Enabling Deep Conceptual Learning in Computing Courses through Conflict-based Collaborative Learning

Swaroop Joshi and Neelam Soundarajan
Dept. of Computer Science & Engineering
The Ohio State University
Columbus, Ohio 43210. USA
Email: {joshis, neelam}@cse.ohio-state.edu

***Abstract*— Piaget's classic work on cognitive development showed that engaging learners in critical discussions with peers about ideas that are *different* than theirs leads to deep conceptual understanding. Implementing such an approach in computer science and, more generally, STEM, courses has some specific challenges. Based on Piaget's theory, we have developed a highly innovative collaborative learning approach that exploits specific affordances of web technologies to address these challenges. It allows small groups of students with different ideas about the topic in question, to engage in a highly-structured discussion that enables each student in the group to develop deep understanding. While a number of researchers have explored approaches to collaborative learning, a key difference with our work is that our focus is helping individual students develop deep understanding, whereas the focus of much of this other work is on developing students' team skills, effective communication abilities, and the like. We have tested our approach in an undergraduate CS course and the effect on student performance was encouraging; moreover, about two-thirds of the students in the course, based on a post-activity survey, felt that the approach was effective in helping them develop conceptual understanding.**

I. INTRODUCTION

One of the key ideas in Piaget's classic theory [1] on the cognitive development of children is that of *equilibration*. Equilibration is the process by which children resolve conflicts between their current internal pictures (Piaget's "schema") and new information that may be presented to them. Piaget's work as well as that of later researchers (e.g., [2]) showed that this type of conflict resolution is especially effective in promoting learning when the conflicting information comes not from someone like a teacher but, rather, from *peers*. This is because if the conflicting information were to come from a teacher – or someone similar– the learner is likely, given the authority of the teacher, to simply accept the new information without critical analysis; by contrast, if the conflicting information is from a peer, the learner is far more likely to analyze the new information as well as her own current understanding critically since, for all the learner knows, the peer may be *wrong* and she may be right! Or, indeed, the correct concept may be something that combines hers as well as the peer's position.

The basic thesis underlying our work is that an approach, which is based on identifying different conceptions of key ideas among small groups of students in computer science (more generally, STEM) courses, followed by having the students engage with each other in a focused and structured discussion to critically analyze these different conceptions, will enable each student in the group to develop a deeper understanding of the concepts. A second key idea underlying our work is that such discussions are best mediated by a carefully designed online system that ensures that the students in the group do, in fact, engage with each others' ideas, that all students participate effectively in the discussion, and that the discussion is not compromised by such things as students' stereotypical biases concerning other students' knowledge or abilities.

The idea of *collaborative learning* in college courses is hardly new. Much of the work in this area, however, has focused on *team success* rather than ensuring that *individual* students in the group develop as deep an understanding of the topic as possible. Thus the work of Johnson et al. [3] emphasizes such factors as *positive interdependence* for the success of the team; positive interdependence is the 'swim together, sink together' feeling wherein team members feel that their success is dependent on other team-members' success. Michaelsen and Sweet [4] talk about a number of similar factors that they argue are key to successful "team-based learning". We discuss some of this work in greater detail in Section II. For now, the key point to reiterate is that the focus of our approach is to ensure that individual students in each group develop as deep an understanding of key concepts as possible with collaborative learning among small groups of students, with each group consisting of students who have *different conceptions of the topic in question*, serving as a *means* of achieving that main goal.

Many course instructors, of course, engage students in in-class discussions concerning subtle aspects of the topic under discussion. While such discussions do help, they pose a number of problems as well. First, since many students may be seeing the topic for the first time, they are likely not to have developed sufficient understanding of the topic to be able to appreciate the subtleties, let alone engage in effective discussions concerning them. Second, some students find it easy to express opinions in a classroom while others, even though they may have a better understanding of the issues, tend to remain silent. Third, some students may be quick on their feet and come up with –possibly incorrect!– answers while others, possibly more thoughful and deliberate, need to time to formulate their answers by which time the class, given the pressure that instructors are under to ensure that all topics receive adequate coverage, may have moved on to other topics. Last and by no means least, some students may harbor preconceived biases or stereotypical views concerning the abilities of other students which tends to seriously compromise the effectiveness of such discussions.

We have developed a highly innovative approach and system that exploits specific affordances of web technologies to address these and a number of related problems. It allows *small* groups of students with *different* ideas about the topic in question, to engage in *highly-structured* discussions that enable each student in the group to develop deep conceptual understanding. The approach effectively addresses each of the problems listed in the last paragraph and has numerous other advantages as well. (See Sec. II and III for more details.)

Some key features of our system, explained in detail in Sec. III, are: (1) asynchronous, rounds-based structure for the discussion in each group; (2) anonymous posting by students in the group; and (3) the requirement that, in each post, the student explicitly specify how his/her position relates to that of each of the other students in the group, as explained in their posts. The first feature ensures that every student gets enough time and opportunity to deliberate on and respond to group members' posts and any *flame wars*, etc. are avoided. The second mitigates any preconceived notions the group members may have about each other, thus allowing a free and open discussion. The last feature is designed to have students engage in a careful analysis of others' ideas, and, in turn, sharpen their own. Then in Sec. IV, we discuss the details of our use of the system in an undergraduate CS course, and present the results in Sec. V. We conclude with a brief summary of our work and some directions for future work (Sec. VI).

