

AGENT Teamwork

Milind Tambe

Computer Science Dept & Information Sciences Institute

University of Southern California

4676 Admiralty Way, Marina del Rey, CA 90292

www.isi.edu/teamcore/tambe

tambe@isi.edu

Tel: 310-448-8720; Fax: 310-822-0751

agents@USC

Lecture Plan

PART 1: Theory

- Basics: What is an agent?
- Why teams of agents?
- What is teamwork?
 - *Who needs teamwork theory?*
 - *Teamwork theories*

Break

PART 2: Practice

- Individual agent architectures: reactive planning
- Practical agent teams: Teamwork models

PART 1

agents@USC

Basics: What is an Agent?

PROPERTY

MEANING

- *Situated* Sense and act in dynamic/uncertain environments
- *Flexible* Reactive (responds to changes in the environment)
Pro-active (acting ahead of time)
- *Autonomous* Exercises control over its own actions
- *Goal-oriented* Purposeful
- *Persistent* Continuously running process
- *Social* Interacts with other agents/people
- *Learning* Adaptive
- *Mobile* Able to transport itself
- *Personality* Character, Emotional state

Motivation: Need for Agent Teamwork

Applications:

- *Synthetic agents*: Training, entertainment, education...
- *Software agents*: Information gathering, logistics planning
- *Robotic agents*: Spacecraft constellations, space-station, rovers...
- *Assistants*: Support human collaboration, groupware...

Understanding Teamwork: Motivators

Coming together is a beginning

Keeping together is progress

Working together is success

Teamwork is the ability to work together towards a common vision.

The ability to direct individual accomplishment toward organizational objectives. It is the fuel that allows common people to attain uncommon results

Great motivators: Perhaps too high level?

- *Lets look through some examples ourselves*

Understanding Teamwork

- Ordinary ~~traffic~~
 - Driving in a convoy
 - *Two friends A & B together drive in a convoy*
 - *B is ~~secretly~~ following A*
 - Pass play in Soccer
 - Contracting ~~with~~ a software company
 - Orchestra
- *Not just a union of simultaneous coordinated actions*
 - *Different from contracting*

What is Teamwork?

Cooperative effort by members of a team to achieve a common goal

American Heritage Dictionary

- Common goal, teammate cooperate, help each other, *team spirit*?

Research questions

- Build robust teams with flexible coordination for complex domains
- Team formation
- Teamwork monitoring, diagnosis, analysis, explanation
- Collaborative negotiation
- Adjustable autonomy in teams
- Team plan-recognition
- Learning in teamwork...

How to Build Effective Agent Teams

agents@USC

Given: Multiple agents with *domain-level* skills & team goals

- Appropriate teamwork/coordination to achieve team goals?
 - *When/with whom/what to communicate? Take up others' tasks?*

Difficulties:

- *Members' dynamically obtain differing, incomplete information*
- *Communication time-consuming, costly, risky*
- *Members unexpectedly fail in fulfilling tasks*
- *Uncertainty about what other team members know*

Previous implementations: Often no general teamwork approach:

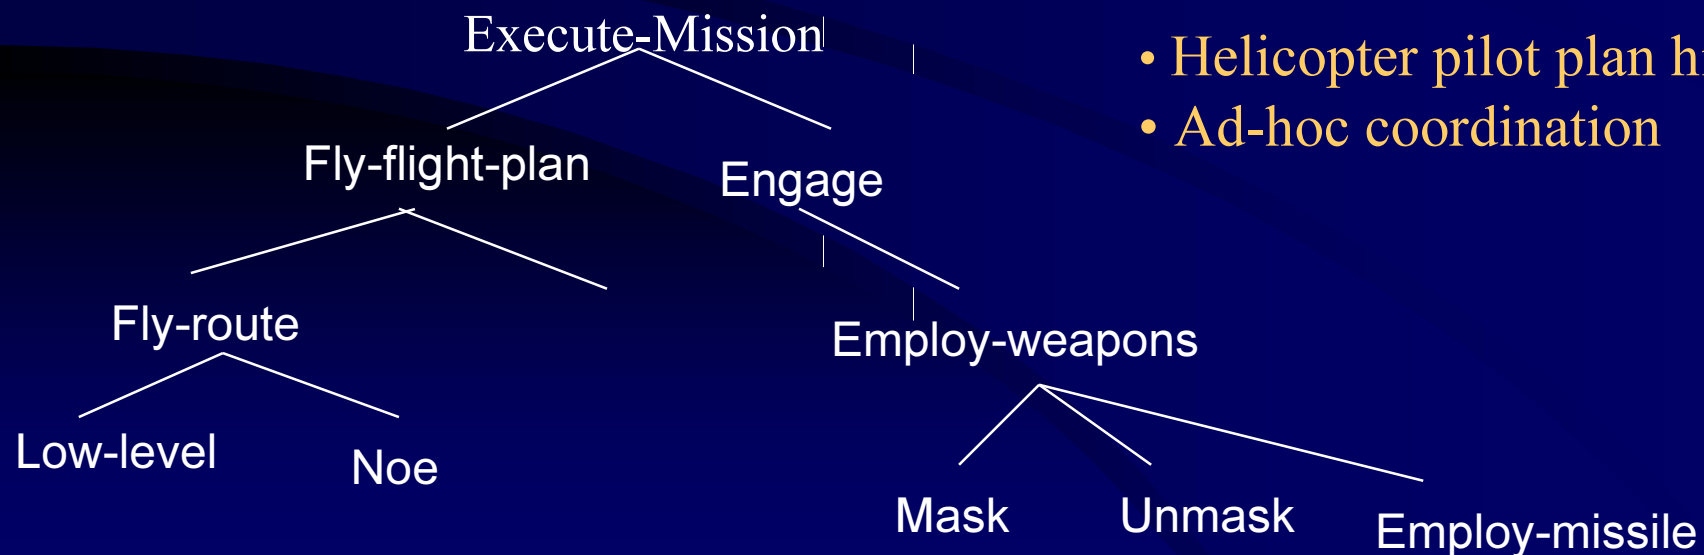


- Individual agents provided individual goals and plans
- Ad-hoc, domain-specific, difficult to write coordination plans
- Results in: Inflexibility (unanticipated failures), lack of reuse

Teamwork Theories

- Theory: Fundamentally understand teamwork
 - *Is there one underlying method for working in teams?*
 - *Is the teamwork in Soccer the same as in for the pilots?*
 - *Are there some fundamentally different types of teamwork?*
- Practical benefits: Specification for teamwork applications
 - *Build reusable teamwork code/libraries*
 - *Teamwork libraries: High-level team-oriented programs*
 - *Build robust applications*

Experience with Ad-Hoc Coordination

A Real-World, Multi-Agent Implementation



- Helicopter pilot plan hierarchy
- Ad-hoc coordination

Experience with Ad-hoc Coordination

A Real-World, Multi-Agent Implementation II

- Initially two-three agents, carefully controlled interactions
- Later, bigger teams, complex scenarios, *unanticipated failures*
 - *Company waited indefinitely, when scout crashed*
 - *Commander returned to home base alone*
 - *Flight leader started mission alone, when others not ready*
 - *Company waited indefinitely when no scout specified*

Problems with fixed, domain-specific coordination plans:

- *No framework to anticipate failures; numerous ad-hoc plans*
- *Coordination code/libraries cannot be reused across domains*
- *Difficult to build effective teams!*

Different Theories of Teamwork

Explicitly based on commitments, beliefs, desires, intentions:

- *Joint Intentions*: Cohen and Levesque, 1990
 - *SharedPlans*: Grosz and Kraus, 1996
 - *Planned team activity*: Sonenberg et al, 1992
 - *We-intentions*: Searle, 1990
- *Basis for effective team building*
- *Don't cover team formation, negotiation, learning,*

Based on decision theory

- *Economic theory of teams*: Marschak & Radner, 1971
- ...

