# Branch-Explore-Merge: Facilitating Real-Time Revision Control in Co-Located Collaborative Visual Exploration


**Will McGrath,[1] Brian Bowman,[1] David McCallum,[2] Juan David Hincapié Ramos,[2] Niklas Elmqvist[1], Pourang Irani[2]**

[1]Purdue University
West Lafayette, IN, USA
{wbmcgrat, bowmanb, elm}@purdue.edu

[2]University of Manitoba
Winnipeg, MB, Canada
@umanitoba.ca



Figure 1: Our BEMVIEWER implementation of the branch-explore-merge protocol showing tabletop (public view) and tablet (private view). Two users are exploring a real estate dataset using a visual query formula as the shared state.


## ABSTRACT

Collaborative work is characterized by participants seamlessly transitioning from working together (coupled) to working alone (decoupled). Groupware should therefore facilitate smoothly varying coupling throughout the entire collaborative session. Towards achieving such transitions for collaborative exploration and search, we propose a protocol based on managing revisions for each collaborator exploring a dataset. The protocol allows participants to diverge from the shared analysis path (branch), study the data independently (explore), and then contribute back their findings onto the shared display (merge). We apply this concept to collaborative search in multidimensional data, and propose an implementation where the public view is a tabletop display and the private views are embedded in handheld tablets. We then use this implementation to perform a qualitative user study involving a real estate dataset. Results show that participants leverage the BEM protocol, spend significant time using their private views (40% to 80% of total task time), and apply public view changes for consultation with collaborators.




*ITS'12,* November 11–14, 2012, Cambridge, Massachusetts, USA.




**ACM Classification:** H.5.2 [Information Interfaces and Presentation]: User Interfaces—*Interaction styles*

**Keywords:** Public views, private views, collaborative visualization, coupling, computer-supported collaborative work.


## INTRODUCTION

In collaborative settings, multiple users working together routinely engage in both *coupled* (working together) and *decoupled* (working alone) activities to solve a problem at hand [2, 21]. The degree of coupling varies along a continuum where on one end, the participants are working together in lockstep on a task, and on the other end, they are solving it independently. These coupling transitions happen constantly and seamlessly in any collaboration [11, 12, 26, 28].

These findings are particularly interesting for the field of *visualization* (the use of interactive graphics to convey large, complex, and/or abstract data [5]), where collaboration has recently been named a grand challenge for research [27]. Of course, collaborative visualization must overcome the same challenges in supporting varying degrees of coupling as discussed above [8]. As a result, designers of collaborative visualization tools have considered various methods of integrating smooth transitions in coupling styles into the workflow. For example, coordinated multiple views [28] can be used to create a link between a private workspace for decoupled activity, and the public shared view for coupled collaboration. Another approach is to simply make the other participant aware of common work and similarities [9]. Other visualization systems have provided restricted access control to parts of the data to enforce a specific degree of coupling [13].

While such solutions can greatly enhance collaborative efforts, they are limited in a number of ways. Adding private views to a tabletop display consumes valuable screen real-estate and thereby counters the collaborative nature of this medium which is already taxed by displaying crucial information. More importantly, private views on shared displays have to be managed as the user moves around, and is even more challenging in the context of multiple users, each in need of a private view consuming space on the shared display. The overhead in managing coordinated views can be significant [28], suggesting that new styles of linking private and public views are necessary. Additionally, private views on a shared workspace are still visible by all users, thus limiting the degree of decoupling (i.e. users around the display can still interfere with actions on the private views of others).

To allow for seamless transitions between coupled and decoupled work styles, we propose an approach where these two forms of interaction are interlinked but distinct. Our approach, referred to as Branch-Explore-Merge (BEM) is inspired by solutions used for asynchronous software subversion systems, such as CVS, Subversion, or RCS. With BEM, participants can view and work together on the shared workspace, but can also branch away from the current state of the visualization, allowing them to explore the data separately and independently from other participants. Results obtained from working on their private decoupled workspace can then be merged back into the shared display. Branching allows participants to deviate from the shared state, while merging is an operation that changes the shared state of the visualization. To make a clear distinction between coupled and decoupled activity, to address real-estate constraints of private views on a shared display, and to alleviate concerns of combining multiple private views on the shared view [28], BEM offloads decoupled work and places it on private views on mobile devices such as smartphones or tablets, leaving the collaborative surface to be used exclusively for shared work.

We have implemented a multivariate visualization system for tabletop displays and multitouch tablets that support the BEM model. We demonstrate the capability of the system through a qualitative evaluation involving a collaborative analytic activity. In the study, participants are asked to explore a multidimensional dataset of real-estate properties using a scatterplot. They are assigned a distributed task for finding a home with specific pre-defined constraints split between participants. A map view supports progressive filtering through faceted browsing, where the collaborators can iteratively filter out items in the dataset using range selections for each attribute in the dataset. The resulting visual query consists of a conjunction of such set selections and forms the shared state that the BEM model operates on. Users can branch off from any state of the public display to explore the dataset without affecting the work of the entire group. Updates can then be merged back onto the shared space.

The study yielded several qualitative observations: (a) users tend to branch very early on in the collaborative process; (b) participants spend considerable time working in decoupled style on the tablets; (c) in scenarios where the results are ambiguous, i.e. no homes found for the collective constraints, users return to the table for more inquiry; and (d) participants almost always end the search process by discussing results in the shared space over the tabletop.

## BACKGROUND

Our work lies at the intersection of computer-supported cooperative work (CSCW), information visualization, and novel computing platforms. Below we review each of these areas and highlight the important messages for each of them.