II. Related Work

Piaget's theory has influenced learning scientists throughout the late twentieth century, and many researchers have employed the ideas of inducing a cognitive conflict to help students learn in variety of settings. For example, Doise and Mugny [2] focused on children and demonstrated that individual learning can be aided by exposing the learner to conflicting ideas from peers. They made an important observation that it was not necessary for the peers with conflicting ideas to actually, physically interact with each other, as long as it was *perceived* that the conflicting viewpoint was that of a peer.

One of the distinctions between the works of Piaget and Vygotsky [5] is that the latter stresses the importance of a "more knowledgeable other" (MKO) in the interaction. In other words, according to Vygotsky, interaction among peers is most fruitful when one of the members of the group is more knowledgeable than the others. Interestingly, while some researchers suggest the importance of Vygotsky's MKO, the results of other researchers suggest that what seems to matter most is the cognitive conflict that a student experiences because of disagreements with other students' conception of the same problem or topic [2, 6, 7]. The focus of our work is interactions among *peers* based on cognitive conflict in the group.

Next we review some collaborative learning techniques used widely at school as well as college level, including in engineering. In Jigsaw [8], each student is placed in a *home* group and in an *expert* group. Each student in a home group is assigned a distinct topic. Students leave their home groups and join other students with the same assigned topic, forming the expert group on the topic. They explore their topic thoroughly and then return to their home groups; the student is then responsible for teaching his or her home group the particular topic. In the Think-Pair-Share (TPS) approach [9], the instructor poses a conceptual question and asks students to think individually about their responses. Then the students pair up with a neighbor and discuss each others' responses. Finally, the instructor calls on some students to share their answers with the entire class.

In team-based-learning (TBL) [4], students are organized into teams of five or six each, and remain in teams throughout the course. The course is organized into units, each two to three weeks long. Before the start of a unit, students are assigned readings. On the first day of the topic, students complete, as individuals, a short test on the topic. Immediately after, they take the same test as teams, coming to consensus on answers. The final step is a short lecture by the instructor focusing on common problems shared by many teams. The rest of the two to three week period is spent on activities that require the teams to apply the concepts and techniques to increasingly challenging problems.

These three approaches do not necessarily use cognitive conflict to group students, but it plays a more central role in peer-instruction (PI) [6, 10]. In PI, each student individually answers a conceptual multiple choice question in class, submitting the answer via a *clicker* or other similar device; then the students turn to their neighbors and, in groups of 3 or 4, discuss the question; after a few minutes of discussion, each student again answers the same question. During the discussion time, the instructor walks around the room, just observing the discussions, instead of participating in them. Crouch and Mazur [10] report that the percentage of students who, following discussion with their peers, change their answer from a wrong choice to the correct one far exceeds the percentage who change from the correct choice to a wrong one.

There are a number of limitations with all these approaches, mostly related to the fact that these are in-class techniques:

- In PI, since the multiple-choice question is about the topic discussed in the lecture, students may not have had enough time to think about it deeply;
- The groups, pairs, teams, etc. are formed randomly (TBL, Jigsaw), or mainly based on which students happen to be seated next to which other students (TPS), rather than on the basis of ensuring cognitive conflict in each group;

- Some students may dominate their groups irrespective of whether they have the right answers or not;
- The amount of time spent in the discussion is limited; hence, students who need time to formulate their arguments may not contribute effectively;
- Since the discussion activity happens entirely in-class, classes with a large number of students (40+) and small meeting times (less than 1 hour) –which is typical in CS– may find it difficult to engage the students in any deep, serious discussion at all;
- The discussions remain ephemeral as there is no record of the in-class discussions and any detailed, technical point made –if at all– is likely to slip out of the participants' memory pretty soon.

Some of these problems can be solved by using online technology for the small-group discussions. For instance, online discussions happening outside the classroom will avoid the problem of disrupting other groups in the classroom; they can also keep detailed records of discussion transcripts which the students can refer to later; etc. Next we review some literature that uses online technologies for collaborative learning.

Computer-Supported Collaborative Learning, abbreviated CSCL, is a branch of the learning sciences that is "concerned with studying how people can learn together with the help of computers" [11]. Computer-Supported Intentional Learning Environments (CSILE) [12] (now KnowledgeForum) was one of the earliest CSCL systems. A group of (middle-school) students using CSILE focus on a specified relatively broad problem and begin to build a database of information about the topic. They raise questions, suggest hypotheses, propose possible solutions, and, most importantly, contribute information obtained from outside experts. There is opportunity for reflection and peer review of each others' contributions by students. The focus of such systems, like most CSCL systems, is on the group as a whole synthesizing/analyzing knowledge.

Some authors, e.g., Cress and Kimmerle [13], have proposed using *wikis* to allow users to add, modify, or delete content using a standard browser, to create a site that thoroughly explores a topic. This is similar to CSILE ([12]) but, as Larusson and Alterman [14] note, "wikis are plastic" and can support a variety and range of learning activities and types of interactions among students. Unfortunately, however, wikis have failed to live up to their promise of enabling cooperative learning. Cole [15] organized his course on information systems with 75 students so that lectures were in alternate weeks, the other weeks being intended for students to discover new material and post to the class wiki. Fully one quarter of the questions on the final exam were to be from the material that students posted. The expectation was that students would post content, edit each other's posts, and engage in collaborative learning. Halfway through the course there had been no posts to the wiki! Leung and Chu [16] in a course on knowledge management and Judd et al. [17] in a large course on psychology report equally poor results of the use of a wiki. Although they obtained positive results using wikis in architecture and English composition classes, Rick and Guzdial [18] report that the results in STEM classes were "overwhelmingly disappointing".

