

Fan-out: Measuring Human Control of Multiple Robots

Dan R. Olsen Jr., Stephen Bart Wood

Brigham Young University

Computer Science Department, Provo, Utah, USA

olsen@cs.byu.edu, bart_wood@yahoo.com

ABSTRACT

A goal of human-robot interaction is to allow one user to operate multiple robots simultaneously. In such a scenario the robots provide leverage to the user's attention. The number of such robots that can be operated is called the fan-out of a human-robot team. Robots that have high neglect tolerance and lower interaction time will achieve higher fan-out. We define an equation that relates fan-out to a robot's activity time and its interaction time. We describe how to measure activity time and fan-out. We then use the fan-out equation to compute interaction effort. We can use this interaction effort as a measure of the effectiveness of a human-robot interaction design. We describe experiments that validate the fan-out equation and its use as a metric for improving human-robot interaction.

Author Keywords

Human-robot interaction, multiple robots, fan-out

ACM Classification Keywords

H5.2 – User Interfaces, I2.9 - Robotics

INTRODUCTION

As computing becomes smaller, faster and cheaper the opportunity arises to embed computing in robots that perform a variety of “dull, dirty and dangerous” tasks that humans would rather not perform themselves. For the foreseeable future robots will not be fully autonomous, but will be directed by humans. This gives rise to the field of human-robot interaction (HRI). Human-robot interaction differs from traditional desktop GUI-based direct manipulation interfaces in two key ways. First, robots must operate in a physical world that is not completely under software control. The physical world imposes its own forces, timing and unexpected events that must be handled by HRI. Secondly, robots are expected to operate independently for extended periods of time. The ability for humans to provide commands that extend over time and can accommodate unexpected circumstances complicates the

HRI significantly. This problem of developing interfaces that control autonomous behavior while adapting to the unexpected is an interesting new area of research.

We are very interested in developing and validating metrics that guide our understanding of how humans interact with semiautonomous robots. We believe that such laws and metrics can focus future HRI development. What we are focused on are not detailed cognitive or ergonomic models but rather measures for comparing competing human-robot interfaces that have some validity. In this paper we look at a particular aspect of HRI, which is the ability for an individual to control multiple robots simultaneously. We refer to this as the *fan-out* of a human-robot team. We hypothesize that the following *fan-out equation* holds,

$$FO = \frac{AT}{IT}$$

where

- *FO*=fan-out or the number of robots a human can control simultaneously,
- *AT*=activity time or the time that a robot is actively effective after receiving commands from a user,
- *IT*=interaction time or the time that it takes for a human to interact with a given robot.

In this paper we develop the rationale for the fan-out equation and report several experiments validating this equation. We show that the equation does describe many phenomena surrounding HRI but that the relationships are more complex than this simple statement of fan-out implies. We also describe the experimental methodologies developed in trying to understand fan-out. We present them as tools for evaluating and measuring design progress in HRI systems.

The robotic task domain that we have focused on is search and rescue where robots must cover an indoor, urban or terrain environment in search of victims, threats, problems, or targets. Although we have restricted our work to this domain, we are hopeful that our methods and metrics will extend to other HRI domains.

PRIOR WORK

Others have done work on human-robot interaction. Sheridan has outlined 5 levels of robot control by users

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

[14]. The levels range from teleoperation, where the user is directly engaged with the actuators of robot, through various levels of computer intervention between the user, the sensors and the actuators, to full autonomy with users merely setting very high-level goals. Fong and Thorpe [8, 9] demonstrated collaborative control where the human and the robot share the initiative, with the robot seeking guidance when needed. These with a variety of other approaches are characterized by their system architecture. Although human-robot interfaces are provided, there is little study of the nature of that interface nor on how to evaluate the quality of the interface.

There have been a number of proposals for new modalities for controlling robots including haptics, gestures, PDAS [7]. Others have looked at the visualization and context memory problems that arise when driving robots. The Egosphere is one such solution [6].

