

CDM: Teaching Discrete Mathematics to Computer Science Majors

KLAUS SUTNER *

Carnegie Mellon University

Pittsburgh, PA 15213

May 17, 2005

Abstract

CDM, for Computational Discrete Mathematics, is a course that attempts to teach a number of topics in discrete mathematics to computer science majors. The course abandons the classical definition-theorem-proof model and instead relies heavily on computation, as a source of motivation and also for experimentation and illustration. The emphasis on computational issues is particularly attractive to computer science majors and increases their involvement and participation.

Categories: K.3.1 Computers and Education : Computer Uses in Education, F.3.1 Logics and Meanings of Programs : Specifying and Verifying and Reasoning about Programs, G.2.1 Discrete Mathematics : Combinatorics, I.1.3 Symbolic and Algebraic Manipulation : Languages and Systems

1 History and Motivation

In the mid 1990's Dana Scott decided that the time was ripe to bring modern computational software into the computer science curriculum at Carnegie Mellon. He started the "ModMath Project" whose goal was to systematically exploit software to teach discrete mathematics to computer science majors in a novel way that was custom-designed for the specific needs of this group of students. The central idea was to push computation into the foreground, both in the sense of presenting mathematical content from a computational point of view and by using computer algebra as a mathematical sandbox in which the students can experiment, find examples and counterexamples and sharpen their intuition.

The classical definition-theorem-proof cycle is replaced by a different loop that begins with some problem or question, based on the current state of mathematical

*Supported in part by NSF CCF-0113919.

knowledge of the student. Computational examples and experiments lead to a better understanding of the problem and, hopefully, to ideas and conjectures how the issue might be tackled. The conjectures are then rigorously proven, added to the current level of knowledge and the cycle begins anew. Of course, often some of the ideas will not work out as planned and there will be some back-and-forth between conjecture formation and proofs, always with the potential aid of computation as a source of inspiration, examples and counterexamples. The role of the proof here is two-fold. First, it helps to ensure that the observations hold in general and are not just some peculiar coincidence. Second, the proof should provide some insight into why the observations hold, insight that could not be gained by just observing the data. This empirical and experiential approach to mathematics can help greatly in coming to a solid understanding of the relevant concepts. Moreover, it is fun to make small discoveries of one's own and then try to formalize and establish them in a rigorous way.

One example from trenches: some years ago a homework assignment asked students to determine the number of binary sequences of length n that contain no two consecutive 0's. The students knew how to tackle simple recurrence equations and were in particular familiar with Fibonacci number. However, no further hints were given as to how the problem might be solved. Most students used a program to compute the number of such 00-free sequences for a few small values of n by brute force, then correctly conjectured the answer (the Fibonacci number F_{n+2}) and gave a proof by induction. Brute-force here means that one first generates all binary sequences of length n and then eliminates the ones containing 00, a very simple task in our standard computational environment. Indeed, in Mathematica, augmented by a macro package for the course, the whole program can be written as a "one-liner":

```
ZZfree[n_] := DeleteCases[Tuples[{0,1}, n], {___,0,0,___}];
```

While this solution is perfectly acceptable and was indeed what the instructor had intended, a few students worked a bit harder initially and wrote a recursive program that directly generates only 00-free sequences. The crucial insight here is that sequences of length n can be obtained by appending suffixes 1 and 10 to shorter sequences of lengths $n - 1$ and $n - 2$, respectively. The induction proof is then essentially replaced by a correctness argument for the generating program: one has to make sure that all 00-free sequences are so generated and that there are no duplicates. The connection between proofs and programs – where the latter notion is perfectly palatable to the students while the former often is not – is one of the major themes of the course.