As we will see in the next section, our approach, in some sense, exploits the plasticity of web systems to address these challenges.

In answering this lead-in question, pick the one answer that you think is most *correct* and *complete*; and provide a brief justification of your choice.
The static mechanism, when used inside a Java class, is used for the following reason:

(a) The "static" keyword is used for only one purpose: to flag the main() function of the Java program so that the system will know that is where the execution should begin. The "static" mechanism is not used for anything else in Java (unlike in C++ which uses it for other purposes).
(b) In some sense, "static" is essentially equivalent to declaring something to be "public" so that a variable or method of the class that is flagged as static can be used anywhere in the program.
(c) Part of what (b) says is correct; when a variable or method of a class is flagged as "static", it is indeed *potentially* usable from anywhere in the program; but only *if* it is also flagged as "public". If it is flagged as "private", it is entirely useless since the rest of the program cannot use it.
(d) Part of what (c) says is correct but only part of it. If a class variable is flagged as "static", there is only one copy of that variable and that will be shared by all instances of that class rather than each instance having its own copy. The variable will be accessible anywhere in the program if it is flagged as "public"; but if it is flagged as "private", it will be only accessible by *static* methods of that class.
(e) Oh, (d) is so close but not quite right! A static, private variable of a class may be accessed by *any* method of the class, not just static methods.
(f) It is a useless mechanism. There are no situations in practice where we would need to use it. It should be removed from the language!

Fig. 1. CONSIDER Phase-1 (Initial Question)

III. Approach

Our approach, named CONSIDER (an acronym for CONflicting Student Ideas to be Discussed, Evaluated, and Resolved), works as follows. Following standard class lectures on a given topic, the instructor will create an assignment that students will complete in *three phases*. In Phase-1, the instructor will post, on the CONSIDER web app, a conceptual question – ideally, a multiple-choice question with distractors chosen to correspond to common misconceptions about the topic. An example appears in Fig. 1.

Each student in the course will have an account on the CONSIDER system, and will be required to log into her

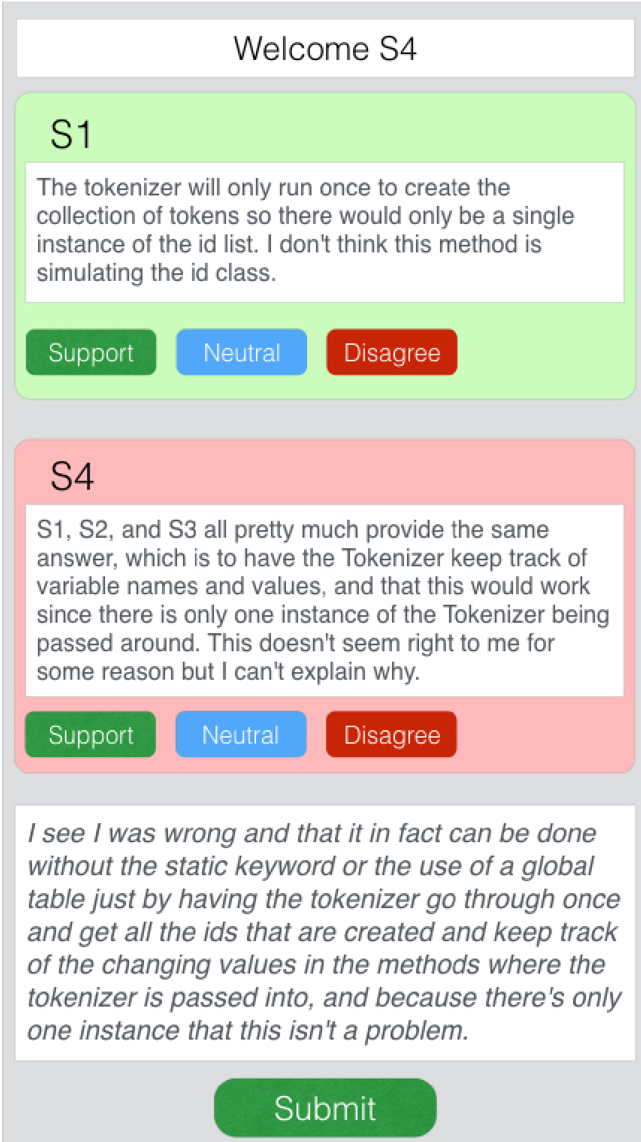

Fig. 2. CONSIDER Phase-2 (Discussion Rounds)

account, and *individually* submit her answer to the question by the specified deadline, typically, 24 or 48 hours from when it was posted. The student will not only have to choose one of the options but also provide a justification for the choice. The student can come back and modify the answer any time before the deadline. After the deadline, the system, possibly with help from the instructor, organizes the students into groups of 4–5 ensuring that each group contains students who chose different answers. If the students are not distributed evenly across the choices, the instructor may analyze the text justifications to form the groups.

Phase-2, the main phase of the CONSIDER activity, starts after the groups are created. In this phase, the students in each group will engage in a series of rounds, $R1, R2, \ldots$, of discussion, each round lasting 24 hours. The question or problem being discussed in this phase may be the same as in Phase-1 or an extension, possibly including a substantial problem-solving component. The goal of the discussion is to help each student in the group sharpen her understanding of the topic and arrive at an answer to that problem. Unlike most other collaborative learning approaches, the goal is *not* for the *group* to arrive at a consensus answer. Instead, the goal is to have each student in the group arrive at her own answer to the question after careful consideration and analysis of the ideas of all the students in the group and the round-structure is intended to ensure this.