There is also a great deal of work on using multiple robots on a task. There are fully autonomous swarming approaches such as Bruemmer, et al [3]. These have very little human intervention because the desired task is preprogrammed. Other autonomous robot teams have done janitorial tasks, box pushing and earth moving [12, 13]. All of these teams have used very little human intervention. Other multi-robot systems have robots operating in formations [2, 4, 16] or according to predefined deployment behaviors [15] These approaches allow users to direct the work of a number of robots simultaneously. Fong et. al. [10] point out the problems with dividing human attention among multiple robots and propose a collaborative control model for driving. In essence their proposals increase the neglect and activity time of the robots to achieve higher fan-out. Others have used a “select and command” model for controlling multiple robots [11].

However, none of these have been carefully evaluated as to the advantages or decrease in effort afforded by the various user interface designs. In most cases the control architecture is intertwined with the human-robot interface making it hard to distinguish which part of the solution is contributing to progress. In this paper we describe a model for isolating and measuring the human-robot interface for teams of robots.

A SAMPLE ROBOT WORLD

To explain our fan-out ideas, we pose the example robot world shown in figure 1. In this world there are robots, targets and obstacles (trees & rocks). The task is for all targets to be touched by robots as soon as possible. This is an abstraction of a search task.

We can assume a simple-minded robot that accepts a direction and a distance from its user and will move in that direction until it either travels the indicated distance or encounters an obstacle, in which case it stops. In figure 1 the robot has three legs to its journey each characterized by a different user command. However, the robot's guidance

may not be perfect. It may drift to the left on the first leg, run into the trees and stop early. Its odometry may be faulty and it may overrun the end of leg one necessitating additional commands from the user to extricate it from the dead-end of rocks and trees.




Figure 1 – Simple Robot World

This example illustrates two measures that are important to our model of fan-out. The first is *neglect-time*. That is the time the robot can run while being ignored by its user. Neglect time is a measure of a robot's autonomy. This is very similar to Crandall's neglect tolerance [5]. Unlike Crandall's work, we are interested in multiple robots rather than efficient interfaces to a single robot. The second measure is *activity-time*, which is the time the robot will make effective progress before it either gets a new command from the user, it stops making effective progress or it completes the command the operator gave it. Neglect time and activity time are not the same. For example, if the user does not trust the odometry, he may watch the robot to make certain it does not overshoot the end of leg 1. The robot is independently active, but is not being neglected. This difference has an important impact on multiplexing human attention among multiple robots.

The relationship between activity time (AT) and neglect time (NT) is determined by the amount of overlap (O) between robot activity and interaction time (IT). Overlap is the percentage of the interaction time where the robot is also active.

$$AT = O * IT + NT$$

This relationship is illustrated by driving a car. The interaction time and the activity time of a car are almost completely overlapped (O=1.0). A car is almost always moving when the driver is steering it. In addition, the neglect time for a car is very small, therefore AT is not much larger than IT. Plugging this into the fan-out equation we see that a person cannot drive more than one car at once. In the case of a manufacturing robot, the robot is not at all active during setup (O=0.0) but the robot will run for days or months after a day of setup. Thus AT is many times larger than IT and the fan-out is quite high. The experimental models that we finally used are based on AT and IT. The relationship between O, NT and AT does not impact our comparisons of various human-robot interfaces.

In our simple robot world we can give the robot more intelligence. If it encounters an obstacle it can bounce off and continue trying to make progress towards its next check point. Thus the robot will operate longer without intervention (increased AT) and the user can trust it more (increased NT). Adding some local vision and planning, the robot might also detect the cul-de-sac of trees and rocks and not enter there without more explicit instructions. Again the robot can be trusted more and NT can increase. Increasing a robot's trusted intelligence can increase its neglect time and thus increase fan-out.

RATIONALE FOR FAN-OUT

The primary reason for our interest in fan-out is that it is a measure of how much leverage of human attention is being provided by a robot's automated capabilities. Autonomy is not an end unto itself. We may study it in that way, but in reality we want automation because it expands human capacity. Fan-out is a measure of the leverage of human attention.