Commitments Fundamental to Agents

- Agents must form and **commit** to plans (*Bratman*)
 - Commit = settle/hold on to a plan, don't drop it easily
- Why commit? Agents are resource bounded
 - Commitments constrain reasoning:
 - *Frame problems for an agent (what to reason about)*
 - Agent needs to determine how to fulfill commitments
 - *Provide filter of admissibility (what not to reason about)*
 - Agent filters out options incompatible with commitments
- Commitments must be balanced: reconsider under key conditions

Team Commitments: Formal Framework

Step by step exploration, including failed attempts!

- *Failures in formalization help our understanding*

Notation (Based on Cohen & Levesque)

- **Bel (x P):** x has P as a belief
- **Goal (x P):** x has P as a goal
- **(Eventually P):** Somewhere in future P becomes true
- **(Until P Q):** Q is TRUE until P TRUE
- **(Always P):** (NOT (Eventually NOT(P)))
- **(Never P):** (Always (NOT(P)))
- **AND, OR, NOT:** Connectives

Commitments to Achievement Goals

- Achievement goals:
 - P is the goal to be achieved, e.g., *inform user about meeting delay*
 - “x” is the agent involved, then:
 - $(Goal\ x\ (Eventually\ P))\ AND\ (Bel\ x\ NOT(P))$
 - Beliefs and goals must be consistent
 - *If $(Goal\ x\ (Eventually\ P))$ then $NOT(Bel\ x\ (Never\ P))$*
- Commitment to P: Don't drop P as a goal easily
 - Agent x will keep P, I.e., $(goal\ x\ P)\ UNTIL$
 - **P Achieved** $(Bel\ x\ P)$ e.g., *informed user*
 - **P Unachievable** $(Bel\ x\ (Never\ P))$ e.g., *impossible to inform*
 - **P Irrelevant** $(Bel\ x\ NOT(Q))$: *Q is possibly the reason for P*

Commitments:

PGOALS or Persistent Goals

- (PGOAL x P Q) is defined as *(Cohen & Levesque)*:

(Bel x NOT(P)) AND (Goal x (Eventually P))

AND

(UNTIL ((Bel x P) OR
((Bel x (Never P)) OR
(Bel x NOT(Q)))

(Goal x (Eventually P)))

- If (Bel x P), would an agent have (PGOAL x P Q)?
- If (Bel x (Never P)) would an agent have (PGOAL x P Q)?

Commitments in Teamwork

Attempt #1: Sum of Individual Commitment to Own Effort

- Team $T = \{x, y\}$ (Two individuals)
- Team searching for a lost object: CSLB
- $P = (\text{Done } T \text{ CSLB})$; $Q = \text{High-level cause for } P$

Team commitment: *(PGOAL x X's-part(P) Q)*

AND

(PGOAL y Y's-part(P) Q)

Remember: Someone is going to build an agent team using this spec!

Problems:

- X and Y don't even know they are a team
- Just individual simultaneous effort

Defining Mutual Beliefs

- $(MB_{x y P})$
 - $(Bel_x P)$
 - $(Bel_y P)$
 - $(Bel_x (Bel_y P))$
 - $(Bel_y (Bel_x P))$
 - $(Bel_x (Bel_y (Bel_x P)))$
 - $(Bel_y (Bel_x (Bel_y P)))$
 -
- Infinite levels of nesting in principle
- Practical systems find compromises
- Remains an active area of research

Formalizing Teamwork

Attempt #2: Add MB to Individual Commitments

Team joint commitment:

$(\text{Joint-PGOAL } T \ P \ Q)$

$(\text{MB } T$

$(\text{PGOAL } x \ X\text{'s-part}(P) \ Q)$

AND

$(\text{PGOAL } y \ Y\text{'s-part}(P) \ Q)$

- MB is Mutual Belief

Problems:

- One team member can drop P and leave, e.g., if $(\text{Bel } x \ P)$
- Glue to hold team together missing

Inspiration from Searle?

agents@USC

- Collective intentional behavior is a primitive phenomenon
- Cannot be analyzed in terms of individual intentions
 - *Puppies playing in a lawn, executing pass play, ...*
 - *No single individual acting alone or accomplish results alone*
- “We” are doing the actions “together”: collective intention
- Searle says this is a primitive, but define in terms of our Cohen/Levesque formulation?
- **Collective Intentions and Actions, John Searle,**
- **Chapter 19, “Intentions in Communication”, MIT Press**

Formalizing Teamwork

Attempt #3: Team as a Single Agent

- Try simple substitution in original formula of PGOAL

(Joint-PGOAL T P Q

(MB T NOT(P)) AND (MG T (Eventually P))

AND

(UNTIL [(MB T P) OR
(MB T (NEVER P)) OR
(MB T (NOT Q))]

(MG T (Eventually P)))

- MG is Mutual Goal: (MB T ((Goal x P) and (Goal y P)))

Problems in Attempt #3: Too Strong

Team is not an individual and hence attempt fails!

- Individuals may diverge in their beliefs
- For instance, x privately believes P, but y believes NOT(P)

Divergence in private beliefs cannot be accomodated

Example:

- Suppose (bel x P), then x drops the achievement (goal x P)
 - *Otherwise, beliefs and goals are inconsistent*
- So (MG T P) false, even though UNTIL condition is false
- So *we were never in Joint-PGoal* to begin with!

Teamwork

Joint Persistent Goal (Cohen & Levesque)

(JPG T P Q

(MB T NOT(P)) AND (MG T (Eventually P))

AND

**(UNTIL [(MB T P) OR
(MB T (NEVER P)) OR
(MB T NOT(Q))]**

(WMG T (Eventually P) Q))

WMG = weak mutual goal, WG = weak goal

(WMG T P Q): (MB T (WG_x T P Q) AND (WG_y T P Q))

- Team mutually believes that each member P as a weak goal

Weak Goals

- UNTIL clause only contains weak goal!

