

The Use of Token-based Protocols in CSCW Tasks - An Empirical Examination

Christopher G. Harris¹

¹Computer Science Department, SUNY Oswego, Oswego, NY 13126
christopher.harris@oswego.edu

Abstract. In Computer Supported Cooperative Work (CSCW), many tasks require exclusive access to a shared resource by a single collaborator. Similarly, in distributed systems, mutual exclusion is required to ensure concurrency in a resource shared among several processes. These resource allocation algorithms can be divided into two genres: token-based and permission-based. To date, few empirical studies have evaluated token-based collaborative behavior in CSCW tasks. We examine four token-based protocols on a task which requires participants to properly order a series of screenshots obtained from ten short films. Using teams of 3, 4, and 5 participants who are collectively incentivized to perform the task as quickly as possible, we evaluate the effects of team size and token based protocol on task completion and participant satisfaction across 600 sessions. Our study determined that task satisfaction was negatively correlated with team size and positively correlated with the perception of “fairness”, or lack of potential bias, of each protocol.

1 Introduction

The use of purely virtual teams is growing. A 2012 RW3 CultureWizard survey of 3300 employees in 103 countries found that 87 percent of respondents belonged to a virtual team and 41 percent of these respondents had never met their teammates face-to-face [16]. Those who manage virtual teams have challenges that often differ from those who handle face-to-face teams, such as a heavier reliance on collaborative communication tools to interact. When virtual teams are static, teammates become aware of the strengths and weaknesses of their fellow collaborators and trust is established [2, 3]. However, with the increase in the use of semi-anonymous freelancers on websites like eLance¹ and Freelancer², the nature of virtual work frequently comprises ad hoc teams brought together to examine a specific problem or work on a specific task, and is temporary by design [12]. The use of temporary work teams are anticipated to increase in the next decade and beyond

¹ <http://www.elance.com>

² <http://www.freelancer.com>

due to a number of favorable factors, including more efficient virtual access to experts through advances in communication tools, a growing disparity in employment costs between the developed and the developing world, and the ease of transmitting payments across the globe [14, 21]. This virtual team approach to accomplishing ad hoc tasks has some shortcomings; as indicated in the 2012 RW3 CultureWizard survey, common problems encountered with virtual teams are that they fail to provide sufficient time to build relationships, require speedy decision making due to time zone differences, involve working with colleagues who do not actively participate, and require integration of different leadership and decision making styles. Naturally, these obstacles raise issues of trust; Jarvenpaa and Liedner investigated how trust develops in temporary virtual teams [10]. When faced with short deadlines and no face-to-face time to establish trust, the team members relied on expectations of trust from other settings that were familiar to them. However, virtual teams with low levels of initial trust and no actions to affirm trust continue to operate in a low-trust environment, which Jarvenpaa and Liedner found negatively impacted the team's performance.

It is therefore important to identify protocols that encourage efficiency by small virtual teams who have little or no previous working knowledge with each other. Distributed algorithms that provide mutual exclusive access to a shared resource, which we discuss in the next section, may provide some guidance. However, unlike machines, humans need some inducement to perform at their best. Thus, one approach to the rapid alignment of teams is to offer a team-based incentive. As with conventional teams, incentives that depend on team performance can play a role in aligning team members toward a specific goal [9]. However, some researchers have found that team-based rewards can also reduce motivation, particularly for more experienced team members [4]. In our study, we use team-based incentives for a task where each participant has an equal opportunity to contribute and examine if extrinsic incentives are effective in aligning collaborators toward a team-based goal.

The remainder of this paper is organized as follows. In the next section, we discuss the background and motivation behind our experiments. In Section 3, we provide a description of our experimental methods. In Section 4, we provide our results. Section 5 provides some additional analysis and is followed by a discussion of our general findings. We conclude and briefly discuss future directions of our work in Section 6.

2 Background and Motivation

Although the number of empirical studies in CSCW is limited, there are parallels with machine-based distributed systems, which we describe in this section.

2.1 Experiments with Virtual Teams

Despite the growth of ad hoc virtual teams over the last decade, there have been only few studies that have touched on collaboration in such environments; moreover, these have been focused on other aspects of task coordination rather than the protocols employed. Morris and Paepke performed a token-based study called Teamsearch using a tabletop device [13], but its examination of token passing and team size was not explored. Antle et. al. performed an empirical study on sustainable development using tokens in [1], but the focus was on sustainability, not on different mechanisms on concurrency. We found no other empirical examples in the CSCW literature examining the effectiveness of token-based protocols or number of collaborators.

