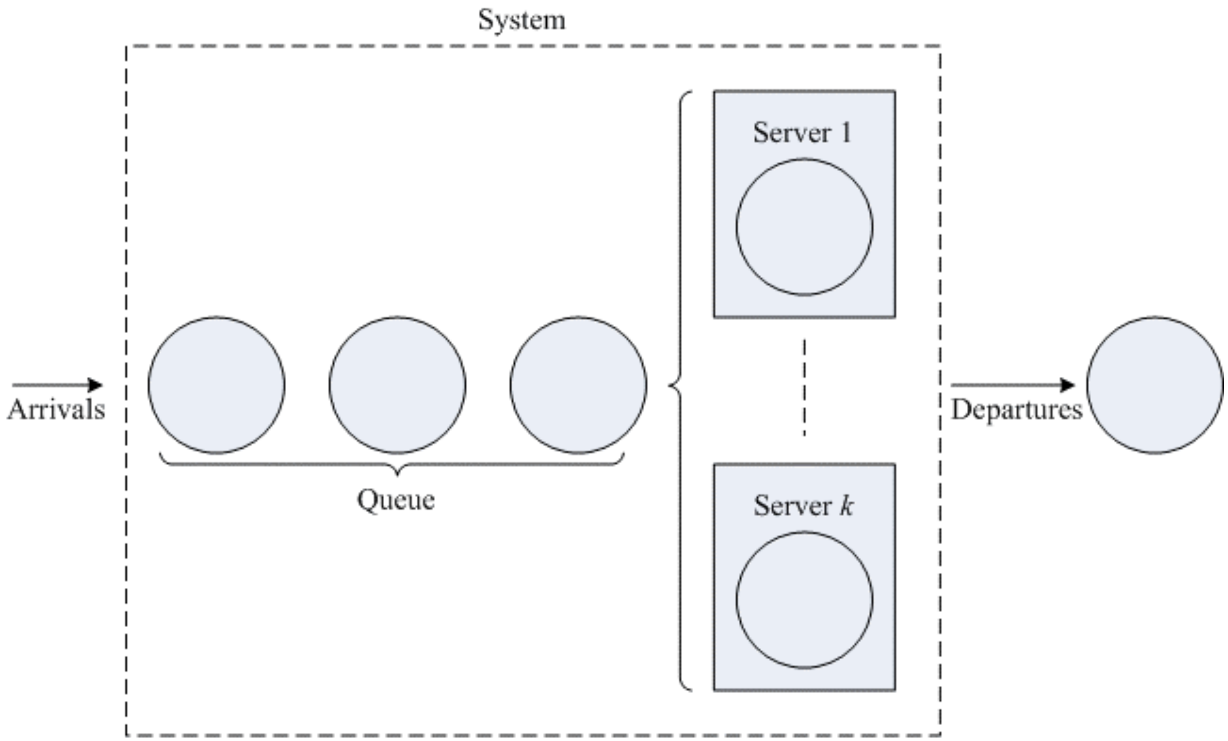

$$\text{Number of servers required} = \frac{\lambda}{\mu}$$

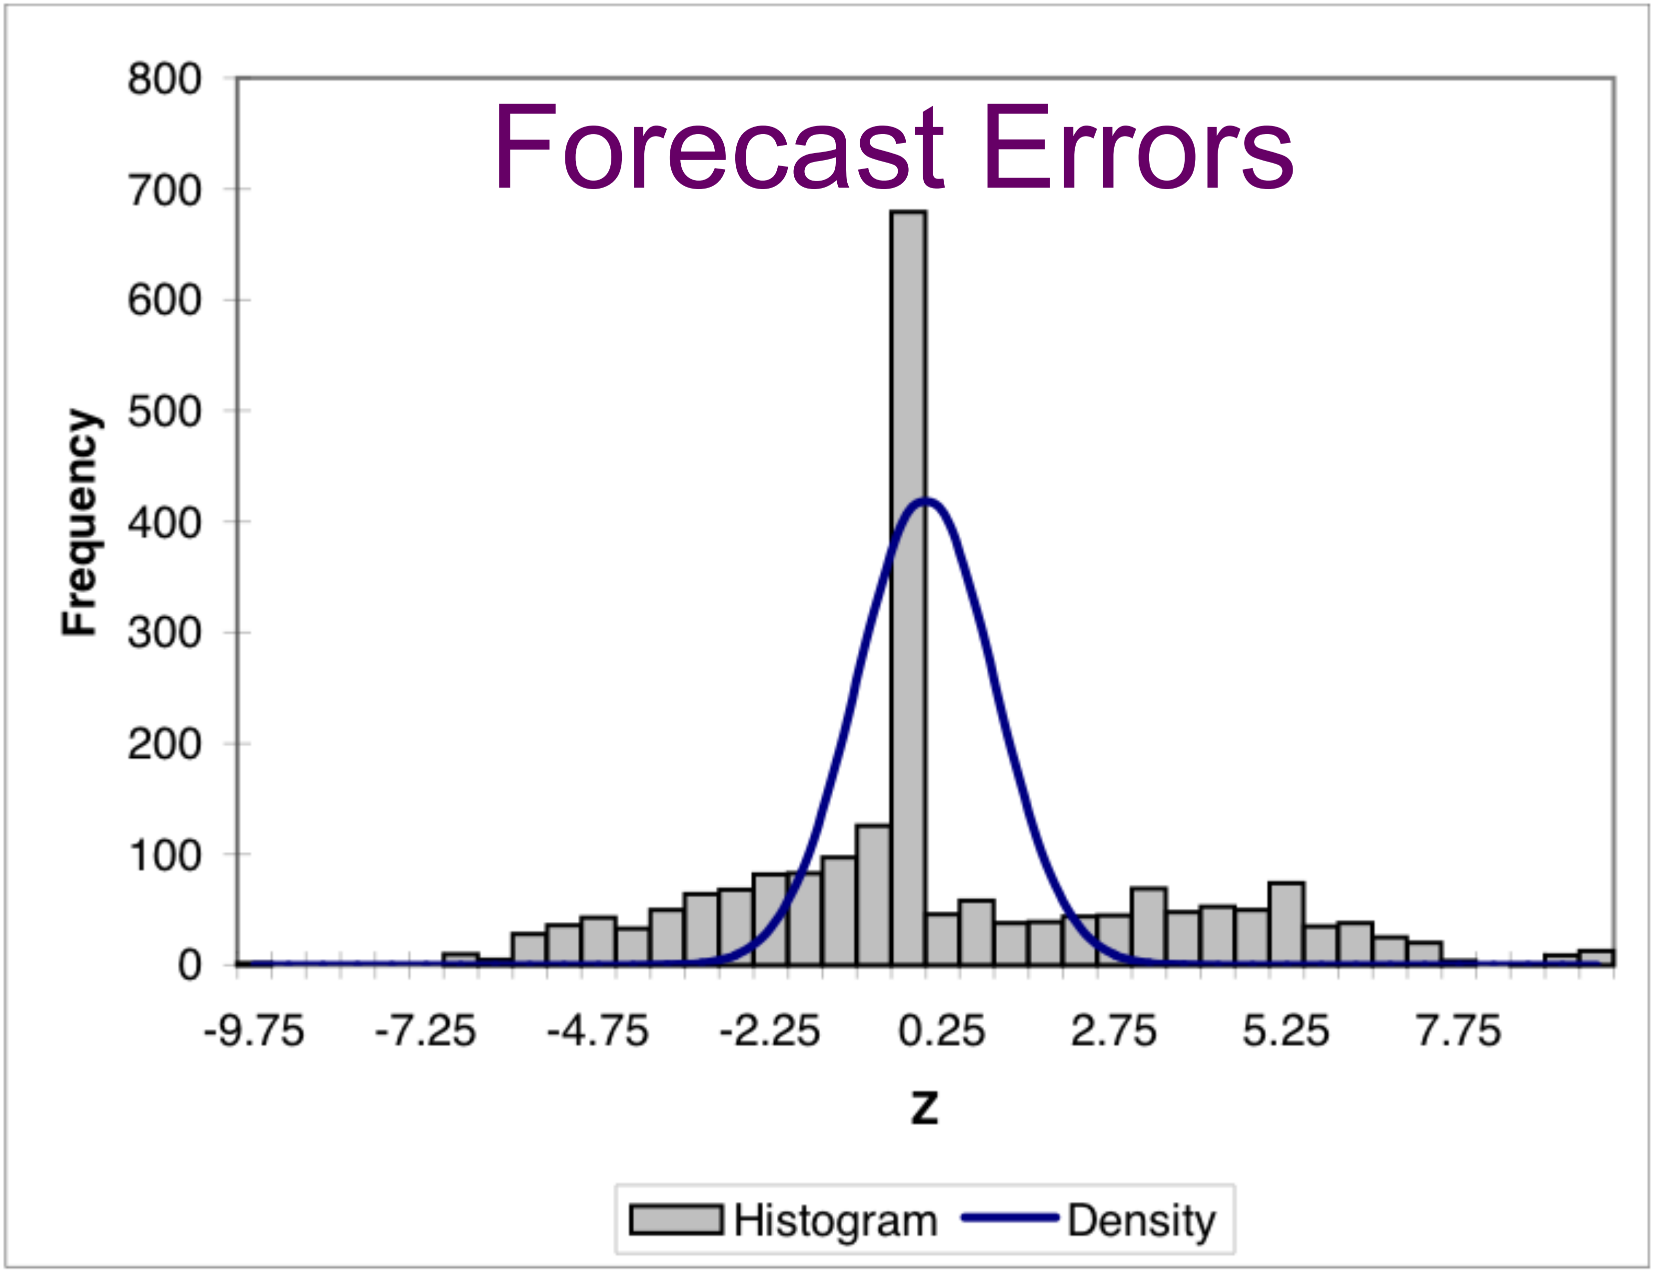

Steckley, H. and Mehrotra. PEIS 2009

And just as bad…

$$\text{Number of servers required} = \frac{\lambda}{\mu} + 2\sqrt{\frac{\lambda}{\mu}}$$

- Service rate varies between servers
- Tremendous levels of agent absenteeism

Consequences

- Forget queueing theory at CLT scale?
- Just increase number of servers?
  - Most days you have agents sitting idle
  - Some days agents cannot keep up
  - Good customer service?
- Real-time control of number of servers
  - Assuming you can get a contract, connect servers in as needed (outsourcing, other)
  - Easy control policy for this queue, but…

What About This One?