(WG x T P Q) defined as:

[(Bel x NOT(P)) AND (Goal x (Eventually P))]

OR

[(Bel x P) AND (Goal x (Eventually (MB T P)))]

OR

[(Bel x (NEVER P)) AND (Goal x (Eventually (MB T (Never P))))]

OR

[(Bel x NOT(Q)) AND (Goal x (Eventually (MB T NOT(Q))))]

Discussion of JPG

- Joint commitment (JPG) is key to teamwork
 - Team members must hold this mental attitude
 - Binds the team together
- Three parts to JPG definition
 - Team mutually believes the goal P is not achieved
 - Team has the mutual goal to achieve P
 - Until Team mutually believes P *achieved, unachievable or irrelevant*:
 - Team has P as a weak mutual goal, i.e., MB in weak goal

Discussion of Weak Goal

- Commitment towards team in weak goal
 - Team members will attempt to achieve P
 - If x privately believes P is **achieved, unachievable or irrelevant**
 - *Commit to make P mutually believed (e.g., by communication)*
- For instance: Initially, JPG exists, but now (Bel x P)
 - One of the four clauses of WMG must hold
 - **[(Bel x P) AND (Goal x (Eventually (MB T P)))]**
 - X must make P mutually believed (communicate with Y)
- Team members will keep each other informed
 - Team member cannot just walk away, must inform others
 - Predicts communication in teamwork!

Castelfranchi's Counter-example (ICMAS'95)

- Two scientist, one French (F) and one American (A)
- Both searching for the AIDS vaccine
- Mutual beliefs among F and A that searching for AIDS vaccine
- Both have identical goals
- Both will let each other know if AIDS vaccine located
- However, not a JPG, because they compete with each other?

Is the JPG definition somehow incorrect?

Joint Intentions

- Joint commitment basis of *joint intention*:
 - Joint commitment to a joint activity while mutually believing that agents are about to do the joint action
 - Must deliberately do it (jointly)

- *Joint intention is not a primitive concept*
 - Differs from Searle's "we intentions" for instance

SharedPlans (Grosz & Kraus, 96)

No joint mental attitude, instead “intention that” for helpful behavior

SharedPlan of a group GR requires that:

- Mutual belief that each member intends that GR achieves joint goal
- Mutual belief (agreement) in the joint recipe
- For each step in the recipe:
 - Some individual/subteam forms SharedPlan for that step
 - Other members intend that individual/subteam perform the step

SharedPlans may be partial (e.g., recipe not fully elaborated)

- Entire web of intentions and beliefs for teamwork defined

BREAK

agents@USC

PART 2

Practical Teamwork

agents@USC

Situated Plans (Reactive Plans)

- Situated/reactive plan consists of:
 - *Preconditions*, matched with agents' beliefs
 - *Termination conditions*, to terminate plan when matched
 - *Plan body* to execute when plan activated
 - May invoke external or internal or no action
- Example: Plan Attend-Agents-Workshop
 - Precondition: *Saw agents workshop call for participation*
 - Body: *Register for workshop, fly, attend sessions, fly-back..*
 - Termination condition: *Attended agents workshop*

Hierarchical Reactive Plans

agents@USC

Evacuation (Rescue) Domain



Execute-mission

Precondition: Mission defined
Termination: Mission ended
Plan body: (Null)

Fly-flight-plan

Precondition: Flight plan
Termination: Reach destination
See enemy
Plan body: (Null)

Fly-to-waypoint

Precondition: Next waypoint
Termination: Reach waypoint
Plan body: Turn helo, fly helo

BDI Architecture

agents@USC

- PRS, SOAR, IRMA....

INTENTIONS

Attend agents
workshop

Fly to
wkshp

Present
At
Wkshp

PLAN LIBRARY

PLAN: Attend Agents Workshop

Invocation: Agents wkshp declared

Body:

Postconds: Workshop attended

GOAL:
Gain agent
expertise

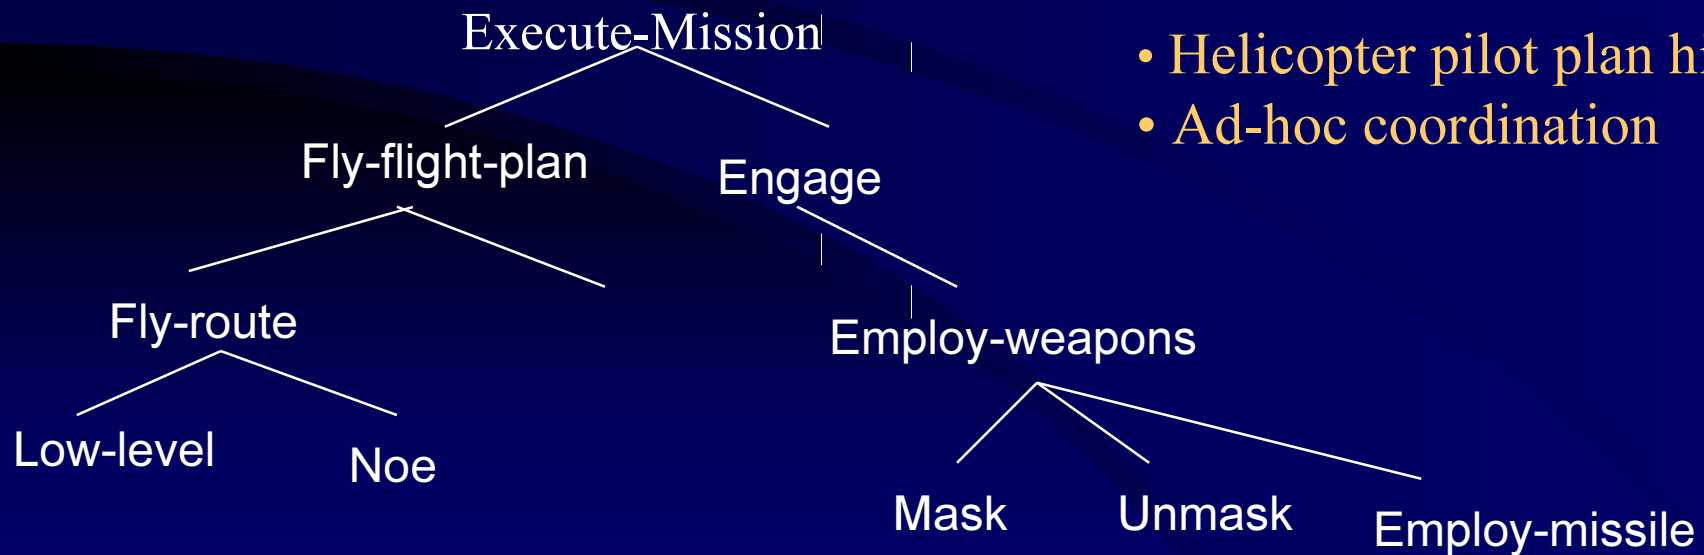

BELIEFS:
Agent wkshp
In Boston

REASONER



Original Pilot Implementation: Individual Plans for Individual Pilots

agents@USC



- Helicopter pilot plan hierarchy
- Ad-hoc coordination

Teamwork: Current Status and Issues

Multi-Agent Implementations

In previous implementations, agents often provided with:

- Individual goals/plans (no explicit team goals/plans)
- Domain-specific, ad-hoc coordination plans

Problems in complex, dynamic domains:

- Inflexibility: difficult to address unanticipated events
- No reusability: within or across domains
- Difficult to build teams: labor-intensive

Teamwork: Novel Approach

Explicit Models of Teamwork

Provide agents with explicit, general model of teamwork:

- Outlines team members' commitments, responsibilities

Advantages:

- Robustness in the presence of unanticipated events
 - *Agents themselves plan responses to teamwork contingencies*
- Reuse/transfer to other applications
 - *Why build coordination from scratch each time?*

Basis in teamwork theories:

- Joint Intentions (Cohen & Levesque), SharedPlans (Grosz & Kraus)
- Few implementations: (Jennings, 95), (Rich & Sidner, 97), (Tambe, 97)

STEAM: A Shell for TEAMWORK

agents@USC

STEAM: General teamwork model for flexible teamwork reasoning

- Algorithm: Agents to reason about team communication & retasking
- *Does not solve ALL of effective team building, e.g., deceptive agents*

Key novelties:

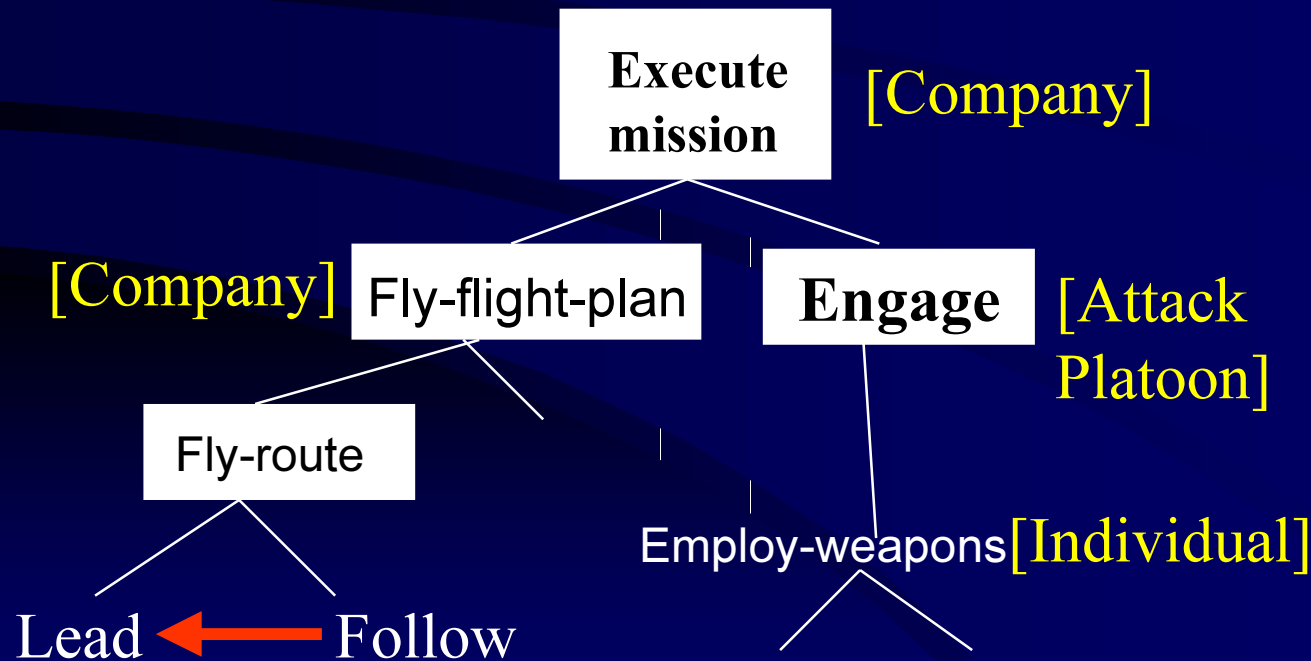
- **Apply & reuse** teamwork model in real-world domains
 - *Empirically illustrates cross-domain general teamwork principles*
- Novel operationalizing of teamwork theories in agent architectures
 - *Modal-logic specifications operationalized in tractable algorithm*
- Theories *starting points*: *STEAM attacks* major unaddressed issues:
 - Communication costs, uncertainty; Team member failure
- Enables team-oriented programs: Agents automatically coordinate
 - STEAM generates appropriate communication & retasking
 - Make it easier to build teams

STEAM Overview

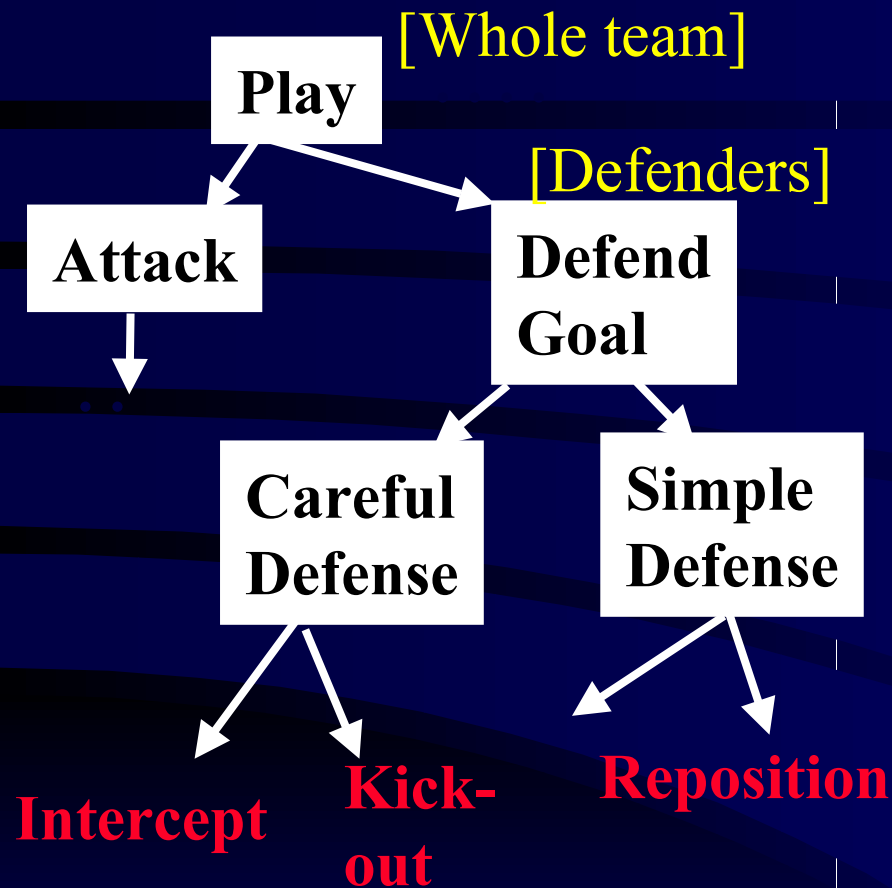
agents@USC

Team-oriented Programs: Explicit team reactive plans

- Hierarchically expand into individual/subteam plans
 - *Each plan has preconditions, body, termination conditions*
- Roles, e.g., *lead role* in formation flying, with constraints
- Assign teams/subteam to team plans based on capability



