

Integrating Components, Contracts, and Reasoning in CS Curricula with RESOLVE: Experiences at Multiple Institutions

Wayne Heym, Paolo A. G. Sivilotti, Paolo Bucci
Computer Science and Engineering
The Ohio State University
Columbus OH, USA
{heyw, paolo, bucci}@cse.ohio-state.edu

Murali Sitaraman, Kevin Plis
School of Computing
Clemson University
Clemson SC, USA
{murali, kplis}@clermson.edu

Joseph E. Hollingsworth
Computer Science
Indiana University Southeast
New Albany IN, USA
jholly@ius.edu

Joan Krone
Math and Computer Science
Denison University
Granville OH, USA
krone@denison.edu

Nigamanth Sridhar
Electrical Engineering and Computer Science
Cleveland State University
Cleveland OH, USA
nsridhar1@csuohio.edu

Abstract— Analytical reasoning is central to code correctness, and every computer science curriculum aims to teach students how to achieve this objective in one form or another. With the acceptance of object-based computing and component-based software engineering, the need for analytical reasoning that is based on formal contracts to establish correctness of software across module boundaries has become ever more obvious. Yet there are few institutions that have integrated modular, analytical reasoning principles into their undergraduate curriculum. Among many reasons for this shortcoming are: the effort it takes overloaded faculty to integrate new ideas of any kind in their courses, the challenge of institutionalizing ideas within a specific context, and constraints of a particular college. This paper presents our experiences over nearly two decades at five different institutions with the hope that they will serve as useful curriculum examples for like-minded educators at other institutions.

Components; correctness proofs; specification; verification; design-by-contract

I. INTRODUCTION

The courses discussed in this paper share two common objectives that we believe are shared by many CS educators: Students must learn to reason about the correctness of code (within components) and students must learn to build and reason about component-based software using contracts. The first one is a typical objective of earlier courses and the second one is a common objective of later courses. Typically, the reasoning process that is taught and the contracts that are used to describe components are presented

informally. We believe the reason for absence of formality in code reasoning and component contracts is multi-fold:

- Commonly used programming languages in CS education, such as C++ and Java, do not have formal syntactic slots for expressing contracts or reasoning.
- Instructors do not have the time to develop formal contract specifications and integrate them seamlessly into their existing courses.
- Fellow faculty at institutions are still skeptical that undergraduate students can actually learn to understand and use formal contracts in reasoning, and that there are benefits for such reasoning in developing high quality software.

We have conducted numerous assessments to show that students can successfully learn formal reasoning principles. A summary of one such assessment spanning a 5-year period in two courses may be found in [1, 2]. More importantly, collectively we have successfully taught nearly 10,000 students at five institutions (four of which are public) over the last 15 years, and have witnessed firsthand the ability of typical undergraduates to learn formal reasoning.

The structures of the computer science curricula at the five institutions are all different and the places in the curriculum where we introduce formal reasoning also differ. Our hope, then, is that through this experience report educators will see that formal reasoning ideas can be introduced in a variety of ways, in a variety of courses at

different depths and using different example sequences. Materials and software used in these courses are either available publically at our web sites or available freely upon request.

At Ohio State, the ideas are introduced in the introductory CS course sequence where Java is the programming language of choice. At IU Southeast, the principles are taught in a second-year data structures and algorithms course, using C++ as the language of instruction. At Clemson, a required two-course software engineering sequence presents the ideas with support from a web-integrated environment to facilitate reasoning. At Cleveland State, the ideas are used throughout the software lifecycle process in a required software engineering course. At Denison, some of the principles underlying verification are presented in a software engineering course. In all cases, analytical reasoning is one of the topics covered, but not the only one. At all institutions, the ideas have gone through extensive assessment before they were institutionalized, though details of these assessment efforts are not the topic of the present paper.

The common element of the offerings of these institutions is the use of RESOLVE specification notation—conceived and refined over a period of 30 years especially with undergraduate computing education in mind [3, 4]. While there are several specification approaches in the literature, RESOLVE notation has been employed and assessed for use in undergraduate audiences for decades. The notation is intended to make a variety of core reasoning concepts taught in undergraduate computing easily accessible. At the same time, RESOLVE is sufficiently developed as a research vehicle that it is possible to specify and verify non-trivial software components automatically using a push-button verifying compiler that we have developed.

It is important to note that the ideas of formal contracts and reasoning are not the only content of the courses discussed in this paper. They are merely integrated with other topics normally taught in the courses. A part of the challenge is how to incorporate reasoning topics without necessarily displacing other topics that need to be covered. For example, in the software engineering course at Clemson, only five weeks are devoted to the topics covered in this paper. The rest of the course is concerned with other software lifecycle topics, such as requirements analysis.

The rest of the paper is organized into the following sections. Section II is a summary of the introductory course sequence at Ohio State, a novel aspect of which is a simultaneous introduction to specification, programming (Java), and reasoning. Section III is a summary of IU Southeast offering of data structures and algorithms in C++, a novelty of which is its online component. Section IV is a summary of the Clemson two-course sequence with focus on the novel usage of a web-integrated reasoning environment for developing verified components in a software engineering course. Section V contains a summary of the Cleveland State and Denison software engineering courses. In the discussion of our experience at each institution, we give an overview of the courses and outline the reasoning topics covered in the course. We also

summarize assignments, labs, or projects that are used to reinforce the ideas. In each case, we present illustrative examples to highlight key points. Sections VI and VII outline some related work and summarize.

II. SUMMARY OF THE INTRO SEQUENCE AT OHIO STATE

Introductory computer science at Ohio State is a two-course (one academic year) sequence: Software I and Software II. About 1000 students, mostly CS or EE majors, take the sequence each year. Each course meets four times per week in forty-student sections: twice in a lecture hall and twice in a computer lab. Pair programming is encouraged during the lab activities. Students complete frequent small homework assignments, in-lab programming activities, and larger multi-week projects [5].