2.2 Similarity with Mutually Exclusion Algorithms

One of the most studied areas in distributed systems is mutual exclusion (Mutex), where processes communicate by asynchronous message passing to coordinate mutually exclusive access to resources. These resource allocation algorithms, first introduced by Dijkstra [5], share a number of similarities with coordinating access among virtual human teams. First, machine processes and human collaborators can be considered as agents, and the motivation for each must be understood. In some cases, trust is unknown for each agent and cannot be assumed. Second, each protocol must determine the order in which agent (human or machine) can access a specific resource. Third, both protocols must handle issues of starvation, race conditions, and bias between participants. Last, both types of agents are evaluated on the same metrics, namely speed and task accuracy.

Broadly speaking, these Mutex algorithms can be divided into two families: *permission-based* and *token-based*. Permission-based algorithms, which typically require additional communication between each collaborator when permission to access a resource is required, are best suited for systems with infrequent demand for a resource; token-based systems, in which a privilege message, or *token*, is shared among all processes in a system (or collaborators in a task), are better suited to systems with resources that are in high demand. Furthermore, another advantage of token-based systems, unlike permission-based systems, is that they rely on a unique token and thus are deadlock-free. In our study, we limit our evaluation to token-based systems.

Token-based systems can be further divided into two different genres: *centralized* and *decentralized*. Centralized systems have a single decision maker that determines which process or collaborator has access to the token (and thus the resource), but may become a bottleneck if the token is not adequately managed. Centralized systems may also suffer from bias towards or against a particular collaborator, potentially leading to starvation. We examine one centralized system as part of our study. On the other hand, decentralized systems provide an established set of rules to determine which collaborator gets the token and which collaborator(s) must wait. Implicit and explicit token-based systems are described in more detail in [6].

Although a number of decentralized algorithms have been established in resource allocation, e.g., [18-20], we examine three that show the most promise for applying to human collaboration tasks: *time based*, where token requests are queued in the order they are requested, *last user determined*, where a token is sent from the current collaborator (user) to another collaborator based on the current token holder's sole decision, and *round robin*, where the token is sent in an established pattern (usually clockwise or counter-clockwise) between all collaborators who have made a request for the token.

Protocols can be *explicit*, with a predefined turn-taking rules, as with the time based and round robin protocols or *implicit*, where the freedom of choice can be made by a collaborator, as with the last user determined protocol and the centralized protocol.

We found no studies in the literature which have empirically compared these centralized and decentralized protocols for virtual teams in CSCW. Dommel and Garcia-Luna-Aceves describe a concept called floor control in [6] which is a temporary permission to access a specific resource; however, they do not conduct any empirical studies on their methods. Prasad et. al. describe an empirical study on floor control in [15], but their objective of their study is different from the one we examine here. Our objective is to obtain a better understanding of which token-based protocols work best for improving the performance of small teams. We measure this by the time taken to perform a task and by a self-reported collaborator satisfaction score. A better understanding of which protocol to use in a given situation can lead to better team performance and greater team satisfaction, particularly when teams are newly established or temporary in nature.

Our contributions are fourfold. First, we empirically examine different token-based systems with different numbers of randomly-assigned collaborators in a task in which each collaborator has an equal opportunity to contribute. This equal opportunity condition allows us to observe the contention between participants for a single resource. Second, in addition to measuring the amount of time taken to complete a task, we also measure each participant's satisfaction with the collaboration protocol. This metric gives us an opportunity to evaluate the perception of fairness, even when there is no evidence that any bias has actually occurred. Third, because ad hoc collaborators need to be quickly aligned towards a single objective, we provide an extrinsic incentive in the form of a monetary bonus for those teams who are able to complete the task quickly and accurately. Last, we examine demographic information obtained from the collaborators and examine if collaboration systems are favored by certain groups of people.

2.3 Research Questions

We examine the following research questions.

1. Are there any main or interaction effects with number of collaborators on task satisfaction?
2. Are there any main or interaction effects with type of token-based protocol on (a) time taken or on (b) task satisfaction?

