Negotiated Decommitment in a Collaborative Agent Environment

Overview

Introduction

Theoretical Framework

Implementation Architecture

Experimental Design

Results and Analysis

Conclusions

Introduction

Significance of the Research Problem

Research Issues

Hypotheses

Significance of the Research Problem

- BDI Model, Rao and Georgeff (1991, 1995)
- Intention Reconsideration, Schut and Woolridge (2001)
- Formalization of Commitment, Norman, Sierra and Jennings, (1998)
- Conditional Commitment, Andersson and Sandholm (1998)

Significance of the Research Problem (cont.)

- Degrees of Commitment, Excelente-Toledo, Bourne and Jennings (2001)
- Decommitment in Self-Interested Societies, Xing and Singh (2001), Sandholm and Lesser (1996)
- Cancellation, Sen and Durfee (1996)

Significance of the Research Problem (cont.)

- Decommitment addressed:
  - Formalization of individual agent architecture
  - Self-interested agent societies
- Decommitment not addressed:
  - Cooperative agent societies
  - Negotiated decommitment

Research Issues

- Intuitive Definition
  - Commitment as intention
  - Rational or accidental decommitment
- Why Decommit
  - Higher priority of a competing potential commitment
  - Previous commitment no longer productive

Research Issues (cont.)

- Repercussions
  - Impact of decommitment
  - Commitment *value*: Estimate of utility
  - Commitment *strength*: Impact on the system

Hypotheses

- Decommitment will improve overall goal achievement of the system
- Negotiated decommitment will be more beneficial than unilateral decommitment
- Overall goal achievement will degrade gracefully as system constraints increase

Theoretical Framework

Distributed Task Scheduling

Individual Agents

Agent Society and Interaction

Negotiation

Commitment and Decommitment

Distributed Task Scheduling

- Sen and Durfee (1996)
- *S* = (A,T)
  - A = $\{a_1, a_2, \ldots, a_k\}$, the set of agents with control of resources, and
  - T= $\{ \tau_1, \tau_2, \ldots, \tau_n\}$, the set of tasks which may be scheduled.

Distributed Task Scheduling (cont.)

- $\tau_i = (A_i, h_i, l_i, w_i, S_i, a_i, d_i, T_i)$
  - $A_i \subseteq A$, set of agents controlling resources;
  - $h_i \in A_i$, the agent requesting performance of a task;
  - l_i is the requested duration of the task;
  - w_i is the priority assigned to the task;
  - S_i is the set of possible starting times for the task;

Distributed Task Scheduling (cont.)

- $\tau_i = (A_i, h_i, l_i, w_i, S_i, a_i, d_i, T_i)$
  - a_i is the timestamp at which h_i requested the task be performed;
  - d_i is the deadline by which time the task must be scheduled;
  - T_i is the time at which the task is actually scheduled.

Distributed Task Scheduling: Commitment Value

- $w_i = (p_i, v_{hi}, c_i, w_{hi}, dt_i)$
  - p_i is the default priority of that type of task;
  - v_{hi} is A_i's assessment of the validity of h_i's information;
  - c_i is the constrainedness of the task, comprised of the number of other agents also asked to perform the task and the duration (l_i) of the task;

Distributed Task Scheduling: Commitment Value (cont.)

- $w_i = (p_i, v_{hi}, c_i, w_{hi}, dt_i)$
  - w_{hi} is h_i's assessment of the value of the task;
  - dt_i is the difference between the time the request was made and the requested start time, or $(S_i - a_i)$

Individual Agents

- Characteristics:
  - Collaborative and benevolent
  - Rational
  - Autonomous
  - Communicative, Capable of Negotiation
  - Multitasking
  - Capable of Time Dependent Planning
  - Capable of Learning

Agent Society and Interaction

- Soh and Tsatsoulis (2001)
- Ω – a multi-agent system
- Ψ – a “neighborhood” in the system
- $\lambda(\alpha,\beta)$ – predicate indicating agent α knows about agent β

Agent Society and Interaction (cont.)

$$\Psi \subseteq \Omega,\ \Psi \neq \varnothing$$

$$\lambda\left(\alpha_i, \alpha_j\right) \forall i \forall j \alpha_i, \alpha_j \in \Psi$$

$$\Omega = \left\{\Psi_1, \Psi_2, \ldots, \Psi_N\right\}$$

Negotiation

- Restricted to neighbors
- Request to perform task or request to decommit
- Local estimate of global utility of commitment used to determine agreement
- Information stored on interactions
- Time bounded

Commitment and Decommitment

- Commitment value = w_i, the priority, or weight, of the potential commitment
- Commitment strength = str_i, the estimated effect of dropping a commitment

Commitment and Decommitment

- $str_i = (n_i, r_{hi}, dnow_i)$
  - n_i is the number of agents potentially affected by the decommitment;
  - r_{hi} is the perceived reliability of the neighbor to whom the commitment was made, that is, the number of times that neighbor honored commitments to A_i;
  - $dnow_i$ is the difference between the scheduled start time of the task and the current time.

Implementation Architecture

Problem Domain

Agent Architecture

Agent Interaction

Local Estimate of Global Utility

Decision Criteria

Problem Domain

- Autonomous Negotiating Teams (ANTS)
  - DARPA funded research effort led by Drs. Tsatsoulis, Niehaus and James of ITTC
- Multi-sensor target tracking
- Radar simulator (Radsim)
- Agents and external software

Radsim

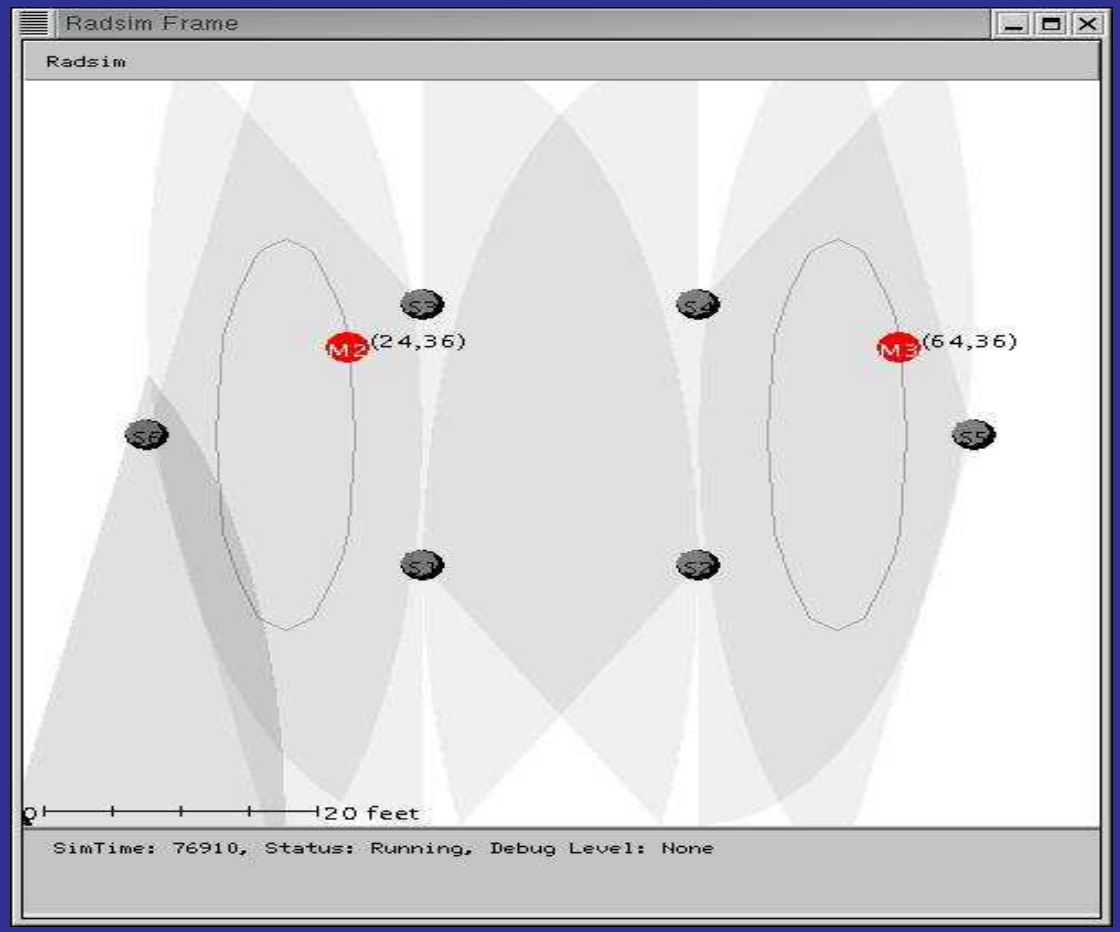
Radsim Frame
Radsim
(24,36)
(64,36)
20 feet
SimTime: 76910, Status: Running, Debug Level: None

Agents and External Software

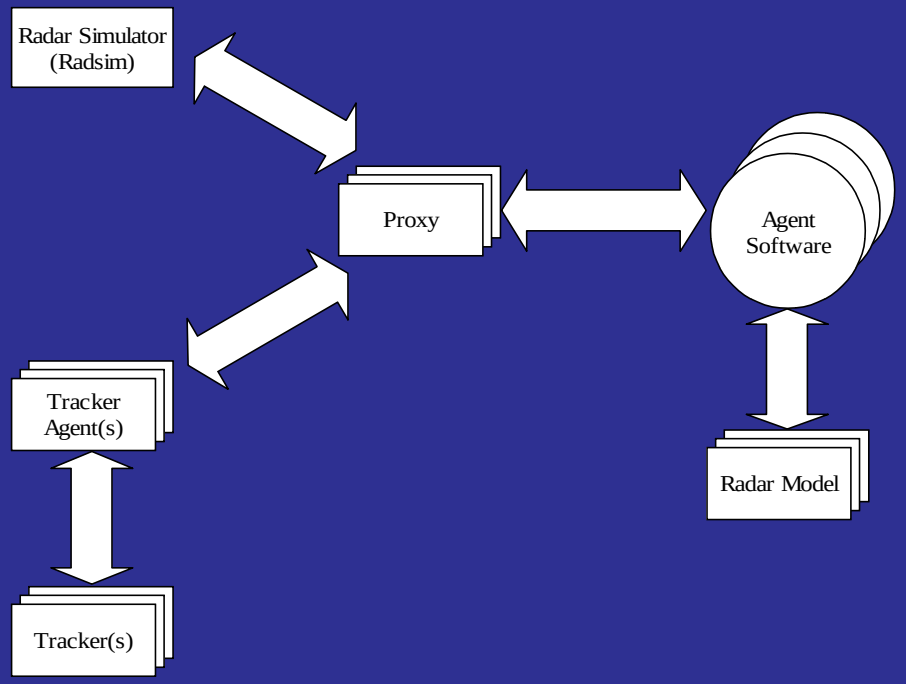

Agent Architecture

- Multithreading
- Communicator
- Scheduler
- Negotiator
- Agent Thread (Decision Maker)