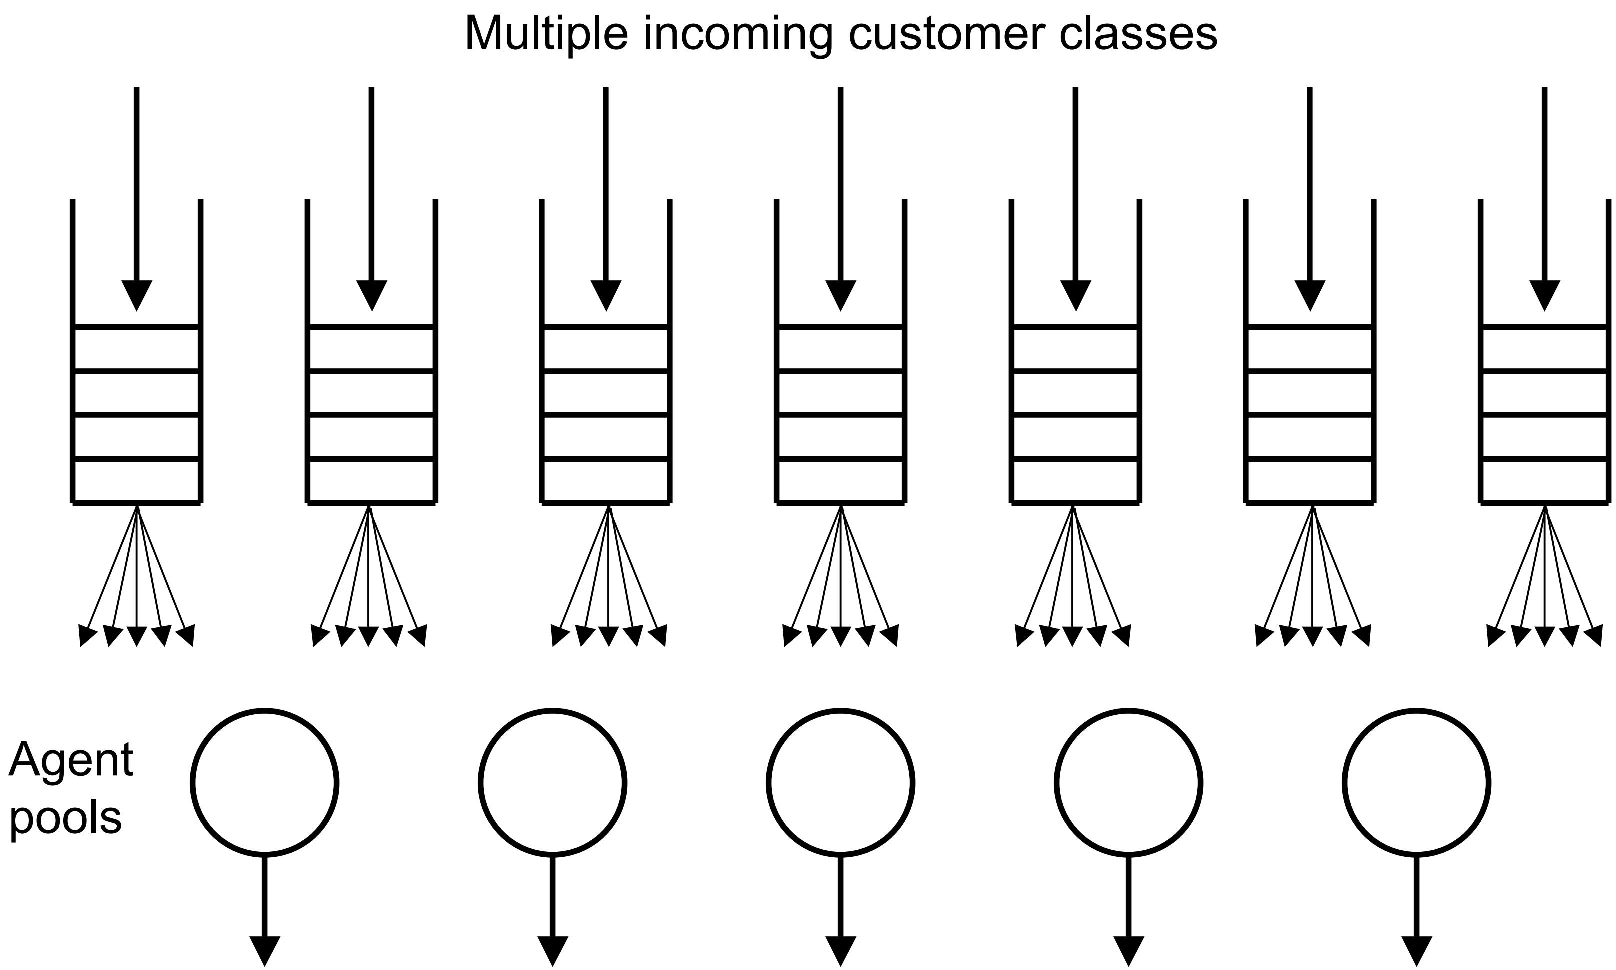

Lots of Service Systems…

- Require real-time interaction between customers and servers
- Have large forecast errors in customer arrival rates
- Have high levels of service capacity variability (both numbers and service rates)
- Require high levels of customer service
- *Real time control via parameterized policies?*
- *And how to do staffing knowing you will use that policy?*

Research Challenges

- Work with real organizations to try to help them (too often overlooked)
- Formulate as finitely parameterized policies - problem specific
- Search coefficient space for good policies
  – Customized sim-opt methods for ADP and other policy tuning
  – Careful statistics needed for real-time control; don't chase noise
- Optimality gap bounds
  – Brown, Smith and Sun (2010), or ad-hoc

References

- Brown, D.B., J. Smith and P. Sun (2010). Information relaxations and duality in stochastic dynamic programs. Operations Research, 58(4), p. 785-801.
- Henderson and Mason. 2004. Ambulance service planning: simulation and data visualization. In M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, eds, Operations Research and Health Care: A Handbook of Methods and Applications, 77-102. Kluwer Academic, Boston.
