Robust Human Interaction with Robotic Swarms

Katia Sycara
The Robotics Institute
Carnegie Mellon University
katia@cs.cmu.edu

Outline

- Introduction
- Framework for human supervisory control
- Human control of independent robots
- Scheduling of operator attention
- Scalable multi-robot displays
- Human control of swarms
  - Control via Leader
  - Optimal human control input timing
- Conclusions

Overview

Research Motivation: As robotic platforms become more reliable and low cost, many tasks can benefit from deployment of multiple robots. Humans and robots must coordinate and interact.

- **Applications**
  - **Military**
  - **Emergency response/Urban Search and Rescue**
  - **Service**
    - **Geriatric (intelligent wheel chairs, companions, etc.)**
    - **Public/social use (hospital drug carts, nursebot, etc.)**

Research Challenge: Validated schemes for human control of robotic teams are currently lacking

Research Objective: Develop and validate theories, models and techniques for *scalable, flexible* and *adaptive* control of multiple robots and system organizations so as to exploit their synergies to achieve effective and predictable performance

Carnegie Mellon

Many forms of interaction are needed

Robots

- must be individually controllable
- must function autonomously for long periods
- must be commandable as cooperating teams
- must adapt to absence of human attention
- must incorporate humans in autonomous plans

Key Idea: Look at HRI from viewpoint of complexity of operator's cognitive complexity of command

This framework allows systematic study of human control of multi-robot systems

•Neglect Tolerance Model (NTM)—for O(n) control

•Bio-inspired swarms—for O(1) control

As size grows, complexity of command dominates

Carnegie Mellon

THE ROBOTICS INSTITUTE

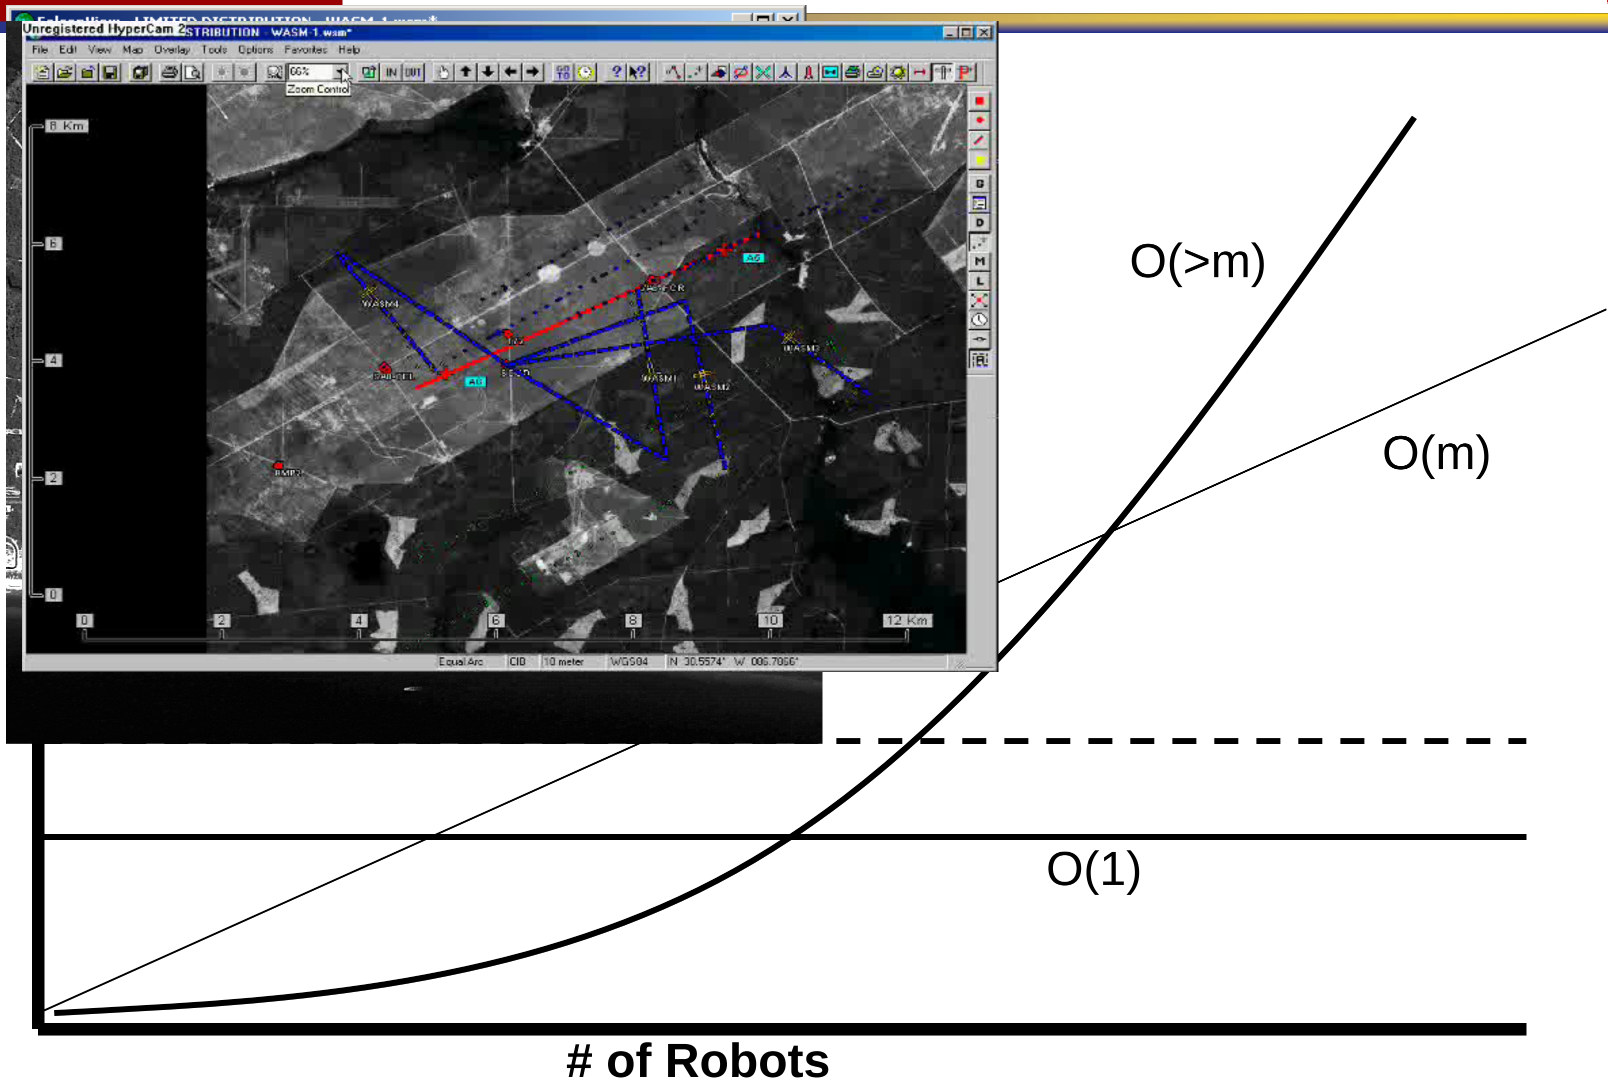

Carnegie Mellon

Research Questions

- **How can people control multiple robot teams of increasing size?**
  - **What is the density of robots a human(s) can control?**
  - **What kinds of command are possible for a particular density?**

Key Constraint: Human attention is limited; attention is the budget

- **How many “things” can a human manage?**
  - Robots
  - Tasks
  - Other people
  - Sources of information

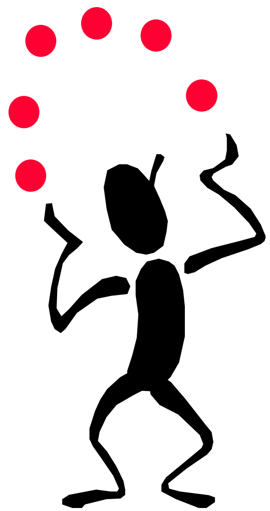

Scheduling Human Attention to Improve HRI Performance

Improving O(n) performance

Robots act independently of one another

Background

- Most autonomous systems require some degree of interaction with a human operator to achieve a desired behaviour, e.g. due to failures that require repair by the human, or due to change in mission, where the human may want to impart a new goal to the system

- Concept of **Neglect Tolerance** for these systems has been well studied in the literature
  - Idea: How well does the system perform when neglected by a human operator?
  - Assumption: If system is neglected, its performance will decrease.

Note, however that we have observed that for robot swarms, neglecting the swarm, i.e. delaying control input, may increase swarm performance. We call this **Neglect Benevolence.**

- Nagavalli, S., Luo, L., Chakraborty, N., Sycara, K., Neglect Benevolence in Human Control of Robotic Swarms, International Conference on Robotics and Automation (ICRA), Hong Kong, China, May 31-June 7, 2014.

Carnegie Mellon

Span of Control: The *fan-out* Metric

THE ROBOTICS INSTITUTE

- For independently operating robots, an operator can control other robots during robot1's NT

- So.. The number of robots that might be controlled is

- Fan-out (homogeneous teams)

$$\mathsf{FanOut} = \mathit{floor}\left(\frac{\mathsf{NT}}{\mathsf{IT}}\right) + 1$$

- This is an upper bound, many practical limits

Research Issues for O(n) control

- **Control of independent robots**
  - Extend Fan-out and the Neglect Tolerance model to more realistic assumptions
- **Can human attention be scheduled?**
  - Can performance be improved by directing attention?
  - Can operator utilization be optimized?
- **Scheduling models for HRI**
  - Using scheduling theory for HRI architectures and prediction
    - Models for directing operator attention to most effective task
- **M human- N robot teams**
  - Call center & other architectures

Assumptions for Aiding HRI through scheduling

- **A subset of HRI tasks can be performed independently**
- **Queuing Model**
  - **Operator can be treated as a server**
  - **Robots in need of attention can be treated as jobs**
- **Operator attention can be "scheduled" to service specified robots without loss of efficiency**
- **Scheduling algorithms could specify improved allocation of operator attention by choosing the order in which jobs are serviced**