### Computer-Supported Cooperative Work

CSCW research centers on theory, design, and implementation of *groupware*: software used concurrently by multiple users. The CSCW field has formally existed since the 1980s, and numerous groupware systems have been developed over the years. An exhaustive survey of this body of work is outside the scope of this paper, and we will instead describe only directly relevant research in the discussion below:

- **Coupling:** A group's *collaborative coupling style* has been defined as the mutual involvement of its participants in each other's work [26]. Research has shown that coupling ranges from fully coupled to fully uncoupled modes, and that groups transition between different degrees of coupling throughout a collaborative session [11, 26, 28].
- **Coordination:** Coordination mechanisms are used to effectively manage shared resources, and are critical for efficient collaboration and for minimizing the impact of interference [22, 26]. The multi-user coordination policies proposed by Morris et al. [18] are particularly useful here.
- **Territoriality:** People collaborating in the same physical space tend to automatically form *territories* that are used for personal work, shared work, or for storage [22]. However, short of systems that explicitly support *personalized views*, these territories are generally ad hoc in nature.

### Collaborative Visualization

Visualization is the graphical representation of data to aid human cognition [5], and collaborative visualization is accordingly the shared use of visualizations by multiple users (often called *analysts*). Studies have shown that involving multiple analysts in the analysis process generally improves the results (in quality or in time, or both). For example, Mark et al. [16] found significant benefits for collaborative visualization in both distributed and co-located settings. Balakrishnan et al. [3] show that a shared visualization significantly improves performance compared to alternative settings. Isenberg et al. [8] give a comprehensive overview of collaborative visualization and its connections to CSCW and HCI.

Several frameworks and toolkits for collaborative visualization exist. Scientific visualization has devoted much effort towards collaboration, mainly for distributed settings (see the survey by Brodlie et al. [4]). For information visualization, the Command Post of the Future [27] was one of the earliest examples. In summary, we take away several findings:

- **Flexible organization:** Isenberg et al. [11] study individuals, pairs, and triples in paper-based analysis tasks. Their results highlight the dynamic nature of collaborative visualization that demand highly flexible tools which support many different strategies for reaching a particular goal.

- **Common ground:** A key coordination task in collaborative visualization is to establish *common ground* [6], i.e., a shared understanding of the state of the analysis. Robinson [20] observed analysts working in pairs to complete an information synthesis task using paper artifacts and note that a significant time during the beginning of each session was spent establishing such a common understanding.

### Interactive Surfaces for Visualization

Novel input and output surface technologies are poised to make a significant impact on visualization research. We can derive the following messages from this topic:

- **Tabletop displays:** Digital tabletop displays have been shown to be particularly well-suited to collaborative information visualization; examples include Isenberg and Carpendale's tabletop visualization system for tree comparison [7], the Cambiera system [9] for face-to-face collaborative analysis of documents, and the Hugin toolkit [13] for mixed-presence visualization. However, none of these systems distinguish private and public views.
- **Public and private views:** The Lark system [28] extends the concept of coordinated multiple views (CMV) to multiuser collaboration. This practice is supported by general knowledge about territoriality on collaborative surfaces. However, Lark's private and public views are still virtual spaces that compete for space on the tabletop surface.
- **Multi-device environments:** Combining several colocated devices facilitates both solitary and collaborative work, such as in the i-LAND [24] system. ConnecTables [25] explicitly supports dynamic coupling through pen-based displays for both individual or group work. The UbiTable [23] is a "scrap display" where documents and laptops can be bridged using visual means. However, they all rely on social protocols [18] for coordinating the work.

### Collaborative Browsing, Filtering and Search

People working together on a common problem often have shared information needs, and so collaborative browsing and search is a common task in many collaborative sessions [17, 29]. With the exception of the Cambiera [9] system reviewed above, very few collaborative visualization systems are designed for collaborative search. The system perhaps most relevant to our work is Facet-Streams [12], where multiple participants use physical widgets on a tabletop display to construct visual queries. However, while the Facet-Streams system does support independent work, the private workspace of each participant is limited by the overall size of the tabletop. Regardless, we can learn the following:

- **Personal and shared devices:** Most systems for colocated collaborative search provide participants with their own device, often supplemented with a shared display [17]. This provides a natural physical partitioning for uncoupled and coupled work. For example, Maekawa et al. [15] present a collaborative web browsing system that partitions a webpage across several mobile devices. Paek et al. [19] propose a system for control of shared displays using PDAs, such as for web browsing, polls, and games. CoSearch [1] provides each user with a mobile phone as well as a shared PC display. Again, while supporting both private and public views, none of these existing projects explicitly tackle merging and revising collaborative state.

### Summary: Contributions

As can be seen from the above literature review, much work has been done on the intersection of these fields. We can summarize the unique contributions of our work as follows: (a) a novel model for clearly separating as well as merging coupled and decoupled activity around a shared dataset; (b) an implementation of this model for collaborative visual search using tablets as private views and a tabletop display as a public view; (c) improved understanding of how coupling patterns affect performance in a collaborative search task; and (e) initial qualitative results supporting the new model.

## BRANCH-EXPLORE-MERGE

The motivation for the Branch-Explore-Merge (BEM) protocol is to embrace the continually changing coupling in collaborative work to allow for participants to seamlessly move from closely to loosely coupled work. This is achieved by adopting a revision control mechanism for interactive visual exploration where the metaphor is that of the user's exploration taking a path that iteratively veers off, travels alone, and joins traffic. Below we discuss the BEM data model and the three protocol components (Figure 2).

### Model

Branch-Explore-Merge (BEM) assumes a collaborative application with a shared state $S$, a visual representation $V(S)$, and a set of interaction techniques $I$ that operate on the state to produce a modified state ($I(S) \rightarrow S'$). In a concrete implementation of the BEM protocol, the visual representation $V(S)$ would be rendered on a shared display that all participants can see (and typically interact with). Meanwhile, all participants also have a private display that shows the visual representation and supports interacting with the state.

Upon starting a BEM session, all private displays are *synchronized*, meaning that they are using the same shared state $S$, and the private displays will update as that state is changed.