Suppose a given group has four students. One important feature of CONSIDER is that the students in a group will *not* know the identities of the other students in the group. The system will simply refer to them as S1, S2, S3 and S4. When S4 logs in for, say, round $R3$, she will see the posts made by all four students in the previous round, $R2$. In her post for $R3$, S4 will be required to indicate (by clicking a *green, red or blue* button on the app) whether she *agrees with, disagrees with, or is neutral/unclear about* the posts made by *each* of S1,. . . ,S4 in $R2$ along with an explanation (especially if she disagrees); and also include her *current*, possibly revised, approach to the problem. Note that S4 has to indicate, in her post for $R3$, whether she agrees/disagrees with her *own* post from $R2$; the point is that, she may have found the $R2$ post from, say, S1 so compelling that she no longer agrees with her own position in $R2$! Indeed, this is *precisely* the point of peer discussion based on different conceptualizations of a problem. An example appears in Fig. 2.

The number of discussion rounds will be decided, in advance, by the instructor depending on the complexity of the topic, the time available, etc. After completion of the discussion rounds, Phase-3, the final phase, begins. In this phase, each student will be required to individually submit his/her final answer to the assignment along with a brief *summary* of the discussion in his/her group. As in the case of the earlier phases, the student may log back into the system and revise her answer, if she wants to, any number of times before the deadline for this phase. This allows the student time to think over her response in each round and keep refining it as required; and also avoids knee-jerk reactions that are common in electronic *discussion forums*. S1's grade for the assignment will depend *only* on the correctness of her final answer and the quality of her summary; there is *no penalty* for changing the answer from an earlier round to the final round; thus the student will focus on arriving at the correct answer by analyzing the other students' positions and her own originally-held position.

While the notion of exploiting differences in understandings among students in a small group and discussions within the group to help all students develop their understanding is based on earlier work, our approach, by careful use of the power of on-line systems, not only addresses the challenges to collaborative learning in college-level STEM courses listed earlier, but also offers a number of other important advantages. The fact that students in a group do not know each others' identities helps ensure free participation and mitigates any prejudices or biases some students may have about others. The structure imposed on the discussion by the carefully defined concept of *rounds* with each student making exactly one post per round ensures that every student contributes effectively to the discussion and helps each student in the group develop his/her understanding. None of this is possible, or at least practical, without the use of technology.

We have implemented the approach as a scalable, platform-independent web application, using Google App Engine and Python, making it ubiquitous, accessible from any net-connected device of choice of the user. Fig. 2 shows a (recreated) user interface of the app as it will look on a mobile device. Note that the students are being identified with aliases (S1, S4, etc.) and there are green/blue/red buttons to indicate whether this student supports, is neutral towards, or disagrees with the previous round post of every student in the group; the text input box at the bottom lets the student enter his/her text response for the current round. Posts by S2 and S3 in this group are not shown due to space constraints.

IV. CONSIDER in a CS Classroom

In this section, we report on our experience with using the CONSIDER approach and system in Autumn'15 in a standard junior/senior-level Computer Science on programming language principles at the Ohio State University. This course, similar to the ones in other programs, focuses on different programming paradigms, on key concepts underlying important classes of programming languages, and on questions/problems related to implementations of these languages.report on the results of this use.

The assignment, conducted as a CONSIDER discussion, was about the "static" mechanism in languages such as C++ and Java. It was conducted as a regular, graded homework activity for the course. All 43 students in the course participated, and 31 of them consented to let us use their data as part of the research. The question shown in Fig. 1 was used in Phase-1 and the answers submitted by the individual students were used to form small groups (typical size: 4) of students with conflicting ideas. This being the first course that focused on the *conceptual* ideas underlying programming languages for most of these students, there is a possibility of misconceptions about the precise nature of these constructs, as well as about their intended usage. The goal of the assignment was to help improve student understanding with respect to both of these aspects of the *static* mechanism.

In Phase-2, consisting of two discussion rounds, the students had to address a more detailed question that asked for strategies to implement a *tokenizer*, which is a key component of implementing any programming language. In an earlier project in the course, the students had already implemented a tokenizer for a simple programming language using the `static` mechanism. The question here was to come up with an approach to implementing the same tokenizer *without* using the static mechanism (and without using a global table of Identifiers which, in this particular context, would amount to using the static mechanism).

In almost every group, some students who started with a wrong notion about how to effect such an implementation ended up rectifying their approach; indeed, even those who started with the correct basic idea, were able to refine their strategy based on the CONSIDER discussions. Fig. 2 shows an example. The student identified as S4 had started with a position that the problem could *not* be solved under the given constraints (not using static mechanism or global table of Id objects) and expressed this in the previous round (box in the middle) and that what others suggested in earlier rounds "doesn't seem right", but he was not able to explain why. In the current round, he sees S1's previous round post (box at the top) which explains: "The tokenizer will only run once . . . there would only be a single instance of the id list". Upon reading this, S4 *agrees* with S1 (by clicking the green Support button, which makes the background for S1's post appear green), and *disagrees* with himself (red background for S4's post), and describes his new position in the input box at the bottom: ". . . I was wrong and . . . it in fact can be done without the static keyword. . . ".