Carnegie Mellon

Typical MrCS interface

THE ROBOTICS INSTITUTE

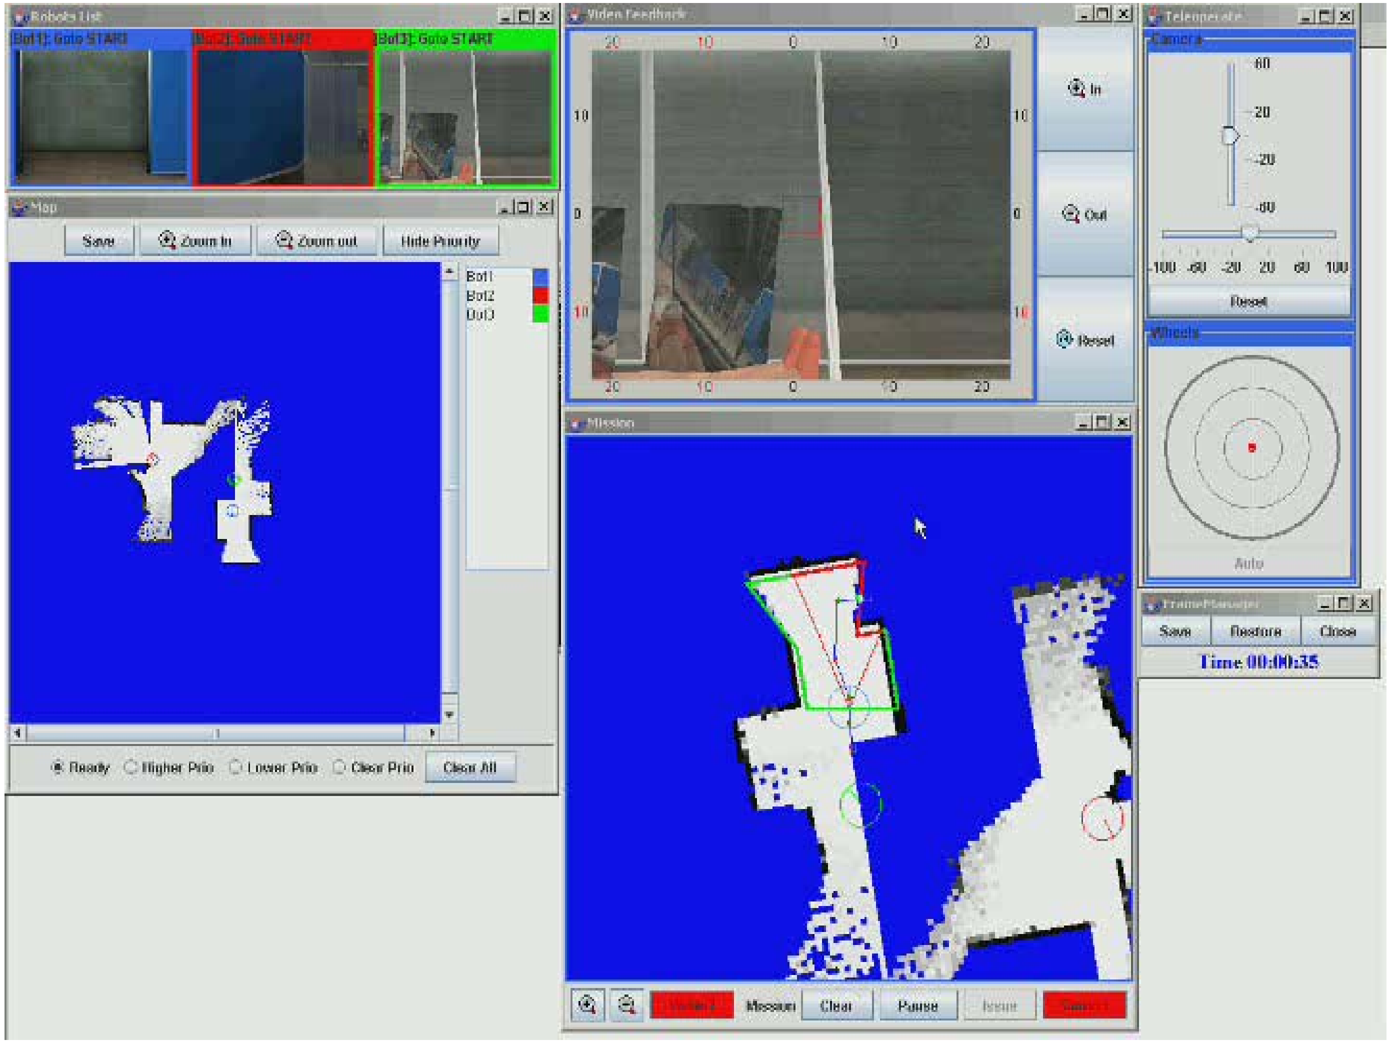

Can Operator Attention be Scheduled?

Carnegie Mellon

THE ROBOTICS INSTITUTE

+ 1

Experiment with **homogeneous** failures robot comparing:

•**Control** condition, i.e. unaided operator

•**FIFO** queue directing operator to “best” robot to service next

•**Alarm** panel showing all robots reporting failures

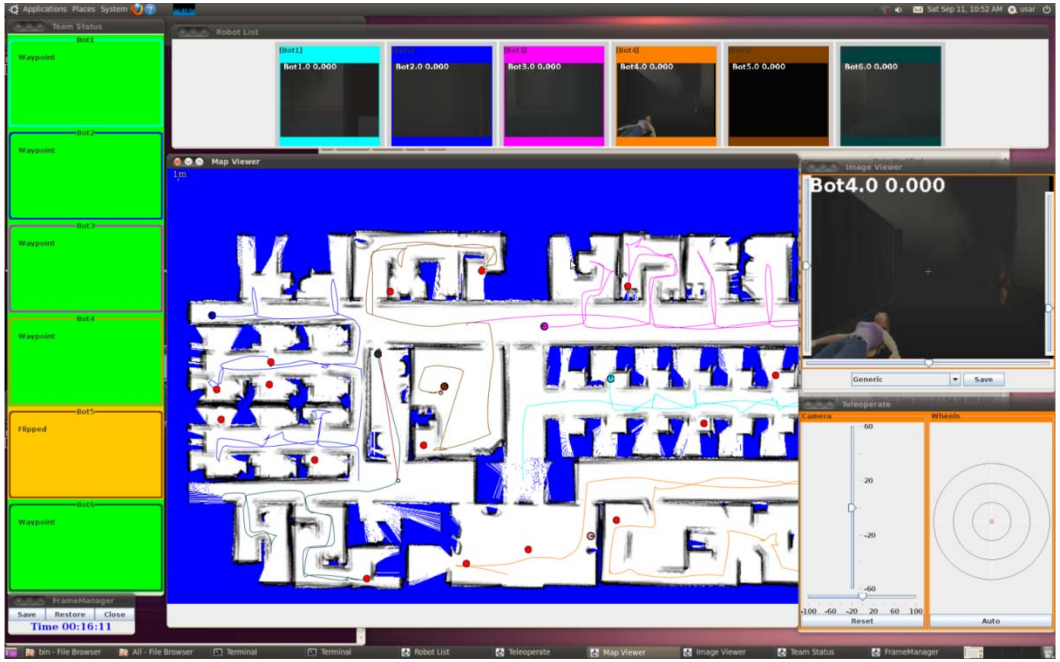

Operator must **both** search and mark victims (primary task) and monitor and maintain robots’ status (secondary task)

- Attention direction via FIFO or alarms
- Monitoring for robot status could be eliminated

Attention Guidance Experiment 1

Select robot-to-repair sooner

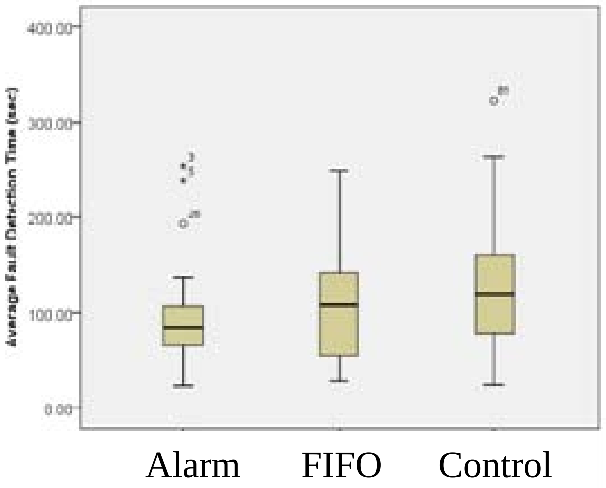

Ordering:

1. Alarm
2. FIFO
3. Control

Faults "repaired" sooner

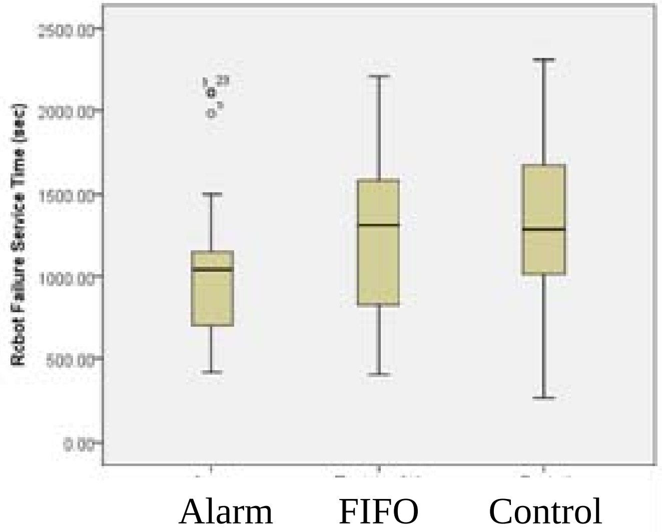

Victims detected sooner

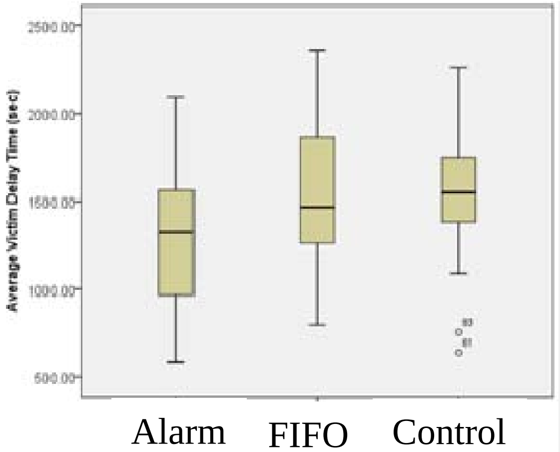

Because failures are homogeneous FIFO and Alarm conditions should be equivalent if operator's attention can be effectively directed

Conclusions

- Alerting operator to failures improved performance
- Directing attention (FIFO) to a particular robot was less effective than merely alerting

Experiment 2-Operator attention CAN be scheduled to improve performance

Heterogeneous Failures

Failure	Description	Time to Resolve
Stuck	Robot was stopped by approaching obstacles	short
Teleop Lagged	Robot executed operator's command with 2~3 seconds delay	intermediate
Camera Sensor Failed	Robot's video feed was frozen right before the failure happened	short
Map Viewer Failed	**Robot's position on the map viewer unable to update**	**Long**

Shortest Job First (SJF) discipline yields maximum throughput

Experiment compared:

- SJF queue
- FIFO queue
- Alarm

Chien, S., Mehrotra, S., Brooks, N., Lewis, M. & Sycara, K. Scheduling operator attention for multi-robot control, 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'12), October 7-12, Villamoura, Portugal, 2012

SJF performance did not differ significantly from the BEST performance on measures

- **SJF and Alarm best on:**
  - **Latency to resolve failures**
  - **Time between selecting and fixing failure**
  - **Failures resolved**
- **SJF and FIFO best on:**
  - **Missed victims**
  - **Mental workload**

Evidence for Cognitively Efficient Task Switching

Carnegie Mellon

THE ROBOTICS INSTITUTE

- Experiment 1:
- If attention can be directed, because time/effort are the same across failures, then Alarm = FIFO
    - However, experiments show Alarm > FIFO
- Experiment 2:
    - If SJF discipline is followed, more robots will be repaired for a given level of interaction. Therefore we expect SJF > Alarm or FIFO
    - Operators in Alarm condition did not follow SJF
        - However, Result SJF~Alarm, and SJF or Alarm > FIFO
- Scheduling attention improved performance but NOT to the extent predicted

Developing Scalable Displays for Robot Teams

O(1) displays for multiple moving cameras

As Number of UVs grows large

- Control must revert to O(1) or at most O(n) schemes

BUT

- Data acquisition/analysis cannot yet be fully automated
  - Many operators are currently needed to control UV platforms
  - Information exploitation accounts for 77% of human effort during Predator operations
- So,.. Bottleneck is aggregating & filtering information
  - Cognitive overload

Image Queue

- priority queue of images sorted by *uniqueness*
- robot captures a video frame
  - frame is added to database along with
    - robot pose
    - laser scan
    - uniqueness score
      - subtract the occupancy grid from the camera view estimate obtained from the laser scan data
      - calculate the remaining area, which is the map coverage unique to that image.
  - image with the highest score is added to the update priority queue
  - its camera view is permanently subtracted from the occupancy grid

Experiment compares MrCS with/without Image Queue

Calculate Redundant Imagery

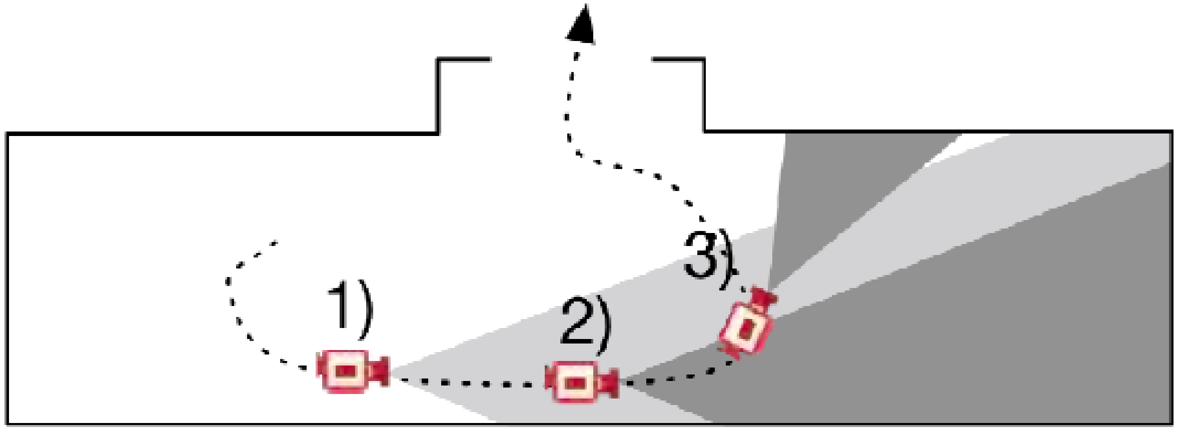

- **Image 1 has the highest utility as it has the largest visual coverage**
- **Image 2 has no utility as its coverage is entirely overlapped by Image 1**
- **Image 3 has a small utility as it provides some coverage Image 1 does not**

Significant Reduction in Errors without loss in performance

Teleoperation Workload

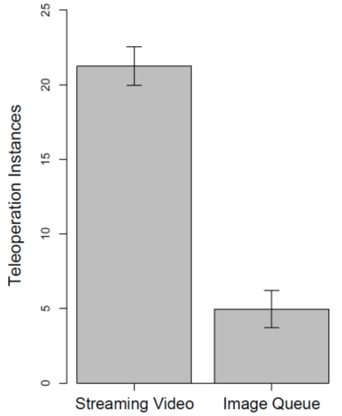

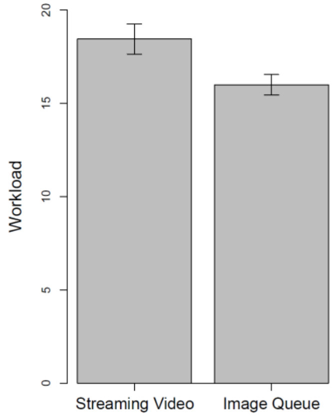

False Alarms

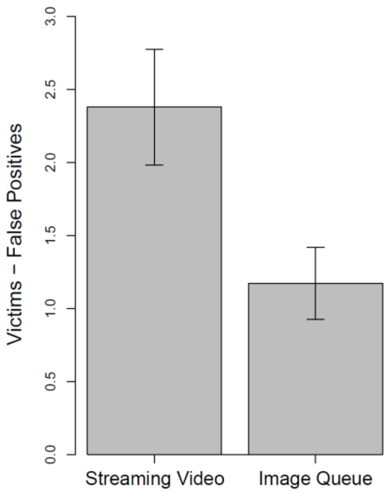

Misses

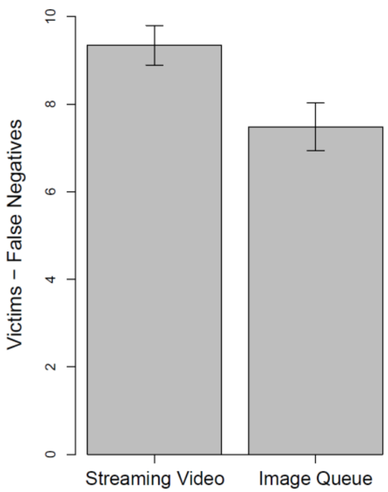

70% coverage 1st 10 frames
99% coverage 100 frames

4 orders of magnitude reduction in bandwidth @ 70% coverage

(filters six 15 min streaming videos into 10 static frames)

Can add more robots without overloading operator

Scalable Control Regimes for Robot Swarms

Carnegie Mellon

Swarm Robotics

- Large collection
- Limited capabilities
- Simple local rules (attraction, repulsion, cohesion)
- Complex emergent behavior (e.g flocking)
- Applications
  - Bridging
  - Search and rescue
  - Surveillance and Reconnaissance
  - Exploration
- Main benefits
  - Simplicity of implementing control laws
    - Robots only use information about direct neighbors
  - Scalability of system
    - Adding robots requires minimal system reconfiguration

Why human control of swarms?

Recognize & mitigate shortcomings of autonomy

Utilize information inaccessible to autonomy

Human Role

Convey changes in intent

Carnegie Mellon

Neglect Benevolence in Human-Swarm Interaction

THE ROBOTICS INSTITUTE

- **Human experiments with simulated swarms showed that delaying the human input could improve performance**
  - **Walker, Nunnally, Lewis, Kolling, Chakraborty, Sycara, 2012**
  - **We termed this phenomenon "Neglect Benevolence" (NB)**
- **Formalized the notion of Neglect Benevolence and applied it control of robot swarms**
  - **Nagavalli, Luo, Chakraborty, Sycara, 2014**
- **Examples of NB implications for Human-Swarm Interaction**
  - **Delaying the human input may be beneficial when the desired performance objective is:**
    - **Minimize time for robotic swarm to reach a goal**
    - **Have robotic swarm perform a task by a given deadline**

Control via Leaders

1. **Explicit Leadership**
   Human influences swarm through control of a leader whose status is recognized by other swarm members and propagated (token) along with the influence
   switching, behavioral parameters, triggers

Distributed Robotics

influence through recognized leader(s)

2. **Tacit Leadership**
   Human influences swarm by controlling member(s) who then influence others but without any indication that influence originated with a "Leader"

Bio-Inspired

Influence without leader recognition

Carnegie Mellon

Research Study

THE ROBOTICS INSTITUTE

A comparison between Explicit and Tacit leaders in influencing swarm consensus

Quality Measures

- **Convergence time**
- **Robustness (noise tolerance)**
- **Effects of graph structure (size and connectivity) on convergence time**
- **Effects of swarm's movement (changing connectivity graph) on convergence conditions**

Amirpour S., Walker, P., Lewis, M., Chakraborty, N., Sycara, K. Explicit vs Tacit Leadership in Influencing the Behavior of Swarms, International Conference on Robotics and Automation (ICRA), Hong Kong, China, May 31-June 7, 2014.

Flooding (Explicit) vs. Consensus (Tacit)