Team-oriented-Program in RoboCup Soccer



Example (defender subteam):

Simple-defense

- Jointly select plan
- Individuals reposition
- Jointly terminate if ball close

Careful-defense

- Jointly select plan
- Defenders intercept/kick-out...
- Jointly terminate if ball away

TOP Interface

The screenshot displays the TOP Interface with three main panels:

- Reactive Plan Hierarchy:** A tree view showing a mission plan. The root is "Team Plan", which includes "Evacuate [Big Team]". Under "Evacuate", there are tasks like "Obtain orders", "Prepare to execute mission", "Execute mission [Big Team]", and "wait at point [Flight team]".
- Organization Hierarchy:** A tree view showing the organizational structure. The root is "Organization", which includes "Big Team", "Flight team", and "Transport". "Big Team" has sub-tasks like "Obtain orders [RPlan]" and "Plan route". "Flight team" includes "Escort" and "Helo 1 [wolf101]". "Transport" includes "Helo 1", "Helo 2 [knight11]", and "Helo 3".
- Domain-level Agents:** A list of available agents. The list includes "RPlan", "teamquickest", "wolf101", "wolf102", "wolf103", "wolf104", "wolf105", "knight10", "knight11", "knight12", "knight13", "knight14", "knight15", "knight16", "knight17", and "knight18".

Team Members at Run-Time

agents@USC

Each agent maintains a copy of the team-oriented program

- Beliefs: Team beliefs, subteam beliefs, individual beliefs
 - *No Shared memory!*
 - Individual plans modify individual beliefs, but not team beliefs
- *STEAM helps execute team-oriented program*

Helicopter pilot



Helo pilot's Beliefs

Team beliefs

Subteam beliefs

Individual beliefs

STEAM Algorithm Overview

agents@USC

Execute team-oriented program (α : team plan, Θ : team(α), π : parameter)

- *If $EU(C) > EU(NC)$ execute establish-joint-commitment protocol*
- *Install team plan α as a joint commitment with team Θ*
- *While team-state(Θ) does not satisfy AUI conditions of α*
 - *If $F \in$ private state, satisfies AUI conditions of α*
 - *If $EU(C) > EU(NC)$ Communicate (Terminate α , Reason: F)*
 - *If receive message (Terminate α , Reason: F) update team-state (Θ)*
 - *If child plan β_i applicable & team(β_i) = Ω & self $\in \Omega$*
 - *In parallel execute team-oriented-program(β_i , Ω , ...)*
 - *Instantiate role constraints*
 - *If teammate failure, evaluate role constraints*
- *If status(α) unachievable, team-oriented program (REPAIR, Θ , α)*

STEAM Overview

Team Plan Execution: Communication

All team plans executed by forming & terminating joint commitments:

- **Request-confirm** exchanges so all team members select & commit
- Establish mutual belief for achieved, unachievable,... to terminate
- *Forming & terminat team plans: All communication in STEAM*

Example: Team of helicopters jointly commit to “execute mission”

- If commander privately believes *mission unachievable*...
- Commander must establish mutual belief in termination condition
- It communicates *mission unachievable*: no one left behind

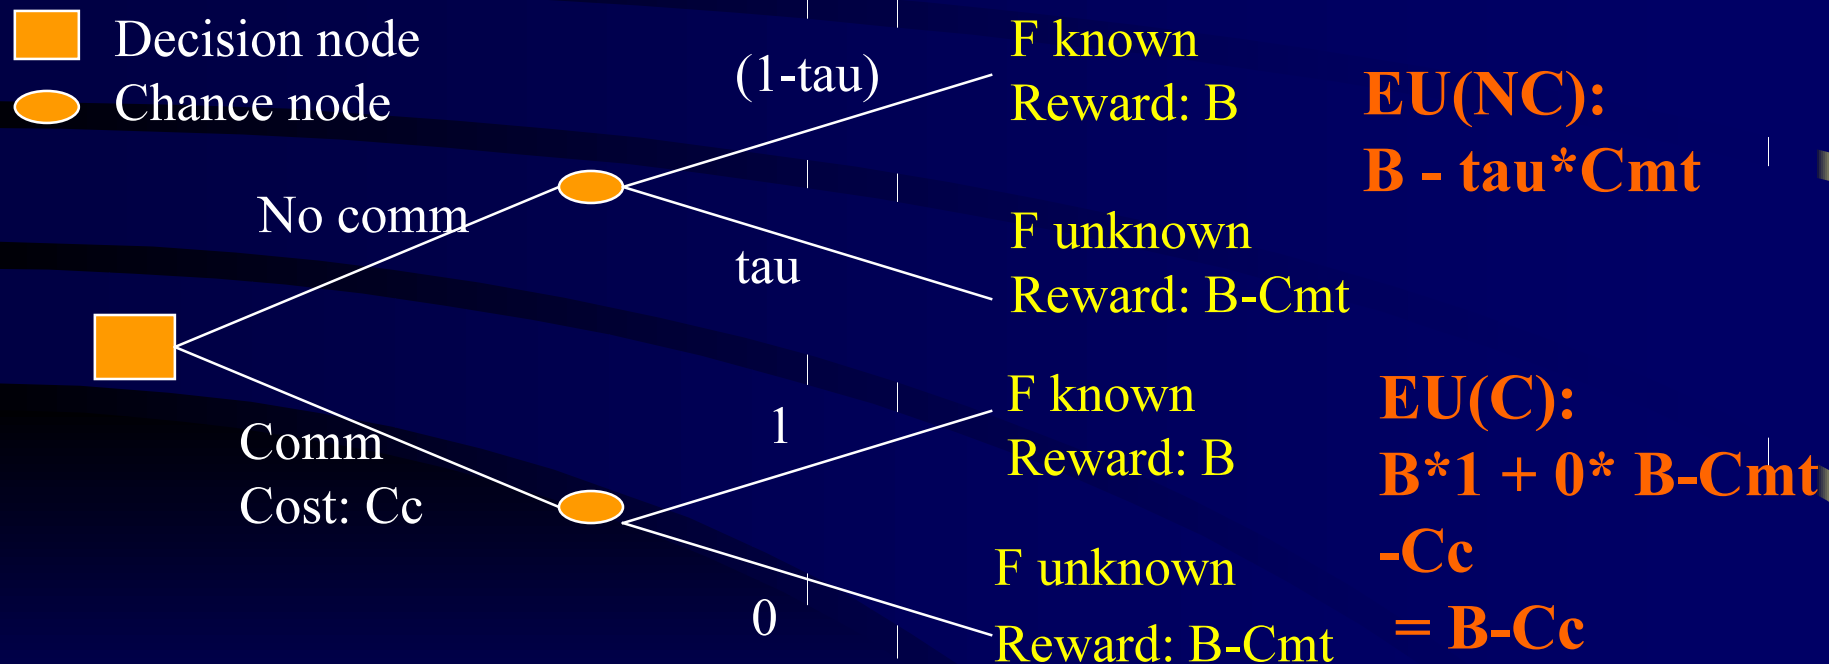
Hierarchy of jointly committed team plans and subteam plans:

- Team coherent when executing & terminating team plans

STEAM: Communication Selectivity

Addresses communication costs & risks:

- Incorporates *decision-theoretic communication selectivity (MEU)*
- Tradeoff communication vs team incoherence costs



Communicate if $EU(C) > EU(NC)$: $\tau * C_{mt} > C_c$

- More complex reasoning being explored: MDPs

Implementing Models of Teamwork

STEAM Overview: Monitor and Repair

Addresses unanticipated team member or subteam failure:

- Monitoring & replanning capabilities
 - Explicit constraints individual/subteam roles & team goal
 - *AND: All roles must be fulfilled*
 - *OR: At least one rule must be fulfilled*
 - *Role-dependency \rightarrow : Role R1 dependent on R2*
 - Constraints may be combined, e.g., $((A \text{ OR } B) \text{ AND } (B \rightarrow C))$
- Scouting failure example:
 - *Wait-for-battle-position-scouted is the team plan*
 - *AND-combination: Scout and Non-scout roles in team plan*
 - If scout crashes, the scout role is not fulfilled
 - AND-combination implies that the team plan fails

STEAM Overview (Continued)

- Joint commitment to replan by reorganization, if critical failures
 - Determine candidates for roles via capability matching
 - Candidates for roles ensure no conflicting critical commitments
 - Individual/subteam may volunteer
 - *If multiple candidates, compare based on capability*
 - Highest capability agent wins
- Scouting failure example continued
 - *Wait-for-battle-position-scouted is the failed team plan*
 - Locate other pilots capable of scouting
 - New candidate scout ensure no conflicting commitments
 - Candidate scout(s) volunteer
 - Best capability scout wins

Operationalization of STEAM

agents@USC

Currently, in the *Soar integrated architecture*

- Set of about 280 rules, available on WWW
- www.isi.edu/teamcore/tambe/agent.html
- Ideally, enhancements in agent architecture (*currently conventions*)
 - Team operators (reactive team plans); team state (mutual beliefs)

Applications:

- Pilot teams for synthetic **Attack** helicopters [1576 rules]
- Player teams for **RoboCup** Soccer agents [~ 550 rules]
- Pilot teams for synthetic **transport** helicopters [1333 rules]
- Heterogeneous teams in TEAMCORE [~1000 rules]
- *Emerging applications (Univ of Portsmouth, UK, by A. Kalus)*

STEAM Rules Example I

Create-communicative-goal-on-achieved

If

Agent A1's private state contains a fact F

AND

Fact F matches an achievement condition AC of a team operator
OP

AND

Fact F is not currently mutually believed

AND

A communicative goal for F is not already generated

THEN

Create possible communicative goal CG to communicate fact F to
team to terminate team operator OP

STEAM Evaluation: Applications

- Helicopter teams in large-scale synthetic exercises (STOW'97)
 - Human pilots set tests, and favorably evaluated
 - *Three companies of Army helicopters, flights of Marine helos*
- ISIS team participated in RoboCup'97 & 98 simulation league
 - Won 3rd place at RoboCup'97, fourth place at RoboCup'98

STEAM reuse in (% of rules) in complex domains:

Attack	RoboCup	Transport	Teamcore	T. Kalus
First use	45%	100%	100%+	50% ?

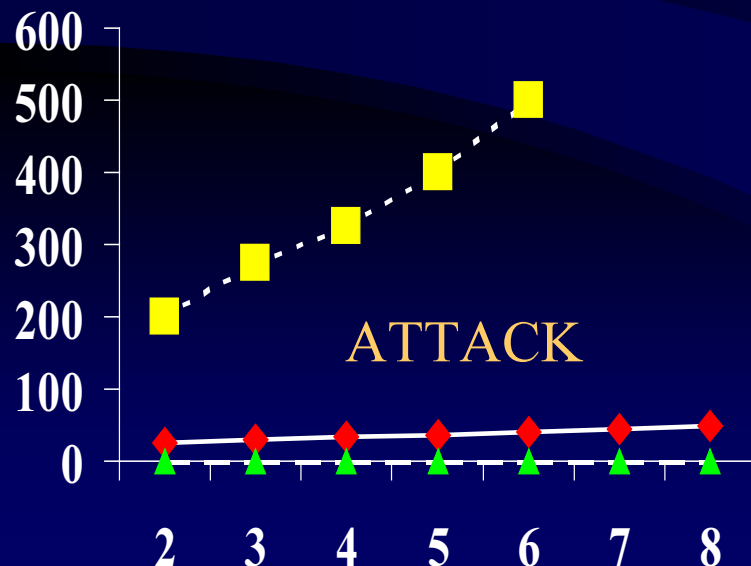
- ➔ *Estimate reuse saves writing 100s of coordination plans*
- ➔ *Reuse saves effort, but any benefits?*
 - ➔ *STEAM vs NO STEAM, over 100s of simulation runs*

STEAM: Evaluation II

	STEAM	No STEAM
ISIS97-CMUUnited97	3.27	1.73
ISIS97-Andhill97	-3.38	-4.36
Transport/Attack	Succeed	Fail
ISIS98-CMUUnited97	4.04	3.91
ISIS98-Andhill97	-1.53	-2.13

- Not compare STEAM with newly written coordination or else must pay rewrite cost
- Goal avg diff for ISIS97 significant but not in ISIS98
- More speedup in tactics in ISIS97: mutual dependence

Communication selectivity?



Teamcore: Heterogeneous Teams

Rapidly bring together legacy agents, specialized agents & humans?

