

# **Analysis of Multiagent Teams**

**Ranjit Nair, Milind Tambe**  
**Computer Science Department**  
**University of Southern California**

# Motivation

- Agent teamwork applied to several realistic domains
  - Framework of beliefs, desires and intentions (BDI)
- How do we analyze the performance of these teams?
  - *Performance critical: linked to loss of human life, etc.*
  - Suggest improvements to the team plan?
  - In particular, improvements to role (re)allocation

## Disaster Rescue simulations



## Battlefield simulations



- Teams operate in uncertain dynamic domains
  - Uncertainty sources: non-determinism, partial observability, multiple agents
  - Hence, use decentralized POMDP model

## Key Contributions:

- *Analysis focused on only communication*
- **Approach:** Role-based Multiagent Team Decision Problem (RMTDP)
  - Techniques for analysis of role allocation and reallocation
- *Analysis using decentralized POMDP model is difficult*
  - Finding optimal Dec-POMDP policy is NEXP-Complete
  - Even evaluating a policy is costly
- **Approach:** Methods for making analysis scalable
  - Decomposition technique based on plan structure
  - Heuristics for improving search time

# Overall Approach

- Applying decentralized POMDPs to the analysis of real world systems

*Team-oriented Programs:  
Team plans, organizations, agents*

The screenshot shows a software interface with three main panels. On the left, a silhouette of a person pointing at a screen is connected to the interface by a double-headed arrow. A blue arrow points from the person to the 'Team Plan' panel. The interface has a menu bar with 'File', 'Edit', and 'Debug'. The 'Team Plan' panel contains a hierarchical list of tasks: 'Evacuate [Big Team]', 'Obtain orders' (with sub-tasks 'Determine number of helos [Obtain orders]' and 'Determine routes [Plan route]'), 'Prepare to execute mission' (with sub-task 'Check route safety [Obtain orders]'), 'Execute mission [Big Team]', 'Fly flight plan' (with sub-tasks 'Travelling overwatch' and 'Fly control routes'), 'Landing zone maneuvers [Flight team]' (with sub-tasks 'Escort operations [Escort]' and 'Transport operations [Transport]'), and 'wait at point [Flight team]' (with sub-task 'Formation landing [Transport]'). The 'Organization' panel shows a tree structure: 'Big Team' (with sub-tasks 'Obtain orders [RPlan]' and 'Obtain safety info'), 'Plan route', 'Flight team' (with sub-task 'Escort'), and 'Transport' (with sub-tasks 'Helo 1 [wolf101]', 'Helo 2 [wolf102]', 'Helo 1', 'Helo 2 [knight11]', and 'Helo 3'). The 'Available Agents' panel lists: 'RPlan', 'teamquickset', 'wolf101', 'wolf102', 'wolf103', 'wolf104', 'wolf105', 'knight10', 'knight11', 'knight12', 'knight13', 'knight14', 'knight15', 'knight16', 'knight17', and 'knight18'. The status bar at the bottom says 'Welcome to Topi...'. A small image strip at the top right shows a 3D simulation of a rescue scene with a helicopter and ground units.

- Other
- Off-line
- As

# Multiagent tiger problem

What did 2 hear?

open left, open right or listen closely?

What did 1 hear?

open left, open right or listen closely?

- Shared reward

- Listen has

- small cost

- unreliable

- No communication

- Reset on open



1



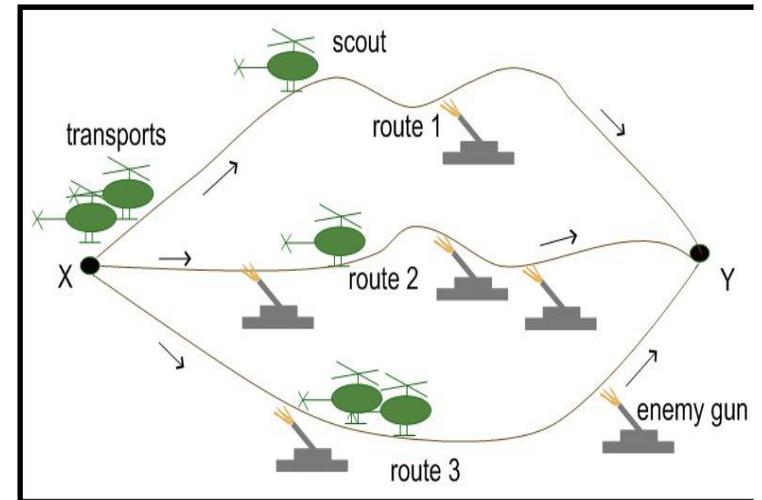
2

What is the best joint policy over horizon T?