Flooding

20 iterations

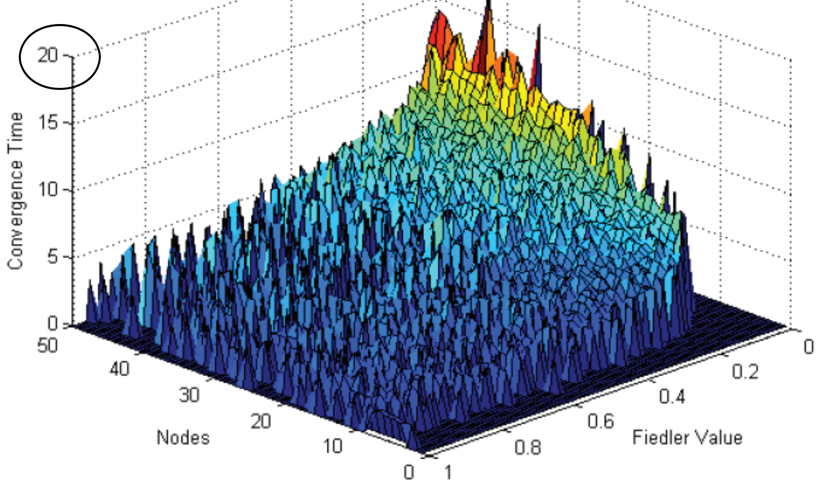

Consensus

4,000 iterations

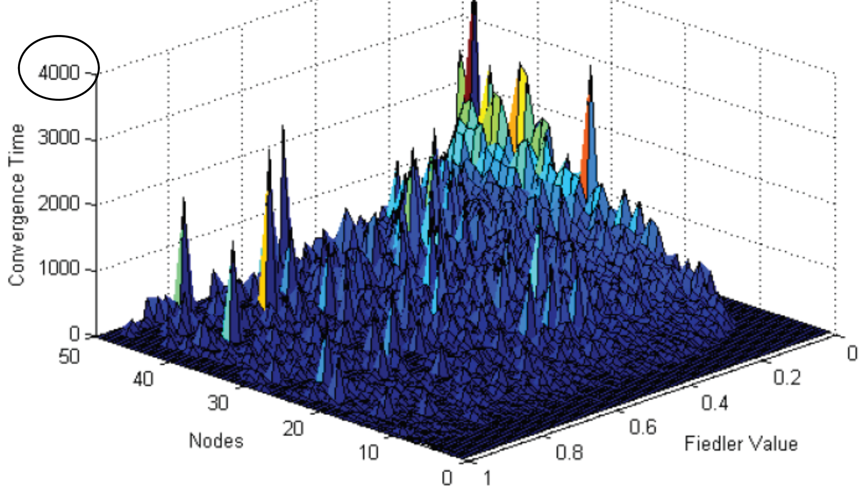

Experiment Environment

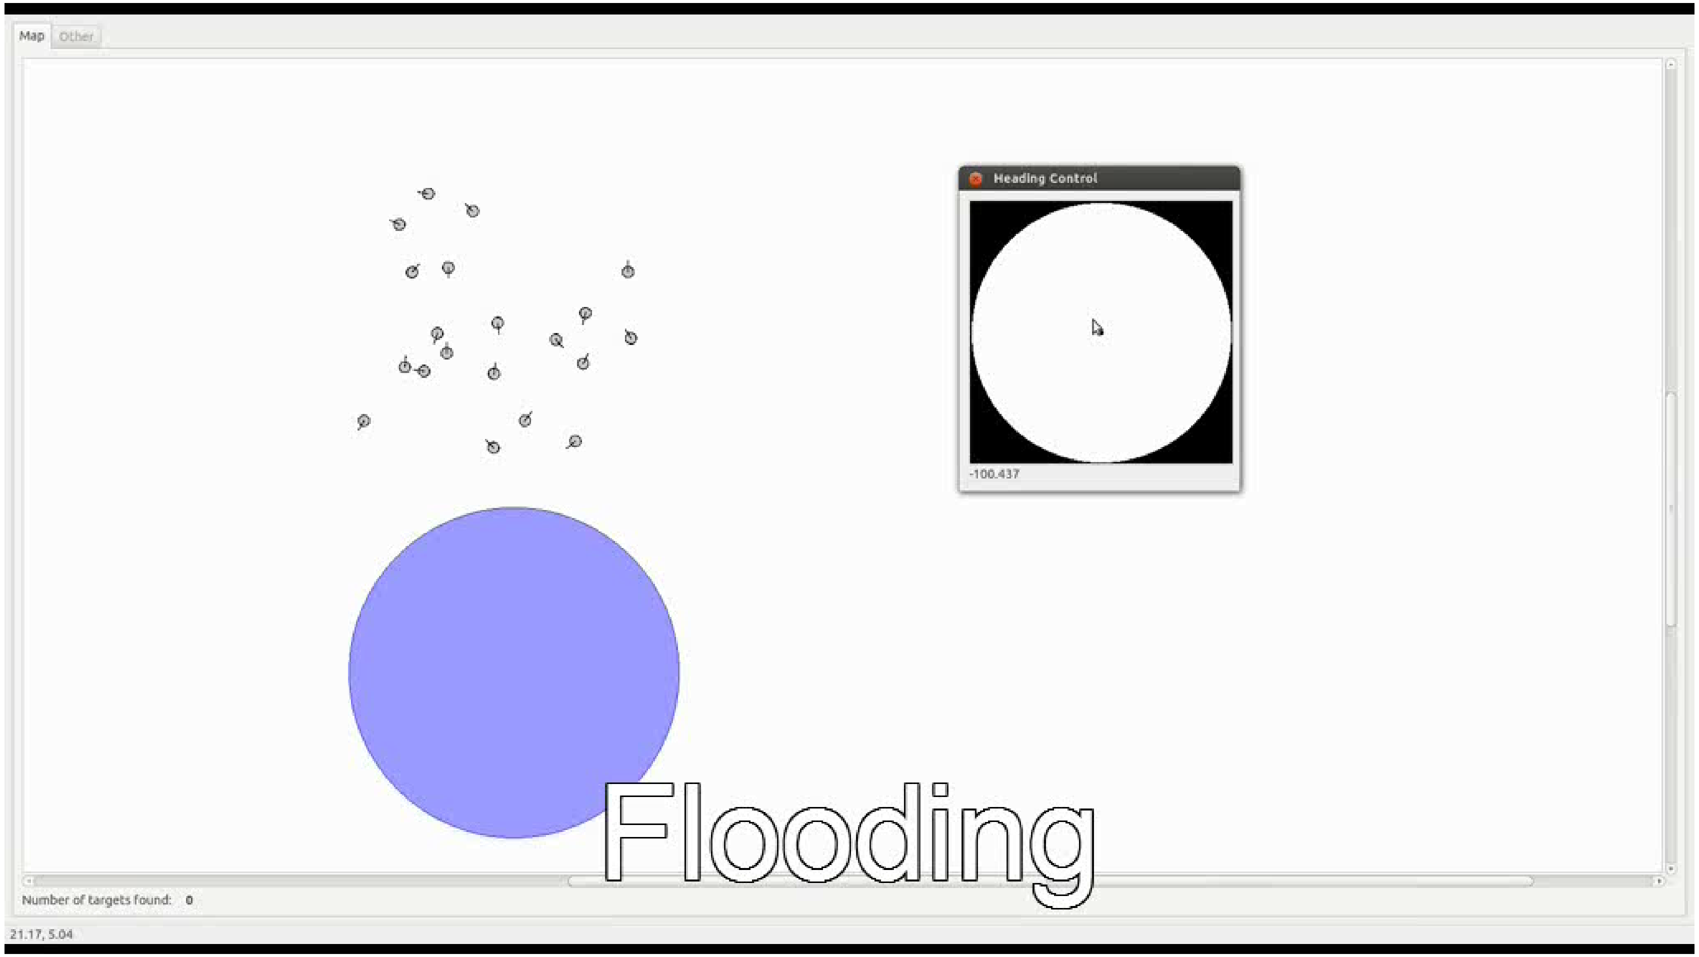

Goals Reached

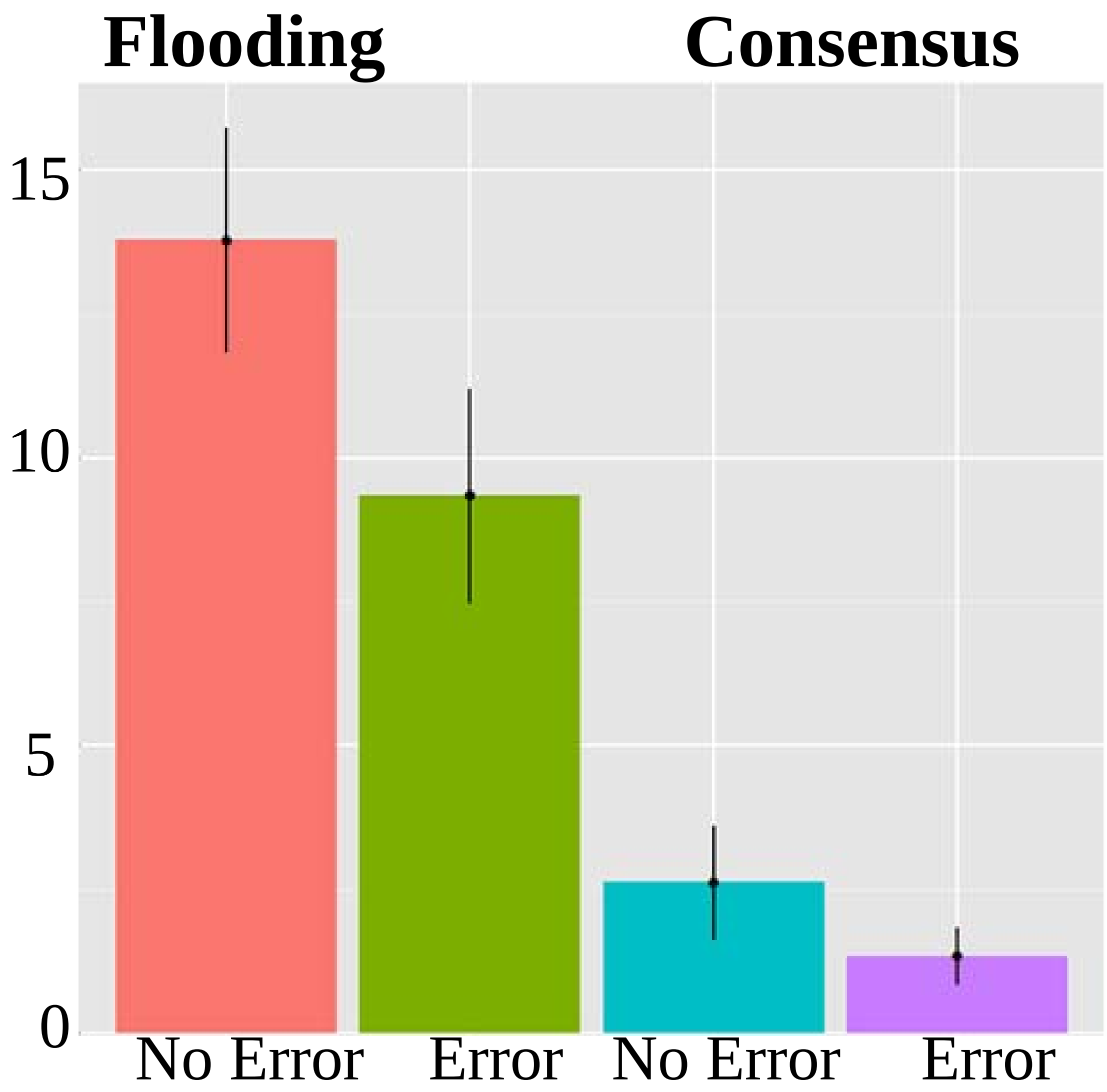

Swarm Diameter

Diameter is significantly larger for Flooding with Error and smaller for Consensus without error