The ability for a human to control multiple robots depends upon how long that robot can be neglected. If a robot makes no progress while being neglected, the human will have no attention to devote to other robots. However, as will be shown, it is difficult to measure neglect time. Instead we measure activity time, which is an average amount of time that a robot functions between instructions from the user. If we divide the average activity time by the amount of time a user must interact with each robot, then we get the fan-out equation.

$$FO = \frac{AT}{IT}$$

However, the relationships are not as simple as this analysis might indicate. We will discuss these interrelationships along with the experimental data. The key point, we believe, in understanding these complexities is that IT is not monolithic. Our current hypothesis is that there are at least 4 components to interaction time. They are:

1. *Robot Monitoring and Selection* – reviewing the state of all robots and deciding which robot needs the user's attention next.
2. *Context Switching* – when switching attention between robots the user must spend time requiring the goals and problems of the new robot.
3. *Problem Solving* – having required the situation the user must analyze the problem and plan a course of action for the robot.
4. *Command Expression* – the user must manipulate input devices tell the robot what to do.

Traditional direct-manipulation/desktop interfaces generally exhibit only components 3 and 4. The experiments that we have performed show the effects of some of these

components in the ways that the data deviates from the predictions of the fan-out equation.

Having broken down IT, it seems that IT should increase with FO. This is because the more robots there are to control, the greater the monitoring and robot selection time. Also the more diverse situations the robots find themselves in the greater the context-switching time. As we will see in the data smarter robots can offload some of the problem solving time from the user and thus reduce IT.

MEASURING HRI

Our hypothesis is that the fan-out equation provides a model for how humans interact with multiple robots. The challenge in validating the fan-out equation is that interaction time (consisting of planning, monitoring and solving) occurs mostly in the user's mind and is therefore difficult to measure directly.

Measuring Neglect Time(NT) and Activity Time(AT)

The properties of NT and AT are characteristics of a robot's ability to function in a world of given complexity. These times are functions of robot ability, task complexity and the user's understanding of the robot's ability.

In measuring either NT or AT we must control for the complexity of the task being posed. If the task posed in figure 1 had half as many trees, or no rocks, the task itself would be simpler and the robot could be safely neglected for a longer time. In essence, the more challenges a robot must face, for which it does not have sufficient autonomy, the lower NT and AT will be. The nature of the challenges will also have an impact. Therefore any measurements of NT or AT must control for the nature of the tasks being posed. We term this *task complexity*.

Our first approach to measuring NT ignored the role of the user in determining the robot's activity. We assumed that there was some measurement of NT in the context of a given task complexity. To measure NT we would randomly place a robot at some location in the world, give it a random goal and then measure the average time that the robot would operate before reaching the goal or failing to make progress.

However, this approach failed to produce data that was consistent with the fan-out equation. After reviewing the videotapes of actual usage we found that this a priori measurement consistently overestimated NT. We identified three problems with this approach. The first is demonstrated on leg 1 of the robot route in figure 1. The robot could feasibly be neglected until it ran into the cul-de-sac of trees and rocks. However, users regularly saw such situations and redirected the robot early to avoid them. The second reason was that users frequently did not trust the robot to work out low-level problems. Users would regularly give the robots shorter-term goals that the user believed were possible. Thirdly, we did not have a good measure for how much a robot's activity overlapped the user's attention to the robot.

All of these failing led to NT predictions that were much larger than actual usage.

A simpler and more accurate measure we found was activity time (AT). We measure the time from a user's command of a robot until that robot either stops making progress or another use command is issued to it. We average these times across all robots in an experiment. This measure of activity time fit better with the fan-out equation and was much easier to measure.