Multithreading

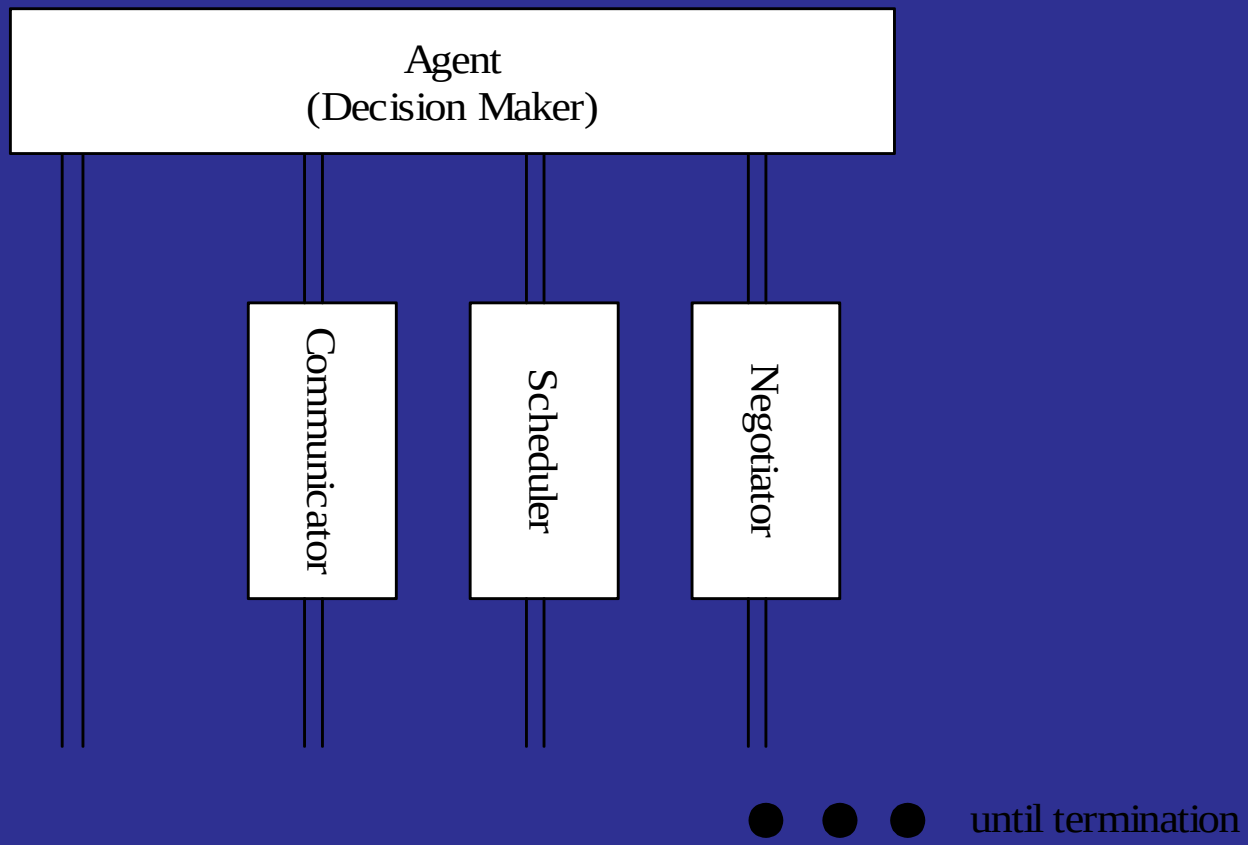

Communicator

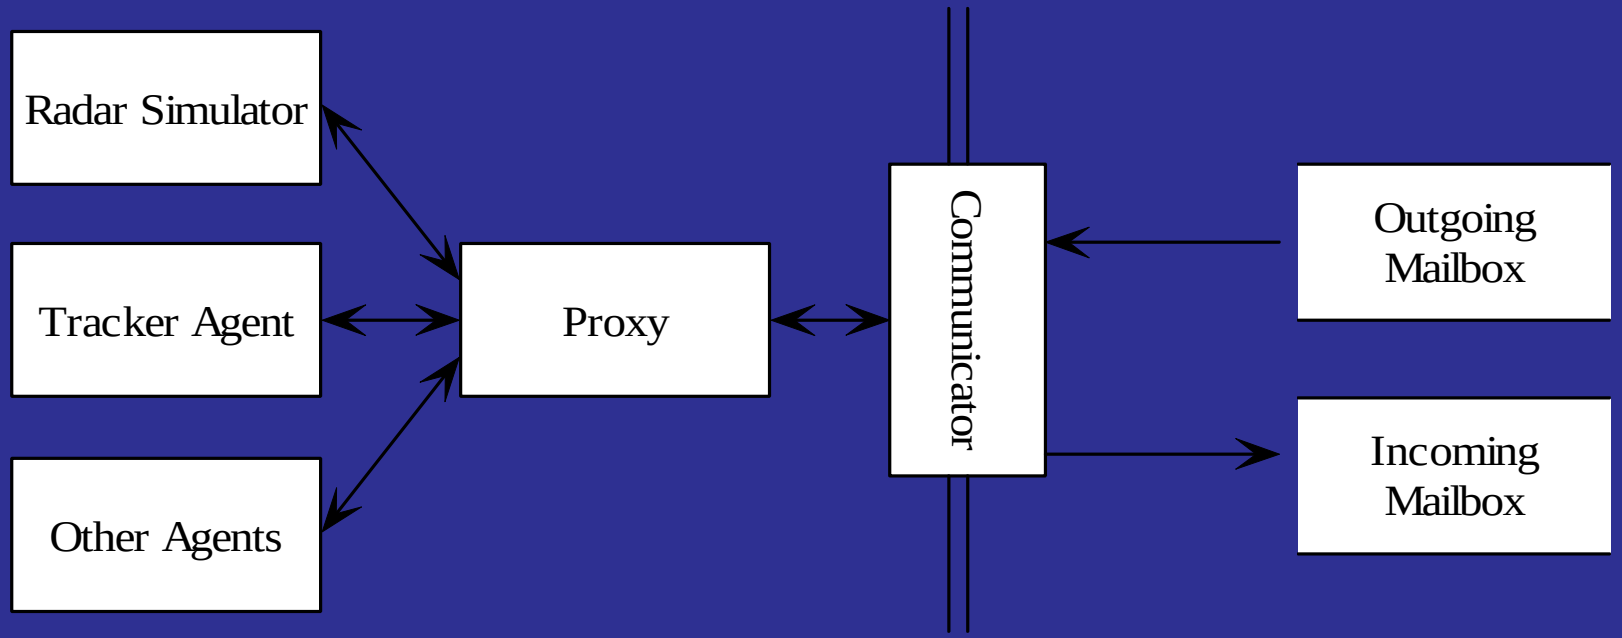

Scheduler

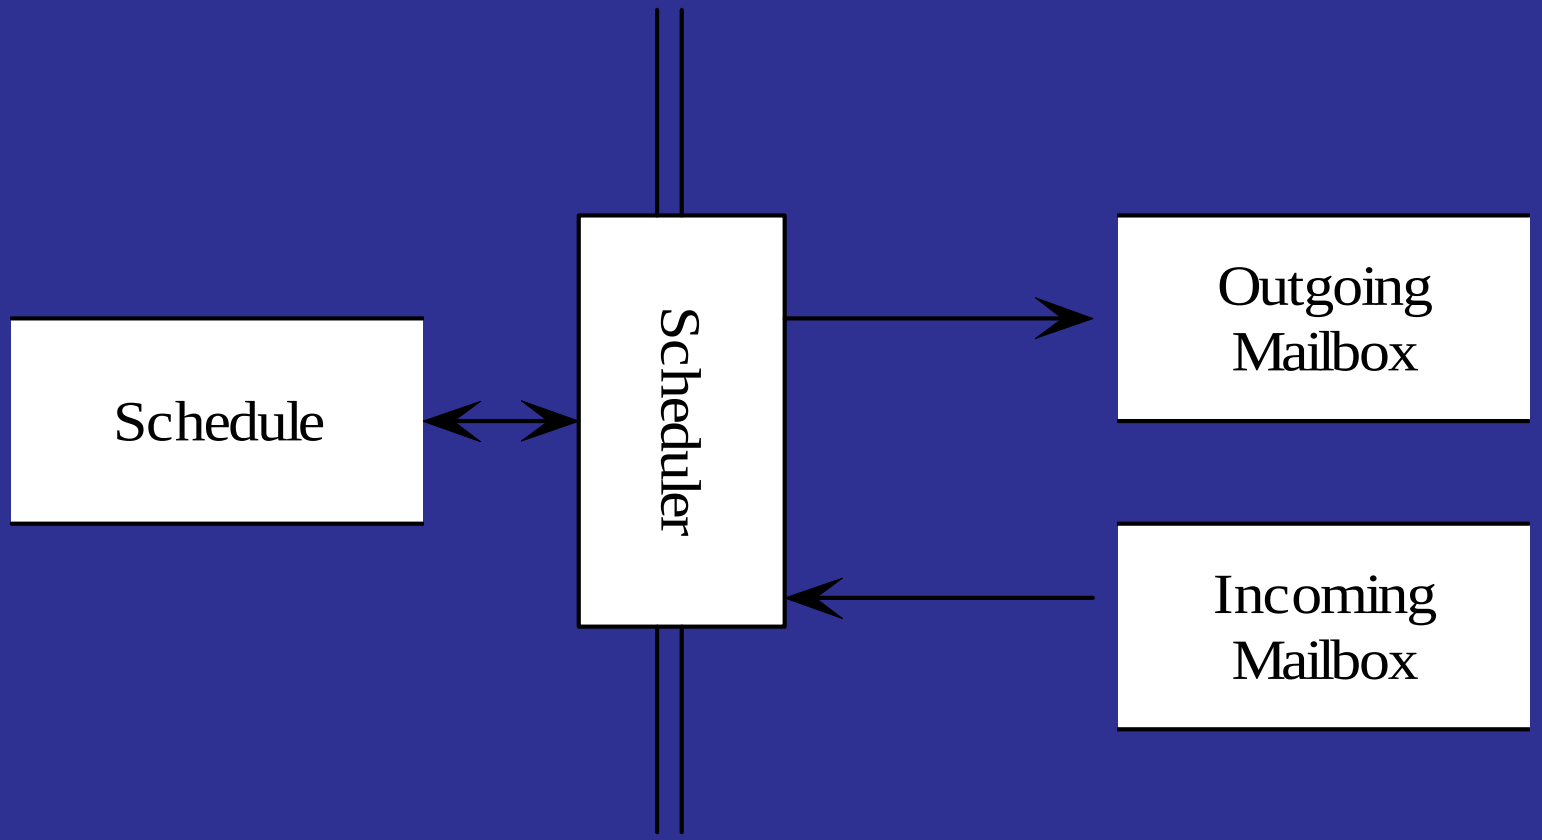

Negotiator

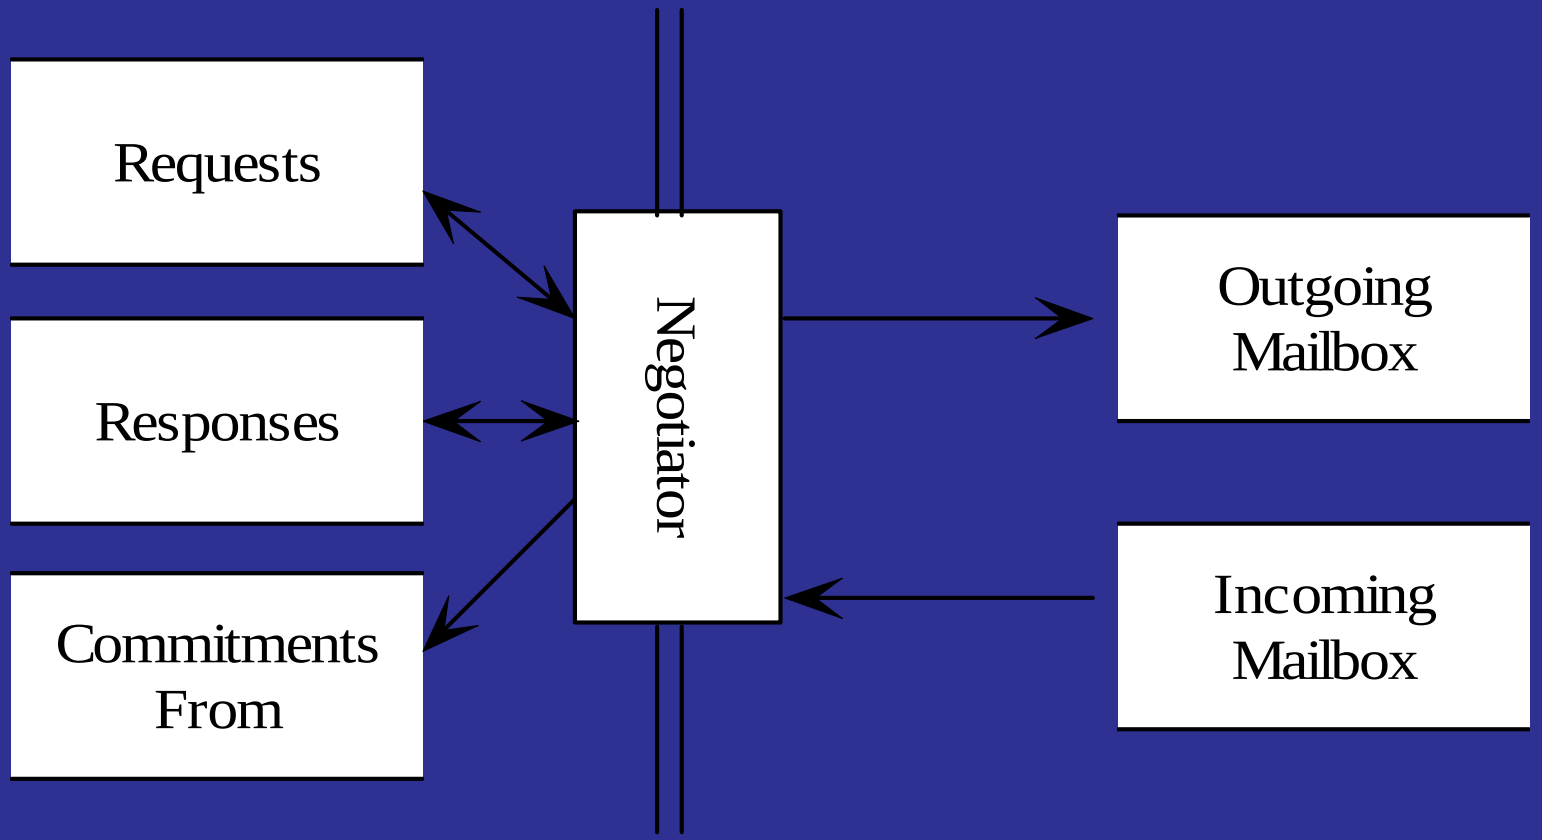
Requests
Responses
Commitments From
Negotiator
Outgoing Mailbox
Incoming Mailbox