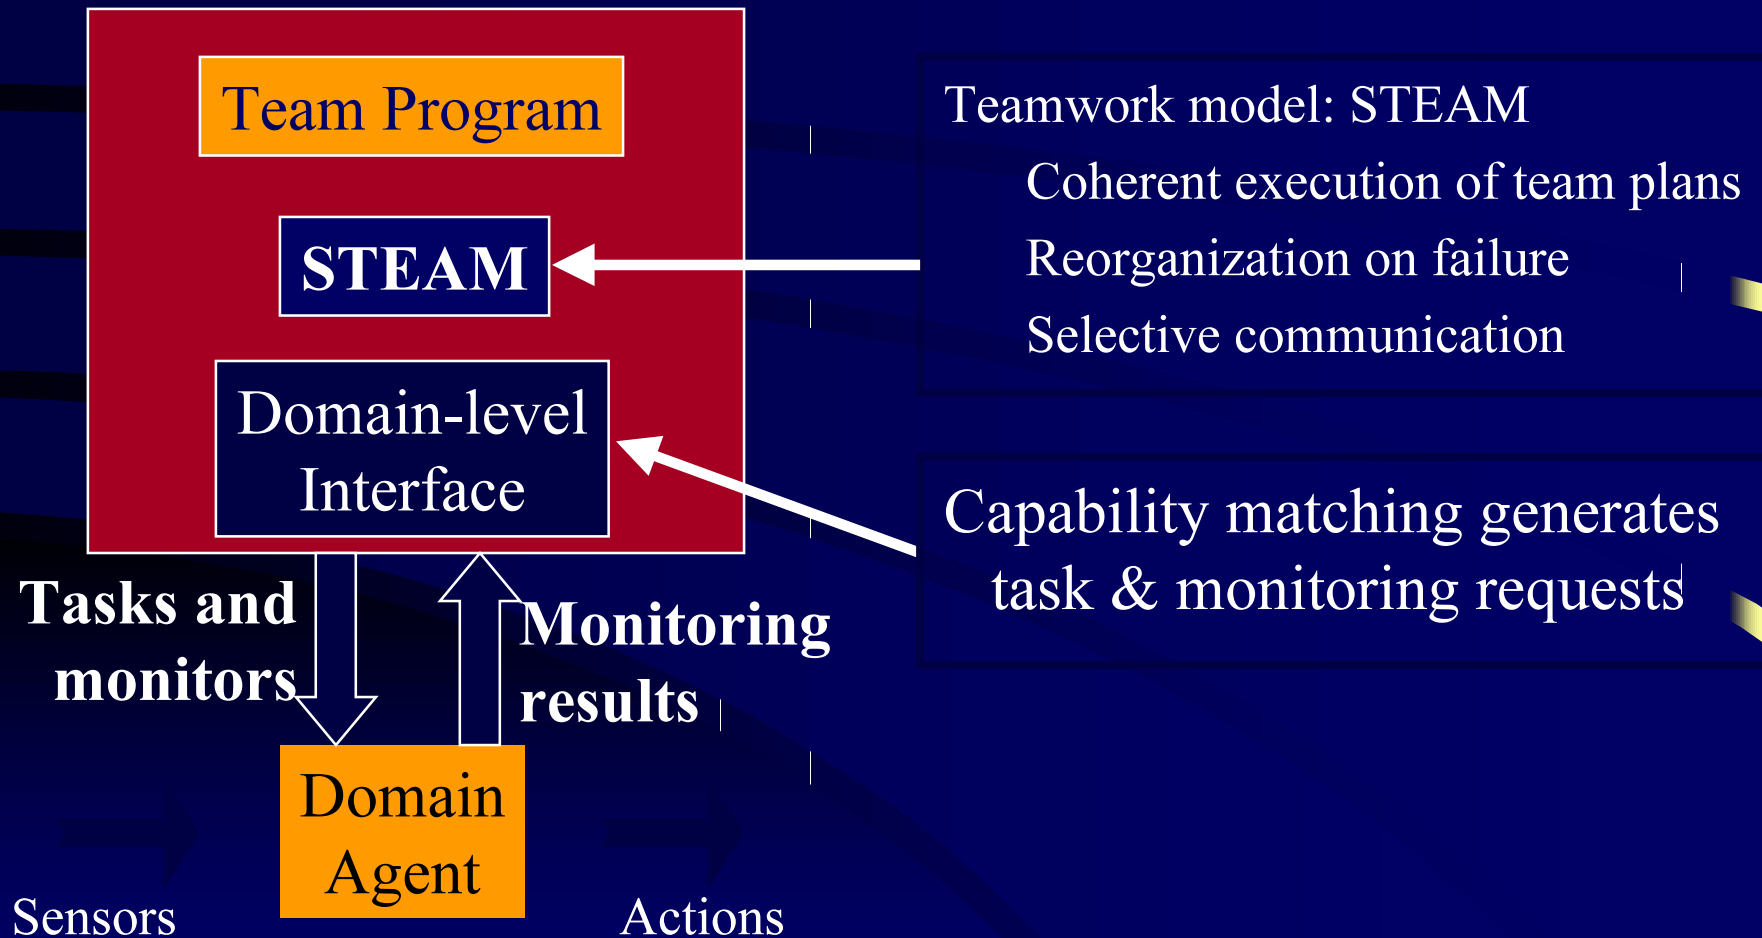
- STEAM had to be built-in in agents
- Teams tend to be homogeneous

Teamcore: Distributed integration architecture, built-in teamwork

- Teamwork-based integration:
- Each agent/human TEAMCORE proxy (with built-in STEAM)
 - *Robustness: agents cover for each others' failure*
 - *STEAM reuse alleviates need to write coordination from scratch*

Previous: Centralized integration architectures, w/o coordination (OAA)

TEAMCORE Proxy: Zooming in on a Single Proxy



TEAMCORE Application: Electric Elves

agents@USC

Team of personal assistants to support a human organization:

- Agent proxies for 9 researchers (called “Friday”)
 - Interfaces: PDA/GPS, WAP phones, workstation, fax, speech
- Agent proxy for a project assistant
- Information agents, schedulers, matchers...


Agent proxies run 24/7

- **Help us with real tasks**
 - Coordinate meetings (reschedule if delays, cancel)
 - Decide presenters at research meetings (via auctions)
 - Track people (www.isi.edu/teamcore/info.html)
 - Order our meals


Electric Elves: Teamcore Proxies



Teamcore proxy



Teamcore proxy



Teamcore proxy

- Reschedule meetings
- Decide presenters
- Track people
- Order our meals

Meet Maker

Papers

Teamcore proxy
Interest Matcher
Research Publication agent



Teamcore proxy
Schedulr agent

Teamcore proxy
Capability Matcher

Team Program
STEAM
Domain-level Interface

Teamcore: Robust teamwork with others, e.g., replace critical role
• Adjustable autonomy raised as a major concern


Elves in Use: Asking Users

RoboCupRescue27: Delay Meeting	Order Dinner
 <p>RoboCupRescue27: Namaste Milind Would you like me to do any of the following for you?</p> <p>Delay meeting by 5 minutes. Delay meeting by 15 minutes. Delay meeting by 30 minutes. Delay meeting by 1 hour. Delay meeting by 2 hours. Delay meeting by 1 day. Say that you are now attending. Say that you will attend later but start the meet Say that you will not attend. Cancel the meeting.</p> <p>OK Show Details</p>	 <p>What would you like to eat today?</p> <p>Do not order now Choose something new for me Choose an order from my usual California Pizza Kitchen JERRY Famous Deli Subway Zebra Cafe</p> <p># of Meals: 1</p> <p>If ordering: Take Out Delivery Dine In</p> <p>OK Show Details</p>