Such changes, which were observed in every group to a varying degree, demonstrate the power of CONSIDER approach. In particular, requiring each student to specifically state whether she agrees with/disagrees with/is neutral about the position expressed by each of the other students and, of course, possibly revising one's own position in the process while, at the same time, having the student's final grade depend only on the correctness of the student's final answer and not, for example, whether the student "borrowed" ideas from any of the other students, truly makes it possible for all students to learn from their peers.

The entire assignment ran for 5 days: Phase-1 was over two days; followed, in Phase-2, by two rounds of discussion, each lasting 24 hours; and Phase-3 lasted 24 hours. Students did all this by logging into the CONSIDER system at their convenience using a web browser of their choice. In the final exam, which happened after two weeks of this activity, a question related to the topic discussed on CONSIDER was asked to evaluate longer-term retention of the concepts due to CONSIDER activity. Following the activity, students were asked to complete an anonymous, online survey seeking their opinions on the CONSIDER approach and the system. 21 of the 31 students who agreed to participate in the research responded to the survey.

V. Results

A. Analysis of the discussion data and Final exam scores

Student participation was measured using mean number of words in a message. Such a surface level measure, though not an indication of the *quality* of the message, serves as a primary indication of the degree of participation of a student in the discussion [19]. While it is not a direct measure of learning, some authors believe that "it is necessary in order for a discussion activity to be successful and result in learning" [20], and that if students feel they have participated effectively, they tend to be more successful in online environments [21]. More recently, some authors have also used 'thread-length' as a measure of participation; see, e.g., [22]. But, in CONSIDER, the 'thread-length' or the number of hierarchical posts in a conversation is more or less constant, since each student makes exactly one post per round. Due to the lack of variance in this factor we exclude it, and use mean number of words alone as the measure of participation. For the student posts we analyzed, it ranges from 31.33 to 283, with 110.09 as the mean.

We acknowledge that reviewers of this paper pointed out that the message length is not a good measure of participation. We are investigating alternative metrics that would account for such factors as whether a student, in disagreeing with another student, simply restates her original position or, in fact, provides cogent arguments to show why the other student's position is incorrect; etc. We believe that some of the existing work on assessing quality of argumentation (e.g., [23, 24]) can serve as the basis for this.

The question in the final exam that was related to the topic discussed in the CONSIDER activity was graded by the course instructor, also an author of this paper, on a 5-point scale. The score on this question ranges from 0 (not attempted) to 5 (correct answer), and the mean is 3.55. We found a positive correlation between the student participation and their performance on the relevant question in the finals (Pearson's $r = .26$), indicating that more involved participation in CONSIDER leads to a development of deep conceptual understanding about the topic and helps with long-term retention of the concept.

Analysis of the discussion logs shows that, in each group, students modified their answers as a result of the discussion. For instance, this student mentions in his summary: "[My solution] was overly complicated…My tokenizer tried to do everything the ID class would do, but need not", and after the discussion he made it simpler. A second student mentioned that he did not come up with a complete solution in the first round, and only after discussing with others could he improve his solution. Another student said, "During several rounds of discussion I was able to get a deeper understanding…[and] ensure that my method is working."

In our observations, such changes in position were typical in each discussion group. These comments highlight the benefits of the unique features of CONSIDER which are essential to enable such discussions on specifics of a student's solution, which in turn lead to learning of deeper concepts of the topic.

B. Analysis of the Survey

21 out of the 31 students participated in the post-activity survey. 15 of them were males, 5 females; one chose not to disclose the gender. 6 of the respondents were Asians, 12 Caucasians and two of mixed-ethnicity; one chose not to disclose ethnicity. All of them were Computer Science majors.

The survey consisted of some quantitative questions and some open-ended comments. The quantitative results are discussed using bar charts. Student comments were grouped into supportive and critical (of the approach). It was observed that

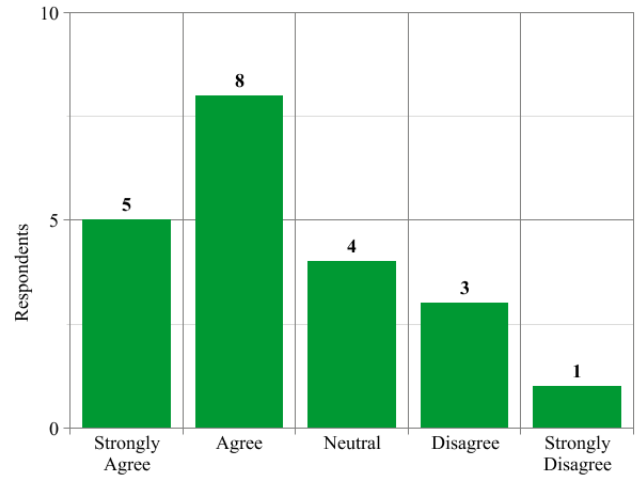

Fig. 3. Survey results: CONSIDER provides a better opportunity to learn compared to in-class discussions.

many students made similar points in their posts. The following discussion covers all the points made in the comments pertaining to a question, but quotes only the most articulate one due to space limitations.