Agent Thread (Decision Maker)

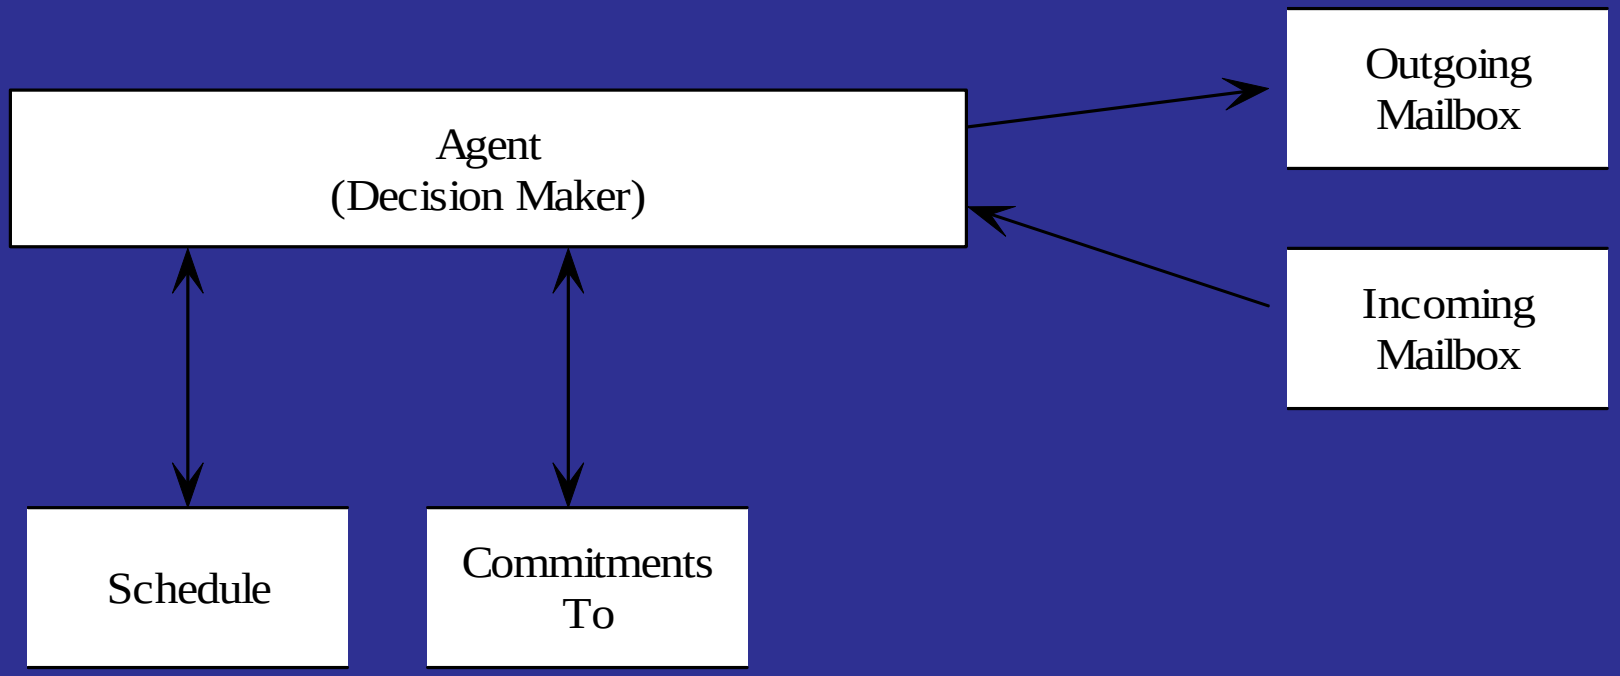

Agent Interaction

- Negotiated Request
- Negotiated Decommitment

Negotiated Request

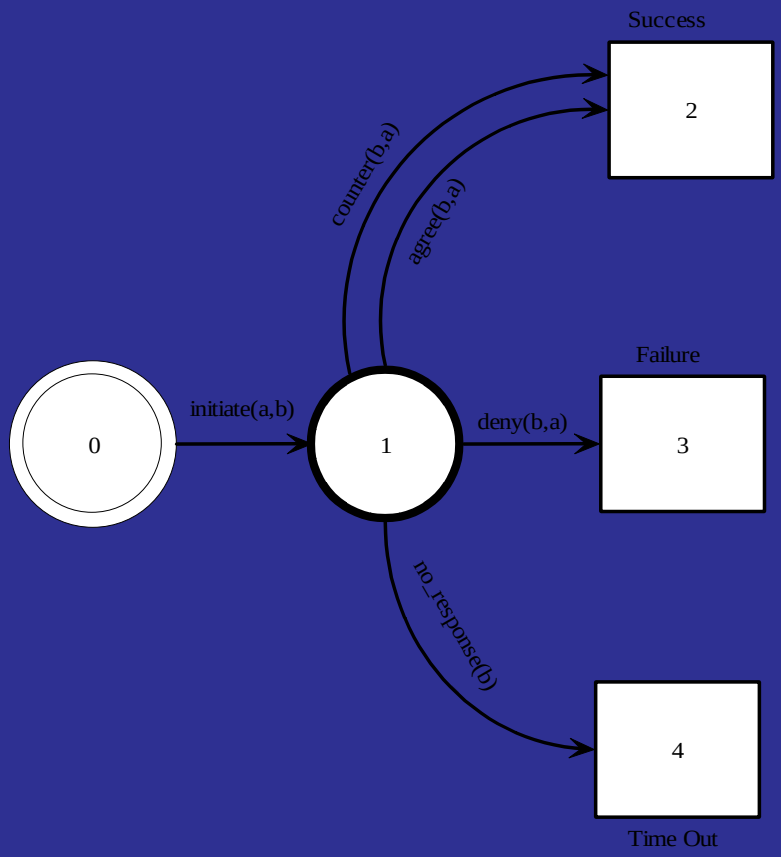

Negotiated Decommitment

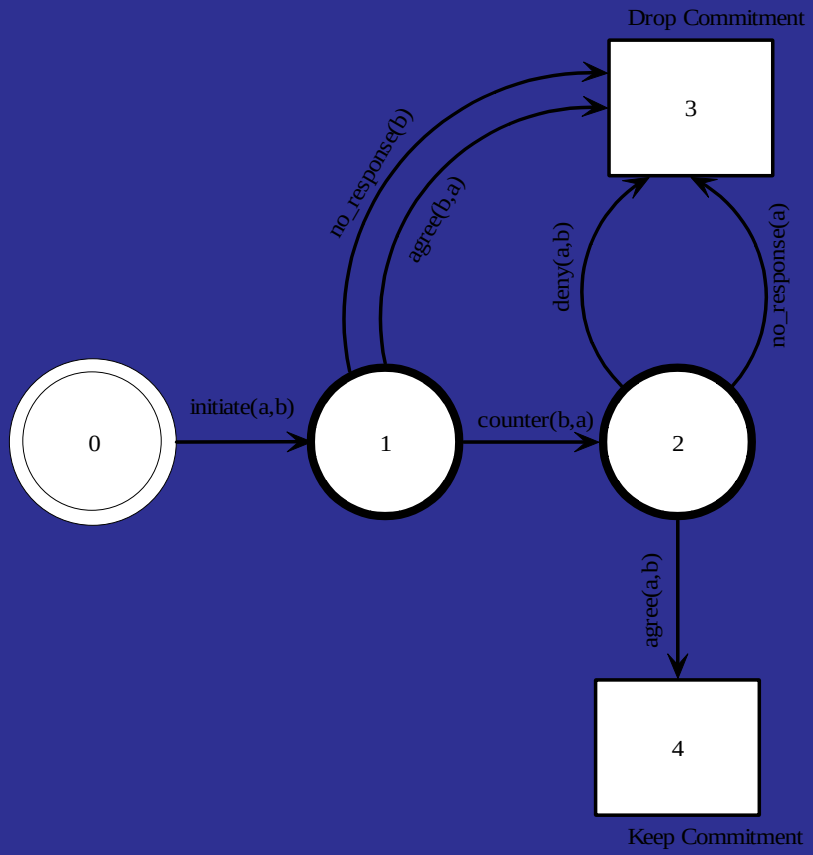

Local Estimate of Global Utility

- Commitment Value
  - $w_i = (p_i, v_{hi}, c_i, w_{hi}, dt_i)$
- Commitment Strength
  - $str_i = (n_i, r_{hi}, dnow_i)$

Decision Criteria

- Three Modes of Operation:
  - Baseline
  - Unilateral Decommitment
  - Negotiated Decommitment

Decision Criteria

- Incoming Sensor Information: Track Now
- Incoming Tracker Information
- Incoming Agent Information
  - Request to:
    - track now; assist later; decommit
  - Response to request
  - Notification of unilateral decommit

Incoming Sensor Information: Track Now

- Highest priority
- Operation Mode:
  - Baseline
  - Unilateral Decommitment
  - Negotiated Decommitment

Incoming Tracker Information

- Send “Track Now” request to agents with current visibility
- Send “Assist Later” requests to agents with projected visibility

Incoming Agent Information

- Requests
  - Track Now
  - Assist Later
  - Decommit
- Responses
- Notification of Unilateral Decommitment

Requests to Track Now

- Similar to “Track Now” task resulting from incoming sensor information except *value* re-assessed
- Operation Mode:
  - Baseline
  - Unilateral Decommitment
  - Negotiated Decommitment

Requests to Assist Later

- Based on projected target location
- Operation Mode:
  - Baseline
  - Unilateral Decommitment
  - Negotiated Decommitment

Requests to Decommit

- Negotiated Decommitment only
- Re-evaluation of initial commitment
- If lower, agree to decommitment
- If higher, make counter offer

Responses to Decommitment Requests

- Negotiated Decommitment only
- If all affected agents agree to decommitment, reduces to unilateral decommitment
- If any affected agent makes counter offer, re-evaluate commitment. If higher, then agree not to decommit

Experimental Design

Performance Evaluation Criteria

Experimental Conditions

Performance Evaluation Criteria

- Planned Measurements per Target
- Three or More Measurements in a Two Second Window per Target
- Balanced Measurements Across Multiple Targets
- Total Number of Measurements Taken
- Average Tracking Error

Experimental Conditions

- Variable:
  - Number of agents, number of targets, target speed
- Constant:
  - Sensor Placement
  - Target Placement
  - Target Path