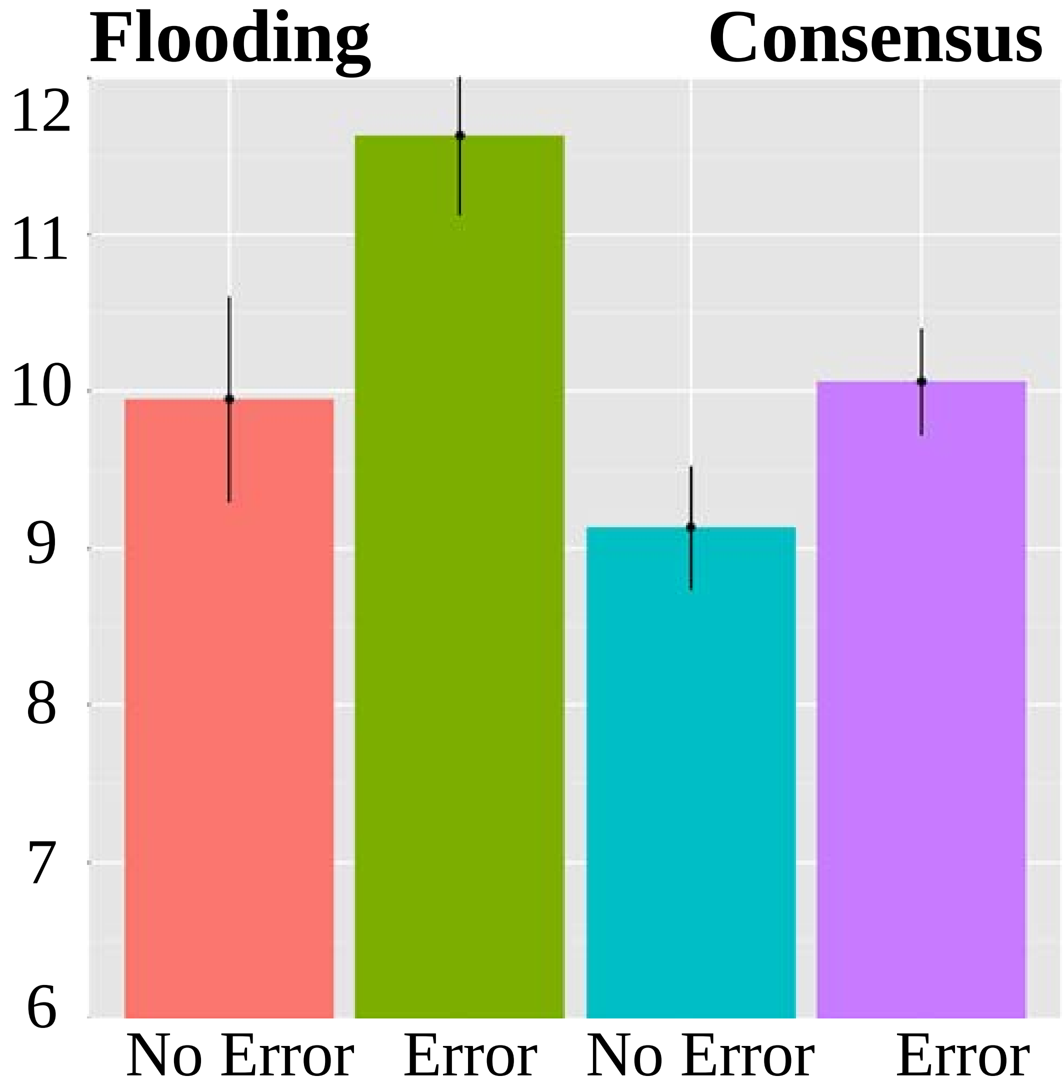

Advantages/Disadvantages

Tacit Leader(s)

- **Closer adjustment to local conditions**
- **Error tolerance**

Explicit Leadership

- **More accurate alignment with intent**
- **More rapid convergence**
- **May loose resilience associated with distributed coordination**

Neglect Benevolence-Timing of control input

Nagavalli, S., Luo, L., Chakraborty, N., Sycara, K., Neglect Benevolence in Human Control of Robotic Swarms, International Conference on Robotics and Automation (ICRA), Hong Kong, China, May 31-June 7, 2014

Neglect Benevolence in Human Control of Robotic Swarms

- **Challenge:** Very often, human input harms the performance of autonomously coordinating multi robot swarms

- **Goal:** Determine the optimal time that human control input should be given to a dynamically evolving multi-robot swarm to improve mission performance.

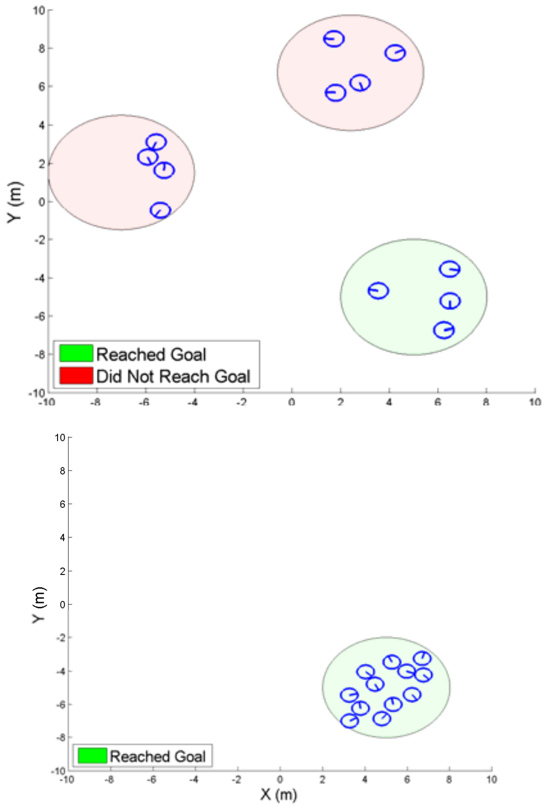

(Top) When the human input is applied immediately as the new goal arises, the swarm splits and only 4 robots reach the goal (green region). (Bottom) When the human input is delayed all robots reached the goal (green region).

Motivating Example

Early Input **Appropriate Input** **Late Input**

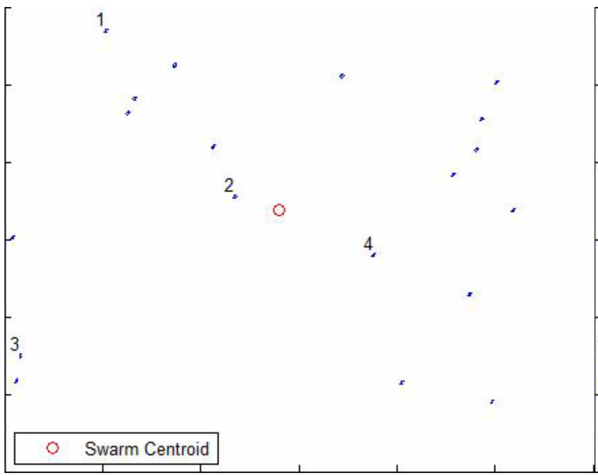

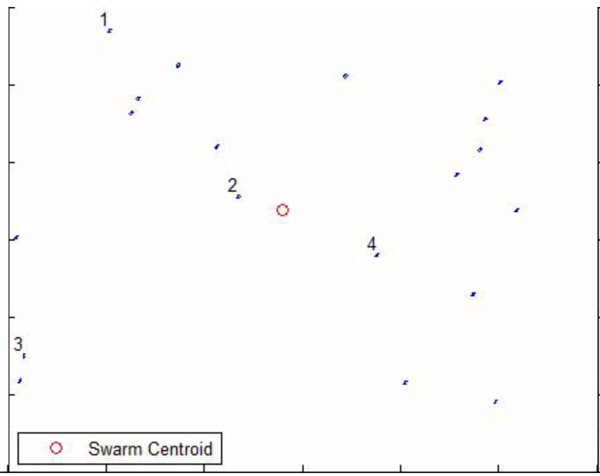

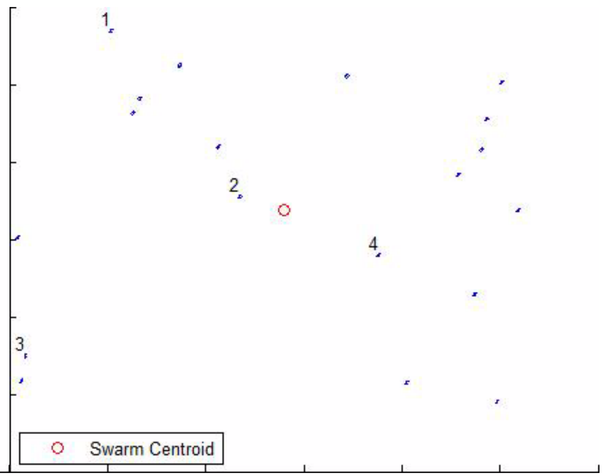

Motivating Examples

Formalization – Neglect Benevolence

- Given
  - System dynamics (including automatic controllers)
  - Performance criterion (e.g. time-to-goal)
  - One human input : U
  - Desired goal states: X
- There may be a subset of states, Xsub where applying the input will decrease system performance
- Each state in Xsub is **Neglect Benevolent**
- If Xsub is not null, the system exhibits **Neglect Benevolence**

