

Proposal Deadline

Email me a brief description of your choice of project and some preliminary thoughts about what you are going to do by **Monday, November 25**.

If you plan to partner with someone, submit just one solution.

AFS Directories

All submissions will be to directories

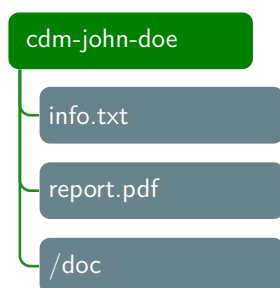
</afs/andrew.cmu.edu/course/15/354/handin/<andrew-id>>

There is plenty of space, but do not submit gigabytes of data.

What To Do


By university policy, the official deadline for all projects is the last day of classes, **December 7, 24:00**. Grades are due 12/18 and the class is small, so I can handle late submissions. Talk to me if 12/07 does not work for you. Place your `tgz` file into the AFS directory above, see below for specs. Your submission will contain multiple files, possibly including `pdf`, `code`, `sample data files`, and so on—make sure that everything is nicely organized. Submit a file `cdm-john-doe.tgz` and include a text file `info.txt` that explains the file structure of your submission (and anything else you want me to know). If there is code, make sure it compiles under a standard Andrew linux environment—no `compile`, no `credit`. I am not going to edit code in order to get it to compile, and I am not going to install special software to get things to work.

If your project is a paper report, your `tgz` file should unpack like so:



Supporting documents (and in particular the **paper under review**) go into `/doc`.

If your project involves code, adhere to the following subdirectory structure. Make sure to provide compelling examples in `/dat`, if it is not obvious how to construct more examples I will assume that your code fails in general. Your `.tgz` file must unpack like this:



Put a ready-made executable into `/bin`, just in case compilation somehow fails.

Some of the projects below can be handled essentially without recourse to the literature, others involve some reading. Do not procrastinate.

There is no need to prove new results. Of course, it's fantastic if you do come up with something new, but it is **not** required. You are expected to demonstrate a clear understanding of the issues and the ability to describe the problem and its solution.

For some of the larger projects you may think about pairing up with someone. For example, the universal Turing machine problem and the Ehrenfeucht-Mycielsky sequence are best handled with a partner.

If you have a proposal for an alternate topic, make sure to talk to me before you start working—not all topics are suitable.

As you will see, the projects have inherently different levels of difficulty; obviously, if you choose a simple project, you will be expected to produce a particularly nice solution. On the other hand, if you go after one of the hard questions, don't worry if there is no breakthrough, just do the best you can (welcome to the world of research).

If you have any question, talk to us, don't just start out on a wild goose chase.

Paper Reports

For a report, you can read any of the papers I mentioned in class; for example, the Elgot and Mezei paper on transducers, the Rabin/Scott classic, Kleene's paper on regular events. Take a look at [projects](#). Explain the central ideas and contributions in the paper. In particular, explain the proofs and, wherever appropriate, provide examples for the given constructions. Think of this like being on a program committee: someone who reads your report should have a good idea of what is going on in the paper, and should be well-prepared to read and understand the actual paper.

In addition to the papers mentioned in class, you can also pick a recent research paper. To keep things under control, let us limit the range of papers to conference proceedings of ICALP, CIAA, DCFS, LATA, FOCS and STACS for the last 8 years. Not all the papers here relate to CDM, pick one that does. If you have some other paper in mind, talk to me asap.

These projects are solo exercises, pick a paper that is manageable. If there is a clash (highly unlikely), it will be resolved on first-come first-serve basis. It is a good idea to start looking for a suitable paper right now, and spend a bit of time making a good choice.

Your paper should roughly be organized like so:

- Introduction: what is the topic, why is it interesting, what is the state of the art.
- Contribution: what is the specific contribution of this paper.
- Theorems & Proofs: what are the technical results and how are they established, in particular: is everything correct?
- Critique: is the paper well-written, could it be easily improved (e.g. by including examples), are the computational parts reproducible, is the bibliography sufficient?

Take item (3) seriously, just because a result is published does not necessarily mean that it is quite correct. After reading your analysis of the paper, it should be quite straightforward for a mathematically literate reader to go through the actual paper.

Specific Projects

- Cycle Finding Algorithms

Floyd's algorithm is just one example of a memoryless cycle finding algorithm. Implement a number of these algorithms and compare their performance. A good test case are the 256 elementary cellular automata. Try to compute transients and periods for n -bit configurations for n as large as possible and try to classify ECAs accordingly. Long transients and periods are particularly interesting (e.g., the pseudo-random number generator ECA 30 should have very long periods).

- Universal Register Machine

Construct a complete URM, without using the macros from lecture. You might want to write a small compiler that translates the version using macros into a real register machine. Try to make the resulting machine as small as possible and demonstrate its correctness by simulating a few small register machines.

To make interesting simulations feasible, you need to write an "intelligent" interpreter: analyze cycles in the diagram of the machine and speed up their execution by performing giant-steps rather than baby-steps.