# Background: Example Domain

## Task:

- Move cargo from X to Y along any route
- Helicopters need to be assigned to transport or scout role
- Scouts make assigned route safe while transports wait
- Once assigned, transport can become scout but not vice versa



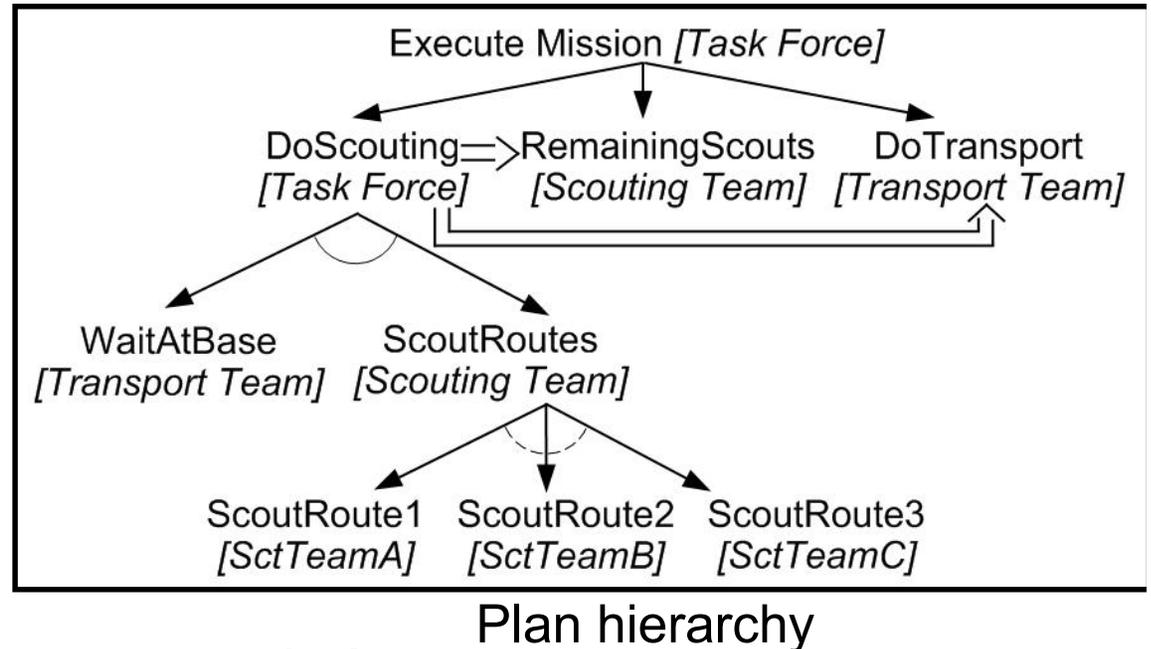
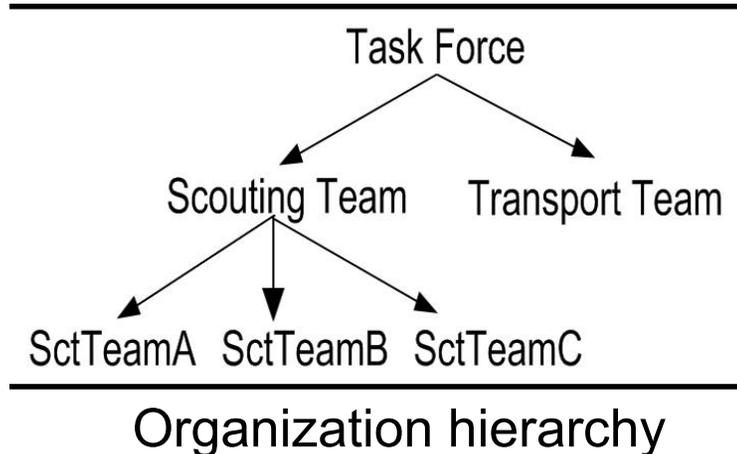
## Uncertainty:

- Each unscouted route: different failure probability and observability.
- Probability of failure depends on number of scouts

## Goal:

- Best role allocation? How many transports? How many scouts on each route?
- How should agents reallocate?