For a number of years Scott, Ed Clarke, Phil Miller and myself taught a variety of ModMath courses at an introductory level that were organized around these ideas, see [9]. Alas, the students' response was decidedly mixed: while the stronger students embraced the new technology and used it to great effect, a large number of weaker students felt intimidated and overwhelmed by the machinery. For them, the extra burden of having to learn how to cope with a specific computing environment never seemed to be quite compensated for by the benefits that all of us envisioned. Ultimately, the project went into hibernation and the curriculum returned to its traditional patterns. Predictably, the computer science majors resumed their standard gripes about too much abstraction and a general lack of connection between the mathematics and their core interests.

Motivated by the enormous progress in symbolic computing software and in partic-

ular the emergence of document-centric computing environments we decided in 2002 that another attempt should be made to redesign the discrete mathematics component of the computer science curriculum. We created a new course titled CDM, for Computational Discrete Mathematics. Unlike with its ModMath predecessors, CDM is not required but rather is offered as a choice towards the fulfillment of the mathematics requirement. The students typically are in their second year and have taken two prerequisite courses in discrete mathematics. The course is presently capped at 20 students, most of whom are computer science majors though some mathematics majors also enrolled and acquitted themselves very well. CDM will be in its fourth incarnation in the Fall of 2005.

In a drastic change from previous policies, the use of computing equipment is also no longer tightly controlled. By the time they take CDM, the computer science majors at Carnegie Mellon tend to be extremely opinionated about the relative merits and demerits of various systems, programs and languages. There appears to be little benefit in forcing them into one particular environment when they are already more comfortable and experienced in another. Specifically, while Mathematica [15] is used as the main workhorse in the course from the perspective of the instructor, and some of the supplementary material is posted in the form of Mathematica notebooks, there is no requirement for its use. In the past, students have successfully employed C, C++, Java, Perl and Maple to tackle the computational aspects of various assignments. As a matter of fact, not only is there no requirement for the use a specific software system, there is not even a requirement for the use of computation at all. The students are free to handle the given assignments any which way they like—as long as the results are elegant and correct. In principle, a purely theoretical approach might succeed, though it often would be too time-consuming and tedious for serious consideration. The students are quick to realize this, and choose their weapons accordingly. Note that this freedom on the side of the students places somewhat of a burden on the course staff, in particular in office hours and in the assessment of student work, but in our experience the results are well worth the effort.

One last comment about the role of programs in CDM. It was pointed out by Bill Thurston in his very thoughtful article [14] that writing correct programs is a difficult task, more difficult in some ways than pure mathematics:

The standard of correctness and completeness necessary to get a computer program to work at all is a couple of orders of magnitude higher than the mathematical community's standard of valid proofs.


We largely agree with this assertion, though we believe that it actually works in our favor in this course. Unlike with Thurston's work in low-dimensional topology, the computations required in CDM are mostly rather basic. Much of the burden of writing correct programs is alleviated by the ready availability of functioning, well-tested code for these fundamental operations. Hence, the challenge Thurston refers to can actually be exploited in a constructive way: dealing with programs and the results of their computations produces better insight into the underlying mathematical principles. Rather than removing precision and careful reasoning it emphasizes them. It leads to a concrete, experiential and integrated understanding of the concepts and their applications, rather than the more abstract, mechanistic and isolated approach fostered by the traditional definition-theorem-proof method.

2 Exploiting Computation

In this section we will give a few characteristic examples that highlight the CDM approach to discrete mathematics. Most of the material is of considerable sophistication and would be difficult to transmit in a more traditional setting. A more complete listing of topics can be found at the website, complete with lecture slides and accompanying material, see [13]. While some of the material at the site can be expected to undergo changes and adjustments the core of the course is now stable. At any rate, it will give a good indication of the philosophy of the course.

2.1 Closure Operations and Invariants

Here is a simple puzzle based on two-colored pebbles that are placed inside three overlapping circles as shown below. Originally, all pebbles are blue side up, and the goal is to determine whether one can bring them into the shown configuration where the central pebble is red side up.