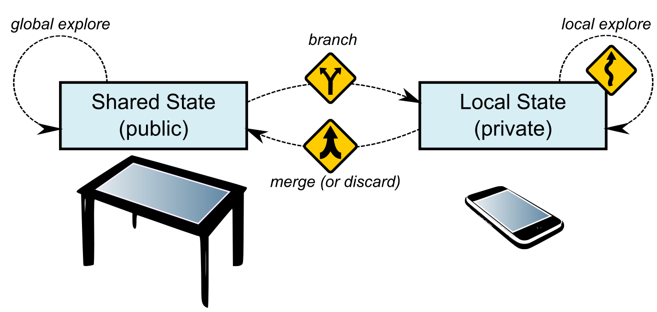


Figure 2: State diagram for branch-explore-merge.

### Branch Operation

Branching in the BEM protocol is a local (i.e., nonconflicting) operation that only affects the particular participant who initiated the branch. The result of a branch operation is simply to *desynchronize* the participant's private display from the shared display; in other words, the global shared state $S$ is copied to a local state $S_i$ for that participant.

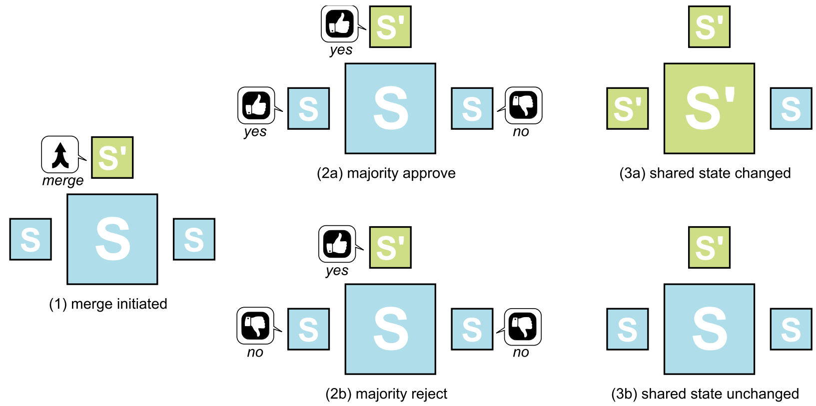


Figure 3: Example of two possible outcomes of a merge operation for three collaborators using a shared surface.

Branching is explicitly invoked by, e.g., clicking on a button, but an implementation may also allow implicit branching; for example, automatically branching (and desynchronizing) a participant who starts to modify the visual representation on their private view instead of on the shared view.

### Explore Operation

After having desynchronized from the shared state using a branch operation, the participant's private view will henceforth render the visual representation $V(S_i)$ instead of $V(S)$. Furthermore, any interaction performed on $S_i$ is specific to that view only. In other words, the participant is now free to independently explore data in uncoupled mode.

### Merge Operation

Merging is invoked by a participant with a desynchronized private view when the participant wants to add his or her results back to the shared state (alternatively, the participant can always revert back to the shared state if they decide that a branched exploration is not worth pursuing further). During the merge operation, the shared display is transformed into a *visual merge* which shows the proposed changes to the shared state. Designing the visual merge is an important aspect, but is specific to the particular visual representation. One useful approach may be to take inspiration from visual diffs for source revision control systems such as SVN or CVS; here, items who are to be removed are colored red, whereas items to be added are colored green. Invariant items that do not change have no color highlighting.

Unlike branching, merging is a global operation that may potentially cause conflicts with other participants as well as the shared state. For this reason, we need to introduce a conflict resolution mechanism to handle this situation.

There are several ways to resolve conflicts, ranging from pure social protocols [18] to system-supported mechanisms; examples of the latter include splitting the screen to show alternatives, allowing people to vote on a particular alternative, or simply overwriting shared state but maintaining an easily accessible history to revert back. Although the BEM protocol is not tied to any one mechanism, we will discuss voting here (Figure 3). The voting policy can vary on the application; simple majority is the most straightforward and useful one. Of course, even if a proposed visual merge is voted down, the participant who initiated the merge will retain that proposed state on his or her own private view. The participant can then choose to refine the state or persuade other participants to accept the change, or simply discard that state. Furthermore, for situations where a visual merge is accepted by simple majority, any naysayers may choose to automatically desynchronize from the new shared state and receive the old state on their private views to continue their work.

## IMPLEMENTATION

We have implemented the BEM protocol in a prototype collaborative visualization system called BEMVIEWER (Figure 1 and 4) for collaborative visual search in multidimensional datasets using geospatial displays. The system consists of a tabletop display, representing the public view, and individual tablets for collaborators, representing private views. Below we describe the details of the system.

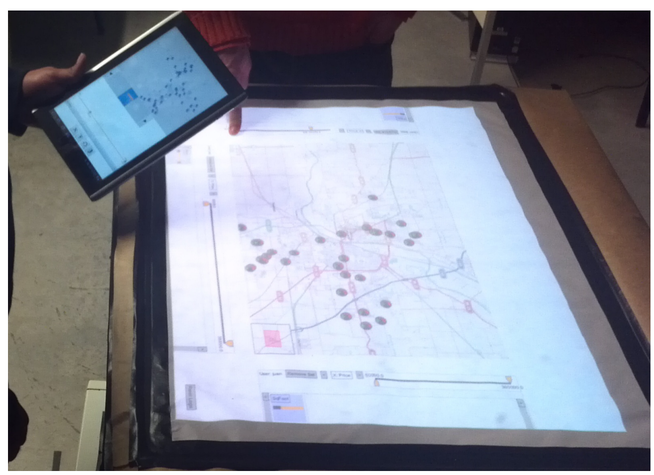

Figure 4: Our BEMViewer prototype setup.

### Visual and Interaction Design

The shared visual representation in the BEMViewer is a geographical map that can be panned and zoomed on both the tabletop display and the tablets using multitouch gestures such as pinch and drag. The data visualized on the map is represented by a multidimensional dataset where two of the dimensions specify a geospatial position using longitude and latitude. These dimensions are used to plot each item in the dataset on the map. Any geospatial dataset can be visualized using the tool, such as real estate, restaurants, hotels, etc.