This activity time measure is dependent on determining when a robot is making effective progress and when it is not. In our simple robot world, robots stop when they reach an obstacle. Thus a robot is active when it is moving. Because of the nature of our test user interface robots are always getting closer to the goal if they are moving. Therefore our simplistic measure of effective progress works in this case. For more intelligent robots this is more complicated. An intelligent robot must balance goal progress with obstacle or threat avoidance. This can lead to interesting feedback and deadlock problems, which cannot always be detected. These issues form the basis for many of the conundrums of Isaac Asimov's robopsychology[1]. In many situations, however, we can detect lack of progress and thus the end of an activity.

Measuring FO

Our next challenge is to measure fan-out (FO). This, of course, is the measure that we want to maximize because it is an estimate of the leverage provided by our human-robot teams. Our first approach to fan-out measurement was the fan-out plateau. For a given task, we must have a measure of task effectiveness. In our simple robot world, this might be the total time required to locate all N targets. In other scenarios it might be the total amount of terrain surveyed or the total number of purple hummingbirds sighted. Parker has identified a variety of task effectiveness metrics for use with robot teams[13]. The fan-out plateau is shown in figure 2. As we give a user more robots to work with, the task effectiveness should rise until it reaches a point where the user's attention is completely saturated. At this point adding more robots will not increase task effectiveness and may in some situations cause a decrease if superfluous robots still demand user attention.

The attractiveness of the fan-out plateau measure is that it directly measures the benefits of more robots on actual task accomplishment. The disadvantage is that it is very expensive to measure. We might hypothesize that for a given HRI team, the fan-out plateau would be between 4 and 12 robots. We then must take 8 experimental runs to find the plateau (if our hypothesis was correct). Individual differences in both users and tasks require that we must take many runs at each of the 8 possibilities in order to develop a statistically significant estimate of the fan-out plateau. Since realistic tasks take 20 minutes to several hours to accomplish, this measurement approach rapidly consumes an unrealistic amount of time.




Figure 2 – Fan-out Plateau

An alternative measure of FO comes from our ability to determine which robots are active and which are not. If we have knowledge of activity as needed by our AT measurement, then we can sample all of the robots at regular time intervals and compute the average number of active robots. What we do to measure FO is to give the user many more robots than we think is reasonable and then measure the average number of active robots across the task. This gives us a strong estimate of actual fan-out that is relatively easy to measure. Note that this is only an estimate of fan-out because the large number of robots introduces its own cognitive load. We believe, however, that ignoring unneeded robots will not impact the value of the metrics for comparisons among competing HRI solutions.

Task saturation

A key problem that we have discovered in measuring fan-out is the concept of task saturation. This is where the task does not warrant as many robots as the user could effectively control. A simple example is in figure 1. If we add another robot to the task, the effectiveness will not go up because one robot can reach the target just as fast as two or three. The problem is that the task does not justify more workers. We will see this effect in the experiments.

Measuring IT

To improve human-robot interaction (HRI) what we really want is a measure of the interaction time (IT). IT is the measure that will tell us if the user interface is getting better or just the robotic automation. Our problem, however, is that we do not have a way to directly measure IT. There are so many things that can happen in a user's mind that we cannot tap into. To measure the Fitt's law effects or keystroke effects will only measure the command expression component of the interface. Our experience is that in a multi-robot scenario, command expression is a minor part of the interaction time.

Solving the fan-out equation for IT can give us a method for its measurement.

$$IT = \frac{AT}{FO}$$

However, this measure of IT is only valid if the fan-out equation is valid and the FO and AT measures are true measures. As has been shown in the preceding discussion we have good estimates for both AT and FO but it would be too strong of a claim to say that we had accurately measured either value. Our approach then is to replace IT with what we call *interaction effort* (IE). Interaction effort is a unitless value that is defined as:

$$IE = \frac{AT}{FO}$$

This is obviously derived from the fan-out equation but it makes no claims of being an exact time. It is a measure of how much effort is required to interact as part of a human-robot team. Unlike interaction time, interaction effort does not give us an absolute scale for measuring interaction-time. The interaction effort measure does give us a way to compare the interactive efficiency of two HRI designs. A comparison tool is sufficient as a measure of progress in HRI design.

