Nicholas J. Hopper

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 216 5023
hopper@cs.cmu.edu
http://www.cs.cmu.edu/~hopper/

Research Interests

Theoretical Computer Science, specializing in Theoretical Cryptography, with interests in Complexity Theory and Computational Learning Theory;
Computer and Network Security, focusing on security protocols, humans in computer security, and the application of cryptographic tools and analyses to privacy concerns.

Education

Carnegie Mellon University — Pittsburgh, PA

- Ph.D. Student, Computer Science
- Thesis title: *Toward A Computational Theory of Steganography*
- Thesis Advisor: Manuel Blum
- Expected Completion: July 2004

University of Minnesota — Morris, MN

- B.A., with high distinction, majors in Mathematics and Computer Science, 1999
- GPA: 4.0/4.0

Publications

[AH02] Luis von Ahn and Nicholas J. Hopper. "Public Key Steganography." To appear in: Proceedings of Eurocrypt 2004, May 2004.

[ABoH03] Luis von Ahn, Andrew Bortz, and Nicholas J. Hopper. "k-Anonymous Message Transmission." In: Proceedings of the 10th ACM Conference on Computer and Communications Security (CCS 2003), October 2003.

[ABHL03] Luis von Ahn, Manuel Blum, Nicholas J. Hopper, and John Langford. "CAPTCHA: Using Hard AI Problems for Security." In: Proceedings of Eurocrypt 2003, May 2003.

[HLV02] Nicholas J. Hopper, John Langford, and Luis von Ahn. "Provably Secure Steganography." In: Proceedings of CRYPTO 2002, August 2002.

[HB01] Nicholas J. Hopper and Manuel Blum, "Secure Human Identification Protocols." In: Proceedings of Asiacrypt 2001, December 2001.

[HSW00] Nicholas J. Hopper, Sanjit A. Seshia and Jeannette M. Wing, "Combining Theory Generation and Model Checking for Security Protocol Analysis." At: Workshop on Formal Methods in Computer Security, July 2000.

[MH99a] Nicholas Freitag McPhee and Nicholas J. Hopper, "Analysis of genetic diversity through population history." In: GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference, July 1999.

[MH99b] Nicholas Freitag McPhee and Nicholas J. Hopper, "AppGP: An alternative structural representation for GP." In: Proceedings of the 1999 Congress on Evolutionary Computation, June 1999.

[MHR98] Nicholas Freitag McPhee, Nicholas J. Hopper and Mitchell L. Reierson, "Impact of types on essentially typeless problems in GP." In: Genetic Programming 1998: Proceedings of the Third Annual Conference, July 1998.

In Progress

[H03] Nicholas J. Hopper "Covert Public-Key Infrastructures." Manuscript, 2003.

[AH03] Luis von Ahn and Nicholas J. Hopper. "An Empirical Study of WWW Password Security." Manuscript in preparation, 2003.

[ABe+03] Luis von Ahn, Alina Beygelzimer, Nicholas J. Hopper, and John Langford. "Covert Multiparty Computation." Manuscript in preparation, 2003.

[AHL04] Luis von Ahn, Nicholas J. Hopper, and John Langford. "Covert Two-Party Computation." *Submitted to CRYPTO 2004.*

Academic Honors and Awards

- National Merit Scholar (1995-1999)
- University of Minnesota Presidential Scholarship (1995-1999)
- ACM World Programming Contest Finalist (1997)
- McCree Award for Achievement and Potential in the Mathematical Sciences (1998-99)
- Scholar of the College, University of Minnesota, Morris (1999)

- NSF Graduate Research Fellowship Program, Honorable Mention (1999)
- NSF Graduate Research Fellowship (2000-2003)
- Siebel Scholar (2004)

Teaching Experience

- *Algorithms*, Fall 2000: Teaching Assistant, taught a weekly recitation and designed and graded homework and exam problems.
- *Security and Cryptography*, Fall 2001: Teaching Assistant, gave several lectures, designed and graded all homeworks and exams.
- *Foundations of Theoretical Cryptography*, Spring 2003: Instructor, developed and taught graduate course on theoretical cryptography.
- *Supervised undergraduate research*:

  Preston Tollinger: *A Secure, Device-Free Challenge-Response Protocol*, CMU Senior Thesis, 2000.

  Ann Lewis: *A Steganographic Text Editor*, NSF Aladdin REU Project, 2002

  Andrew Bortz: *Anonymous Communications*, NSF Aladdin REU Project, 2003

  Adam Bender: *HumanAut*, NSF Aladdin REU Project, 2003

Academic & Professional Service

- CMU Computer Science Department Doctoral Review Committee, 1999-2002.
- CMU Computer Science Department Ph.D. Admissions Committee, 2002, 2003.
- Program Committee, CEC Special Session on Security (2003).
- Paper referee for: *Science* (2000), STOC 2004, Eurocrypt 2004, IH 2004, *M&SOM* (2004).

Conference and Workshop Presentations

- "Secure Human Identification Protocols," *Asiacrypt 2001*, Gold Coast, Australia, December 2001.
- "Human Interactive Proofs." *NSF ALADDIN Workshop on Human Interactive Proofs*, Palo Alto, January 2002.
- "Provably Secure Steganography," *Crypto 2002*, Santa Barbara, August 2002.