1) Reflections on Learning using CONSIDER: More than 60% of the respondents answered 'Agree' or 'Strongly Agree' on a 5-point likert scale when asked whether they thought that CONSIDER provided a better opportunity to learn compared to in-class discussions (Fig. 3). Many of them also recommended, in the text comments, that the app should be used more frequently in that course, and should also be introduced in other CS courses. One student suggested this should be made a regular, weekly activity. Students are not accustomed to working on the same homework over a few days, and making this a weekly activity will form that habit leading to more effective participation in CONSIDER, he said. Some students who did not like the approach suggested that an open ended, challenge-type question would be better than the multiple choice question. Some others said they would prefer an intervention like this used as an extra credit assignment than a graded homework. In future, we would like to start the CONSIDER exercise early in the semester and, indeed, conduct an ungraded 'practice sessions', so that the students are familiar with the interface as well as the approach.

2) Anonymity: Student responses on two key features of CONSIDER were requested through two other questions, each of which also used a 5-point Agree–Disagree Likert scale. On the question of *anonymity*, more than 70% respondents either 'Agreed' or 'Strongly Agreed' with the statement: "Not knowing the identities of the other students in the group had a positive impact on the quality of the discussion." (Fig. 4; green bars). Some of the text comments explained why they agreed with the statement. For example, one student said: "I liked the anonymity of the app, it made me feel *safer* when submitting answers". Another comment: "I honestly really love this idea. Especially for someone like myself, who is always afraid to

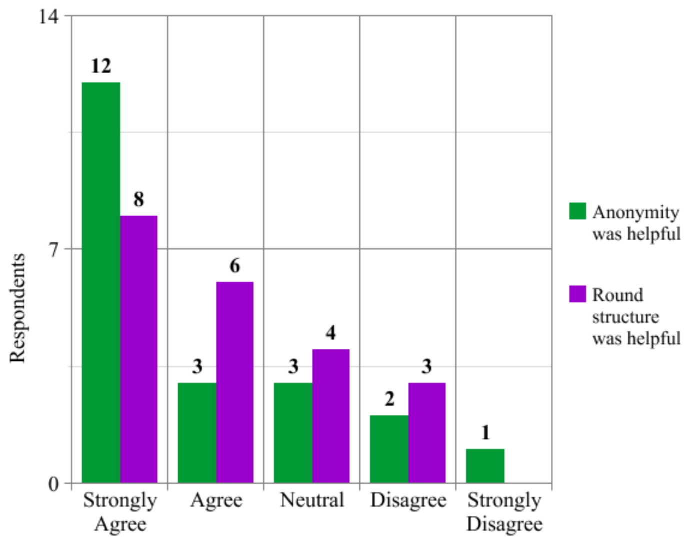

Fig. 4. Survey results: Features of CONSIDER

ask questions for fear of embarrassment. The anonymized piece of it was probably my favorite part." – is representative of the section of the student population that does not readily participate in discussions. As mentioned in Section III, our approach claims to address that issue, among other things, and this comment corroborates that claim.

On the other end of the spectrum, we saw the following comment: "You don't get to be anonymous in the real world. I want people to own their answers and opinions. There's nothing shameful about being wrong, and I think people should do it more often." Although, as educators and/or education researchers, we would agree with the underlying sentiment that it is okay to make mistakes in a learning environment, it is not very easy to persuade the 'shy' students to start participating in class discussions overnight. It is better to design techniques where such students feel "safe" about voicing their opinions, and are not "embarrassed" about making mistakes. Also, if participating in discussions involves critiquing a peer, the social cost of threatening a good relationship with your classmates is sometimes so high that people tend not to be seen as arguing in such situations [7]. The same student who was not in favor of anonymity, however, continued: "…if you haven't discussed with your group in-person, then anonymity online is better." Since the CONSIDER setup does not precede with an in-person discussion, we believe, even this student agrees with the usefulness of the anonymity feature in CONSIDER.

3) Asynchronous, Rounds-based Discussions: The other unique feature of CONSIDER is its asynchronous, rounds-based structure. Two-thirds of the respondents either 'Agreed' or 'Strongly Agreed' with the statement: "Organization of the discussion into a series of rounds had a positive impact on the quality of the discussion." (Fig. 4; purple bars.) The fact that the rounds structure ensures that everyone contributes to the discussion was highlighted by student comments like this one: "I like how nobody can see someone else's first answer until they have answered themselves…" This (indirectly) refers to the free-rider problem, where it is perceived that, in collaborative learning, there is a possibility of one or more students simply peeking into the answers of others from his/her group and get the credit and/or grades without really doing the work (see [22, 25, 26]). In CONSIDER, since everyone has to come up with his/her own initial answer, and then comment on *every* group member's comment, this problem is somewhat mitigated.

Some other students had some concerns about the rounds-structure, the main one being the 'turn-around time' of a response. Consider a case where a student, say S1, reads S2's Round-2 post which, he thinks, is not clear on one or more points, so he includes a question asking S2 to clarify those points in his $R3$ post. Now S2 sees this question in $R4$ and (hopefully) includes the clarification in his post for that round. But that post is not visible to S1 till Round-5 begins! Thus, for S1 to get the answer back from S2, it has taken from $R3$ to $R5$, which could be a gap of up to 48 hours assuming each round is of 24 hours. We are planning to add a 'quick-response' feature in the next version of the app to address this problem. This feature will allow a student to ask at most one clarifying question to each member of his group per round. A notification about the question will be *emailed* to that student immediately. That student is expected to, but not required to, send the clarification before the round ends. If she does, the clarification will be posted on the app, in their group, so that all the students in that group can see it right away and formulate their current-round posts accordingly.