### Query Formula

Because the BEMViewer is a tool for collaborative visual search, its shared state is actually **not** the dataset, but instead the *query formula* used to filter the dataset. Our query formula is a simple conjunction of terms, where a term can be any logical expression; similar to the classic Dynamic HomeFinder [30], our implementation supports simple intervals of the format $[\min, \max]$, i.e., which constrain item values for a given dimension to be within a particular interval.

Of course, changing the query formula for the BEMViewer geospatial visualization will cause houses to potentially appear and disappear on the map. For this reason, it is the query formula that is branched, explored, and merged using our BEM protocol. We implement the operations as follows:

- **Branch:** Copy the user's query formula to the user's tablet and detach them from the shared state.
- **Explore:** Allow the user to independently add, modify, or remove terms from the query formula on their private device without affecting the shared display.
- **Merge:** Upload the proposed changes to the query formula to the shared tabletop display.

In our implementation, we sidestep the need for an explicit vote upon merge by using *consensus icons* that convey the filters of all participants simultaneously (see below). For branching, there are several ways to implement the copying of shared state from tabletop to tablet; in the interest of simplifying the interface, we chose to copy only the user's own query formula to the tablet, but users were still able to copy the formula of any of their collaborators as well.

### Interactions

Upon joining a collaborative session, each participant gets assigned a *personal search panel* on the tabletop (Figure 1 shows three such panels, one per connected user). This search panel is also replicated on their tablet, i.e. their private view (Figure 1). The panel consists of a visual representation of the current query formula, buttons to branch, discard, and merge your personal state, and a simple filter interface. The filter interface is a drop-down menu where the user can choose which dimension in the dataset to operate on, and a dynamic query range slider [30] for defining an interval. The panel also allows for selecting, modifying, and deleting query terms in the query formula.

Interacting with the personal search panel on the shared tabletop will cause the global query formula to change for all users. This is implied in the fact that the panel is located on the shared surface. Interacting with the panel on the tablet, on the other hand, carries a more personal connotation and will thus only affect the local device.

### Consensus Icons

Our BEMViewer has two modes of managing the shared state: (1) a master query formula, as described in the BEM protocol, which is reached by consensus for all participants, or (2) individual queries, conveyed using consensus icons. A consensus icon is simply a visual glyph representing an item's filter status in the dataset where different parts of the glyph convey whether a given collaborator has filtered out the icon or not. Ideally, the different parts are oriented towards each user as well. Figure 5 shows this for four collaborators; here, the presence of a circle sector indicates that the item is within the query formula of a that user (i.e. not filtered).


(a) All agree. (b) Even split.

Figure 5: Consensus icons for four collaborators.

### Implementation Notes

We have implemented the tabletop display version of the BEMVIEWER as a Java application that uses the Piccolo2D structured vector graphics library. The map on the tabletop is implemented using the OpenStreetMap API. The application also provides a TCP/IP server port that accepts connections from BEMViewer clients on the network, lets them choose a seat (the system currently provides four sets, one on each side of the table), transmits the dataset to the client over the network, and then adds them to the current session.

The BEMVIEWER client is designed to be used on an Android tablet and uses the Android Canvas and Drawable APIs for graphical rendering. The client can connect to any BEMViewer server, support choosing a seat, and then provides the full set of interactions for collaboratively branching, exploring, and merging the exploration. The map view here is implemented using the Google Maps API.

## QUALITATIVE EVALUATION

A standard approach used in prior systems to evaluate collaborative search systems is to ask groups of two or more participants to solve a set of collaborative search tasks [1, 12, 17]. Assessment is generally conducted through a combination of observational and objective metrics (often using datasets with ground truth [10]). To stimulate discussion and consensus, participants may each receive slightly conflicting search criteria that requires negotiation [12].

We applied a similar methodology to gain insight about the usability and utility of BEM and to collect data about usage patterns of this model. In general, our study was inspired by the methods used to evaluate the Facet-Streams [12] system. We observed the use of BEM during a realistic collaborative search task using a geospatial real estate visualization. Group members were required to collaborate to negotiate on compromises we introduced in the study (compromises were controlled for the experiment). In particular, we wanted to gauge how participants would engage in collaborative search strategies on both an individual and group level.

### Participants

We recruited 4 groups of 3 participants each from an advanced HCI course at our university. All participants were experienced computer and touch interaction users.

### Task and Dataset

We chose a multidimensional dataset consisting of 150 real-estate properties organized into 8 dimensions: price, square footage, acreage, number of bedrooms, number of bathrooms, distance to university, distance to shopping center, and distance to factory (i.e., a workplace). We presented participants with the scenario wherein they were asked to purchase a house in which they would live in together. Each participant was given a different set of constraints. For example, one would be concerned about the price, the other about the size and distance to the factory. While each person's constraints were personal, we did not prevent them from sharing their constraints with other participants. However, participants were told to treat their own personal constraints as more important than that of other members in the triad.

The objective of the study was to solve the collaborative search tasks using our BEMViewer software on a tabletop and three tablets, one per participant. Participants were asked to complete 5 tasks of increasing difficulty. For the first task, participants would be able to locate a home that met all criteria. In the second task, several homes would match their constraints and participants were asked to choose one from the list. In the last three tasks, no one house matched all criteria. Participants were required to relax their constraints to satisfy the conditions. Table 1 shows all five tasks.

| # | User 1 | User 2 | User 3 |
|---|---|---|---|
| 1 | BR $> 3$<br>Univ $< 4.5$ mi | SqFeet $< 1900$ | Price $<$ \$175k |
| 2 | Price $>$ \$100k<br>Price $<$ \$200k | BR $\geq 3$<br>Acreage $> 2$ | Shops $< 3$ mi<br>Factory $< 4$ mi |
| 3 | Price $<$ \$240k<br>Acreage $> 3$ | Shops $< 3$ mi<br>Univ $< 3$ mi | SqFeet $> 1900$ |
| 4 | $2 \leq$ Bedrooms $\leq 6$<br>Acreage $\leq 1$<br>Price $<$ \$105k | SqFeet $> 450$<br>Shops $< 3.5$ mi | Univ $< 4.5$mi<br>Factory $< 2$ mi |
| 5 | Less than \$150k, maximize bedrooms for cost ($\geq 2$). | | |

Table 1: The real estate search tasks given to users in our qualitative evaluation on collaborative search.