- “k-Anonymous Message Transmission.” *NSF ALADDIN Workshop on Privacy in DATA*, Pittsburgh, April 2003.
- “Machine Learning and Reductions in Cryptography.” (Invited) *Workshop on Machine Learning Reductions*, Chicago, September 2003.

References

Professor Manuel Blum
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268 3742
`mblum@cs.cmu.edu`

Professor Michael Reiter
Carnegie Mellon University
Hammerschlag Hall, Room
D208
Pittsburgh, PA 15213
(412) 268 1318
`reiter@cmu.edu`

Professor Steven Rudich
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
(412) 268 7885
`rudich@cs.cmu.edu`

Asst. Professor David Wagner
EECS Computer Science Division
University of California, Berkeley
Berkeley, CA 94720-1776
(510) 642 2758
`daw@cs.berkeley.edu`

Other Professional Experience

University of Minnesota Morris, MN
1997-99 Research Assistant, Computer Science Department. Investigated impact of program representation on performance of Genetic Programming for learning simple concept classes.

University of Minnesota Morris, MN
1997-98 Grader, Computer Science Department. Graded papers and assisted with labs for upper-division courses in Networks, Databases, and Software Design.

Andersen Consulting Enterprises Minneapolis, MN
1998-1999 Summer Intern. Performed and automated various systems and database administration tasks related to new and upgraded installations of a production airline accounting system, resulting in annual savings of over $100K.

Abstracts of accepted and submitted papers

1. Luis von Ahn, Nicholas J. Hopper, and John Langford.

   "Covert Two-Party Computation." Submitted to *CRYPTO 2004*.

   Abstract. We introduce the novel concept of *covert two-party computation.* Whereas ordinary secure two-party computation only guarantees that no more knowledge is leaked about the inputs of the individual parties than the result of the computation, covert two-party computation employs steganography to yield the following additional guarantees: (A) no outside eavesdropper can determine whether the two parties are performing the computation or simply communicating as they normally do; (B) before learning $f(x_A, x_B)$, neither party can tell whether the other is running the protocol; (C) after the protocol concludes, each party can only determine if the other ran the protocol insofar as they can distinguish $f(x_A, x_B)$ from uniformly chosen random bits. Covert two-party computation thus allows the construction of protocols that return $f(x_A, x_B)$ only when it equals a certain value of interest (such as "Yes, we are romantically interested in each other") but for which *neither party can determine whether the other even ran the protocol whenever $f(x_A, x_B)$ does not equal the value of interest.* We introduce security definitions for covert two-party computation and we construct protocols with provable security based on the Decisional Diffie-Hellman assumption.

2. Luis von Ahn and Nicholas J. Hopper.

   "Public-Key Steganography." To appear in: *Advances in Cryptology : Proceedings of Eurocrypt 2004*, May 2004.

   Abstract. Informally, a public-key steganography protocol allows two parties, who have never met or exchanged a secret, to send hidden messages over a public channel so that an adversary cannot even detect that these hidden messages are being sent. Unlike previous settings in which provable security has been applied to steganography, public-key steganography is information-theoretically *impossible.* In this work we introduce computational security conditions for public-key steganography similar to those introduced by Hopper, Langford and von Ahn for the private-key setting. We also give the first protocols for public-key steganography and steganographic key exchange that are provably secure under standard cryptographic assumptions. Additionally, in the random oracle model, we present a protocol that is secure against adversaries that have access to a decoding oracle (a steganographic analogue of Rackoff and Simon's attacker-specific adaptive chosen-ciphertext adversaries fr om CRYPTO 91).

3. Luis von Ahn, Andrew Bortz, and Nicholas J. Hopper.

   "k-Anonymous Message Transmission." In: *CCS 2003: Proceedings of the 10th annual ACM Conference on Computer and Communications Security*, October 2003.

   Abstract. Informally, a communication protocol is *sender k - anonymous* if it can guarantee that an adversary, trying to determine the sender of a particular message, can only narrow down its search to a set of k suspects. *Receiver k-anonymity* places a similar guarantee on the receiver: an adversary, at best, can only narrow down the possible receivers to a set of size k. In this paper we introduce the notions of sender and receiver k-anonymity and consider their applications. We show that there exist *simple* and *efficient* protocols which are k-anonymous for both the sender and the receiver in a model where a polynomial time adversary can see all traffic in the network and can control up to a constant fraction of the participants. Our protocol is provably secure, practical, and does not require the existence of trusted third parties. This paper also provides a *conceptually simple* augmentation to

Chaum's DC-Nets that adds robustness against adversaries who attempt to disrupt the protocol through perpetual transmission or selective non-participation.

4. Luis von Ahn, Manuel Blum, Nicholas J. Hopper and John Langford.

   "CAPTCHA: using Hard AI problems for security." In: Advances in Cryptology : Proceedings of Eurocrypt 2003, May 2003.