This sequence is not an introduction to programming. Before taking Software I, students must demonstrate programming proficiency in an imperative language of their choice. AP Computer Science A serves as a possible prerequisite, as do various one-semester courses offered at Ohio State such as “Intro to programming in Java” and “Intro to programming in C++”.

Java is used as the programming language and Eclipse as the development environment for the entire sequence. Eclipse is configured with some plugins (e.g., Checkstyle and FindBugs) as well as specific preference settings (e.g., format-on-save) to encourage adherence to specific coding conventions. Other standard tools, such as JUnit, Javadoc, and the Eclipse debugger, are also used throughout the sequence.

The key characteristic of this sequence is the order of presentation: Students are first taught how to *use* components, then how to *implement* them. We call this approach “client-view first” [6, 7]. This deliberate decomposition cleanly partitions the two courses: Software I is entirely client-view, while Software II is from the implementer’s perspective. In the first course, students learn to read component specifications, including mathematical models of state and behavioral contracts for methods. They write procedural code using these components and reason about the behavior of their code based on the component specifications. In the second course, students peel back the cover and look inside the components they used in Software I. They learn the relationship between a component’s internal concrete realization and its external abstract state. They implement method bodies that satisfy given contracts and they learn how to create layered, modular implementations.

The client-view first approach has many benefits. Firstly, students quickly get to use a rich collection of components to do engaging projects. For example, in the first few weeks of Software I, students use an XMLTree component to scrape interesting information from real-time RSS feeds. Secondly, students gain an appreciation for clean, well-documented interfaces. For example, tracing the step-by-step execution of a program using a Set component is possible only if students fully understand the client-side documentation of that component. Finally, students learn best practices for building modular systems. These best practices emerge as natural consequences of the strict

decomposition of client view and implementer’s view, rather than simple ad hoc coding idioms.

We have found that formal methods, including specifications and mathematical models, are invaluable in supporting the client-view first structure of this course sequence. Furthermore, we have found that formal methods can be tightly integrated throughout the introductory sequence and can support learning of traditional topics such as iteration, recursion, argument passing, reference types, etc.

A. Abstraction: Mathematical Models for the Client

Information hiding is an important topic in any course that teaches modular programming. Equally important, but often omitted, is the other side of the coin: Abstraction. Since information hiding restricts what a client knows about the internals of a component, there is a complementary need for a client-visible explanation of a component’s behavior. Such an explanation should not expose internal implementation details, so it is written in terms of a component’s *abstract* state.

Abstraction is pervasive in computer science. It is used so often, in so many ways, and in so many different courses, that it hardly draws any particular attention. In our client-view first introduction to computer science, however, we highlight abstraction as a first-class citizen of software design [8]. For each programming type, a corresponding mathematical model is introduced or reused. In Software I, the description of a component is always done in terms of its mathematical model.

There are many opportunities for confusion between an implementation and its abstraction. There is overlap in vocabulary; for example, “type” could refer to a program type (`int`) or a mathematical type (`integer`). There is also overlap in notation; for example, “=” could be assignment in Java or a statement about equality in mathematics. In some cases, there is even overlap in semantics, when the mapping between one and the other is trivial, for example the mathematical model for an `UnboundedInteger` being an integer.

Therefore, it is important that the distinction between implementation and abstraction be made clearly and consistently. Java provides a convenient syntax for supporting this distinction: We use interfaces to define mathematical models and method contracts written in terms of those models, while classes contain implementations. By decomposing each component into an interface and a class, the discipline of writing an implementation-neutral cover story follows per force. The best practice of “coding to the interface” also follows as a natural consequence of a client-view first approach and is supported with interfaces. More details about this use of Java interfaces are in [7].

While interfaces and abstract classes are used in the standard Java libraries, their use is not consistent and they are never fully abstract. Therefore, we provide students with a library of components with clear abstractions defined in distinct interfaces [9]. Some of these components parallel parts of the Java collections framework (e.g., `Queue`, `Stack`, `List`), some are illustrative of basic concepts (e.g., `AMPMClock`), some are utilitarian (e.g., `XMLTree`,

`SimpleReader`, `SimpleWriter`), and some are exemplars of best practices (e.g., `NaturalNumber`).

The building blocks for mathematical models for all of these components are simple mathematical types familiar to all freshmen, such as integers, reals, booleans, sets, and strings (i.e., lists of elements). Each type has a few operators, such as \wedge for exponentiation and $| |$ for the length of a string. Importantly, these operators are distinct from operators in the programming language. This difference between math syntax and Java syntax reinforces for students the difference between a client-side cover story (written in the former) and the artifact itself (written in the latter).

1) Example: Strings for Queues and Stacks

Mathematical strings—finite sequences of entries—are used extensively to model container classes. For example, both `Queues` and `Stacks` are modeled by strings.

```
type Queue is modeled by string of T
type Stack is modeled by string of T
```

This documentation is part of the Javadoc for each component’s core interface which also contains formal contracts for each method. The contracts are written in terms of pre and post-conditions, with $\#$ in post-conditions to indicate the initial values of formal parameters and the name of the method to indicate the returned value. Parameter modes (such as “clears” meaning the argument will be reset to an initial value for its type and “updates” meaning the argument may be modified as specified by the ensures clause) are also used as shorthand for common conditions. Custom Javadoc tags are used to structure these contracts, integrating them with standard Java documentation and including them in generated documentation.