“ More & More computers are ordering food, ...we need to think about marketing” local Subway owner

Elves in Use: Auctions

TEAMCORE20		presenter	
team-team			
Agent	capability	willingness	Overall
Paul Scerri	1.0	1.0	1.0
David Pynadath	1.0	0.0	0.3
Milind Tambe	1.0	0.0	0.3
Jay Modi	1.0	0.0	0.3
Shriniwas Kulkarni			0.0
Hyuckchul Jung	0.0	0.0	0.0
Lei Ding		0.0	0.0
Takayuki Ito		0.0	0.0
Ranjit Nair		0.0	0.0
other-friday			0.0



Jay Modi

ASSian

Adjustable Autonomy in Teams

Proxies for users: *Teamwork with others, while serving human users*


Adjustable autonomy: “*Dynamically adjust agent’s autonomy*”

- Autonomous action on behalf of humans reduces burden, but...
 - *Proxies face significant uncertainty, e.g., how hungry?*
 - *Errors in autonomous actions may be costly*
- Reduce autonomy, transfer control to humans in critical situations

Teams raise novel challenges for adjust autonomy!

- Previous work: Individual agent/user interactions
- With teams, an agent must serve the user AND the team
 - *E.g., Cannot wait for user input: causes team miscoordination*
- Approach based on **Markov Decision Processes (Agents’2001)**

Electric Elves: 24/7 Run in a Real Organizational Environment



Teamwork Research: Monitoring

agents@USC

Monitor teammates' critical, but how w/o communication (w Kaminka)?

Compare behavior with teammates exploiting plan-recognition:

- *Infer a teammate's team plan hierarchy given its actions*
 - If ambiguity, resolve via team coherence heuristic
 - If A1 in Engage, A2 in {Hide, Engage}, prefer A2: Engage
 - *Start from top, compare own plan hierarchy with inferred hierarchy*
 - *Failure if detect a difference in team plans*
- E.g., if {A1: Engage, A2: Engage, A3: Hide} then declare failure

Experiments: From 14 helo test cases, 12 failures detected (AAAI'98)

Theory:

- *Theorem: Central monitoring: No sound & complete failure detection*
- *Theorem: Distributed monitoring: Sound & complete failure detection*
 - Monitor *key* agents, use only team coherence heuristic

Summary

Teamwork

Teamwork increasingly critical in a vast range of applications

- Training, education, entertainment, manufacturing, design...

Fundamental understanding of teamwork important

- Teamwork theory
- Not all aspects of teamwork covered

Practical teamwork models based on theory

- Reuse teamwork capabilities, flexibility in teamwork
- Team-oriented programming

Lecture Readings

- Barbara Grosz and Sarit Kraus. 1996. "Collaborative Plans for Complex Group Action." In Artificial Intelligence. 86(2), pp. 269-357.
- Cohen, P. R. & Levesque, H. J. (1991). Teamwork. Nous 25(4), pp. 487-512. Special Issue on Cognitive Science and Artificial Intelligence.
- Smith, I. A. & Cohen, P. R. (1996). Toward a Semantics for an Agent Communications Language Based on Speech-Acts. Proceedings of AAI-96.
- Levesque, H. J., Cohen, P. R. & Nunes, J. T. H. (1990) On acting together. Proceedings of AAI-90. Boston, Massachusetts. (May not be available on-line).
- Tambe, M. "Towards Flexible Teamwork", Journal of Artificial Intelligence Research, Volume 7, Pages 83-124
- Jennings, N.R., Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems using Joint Intentions, Artificial Intelligence, 75 (2), 1995, 195-240.
- Rich, C. and Sidner, C.L. COLLAGEN: When Agents Collaborate with People, In First Int. Conf. on Autonomous Agents, Marina del Rey, CA, Feb. 1997
- Decker, K. and Lesser, V., "Designing a Family of Coordination Algorithms," In Proceedings of the First International Conference on Multi-Agent Systems.

Thank YOU

agents@USC

STEAM Evaluation in RoboCup

Reuse of general teamwork model

- 35-45% of STEAM rules reused (origin in battlefield simulations)
- Alleviate effort for encoding domain-specific coordination plans

Performance impact of STEAM (of the part reused)

- ISIS97 vs. CMUnited97
 - **1.5 goals (out of 3.2)** per game on average
- ISIS97 vs Andhill97
 - **1 goal (out of 3.4)** per game on average

Lesson learned: General teamwork models can be useful

- *Reduce development time & improve performance*

Establishing Joint Commitments (Smith and Cohen, 96)

agents@USC

Explicit exchange of “Request” then “Confirm” for JPG

- PWAG (PWG): Persistent weak achievement goal
 - *Similar to WG, but involves more commitment*
 - *“Half of a JPG”*

$[(\text{Bel } x \text{ NOT}(P)) \text{ AND } (\text{PGoal } x \text{ (Eventually } P))]$

OR

$[(\text{Bel } x P) \text{ AND } (\text{PGoal } x \text{ (Eventually (MB T } P)))]$

OR

$[(\text{Bel } x \text{ (NEVER } P)) \text{ AND } (\text{PGoal } x \text{ (Eventually (MB T (Never } P))))]$

OR

$[(\text{Bel } x \text{ NOT}(Q)) \text{ AND } (\text{PGoal } x \text{ (Eventually (MB T NOT}(Q))))]$

Establishing Joint Commitments (Smith and Cohen, 96)

agents@USC

How the “protocol” works:

- Sender requests others, implies it has a PWAG
- Others may confirm, which implies they have a PWAG
- Request followed by confirm establishes JPG
- If at anytime *during such an exchange* the goal is achieved, unachievable or irrelevant, then must inform others
 - *Don't request and then leave others hanging*

Example of Cohen/Levesque's JPG

Example

- *Assume (JPG T P Q) exists*
 - Suppose x believes P, i.e., (Bel x P)
 - No longer a JPG, since (MB T NOT(P)) fails
- *However, for JPG to exist, WMG must have been true*
 - One of the four clauses in (WG x T P Q) must be true
 - [(Bel x P) AND (Goal x (Eventually (MB T P)))]
 - X must make P mutually believed (communicate with Y)

Influence of Teamwork Theories on STEAM

- Joint commitments as basic building blocks of teamwork
 - Commitments guide coordination, communication
- Specify entire hierarchy of intentions as in SharedPlans:
 - Joint intention for high-level team goal & joint actions in its service
- Knowledge of other team members (SharedPlans)
 - Domain-specific coordination relationships explicit to monitor
 - Role-performance ability of individual/subteam monitored
- Joint intention to replan upon unachievability (Partial SharedPlans)
 - Individuals may volunteer or requested to take up tasks