Sensor Placement

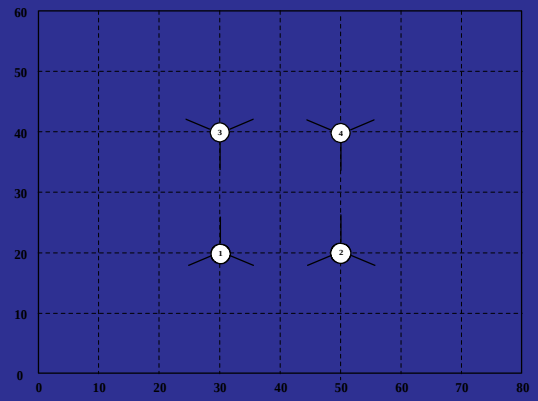

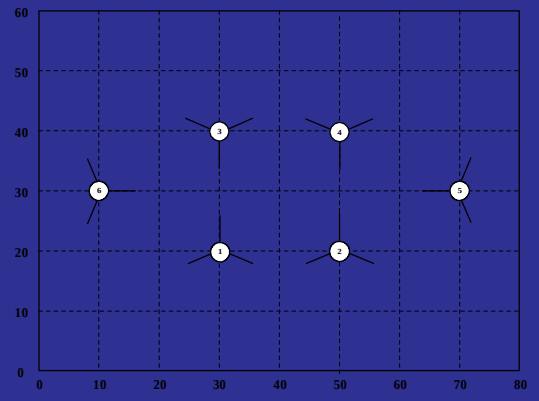

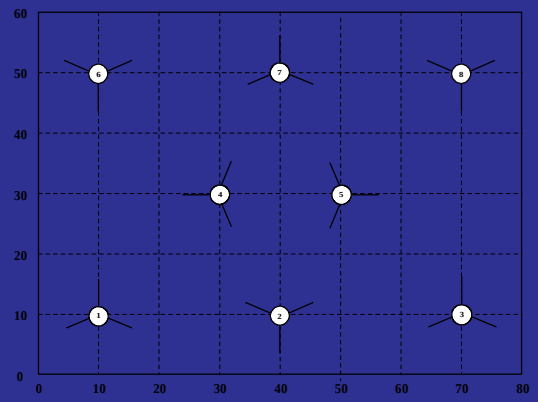

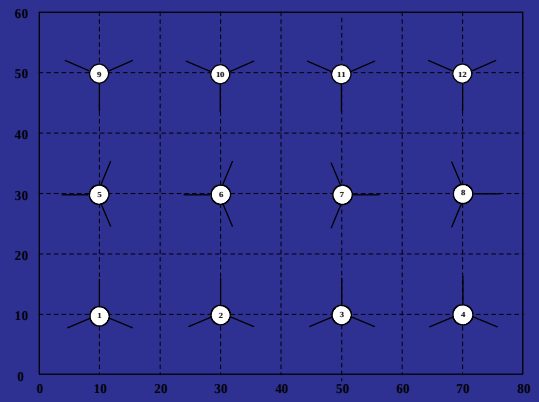

Target Placement

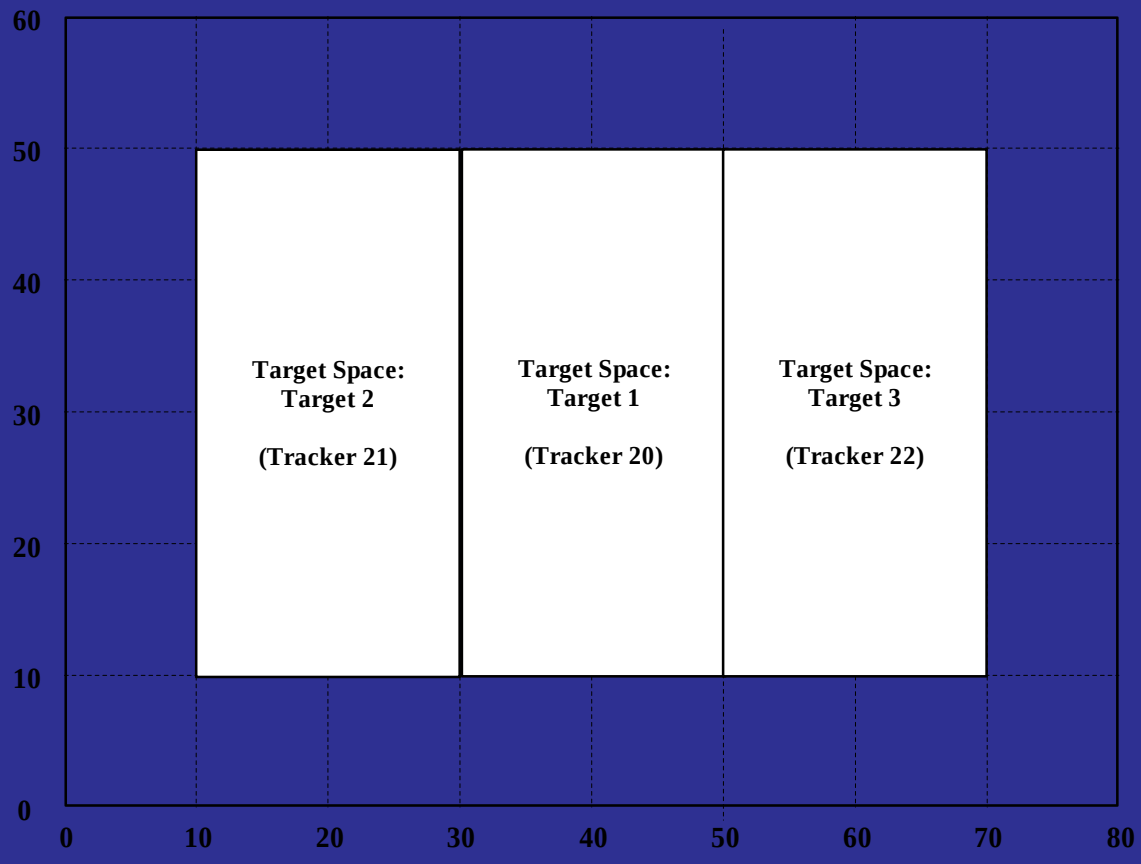

Target Path

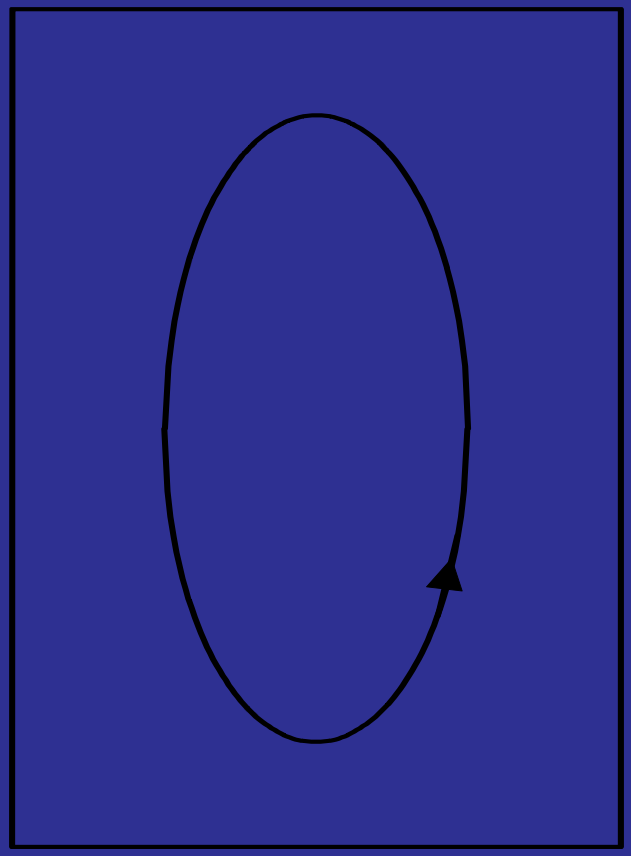

Results and Analysis

Agent Decisions

Overall Goal Achievement

Graceful Degradation of Performance

Discussion of Results

Agent Decisions

- Results based on a total of 134,096 agent decisions
  - Baseline: 46,722 agent decisions
  - Unilateral Decommitment: 44,712 agent decisions
  - Negotiated Decommitment: 42,662 agent decisions
- Average of 1241.63 decisions per condition

Overall Goal Achievement: Results for each of the Performance Evaluation Criteria