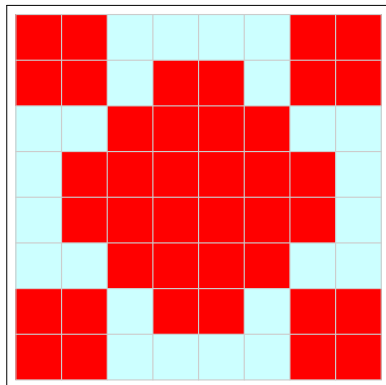


The only permissible operations are to flip all four pebbles in one circle over, or to arrange all pebbles in one circle red side up. Since there are only $2^7 = 128$ possible configurations it is easy to compute by brute force all reachable configurations and to verify that the target configuration is not among them.

However, the assignment requires an explanation, not just a yes/no answer. So the program alone will not do: it provides no insight whatsoever as to why the target configuration fails to be reachable. In the second step, it is necessary to inspect the reachable configurations and to determine some appropriate invariant: some property that holds of the initial configuration, is preserved by the admissible operations and is violated by the target configuration. A slightly more ambitious project is to give a full characterization of the reachable configurations, a task that seems to require some linear algebra over the two-element field.

The operation of flipping a collection of pebbles easily leads to a class of related but more difficult problems. For example, we can think of the nodes of an arbitrary graph as the pebbles, and “flipping” one node will affect not just the node but also all its neighbors. Given some starting configuration we would like to understand all reachable configurations. Even if we limit our attention to grid graphs the questions become quite complicated. The nodes of the graph here are the squares of a generalized chessboard and each square is considered adjacent to its north, south, west and east

neighbor. The following picture shows which pebbles have to be flipped to turn the all-blue configuration on an 8 by 8 chessboard to a an all-red configuration. Note how the nodes are chosen so that each node of the graph has an odd number of neighbors, including itself, in the chosen set.



Needless to say, the solution was not found by searching through all possible 2^{64} configurations but is based on, firstly, a reduction of the search space to dimension 8 and, secondly, the use of linear algebra on the smaller space to avoid search altogether. The solution is essentially unique in this case. To determine the number of solutions for other grids leads to rather interesting problems related to divisibility properties of Fibonacci polynomials over the two-element field, see [10].

2.2 Pólya and Burnside

Students in CDM all have encountered basic group theory in previous courses, but never seem to have found the topic particularly relevant. A concept that they are better acquainted with and more appreciative of is the representation of equivalence relations in a way suitable for computation. For example, they are all familiar with the use of the union/find algorithm to construct mazes. We connect the two ideas and use symmetry groups to limit the size of search spaces in some games such as checkers or nine men's morris, apparently an excellent motivation for computer science majors. Group actions are introduced as yet another way to represent certain types of equivalence relations, albeit in a somewhat abstract fashion. It comes as a pleasant surprise to most students that the abstraction has an immediate computational payoff: otherwise difficult question about the number of equivalence classes can be reduced to questions about polynomials. The brute-force approach where all possible configurations are generated first and the equivalent ones identified later becomes demonstrably useless for all but the smallest examples – though it does help greatly to build up some intuition for the basic ideas in Pólya-Redfield counting.

The students are given access to software that performs the necessary cycle-decomposition of elements of a permutation group and generates the cycle index polynomials. For example, the invocation

```
P4 = CycleIndexPolynomial[DihedralGroupSG[4]];
SubstitutePolynomial[P4, 6, 2, Full -> True]
```

will duly produce the polynomial

$$\frac{1}{8} \left((c_1 + c_2)^4 + 2(c_1 + c_2)^2 (c_1^2 + c_2^2) + 3(c_1^2 + c_2^2)^2 + 2(c_1^4 + c_2^4) \right)$$

In class, only two-dimensional examples such as the dihedral group above are discussed. For the assignments, the students are required to deal with symmetries in three-space. In particular, they have to discover the structure of the octahedral group and find a nice set of generators and relations. The actual calculations are easy to carry out – if one has the ability to manipulate polynomials of several hundred terms. Somewhat unexpectedly, a number of students elected to use systems other than our standard environment for polynomial arithmetic. In fact, some even implemented the requisite symbolic algorithms on their own, exploiting work from other courses.

