ALIDADE
INCORPORATED

Decentralized Command and Control: Self-Organization in a Simple Model for Emergency Response

Dr. Michael Bell

14th International Command and Control Research and Technology Symposium

16 June 2009

Introduction

- Military organizations traditionally require *unity of command:*

  "All forces operate under a single commander with the requisite authority to direct all forces employed in pursuit of a common purpose"*

- In operations requiring cooperation with or support from other agencies, the private sector, or foreign nations ("complex endeavors"), unity of command may not be possible
- In such cases, doctrine focuses on *unity of effort:*

  "coordination and cooperation toward common objectives, even if participants are not necessarily part of the same command or organization"*

*US DoD, Joint Publication 3-0 (2008)

Civilian Agencies

- The National Incident Management System (NIMS) provides for a *unified command:*

  "agencies with different legal, geographic, and functional authorities and responsibilities… work together effectively without affecting individual agency authority, responsibility, or accountability"*

- NIMS and the National Response Framework

  "are designed to ensure that local jurisdictions retain command, control, and authority over response activities for their jurisdictional areas"*

*DHS, National Incident Management System (2008)

Possible Solution

- A network-centric approach:
  1. A robustly networked force improves information sharing.
  2. Information sharing and collaboration enhance the quality of information and shared situational awareness.
  3. *Shared situational awareness enables self-synchronization.* [emphasis added]
  4. These, in turn, dramatically increase mission effectiveness.

– D.S. Alberts, “Information Age Transformation,” 1996

- Self-synchronization (temporal) + self-assembly (spatial) = self-organization (complex systems)
- Is it possible to “self-organize” an operation by insuring common intent and purpose and shared situational awareness?

Incident Response Model

- Scenario
  - An event has has caused a number of simultaneous incidents, randomly distributed over a geographical area *(e.g.,* storm-related power outages)
  - A force of first-responders (*e.g.,* utility company service trucks) is available, initially distributed randomly across the district
  - Incidents and responders are identical; service time is negligible compared to transit time
- Shared awareness
  - Each responder has timely information on the location of all unresolved incidents
- Common intent
  - Service all incidents in the shortest possible time
- Decentralized command and control
  - No central planning or command; no communication between responders
- Concept of operations
  - Each responder deals with the nearest unresolved incident

Technical Note

- For one responder, this is just the Traveling Salesman problem (TSP)
- TSP is “hard” (NP-complete)
- Optimization is impractical for more than a few incidents
- The proposed “greedy” algorithm (heuristic) is
  - Easy to implement
  - Known to produce the worst possible result for certain cases
- For multiple responders, less is known
  - Related to the vehicle scheduling problem (VSP)
  - Hard for a central planner (not known to be NP-C)
  - Individual responders cannot optimize without knowledge of other responders

Agent-Based Simulation

QuickTime™ and a
decompressor
are needed to see this picture.

- Environment: NetLogo
- Area = 33 x 33 grid
- 1089 potential incident locations
- 50 incidents (density = 50/1089 = 0.046)
- Responder speed = 1 grid site/time step
- Zero time required to service an incident
- Experiment: 1000 replications with random initial conditions