3. Do demographics impact the task satisfaction of a given protocol?

3 Experimental Methods

To study collaborator behavior among ad hoc teams, we wanted a task that was easy to understand, provided no single collaborator with a specific advantage, and which could be easily measured. We developed a game that allowed us to examine two different factors: token-based approach (centralized, round robin, last-user determined, or time-based) and number of collaborators (3, 4 or 5).

3.1 Experimental Setup

We developed a game that required participants to complete a task in small groups. For each task, called a *session*, the type of token management was randomly assigned. Since most players participated in multiple sessions, participants were provided a new identity for each session. This reduced the possibility that any disagreements between collaborators in one session would carry over to future sessions. Each collaborator could only participate once for each short film, or in a maximum of 10 sessions.

We examined our factors across 10 short films found on YouTube; each film was between 3.5 to 8.5 minutes in length, with an average length of 5.5 minutes. These short films were selected through a poll conducted on Amazon Mechanical Turk (MTurk)³ prior to the commencement of our study. Twenty screenshots, called tiles, were taken from each short film, with considerable care taken so that the tiles could not be ambiguously ordered and that the tiles were associated with significant events in each short film.

Four tiles were randomly distributed to each collaborator (e.g., groups of 3 collaborators were assigned a total of $3 \times 4 = 12$ tiles; groups of 4 collaborators were assigned a total of $4 \times 4 = 16$ tiles). Participants watched the assigned short film and were then instructed to collaborate with others to order all assigned tiles in the same order as the short film. Participants could only add their own assigned tiles to the ordered list, called a *storyboard*, when they had possession of the token. Because each collaborator could only see the tiles assigned to him or her and those already added to the storyboard, no single participant could direct others on when or where to place their assigned tiles; each participant had to rely on efficient assignment of the token in order for the team to obtain a low time. Participants could type messages to each other using a simple messaging service. A session timer was started when all participants had finished watching the short film and a token was randomly assigned to one participant. The timer ended when all tiles were put in the correct order and at least one of the collaborators hit a “submit answer” button. The films used in our study and the URLs for each are provided in Table 1.

³ <http://www.mturk.com>

Table 1. Title, length and URL of each of the short films examined in our study, and the user rating obtained in a pilot study (10 = best, 1 = worst)

Film Title	Length	URL	User Rating
Lovefield	5:29	https://www.youtube.com/watch?v=4meeZifCVro	9.07
The Elevator	3:37	https://www.youtube.com/watch?v=Q-TQQE1y68c	8.87
My Shoes	3:54	https://www.youtube.com/watch?v=Y_F4-hV0iPM	8.80
Inbox	8:37	https://www.youtube.com/watch?v=75wNgCo-BQM	8.69
Pigeon: Impossible	6:14	https://www.youtube.com/watch?v=jEjUAnPc2VA	8.47
The Exam	7:00	https://www.youtube.com/watch?v=1HC1ANf4L6s	8.20
Fireflies	5:53	https://www.youtube.com/watch?v=1d_mCmMdlLY	8.07
The Date	3:48	https://www.youtube.com/watch?v=t7PQJ2yylvI	8.00
Carrot Crazy	3:28	https://www.youtube.com/watch?v=7V7MOk0FZrg	7.93
Thirst	6:30	https://www.youtube.com/watch?v=ck0028dgUmA	7.87


3.2 Assigning Participants to Sessions

The single-greatest challenge we faced in our game design was the ability to queue 3-5 players to be available at the same time to play a synchronous game. We eventually resolved this issue using a slightly modified version of TurkServer [11] to queue participants until we had the necessary number of players available. Even with this queuing method in place, the attrition rate in our study was 17.9 percent, illustrating the challenges of multi-player synchronous games.

Our study was conducted between June 2, 2013 and October 15, 2013. Participants were initially solicited through MTurk and via word-of-mouth and were paid US \$0.10 to provide demographic information (age, gender, self-reported location), to participate in a session, and to provide satisfaction on the protocol used for that session. We recorded their IP address to verify their self-reported location, and did not include participants in our demographic evaluation if the locations differed.