   Abstract. We introduce CAPTCHA, an automated test that humans can pass, but current computer programs can't pass: any program that has high success over a CAPTCHA can be used to solve an unsolved Artificial Intelligence (AI) problem. We provide several novel constructions of CAPTCHAs. Since CAPTCHAs have many applications in practical security, our approach introduces a new class of hard problems that can be exploited for security purposes. Much like research in cryptography has had a positive impact on algorithms for factoring and discrete log, we hope that the use of hard AI problems for security purposes allows us to advance the field of Artificial Intelligence. We introduce two families of AI problems that can be used to construct CAPTCHAs and we show that solutions to such problems can be used for steganographic communication. CAPTCHAs based on these AI problem families, then, imply a win-win situation: either the problems remain unsolved and there is a way to differentiate humans from computers, or the problems are solved and there is a way to communicate covertly on some channels.

5. Nicholas J. Hopper, John Langford, and Luis von Ahn.

   "Provably Secure Steganography." In: Advances in Cryptology: Proceedings of CRYPTO 2002, August 2002.

   Abstract. Informally, *steganography* is the process of sending a secret message from Alice to Bob in such a way that an eavesdropper (who listens to all communications) cannot even tell that a secret message is being sent. In this work, we initiate the study of steganography from a complexity-theoretic point of view. We introduce definitions based on computational indistinguishability and we prove that the existence of one-way functions implies the existence of secure steganographic protocols.

6. Nicholas J. Hopper and Manuel Blum.

   "Secure Human Identification Protocols." In: Advances in Cryptology: Proceedings of Asiacrypt 2001, December 2001.

   Abstract. One interesting and important challenge for the cryptologic community is that of providing secure authentication and identification for unassisted humans. There are a range of protocols for secure identification which require various forms of trusted hardware or software, aimed at protecting privacy and financial assets. But how do we verify our identity, securely, when we don't have or don't trust our smart card, palmtop, or laptop?

   In this paper, we provide definitions of what we believe to be reasonable goals for secure human identification. We demonstrate that existing solutions do not meet these reasonable definitions. Finally, we provide solutions which demonstrate the feasibility of the security conditions attached to our definitions, but which are impractical for use by humans.

7. Nicholas J. Hopper, Sanjit Seshia, and Jeannette M. Wing.

   "Combining Theory Generation and Model Checking for Security Protocol Analysis." At: *Workshop on Formal Methods in Computer Security*, July 2000.

Abstract. This paper reviews two relatively new tools for automated formal analysis of security protocols. One applies the formal methods technique of model checking to the task of protocol analysis, while the other utilizes the method of theory generation, which borrows from both model checking and automated theorem proving. For purposes of comparison, the tools are both applied to a suite of sample protocols with known flaws, including the protocol used in an earlier study to provide a baseline. We then suggest a heuristic for combining the two approaches to provide a more complete analysis than either approach can provide alone.

8. Nicholas Freitag McPhee and Nicholas J. Hopper.

   "Analysis of genetic diversity through population history." In: GECCO-99: Proceedings of the Genetic and Evolutionary Computation Conference, July 1999.

   Abstract. The idea that diversity in the population of a genetic algorithm affects the algorithm's search efficiency is widely accepted. However, little is known about the amount of node level diversity present in Genetic Programming (GP) runs. In this paper, we introduce several techniques for measuring the diversity of a population based on the genetic history of the individuals. We then apply these measures to the genetic histories of several runs of four different problems. The results of this analysis show that a surprisingly small amount of diversity is present in the final population of a GP run. We conclude by suggesting a variety of other potential applications of these measures.

9. Nicholas Freitag McPhee and Nicholas J. Hopper.

   "AppGP: An alternative structural representation for GP." In: Proceedings of the 1999 Congress on Evolutionary Computation, June 1999.

   Abstract. It has been shown that Standard Genetic Programming using standard subtree crossover is prone to a form of structural convergence which makes it extremely difficult to make changes near the root, occasionally causing runs to become trapped in local maxima. Based on these structural limitations we propose a different tree representation, AppGP, which we hope will avoid this problem in some cases. In this paper, we describe this representation, and compare its performance to the performance of Standard GP on a suite of test problems. We find that on all of the test problems, AppGP does no worse than Standard GP, and in several it does considerably better, suggesting that the representation warrants further study.

10. Nicholas Freitag McPhee, Nicholas J. Hopper and Mitchell L. Reierson, "Impact of types on essentially typeless problems in GP." In: Genetic Programming 1998: Proceedings of the Third Annual Conference, July 1998.

    Abstract. Several researchers have shown type systems to be valuable in extending the range of problems (conveniently) addressed by Genetic Programming. There are other possible benefits of type systems, however, that derive from the new kinds of structural representation they make possible, and the effects that this has on the performance of recombination operators like the crossover operator. In this paper we compare the performance of Standard (untyped) Genetic Programming (SGP) and Hindley-Milner (typed) Genetic Programming (HMGP) on a suite of problems where an untyped representation (satisfying the closure property) is quite natural. We find that on several problems HMGP significantly outperforms SGP, while on other problems the performance of SGP and HMGP are essentially the same. We also suggest an intermediate representation that should provide many of the benefits of HMGP on these problems without requiring the complexity of a powerful type system.