There is some anecdotal evidence for the strength of this approach: after last year’s notes were posted on the web a researcher in the Netherlands working on a sophisticated checkers problem sent email to inquire about a number of problems in the actual implementation of search space reducing techniques. Some of the questions will be handed off to the students in the next incarnation of the course.

2.3 Logic and Provers

Finding convincing motivation for the use of a formal languages such as predicate logic and formal systems that describe and regulate proofs is always a major challenge in any discrete mathematics course. While most students would concede that Fermat’s original description of his infamous theorem

There do not exist four numbers, the last being larger than two, such that the sum of the first two, both raised to the power of the fourth, are equal to the third, also raised to the power of the fourth.

is inferior to, say,

For no positive integers x, y, z and n , where $n > 2$: $x^n + y^n = z^n$.

very few would agree that the step to

$$\neg \exists x, y, z, n (n > 2 \wedge x^n + y^n = z^n)$$

provides any further advantages. However, the last version is an expression in a formal language, albeit one whose validity over the natural number turned out to be exceedingly difficult to verify. But for many other, less complicated assertions, perhaps over other domains, validity can be established with the help of programs that operate on formal expressions. Indeed, (semi-) automatic verification of hardware and software system using formal specifications in appropriate logics has become a major industry and commercial programs use rather sophisticated tools from proof theory and model theory, see [4].

As it turns out, even in propositional logic there are some challenging computational problems to be solved. In CDM we begin with a discussion of various normal forms (negation normal form, conjunctive and disjunctive normal form) and algorithms that convert a given propositional formula into these forms. Then we introduce the Davis-Putnam algorithm and discuss a few of the more obvious heuristics that can be used

to guide the selection of appropriate literals in the splitting step. It is remarkable that even a simple implementation of the algorithm in a high-level environment such as Mathematica can typically handle formulae with thousands of variables. This also provides an opportunity to put the theory of computational hardware into perspective.


Of course, for predicate logic more powerful tools are needed and the algorithms tend to be burdened with numerous technical details. We use Analytica II, a natural deduction based theorem prover that exploits algebraic simplification to establish results in what might be loosely called 19th century mathematics. Analytica produces eminently human-readable proofs, with formulae typeset in near publication quality, see [1, 2]. For example, the system can establish the Bernstein approximation theorem with only minimal help from the user. More importantly, Analytica has extensive tracing facilities that make it possible for the students to observe its attempts at proving a theorem. For example, the following command will produce a proof of the fact that a function g is surjective whenever $g \circ g$ is.

```
Prove[ imp[ all[ x, some[ y, x == g[g[y] ] ] ],
         all[ x, some[ y, x == g[y] ] ] ] ]
```

We hasten to point out that the purpose of the use of Analytica is to demonstrate the importance of formal systems. It is not intended as a mathematics tutor of any kind and would certainly fail miserably in this capacity. It does, however, provide many examples of computational problems in propositional and predicate logic that would otherwise appear hopelessly artificial and pointless. Skolemization of formulae in predicate logic is one case in point here. In the example above, the given input formula $\forall x \exists y (x = g(g(y))) \Rightarrow \forall x \exists y (x = g(y))$ is first converted by the system into negative Skolemized form: $v_1 = g(g(f_{v_1})) \vdash c_1 = g(v_2)$. The latter is then established by finding suitable values for the variables v_1 and v_2 while the Skolem function c_1 and f_{v_1} remain uninterpreted, a process rather similar to natural reasoning. Analytica also has a mechanism to “un-skolemize” a formula which adds greatly to the human-readability of the generated proof. Another example is a simplification strategy in Analytica which attempts to reduce the logical complexity of a sequent by eliminating “superfluous” parts. It is a very healthy exercise to discuss implementations of this method and to rigorously establish its soundness.