Planned Measurements per Target

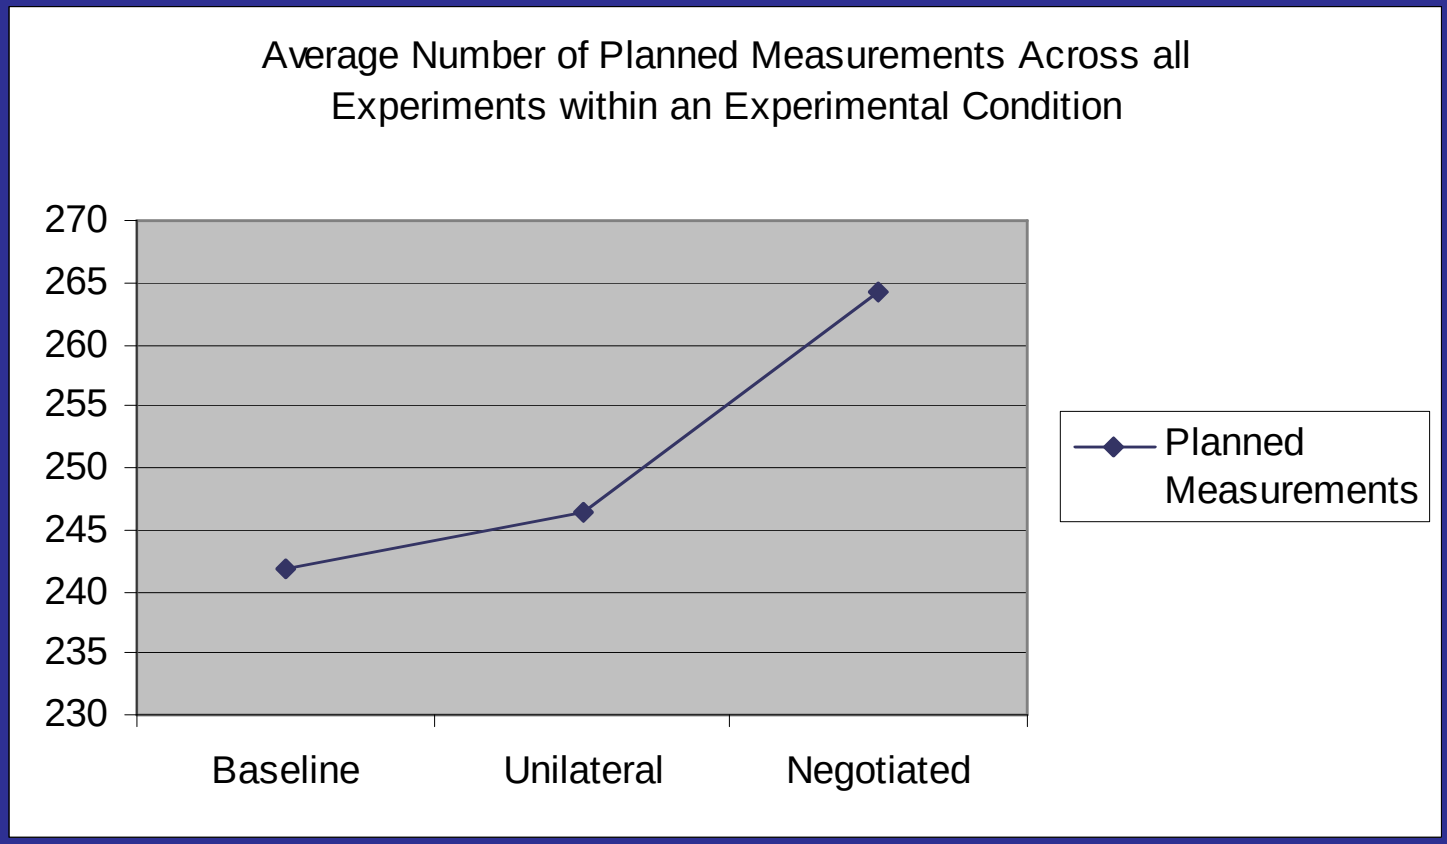

Three or More Measurements in a Two Second Window per Target

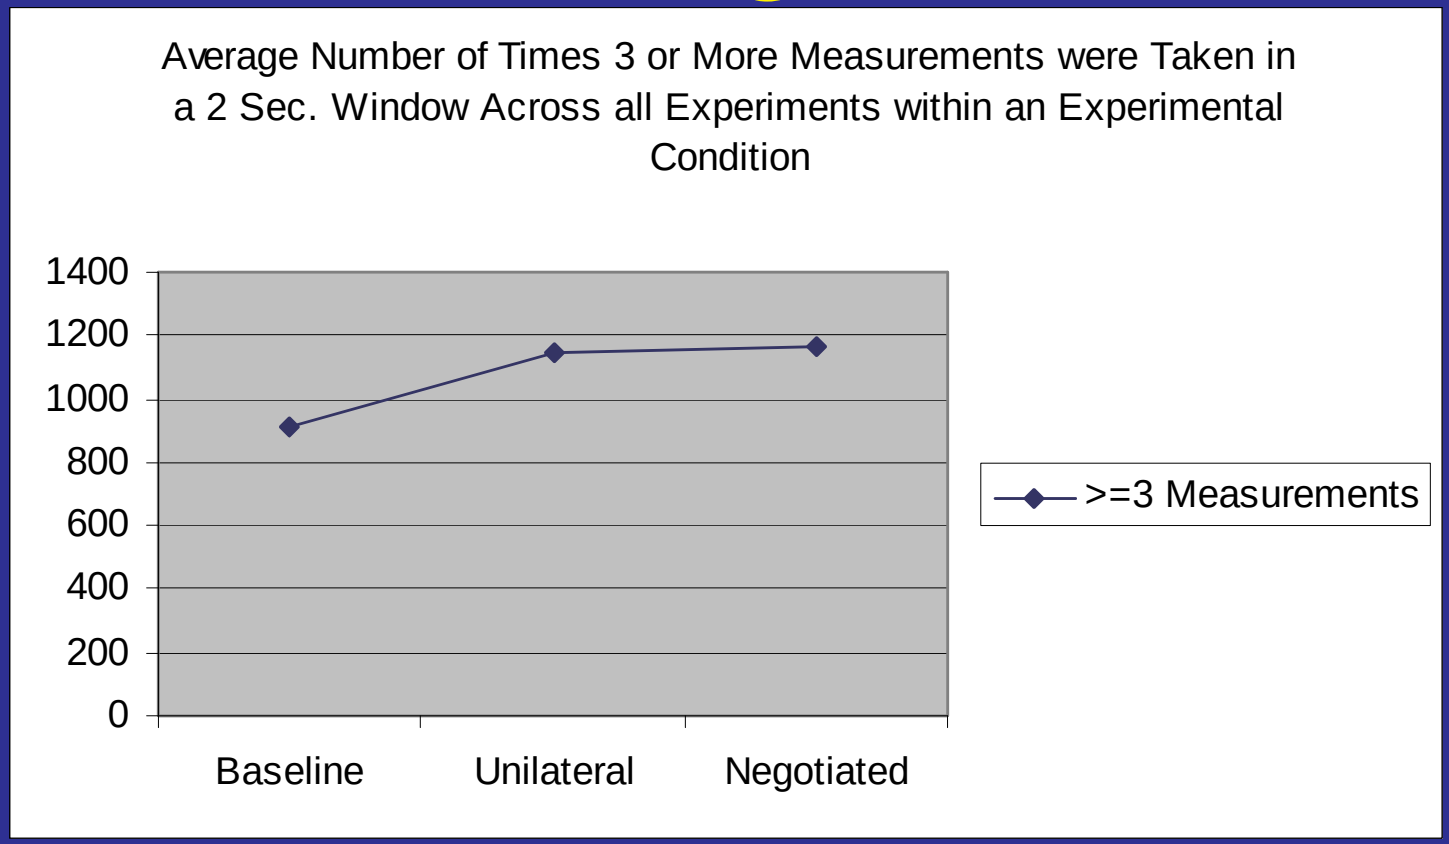

Balanced Measurements Across Multiple Targets

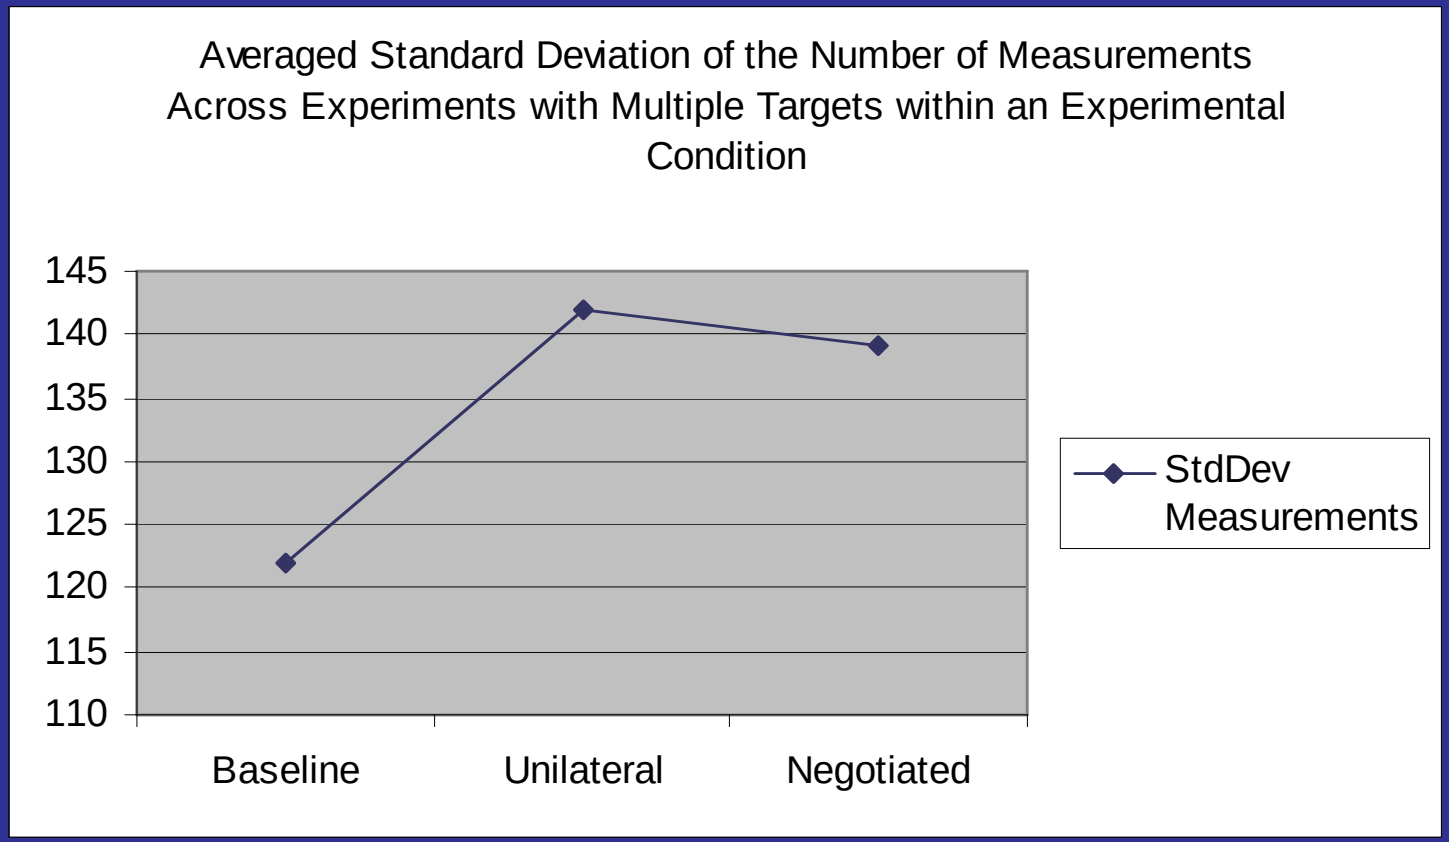

Total Number of Measurements Taken

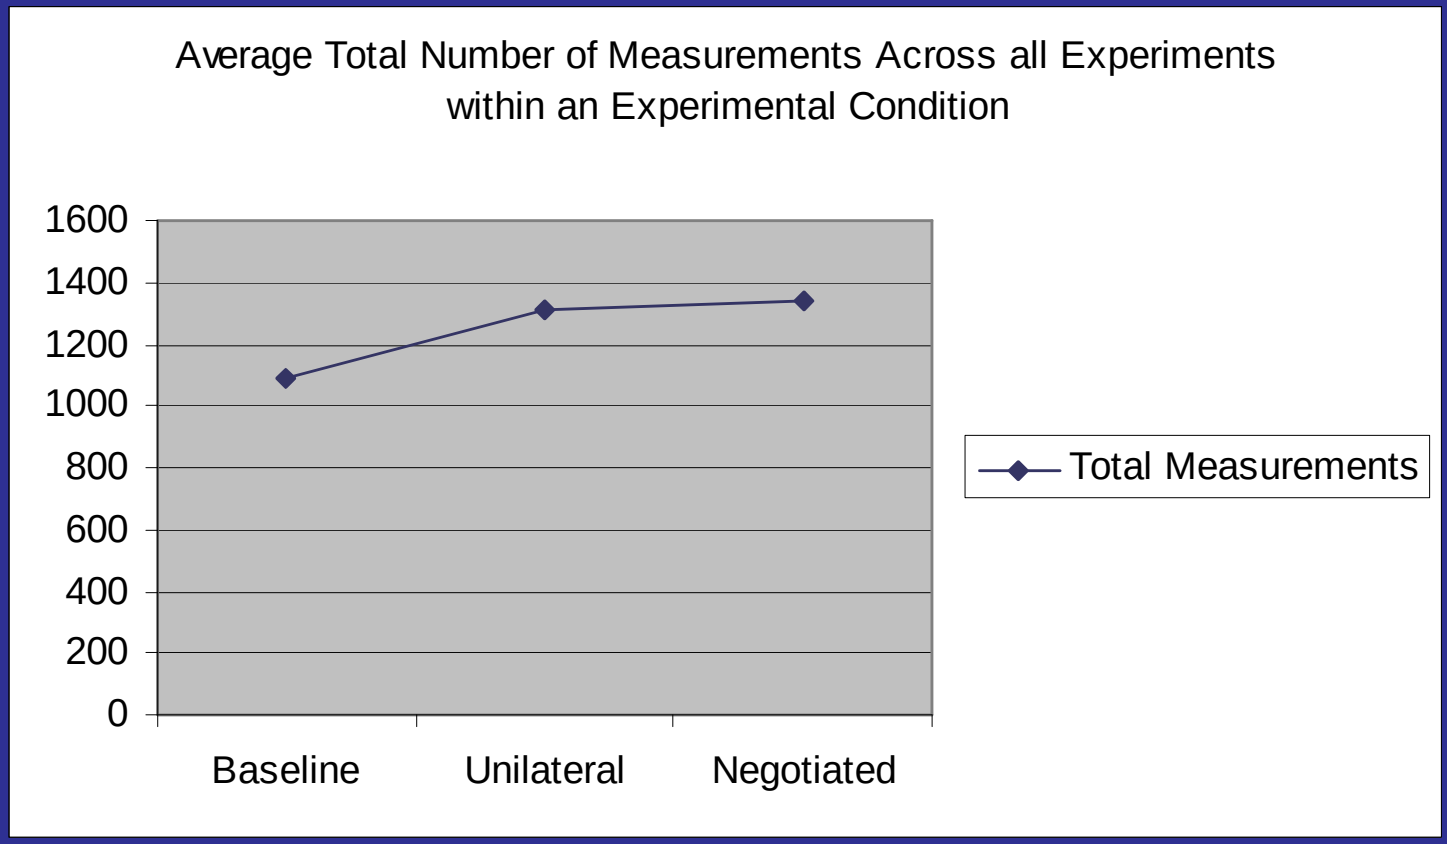

Average Tracking Error

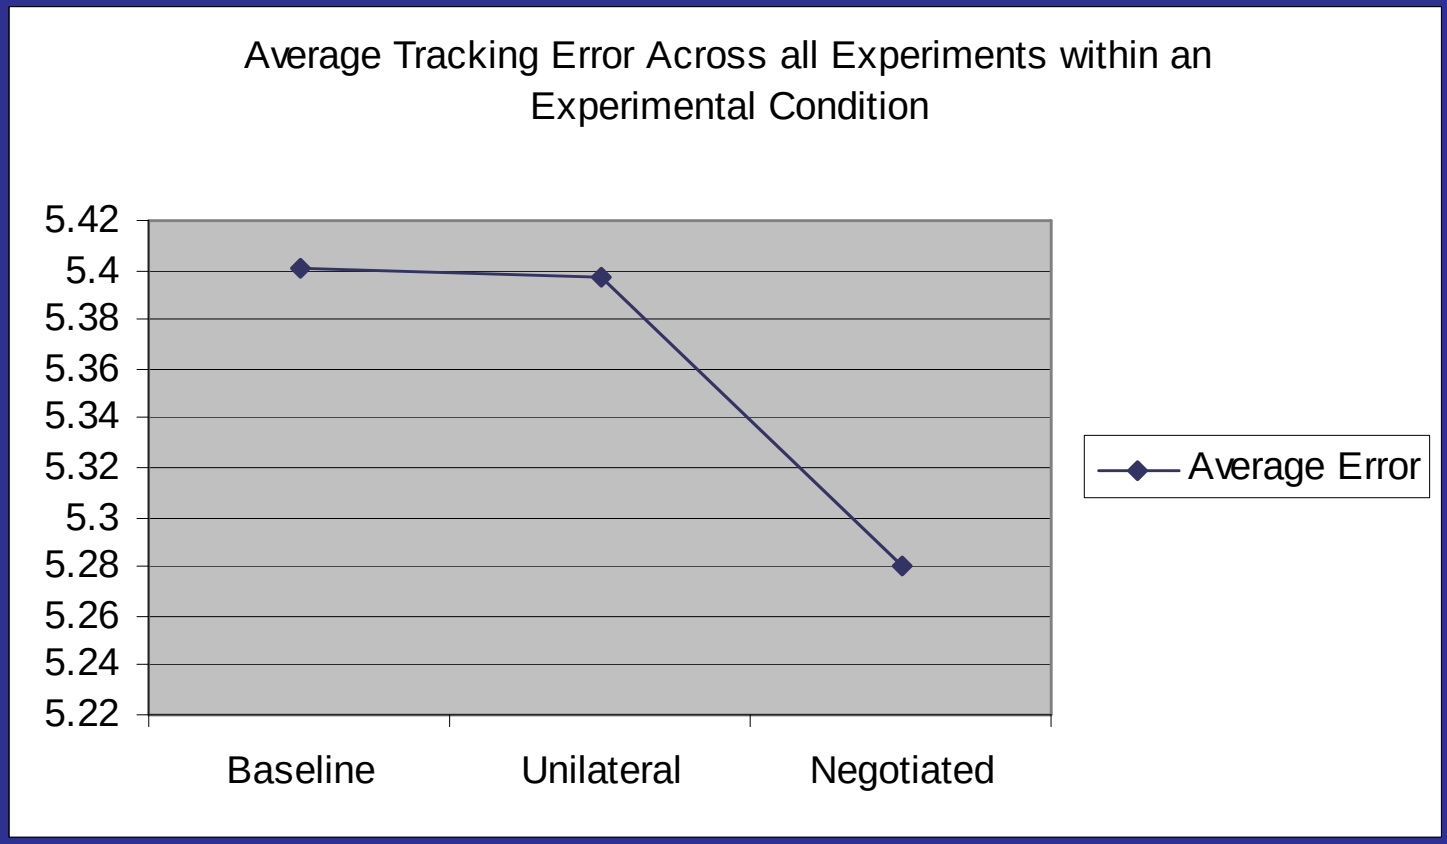

Graceful Degradation of Performance

- Constrainedness of Condition
- Average Tracking Error by Constrainedness
- Average Tracking Error by Target Speed
- Re-evaluation of Constrainedness
- Average Tracking Error by Constrainedness
- Evaluation Criteria by Constrainedness

Constrainedness of Condition

Agents	Targets	Target Speed	Constraint Value	Agents	Targets	Target Speed	Constraint Value
