Crowdsourcing Convergence: Aggregating Partial Clusters to Facilitate Collaborative Convergence

Joel H. Helquist
Utah Valley University
joelh@uvu.edu

Christopher B.R. Diller
University of Nebraska
cdiller@unomaha.edu

John Kruse
MITRE
john@kruser.org

Abstract

This paper is an exploratory effort to investigate the possibility of using crowdsourcing to execute part of the collaborative convergence process. Participants were assigned with creating buckets or clusters from a random subset of the overall pool of brainstorming ideas. These sub-sorts were aggregated into a weighted graph and partitioned into discrete buckets. Analysis of this aggregated, consensus sort provides support that crowdsourcing may be a feasible option when organizing brainstorming ideas into discrete categories or buckets.

1. Introduction

Considerable research has examined how groups can work together collaboratively to achieve collective goals or objectives. The literature shows that collaborative GSS tools can facilitate a group as it moves through the various stages of collaborative work [1,2].

Generally, the activities of group collaboration can be characterized as being either divergent or convergent. The majority of this research has focused on the idea generation process or divergence stage of collaboration [3,4]. During this stage of collaboration, due to the parallel nature of the work, a group can generate a large number of ideas as each group member is working independently and has little need for process coordination [5].

Convergence represents the activities required to move the group from having many ideas to refining and focusing on the few ideas that are more valuable or worthwhile [6,7]. Convergence is a much more difficult activity than divergence for a number of reasons. First, the activity is largely serial in nature; the group must work together, rather than individually, to synthesize the ideas. The process coordination requirements are, as a result, much higher [7]. Also, these activities typically require more cognitive effort. Users must compare, contrast, and evaluate the brainstorming ideas in order to yield a refined set of ideas worth of further consideration. Typically, the products of a convergence stage will be more organized, concise, or actionable.

One of the ways a group can converge on ideas is by examining each brainstorming idea and grouping similar ideas together. By organizing the overall pool of brainstorming ideas into buckets of like or related ideas, the group is able to synthesize the ideas and identify those areas and ideas that are worthy of more attention or consideration. Typically, this process is conducted in a face-to-face, synchronous environment, with an expert facilitator guiding the discussion and action.

This paper furthers collaboration research by examining how to effectively converge through the use of a distributed, crowdsourced approach that can potentially lower cognitive costs and speed the convergence process for the group while still achieving acceptable results. Specifically, each member of the group is given a subset of the overall brainstorming ideas to organize into buckets. After the subsets are created by individual participants, they are then automatically aggregated to yield the end product of a consolidated, sorted set of brainstorming ideas.

The paper is organized as follows. First, a review of the existing literature is presented. Second, the aggregation of the distributed, crowdsourcing approach is discussed. Third, the metrics and measures are presented followed in the results section. The paper concludes with the implications, limitations, and future research.

HiCSS

2. Literature review

Collaborative systems, like Group Support Systems (GSS), have been widely used and researched. These tools aim to improve group productivity by enabling more efficient collaboration. These process improvements are afforded by the technology and by the process structuring or collaborative workflow [8].

Considerable research has been conducted on GSS tools through the years. Collaborative tools can bring a variety of efficiency gains over traditional, face-to-face collaboration [1]. For example, participants are able to work anonymously, allowing ideas to flow freely from all participants and each idea is evaluated on its own merit rather than the person who wrote the idea. Members of the group are able to work independently and in parallel, creating brainstorming input.

Besides the benefits of the technology tools, collaborative research has focused on collaborative workflows or ways to structure the collaborative work in order to yield the best output [9]. Two high-level categories or patterns of collaboration have emerged from this literature: diverging activities and converging activities [10].

Diverging activities are those that create more brainstorming ideas for the group to consider. As mentioned, GSS tools enable quick and efficient generation of collaborative ideas.

The convergence stage consists of activities whereby the group processes the overall pool or set of brainstorming ideas that was generated during the divergence, or idea generation, stage. There are a variety of activities that can take place during convergence; however, the end goal is to reduce the number of ideas that are worthy of further consideration or evaluation [7]. Oftentimes, these activities include mundane activities like removing duplicate or off-topic comments that were entered to more complex activities like grouping similar brainstorming ideas by theme. Grouping ideas into similar buckets or groups enables the collaborative participants to more easily process the overall pool of ideas and to make progress toward the end goal of the collaborative effort [10].

