PRINCETON ACM / IEEE-CS CHAPTERS
APRIL 2014 JOINT MEETING

Accountable Algorithms

Important decisions about people are increasingly made by algorithms: votes are counted; voter rolls are purged; financial aid decisions are made; taxpayers are chosen for audits; air travelers are selected for search; credit eligibility decisions are made. Citizens, and society as a whole, have an interest in making these processes more transparent. Yet the full basis for these decisions is rarely available to affected people: the algorithm or some inputs may be secret; or the implementation may be secret; or the process may not be precisely described. A person who suspects the process went wrong has little recourse.

To address this problem, we need to use **accountable algorithms**, which provide both an output and a proof that can convince a skeptical party that the algorithm was applied correctly to a given set of inputs to produce the announced output. Critically, the proof can convince an observer while maintaining the secrecy of the algorithm, the inputs, or both.

As an example, consider a government tax authority that is deciding which taxpayers to audit. Taxpayers are worried that audit decisions may be based on bias or political agenda rather than legitimate criteria; or they may be worried that the authority's code is buggy. The authority does not want to disclose the details of its decision algorithm, for fear that tax evaders will be able to avoid audits. The accountable algorithms framework will allow the tax authority to maintain the secrecy of its algorithm (in the sense that any observer learns nothing about the algorithm beyond what is conveyed by whatever input-output pairs that observer can see), while allowing each taxpayer to verify that:

- the authority committed to its secret algorithm in advance,
- the result asserted by the authority is the correct output of the authority's algorithm when applied to the individual taxpayer's data, and
- the authority can reveal its algorithm to an oversight body (such as a court or legislature) for examination later, and taxpayers can verify that the revealed algorithm is the same one used to make decisions about them.

Joshua A. Kroll is a PhD candidate in Computer Science at the Center for Information Technology Policy at Princeton University, where he is advised by Edward W. Felten and Andrew W. Appel. His research spans computer security, privacy, and the interplay between technology and public policy. He received the National Science Foundation Graduate Research Fellowship in 2011.

Date:	Thursday, April 17, 2014, 8:00 pm. (Refreshments and networking at 7:30 pm.)
Place:	Small Auditorium, Room CS 105 Computer Science Building, Princeton University
Information:	Dennis Mancl (908) 582-7086
On-line info:	**http://PrincetonACM.acm.org**

All Princeton ACM / IEEE-CS meetings are open to the public. Students and their parents are welcome. There is no admission charge, and refreshments are served.

A pre-meeting dinner is held at 6:00 p.m. at Ruby Tuesday's Restaurant on Route 1. Please send email to **princetonacm@acm.org** in advance if you plan to attend the dinner.

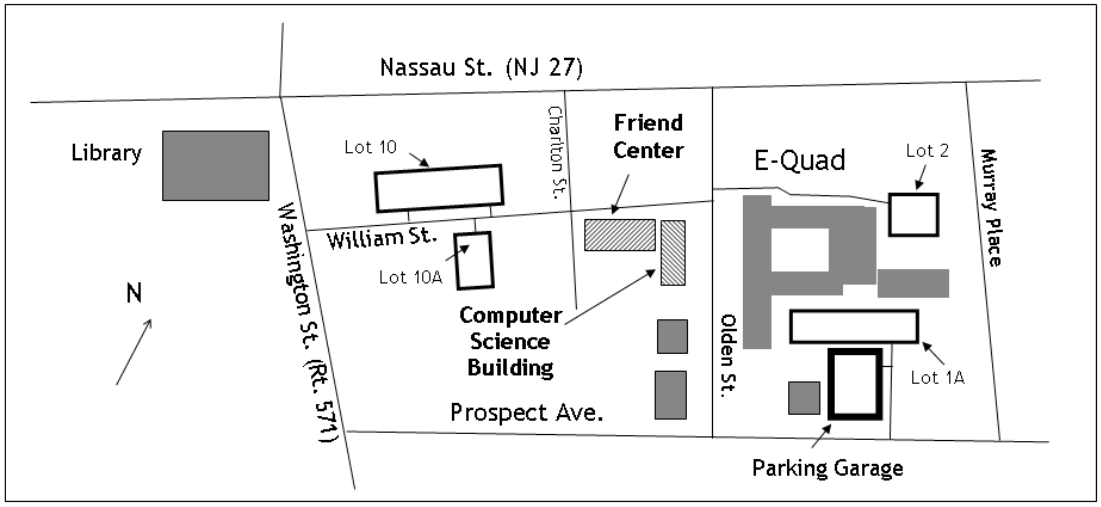
Nassau St. (NJ 27)
Library
Lot 10
Charlton St.
Friend Center
E-Quad
Lot 2
Murray Place
William St.
Lot 10A
N
Washington St. (Rt. 571)
Computer Science Building
Olden St.
Lot 1A
Prospect Ave.
Parking Garage