12	1	0.1	0.0083	8	1	1	0.1250
8	1	0.1	0.0125	8	2	0.5	
6	1	0.1	0.0167	12	3	0.5	
12	2	0.1		6	1	1	0.1667
4	1	0.1	0.0250	6	2	0.5	
8	2	0.1		12	2	1	
12	3	0.1		8	3	0.5	0.1875
6	2	0.1	0.0333	4	1	1	0.2500
8	3	0.1	0.0375	4	2	0.5	
12	1	0.5	0.0417	6	3	0.5	
4	2	0.1	0.0500	8	2	1	
6	3	0.1		12	3	1	
8	1	0.5	0.0625	6	2	1	0.3333
4	3	0.1	0.0750	4	3	0.5	0.3750
6	1	0.5	0.0833	8	3	1	
12	1	1		4	2	1	0.5000
12	2	0.5		6	3	1	
4	1	0.5	0.1250	4	3	1	0.7500

Average Tracking Error by Constrainedness

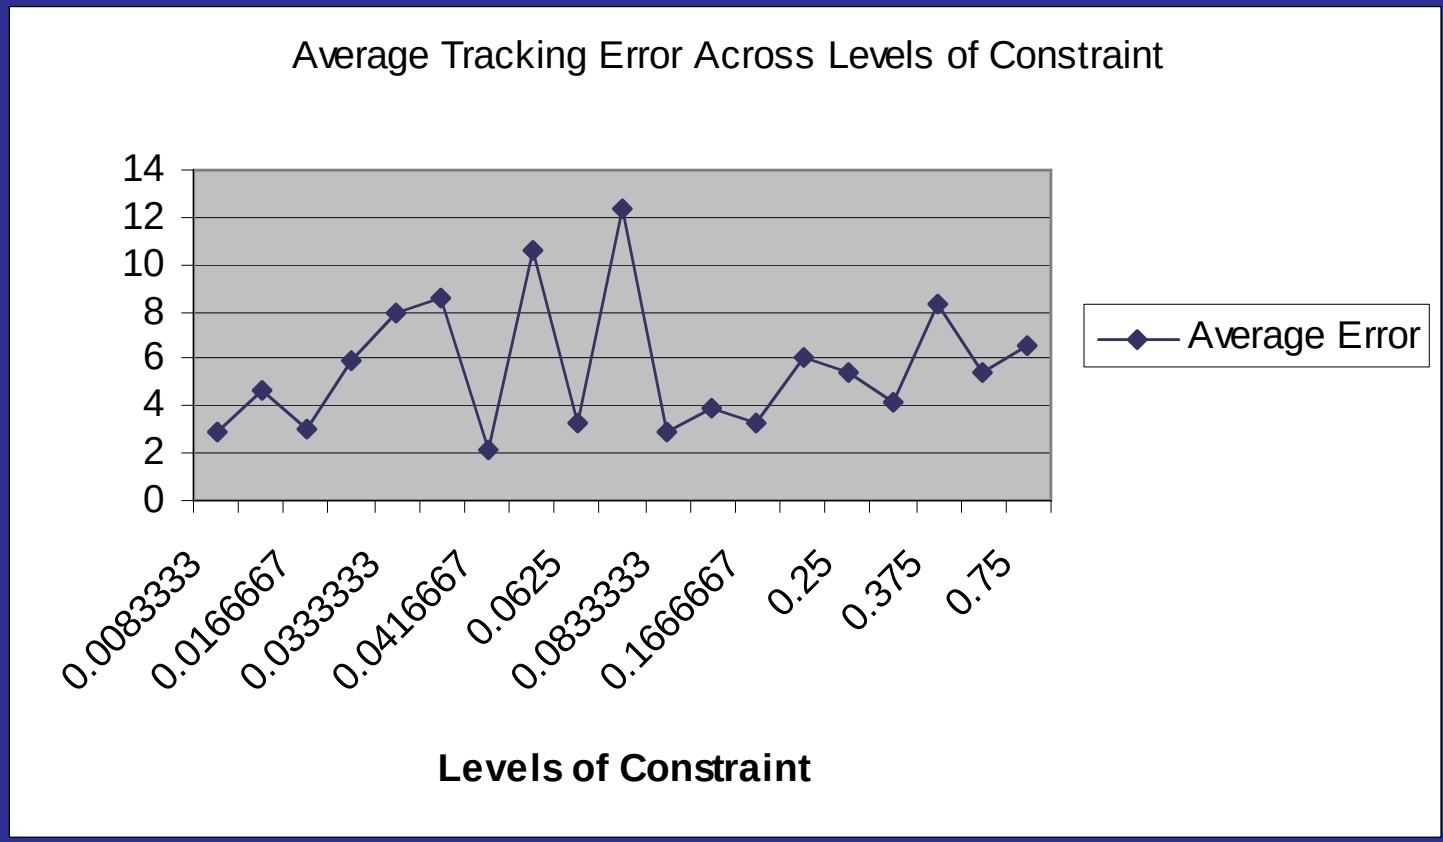

Average Tracking Error by Target Speed

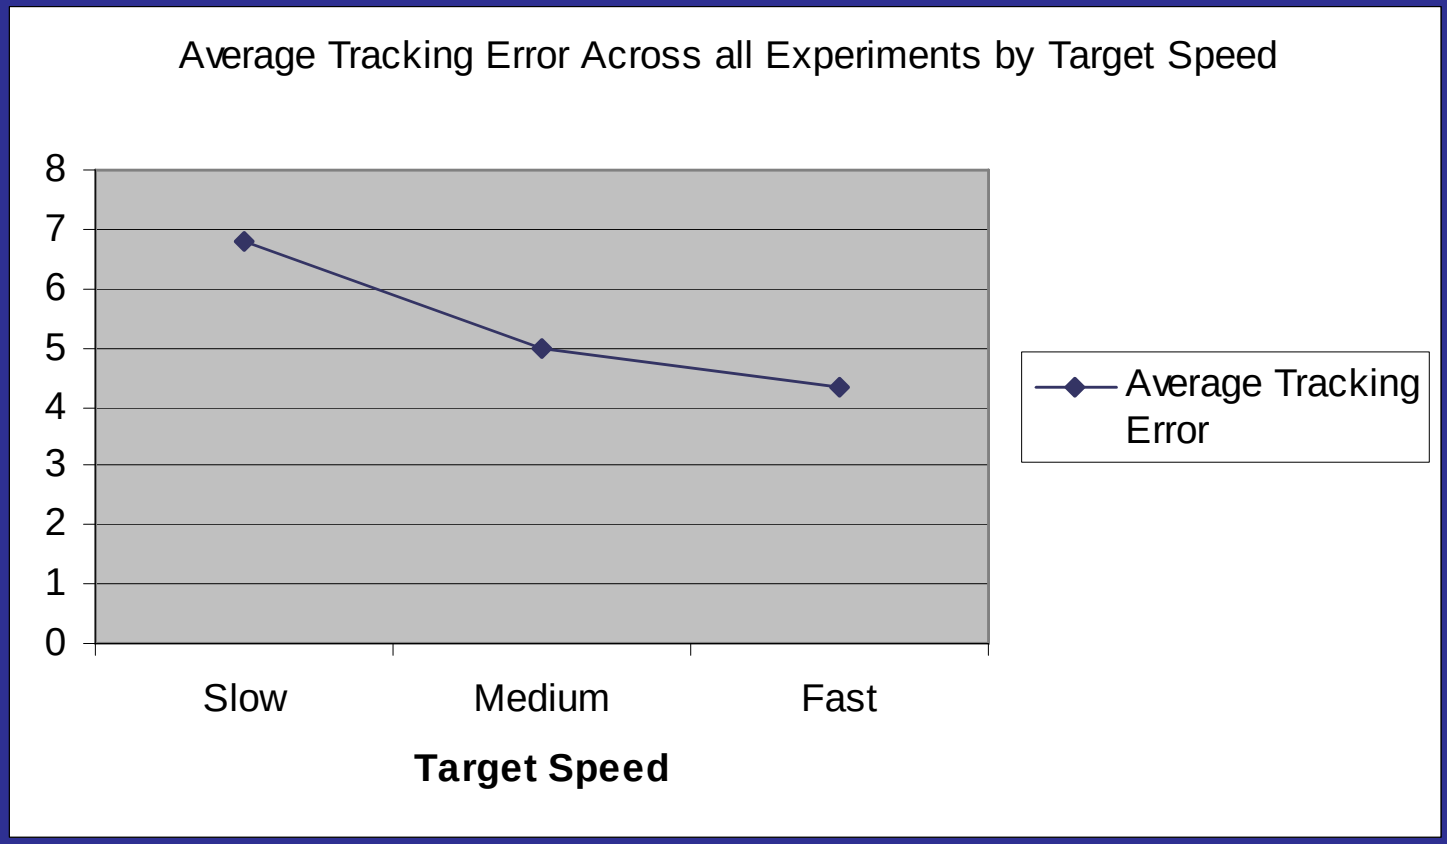

Re-evaluation of Constrainedness

Agents	Targets	Level	Agents	Targets	Level
12	1	.083	12	3	.25
8	1	.125	6	2	.33
6	1	.167	8	3	.375
12	2	.167	4	2	.5
4	1	.25	6	3	.5
8	2	.25	4	3	.75