Validating the fan-out equation

Our model of IE depends upon the validity of the fan-out equation, which is difficult to prove without measuring IT or IE directly.

Our approach to validating the fan-out equation is as follows. If we have 1) a set of robots that have varying abilities and thus varying neglect times, 2) all robots have identical user interfaces and 3) we use the various types of robots on similar tasks of the same task complexity, then if the fan-out equation is valid, the measure of IE should be constant across all such trials. This should be true because the user interface is constant and IE should be determined by the user interface. The experiments described in the remainder of this paper will show where this does and does not hold.

ROBOT SIMULATIONS

As a means of validating the fan-out equations we chose robot simulations rather than actual robots. We did this for several reasons. The first is that it is much easier to control the task conditions. When trying to validate the fan-out equation we need careful controls. These are hard to achieve with real robots in real situations. Secondly we are trying to discover laws that model how humans interact with multiple independent robot agents. The physical characteristics of those agents should not change the laws. Third we want to test robots with a variety of levels of intelligence. Changing a simulated robot's sensory and reasoning capacity is much simpler than building the corresponding robots. To perform the experiments that we did, we would have needed a fleet of 15 robots (5 each of 3 types), with identical interfaces.

There is one way in which the real world differs sharply from our simulated world. In the real world, robots crash into obstacles, fall into holes, and run into each other.

Safety is a real issue and lack of safety reduces the user's trust. As discussed earlier reduced trust leads to reduced activity times. In our simulations, robots never crash or fail therefore trust is higher than reality. However, we believe that this will be reflected in different activity times and should not affect the validity of the fan-out equation.

The task

For our fan-out experiments we chose a maze-searching task. We built a random maze generator that can automatically generate tasks of a given complexity. We defined task complexity as the dimensions of the maze, density of obstacles and number of targets. Using our random maze generator we were able to create a variety of tasks of a given complexity. After random placement of obstacles and targets the maze was automatically checked to make certain that all targets were reachable. Our measure of task effectiveness was the time required for all targets to be touched by a robot.

All robots had the same user interface as shown in figure 3. The user controls are quite simple and the same for all experiments. Each robot has a goal represented by a small square that the user can drag around. The robot will attempt to reach the goal. The variation in robots is in how they deal with obstacles. For less intelligent robots the user can set a series of very short-term goals with no obstacles. For more intelligent robots more distant goals can be used with the robot working out the intervening obstacles.




Figure 3 – Dragging Robots

A major variation of this user interface, that we used in most of our tests, obscures all regions of the maze that have not been visited by robots, as in figure 4. The idea is that until a robot reaches an area and broadcasts what it finds, the terrain is unknown.




Figure 4 –Obscured World

Three types of robots

To test our fan-out theory of constant IE for a given user-interface we developed three types of simulated robots. The first type (*simple*) heads directly towards its current goal until it reaches the goal or runs into an obstacle. This is a relatively simple robot with little intelligence.

The second type (*bounce*) bounces off obstacles and attempts to get closer to the goal even if there is no direct path. It never backs up and thus gets trapped in cul-de-sacs. The bouncing technique solves many simple obstacle avoidance problems but none that require any global knowledge. This robot stops whenever it cannot find a local movement that would get it closer to the goal than its current position.

The third type of robot (*plan*) has a “sensor radius”. It assumes that the robot can “see” all obstacles within the sensor radius. It then uses a shortest path algorithm to plan a way to reach the point on its sensor perimeter that was closest to the goal. This planning is performed after every movement. This robot stops whenever its current position was closer to the goal than any reachable point in its sensor perimeter. This robot can avoid local dead-ends, but not larger structures where the problems are larger than its sensor radius.

We measured average neglect time for each of the types of robots using the random placement/task method. As robot intelligence increased, neglect time increased also. This gave us three types of simulated robots with identical tasks and user interfaces.