For example, Listing 1 gives the contracts for `Queue`’s `append` and `dequeue` methods. These are read by the students and used to reason abstractly about their behavior.

```
/**
 * @updates this
 * @clears q
 * @ensures this = #this * #q
 */
void append(Queue<T> q)

/**
 * @updates this
 * @requires this /= <>
 * @ensures #this = <dequeue> * this
 */
T dequeue()
```

Listing 1: Javadoc Contracts for Queue Methods

Notice the mathematical notation and operators used in contracts, where $=$ is equality (not assignment) and \neq is inequality (not $!=$). The $*$ operator is string concatenation, the term $\langle \rangle$ is the empty string, and $\langle dequeue \rangle$ is a string consisting of a single element (the value returned by the `dequeue` method).

B. Example: Unbounded Natural Numbers

Other components require slightly refined models. For example, the `NaturalNumber` component stores an unbounded non-negative integer and involves an invariant on the abstract state (*i.e.*, a *constraint*), namely that the integer is not negative:

type `NaturalNumber` **is modeled by** `NATURAL`

```
NATURAL is integer
  exemplar n
  constraint n >= 0
```

The interface for this component contains the method contracts which use this mathematical abstraction for the value of `this`. For example, in the method contract for `divideBy10` (below), the (abstract) value of `this` is a `NATURAL`:

```
/**
 * @updates this
 * @ensures #this = 10 * this + divideBy10 and
 *         0 <= divideBy10 < 10
 */
int divideBy10()
```

C. Tracing Tables

Tracing tables employ strictly a client-view when “hand tracing” over any method call. Such tracing is in service of reasoning about the correctness of a method body under development. At this stage, each method to be implemented is done so in layered fashion. That is, it is not cooperating with other methods to represent a data type; instead, it provides an enhanced capability for an existing data type or is otherwise a “stand-alone” method. Such a method’s contract presents its client-view; the implementer reasons whether the proposed method body always satisfies its contract under the assumption that all methods called satisfy their contracts. The tracing table is similar in spirit to the reasoning table discussed in [10] for a Java-based course at Clemson.

One pedagogical advantage of this approach comes in teaching recursion. By learning to trace over all method calls, including recursive ones, students can see the implementation of a recursive method body as a problem of the same kind as the foregoing, with just a couple of additional constraints on the reasoning [6].

D. Implementing a Data Type

The task of arranging instance variables and bodies of constructors and methods to cooperate in representing a data type is significantly complex. For example, just the task of reasoning about the correctness of one method body involved in such a representation is significantly more complex than about that of one layered method body. Fortunately, techniques exist to organize and make explicit these additional demands in order to manage and simplify this complexity. Several of these techniques are involved in the art of designing the client-view of a data type; some more are involved in the art of designing a plan of implementation of such a data type. Our pedagogical approach supposes that students are well-served when they

are asked to examine and work in the presence of many examples of good design long before they are asked to perform overall design tasks themselves. Therefore, Software II focuses on the practice of implementing a representation of a component and has students practice implementing the component’s constructors and methods within the constraints of an already-complete design.

The primary artifact of such a completed design is, of course, the client-view of the component. In the currently-used programming language, this view is presented in Java interfaces, which, typically, have one or more generic type parameters and can be seen in [5]. Specification reuse is achieved via multiple inheritance. A salient feature of this design arises from the recognition that layered implementations of methods are significantly easier than implementations involving representation. Hence, the latter are kept to a small number by finding a minimal *kernel* set of methods that capture the primary features of the component and make its values observable and controllable via these kernel methods. The contracts for these methods are gathered together in the interface that also has the key job of describing the client-view of the values of objects of this component type. All other useful methods for this component have their contracts presented together in an “enhanced” interface, which extends the kernel interface. The enhanced interface is implemented by an abstract class, which doesn’t implement any of the kernel methods, but which does implement the enhanced methods in a layered fashion based, ultimately, on calls to the kernel methods, whose implementations are deferred to “kernel” classes that extend this “secondary” abstract class.

It is valuable for a given component to have multiple kernel implementations because these can differ in their time- and space-performance characteristics. As an example of engineering tradeoffs, they also vary in ease of implementation and/or understanding. A kernel class that is easy to implement can be built before one that has better performance characteristics and be used as a reference implementation in unit testing of the latter implementation. We have our students implement kernel classes in an in-class laboratory setting and as outside-of-class pair programming projects. We present our design decisions for each kernel class in three parts. A primary part is the choice of the types and names of the instance variables used in the representation. The next two parts, together, constitute the *internal contract* among the kernel methods. Often, not all (combinations of) possible values of these variables are sensibly used in the representation: a stated *representation invariant* (presented in a custom Javadoc tag `@convention`) captures this design decision. The intended interpretation of the instance variables’ values cooperatively representing a client-view value is stated in the *abstraction relation* (presented in a custom Javadoc tag `@correspondence`). The abstraction relation need only be sensible for instance variables’ values that satisfy the representation invariant. If each constructor and method body respects these three aspects of the design with respect to the corresponding contract in the interface, then the kernel class implementation will be correct. Hence, each operation’s body can be developed independently of each other body.

Additionally, we may wish our students to satisfy certain performance criteria, so we state those as well and may also require that certain algorithms be used.