When comparing these two stages, participants enjoy the brainstorming more than the convergence of the ideas. As shown by Chen et al [11] in Figure 1, the user satisfaction levels increase during divergence (Idea Generation) but fall during convergence (Idea

Organization). Additionally, the length of time is considerably longer in convergence as the users must work together, oftentimes through an expert facilitator, to review, synthesize, and organize the ideas [12].
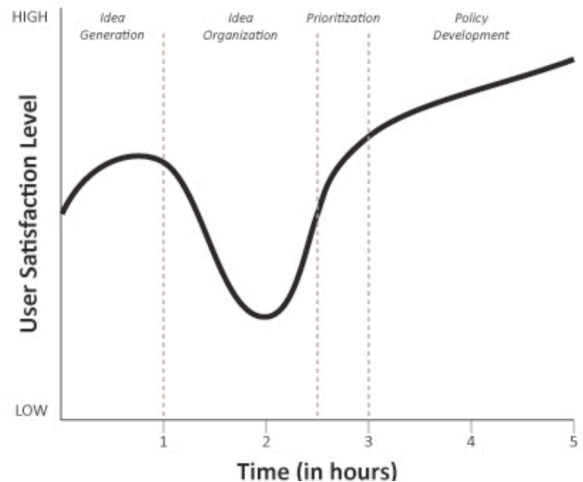

Figure 1: User satisfaction levels

Considerable research has examined how to effectively and efficiently execute divergent collaborative activities [7,13]. Numerous articles exist that examine a wide variety of factors impacting divergence, including such factors as group size, group proximity, and task composition [14,15]. However, far less research exists in the literature addressing how to most effectively converge. Recent work has increased in this area but considerable research is still needed.

One line of research in this area investigated using artificial intelligence (AI) to automatically cluster the ideas without the help or input of the collaborative participants [16,17,18]. The AI approaches have met with mixed results; the brainstorming ideas are such short snippets of text that automatic classification becomes a challenge. This line of research is certainly promising as technology and AI techniques improve.

This research paper looks at the process of how to effectively take the entire pool of brainstorming ideas and group them into similar buckets for later processing using human power. As mentioned previously, groups typically conduct this activity together, in a serial fashion. Oftentimes a facilitator will lead the group through the activity. This serial activity slows down the productivity of the group as all members work together. One of the goals of this area of research is to improve the dip experienced during the convergence activities by the participants, see Figure 2 [19].
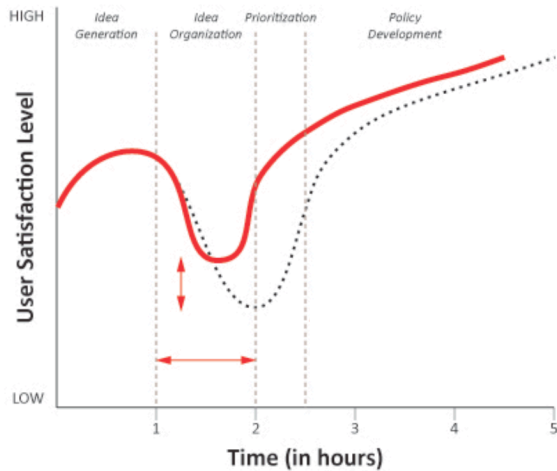
Figure 2: Improvement in user satisfaction levels [19]

This paper is an exploratory effort into the use of a crowdsourcing approach to organize the brainstorming idea into buckets. Each member of the collaborative group receives a subset of the overall brainstorming population to cluster or group. In this fashion, group members would each be able to work independently and in parallel again. Once enough subsets of the clustering have been completed, the system aggregates these partial sorts in order to find an overall sorting for the entire pool of brainstorming ideas. This aggregate sort represents a certain level of consensus. Each participant may potentially spend less time in the clustering activity than if the group completed the exercise together.