Illustration of Neglect Benevolence in State Space

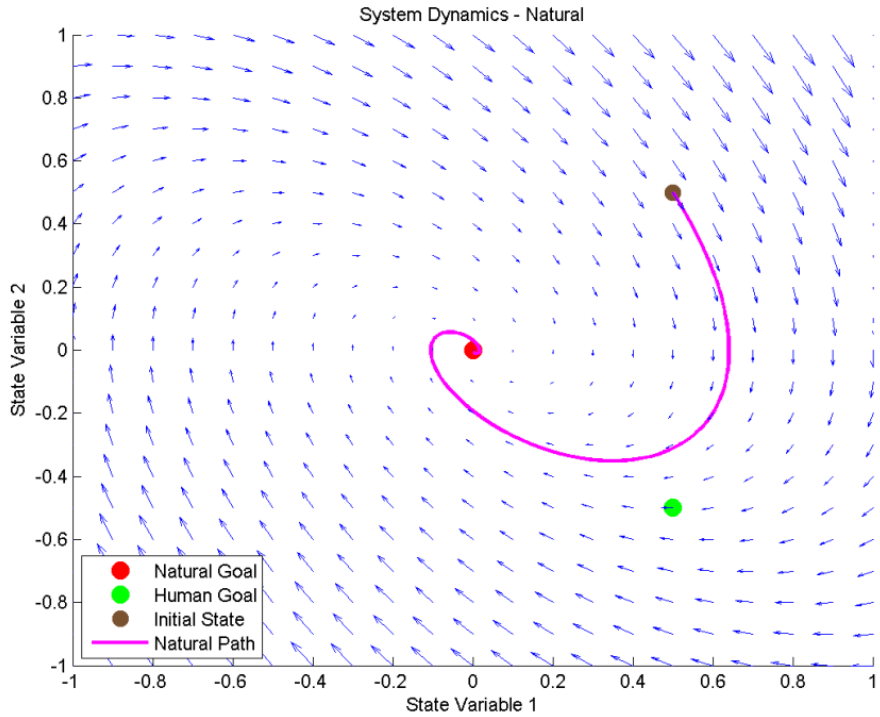

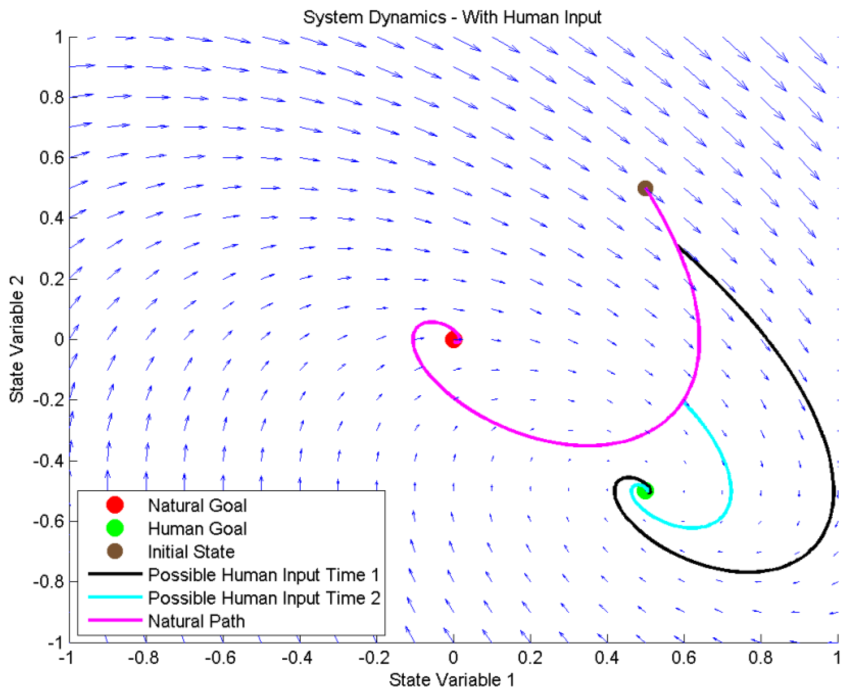

NB in Formation Control of Robotic Swarms

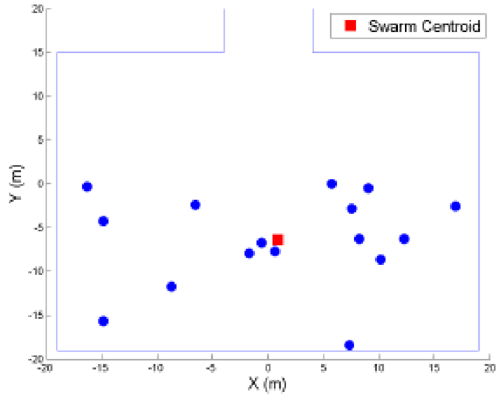

(a) Initial Positions of the Robots

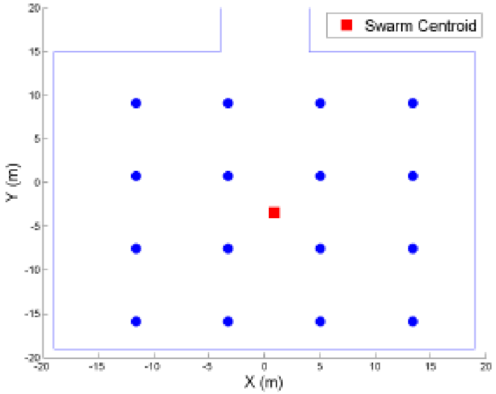

(b) Robot Formation Achieved when No Human Input is Applied

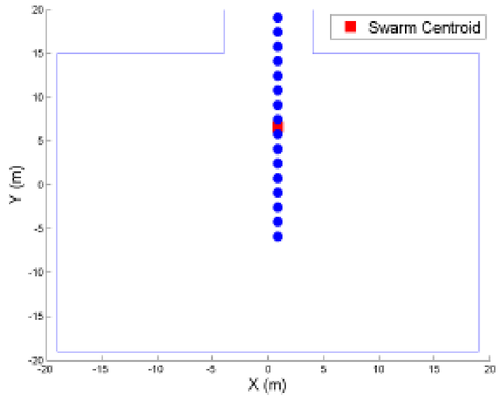

(c) Robot Formation Desired by the Human and Generated when Human Input is Applied

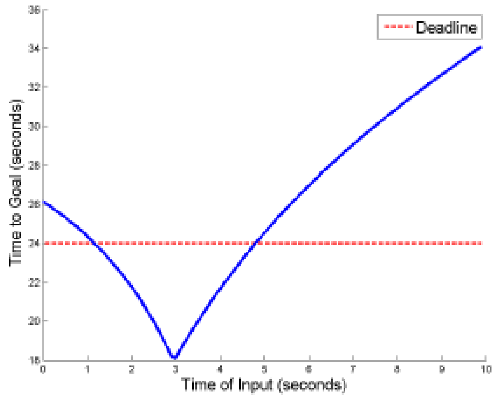

(d) Time at which the Desired Formation is Achieved vs. Time at which Human Input is Applied

Experiments – NB in Formation Control of Robotic Swarms

Input Applied Too Early (10.0 Seconds)

Input Applied at Optimal Time (20.6 Seconds)

Robots in a Room

Neglect Benevolence

- **Introduced a formal notion of Neglect Benevolence, i.e. delaying human control input, as mission goals change, may benefit system performance**
- **Proved that all stable Linear Time Invariant systems exhibit Neglect Benevolence**
- **Designed an algorithm for computing the optimal input time**
- **Applied Neglect Benevolence analysis to formation control of robotic swarms**
- **Demonstrated that some deadlines can only be achieved if the human input to the robotic swarm is delayed**

Carnegie Mellon

HSI Reference Task

THE ROBOTICS INSTITUTE

- **Neglect Benevolence leads to a formulation of human control of swarms as:** *Diverting a swarm's trajectory through state space to some new desired trajectory*

- **Swarm is evolving towards its natural goal (original configuration)**
- **Operator desires a new goal (different configuration)**
- **Operator supplies input to divert swarm to new goal**
- **Can humans choose the optimal time?**
- **Human performance is compared with optimal performance (timing of input for NB task)**

Intelligibility of Swarm Behavior

- **Ability to recognize swarm behaviors depends upon human perception of regularities (Gestalt principle of common fate) in interactions between swarm members**
  - **Heading, velocity, proximity, etc.**

- **Because swarm behavior such as switching between formations provides few of the Gestalt cues people use to recognize coherent behaviors**
- **Could we augment the display to help them recognize the closest point between trajectories?**