An illustrative example comes from the third week of Software II. One in-class laboratory is dedicated to students implementing a kernel for the Set interface using one Queue value as the representation. The class they complete has the name Set2 in our catalog [5]. The one instance variable is named *elements*. The client-view for Queue (presented in QueueKernel) informs us that any value of *elements* is a *string of T*. It is to be used to represent a *set of T*, the client-view found in SetKernel. By this point in the two-course sequence our students are well-versed at taking the client view, and they reason about a component, e.g., Queue, using its math model rather than a concrete representation such as a linked-list or an array. The obvious interpretation of a string as a set is stated in the abstraction relation using a mathematical function from *string of T* to *set of T* named *entries*. This function “gathers together” every element present in the string to be a member of the set, disregarding any duplicates, of course. Because we designers recognized that the algorithms for methods *remove* and *removeAny* would be made more complicated if duplicates were permitted in the representation string, the representation invariant disallows duplicates (by insisting that the length of the string equals the size of the set obtained from that string via the mathematical function *entries*):

```
|$this.elements| = |entries($this.elements)|
```

We use the dollar sign (\$) as a prefix to “this” to indicate the representing class so that we can use “this” by itself to indicate the client-view value. Hence the abstraction relation here is expressed as:

```
this = entries($this.elements)
```

Consider, then, reasoning about the correctness of an implementation of the SetKernel method *removeAny*. The contract for this method is shown below:

```
/**
 * @updates this
 * @requires |this| > 0
 * @ensures removeAny is in #this and
 *           this = #this \ {removeAny}
 */
T removeAny();
```

The precondition states that the set is non-empty. The postcondition says that the return value should be some element of the incoming value of the set, and that the outgoing value of the set should contain all values that were in the incoming value, except the value that is returned. The implementer can reason by the abstraction relation that every element in the queue is a member of the represented set, so that dequeuing the front element of the queue will provide an element of the incoming set, and this value can then be returned. It is only by assuming, however, that the queue has no duplicates that one can reason that dequeuing an element has caused the queue to represent a set that no longer has in it the element that came out of the queue. It is the representation invariant that justifies this assumption.

The implementer, then, also has the obligation to satisfy this same representation invariant upon return. Obeying this obligation, of course, is easily done, for example, by not gratuitously duplicating another element within the queue, even though doing so would still, through the abstraction function, represent the contractually expected outgoing value of the set. The additional involvement of the abstraction relation and the representation invariant is the reason why implementing a kernel method is more difficult than implementing a method without using direct access to a representation. For this in-class laboratory, we asked students to implement one private static helper method (a typical layered implementation) named “moveToFront” as preparatory homework. We provided the implementation of the one constructor specified in comments in the SetKernel interface, but asked students to implement the five SetKernel [5] methods during the laboratory session.

The topic of this subsection is revisited in Section IV.B with a detailed code example in the RESOLVE notation.

III. SUMMARY OF THE DATA STRUCTURES COURSE AT IU SOUTHEAST

IU Southeast’s B.S. in computer science includes CS1 and CS2 as the first two courses for majors. Data structures and algorithms are introduced in CS2 as is common; however, a more in-depth coverage is permitted by our third course CSCI C343, which is called Data Structures. The introduction of data structures and algorithms in CS2 frees up time in C343 to integrate principles of engineering software that can lead to higher quality software development. One of the principles that we hit hard is design-by-contract, [11] which for us includes contracts consisting of requires and ensures clauses and the assignment of obligations and benefits between the client (calling operation) and service supplier (called operation). This one-semester course typically includes software development labs that include, first, client-view programming, followed by implementing container components by directly using C++’s built-in types, e.g., pointer types, and also by layering a new component on other existing components. Typical component abstractions implemented are: lists, sequences, maps, and sorting machines [12]. For example, the students often implement Map three different ways using the technique of layering – layered on list using linear search, layered on an array of lists using hashing, and layered on a binary tree component using binary search. All of these C++ components [13] have contracts embedded as comments where the contracts are written in the RESOLVE specification language. Students use the components’ contracts to reason abstractly about their behavior—this is taught early in the semester.

A. Teaching Students to Evaluate Software Design

One of the primary learning outcomes is: students will be able to evaluate an operation’s design (including its implementations, this course’s primary concern for this learning outcome) for adhering to the design-by-contract assignment of obligations and benefits. Instructional materials have been developed to support this learning outcome and time is spent in class examining various

operations that adhere and do not adhere to the design-by-contract engineering technique.

Below are two examples of operations that might be given to students for evaluation. During grading of lab assignments we collect examples of student submitted operations that violate design-by-contract and use these to drive discussion or test questions—many are more complex and nuanced than what follows.

The Queue template that we supply to the students is an unbounded queue and has specifications written in RESOLVE that are embedded as C++ comments. Specifications of Queue operations are similar to those discussed in the previous section, except that they are presented in a C++ context.

Now examine the operation *addToAll* in Listing 2 which has a *requires* clause of *true*, i.e., the client of *addToAll* is permitted to call it with a queue of any length. In English, *addToAll* ensures that it will add the value found in parameter *x* to all integers found in the parameter *q*. *addToAll* is an example of an operation to be evaluated by the students in order to determine if it is a “good” client of Queue, i.e., that *addToAll* never violates (under any circumstances) any of the *requires* clauses of the operations it calls.

```
typedef Queue<Integer> IntegerQueue;

void addToAll(IntegerQueue& q, Integer x)
// requires: true
// ensures: for all k: integer (0 <= k < |q|
//          implies q[k,k+1] = #q[k,k+1] + x) and
//          x = #x
{
    IntegerQueue qTemp;

    do {
        Integer y;
        q.dequeue(y);
        y = y + x;
        qTemp.enqueue(y);
    } while (q.length() > 0);
    q.transferFrom(qTemp);
} // addToAll
```

Listing 2: A Purported Implementation of the addToAll Operation

Examining *addToAll*'s code quickly reveals that it is a defective client in that it cannot guarantee that it will satisfy *dequeue*'s *requires* clause under all circumstances. (Recall, from Listing 1, that the *dequeue* operation's precondition *requires* that the queue be non-empty.) To computing instructors, this is an obvious “rookie” mistake, but to the students it is not so obvious. If this example is used as a teaching tool, then once the class agrees that *addToAll* is a defective client, we can then turn our attention to correcting the defect, e.g., by using a zero-trip loop such as *for* or *while*. Training the student's eye to catch these types of defects is paramount for us.