Completion Time

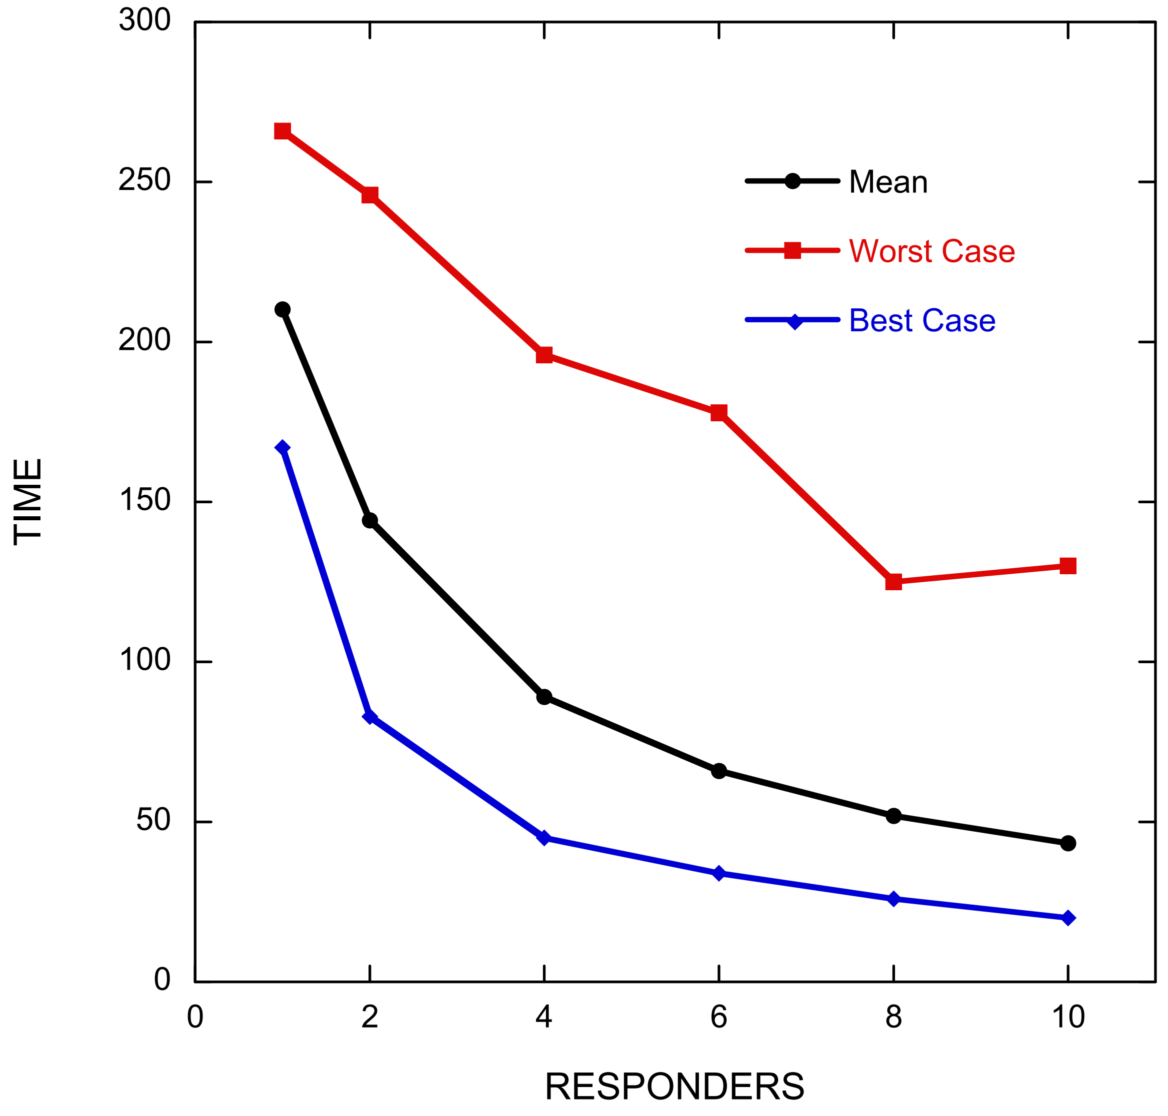

Effectiveness

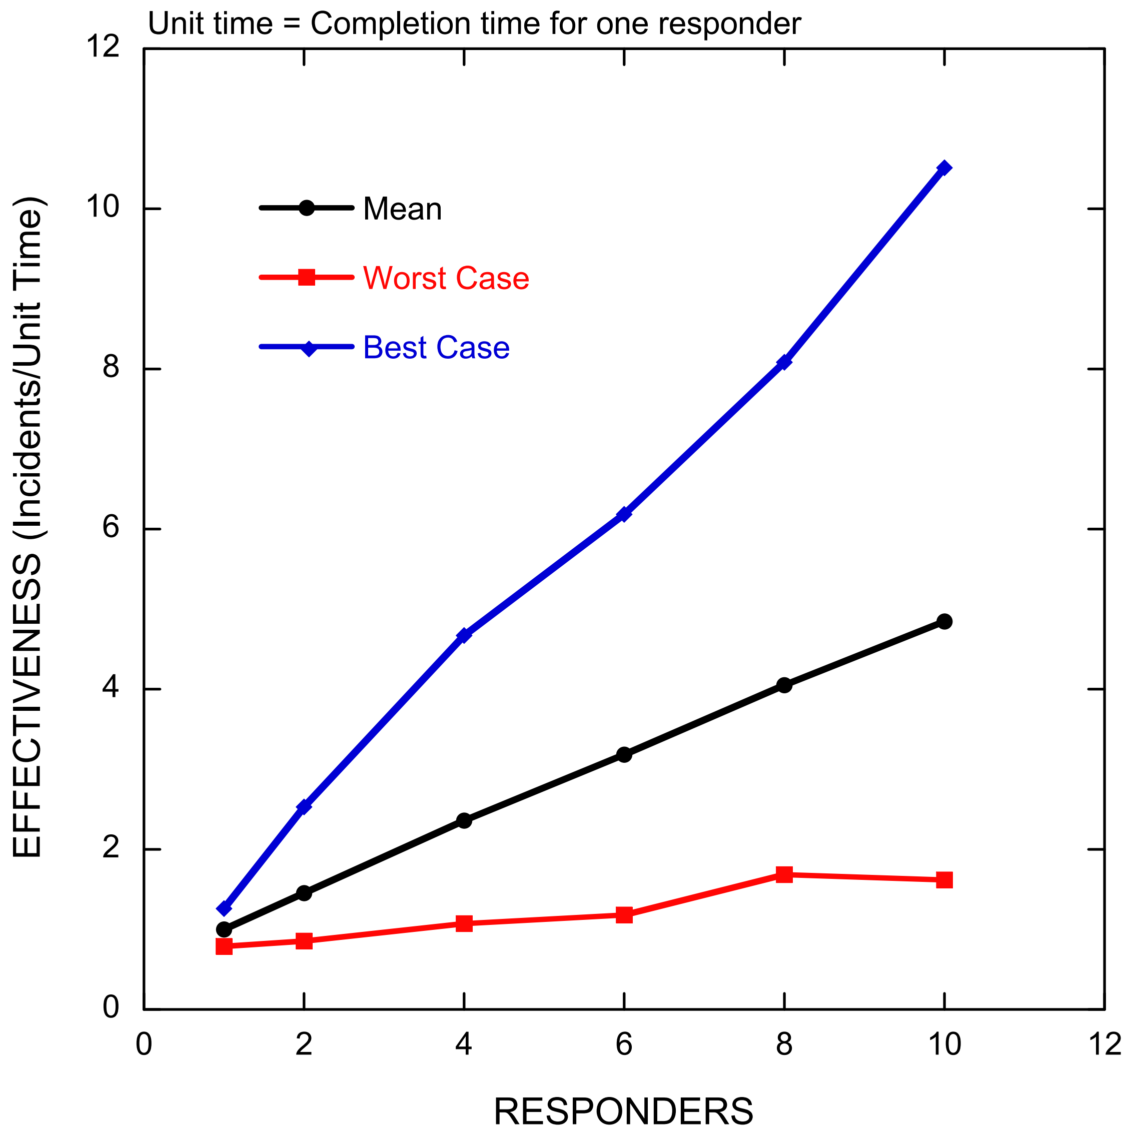

Marginal Effectiveness

- (Mean effectiveness) vs. (number of responders) has slope of roughly 0.47
- Bad news:
  - 10 responders have about 4.9 times the effectiveness of one
  - 53% of the effort of additional responders is wasted
- Good news:
  - Best case (minimum time) shows improvement of roughly 10x for 10 responders
  - Constant marginal effectiveness implies no evidence of diminishing returns