# Background: Team-oriented Program



Given an assignment of subteams to subplans:

- Role Allocation: Best allocation of agents to roles (in organization hierarchy)?
  - E.g. How many helicopters to each subteam?
- Role Reallocation: When and how should agents reallocate?
  - Compare different reallocation strategies (see paper)

# MDPs and POMDPs (background)

**Markov Decision Problem (MDP) =  $\langle S, A, P, R \rangle$**

S: States

A: Actions

Agent's actions have non-deterministic effects

P: Transition function

Obeys Markovian property

$P(s,a,s') = \Pr(s'|a,s)$

R: reward function  $R:S \times A \rightarrow \mathcal{R}$

Goal: Find best action for each state (policy)

**Partially observable Markov Decision (POMDP) =  $\langle S, A, P, O, \Omega, R \rangle$**

Agent has partial knowledge of state  $O(s,a,\omega) = \Pr(\omega|s,a)$

Goal: Find best action for each belief state

## *Role-based Multiagent Team Decision Problem*

- $\langle S, A, P, \Omega, O, R \rangle$ : same as other DEC-POMDP models
- Separate out coordination actions that we wish to analyze i.e. role-taking
  - $A = \times_i A_i$ :  $A_i$  is role-taking  $Y_i$  or role-execution  $\Phi_i$
- R: Reward; sub-divided based on action types
  - Reward for role taking and for execution actions

*Policy  $\pi$* : Action selection of team is specified by joint policy

Joint policy:  $\langle \pi_1, \dots, \pi_n \rangle$

Local policy for agent  $i$ ,  $\pi_i$  :