2.4 Finite State Machines

Students in CDM have usually had some cursory exposure to finite state machines, mostly from the more pragmatic point of view of a programmer or perhaps in some simple string matching applications. While applications of state machines provide ample motivation, they do tend to obscure the mathematics a bit. To exhibit the mathematical structure we start with a simple problem: Construct a deterministic finite state automaton that recognizes numbers written in base B that are divisible by m . One attractive feature of these machines is that they can be represented by labeled diagrams that are easily understandable to humans. The example below shows the automaton in diagrammatic form for $m = 5$ when the input is given in binary. In algebraic form, the transition function is given by $\delta(p, a) = 2p + a \pmod{5}$ where the state set consists of the modular numbers \mathbb{Z}_5 .



The diagram is just one instance of a natural automaton that uses the modular numbers \mathbb{Z}_m as state set and essentially computes the value of the input modulo m by using a Horner scheme to evaluate the corresponding polynomial. However, in many cases the resulting machine fails to be minimal: there is a smaller DFA that decides divisibility by m using fewer than m states. This effect can be observed easily in our environment by constructing the natural Horner automata and minimizing them for some values of m and B . Here is a table for $m \leq 12$ and $B \leq 16$ that demonstrates this effect. The table can be generated easily with two lines of code.

	B															
m	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	3	3	2	3	3	2	3	3	2	3	3	2	3	3	2	3
4	3	4	4	2	4	3	4	2	4	3	4	2	4	3	4	2
5	5	5	5	5	2	5	5	5	5	2	5	5	5	5	2	5
6	4	4	3	4	6	2	6	4	3	4	6	2	6	4	3	4
7	7	7	7	7	7	7	2	7	7	7	7	7	7	2	7	7
8	4	8	8	3	8	5	8	2	8	5	8	3	8	5	8	2
9	9	9	3	9	9	4	9	9	2	9	9	4	9	9	4	9
10	6	10	10	6	3	6	10	6	10	2	10	6	10	6	3	6
11	11	11	11	11	11	11	11	11	11	11	2	11	11	11	11	11
12	5	5	5	4	12	3	12	4	5	7	12	2	12	7	5	4

A lot of information can be extracted from the table. For example, for prime m it appears that the minimal machine always has size m unless B is a multiple of m , in which case the size drops to 2. To establish this and other observations it is best to construct the minimal automaton preserving the natural state set as a partition of the the modular numbers. Incidentally, this provides another opportunity to discuss implementations of equivalence relations, in this case an implementation built on the notion of a kernel relation. At any rate, careful inspection of these partitions together with some thought about the structure of the Horner automaton leads to a connection to solutions sets for a class of modular equations and, ultimately, to a complete characterization of the minimal automaton for divisibility. Needless to say,

our standard computational environment can handle the requisite modular arithmetic as easily as it deals with the finite state machines. Thus we can solve the modular equations, use the solutions to construct a finite state machine and then verify that the machine is indeed minimal. The seamless integration of computation in different areas of mathematics is of great importance to our approach.

Looking for generalizations one can ask about different number representations. For example, we may consider reverse base B where the least significant digit comes first, or Fibonacci representation. We know characterizations for some of these number systems, but mostly there are lots of interesting open problems.

3 The Afterlife

The primary purpose of mathematics-related courses in a computer science curriculum is to provide a foundation for material to be presented later in the curriculum, material that relies heavily on these foundations and that cannot be properly understood without them. Needless to say, the consumer courses may be separated from the preparatory mathematics courses by several semesters. As a consequence, it is very important that the students retain access to the material, at least for the length of their stay at the university – though one would naively hope that they will actually be able to make use of the material beyond the acquisition of a degree. These concerns are addressed in an NSF-funded project with the suggestive title “Course Capsules: Persistent, Personalized Courseware (CCaps)”, see [8]. The goal of the CCaps project is to develop mechanisms that make it possible to package the whole content of a course, including student contributions, into an electronic repository that will be available indefinitely over the the web for use by the student. This is in contrast to course management systems such as the popular Blackboard system which focus on delivery, but largely ignore archiving and long term access. Persistent, intelligent access is a particularly pressing issue for a course like CDM which exists exclusively in computer-held form (there is no textbook) and whose content is supposed to be applied later on in a number of other courses.