The deliverable must be useful for demonstrations: I want to be able to show an actual universal computer perform simple tasks such as addition or perhaps multiplication.

Your code must run on a standard Andrew linux box, specialized software is not acceptable.

- SAT Solver

The Satisfiability problem for propositional logic (rather than first-order logic) is obviously decidable, though it is currently no known whether it can be handled in polynomial time. Implement a simple SAT solver along the lines of the classical Davis/Putnam algorithm [DPLL](#). The point here is not blinding speed but easy traceability. For example, it should be easy to demonstrate the effect of various heuristics in the splitting case.

You can check your algorithm against one of the performers in the SAT competition [SAT Competition](#). This is (for correctness only, not for speed; the latter will only lead to depression on your part).

- Divisibility and Minimal DFAs

Determine the size of the minimal DFA recognizing all multiples of a fixed modulus m in various numeration systems. This has been done for normal radix notation and reverse radix, but the results can probably be improved upon, see [Divisibility](#).

- A Small "Universal" Turing Machine

There is a claim that a 2-state/3-symbols Turing machine can already be universal, see

- [2-3-Turing Machine](#),
- [2-3-TM Prize](#),
- [2-3-TM Proof](#).

As you will see, the alleged proof by Alex Smith has lots and lots of holes, see for example the objections raised at FOM linked at the wiki site.

The ultimate goal here is to determine whether Smith's argument holds water. This is probably too hard, so any significant progress towards a resolution will count. For example, a clean and mathematically precise version of the construction would already be of interest.

- Presburger Arithmetic

Build a decision algorithm for Presburger Arithmetic. There are two parts to this: first, a small library of finite state machine operations that are needed to carry out the constructions on synchronous transducers. Second, an interface that translates a formula into the corresponding machines. This part is optional, if you have a reasonable way to call your library either from the command line or from, say, Python, the user can handle the necessary translation.

As a benchmark, your algorithm should have no problem with the Frobenius problem that we discussed.

- ECA and k -Cycles

Consider the automatic structure $\langle \Sigma^*; \rightarrow \rangle$ where the next step relation is given by an elementary cellular automaton, assuming periodic boundary conditions. Adapt the decision algorithm for automatic structures to handle the problem of finding k -cycles.

This is somewhat easier than Presburger since the corresponding transductions are length-preserving, the challenge here is to handle “large” values of k up to, say, 20.

On input a rule number r , $0 \leq r < 256$, and some cycle length k , your algorithm should return the semilinear set that describes the grid sizes n for which a k -cycle exists.

- Busy Beaver

There has recently been some amazing progress on the infamous Busy Beaver problem for 5 states, see [BB-5](#) and [Busy Beaver](#).

Give a description of the new results, the methods used and evaluate them with a view towards the criticism of earlier efforts in Harland’s paper.

- AES

Implement AES. The goal here is not to produce anything resembling a serious implementation, speed and throughput are irrelevant. Instead, the code has to be useful to demonstrate the behavior of the algorithm, or at least a watered-down version of it. To this end, the implementation must be easily modifiable and produce nice traces, so the only plausible languages are Python and Mathematica.

Again, there is no need to try to streamline things. E.g., the required finite field operations can be as hamfisted as imaginable. Alternatively, use some finite field library as a black box.

- Dropping Balls

There is an old single-player game dating back to the 1960s that has a nice explanation in terms of group actions, see [Trellis Automata](#).

The first three questions in the reference are quite manageable, but the last one (generalize to other trellises) is utterly open-ended.

- The Computational Universe and the Bergers

N. J. Wildberger is a mathematician in Australia, who has a few rather unorthodox ideas that relate to our quest for computational interpretations of mathematics. E.g., he does not believe in the fundamental theorem of algebra, which he likes to refer to as the “fundamental dream of algebra.” His youtube videos seem to be unavailable at this point, take a look at [Wildberger](#).

D. Zeilberger at Rutgers has similarly voiced somewhat contentious ideas, some constructive, some not so much. In particular his paper [Zeilberger](#) titled “What is Mathematics and What Should it Be?” is pushing things a bit too far.

Familiarize yourself with some of their ideas and write a critique.


- Ehrenfeucht-Mycielsky Sequence

The authors defined a deceptively simple infinite binary sequence that apparently has limiting density $1/2$. In fact, computing a few million bits of the sequences shows very rapid convergence. Unfortunately, no proof is known though lots of little facts have been unearthed by a number of people, see [EM-sequence](#) and the web.

Try to make a dent in this problem (don't aim for a full proof, it's probably quite hard).

- ECA 30

Elementary cellular automaton number 30 has amazing properties, it seems to generate random patterns.



You can earn real money (\$30,000) by solving some concrete questions related to ECA 30, see [ECA 30](#). This is probably quite hard, but if you think you might want to take a swipe at it, talk to me first (I am on the prize committee).