Participants could be assigned to different token-based protocols for each of the 10 sessions, but they would always be assigned teams with the same number of collaborators. A total of 784 unique players participated, representing an average of 3.18 sessions per player. Separate leaderboards were maintained for each combination of token-based protocol and number of collaborators, for a total of 12 leaderboards. Collaborators were told that they (along with their team) would be entered in a drawing for a \$20 cash bonus provided their session was listed on the leaderboard at the end of our study. This incentive appeared to motivate a number of participants to play as much as possible, as nearly a third of all collaborators played the maximum of 10 sessions. Participants in the 3-collaborator models had a higher retention rate than those in the 4 or 5 collaborator models. Nearly a third (32.6%) of all participants played the maximum of 10 sessions, which is a far higher retention rate than typical for repeated tasks conducted through crowdsourcing (e.g., [7, 8, 17]). This indicates the monetary incentives we offered worked well and the task provided was sufficiently engaging. The distribution of number of sessions per player is provided in Figure 1.

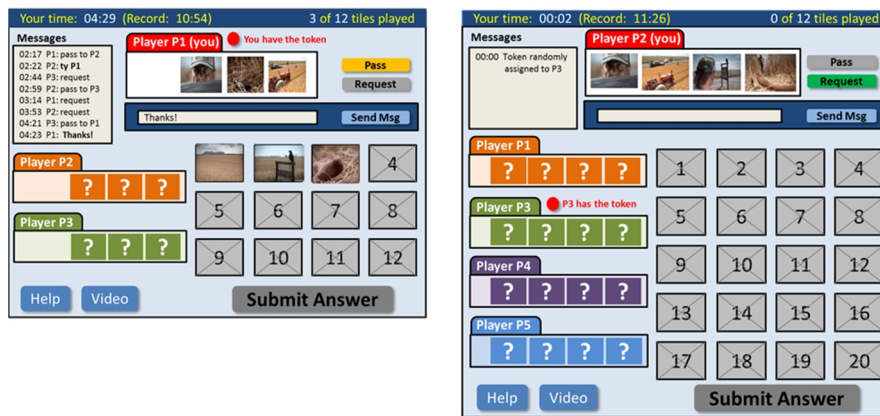
Fig. 1. Participant retention rate used in this study, broken down by number of collaborators.



3.3 Game Interface

The game interface used in our study was made available to participants through a unique link for each of the 600 sessions we conducted. Once the participant clicked on the session link, they entered a queue and waited for the right number of participants to arrive. Queuing time ranged from a few seconds to several minutes.

Fig. 2. Examples of the game screens for the three-collaborator version (left) and the five-collaborator version (right) for the time-based protocol. On the left, the participant currently has the token whereas on the right, another participant has the token.



Next, the participant is provided with the rules for the assigned token-based protocol and the URL for the short film. After they complete viewing of the short film, the partici-

participant clicks a button and waits for the other participants. Figure 2 provides two examples of the game screens provided to each participant with the version with three collaborators on the left and the version with five collaborators on the right (the version with four collaborators is not shown due to space).

At the top of each participant's game screen, they are provided a status bar containing the elapsed time, the best time recorded for the combination of number of collaborators, and the number of tiles currently on the storyboard. Below the status bar, the four tiles randomly assigned to the participant are provided, along with buttons to pass and request the token. Figure 3 shows the different participant areas for each of the four token-based protocols. A messaging window is given to communicate with the other users. The user can also observe who currently has the token along with the other token requests along the lower left-hand side of the screen and the storyboard along the lower right-hand side of the screen.

Fig. 3. Examples of options available to the participant for passing and requesting tokens for our four protocols: (a) centralized (b) time-based, (c) last user determined and (d) round robin.



In order to place one of their four assigned tiles on the storyboard (or to rearrange tiles already played by other participants on the storyboard), the token is required to be in their possession. Therefore only one participant can interact with the storyboard at any one time. The protocols used are:

- a. *Centralized* – One randomly assigned participant is randomly assigned the token and they are given the task to send and request the token for all other participants for the entire session.
- b. *Time Based (decentralized)* – One participant is initially assigned the token by random determination. Other participants who wish to obtain the token can request the token by pressing the “request” button. The token requests are processed in chronological order based on the time the button was pressed.
- c. *Last User Determined (decentralized)* – One participant is initially assigned the token by random determination. Each participant, when they have possession of the token, can explicitly send the token to any other participant, but they do not have the ability to revoke the token from other participants.

d. *Round Robin (decentralized)* – One participant is initially assigned the token by random determination. Once a participant passes the token, token travels in a circular order until it reaches the next participant that has made a token request.