The target document format in CCaps is OMDoc, a specialized XML application that allows one to represent mathematical knowledge structures and that extends the well-known OpenMath standard in several ways, see [6]. We have developed a number of software tools that help in the markup process, notably a PowerPoint to OMDoc converter, a Mathematica to OMDoc converter and a powerful emacs mode to handle the final editing without endangering the mental health of the editor, see [7]. Alas, it has turned out that the effort to perform the actual markup is rather significant, our tools notwithstanding, so progress has been slower than originally hoped for.

Ultimately, the students will be able to turn to their CDM capsule to gain intelligent access to the material. The semantic markup will provide the ability to index and search on content rather than just performing the customary string matching. For example, all uses of the notion of “equivalence relation” will be easily identified and collated. It will be possible to extract all theorems and lemmata about equivalence relations or, alternatively, find their use in concrete algorithms. Moreover, the student will have access to examples and exercises that will her allow to regain mastery of the relevant concepts. The interactive nature of mathematics documents on the web

is rather significant here: the definition of an unknown or forgotten term is only a few mouse-clicks away, hopefully complete with illuminating examples and further pointers. Assessment at least at the superficial level of multiple-choice problems and simple computational exercises can also be fully automated.

Another problem is created by the computational approach of CDM itself: at least during their tenure at Carnegie, the students should be able to rerun the sample computations in the standard, Mathematica-based environment. This seemingly innocent requirement actually causes substantial problems as anyone will recognize who has ever tried to compile and run a, say, C++ program downloaded from some web site. The problem of ever changing compilers and libraries is somewhat ameliorated by using a high-level environment such as Mathematica, where changes to the language between versions are relatively minor (though the same cannot be said for the underlying mathematical algorithms). However, macro packages written even in a stable environment such as Mathematica are notoriously unreliable software artifacts, even when they come from reputable repositories such as Wolfram's Information Center. In our case, the core package contains several hundred functions dealing with combinatorics, graph theory, finite state machines, syntactic semigroups and cellular automata. The package has grown by accretion over the years and is now rather solid, see [12] for a recent application and [11] for a predecessor version of the system.

In light of the persistence problem we have developed a special packaging mechanism that will make it possible to preserve the old functionality should changes be made in the future. The packaging system also provides for automatic testing of the code. More important for teaching applications is the extensive help browser that comes with the system: context dependent help is available for each command and there are numerous sample computations that demonstrate the use of the machinery, from simple one-liners to complicated applications more typical of research applications. Ultimately one would hope that web-based mathematics services will alleviate this problem, but at present local solutions are essential.

4 Some Preliminary Conclusions

As mentioned already, CDM currently has an enrollment cap of 20 students and the course has been taught only three times thus far. One important goal of these first few rounds was to find a compelling selection of material. For example, regarding the central notion of computability we felt obliged to provide some historical background and an introduction into the classical theory of computability. On the other hand, for the actual computations, some understanding of the more recent theory of computational complexity is necessary. Striking a balance between these competing interests without introducing inappropriate levels of technical difficulty turned out to be quite challenging. A typical student comment from last Fall seems to indicate that the proper balance has been found: "This class covers a lot of ground, but it isn't impossible to understand, and it is extremely interesting." Another central concern was to find a mode of presentation that engages the students to the point where they are prepared to invest considerable time and effort in mastering the material. Again quoting a student from last Fall: "[The instructor] could make studying the mathematics of a rock interesting."