One of the benefits of this crowdsourced approach is that it can be applied to a wide variety of contexts. Participants do not necessarily need to meet in a face-to-face meeting or even a synchronous meeting to execute this approach to convergence.

Also, previous research has shown that this crowdsourcing approach reduces the burden on the participants [20]. Dividing up the sorting task reduced the perceived difficulty of the task and improved the satisfaction levels as compared with those doing a traditional sort. This research seeks to builds on this line of research by exploring the possibility of aggregating the results from the partial sorts. A key premise of this approach is that it is truly a crowdsourced approach; the items will be clustered by different participants in order to identify a consensus. Accordingly, this approach to clustering brainstorming ideas is more appropriate for specific collaborative contexts. For example, this approach may only be feasible when working with a large group as there may not be enough human capital to execute this plan otherwise.

This research is an exploratory effort into discovering ways to crowdsource the convergence process. The overarching research question is to investigate if this "divide and conquer" approach can be a viable option to helps group converge.

3. Crowdsourced clustering

The partial sorts used in this research project were selected from a brainstorming experiment using the school of business task. This task requires the subjects to brainstorm solutions to a school of business that is experiencing various issues [21]. These problems include items such as declining enrollment, problems with graduate assistants, and problems with the instructors. The subjects were asked to brainstorm solutions or recommendations for the school of business.

From the total set of brainstorming ideas, a pool of 110 brainstorming ideas was randomly selected to be the entire population for this research. The number of 110 was selected because it represented a reasonably large number of brainstorming ideas to work with while trying to stay manageable to the subjects.

From that population, random subsets of ideas were generated that contained 55 ideas. The number 55 was selected in that it is half of the 110 ideas. In the experiment, subjects were randomly assigned to either cluster the entire population, 110 ideas, by themselves or they were asked to cluster or group the partial set of ideas that they were given.

The partial sorts from the participants were then aggregated together using a weighted graph data structure. For each of the partial sorts, all brainstorming ideas in each group were processed to see which items were marked as being in the same group. The ideas in each group then became connected nodes in the graph. Each time a subsequent partial sort, from another subject, also had a relationship between the same two ideas (nodes), the vertex weight would increase by one.

Figure 3 illustrates the result from this process of creating the weighted graph from the list of partial sorts. The thicker lines between the nodes indicate that more people thought those ideas belonged together.
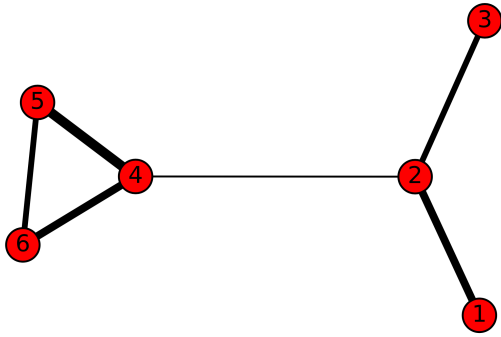
Figure 3: Weighted graph of brainstorming ideas

In this particular example, the weights between the nodes indicate that potentially there are two groups or buckets for these ideas. The first group would contain ideas four through six and the second group would contain ideas one through three.

Once all of the partial sorts were loaded into one consolidated graph, a final grouping needed to be conducted. However, due to the number of partial sorts used and the number of nodes, the weighted graph becomes much more complex and identifying how many buckets should be used and which ideas should be grouped together becomes much more complicated, see Figure 4.
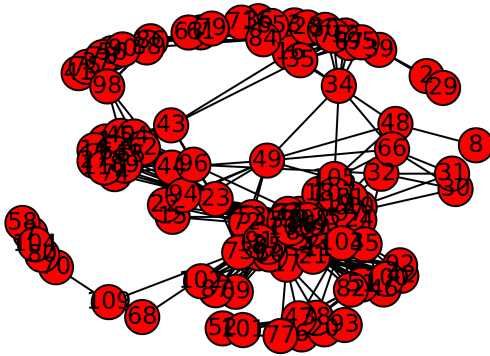

Figure 4: Weighted graph with all brainstorming ideas

In order to divide the graph and identify the groupings or buckets, a program called METIS was used [22]. METIS is an open source suite of tools developed by the University of Minnesota to partition graphs and meshes. The algorithms implemented are based on the multilevel graph partitioning paradigm. Using the weights between the nodes, the METIS program partitions the graph into distinct groups. Essentially, this crowdsourcing method of aggregating and

partitioning the graph yields a sort of group consensus regarding which brainstorming items belong together.