Participants can view the short film in another window by pressing a view button at the bottom left of the game screen. Also at the bottom left of the screen, there is a help button that provides the rules given to each participant at the beginning of the game.

The timer stops once all the tiles are in the correct order on the storyboard and a “submit answer” button below the storyboard (on the bottom right of the screen) has been pressed by one of the participants. Participants are not told which tiles are incorrectly ordered, but a count of incorrectly ordered tiles is provided to all participants once the “submit answer” button has been pressed. The “submit answer” button remains greyed out and unavailable until all tiles have been placed on the storyboard. Once the user has submitted their work, they are asked the following question:

On the task you have just completed, how satisfied were you on the fairness on the sharing of the token? (1 = extremely dissatisfied, 10 = extremely satisfied).

The participant is then presented with a leaderboard containing the top 10 team scores for that combination of token-based protocol and number of collaborators, along with a unique combination of four letters and numbers for that session. A link was provided to each participant to periodically review their standing on the leaderboard. Each participant on the team with the lowest time at the end of our study was paid \$20.

4 Results

The mean and standard deviation for the time taken and the average task satisfaction rating from 600 sessions is presented in Table 2. From these results, we can observe that the task time increases and the task satisfaction decreases as the number of collaborators in the task increase from three to five. Additionally, the round robin protocol results in higher task completion times regardless of the number of contributors used, while the time-based protocol results in the shortest completion times. For tasks with fewer collaborators, the time-based protocol results in higher task satisfaction scores, but as the number of collaborators grow to five, it becomes the protocol with the lowest task satisfaction score. We address this issue further in the discussion section.

A two-way ANOVA was conducted that examined the effect of the number of collaborators and token-based protocol on the *time taken* to perform the task. The results are provided in Table 3. Although no significant interaction effect was found, the simple main effects of the number of collaborators, $F(2, 588) = 1364.548, p < .0001$, as well as with the protocol used $F(3, 588) = 32.023, p < .0001$, were both significant. It is expected that

as the number of collaborators participating increases, the time taken also increases, since the size of the storyboard increases as well. However, the effect of the protocol used on time to perform the task answers part (a) of our second hypothesis and we find protocol does have a simple main effect on the time taken.

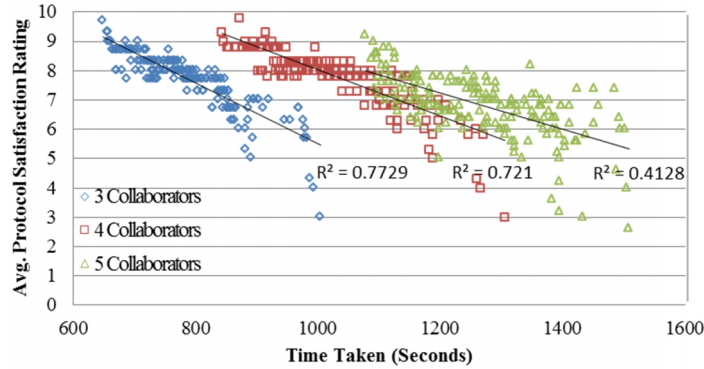
Table 2. Mean and standard deviations for time taken and average satisfaction rating for each factor (number of collaborators, protocol used).

Factor	Time (seconds)		Avg. Satisfaction Rating		N
	μ	SD	μ	SD	
Number of Collaborators					
3	788.6	71.1	7.70	0.81	200
4	1037.8	95.4	7.75	0.81	200
5	1261.3	98.7	6.85	0.90	200
Protocol Used					
Centralized	1008.2	86.3	7.68	0.86	150
Time-based	996.8	85.6	7.61	0.75	150
Last User Determined	1021.9	72.8	7.60	0.76	150
Round Robin	1090.0	108.9	6.83	1.00	150
Number of Collaborators × Protocol Used					
3, Centralized	767.0	63.1	7.89	0.91	50
3, Time-based	757.4	55.8	8.20	0.48	50
3, Last User Determined	783.4	70.2	7.73	0.82	50
3, Round Robin	846.4	95.3	6.99	1.03	50
4, Centralized	1019.3	100.7	7.95	0.92	50
4, Time-based	1008.0	97.6	8.24	0.48	50
4, Last User Determined	1024.2	74.7	7.77	0.82	50
4, Round Robin	1099.8	108.6	7.02	1.03	50
5, Centralized	1238.2	95.1	7.22	0.75	50
5, Time-based	1225.0	103.4	6.39	1.29	50
5, Last User Determined	1258.0	73.6	7.30	0.63	50
5, Round Robin	1323.9	122.7	6.49	0.95	50