VALIDATING THE FAN-OUT EQUATION

To validate the fan-out equation we performed a number of experiments using our simulated robot world. Our experimental runs were of two types. In our initial runs individual university students were solicited and compensated to serve as test drivers. They were each given about 30 minutes of training and practice time with each type of robot and then given a series of mazes to solve using various types of robots.

Task Saturation

Task saturation showed up in the early tests. In our first tests we started all robots in the upper left hand corner of

the world. Since this is a search task there is an expanding “frontier” of unsearched territory. This frontier limits the number of robots that can effectively work in an area. Fan-out was low because the problem space was too crowded to get many robots started and once lead robots were out of the way users tended to focus on the lead robots rather than bring in others from behind as the frontier expanded. Because none of our users worked for more than 2 hours on the robots, there was no time to teach or develop higher-level strategies such as how to marshal additional workers. We resolved the frontier problem by evenly distributing robots around the periphery of the world. This is less realistic for a search scenario, but eliminated the frontier problem.

We originally posited two interface styles, one with the world entirely visible (*light worlds*) and the other with areas obscured when not yet visited by a robot (*dark worlds*). We thought of this as two UI variants. Instead it was two task variants. In the dark worlds the task is really to survey the world. Once surveyed, touching the targets is trivial. In the light worlds the problem was path planning to the targets. Since reaching known targets is a smaller problem than searching a world, task saturation occurred much earlier. Because of this all of our races were run with dark worlds (Figure 4).

Figure 5 shows the relationship between fan-out and remaining targets. The dark thin line is the number of remaining targets not yet touched and the lighter jagged line is the average number of active robots. This graph is the average of 18 runs from 8 subjects using planning type robots. Other experiments showed similar graphs. In the very early part of the run it takes time to get many robots moving. Then as targets are located, the problem becomes smaller and the fan-out reduces along with it. The crossover occurs because in a dark world the fact that one or two targets remain is not helpful because any of the unsearched areas could contain those targets.

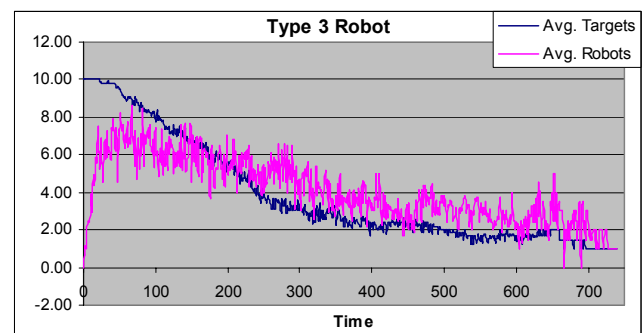


Figure 5 – Task Saturation

Test Data

These individual tests gave us good feedback and helped us refine our ideas about fan-out and interaction time. However, unmotivated subjects distorted the results. We had a number of subjects who just spent the time and

collected their money without seriously trying to perform the tasks quickly. It was clear that attempting to supervise multiple robots is more mentally demanding than only one. In many cases fan-out was not high even though from viewing the videotapes, the subjects were easily capable of doing better. To resolve this issue we held a series of “robot-races”. Groups of 8 people were assembled in a room each with identical workstations and problem sets. Each trial was conducted as a race with monetary prizes to first, second and third place in each trial. The motivation of subjects was better and the fan-out results were much higher and more uniform.

In our first race there were 8 participants all running 8 races using the dark worlds. The density of obstacles was 35% with 18 robots available and 10 targets to find. We ran 2 races with simple robots and 3 races each for the bounce and plan robots for a total of 64 trial runs. The measured fan-out and activity time along with the computed interaction time is shown in figure 6. Analysis of variance shows that there is no statistical difference in the interaction times across the three robot types. This supports our fan-out equation hypothesis.