- Ingolfsson and coauthors at Edmonton – many papers
- Maxwell, Henderson and Topaloglu. 2010. *Proceedings of the 2010 WSC.*
- Maxwell, Henderson and Topaloglu. 2010. Tuning approximate dynamic programming policies for ambulance redeployment via direct search. Submitted.
- Maxwell, M. S., S. G. Henderson and H. Topaloglu. 2011. Equivalence results for approximate dynamic programming and compliance table policies for ambulance redeployment. Submitted.
- Maxwell, M. S., S. G. Henderson and H. Topaloglu. 2011. A bound on the performance of ambulance redeployment policies. Working paper.
- Maxwell, Restrepo, Henderson and Topaloglu. 2010. Approximate dynamic programming-based ambulance redeployment. *INFORMS Journal on Computing* 22 266-281
- W. B. Powell and B. Van Roy. 2004. Approximate Dynamic Programming for High-Dimensional Dynamic Resource Allocation Problems. In *Handbook of Learning and Approximate Dynamic Programming*, edited by J. Si, A. G. Barto, W. B. Powell, and D. Wunsch, Wiley-IEEE Press, Hoboken, NJ, 2004, pp. 261-279.
- Steckley, S. G., S. G. Henderson and V. Mehrotra. 2009. Forecast errors in service systems. *Probability in the Engineering and Informational Sciences* 23 305-332.