4) Reflections on CONSIDER assignment: Other questions in the survey focused on how this particular CONSIDER assignment was designed, in terms of the number of discussion rounds and the duration of each round (Table I). Majority of the students (52%) felt that having two rounds of discussion was just about right. In some groups, however, the discussion converged in the first discussion round itself, and the students felt the second round was not required. This indicated the importance of creating groups based on cognitive conflict. Groups that did not have a good degree of divergence in ideas ended up not having very lively discussions. On the other hand, like we discussed in the previous subsection, if a student's post is not clear enough, others in the group tend to ask clarifying questions which need more number of rounds to get answered. In such cases, the number of rounds would have been 'not adequate' for the students to have a useful discussion.

About 62% respondents indicated that the 24-hour duration for each discussion round was appropriate. The remaining 38% indicated that they would like to have longer rounds. This is, at least partly, influenced by the fact that the students are not used to having some homework task to do every day, as well as the fact that this particular assignment was conducted in the last few weeks of the semester, and students found themselves hard-pressed for time. As mentioned earlier, conducting CONSIDER assignments early on in the semester, and making it a routine exercise would address such issues.

TABLE I
SURVEY RESULTS: CONSIDER ASSIGNMENT

	Not adequate	Just about right	Too much
Num. of Rounds (2)	9	11	1
Round Duration (24 h)	8	13	0

VI. CONCLUSION AND FUTURE WORK

Classic work by Piaget, Doise and Mugny and others has shown that cognitive conflict can contribute effectively to children's learning. Prevalent collaborative learning techniques used in CS and engineering education tend to focus more on aspects such as team work and group cognition, than on individual learning of technical concepts. In CONSIDER, we have developed a novel and effective online approach to use cognitive conflict in college-level computing and engineering courses in collaborative learning activities that will help students develop deep understanding of relevant concepts. Some unique features of our approach and their benefits are as follows:

- Small group formation based on cognitive conflict: The discussion in each group is driven by the conceptual disagreement about the topic among its members; attempts to resolve it would lead to deeper understanding.
- Anonymous posting in groups: Students participate more freely, and the effectiveness of the discussion is not compromised by any gender/ethnic/other pre-conceptions some students may have, or by apprehensions of offending a classmate. Also, shy students feel less pressured when participating in the discussions.
- Asynchronous, structured rounds-based discussions: Each student, whether quick on her feet, or prefers to think through subtle ramifications before posting, or anything in-between, participates equally effectively. Also helps avoid flame wars, i.e., a continuous back and forth of posts which escalate the conflict rather than resolving it, and restricts any member(s) from dominating the discussion.
- Online record of the discussion: Instructors can look at the interactions and decide if further explanation is required for the topic.

We have implemented this approach as a platform-independent, device-independent, scalable web-application using Google App Engine and Python that can be accessed on desktops/laptops, tablets and cell phones with ease. We used it in an undergraduate programming languages course in Computer Science and Engineering where a discussion on "static" mechanism in languages like Java, C++ was conducted as a homework assignment. The discussion went on for 5 days, with 2 discussion rounds in CONSIDER. All 43 students in the class participated in the discussion, and 31 of them have consented to be part of this study. A relevant question was asked on the final exam. We analyzed the data and found a positive correlation between students' participation in the activity and their score on that exam question.

21 students completed to the post-activity anonymous survey. More than 60% of them responded that CONSIDER provided a better opportunity for deep conceptual learning compared to in-class discussions. About two-thirds of them said that the features of anonymity and asynchronous round-based structures had a positive impact on the quality of the discussion.

Limitations

An important precondition for CONSIDER to be effective is optimal group formation. When the initial question is a multiple choice question, ideally, it should be possible to organize the students into optimal 'conflicting' groups, provided the question is framed properly. If it is not, the instructor(s) will have to be involved in group formation, which is a time consuming process. Coming up with suitable questions that will distribute the students evenly across the conflicting choices is not easy either.

In Section V, we discussed the problem with turn-around time: if a student asks a clarifying question to a peer as part of her current round post, she will get the response only after the current and the next rounds are over, which could be up to 48 hours. We are implementing the 'quick-response' facility for a clarifying question in the next version of the app to address this issue.

If some students drop out or are unable to respond in some round(s), the quality of their group's discussion may be impacted. But like any other homework or assignment, it is not possible to *force* any student to participate in the activity.

Students are not used to actively thinking about a problem and participating in a discussion around it over 5–6 days, albeit for only 20–30 minutes per day. Perhaps starting to use CONSIDER early in the semester and making it a part of the culture of the course might help in this regard.

Future work

We have demonstrated the effectiveness of CONSIDER in improving deep conceptual understanding regarding one topic in a CS course. Next we would like to evaluate the efficacy of CONSIDER in comparison with other existing online learning tools such as Piazza (https://piazza.com). These tools are quite popular in college courses, and provide facilities to organize small-group discussions, but lack the unique features of CONSIDER like anonymity (it is optional in Piazza) and round-based asynchronous discussions (discussions in Piazza are threaded). Two independent topics of comparable difficulty will be discussed using Piazza and CONSIDER. An exam question on each topic will be asked in the finals. The students' performance on those questions will be statistically compared, in order to evaluate the effectiveness of CONSIDER on actual learning in comparison to existing tools like Piazza.

Our tool is available as an open source software, which other educators can download and configure to use in their courses. It is highly customizable in terms of features such as number of rounds, duration of rounds, group size, etc., to suit their specific needs. It can be accessed at http://go.osu.edu/consider.