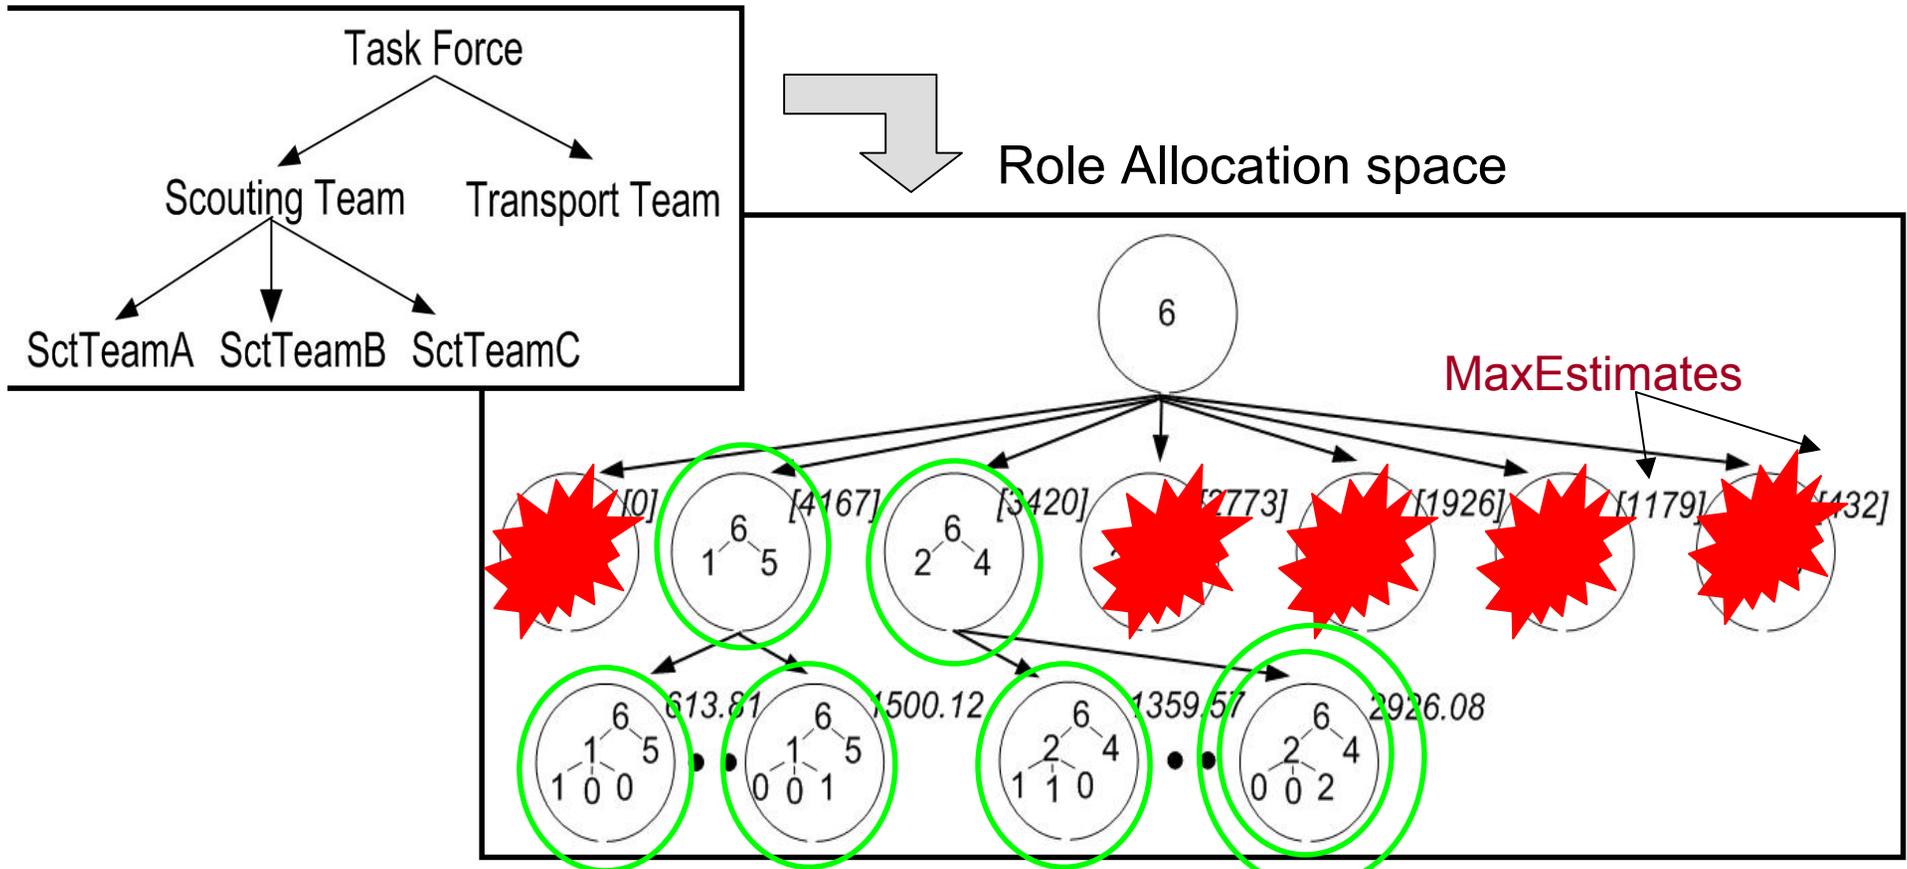
- $\pi_{i\text{Role taking}}$ :  $\omega^1_i, \dots, \omega^t_i \rightarrow$  role-taking action
- $\pi_{i\text{Role execution}}$ :  $\omega^1_i, \dots, \omega^t_i \rightarrow$  role-execution action

# Complexity Issues

- **Theorem1:** The decision problem of determining if there exist role-taking and role-execution policies that yield a reward at least  $K$  over finite horizon  $T$  is **NEXP-complete**. (*Policy Existence Problem*)
- What if we fix the role-execution policy?
- **Theorem2:** Policy Existence Problem for role-taking policy with a fixed role-execution policy is **NEXP-complete**.
- Finding the globally optimal role-taking policy: intractable and likely doubly exponential
  - Brute force search requires  $O\left[\left(|\text{Role-taking}| \frac{|\Omega|^T - 1}{|\Omega| - 1}\right)^n\right]$  evaluations.
- Hence, separate role allocation and reallocation analyses

# Analysis of Role Allocation

- Assumes fixed reallocation policy (e.g. STEAM) and fixed role-execution policy (from TOP)