Completion Time Distributions

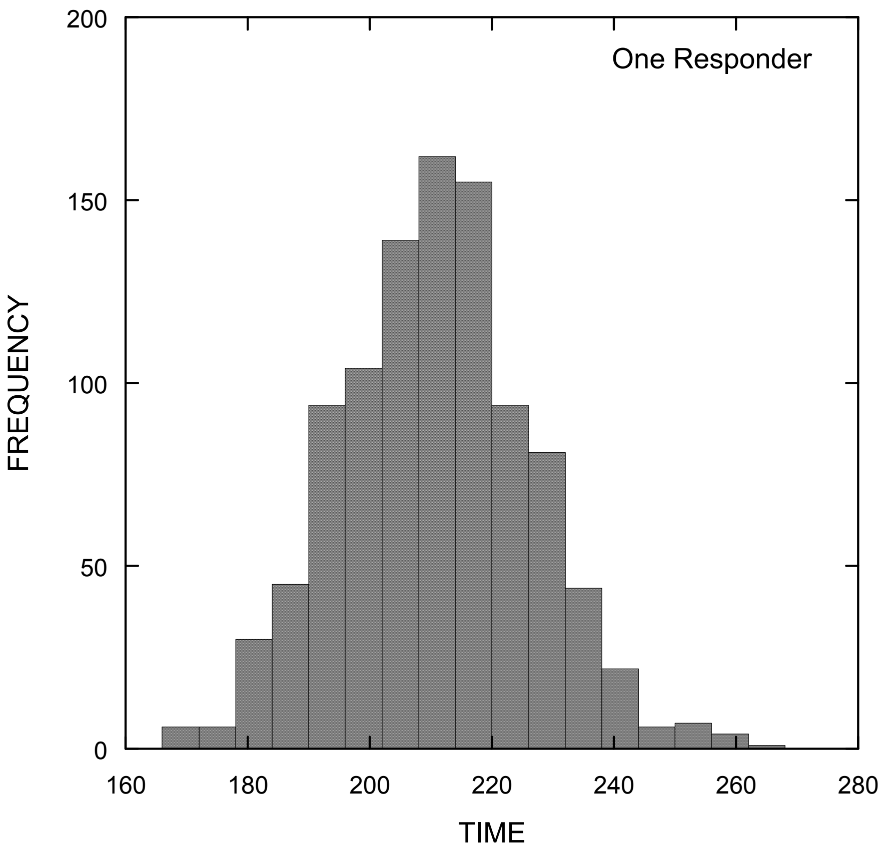

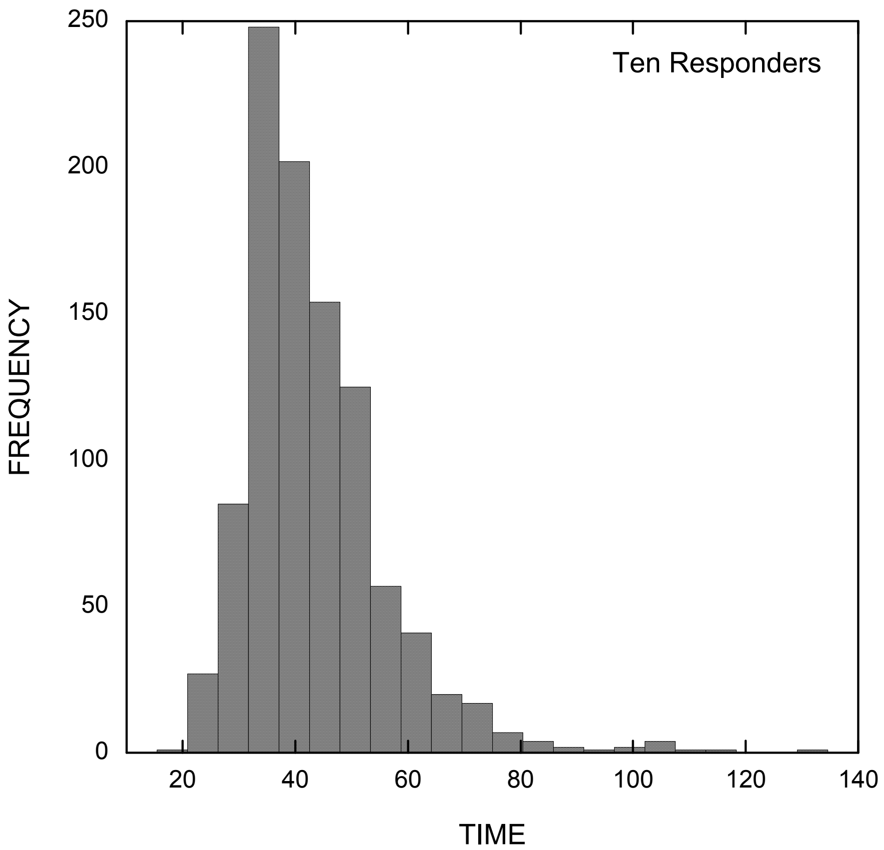

Dysfunctional Self-Organization

- At some point many (even all) responders form a tight cluster that travels together with members competing for the same nearest incident
- Because:
  - Responders that choose the same goal approach one another