Solving a task implied arriving at a consensus on which house to purchase. For the tasks where no perfect house existed, the participants were asked to negotiate amongst themselves and find the house that most closely matched their collective constraints. We informed participants that we would weigh each of the dataset dimensions equally.

### Performance Metrics

Our study was qualitative and observational. As indicated earlier, some of the more salient questions we were investigating concerned such things as "do users comprehend and use the BEM model for solving such tasks?", "If the model is used, how early/late do users branch off to do decoupled work?", "when users merge their results, do they continue to branch off again to filter further, or is further filtering performed on the tabletop?", etc. We used concession distance to measure how closely their results matched ground truth.

*Concession distance* is the normalized distance between the house selected by the group and the subset $S$ of the search space defined by their collective constraints. For simplicity, we use the rectilinear distance between the search subset $S = [s(i)_{min}; s(i)_{max}]$ and the house $h$ in all dimensions $i$ of the dataset, normalizing using the range $D(i)$ of each dimension. We designed most trials so that a concession distance of 0 (all constraints fulfilled) is not possible. In other words, this would model situations where the search subset $S$ contains no house, but where the house that is closest to $S$ is optimal (we assume all dimensions have equal weight).

Note that our concession distance method is different from the discrete concession/fail metric used by Jetter et al. [12], but we think it better captures typical search behavior.

### Equipment

We evaluated the systems using a 40-inch FTIR tabletop display equipped with a projector, having a resolution of $1280 \times 800$ pixels. For the mobile devices, we used three Acer Iconia tablets running the Android 3.0 operating system. The tablets communicated wirelessly with the server. The experimental software used for both tabletop and tablets was the BEMViewer prototype described previously.

### Experimental Design

Given that our approach was qualitative and designed to understand patterns of coupling better, we asked groups of participants to do all the tasks using the full BEMViewer system. Group dynamics can change over time as people get to know each other and this could influence the outcomes and performance. Therefore, we paid attention to how these patterns change over the length of the study. The outcomes were observed as discussed above.

### Procedure

Each session consisted of a group of three participants (a *triad*). Upon arriving in our lab space for an experimental session, participants were asked to read and sign a consent form which broadly outlined the study. The experimenter gave a brief background, answered questions, and allowed participants to make introductions with each other.

The group first performed a test task to get them familiar with the interface and the tasks. The experimenter explained and demonstrated how to use the interface to filter in a sample real estate dataset. Tools for filtering, adding filters, and removing filters were demonstrated to the participants. Participants were also taught how to branch off from the shared display, how to perform queries on their mobile device, and how to discard or merge their findings with the shared display. Participants were allowed to practice until they felt comfortable. Questions about the functionality of the tool were allowed throughout the trials.

All search trials began with the BEM tool initialized with no filters active and all points representing houses visible on the shared geographic map. Participants were given a sheet of paper listing their individual constraints for each task and were allowed to commence the collaborative search task. We did not set a time limit as we were not concerned with efficiency with the task. Instead participants were allowed as much time as possible to complete the tasks. A trial ended when the group reached consensus on a target home.

After completing all tasks, each participant filled out a personal questionnaire about their subjective assessment of the system. Each session lasted about 60 minutes in total, including filling out the post-test questionnaires.

## RESULTS

Here we describe the results in terms of observations, branch-explore-merge patterns, and the concession distances.

### General Observations

All group sessions were characterized by lively interaction, both verbal and gestural (pointing to the tabletop or to tablets), as well as common and consistent switching between the shared tabletop and the private tablets. The transitions were quick and often seamless. Some participants would also use their tablet to show certain aspects of the dataset to other participants. This seems to confirm that collaborative coupling was indeed smoothly transitioning from decoupled to coupled states during the collaboration [26].

We did note that different participants attributed different weights to various dimensions; for example, most participants took price criteria as being serious, whereas dimensions such as distance to work or square footage seemed to have varying importance to participants. A common method of searching was for participants to use the tablets to configure filters, merge their changes, and then look together on the tabletop for houses matching the global criteria.

### BEM Patterns

To further understand how users engage in collaborative practices in a search task, we developed a visualization technique, BEMViz, inspired by visuals such as TimeNets [14] and Minard's map of Napoleon's March. BEMViz displays the branching and merging patterns using a timeline with a specific line for each participant involved in a collaborative session. Color is used to identify each user and line thickness indicates the percentage of currently visible items. The diagram has a baseline, which is set to a central filter state. The distance of a user's line from the baseline is an indication of how many items are in common; if all the user's items exist on the central entity, the lines are overlaid.

We utilize BEMViz with the tabletop as the ground truth, shown in dark grey (Figure 8). As users filter on their tablets, their lines diverge from the tabletop. The lines merge with the tabletop when users interact with their personal tabletop controls, or when they upload filters from the tablet. The line thickness shows the percentage of visible houses for each device. These are the houses that match the criteria in the relevant filters. A change in line thickness indicates that a filter was modified, created or deleted. The line thickness of the ground truth also changes when users merge their filters.

Our results show that the branch-explore-merge protocol seemed straightforward for users to adopt in the context of both public and private views; the ratio of device use for tablet vs. tabletop ranged between 40:60% to 80:20% of the total task time. Analysis of the BEM patterns (see below) clearly indicate that participants often transitioned between coupled and decoupled work when solving a task, which would seem to indicate that the protocol facilitated collaborative search. Whether this translates to a higher-quality solution is difficult to say (see below for concession distance results), but we think that this at least is an indication that the collaborative process is facilitated by the BEM protocol.