Another positive indicator is the quality of the work submitted by the students. Overall, it is remarkably high – though one has to bear in mind that the students are self-selecting and presumably somewhat better motivated than their peers. For example, the best solution to the characterization of reachable pebble configuration from the problem in section 2 is due to a student in the course. Another student found a novel way to establish the fact that over \mathbb{Z}_2 every function is polynomial. A few of the students have gone on to take more advanced courses such as constructive logic and seem to be well prepared, though this evidence is currently anecdotal.

To obtain a slightly less informal understanding of student response to the course we used a standard faculty course evaluation instrument developed at the Eberly Center for Teaching Excellence, see [3]. In a nutshell, the course rating for CDM was 4.52, 4.89 and 4.7, the instructor rating 4.61, 5.00 and 4.95, on a scale from 1 to 5. Both ratings are significantly higher than those for comparable courses in our curriculum.

Overall we feel that the first phase of our project has been completed successfully and that we can now turn to the crucial second step of turning the material into a course capsule as described in the previous section. To this end it will be necessary to rethink and make explicit the learning objectives for each of the course topics. A more elaborate system of completed and partially completed exercises has to be developed, along with better and scalable machinery for assessment. At this point, most of the course material still exists in a strictly presentation oriented format and needs to be converted into a semantically rich format. Since full semantic markup has proven to be an extremely labor intensive task, efforts are currently under way to secure funding for a joint project with the Open Learning Initiative at CMU, see [5]. The OLI project offers in particular a lot of experience in providing reliable, persistent web services as well as assessment and cognitive task analysis.

The most recent version of course notes and a number of notebooks are available at [13]. The Mathematica package used in CDM is also available at the site, and alternatively at <http://library.wolfram.com/infocenter/MathSource/4603>. The code contains extensive on-line help and a number of sample notebooks. Additional material such as assignments, tests, a printed manual and various supplementary notebooks will be made available to interested parties.

References

- [1] A. Bauer, E. Clarke, and X. Zhao. Analytica—an Experiment in Combining Theorem Proving and Symbolic Computation. *Journal of Automated Reasoning*, 21(3):295–325, 1998.
- [2] E. Clarke, M. Kohlhase, J. Ouaknine, and K. Sutner. Analytica 2. In *Proc. Calculemus, Rome*, 2003.
- [3] Eberly Center for Teaching Excellence. <http://www.cmu.edu/teaching/eberlycenter>.
- [4] M. Huth and M. Ryan. *Logic in Computer Science: Modelling and Reasoning about Systems*. Cambridge UP, 2000.
- [5] Open Learning Initiative. <http://www.cmu.edu/oli>.

- [6] M. Kohlhase. OMDoc: An open markup format for mathematical documents. <http://www.mathweb.org/omdoc/index.html>, 2002. Version 1.2.
- [7] M. Kohlhase and A. Kohlhase. Dissolving the author's dilemma. In *MKM 2004*, 2004.
- [8] M. Kohlhase, K. Sutner, P. Jansen, A. Kohlhase, P. Lee, D. Scott, and M. Sudszik. Acquiring mathematical content in an academic environment. In *MathML Conference*, 2002.
- [9] P. Miller and K. Sutner. Exploiting computer algebra systems in computer science courses. *ACM SIGCSE Bulletin*, 29(1):377–378, 1997.
- [10] K. Sutner. On σ -automata. *Complex Systems*, 2(1):1–28, 1988.
- [11] K. Sutner. Implementing finite state machines. In N. Dean and G. Shannon, editors, *Computational Support for Discrete Mathematics*, volume 15, pages 347–365. DIMACS, 1994.
- [12] K. Sutner. `automata`, a hybrid system for computational automata theory. In J.-M. Champarnaud and D. Maurel, editors, *CIAA 2002*, pages 217–222, Tours, France, 2002.
- [13] K. Sutner. Computational Discrete Mathematics. <http://www.andrew.cmu.edu/course/15-354>, 2004.
- [14] W. Thurston. On proof and progress in mathematics. *Bull. Amer. Math. Soc.*, 30(2):161–177, 1994.
- [15] S. Wolfram. *The Mathematica Book*. Cambridge University Press, 2002.