Applying Consensus to Generate Static Formations

Carnegie Mellon

THE ROBOTICS INSTITUTE

Circle | **S-Formation**

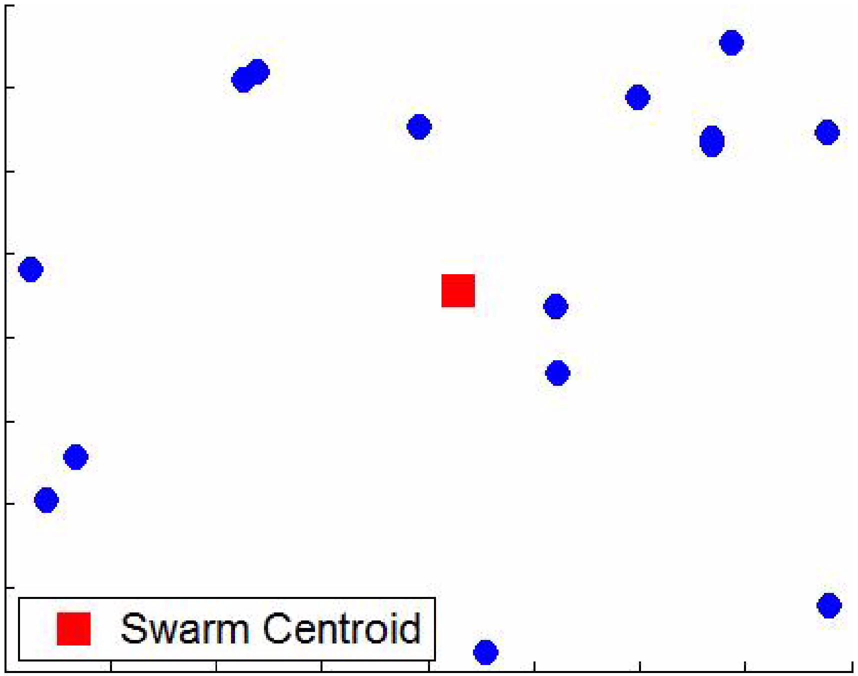

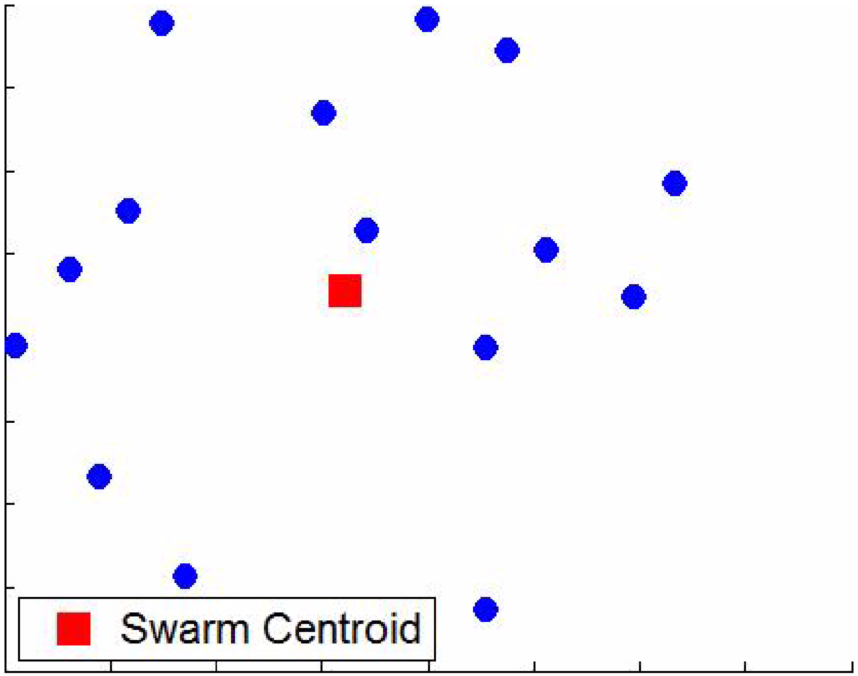

Consensus-Based Control

Visualizing Swarm Behavior

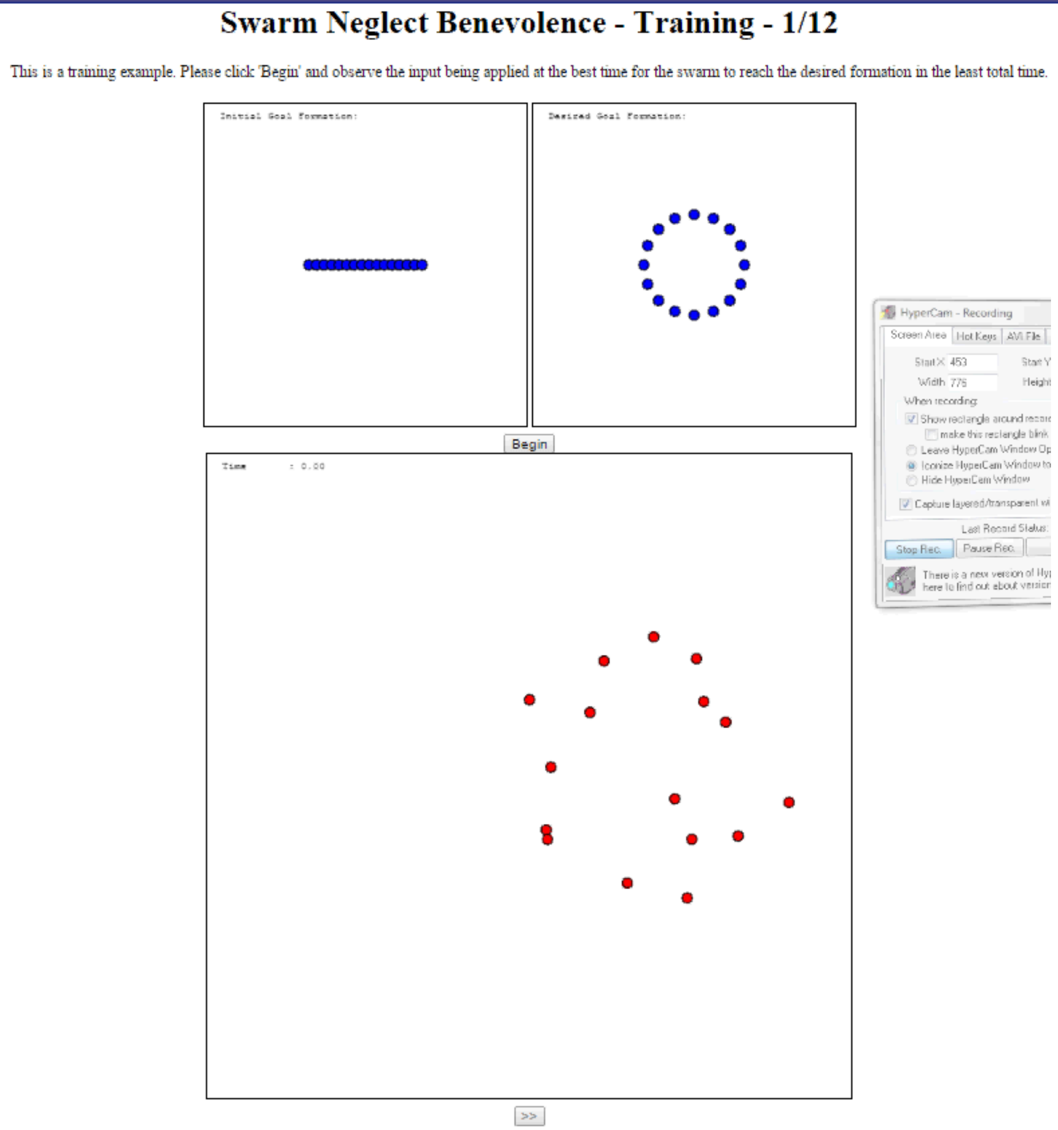

Visualizing Swarm Behavior

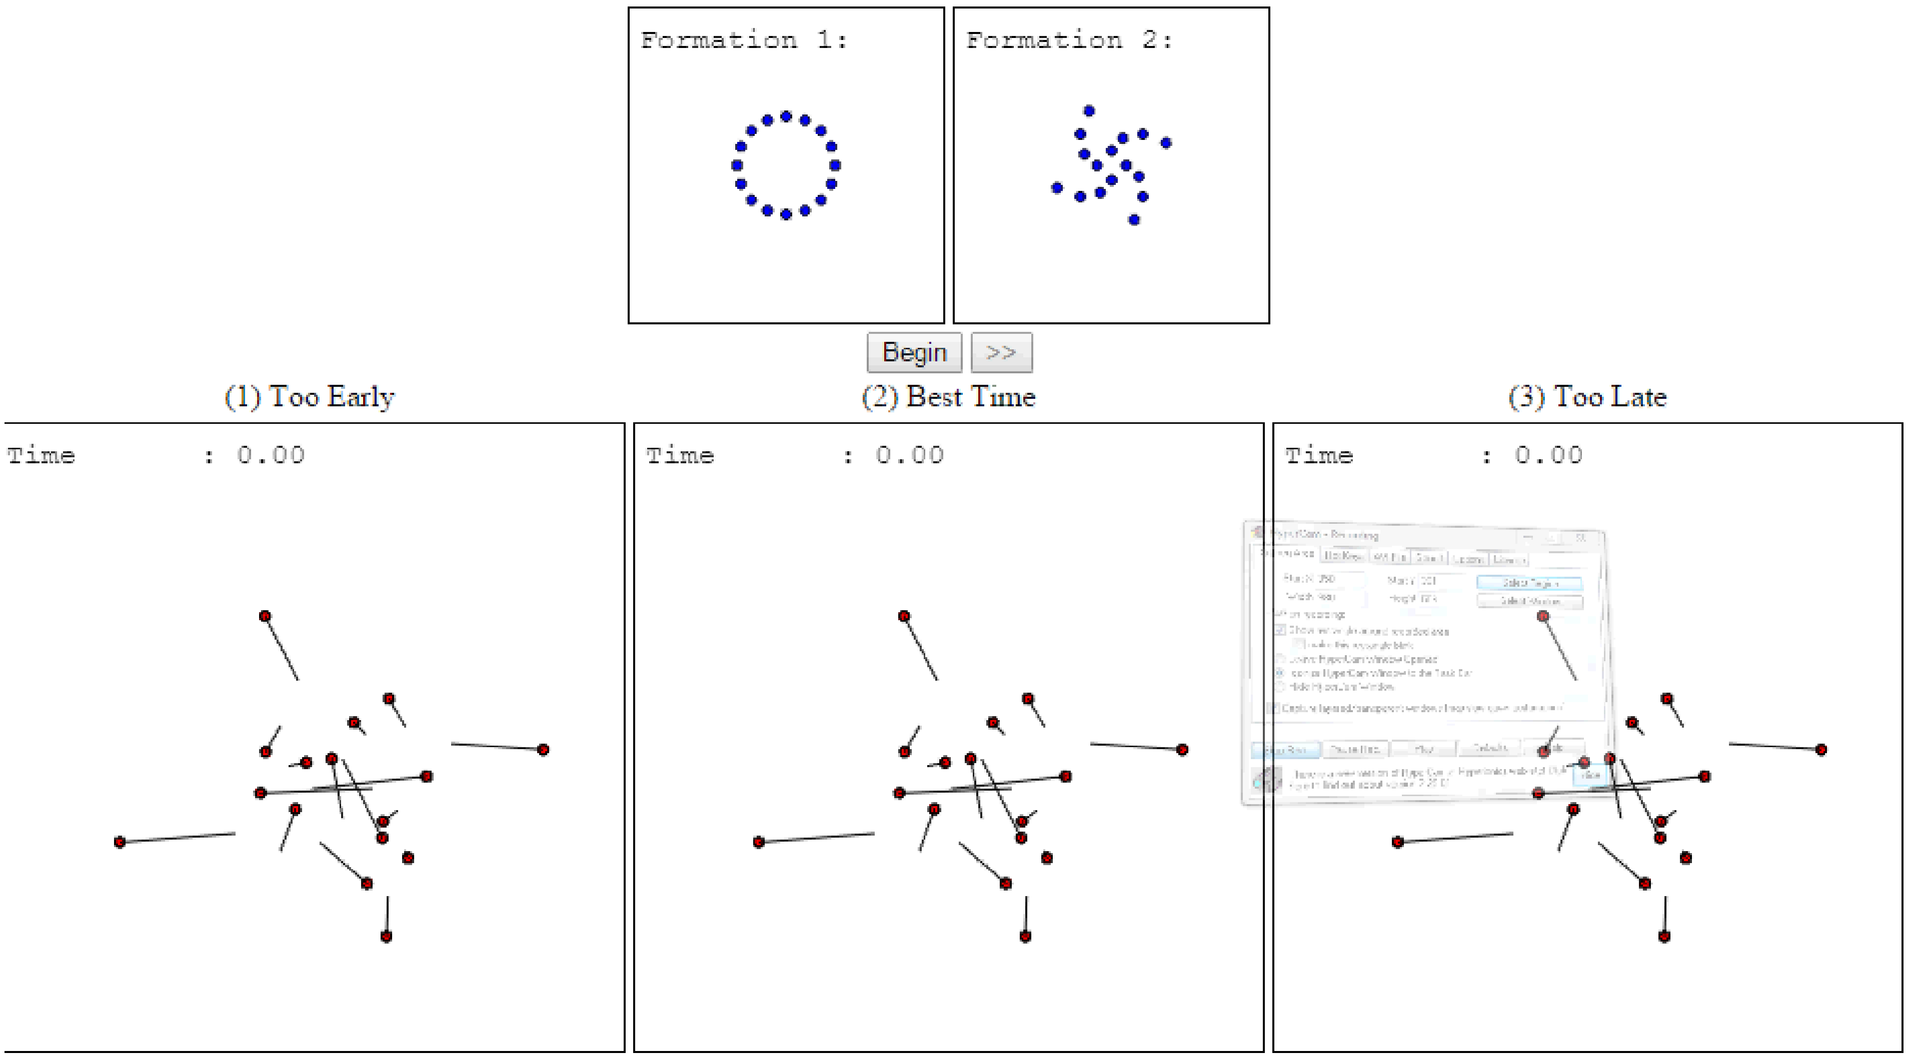

Deviation from Optimal

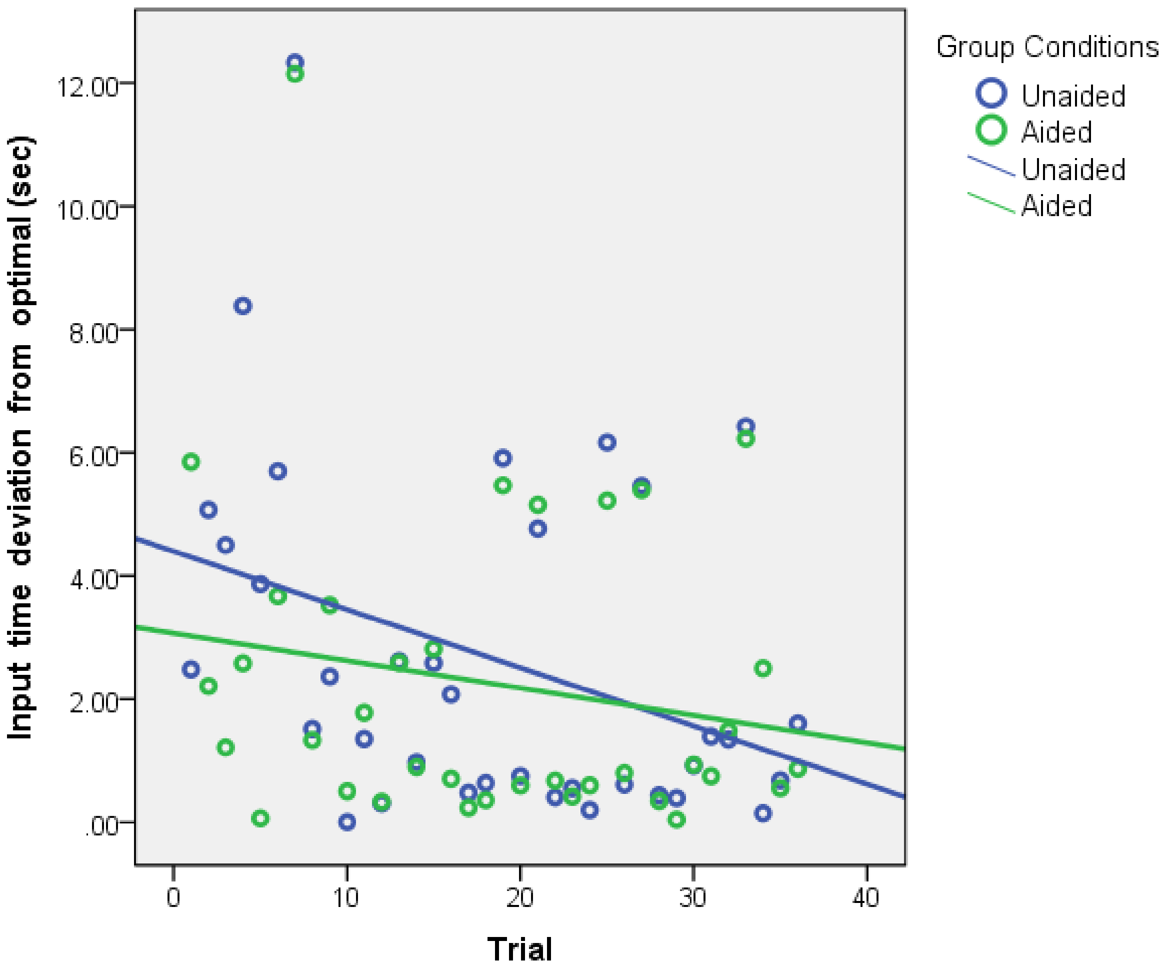

Nagavalli, S., Chien, S., Lewis, M., Chakraborty, N., & Sycara, K. (2015). Bounds of Neglect Benevolence in Input Timing for Human Interaction with Robotic Swarms. Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI 15).

Conclusions

Given our overall goal of making human supervisory control for multiple robots more effective, we

- **Presented a taxonomy of human supervisory multi-robot control tasks based on cognitive complexity**
- **For independent robots, we showed that schemes for scheduling operator attention are effective. This requires:**
  - **higher level of robot automation, e.g self-reflection**
  - **additional information in the scheduling scheme (e.g. spatial constraints)**
- **Developed scalable displays**
- **Presented tradeoffs in swarm control via leaders**
- **Neglect Benevolence concept for optimal human timing**

Gesture Control of Swarm

Future Work

- Leader-Based Swarm Architectures in communication-challenged environments
  - Control leaders
  - Information leaders
  - Knowledge leaders
- Human Swarm Trust
  - Automation vs Autonomy
  - Swarm increases its reliability
  - Self-monitoring and self repair

Carnegie Mellon

Recent Relevant Publications

THE ROBOTICS INSTITUTE

- **Kolling, A., Walker, P., Chakraborty, N., Sycara, K. Lewis, M. Human Interaction with Robot Swarms: A Survey, IEEE Transactions on Systems, Man and Cybernetics, Vol 46, No. 1, February 2016 DOI: 10.1109/THMS.2015.2480801,**

- **Xu, Y., Scheller-Wolf, A., Sycara, K. "The Benefit of Introducing Variability in Quality Based Service Domains",** ***Operations Research,*** **Vol 63, Issue 1, January 2015 Permalink: http://dx.doi.org/10.1287/opre.2014.1330**

- **Nagavalli, S., Chien, S., Lewis, M., Chakraborty, N., & Sycara, K. (2015). Bounds of Neglect Benevolence in Input Timing for Human Interaction with Robotic Swarms. Proceedings of the 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI 15).**

- **Walker, P. Amirpour S. Chakraborty, N., Lewis M., Sycara, K. Control of Swarms with Multiple Leader Agents, International Conference on Systems, Man and Cybernetics, San Diego CA, October 5-8, 2014.**

- **Luo, R., Chakraborty, N., Sycara, K. Supervisory Control for Cost-Effective Redistribution of Robotic Swarms, International Conference on Systems, Man and Cybernetics, San Diego CA, October 5-8, 2014. (*Finalist for Best Student Paper Award).***

- **Nagavalli S., Lybarger A., Luo, L., Chakraborty, N., Sycara, K. Aligning Coordinate Frames in Multi-Robot Systems with Relative Sensing Information, International Conference on Intelligent Robots and Systems (IROS), Chicago, September 14-18, 2014**

- **Walker, P., Amirpour S., Chakraborty, N., Lewis M., Sycara, K. Human Control of Robot Swarms with Dynamic Leaders, International Conference on Intelligent Robots and Systems (IROS), Chicago, September 14-18, 2014**

Carnegie Mellon

Recent Relevant Publications

THE ROBOTICS INSTITUTE

- Chien, S.Y., Semnani-Azad, Z., Lewis, M. Sycara, K. Towards the development of an Inter-Cultural Scale to Measure Trust in Automation, 6th International Conference on Cross Cultural Design part of Human Computer Interaction International (HCII-2014), Crete, Greece, June 22-27, 2014 (*Best Paper Award*)

- Nagavalli, S., Luo, L., Chakraborty, N., Sycara, K., Neglect Benevolence in Human Control of Robotic Swarms, International Conference on Robotics and Automation (ICRA), Hong Kong, China, May 31-June 7, 2014.