Average Tracking Error by Constrainedness

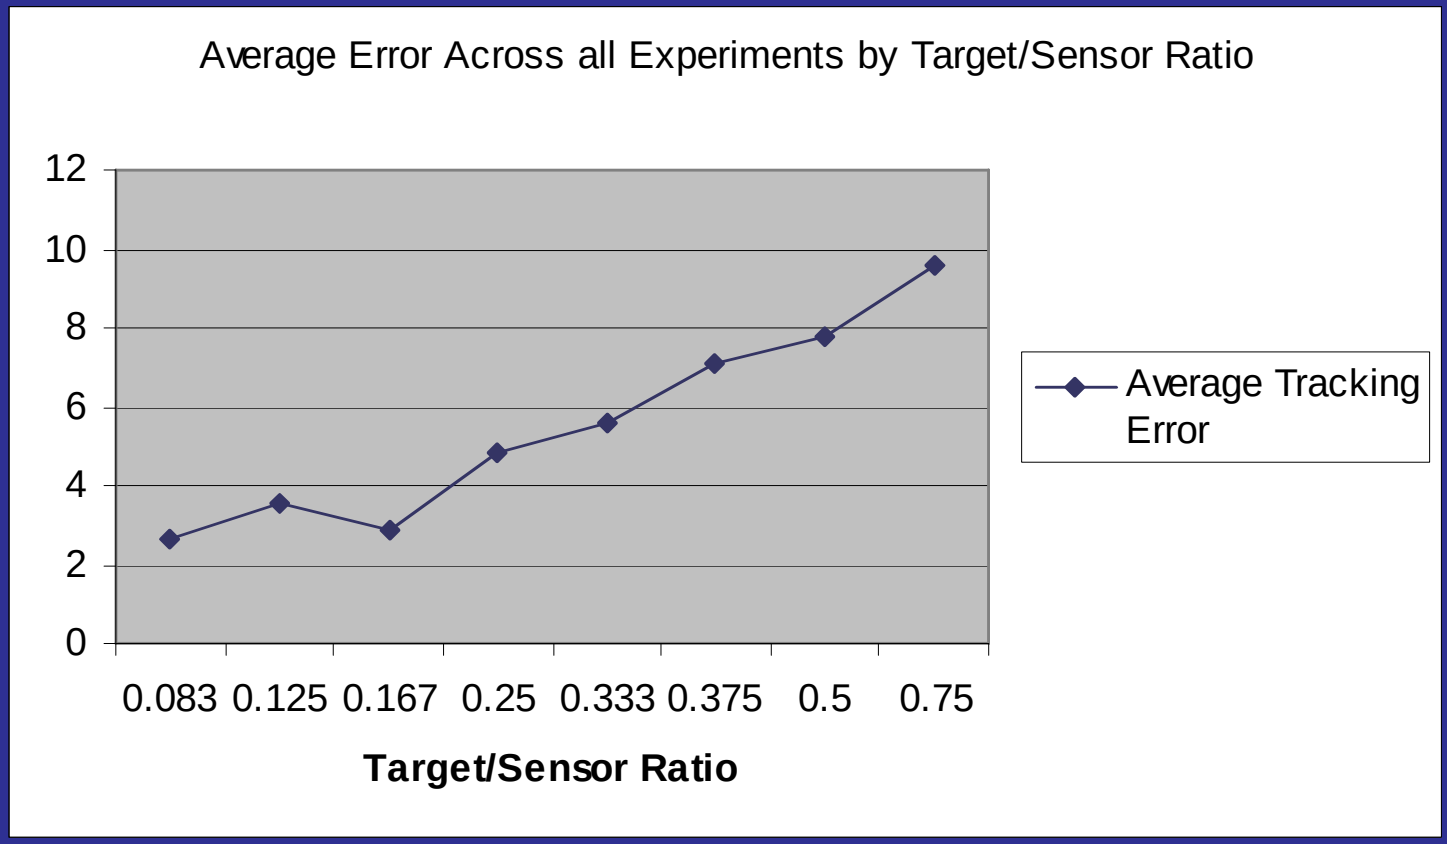

Graceful Degradation of Performance by Constrainedness: Results for each of the Performance Evaluation Criteria