Robot Type	Mean Fan-out	Mean Activity Time	Computed Interaction Effort
Simple	1.46	4.36	3.06
Bounce	2.94	7.82	2.77
Plan	5.11	14.42	2.88

Figure 6 – Test 1 - 18 robots, 10 targets, 35% obstacles

To evaluate our hypothesis that activity time and thus fan-out is determined by task complexity we ran a second identical competition except that the obstacle density was 22%. The data is shown in figure 7. Activity time clearly increases with a reduction in task complexity along with fan-out, as we predicted. The interaction time computations are not statistically different as we hypothesized.

Robot Type	Mean Fan-out	Mean Activity Time	Computed Interaction Effort
Simple	1.84	4.99	2.88
Bounce	3.36	11.36	3.38
Plan	9.09	24.18	2.69

Figure 7 – Test 2 - 18 robots, 10 targets, 22% obstacles

One of our goals in this work was to develop a measure of interaction effort that could serve as a measure of the effectiveness of a human-robot interface. To test this we ran a third competition of 8 subjects in 8 races. Test 3 was the same as test 1 except that we reduced the resolution of the display from 1600x1200 to 800x600. This meant that the mazes would not fit on the screen and more scrolling would be required. This is obviously an inferior interface to the one used in test 1. Figure 8 compares the fan-out and the interaction effort of tests 1 and 3.

Robot Type	Mean Interaction Effort		
	no scroll	scrolled	diff
Simple	3.06	4.48	46%
Bounce	2.77	3.47	25%
Plan	2.88	3.63	26%

Figure 8 – Compare scrolled and unscrolled interfaces

Test 3 shows that inferior interfaces produce higher interaction effort, which is consistent with our desire to use interaction effort as a measure of the quality of a human-robot interface.

However, figure 8 also shows a non-uniform interaction effort across robot types for the scrolling condition. This is not consistent with our fan-out equation hypothesis. Since all three robots had the same user interface they should exhibit similar interaction effort measures. Analysis of variance shows that the bounce and plan robots have identical interaction effort but that the simple robot is different from both of them.

We explain the anomaly in the simple robots by the fact that interaction effort masks many different components as described earlier and fan-out partially determines those components. Figure 9 shows the fan-out measures for test 3. The fan-out for the simple robots is barely above 1 indicating that the user is heavily engaged with a single robot. We watched the user behavior and noticed that the interaction is dominated by expressing commands to the robot with very little planning. With the bounce and plan robots the fan-out is much higher and users spend more time planning for the robot and less time trying to input commands to them.

Robot Type	Mean Fan-out
Simple	1.12
Bounce	2.47
Plan	3.97

Figure 9 – Fan-out for Test 3 (scrolled world)

It appears that when fan-out drops very low the nature of the human-robot interaction changes and the interaction effort changes also. To understand this effect better we ran a fourth competition where we varied the speed of the robots. Varying the speed of the robot will change its neglect time without changing either the robot’s intelligence or the user interface. A slower robot will take longer to run into an obstacle and therefore can be neglected longer. We used the same worlds and interface as in test 1, but we varied speeds across each run with only two robot types. The results are shown in figure 10.

Robot Type	Robot Speed	Mean Fan-out	Mean Activity Time	Computed Interaction Effort
Simple	3	2.54	7.21	3.05
Simple	6	1.21	3.51	3.09
Simple	9	0.89	3.26	3.67
Bounce	3	4.44	13.54	3.51
Bounce	6	3.11	9.60	3.10
Bounce	9	1.97	5.76	2.94
Bounce	12	1.82	4.42	2.51
Bounce	15	1.62	4.04	2.53

Figure 10 – Test 4 - Varying Robot Speed

For each class of robot, increasing the robot's speed decreases activity time and correspondingly reduces fan-out. Again with the fastest simple robots, the fan-out drops very low (the user cannot even keep one robot going all the time) and the interaction effort is quite different from the slower robots. This confirms the change in interaction when fan-out drops. This indicates that the fan-out equation does not completely capture all of the relationships between fan-out and interaction. This is also confirmed when we look at the interaction effort for the bounce robots. As speed increases, fan-out drops, as we would expect. However, interaction effort also drops steadily by a small amount. This would confirm the robot monitoring and context switching effort that we hypothesized. As fan-out is reduced, these two components of interaction effort should correspondingly reduce while other interactive costs remain constant. This would explain the trend in the data from test 4.