Acknowledgment

The authors would like to thank Rohit Kapoor who helped develop the CONSIDER web app, the students of the programming languages course who participated in the activity and responded to the survey, and the anonymous reviewers for their feedback that helped improve this paper.

References

[1] J. Piaget, *The early growth of logic in the child*. London: Routledge and Kegan-Paul Ltd., 1964.

[2] W. Doise and G. Mugny, "Sociocognitive conflict," in *The Social Development of the Intellect*, ser. International Series in Experimental Social Psychology. Amsterdam: Pergamon, 1984, vol. 10, pp. 77–101.

[3] D. W. Johnson, R. T. Johnson, and K. A. Smith, *Active learning: cooperation in the college classroom*. Edina, MN: Interaction Book Co., 1991.

[4] L. K. Michaelsen and M. Sweet, "The essential elements of team-based learning," *New Directions for Teaching and Learning*, vol. 2008, no. 116, pp. 7–27, Dec. 2008.

[5] L. S. Vygotsky, *Mind in Society: The Development of Higher Psychological Processes*. Cambridge: Harvard University Press, 1978.

[6] E. Mazur, "Peer instruction," in *AIP Conference Proceedings*, vol. 399. AIP Publishing, Mar. 1997, pp. 981–988.

[7] J. Andriessen, "Arguing to Learn," in *Handbook of the Learning Sci.*, R. K. Sawyer, Ed. Cambridge University Press, 2006, pp. 443–459.

[8] E. Aronson, *The Jigsaw classroom*. Beverly Hills, Calif.: Sage Publications, 1978.

[9] J. McTighe and F. T. Lyman Jr., "Cueing Thinking in the Classroom: The Promise of Theory-Embedded Tools," *Educational Leadership*, vol. 45, no. 7, p. 18, Apr. 1988.

[10] C. H. Crouch and E. Mazur, "Peer Instruction: Ten years of experience and results," *American Journal of Physics*, vol. 69, no. 9, pp. 970–977, Sep. 2001.

[11] G. Stahl, T. Koschmann, and D. Suthers, "Computer-supported collaborative learning: An historical perspective," in *Cambridge handbook of the learning sci.*, R. Sawyer, Ed. Cambridge, 2006, pp. 409–426.

[12] M. Scardamalia, C. Bereiter, R. S. McLean, J. Swallow, and E. Woodruff, "Computer-Supported Intentional Learning Environments," *J. of Edu. Computing Research*, vol. 5, no. 1, pp. 51–68, Feb. 1989.

[13] U. Cress and J. Kimmerle, "A systemic and cognitive view on collaborative knowledge building with wikis," *International Journal of Computer-Supported Collaborative Learning*, vol. 3, no. 2, pp. 105–122, Jan. 2008.

[14] J. A. Larusson and R. Alterman, "Wikis to support the collaborative part of collaborative learning," *International Journal of Computer-Supported Collaborative Learning*, vol. 4, no. 4, pp. 371–402, Sep. 2009.

[15] M. Cole, "Using wiki technology to support student engagement," *Computers & Edu.*, vol. 52, no. 1, pp. 141–146, Jan. 2009.

[16] K. Leung and S. K. W. Chu, "Using wikis for collaborative learning," in *Proceedings of the 2009 Intl. Conf. on Knowledge Management*, 2009, pp. 1–13.

[17] T. Judd, G. Kennedy, and S. Cropper, "Using Wikis for Collaborative Learning," *Australasian J. of Educational Technology*, vol. 26, no. 3, pp. 341–354, Jan. 2010.

[18] J. Rick and M. Guzdial, "Situating CoWeb: a scholarship of application," *I. J. of Computer-Supported Collaborative Learning*, vol. 1, no. 1, pp. 89–115, Mar. 2006.

[19] R. Benbunan-Fich and S. R. Hiltz, "Impacts of Asynchronous Learning Networks on Individual and Group Problem Solving: A Field Experiment," *Group Decision and Negotiation*, vol. 8, no. 5, pp. 409–426, Sep. 1999.

[20] V. P. Dennen, "From message posting to learning dialogues: Factors affecting learner participation in asynchronous discussion," *Distance Education*, vol. 26, no. 1, pp. 127–148, Jan. 2005.

[21] R.-S. Shaw, "The Relationships among Group Size, Participation, and Performance of Programming Language Learning Supported with Online Forums," *Computers & Education*, vol. 62, pp. 196–207, Mar. 2013.

[22] J. Hewitt, "How Habitual Online Practices Affect the Development of Asynchronous Discussion Threads," *Journal of Educational Computing Research*, vol. 28, no. 1, pp. 31–45, Jan. 2003.

[23] A. Weinberger and F. Fischer, "A framework to analyze argumentative knowledge construction in computer-supported collaborative learning," *Computers & Education*, vol. 46, no. 1, pp. 71–95, Jan. 2006.

[24] K.-H. Yeh and H.-C. She, "On-line synchronous scientific argumentation learning: Nurturing students' argumentation ability and conceptual change in science context," *Computers & Education*, vol. 55, no. 2, pp. 586–602, Sep. 2010.

[25] S. J. Karau and K. D. Williams, "Social Loafing: Research Findings, Implications, and Future Directions," *Current Directions in Psychological Science*, vol. 4, no. 5, pp. 134–140, 1995.

[26] N. L. Kerr, "Motivation losses in small groups: A social dilemma analysis," *Journal of Personality and Social Psychology*, vol. 45, no. 4, pp. 819–828, 1983.