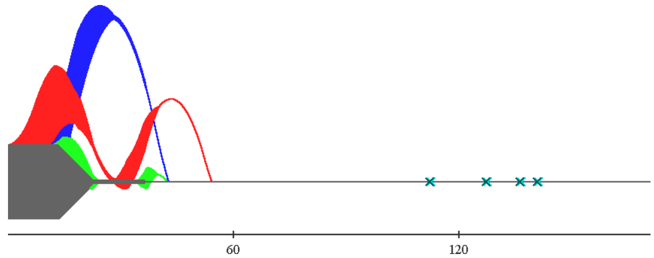


Figure 6: Example of early branching.

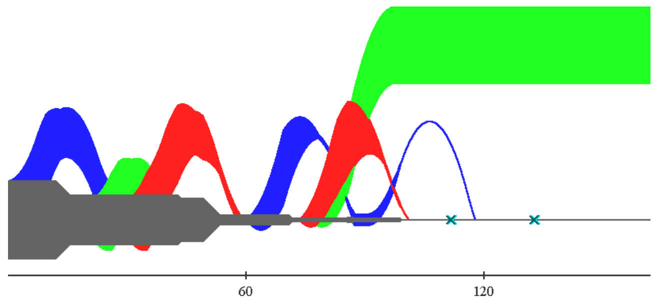


Figure 7: Green user branched off but did not merge.

**Early branching.** In about 80% of the tasks, participants would branch off to work on their tablet interface within the first 10% of the total task time. Figure 6) shows two cases where such early branching takes place. Early branching can be interpreted to suggest that: (a) users are more comfortable querying on the tablet interface; (b) decoupled work for collaborative search takes place very early on; and (c) users may wish to not overwrite the global state of the search, and instead perform local searches until they feel comfortable with their personal results. This outcome suggests that the BEM model's ability to decouple is useful for collaborative search.

**Abandoning the merge.** Our observations suggest that in some cases users merge, but then abandon their merged state without refreshing the global state. This also happens much later in the collaborative process, often nearly toward the end (Figure 7). This seems to also happen when the other active users are mostly operating on the shared display and not on their tablets. One reason for this might be that users may have reached a consensus without needing the local search results of the user that has abandoned the merge operation. Another explanation might be that the user gets involved with refinement on the tabletop and does not need to branch to resolve the query, essentially letting their private view languish. This usage model is well supported in BEM as it provides sufficient flexibility to move seamlessly between devices.

**Resolving conflicts.** To resolve conflicts, all participants invariably discussed potential solutions on the shared display. In over 90% of all tasks, the final decision took place after adjusting filters on the tabletop (Figure 8). Interestingly, for complex tasks we noticed a number of repeated branches and

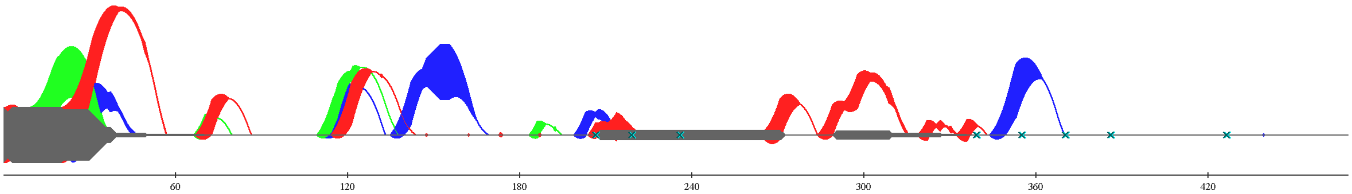


Figure 8: Patterns showing conflict resolution in several branches and merges for two participants (blue and red).

merges among participants. Usually one participant would take an active role in such activity. This was particularly the case with tasks 3 and 4 where the resulting set could lead to a large number of potential solutions and users would need to filter repeatedly to arrive at a satisfactory solution.

In the same pattern above, we notice a number of branches that take place simultaneously, i.e. that all three users have branched off. However, later in the same task, branching happens sequentially, where one user branches off, filters the dataset, and the other user picks up from the filtered dataset. Such patterns emerged in complex tasks where all users were relaxing their constraints, either together or in sequence.

**Minimal concession.** We also observed cases wherein a participant would not relax their constraints in response to other participants. Instead they would rely on the other users to relax their filter options before making changes. They would also view the outcome of the collective results after seeing a user modify their query set but would not necessarily branch off (the blue user in the figure below).

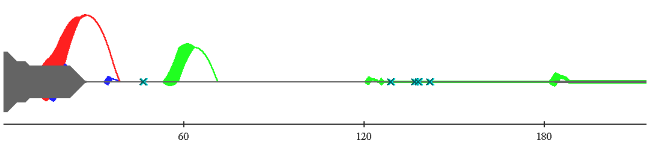


Figure 9: The blue user is not relaxing their constraints despite the task requiring mutual concessions.

### Concession Distance

The purpose of our experiment was not to compare the results to some ground truth of what would be the best tradeoff, particularly since people tend to weigh different dimensions differently. Instead, we use the average concession distance metric as an informal indication of performance. For example, an average concession distance of .059 means that the group ended up relaxing their given constraints by 5.9% when determining an answer. Our data shows that the collaborative search tasks were fairly accurate; on average across all tasks, the participants only had to relax their constraints by approximately 2% in finding an answer to the tasks.

## DISCUSSION

The results from our qualitative evaluation clearly show the utility of the branch-explore-merge protocol. Participants in our study split their time between tabletop and tablet, and they frequently branch, explore, and merge results back to the shared state. We found that the quality overall was very high; solutions were within less than 2% of the search parameters.

### Generalizations

Ample research has shown that collaborative coupling is a continuous scale ranging from uncoupled to tightly coupled work, yet it can be argued that our BEM system only explicitly supports two discrete points along this scale—the two extremes. However, while it is true that our public and private views are implemented using two very different types of devices, we observed that users would freely share the content with their tablets to other participants while solving the common task. In other words, even a private device such as a tablet can be used in a collaboratively coupled fashion.