The final groupings or clusters from the METIS program were then compared to the sort created by a professional facilitator. The comparison was run using the Normalized Clustering Error (NCE) metric [18]. NCE is a metric that analyzes the associations between brainstorming ideas. NCE ranges from zero to one. Zero indicates that the aggregated clustering is an exact match with the professional facilitator cluster. A one indicates that there are no matches between the two clusters.

Formally, the NCE metric is defined as follows:

$$NCE = \frac{E}{A_t}$$

E represents the total number of incorrect and missing associations in the group cluster as compared with the facilitator cluster. The denominator, A_t, represents the total number of associations found in both the group's cluster and the facilitator's cluster. It is the upper bound as the summation of all associations.

4. Results

This section covers the assessment of how well the aggregation and partitioning of the graph worked. Twenty partial sorts, of the random 55 brainstorming ideas, were aggregated into one, weighted graph. The METIS program divided the graph into ten discrete partitions or groups. This final set of buckets were compared against the clustering performed by a professional facilitator using NCE.

Aggregating all twenty of the partial sorts yielded an NCE score of 0.68. This NCE score was then compared against the NCE scores of the subjects that sorted all 110 ideas. In this control treatment, each subject sorted all the ideas into buckets individually. The average NCE score for those in this control treatment was 0.74. The subjects in the treatment condition, that worked autonomously on clustering a random selection of half the brainstorming ideas, were able to produce a result that was better than the average score of the individuals that had to group all of the brainstorming ideas.

The next step in the analysis examined how the NCE scores would fluctuate as more partial sorts were added to the aggregate. The first weighted graph included six of the partial sorts that were randomly selected. After each NCE calculation, an additional, random sort was added to the aggregate. In the end, NCE scores were

calculated for six up to twenty partial sorts being used. Table 1 includes the results from this analysis. Again, comparing these sorts to the average NCE of those that sorted all individually, the crowdsourced approach performed commensurately well. Figure 5 shows the trend of NCE scores as the number of partial sorts are added as compared with the average NCE from the control group of 0.74. The regression from this result shows that there is a significant relationship between number of partial sorts used and the NCE score (β = -.011, $p < .01$).

Table 1: NCE scores by number of sorts used

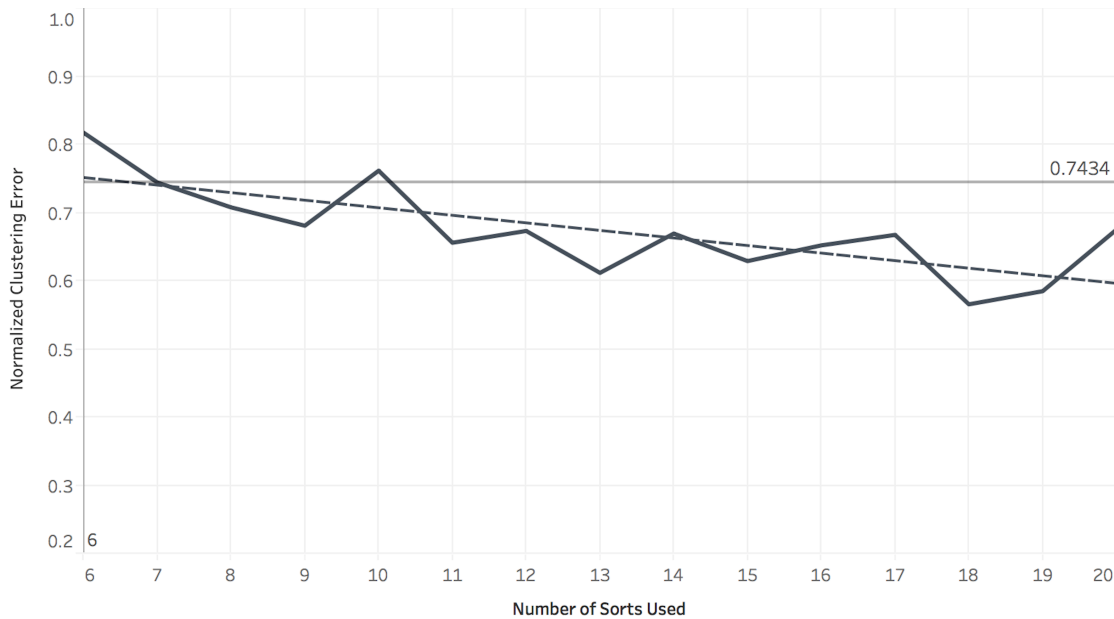| Number of Partial Sorts Used	NCE Score
6	0.82
7	0.74
8	0.71
9	0.68
10	0.76
11	0.66
12	0.67
13	0.61
14	0.67
15	0.63
16	0.65
17	0.67
18	0.57
19	0.59
20	0.68