Planned Measurements

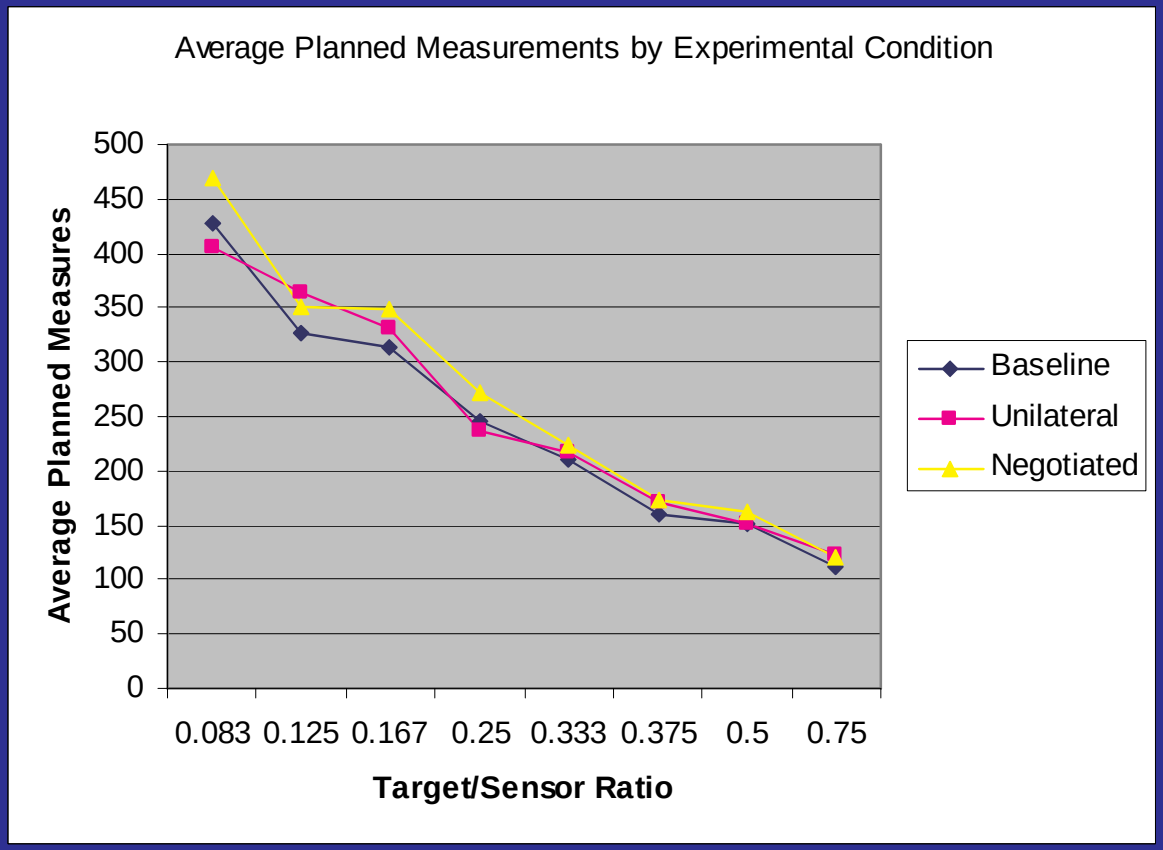

Three of More Measurements in a Two Second Window

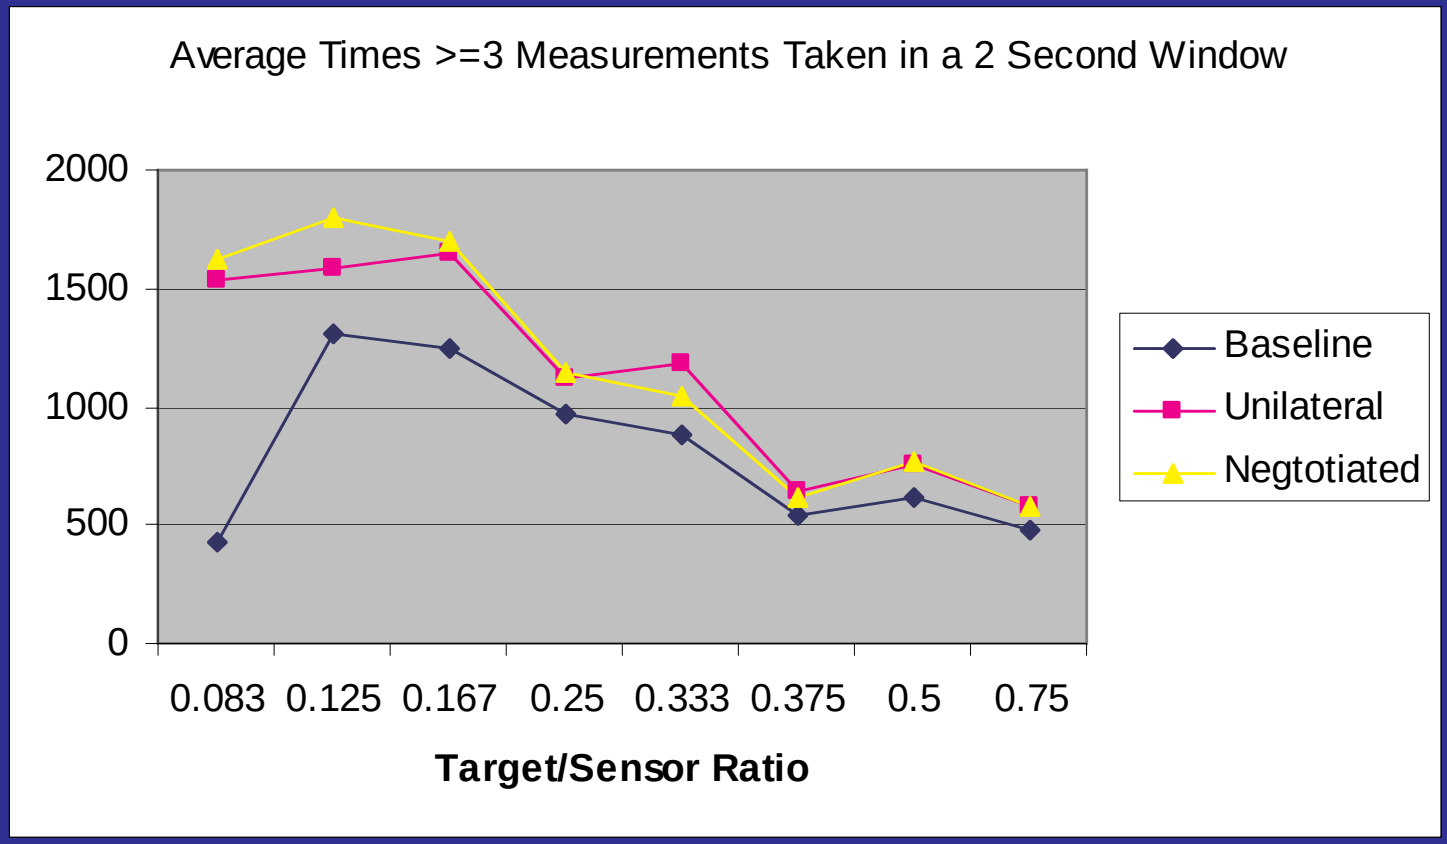

Balanced Measurement Across Multiple Targets

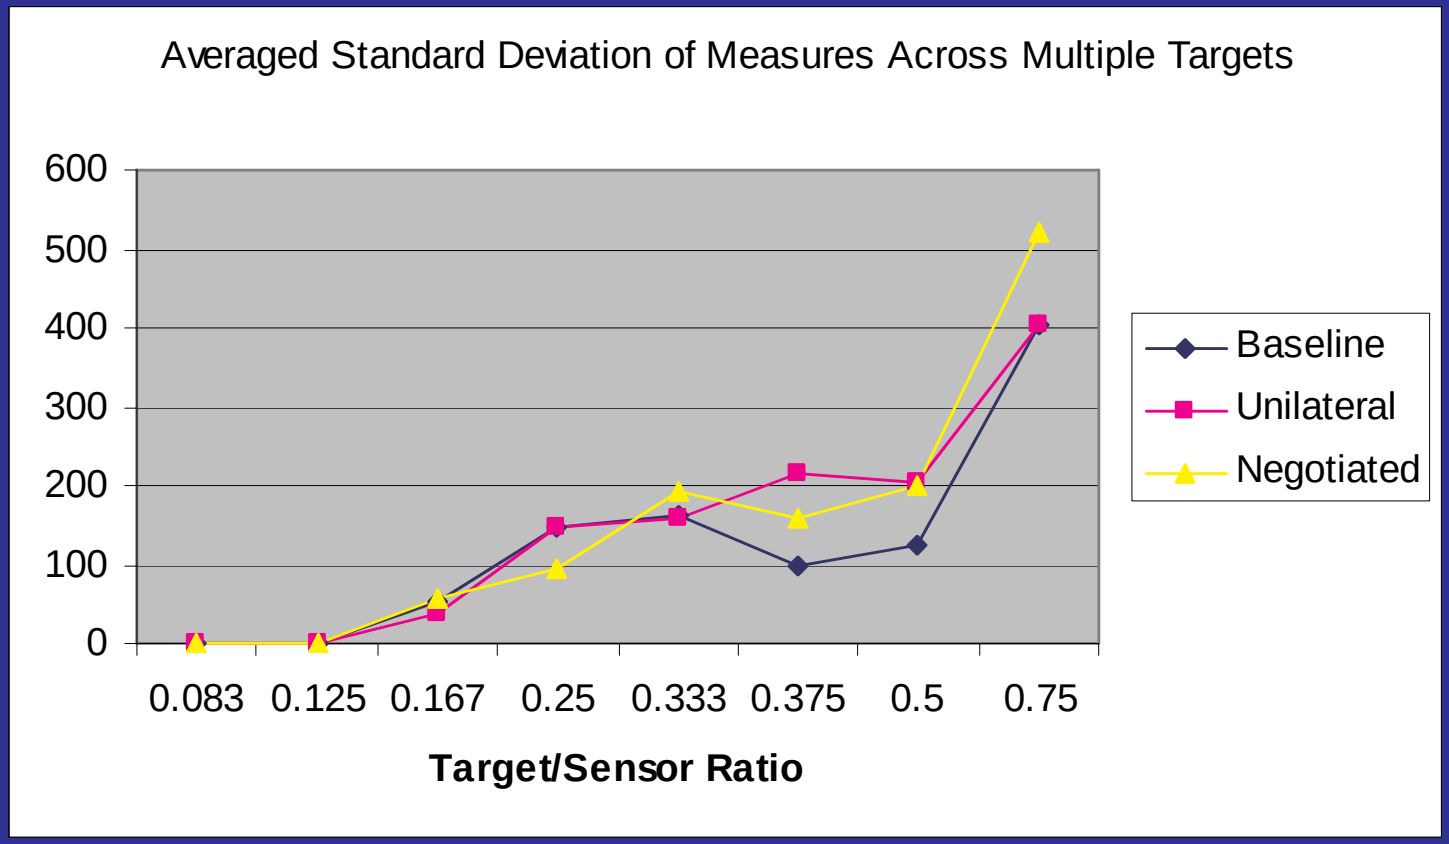

Total Measurements per Target

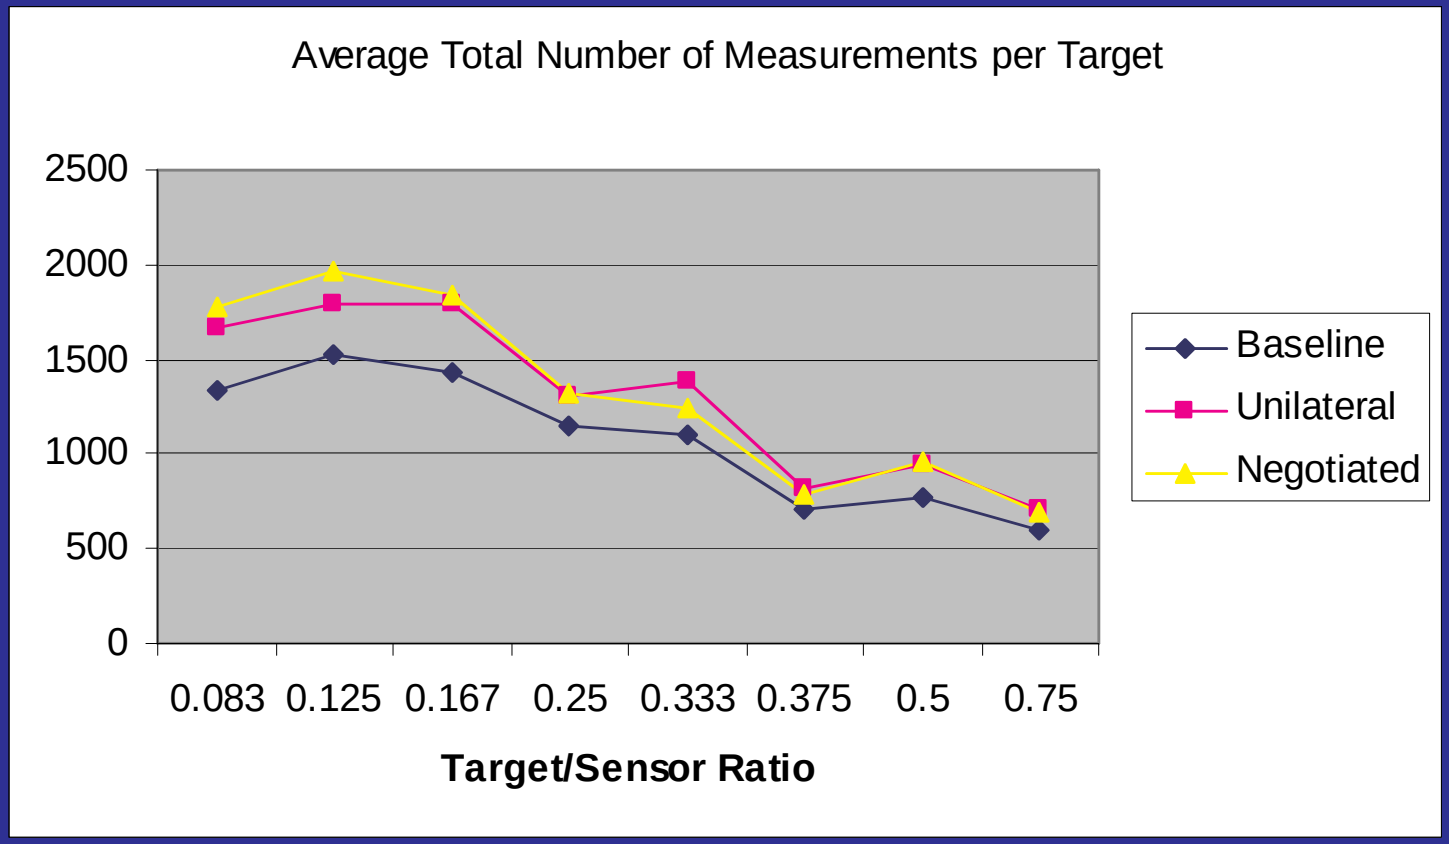

Average Tracking Error

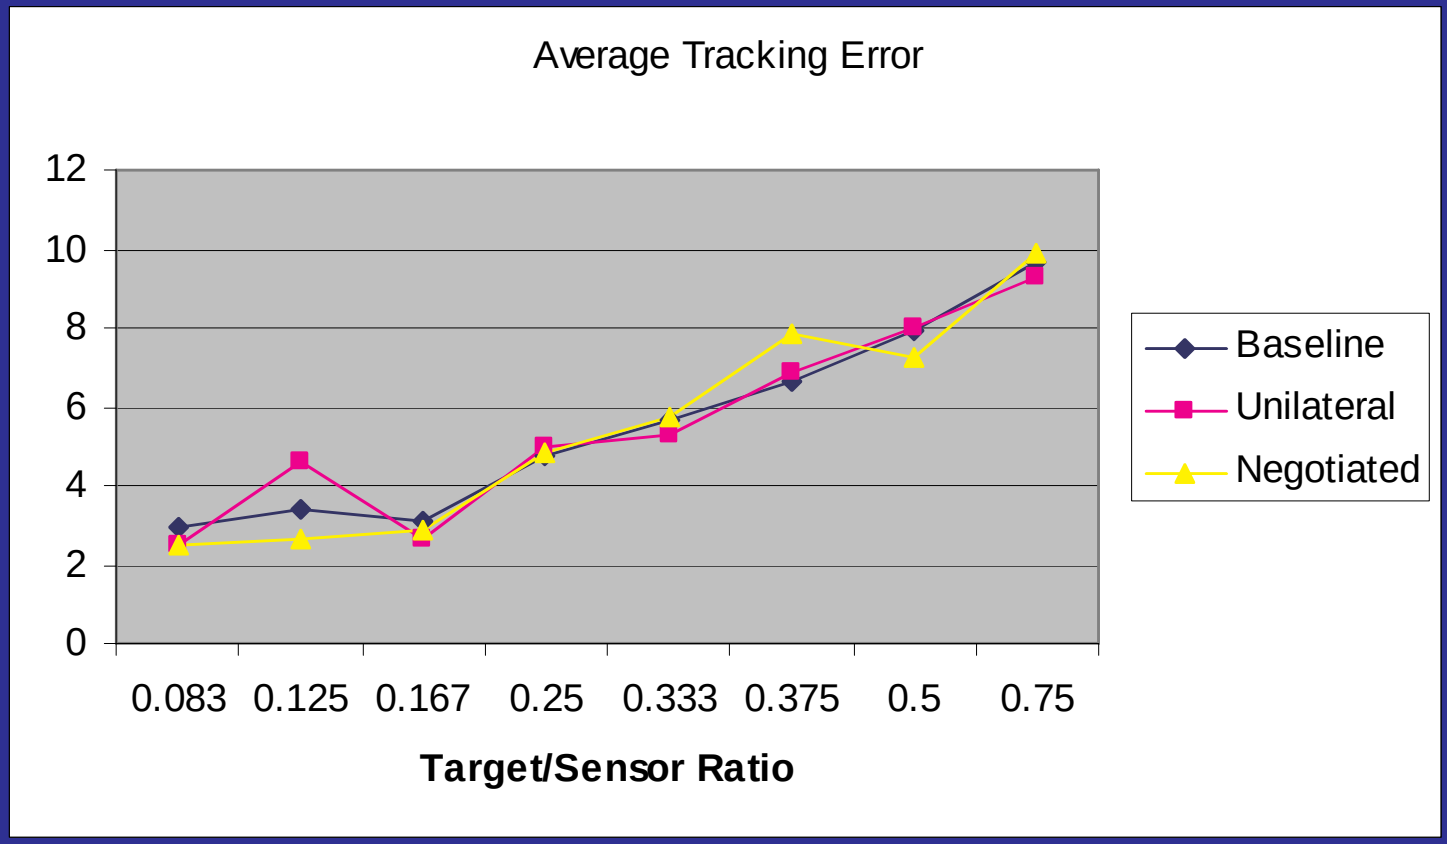

Discussion of Results

Overall Goal Achievement

Graceful Degradation of Performance

Overall Goal Achievement

- Evaluation criteria showed improvement, except balanced measurements
- Magnitude of improvement from unilateral to negotiated decommitment not as high as expected

An Example

- Requested commitment: 6.387
- Scheduled commitments: 6.309
- Baseline - Can’t decommit: 6.309
- Unilateral - Decommit: 6.387
- Negotiated - Received counter offer: 8.907

Graceful Degradation of Performance

- Evaluation criteria showed graceful degradation of performance with increasing constraints, except balanced measurements
- Neither decommitment condition showed improvement over the baseline condition

Conclusions

Significance

Future Directions

Significance

- Negotiated decommitment has not been previously addressed in the literature
- Unilateral decommitment has been studied, primarily in self-interested agent societies

Significance (cont.)

- Research results support all three hypotheses:
  - Unilateral decommitment improves goal achievement over baseline condition
  - Negotiated decommitment improves goal achievement over unilateral decommitment
  - Graceful degradation of performance under increasing constraints

Future Directions

- Domains with different characteristics:
  - Increased reliability of future predictions
  - Reduced communication bottleneck
- Sensitivity testing of commitment *value* and *strength* measures
- Investigation of implications of target speed on system performance