Furthermore, the intention of our collaborative protocol is not to **limit** the users, but to **empower** them. Thus, even if our protocol has certain stipulations on the nature of uncoupled and coupled work, and even if our BEMViewer implementation strictly used physical devices for private views, this is not to say that we think a real-world collaborative application need be this literal. Our intention in this paper was to control the collaborative setting to study certain phenomena related to collaborative coupling, but these controls need not stay in place for a real system. For example, we can easily see such a system supporting "hybrid" private views that still are rendered on the shared display, but which the user can easily bring onto a physical device with a simple operation.

Our work has so far considered only co-located and synchronous collaborative sessions, and the BEMViewer supports only up to four participants. However, there is no reason to believe that—given appropriate visual design—the protocol will not support more than four participants, or that it will not be just as suitable for distributed, mixed-presence, and/or asynchronous collaborative modes. Furthermore, our focus in this paper has been on collaborative search, whereas visual exploration can include many other tasks such as sensemaking, decision-making, and planning. Exploring these parameters is left for future work, however.

We used concession distance to evaluate our results from the user study, summing up the normalized distances for each dimension between a chosen house and the boundaries of the search space given in the constraints. However, this model assumes that all dimensions have equal weight, but our observations from the experimental sessions clearly showed that this is not the case; for example, people tend to be particularly sensitive about money (spending less is often perceived as better). Therefore, it is likely that such biases played a role in choosing what was perceived to be the closest house given a search space where no "perfect" house existed. However, we speculate that the BEM protocol may even have helped participants in identifying and discussing these biases.

Another interesting observation is that device preference given the choice of two devices—tablet vs. tabletop—is a very fickle phenomenon. In early pilot testing, our tablet interface was not fully developed and as aesthetically pleasing as in the current version, and the result was that participants preferred the tabletop interface and basically did not interact with the tablets at all. However, in the evaluation, as can be seen from the data, participants preferred the tablet, most likely because of the higher relative screen resolution (DPI) as well as the more responsive touch detection. This leads us to speculate that small differences (or imbalances) in interaction design between devices in a multi-device environment may have a large impact on the device affinity of the users. This is an interesting claim for future studies of such systems.

Finally, this evaluation was primarily qualitative in nature, and in particular, we did not perform a controlled comparison with a baseline condition. Conceivable baseline comparisons would be a tabletop only, or even a desktop computer interface with a single mouse and keyboard. However, the BEM protocol provides new affordances that simply did not exist previously, making such comparisons potentially lopsided. Instead, our qualitative evaluation provides more useful findings on the protocol's utility for collaborative search.

### Beyond Maps

The visualization components in our BEMViewer prototype are relatively simple: a geospatial map with some graphical embellishments. This was done to let us study the effectiveness of the BEM protocol without visual complexity getting in the way of users understanding the visualizations.

However, there is nothing in the branch-explore-emerge protocol that limits the nature of the visual mapping. The only requirement is the capability to, given two visualization states, render their difference as a visual difference, and then supporting the merge operation. This leaves the space open for using virtually any visualization design in the protocol. For example, a social network visualization tool may support expanding and collapsing the node-link diagram as filtering operations, and a visual difference would highlight the topological changes to the network that a merge would cause. Or, an intelligence analysis tool could be adapted for collaborative work by placing a common report on the shared display and allowing participants to contribute to this report.

As a case in point, we have also implemented a slight variation of the BEMViewer that uses 2D scatterplots with dynamically controllable axis mappings as the main visual representation (Figure 10). This tool can be used for general visual exploration in multidimensional datasets for situations when the data does not have an intrinsic geospatial position. Unlike the map-based version, the scatterplot BEMViewer has an explicit vote protocol and uses a simple visual merge operation where nodes to be added by a merge are highlighted in green, and nodes to be removed are highlighted in red.

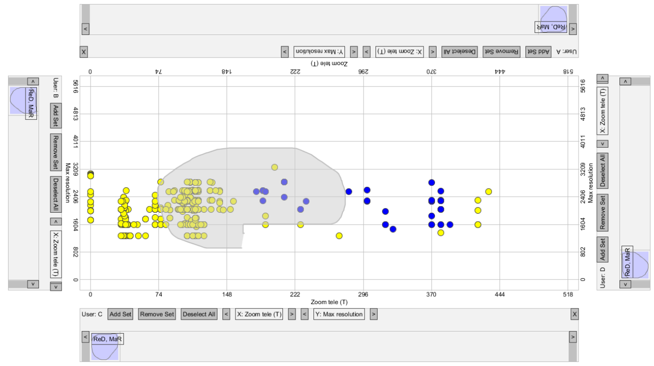

Figure 10: BEMViewer for multidimensional data.

## CONCLUSIONS AND FUTURE WORK

We have proposed Branch-Merge-Explore, a new collaborative protocol for explicitly supporting varying degrees of coupling styles in co-located collaborative visualization settings. Our approach incorporates both public and private views and allows participants to temporarily desynchronize their private views from the main visualization state to perform independent work. Synchronizing the new state to the shared and public display consists of a visual merge and a common vote among all participants. An informal evaluation compared search behavior for the branch-explore-merge protocol using a prototype implementation using a tabletop display as a public view and individual tablets as private views. Results indicate that the protocol appeared to facilitate collaboration, particularly where no perfect result existed.

Our future work will consist of implementing BEM in the context of other visual representations, interactions, and datasets. In general, we think that there is tremendous potential in pairing mobile devices with large displays, be it tabletops or wall-mounted ones, and we are working on software toolkits for supporting such ubiquitous visualization spaces.

## ACKNOWLEDGMENTS

This work is partly supported by the U.S. National Science Foundation grants IIS-1249229 and TUES-1123108. We thank the user study participants for their time and effort.