Another design exercise asks the students to evaluate an implementation of an operation with respect to taking advantage of the operation's *requires* clause. In C343,

students work with both bounded and unbounded versions of our components. The *dequeue* operation in Listing 3 comes from a bounded version where the BoundedQueue's internal representation is an array named *contents*, using a straightforward non-circular implementation.

```
template <class T, int maxLength>
void BoundedQueue1<T, maxLength>::dequeue(T& x)
{
    if(currentLength > 0) {
        x = contents[0];
        int z = (currentLength - 1)
        for (int k = 0; k < z; k++) {
            contents[k] = contents[k + 1];
        } // end for
        currentLength--;
    } // end if
} // dequeue
```

Listing 3: A Purported Implementation of the dequeue Operation

This operation is an example of an implementation that students are asked to evaluate for taking advantage of the *requires* clause. Under design-by-contract, the obligation for meeting the *requires* clause is on the client—in this case the caller of *dequeue* is required to call with a non-empty queue. The corresponding benefit is that the implementer of *dequeue* can assume the queue to be non-empty at the time of the call, therefore there is no need to have extra code to check for the queue being empty. The implementer of *dequeue* in Listing 3 failed to take advantage of this benefit as is seen by the all-encompassing *if* statement that first verifies that *currentLength* > 0. What is illustrated here is often referred to as defensive programming, which in general we agree with but only in specific locations of a software system. Defensive programming should be used when the software system has no control of the incoming data, e.g., data incoming from across the Internet. However, once that data has passed the system's “outer checking wall” and is now moving from one internal operation to another, a switch to design-by-contract must be made, which brings with it the ability to leverage design-by-contract's benefits.

B. Online Course Offering Considerations

Learning can be improved with increased student engagement; this is true for face-to-face classes as well as with online courses. Effective teaching in the online environment [14] must foster student-to-faculty interaction and also engagement among students. We have used the examples above in discussion forums as a way to foster and increase online student engagement in our face-to-face, hybrid, and our totally online classes. Through our learning management system, we post one or more examples where the students are asked to evaluate an operation's implementation with respect to adhering to design-by-contract principles. We are careful to set the discussion forum's options so that a student can only see the posts of other students after he or she has first made a post. These posts can be as simple as determining if design-by-contract has been followed and why not if the answer is “no”. Or, going further, the instructions might require the student to

post a better-engineered version of the operation under evaluation.

IV. SUMMARY OF THE SOFTWARE ENGINEERING SEQUENCE AT CLEMSON

With unanimous consent, the faculty of the School of Computing at Clemson has included the requirement that the CS graduates “will be able to apply design and development principles to develop defect-free software that functions according to its specification” as one of the six outcomes for all graduates. Since 2012, the principles of specification and reasoning are institutionalized in two required courses for CS majors, one, a second year introduction to software development foundations, where the basic concepts are presented in the context of Java, [9, 10] and another, a third year software engineering course. Students use the RESOLVE specification and implementation language in this latter course and develop verified components using a web-integrated environment [15, 16].

A. Software Development Foundations

At Clemson, a prerequisite to the software engineering course (discussed next) is a required course that introduces students to object-based development using Java, and instills basic principles of analytical reasoning using specifications. Contents and organization of this course as an exemplar ACM 2013 curriculum course are detailed in [10]. Some of the topics covered in this course are similar to those detailed for the introductory Java-based course sequence at Ohio State in Section II.

One of the important additions to the course, occurring in all its sections for the 2016-2017 academic year, is the incorporation of a web-based tool to enhance student understanding of formal reasoning. A detailed description of the tool, its experimentation in multiple sections, and results from evaluation are reported elsewhere [17]. Using simple examples, the tool aims to help students gain a general understanding about code, going beyond tracing with specific input values. The goal is for students to learn to predict what a given piece of code, involving assignment and control statements, will do on arbitrary inputs without actually running the code. The tool is supported in the back by the RESOLVE verifier to give feedback to students and to give feedback to instructors regarding where students face obstacles in reasoning.

B. Software Engineering

The software engineering course at Clemson uses the RESOLVE specification and implementation language, and students develop code according to given contract specifications. Moreover, they use a web-integrated and automated RESOLVE verification system, available through a web IDE to verify that their code satisfies formal contracts [18]. Different sections of the software engineering course introduce and use RESOLVE and its supporting tools to varying degrees. Given below is the summary of some of the activities in one recent offering.

One of the simplest exercises students do in this course is to develop code for the secondary Queue Rotate operation. (See Listing 4.) In the ensures clause, `Prt_Btwn` is

a mathematical string notation, and it is the substring between the given positions. The verification system checks that the code (procedure) satisfies the contracts for the operations it calls (requires clauses of Dequeue and Enqueue) and that the code ensures the contract of the operation it implements (ensures clause of Rotate). The verification, of course, assumes the requires clause of Rotate, because it is the responsibility of Rotate’s caller. Any errors are reported and students iterate the process until the code is verified.

```
Operation Rotate(updates Q: Queue);
  requires 1 <= |Q|;
  ensures Q = Prt_Btwn(1, |#Q|, #Q) o
           Prt_Btwn(0, 1, #Q);

Procedure
  Var E: Integer;
  Dequeue(E, Q);
  Enqueue(E, Q);
end Rotate;
```

Listing 4: The Rotate Operation

A more complex activity involves implementing one component using another. For example, implementation of a Set component was discussed in Section II. Listing 5 contains part of an implementation of Queue_Template using List_Template in RESOLVE. The key idea in implementing data abstraction is the use of internal contracts asserting a convention (to capture the representation invariant) and a correspondence (which captures the abstraction function or relation between the abstraction in the specification and the representation in the implementation). Students are given the interface specifications of components and the internal contracts that must be satisfied.