Figure 5: NCE score as number of sorts aggregated increases

5. Conclusion

This paper examines a new approach to collaborative convergence that uses a crowdsourcing methodology. Individual participants are given a subset of ideas to organize into logical clusters or buckets. The system then takes these partial sorts and automatically aggregates them into one consolidated sort. The research question for this paper is whether or not this crowdsourcing can yield a viable option for groups that need to converge. Based on this analysis and the results, it is fair to say that this approach may indeed be a viable option.

This novel crowdsourcing approach produced one expected result; participants that only sorted half the brainstorming ideas were able to do so in less time than those that had to sort all 110 ideas. However, the aggregate time contributed by the partial sort participants was longer than those that sorted all the ideas on their own. Again, crowdsourcing relies on

leveraging smaller amounts of time per person but the gross effect is that more time is spent overall.

One key limitation of this study is that the METIS program, used to create the aggregated results, requires one significant variable to be set - the number of end buckets or clusters desired. Since this research is considered exploratory in nature, this variable was set to be the number of buckets identified by the expert facilitator. In subsequent research, different ways need to be investigated as to how to partition the weighted graph without knowing a priori how many buckets there should be. Research needs to examine a more organic method to identify the number of buckets, rather than METIS, or at least examine potential ways to derive this number using METIS.

Research needs to be conducted to determine the best way to use an approach like crowdsourced convergence. One possible alternative would be to have these crowdsourced buckets produce an initial sort from which more refinement of the buckets is conducted collectively by the group. In other words, convergence may start crowdsourced and then switch to more of a traditional convergence process later on in the workflow. It is conceivable that this approach would potentially yield a high-quality sort and still reduce the overall time spent sorting.

As this is just the first, exploratory paper, substantial limitations exist and considerable future research is needed. One major limitation of this study is that it is relevant only in contexts where a crowdsourcing approach is feasible. More specifically, there need to be a sufficient number of participants willing to contribute partial sorts in order to derive a successful aggregate sort.

Additionally, a more thorough examination of this approach is needed to determine if it can reduce the overhead and user satisfaction issues illustrated in Figure 1 in order to transform the process to one more like Figure 2.

While there are many effective and efficient methods for diverging with collaborative groups, there are fewer proven methods for converging, and most rob the group of the efficiencies and speed provided by parallel work. Ultimately, this approach makes an opening attempt at providing collaborators with new mechanisms for group work and crowdsourcing.

6. References

[1] Nunamaker Jr., J.F., Briggs, R.O., Mittleman, D.D., and Nunamaker Jr., J.F. Electronic Meeting Systems: Ten Years of Lessons Learned. In D. Coleman and R. Khanna, eds., *Groupware: Technology and Applications*. Prentice Hall PTR, Saddle River, NJ, 1995, 149–193.

[2] Nunamaker Jr., J.F., Vogel, D.R., Heminger, A., Martz, B., Grohowski, R., and McGoff, C. Experiences at IBM with group support systems: A field study. *Decision Support Systems 5*, 2 (1989), 183–196.