  - The first responder to reach the incident deals with it, but the other responders are now closer than when they started
  - This makes it more likely that they will again choose the same goal
  - Eventually, groups of responders travel together, reducing effectiveness

Separation

Two Responders, 50 Incidents

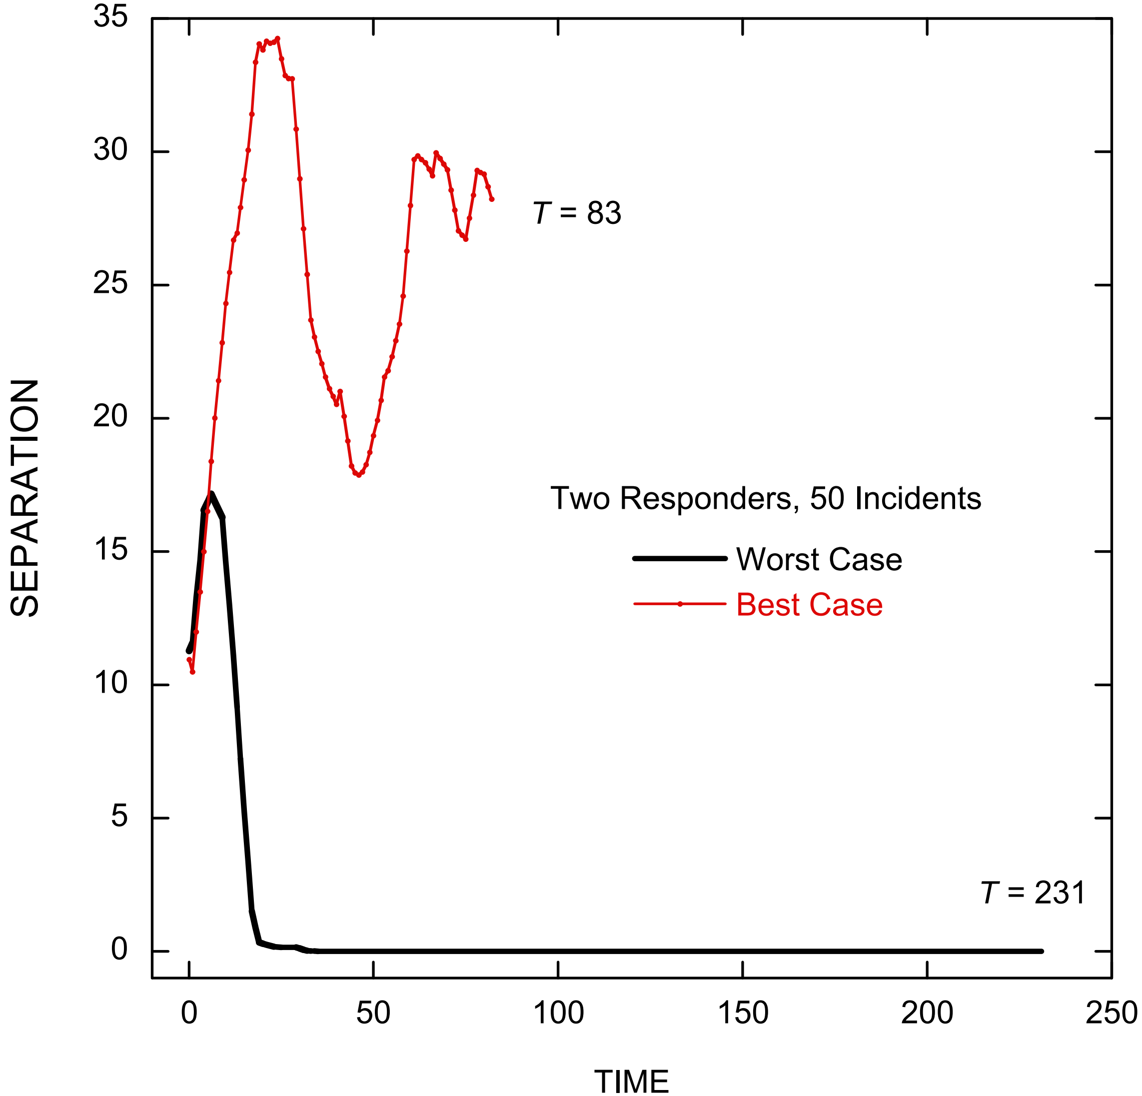

Avoidance Rule

If (current goal drops off incident list) then
(set goal as *second* nearest incident)