In the List_Template, a list is modeled mathematically as an ordered pair of strings of entries: a *preceding* string of entries, and a *remaining* string of entries. This conceptualization allows one to think of a list as having a cursor position in between the two strings. When an element is inserted into the list, it is added to the front of the *remaining* string. The remove operation has the opposite effect, so that an insertion followed by an immediate removal leaves the list unaffected. There is also an operation Advance that allows the cursor position to be moved so that insertions can occur at different places in the list. The Reset operation has the effect of moving the cursor position to the front and Advance_to_End has the effect of moving it to the rear.

For the assignment given in Listing 5, the representation invariant is that the preceding string is always empty. The correspondence is that the abstract queue is what is in the remaining string. This is why the code for Dequeue can simply call the list Remove operation. But the Enqueue code needs to move the cursor position to the end before insertion. To satisfy the representation invariant that code needs to reset the cursor, so that the preceding string is empty.

```

Realization List_Based_Realiz for Queue_Template;
  -- not shown: declaration of an instance
  -- of list of entries, named LF.List
  Type Queue = Record
    Contents: LF.List;
    ...
  end;
  convention Q.Contents.Prec = Empty_String;
  correspondence Conc.Q = Q.Contents.Rem;

  Procedure Enqueue(alters E: Entry;
                   updates Q: Queue);
    Advance_to_End(Q.Contents);
    Insert(E,Q.Contents);
    Reset(Q.Contents);
  end Enqueue;

  Procedure Dequeue(replaces R: Entry;
                   updates Q: Queue);
    Remove(R,Q.Contents);
  end Dequeue;

  -- code for other Queue operations
  ...
end List_Based_Realiz;

```

Listing 5: The Enqueue and Dequeue Procedures

V. SUMMARY OF THE SOFTWARE ENGINEERING COURSES AT CLEVELAND STATE AND DENISON

At both Cleveland State and Denison universities, the RESOLVE notation is used in an undergraduate Software Engineering course, which is a four-credit-hour, overview course that covers all aspects of the software engineering life-cycle.

At Cleveland State, the course has specific emphasis on the correct construction of software systems and it is required for all Computer Science and Computer Engineering students. Students spend five weeks (20 hours of total instructional time) working with the RESOLVE notation. The first use of the notation is in the context of requirements specification. Students are introduced to the idea of specifying requirements in a formal, unambiguous manner. At the same time, students are also made conscious that rigorous formal specifications do not necessitate heavy, detailed, proof systems. Rather, the students are exposed to the modular specification method that is characteristic of RESOLVE—the proof obligations are embedded in the lower-level components, and *clients* of these lower-level components can leverage these proof guarantees in a light-weight manner. Later in the course, when working on design and implementation phases in the software engineering lifecycle, students rely on RESOLVE-style specifications to reason about correctness of implementations.

An example of an exercise that students work on in this course is the detailed specification of software for a portion of an air-traffic control (ATC) system. In particular, students describe how the ATC system can manage a number of aircraft in a given sector of airspace at a particular time, with the following constraints: (a) the airspace may include a number of aircraft, each with a unique identifier, (b) all

aircraft must be separated by at least 300 meters in height, (c) a controller may create a new sector, (d) aircraft may enter, leave, and move in the sector, and (e) a controller may look up a particular aircraft, and check if a particular region is occupied. The students go through an iterative design process, writing specifications for the abstract state of the sector, and each of the operations. Over the course of the iterations, the students use specifications for modeling the Set, PartialMap, and eventually, BijectivePartialMap components. The students then work through a reasoning exercise to construct a correctness argument for the Sector component.

The catalog description for the elective software engineering course at Denison emphasizes the importance of making connections between theory and practice, and it states that students will apply their theoretic background, together with current research ideas to solve real problems and that they will draw from their entire computer science curriculum, noting how theoretic results apply to real problems.

This course covers all aspects of software engineering even as students seek to apply the principles to a real-world problem. For example, in one offering, students completed a webpage that allowed local taxpayers to look at the tax reports of the school system and make predictions according to what level of taxation the village had.

Students learn to design software that is based on the principles that the RESOLVE system promotes: Each software component should be carefully specified so that reasoning about the design can be done abstractly. The overall design must be modular so each component can be checked for correctness based only on its abstract specification. It should be implemented efficiently. Students have access to the RESOLVE web interface provided by Clemson University where they see examples of software specified, implemented, and verified using that system. Before their introduction to RESOLVE, the students examine specifications of library components in C++ to critique those specifications and figure out why they are not adequate for making choices about using those components. Then they examine RESOLVE specifications for stacks and queues and are introduced to the concept of formal, mathematical specifications. They write one or more enhancements of stacks or queues to see the importance of, and ease of building new components using existing components that are formally specified and that have been verified.

Students at Denison have a strong mathematical foundation and so we examine and apply some basic proof rules for RESOLVE constructs so students get a glimpse of what it means to use automated verification. Then they use the RESOLVE web IDE to verify the enhancements they have written. Although students are not RESOLVE experts upon completing CS349, they have gained a sense of how software might be done in the future and they have not only some practical experience in doing their own project, but they have developed and mastered many principles that they can carry over into their graduate programs or jobs.

VI. RELATED WORK

The importance of mathematical thinking in computer science is widely recognized, and there are many approaches for teaching it as part of undergraduate curricula [19, 20]. Educators have explored a variety of ways to teach formal methods, in particular, at the undergraduate level [21, 22, 23, 24, 25]. The assessment of these approaches is typically done in the context of a single institution, and often a single course. In contrast, we have considered how a single approach, one using RESOLVE for formal reasoning, has succeeded at a variety of institutions and at a variety of levels of courses.