Finding best initial role allocation

- Branch-and-bound search using MaxEstimates of parent nodes for pruning

# RMTDP Decomposition

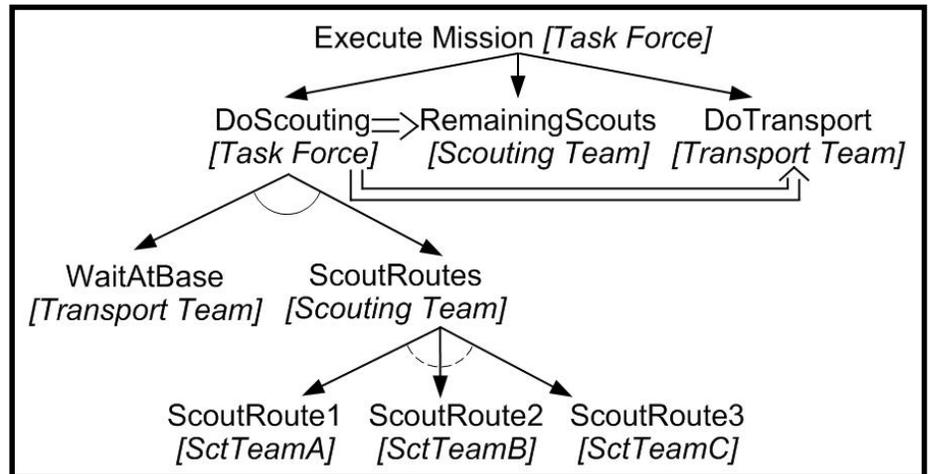
- Identify components in plan hierarchy to decompose RMTDP

- Components with temporal constraints.

- Independent components

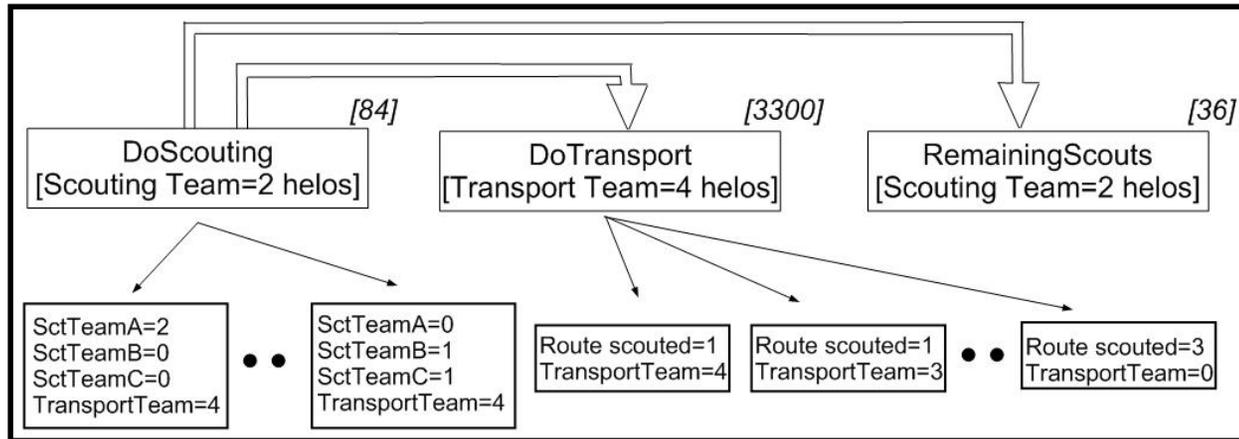
- Obtain smaller RMTDPs for each component

- Provided by domain expert



Decomposition allows fast component-wise computation of max estimates.

# Component-wise Max Estimate



3<sup>rd</sup> parent in allocation space

- Obtain start states and starting observation histories
  - 1<sup>st</sup> component: start states= all possible initial allocations
  - Otherwise: start states = end states of previous component; similarly for observation histories
- Obtain maximum expected utility (MEU) of each component over all start states and observation histories
- $\text{MaxEstimate} = \sum_{\nabla_j} \text{MEU}_j$ 
  - Called *MAXEXP* = 84+3300+36=3420

## Component-wise Max Estimate

### Savings due to decomposition:

- Component-wise evaluation avoids duplication of evaluation
  - Combine end states before determining next component's start states
- Not all variables of a component are relevant to resulting components
  - Remove irrelevant variables from end states
  - Delete resulting duplicate states → fewer start states
- Similarly, with observation histories
  - Irrelevant observations can be removed