[3] Kolfschoten, G.L., Briggs, R.O., Appelman, J.H., and de G.J., V. ThinkLets as Building Blocks for Collaboration Processes: A Further Conceptualization. *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*, Los Alamitos: IEEE Computer Society Press (2005).

[4] Briggs, R.O., De Vreede, G.J., Nunamaker Jr., J.F., and Tobey, D. ThinkLets: achieving predictable, repeatable patterns of group interaction with group support systems (GSS). *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*, Computer Society Press (2001).

[5] Dennis, A.R. and Valacich, J.S. Computer brainstorms: More heads are better than one. *Journal of Applied Psychology 78*, 4 (1993), 531–537.

[6] de Vreede, G.J. and Briggs, R.O. Collaboration Engineering: Designing Repeatable Processes for High-Value Collaborative Tasks. *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*, (2005).

[7] Briggs, R.O., de Vreede, G.-J., and Nunamaker Jr., J.F. Collaboration engineering with ThinkLets to pursue sustained sucess with group support systems. *Journal of Management Information Systems 19*, 4 (2003), 31–64.

[8] Fjermestad, J. and Hiltz, S.R. An assessment of group support systems experiment research: Methodology and results. *Journal of Management Information Systems 15*, 3 (1998), 7.

[9] Briggs, R.O., Kolfschoten, G.L., Albrecht, C.C., Lukosch, S.G., de Vreede, G.-J., and Dean, D.L. A seven layer model of collaboration: Good technology is not enough. *30th International Conference on Information Systems*, (2009).

[10] Davis, A.J., De Vreede, G.J., and Briggs, R.O. Designing thinklets for convergence. *Association for Information Systems - 13th Americas Conference on Information Systems, AMCIS 2007: Reaching New Heights 1*, (2007), 25–36.

[11] Chen, H., Hsu, P., Orwig, R., Hoopes, L., and Nunamaker J. F., J. Automatic concept classification of text from electronic meetings. *Communications of the ACM 37*, 10 (1994), 56–73.

[12] Easton, G.K., George, J.F., Nunamaker Jr., J.F., and Pendergast, M.O. Using two different electronic meeting system tools for the same task: An experimental comparison. *Journal of Management Information Systems 7*, 1 (1990), 85–100.

[13] Santanen, E.L. and de Vreede, G.-J.J. *Creative Approaches to Measuring Creativity: Comparing the Effectiveness of Four Divergence thinkLets*. IEEE Computer Society, 2004.

[14] Griffith, T.L., Fuller, M.A., and Northcraft, G.B. Facilitator influence in group support systems: Intended and unitended effects. *Information Systems Research 9*, 1 (1998), 20–36.

[15] Gallupe, R.B., Dennis, A.R., Cooper, W.H., Valacich, J.S., Bastianutti, L.M., and Nunamaker Jr., J.F. Electronic Brainstorming and Group Size. *The Academy of Management Journal 35*, 2 (1992), 350–369.

[16] Chen, H., Ng, K., Martinez, J., Sutjahjo, S., and Nunamaker Jr., J.F. *An Artificial Intelligence Concept Generator Tool for Electronic Meetings*. 1993.

[17] Chen, H., Hsu, P., Orwig, R., Hoopes, L., and Nunamaker, J.F. Automatic concept classification of text from electronic meetings. *Communications of the ACM 37*, 10 (1994), 56–73.

[18] Roussinov, D.G. and Chen, H. Document clustering for electronic meetings: an experimental comparison of two techniques. *Decision Support Systems 27*, 1–2 (1999), 67–79.

[19] Diller, C.B.R. An Exploration of Distributed Parallel Sorting in GSS. 2013.

[20] Diller, C.B.R., Helquist, J.H., and Kruse, J. Improving Collaborative Convergence through Distributed and Parallel Sorting. *International Journal of e-Collaboration 12*, 2 (2016), 9–26.

[21] Mennecke, B.E. and Wheeler, B.C. Tasks matter: modeling group task processes in experimental CSCW research. (1993), 71–80 vol.4.

[22] Karypis, G. and Kumar, V. A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs. *SIAM Journal on Scientific Computing 20*, 1 (1998), 359–392.