Design-by-contract, in a general sense, is naturally part of any specification and reasoning methodology. One language often associated with this style is Eiffel, which has found use in education, as illustrated in [26]. While languages such as Eiffel enjoy significant tool support, the focus is typically on the specification rather than on the analytical reasoning involved in verification. In contrast, the RESOLVE approach supports both specification and mathematical reasoning through verification conditions, while at the same time including tool support for rapid feedback to students.

One idea that frequently appears in the context of teaching formal methods is that of run-time checking for violations of assertions or contracts. Unfortunately, executable specs can easily blur the distinction between mathematical model and concrete implementation. For example, in [27], the author provides a set of exercises that require students to implement an instructor-created ADT that has been specified axiomatically. The students begin by synthesizing a Java class interface based on the provided axioms and annotate this class with object-oriented assertions translated from the axioms. A third-party open-source tool is used, which processes the assertion-annotated Java source code and generates instrumented code for checking the assertions at run time. Inheritance is used in the Java-class implementations. In contrast, specifications in our approach are model-based and they employ types from mathematical theories as models for object (and variable) values. Feldman points out that, in the context of axiomatic specifications, “it is difficult to express in this formalism the fact that the values corresponding to all other keys [in a map] have remained the same” after the call of a put method. It is an important advantage that, using model-based specification, expressing this fact (and all others we’ve encountered) is easy. While Feldman mentions “class invariants” several times, describing important technical aspects relating to them, the paper provides no specific example as it might appear in an exercise. Hence, there is no example of using: an internal contract that specifies how a variable’s internal representation maps to its external abstract value, or an internal contract for specifying the representation invariant. These internal contracts are central to the assignments given at Clemson or Ohio State, for example, which ask students to implement a given component using model-based specifications of all the involved components.

In [28], the authors describe teaching students how to develop contracts for operations (pre- and post-conditions)

using their tool called ProVIDE. The student is asked to provide a post-condition for an operation and the tool aids the student with constructing an appropriate precondition. We are not asking our students to develop contracts; our course materials provide the contracts, and students evaluate and use these contracts.

VII. SUMMARY

In principle, any formal specification language, such as those summarized in [29] could be used to describe contracts. Some of these languages are tailored to particular programming languages (e.g., JML) and some others (e.g., Z) are more general. However, few specification languages have been designed for specifications to be used in mechanical reasoning about code correctness. Even fewer have been conceived and developed to teach component-based software engineering for undergraduate computing students. Almost none has received multiple decades of undergraduate educational experimentation and tuning as has RESOLVE. This paper illustrates how it can serve as an ideal vehicle in a range of computing courses, starting from introductory courses to advanced software engineering courses with or without using any popular programming language in conjunction. It also shows that the language is conducive to presenting software engineering and reasoning concepts in a few lectures to a few weeks at a variety of educational institutions with differing curricular constraints. The plethora of materials, exercises, activities, and assignments available make it easy for educators to adapt and customize.

ACKNOWLEDGMENT

We gratefully acknowledge the members of the RESOLVE Software Research Groups (RSRG) at Clemson and Ohio State. Many individual members of our research groups contributed significantly to the ideas contained in this paper. We also acknowledge that this research is funded in part by U. S. National Science Foundation (NSF) grants IUSE-1611714, DUE-1022941, SHF-1161196, and DUE-1610957. Any opinions, findings, conclusions, or recommendations expressed here are those of the authors and do not necessarily reflect the views of the NSF.

REFERENCES

- [1] Svetlana Drachova. 2013. Teaching and Assessment of Mathematical Principles for Software Correctness Using a Reasoning Concept Inventory. Ph.D. Dissertation. Clemson University.
- [2] Svetlana V. Drachova, Jason O. Hallstrom, Joseph E. Hollingsworth, Joan Krone, Rich Pak, and Murali Sitaraman. 2015. Teaching Mathematical Reasoning Principles for Software Correctness and Its Assessment. *Trans. Comput. Educ.* 15, 3, Article 15 (August 2015), 22 pages. DOI=10.1145/2716316 <http://doi.acm.org/10.1145/2716316>
- [3] Murali Sitaraman and Bruce W. Weide. 1994. Special Feature on RESOLVE, ACM SIGSOFT Software Engineering Notes, October 1994, 21-68.
- [4] Murali Sitaraman, et al., “Building a Push-Button RESOLVE Verifier: Progress and Challenges,” *Formal Aspects of Computing* 23, 2011, 607-626.