- Logic: “break ties” by giving the first responder on the scene priority to proceed to the nearest incident

- Improvement
  - Dramatic reduction in extremely long-time cases
  - Increased symmetry of the time distribution

Results of Avoidance Rule

(10 responders)

Without Avoidance

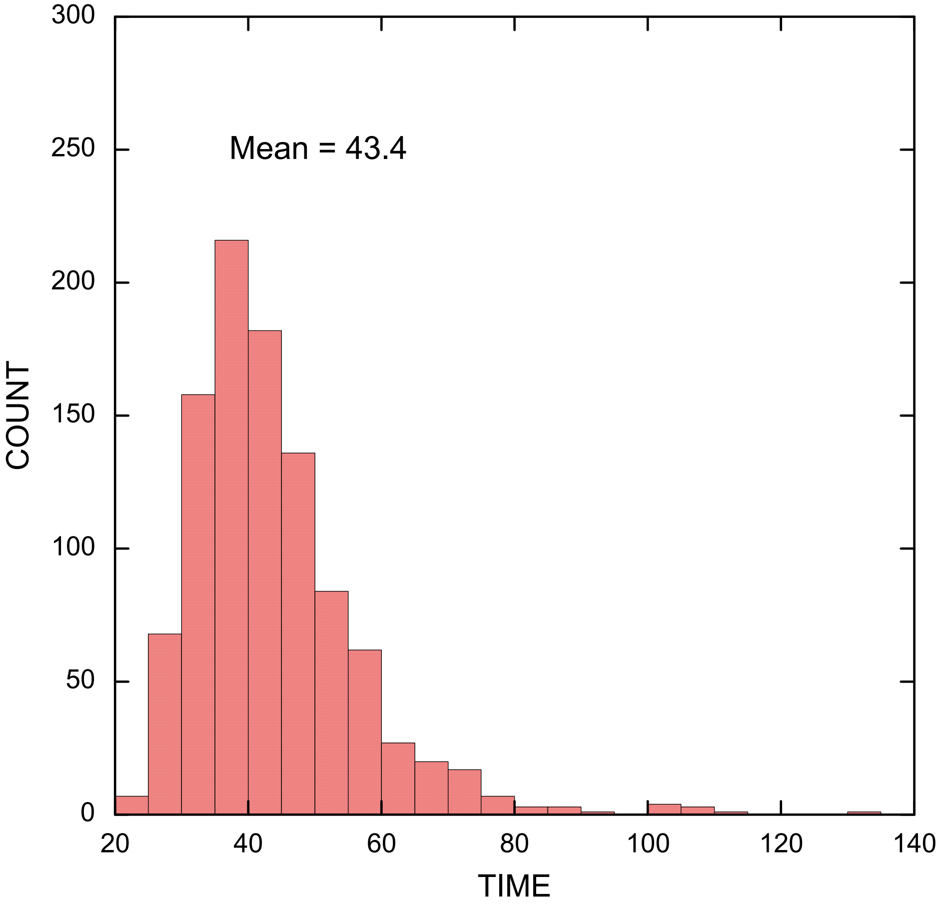

With Avoidance

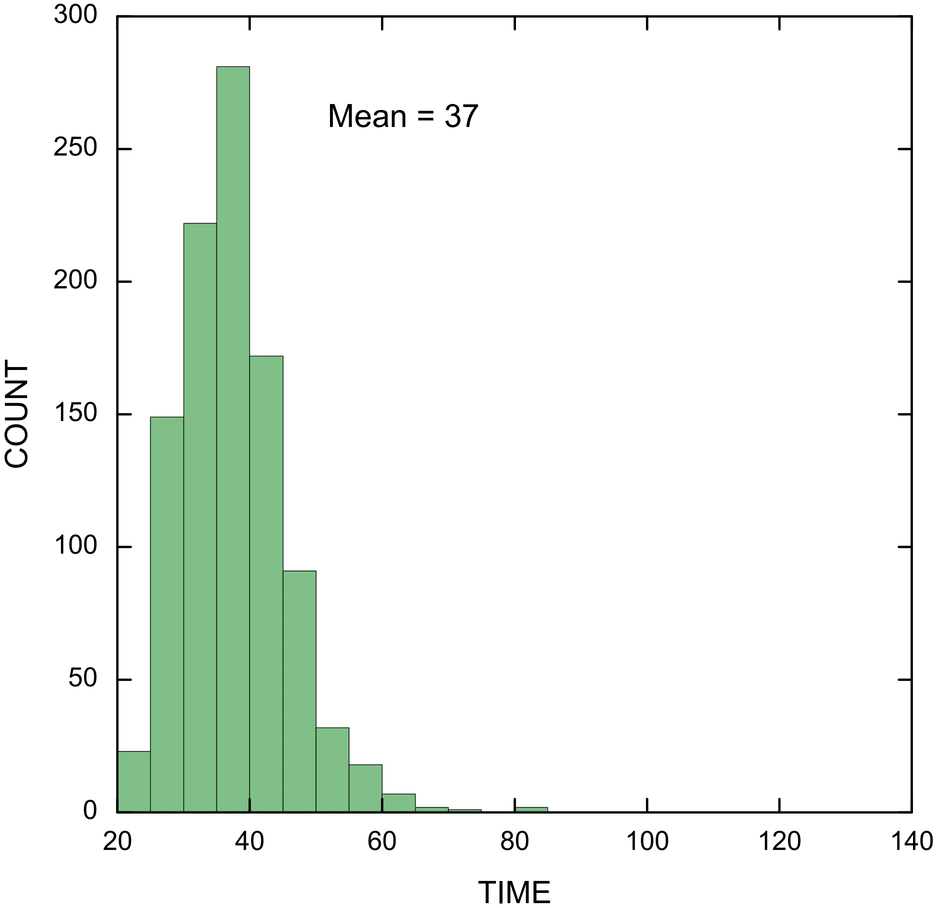

Separation

Two Responders, 50 Incidents, Worst Case

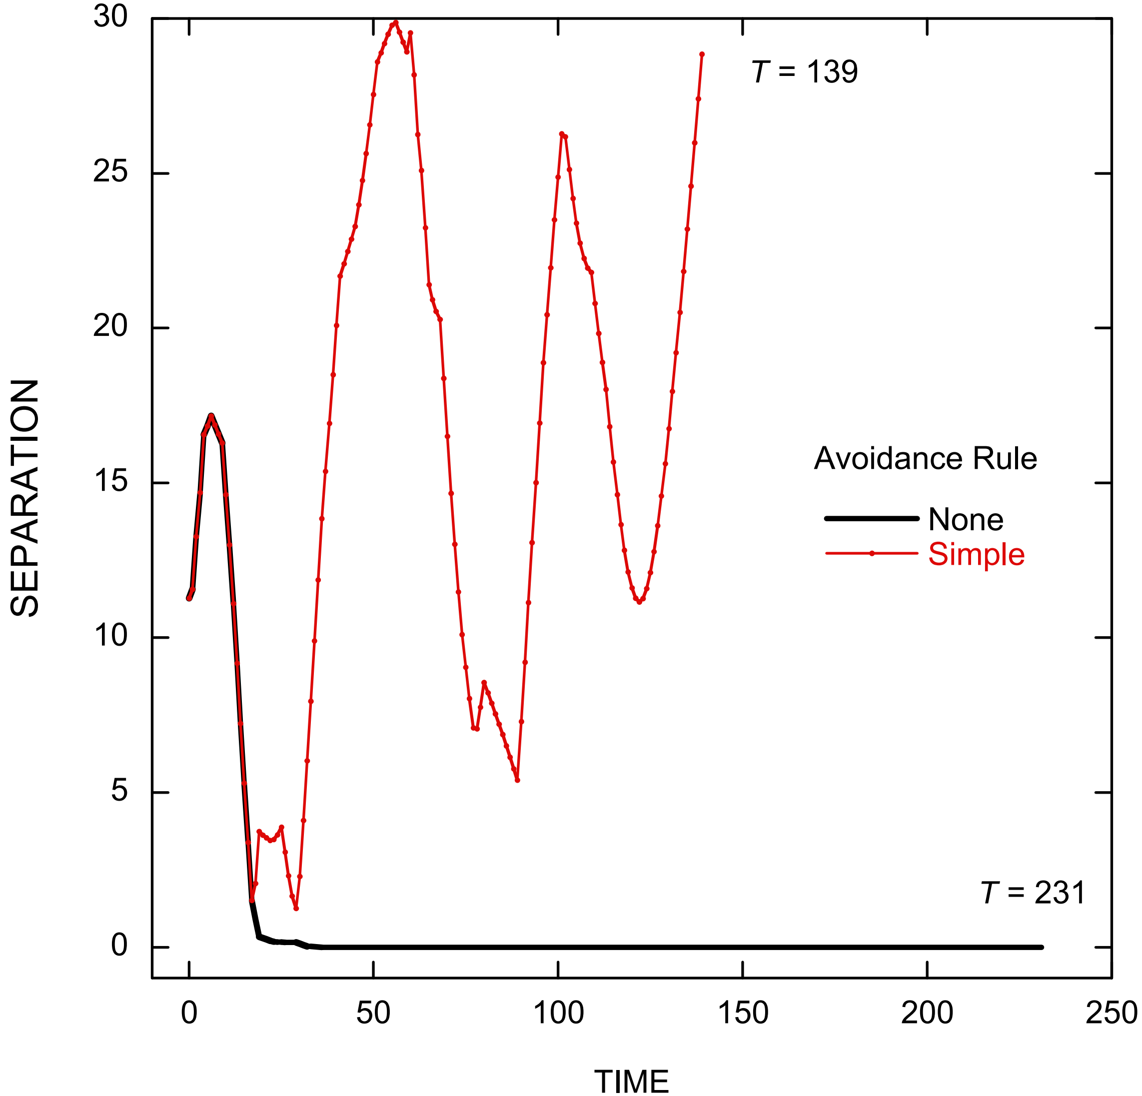

Improvement (Time)