## REFERENCES

1. S. Amershi and M. R. Morris. CoSearch: a system for co-located collaborative web search. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 1647–1656, 2008.
2. K. Baker, S. Greenberg, and C. Gutwin. Empirical development of a heuristic evaluation methodology for shared workspace groupware. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, 96–105, 2002.
3. A. Balakrishnan, S. Fussell, and S. Kiesler. Do visualizations improve synchronous remote collaboration? In *Proceedings ACM Conference on Human Factors in Computing Systems*, 1227–1236, 2008.
4. K. W. Brodlie, D. A. Duce, J. R. Gallop, J. P. R. B. Walton, and J. D. Wood. Distributed and collaborative visualization. *Computer Graphics Forum*, 23(2):223–251, 2004.
5. S. K. Card, J. D. Mackinlay, and B. Shneiderman, editors. *Readings in Information Visualization: Using Vision to Think*. Morgan Kaufmann Publishers, San Francisco, 1999.

6. H. H. Clark and S. E. Brennan. Grounding in communication. In L. B. Resnick, J. M. Levine, and S. D. Teasley, editors, *Perspectives on Socially Shared Cognition*, 127–149. Amer. Psych. Assoc., 1991.
7. P. Isenberg and S. Carpendale. Interactive tree comparison for co-located collaborative information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 13(6):1232–1239, 2007.
8. P. Isenberg, N. Elmqvist, J. Scholtz, D. Cernea, K.-L. Ma, and H. Hagen. Collaborative visualization: definition, challenges, and research agenda. *Information Visualization*, 10(4):310–326, 2011.
9. P. Isenberg and D. Fisher. Collaborative brushing and linking for co-located visual analytics of document collections. *Computer Graphics Forum*, 28(3):1031–1038, 2009.
10. P. Isenberg, D. Fisher, M. R. Morris, K. Inkpen, and M. Czerwinski. An exploratory study of co-located collaborative visual analytics around a tabletop display. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, 179–186, 2010.
11. P. Isenberg, A. Tang, and M. S. T. Carpendale. An exploratory study of visual information analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 1217–1226, 2008.
12. H.-C. Jetter, J. Gerken, M. Zöllner, H. Reiterer, and N. Milic-Frayling. Materializing the query with facet-streams – a hybrid surface for collaborative search on tabletops. In *Proceedings of ACM Conference on Human Factors in Computing Systems*, 3013–3022, 2011.
13. K. Kim, W. Javed, C. Williams, N. Elmqvist, and P. Irani. Hugin: A framework for awareness and coordination in mixed-presence collaborative information visualization. In *Proceedings of the ACM Conference on Interactive Tabletops and Surfaces*, 231–240, 2010.
14. N. W. Kim, S. K. Card, and J. Heer. Tracing genealogical data with TimeNets. In *Proceedings of ACM Conference on Advanced Visual Interfaces*, 241–248, 2010.
15. T. Maekawa, T. Hara, and S. Nishio. A collaborative web browsing system for multiple mobile users. In *Proceedings of the IEEE Conference on Pervasive Computing and Communications*, 22–35, 2006.
16. G. Mark, A. Kobsa, and V. M. González. Do four eyes see better than two? Collaborative versus individual discovery in data visualization systems. In *Proceedings of the International Conference on Information Visualization*, 249–255, 2002.
17. M. R. Morris, J. Lombardo, and D. Wigdor. WeSearch: supporting collaborative search and sensemaking on a tabletop display. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 401–410, 2010.
18. M. R. Morris, K. Ryall, C. Shen, C. Forlines, and F. Vernier. Beyond 'social protocols': multi-user coordination policies for co-located groupware. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 262–265, 2004.
19. T. Paek, M. Agrawala, S. Basu, S. Drucker, T. Kristjansson, R. Logan, K. Toyama, and A. Wilson. Toward universal mobile interaction for shared displays. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work*, 266–269, 2004.
20. A. C. Robinson. Collaborative synthesis of visual analytic results. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, 67–74, 2008.
21. K. Schmidt and C. Simone. Coordination mechanisms: Towards a conceptual foundation of CSCW systems design. *Computer Supported Cooperative Work*, 5(2/3):155–200, 1996.
22. S. D. Scott, M. S. T. Carpendale, and K. M. Inkpen. Territoriality in collaborative tabletop workspaces. In *Proceedings of the ACM Conference on Computer Supported Cooperative Work*, 294–303, 2004.
23. C. Shen, K. Everitt, and K. Ryall. UbiTable: impromptu face-to-face collaboration on horizontal interactive surfaces. In *Proceedings of the International Conference on Ubiquitous Computing*, volume 2864 of *LNCS*, 281–288. Springer, 2003.
24. N. A. Streitz, J. Geissler, T. Holmer, S. Konomi, C. Müller-Tomfelde, W. Reischl, P. Rexroth, P. Seitz, and R. Steinmetz. i-LAND: An interactive landscape for creativity and innovation. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 120–127, 1999.
25. P. Tandler, T. Prante, C. Müller-Tomfelde, N. Streitz, and R. Steinmetz. ConnecTables: Dynamic coupling of displays for the flexible creation of shared workspaces. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, 11–20, 2001.
26. A. Tang, M. Tory, B. Po, P. Neumann, and S. Carpendale. Collaborative coupling over tabletop displays. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, 1181–1190, 2006.
27. J. J. Thomas and K. A. Cook, editors. *Illuminating the Path: The Research and Development Agenda for Visual Analytics*. IEEE Computer Society Press, 2005.
28. M. Tobiasz, P. Isenberg, and S. Carpendale. Lark: Coordinating co-located collaboration with information visualization. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):1065–1072, 2009.
29. M. Twidale, D. M. Nichols, and C. D. Paice. Browsing is a collaborative process. *Information Processing and Management*, 33(6):761–783, 1997.
30. C. Williamson and B. Shneiderman. The dynamic HomeFinder: Evaluating dynamic queries in a real-estate information exploration system. In *Proceedings of the ACM SIGIR Conference on Research and Development in Information Retrieval*, 338–346, 1992.