A two-way ANOVA was also conducted to examine the effect of the number of collaborators and token-based protocol on the average task satisfaction rating. Simple main effects were found for the number of collaborators, $F(2, 588) = 67.503, p < .0001$, as well as with the protocol used $F(3, 588) = 31.743, p < .0001$. An interaction effect between the two factors were also found, $F(6, 588) = 9.229, p < .0001$. This addresses our first hypotheses and part (b) of our second hypothesis. We find that both the number of collaborators and the protocol used does have an effect on task satisfaction rating.

The relationship between the task satisfaction rating and the time taken for the task was strong, particularly for tasks with three or four collaborators, as shown in Figure 4. As the number of collaborators on a task increases, the task satisfaction rating decreases but the number of collaborators becomes a weaker overall predictor for task satisfaction.

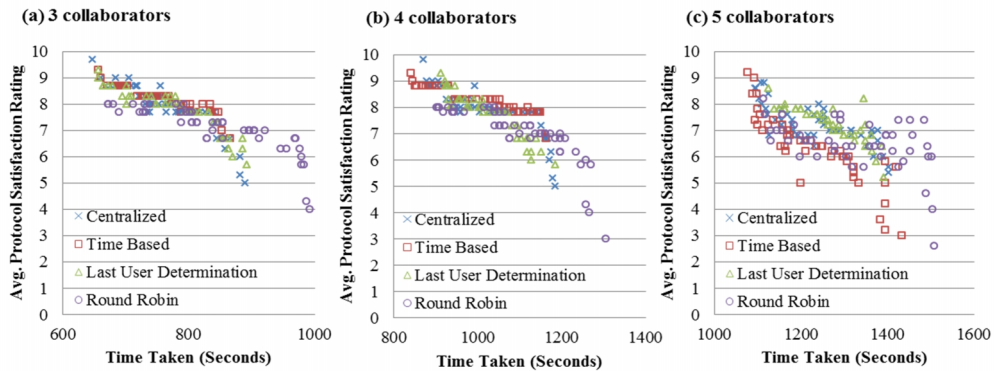
Fig. 4. Relationship between the time taken and the average task satisfaction rating for tasks with 3, 4 and 5 contributors. The R-squared value is given for tasks with different collaborator sizes.



5 Analysis and Discussion

In Figure 5, we examine how our two factors interact for each of our 600 sessions. We break each of the different collaboration sizes (our first factor) into different graphs and examine the average task satisfaction rating (our second factor) as it relates to the time taken.

Fig. 5. Illustration of task satisfaction rating and time taken for each token based protocol, broken into separate graphs by number of collaborators.



We observe that clustering is still fairly cohesive based on the token-based protocol used; however, as the number of collaborators increases from three (left-most graph) to five

(right-most graph), the clustering becomes less cohesive. This indicates that the time taken becomes a weaker predictor of the task satisfaction and other factors, particularly the protocol used may become a better predictor. We also notice that the round robin protocol is our consistently weakest performer whereas the time based protocol is a strong performer when the number of contributors used is three or four, but a poor performer when the number of contributors increases to five. We logged information on each of the 600 sessions, including the messaging between participants, and therefore we infer some reasons why this may be the case. First, many participants in the five-collaborator model using the time-based protocol were requesting the token as soon as possible but took considerable time to play their tiles. Much like a contestant on a game show who presses the response button in order to “jump the queue” before listening to the complete question, these participants were requesting the token before they were prepared to take action, frustrating others who were prepared but further down in the queue. This gave the impression that the protocol was being manipulated by some players, lowering that protocol’s satisfaction rating. We did not observe this behavior in the models with three or four collaborators.

We were also surprised to see that the centralized protocol obtained better times than the last-user selection protocol (since both are implicit, we had initially anticipated them to have similar times). From the logged messages passed between participants, we see that if one participant is identified as the “token master”, they are quicker to respond to token requests made by the other participants, whereas the participant in the last user determination model frequently took longer to pass the token, even when it had been requested by many players. The task satisfaction rating for the centralized model was lower than that of the last user determined protocol, however. This appeared to be skewed by participants in a few sessions; in those sessions, we observed there was an implication that two participants were passing between each other and ignoring all other participants. This implies that bonds between collaborators could be identified and established quickly. This phenomenon was not observed in the centralized model.