CONCLUSIONS

It is clear from the test that the fan-out equation does model many of the effects of human interaction with multiple robots. The experiments also indicate that interaction effort, as computed from activity time and fan-out can be used to compare the quality of different HRI designs. This gives us a mechanism for evaluating the user interface part of human-robot interaction. However, it is also clear that fan-out has more underlying complexity that the equation would indicate. This is particularly true with very low fan-out.

REFERENCES

- Asimov, I, *I, Robot*, Gnome Press, 1950.
- Balch, T. and Arkin, R.C.. "Behavior-based Formation Control for Multi-robot Teams." *IEEE Transactions on Robotics and Automation*, **14(6)**, 1998.
- Bruemmer, D. J., Dudenhoeffer, D. D., and McKay, M. D. "A Robotic Swarm for Spill Finding and Perimeter Formation," *Spectrum: International Conference on Nuclear and Hazardous Waste Management*, 2002.
- Chen, Q. and Luh, J.Y.S.. "Coordination and Control of a Group of Small Mobile Robots." In *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 2315-2320, San Diego CA, 1994.
- Crandall, J.W. and Goodrich, M. A., "Characterizing Efficiency of Human-Robot Interaction: A Case Study of Shared-Control Teleoperation." *Proceedings of the 2002 IEEE/RSJ International Conference Intelligent Robotics and Systems*, 2002.
- Kawamura, K., Peters, R. A., Johnson, C., Nilas, P., and Thongchai, S. "Supervisory Control of Mobile Robots using Sensory Egosphere" *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, 2001.
- Fong, T, Conti, F., Grange, S., and Baur, C. "Novel Interfaces for Remote Driving: Gesture, Haptic and PDA," *SPIE Telemanipulator and Telepresence Technologies VII*, 2000.
- Fong, T., Thorpe, C., and Baur, C., "Collaboration Dialogue, and Human-Robot Interaction," *Proceedings of the 10th International Symposium of Robotics Research*, 2001.
- Fong, T., and Thorpe, C. "Robot as Partner: Vehicle Teleoperation with Collaborative Control," *Workshop on Multi-Robot Systems, Naval Research Laboratory*, Washington, D.C, 2002.
- Fong, T., Grange, S., Thorp, C., and Baur, C., "Multi-Robot Remote Driving with Collaborative Control" *IEEE International Workshop on Robot-Human Interactive Communication*, 2001.
- Jones, C., and Mataric, M.J., "Sequential Task Execution in a Minimalist Distributed Robotic System" *Proceedings of the Simulation of Adaptive Behavior*, 2002.
- Parker, C. and Zhang, H., "Robot Collective Construction by Blind Bulldozing," *IEEE Conference on Systems, Cybernetics and Man*, 2002.
- Parker, L. E., "Evaluating Success in Autonomous Multi-Robot Teams: Experiences from ALLIANCE Architecture," *Implementations Journal of Theoretical and Experimental Artificial Intelligence*, 2001.
- Sheridan, T.B., *Teletrobotics, Automation and Human Supervisory Control* MIT Press, 1992
- Simmons, R., Apefelbaum, D., Fox, D., Goldman, R. P., Haigh, K. Z., Musliner, D. J., Pelican, M., and Thrun, S., "Coordinated Deployment of Multiple, Heterogeneous Robots," *Proceedings of the Conference on Intelligent Robots and Systems*, 2000.
- Wang, P.K.C., "Navigation Strategies for Multiple Autonomous Robots Moving in Formation." *Journal of Robotic Systems*, **8(2)**, pp. 177-195, 1991.