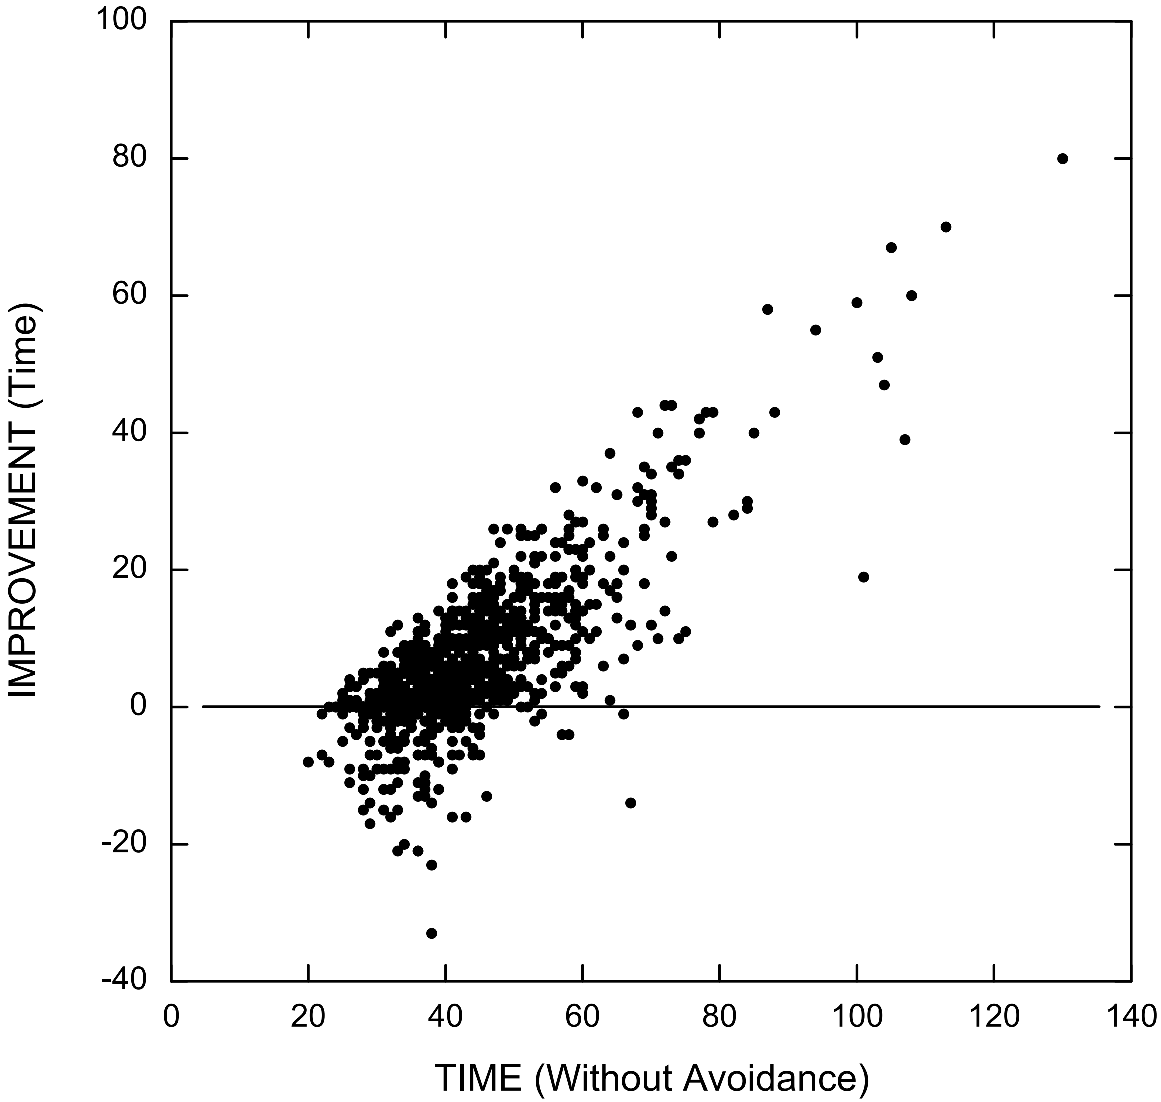

Improved Avoidance Rule

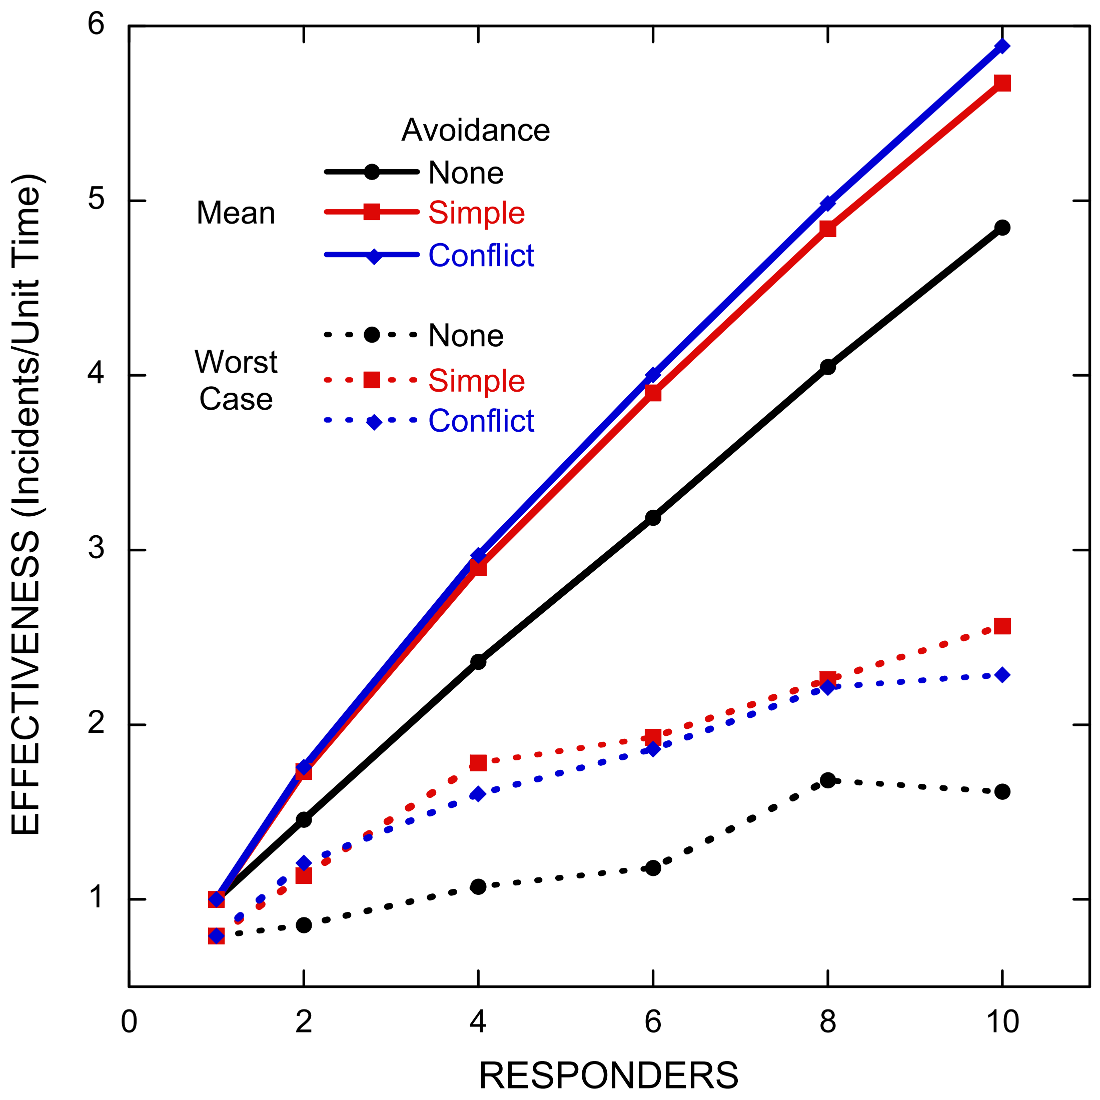

Improved Rule Results

- Performance improvement
  - 18% average
  - 43% worst case (59% for simple avoidance)
- Significant narrowing of distribution (improved predictability)
- Diminishing returns
  - Second responder has 73% marginal effectiveness; ninth has only 43%
- Improved avoidance rule still sometimes makes things worse
- Problem may be caused by competition among three or more responders

Future Work

- Find source of remaining inefficiency (approximately 40% for 10 responders) after an avoidance rule is applied
  - Examine more sophisticated variations of the avoidance rule
  - Try other (“non-greedy”) heuristics
- Measures of performance or constraints other than time (distance traveled, resource efficiency, load balance, etc.)
- Limits on information sharing (delays, errors, general or selective restrictions on distribution)
- Effects of additional information (responder locations and/or goals)
- Variation among incidents (location relative to terrain, time to service)
- Variation among responders (speed over terrain, speed of service, capacity)

Conclusions

- Decentralized C2 can be effective in our model; for 10 responders with simple (greedy) behavior rule
  - Best case: 10x better performance than one responder
  - On average: ≈ 5x better than one
- Perverse (dysfunctional) self-organization
  - Produces a long tail of pathological cases