We recorded the self-reported demographics of each participant, including age range, gender and location. This gave us some insight into how the task satisfaction rating of each protocol was affected by demographics. We broke geographic locations into seven regions and three age groups, which is given in Table 3. We performed chi-squared tests on the ratings each participant provided.

Overall, we find region do have preferences, $\chi^2(6, N=776) = 262.883, p < 0.0001$; task satisfaction ratings for the explicit models (time based and round robin) were higher for participants from Europe, North America and Australia, whereas implicit models (centralized and last player determined) were favored by participants from South Asia, South and Central America and Africa and the Middle East. The centralized protocol was favored more by females, $\chi^2(3, N=321) = 170.833, p < 0.0001$, and by participants over 35 years of age $\chi^2(3, N=154) = 112.40, p < 0.0001$, whereas the time-based protocol was favored by males $\chi^2(3, N=455) = 211.54, p < 0.0001$, and by participants under 25 years of age,

$\chi^2(3, N=303) = 89.481, p < 0.0001$. We therefore find that the protocol used does matter for task satisfaction for different geographic regions, genders and age groups. This reinforces our third hypotheses. We plan to explore the relationship between these demographic factors and protocol preferences in future work.

Table 3. Demographic breakdown of participants by region, gender and age range.

	Female	Male	N
Region			
Africa and Middle East	6	37	43
Australia and Oceania	13	12	25
East and North Asia	18	29	47
Europe	73	69	142
North America	106	90	196
South and Central America	16	37	53
South Asia	89	181	270
TOTAL	321	455	776
Age Range			
Under 25	141	162	303
25-34	120	199	319
35 and Over	60	94	154
TOTAL	321	455	776

6 Conclusion

The number of tasks incorporating virtual collaborations is expected to increase due to favorable conditions in the global workplace. A large percentage of these collaborations will be short-term by nature, which means trust and bonding exercises between collaborators are not practical. In tasks where collaborators must compete for limited a limited set of resources, factors such as bias, concurrency, and deadlocking need to be adequately addressed. Thus, the mutual exclusion protocol and the size of the collaboration teams, when appropriately chosen, can contribute to task success.

An empirical study with 600 different tasks and 2400 participants was conducted that examined these two factors. We designed a game that had participants collaborate in a synchronous task of different sizes and different token-based protocols, four of which we borrow from work in distributed systems. The game used incentives to have participants work on a task to put screenshots of a short film in order on a storyboard. The task was designed to be easy to learn and require no prior external knowledge to participate. Providing monetary incentives aligned teams to focus on a single metric (task completion time) and as a result, our participant retention rate was much higher than expected, with a third of participants completing the maximum number of tasks available.

We explored several hypotheses in this study. We found that the time taken does depend on the protocol chosen. We also found that the self-reported task satisfaction rating

depends on the time taken by the team to complete a task and the protocol used, with an overall preference for the time-based protocol in smaller groups and for the last user determined and centralized protocols in the larger collaboration model. In future work, we hope to expand our evaluation to larger groups to see if what we observed with the five-collaborator model follows a trend or if it was an anomaly.

Since worker satisfaction in these collaboration models correlates highly with greater trust between participants, the choice of protocol can enhance the collaborator effort or detract from it. This is particularly true early in the collaboration process when each participant has little information to use to determine trust, as Jarvenpaa and Leidner had illustrated in an earlier study [10].

Demographics also play a part in satisfaction with the protocol employed, with some regions of the world preferring implicit models (centralized and last user determined) while other regions prefer explicit models, particularly time-based protocols. Females and younger users gave higher task satisfaction scores using the centralized protocol whereas males and younger participants were happier with the time-based protocol; however, we realize the role of demographic in virtual task design needs to be examined in more detail.

In future work, we plan to study the role of communication in token passing, specifically we wish to see if a participant's satisfaction was related more to familiarity with the token-passing protocol than on the protocol's performance. We also plan to examine additional token-based protocols and examine the role of bias in collaborative tasks.