- Amirpour S., Walker, P., Lewis, M., Chakraborty, N., Sycara, K. Explicit vs Tacit Leadership in Influencing the Behavior of Swarms, International Conference on Robotics and Automation (ICRA), Hong Kong, China, May 31-June 7, 2014.

- Nunnally, S., Walker, P., Chakraborty, N., Lewis, M. Sycara, K Using Haptic Feedback in Human Robotic Swarm Interaction, HFES, , San Diego, September 30-October 4 2013

- Walker, P, Nunnally, S., Chakraborty, N., Lewis, M, Sycara, K. Levels of Automation for Human Influence of Robot Swarms, HFES, San Diego, September 30-October 4 2013. (*Best Paper Award*).

- Nunnally, S., Walker, P., Chakraborty, N., Lewis, M. Sycara, K Using Coverage for Measuring the Effect of Haptic Feedback in Human Robot Swarm Interaction, In Proceedings of the Conference on System Man and Cybernetics, Manchester, UK., October 13-16 2013.

- Walker, P., Amirpour S., Lewis, M., Chakraborty, N., Sycara, K. Human Control of Leader-Based Swarms Proceedings of the Conference on System Man and Cybernetics, Manchester, UK., October 13-16 2013.

- A. Kolling, S. Nunnally, M. Lewis, K. Sycara, “Towards Human Control of Robot Swarms”, ACM/IEEE International Conference on Human-Robot Interaction, 2012.

Carnegie Mellon

Recent Relevant Publications

- Chien, S., Mehrotra, S., Brooks, N., Sycara, K., & Lewis, M. Effects of Alarms on Control of Robot Teams, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11)
- Dai, T, Sycara, K., Lewis, M. A game theoretic queuing approach to self-assessment in human-robot interaction systems. IEEE International Conference on Robotics and Automation (ICRA 2011), May 9-13, Shanghai, China, 2011.
- Lee, P., Kolling, A., & Lewis, M. Workload Modeling using Time Windows and Utilization in an Air Traffic Control Task, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11).
- Brooks, N., Wang, H., Chien, S., Lewis, M., Scerri, P.,& Sycara, K. (2011) Asynchronous Control with ATR for Large Robot Teams, Proceedings of the 55th Annual Meeting Human Factors and Ergonomics Society (HFES'11).
- Abedin, S., Wang, H., Lee, P., Lewis, M., Brooks, N., Owens, S., Scerri, P. and Sycara, K. SUAVE: Integrating UAV Video Using a 3D Model IEEE International Conference on Systems, Man, and Cybernetics, (SMC'11)
- Wang, H., Kolling, A., Abedin, S., Lee, P., Chien, S., Lewis, M., Brooks, N., Owens, S., Scerri, P. & Sycara, K. (2011) Scalable target detection for large robot teams, Proceedings of the 6th ACM/IEEE International Conference on Human-Robot Interaction.
- Lewis, M., Wang, H., Chien, S., Velagapudi, P., Scerri, P. & Sycara, K. (2010). Choosing autonomy modes for multirobot search, Human Factors, doi:10.1177/0018720810366859
- Xu, Y, Dai, T., Sycara, K. & Lewis, M. (2010). Service level differentiation in multi-robots control, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), 2224-223
- Chien, S., Wang, H, & Lewis, M. (2010). Human vs. algorithmic path planning for search and rescue by robot teams, Proceedings of the 54th Annual Meeting of the Human Factors and Ergonomics Society (HFES'10), September 27- October 1.

Carnegie Mellon

Recent Relevant Publications

THE ROBOTICS INSTITUTE

- **Wang, H., Lewis, M., Chien, S., Scerri, P., Velagapudi, P., Sycara, K. & Kane, B. (2010). Teams organization and performance in multi-human/multi-robot teams, 2010 IEEE International Conference on Systems, Man, and Cybernetics, (SMC'10), October 10-13.**
- **Scerri, P., Velagapudi, P., Sycara, K., Wang, H., Chien, S. & Lewis, M. (2010). Towards an understanding of the impact of autonomous path planning on victim search in USAR, Proceedings of the 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'10), October 18-22.**
- **Scerri, P., Owens, S., Sycara, K. & Lewis, M. (2010). User evaluation of a GUI for controlling an autonomous persistent surveillance team, In SPIE'10.**