  - Can be corrected with avoidance rules
  - Average performance improves by 20% to ≈ 6x one responder
- Lack of direct communication can be (partially) compensated by detecting changes in the environment (stigmergy)
- Self-organization is not always apparent
  - Perverse behavior (pack formation) is obvious in the simulation
  - Avoidance rules eliminate that pattern, but behavior is equally self-organized
- Better rules and rule development methods are needed
  - Hard problems need heuristics not optimizations
  - Agents cannot always tell if they are part of a self-organized behavior or structure

Conclusions

- Decentralized C2 can be effective in our model
- Lack of direct communication can be (partially) compensated by detecting changes in the environment (stigmergy)
- For 10 responders with simple (greedy) behavior rule
  - Best case: 10x better performance than one responder
  - On average: ≈ 5x better than one
- Perverse (dysfunctional) self-organization
  - Produces a long tail of pathological cases
  - Can be corrected with avoidance rules
  - Average performance improves by 20% to ≈ 6x one responder
- Self-organization is not always apparent
  - Perverse behavior (pack formation) is obvious in the simulation
  - Avoidance rules eliminate that pattern, but behavior is equally self-organized
- Better rules and rule development methods are needed
  - Hard problems need heuristics not optimizations
  - Agents cannot always tell if they are part of a self-organized behavior or structure