- [5] "OSU CSE Components – API Specification." (31 January 2016) Retrieved from <http://web.cse.ohio-state.edu/software/common/doc/>
- [6] Long, T.J., Weide, B.W., Bucci, P., and Sitaraman, M., "Client View First: An Exodus From Implementation-Biased Teaching", Proceedings of the 30th SIGCSE Technical Symposium on Computer Science Education, ACM Press, 1999, 136-140.
- [7] Paolo A.G. Sivilotti and Matthew Lang. 2010. "Interfaces first (and foremost) with Java". In Proceedings of the 41st ACM technical symposium on Computer science education (SIGCSE '10). ACM, New York, NY, USA, 515-519. DOI=<http://dx.doi.org/10.1145/1734263.1734436>
- [8] Bucci, P., Long, T.J., and Weide, B.W., "Do We Really Teach Abstraction?", Proceedings of the 32nd SIGCSE Technical Symposium on Computer Science Education, ACM Press, 2001, 26-30.
- [9] Charles T. Cook, Svetlana Drachova, Jason O. Hallstrom, Joseph E. Hollingsworth, David P. Jacobs, Joan Krone, and Murali Sitaraman. 2012. A systematic approach to teaching abstraction and mathematical modeling. In Proceedings of the 17th ACM annual conference on Innovation and technology in computer science education (ITiCSE '12). ACM, New York, NY, USA, 357-362. DOI=<http://dx.doi.org/10.1145/2325296.2325378>
- [10] Jason O. Hallstrom, Cathy Hochrine, Jacob Sorber, and Murali Sitaraman. 2014. An ACM 2013 exemplar course integrating fundamentals, languages, and software engineering. In Proceedings of the 45th ACM technical symposium on Computer science education (SIGCSE '14). ACM, New York, NY, USA, 211-216. DOI=<http://dx.doi.org/10.1145/2538862.2538969>
- [11] Bertrand Meyer. Applying Design by Contract. Computer, 25(10):40–51, October 1992.
- [12] B. W. Weide, W. F. Ogden, and M. Sitaraman. Sept. 1994. Recasting Algorithms to Encourage Reuse. In IEEE Software, vol. 11, no. 5, pp. 80–88. DOI=10.1109/52.311066.
- [13] "C343 Standard C++ Components." (31 January 2016) Retrieved from <http://pages.iu.edu/~jholly/C343/Notes/ComponentSpecs/StandardComponents.html>
- [14] Marcia D. Dixon. 2010. Creating effective student engagement in online courses: What do students find engaging? Journal of the Scholarship of Teaching and Learning, 10(2):1-13, June 2010.
- [15] Charles T. Cook, Heather Harton, Hampton Smith, and Murali Sitaraman. 2012. Specification Engineering and Modular Verification Using a Web-Integrated Verifying Compiler. Proc. 34th International Conference on Software Engineering, IEEE/ACM, 2012, 1379-1382.
- [16] Charles T. Cook, Svetlana Drachova, Yu-Shan Sun, Murali Sitaraman, Jeff Carver, and Joseph E. Hollingsworth. 2013. Specification and Reasoning in SE Projects Using a Web-IDE. Proc. 26th Conference on Software Engineering Education and Training, IEEE, 2013.
- [17] Michelle Cook, Megan Fowler, Jason O. Hallstrom, Joseph E. Hollingsworth, Matthew P. Pfister, Tim Schwab, Yu-Shan Sun, and Murali Sitaraman. 2017. Pinpointing Student Obstacles to Logical Reasoning about Code. Technical Report RSRG-17-01, School of Computing, Clemson University.
- [18] <http://www.cs.clemson.edu/group/resolve> (5 Sept 2014).
- [19] Douglas Baldwin, Henry M. Walker, and Peter B. Henderson. 2013. The roles of mathematics in computer science. ACM Inroads 4, 4 (December 2013), 74-80. DOI: <http://dx.doi.org/10.1145/2537753.2537777>
- [20] John P. Dougherty, Joseph E. Hollingsworth, Joan Krone, and Murali Sitaraman. 2016. Mathematical Reasoning in Computing Education: Connecting Math We Teach with Writing Correct Programs (Abstract Only). In Proceedings of the 47th ACM Technical Symposium on Computing Science Education (SIGCSE '16). ACM, New York, NY, USA, 707-707. DOI: <https://doi.org/10.1145/2839509.2850491>
- [21] Jeannette M. Wing. 2000. Invited Talk: Weaving Formal Methods into the Undergraduate Computer Science Curriculum. In Proceedings of the 8th International Conference on Algebraic Methodology and Software Technology (AMAST '00), Teodor Rus (Ed.). Springer-Verlag, London, UK, UK, 2-9.
- [22] Doug Baldwin, "Math-thinking-1 – Mathematical reasoning in CS curricula", at: <http://mail.geneseo.edu/mailman/listinfo/math-thinking-1/> (5 Sept 2014).
- [23] Raymond Boute. 2009. Teaching and practicing computer science at the university level. SIGCSE Bull. 41, 2 (June 2009), 24-30. DOI=<http://dx.doi.org/10.1145/1595453.1595458>
- [24] Gene Fisher and Corrigan Johnson. 2016. Making Formal Methods More Relevant to Software Engineering Students via Automated Test Generation. In Proceedings of the 2016 ACM Conference on Innovation and Technology in Computer Science Education (ITiCSE '16). ACM, New York, NY, USA, 224-229. DOI: <http://dx.doi.org/10.1145/2899415.2899424>
- [25] Erik Poll. 2009. Teaching Program Specification and Verification Using JML and ESC/Java2. In Proceedings of the 2nd International Conference on Teaching Formal Methods (TFM '09), Jeremy Gibbons and José Nuno Oliveira (Eds.). Springer-Verlag, Berlin, Heidelberg, 92-104. DOI=http://dx.doi.org/10.1007/978-3-642-04912-5_7
- [26] Paige R.F., Ostroff J.S. (2004) Specification-Driven Design with Eiffel and Agents for Teaching Lightweight Formal Methods. In: Dean C.N., Boute R.T. (eds) Teaching Formal Methods. TFM 2004. Lecture Notes in Computer Science, vol 3294. Springer, Berlin, Heidelberg
- [27] Yishai A. Feldman. 2005. Teaching quality object-oriented programming. J. Educ. Resour. Comput. 5, 1, Article 1 (March 2005). DOI=<http://dx.doi.org/10.1145/1101670.1101671>.
- [28] Timothy S. Gegg-Harrison, Gary R. Bunce, Rebecca D. Ganetzky, Christina M. Olson, and Joshua D. Wilson. 2003. Studying program correctness by constructing contracts. In Proceedings of the 8th annual conference on Innovation and technology in computer science education (ITiCSE '03), David Finkel (Ed.). ACM, New York, NY, USA, 129-133. DOI=<http://dx.doi.org/10.1145/961511.961548>
- [29] Victor Klebanov, et al., The 1st Verified Software Competition: Experience Report. In FM 2011: Formal Methods, Proceedings 17th International Symposium on Formal Methods, Springer LNCS 6664, June 2011, 154-168.