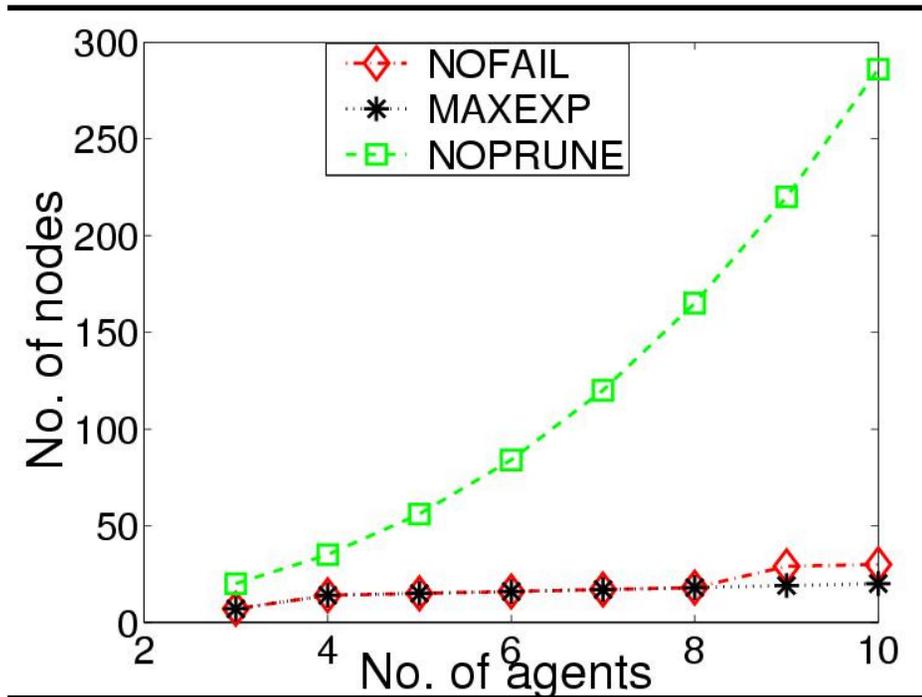
### NOFAIL heuristic:

- Similar to MAXEXP
- Assumes agents don't failure (only for computation of max estimate)
- Results in less branching in evaluation
- Valid for some domains

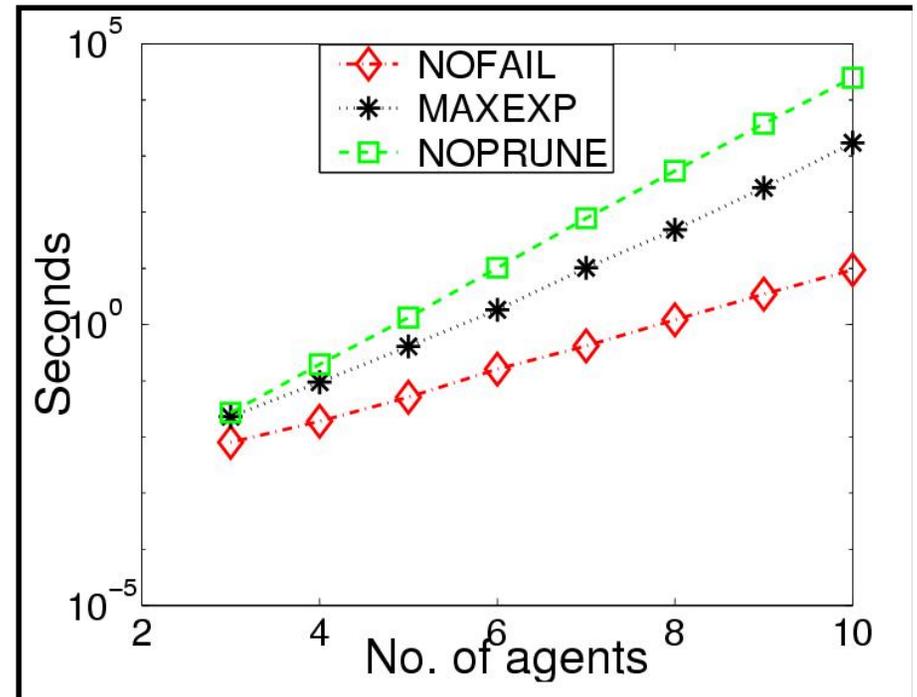
Other heuristics that guarantee correctness possible

# Role Allocation Results

- NOPRUNE: brute force evaluation of all leafs



- 20 fold reduction in number of nodes for 10 agents
- MAXEXP evaluates fewer nodes than NOFAIL



- MAXEXP:14-fold speed up over NOPRUNE for n=10
- NOFAIL: 140-fold speed up over MAXEXP for n=10



# Summary

- BDI-based team plans need analysis tools
  - *Marry BDI and POMDP approaches*
- RMTDP Model for analysis of role (re)allocation
  - Useful for evaluating a TOP
- Finding best initial role allocation
  - Novel decomposition technique
- Comparing role reallocation strategies (in paper)
  - Family of locally optimal perturbations

**Thank You**