References

1. Antle, A. N., Bevans, A., Tanenbaum, J., Seaborn, K., & Wang, S. (2011, January). Futura: design for collaborative learning and game play on a multi-touch digital tabletop. In Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction (pp. 93-100). ACM.
2. Armstrong, D. J., & Cole, P. (2002). Managing distances and differences in geographically distributed work groups. *Distributed work*, 167-186.
3. Binder, J. C. (2007). *Global project management: communication, collaboration and management across borders*. Gower Publishing, Ltd.
4. DeMatteo, J. S., Eby, L. T., & Sundstrom, E. (1998). Team-based rewards: current empirical evidence. *Research in organizational behavior*, 20, 141-183.
5. Dijkstra, E. W. (1965). *Cooperating Sequential Processes*, Technical Report EWD-123.
6. Dommel, H. P., & Garcia-Luna-Aceves, J. J. (1997). Floor control for multimedia conferencing and collaboration. *Multimedia Systems*, 5(1), 23-38.
7. Eickhoff, C., Harris, C. G., de Vries, A. P., & Srinivasan, P. (2012, August). Quality through flow and immersion: gamifying crowdsourced relevance assessments. In Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval (pp. 871-880). ACM.

8. Harris, C., & Wu, C. (2014). Using tri-reference point theory to evaluate risk attitude and the effects of financial incentives in a gamified crowdsourcing task. *Journal of Business Economics*, 84(3), 281-302.
9. Hertel, G., Geister, S., & Konradt, U. (2005). Managing virtual teams: A review of current empirical research. *Human Resource Management Review*, 15(1), 69-95.
10. Jarvenpaa, S. L., & Leidner, D. E. (1998). Communication and trust in global virtual teams. *Journal of Computer-Mediated Communication*, 3(4), 0-0.
11. Mao, A., Chen, Y., Gajos, K. Z., Parkes, D., Procaccia, A. D., & Zhang, H. (2012). Turkserver: Enabling synchronous and longitudinal online experiments. *Proceedings of HCOMP*, 12.
12. Mettler, A., & Williams, A. D. (2011). The rise of the micro-multinational: How freelancers and technology-savvy start-ups are driving growth, jobs and innovation. *Lisbon Council Policy Brief*, 5(3).
13. Morris, M. R., Paepcke, A., & Winograd, T. (2006, January). Teamsearch: Comparing techniques for co-present collaborative search of digital media. In *Horizontal Interactive Human-Computer Systems, 2006*. IEEE.
14. Passerini, K., El Tarabishy, A., & Patten, K. (2012). The Changing Nature of "Workspace" and "Workplace:" What It Means for SMEs. In *Information Technology for Small Business* (pp. 37-46). Springer New York.
15. Prasad, R. V., Jamadagni, H. S., Shankar, H. N., & Pawelczak, P. (2005, December). Fixing Number of Floors for Virtual Voice-Only Conference-an Empirical Study. In *Multimedia, International Symposium on* (pp. 120-127). IEEE Computer Society.
16. RW3-CultureWizard. The Challenges of Working in Virtual Teams RW3-CultureWizard, 2012. Available from: <http://rw-3.com/2012VirtualTeamsSurveyReport.pdf>
17. Rzeszotarski, J. M., Chi, E., Paritosh, P., & Dai, P. (2013, March). Inserting Micro-Breaks into Crowdsourcing Workflows. In *First AAAI Conference on Human Computation*.
18. Singh, M., & Tapaswi, S. (2012). Token Based Mutual Exclusion in Peer-to-Peer Systems. *Technologies and Protocols for the Future of Internet Design: Reinventing the Web*, 214-228.
19. Singhal, M. (1989). A heuristically-aided algorithm for mutual exclusion in distributed systems. *Computers, IEEE Transactions on*, 38(5), 651-662.
20. Suzuki, I., & Kasami, T. (1985). A distributed mutual exclusion algorithm. *ACM Transactions on Computer Systems (TOCS)*, 3(4), 344-349.
21. Thissen, M. R., Myers, S. K., Sikes, R. N., Robinson, J. P., & Grouverman, V. M. (2011). Globalization of Software Development Teams, *The Systemic Dimension of Globalization*, Prof. Piotr Pachura (Ed.), InTech, Available from: <http://www.intechopen.com/books/the-systemic-dimension-of-globalization/globalization-of-software-development-teams>