Candidate's statement: **Phillip Rogaway**

Homepage http://www.cs.ucdavis.edu/~rogaway/ (←purple text for clickable links)

The requested merit is an advancement from Professor Step VII. The following enumerates selected activities and recognition during the review period of 7/1/2013 to 6/30/2016. Let me start with some context.

During this review period my work has undergone a shift occasioned by the Snowden revelations. The disclosures, which began with a *Guardian* article of June 5, 2013, revealed classified government programs to comprehensively monitor the world's population. Exabytes of data, almost all of which is collected outside of Fourth Amendment requirements of warrant or particularity, are used for purposes like law enforcement, counter-terrorism, assassinations, trade negotiations, and to keep tabs on domestic change movements.

Most cryptographers, and most computer scientists in general, reacted to the Snowden disclosures with a yawn. "We know all this," cryptographers said. "This is a political concern." "This is not our job."

But even an outsider might view such claims with skepticism. If cryptography's most *basic* aim is to enable secure communication—and it is—then how could it *not* be a colossal failure of the field when ordinary people lack even a modicum of communication privacy when interacting electronically? And if the Internet is this wonderful gift that computing technology has brought to mankind, then how is it not a failure of computer scientists if it has been instrumented to Collect it All, Process it All, Exploit it All, and Know it All (in the words of one presentation from the Snowden trove)?

Cryptography has many possible aims. *Crypto-for-privacy* aims to protect the right to free speech, movement, and economic engagement. *Crypto-for-security* aims to protect the stream of electronic commerce and ensure that computing systems "work" despite adversarial attack. *Crypto-for-crypto* aims to solve intriguing problems that cryptographers themselves dream up. *Crypto-for-power* seeks to advance the interests of the state or corporations by viewing people as sources of useful information. While crypto-for-security, crypto-for-crypto, and crypto-for-power have flourished, crypto-for-privacy has completely failed.

Most of my past work falls under the mantle of crypto-for-security. But by the end of the "summer of Snowden" I had decided to focus on crypto-for-privacy. I would work to create and popularize an area I call *anti-surveillance research*. I would pursue technical contributions, of course, but also non-technical contributions that directly speak to the social, political, and moral character of my field.

Research I would rather write a modest number of highly influential papers than a large number of papers that won't mean a great deal. I obsess on writing to a degree some might find absurd. I run a small shop, like two Ph.D. students and a postdoc.

I published nine papers during the review period, of which six appear in my field's traditional tier-1 venues: CRYPTO (3×) and EUROCRYPT (3×). The five best are described below. As is my policy, each has a full version (∼30–40 pages, on ePrint) that I release alongside the proceedings version (∼20 pages).

▶ **Subversion-resistant encryption [76].** My first piece of anti-surveillance research explored one way that "big brother" (think NSA/GCHQ) could try to subvert symmetric (shared-key) encryption. We call it an *algorithm substitution attack* (ASA). Within some proprietary software (like a web browser), instead of some standardized encryption algorithm $\mathcal{E}$, big brother manages to get implemented a secret, subverted algorithm $\widetilde{\mathcal{E}}$. It must play well with unsubverted code: $\mathcal{D}(K, C) = M$ when $C \leftarrow \widetilde{\mathcal{E}}(K, M)$. Yet using a secret algorithm $\widetilde{\mathcal{D}}$, big brother can decrypt any ciphertext C it sees (no user key needed): $\widetilde{\mathcal{D}}(C) = M$ if $C \leftarrow \widetilde{\mathcal{E}}(K, M)$. What is more, big brother manages to arrange that ciphertexts produced by $\mathcal{E}$ and $\widetilde{\mathcal{E}}$ are indistinguishable, even in the presence of user keys, so that users can't tell that anything's awry.

Our paper formalizes the problem and shows that all conventional ways to encrypt *can* be subverted by ASAs. Still, we to show a simple way to encrypt that is immune to ASAs.

The paper won the 2015 PETS award for the best privacy paper. (I'm told it also won the PC's vote for best-paper at CRYPTO 2014, but that the Chairs, concerned with the politics, preferred to give no award.)

► **Bigkey encryption [82].** My second piece of anti-surveillance research responds to a (largely valid) criticism of contemporary cryptography: that crypto work is pointless because adversaries prefer to go around it. For example, rather than bothering with some fancy cryptanalysis, a state-level adversary might exploit known vulnerabilities to install an APT (advanced persistent threat) that quietly exfiltrates your keys. We ask if anything practical can be done to stop this. The answer we give is *yes*—if you make your cryptographic keys huge and use them well. If a key is a terabyte, say, it will be hard to exfiltrate unnoticed (exfiltration will take a long time, and might be detected). At the same time, we have to ensure that using long keys is efficient (a practical algorithm can't even read the entire key; it will have to sample from it).

Ideas related to *bigkey cryptography* have been put forward before, as the *bounded-retrieval model.* But the work wasn't practical or concretely analyzed. We make it so by studying and exploiting a cute little game. Give an adversary a long random key of k bits (eg, $k = 10^{13}$). Let it exfiltrate any ℓ bits of information about that key (eg, $\ell = k/2$). Then identify p random positions into the long key and ask the adversary to predict all those bits. What is the maximum probably $f(k, \ell, p)$ that it can? We manage to compute this, using some pretty heavy mathematics. We show how this leads to practical, exfiltration-resistant crypto.

► **Robust-AE [79] and online-AE [80] and generic composition [75].** Authenticated encryption (AE) is symmetric encryption designed to provide both privacy *and* authenticity. It is a goal that I have done more to advance than any other. AE is a notion that has definitely caught on; for example, there is now an annual workshop devoted to it, and a cryptographic competition, CAESAR, drew an unprecedented 57 submissions from teams around the world. AE is now accepted as the "right" way to encrypt.

What makes AE "work" is definitional strengthening: make a notion strong enough, then provably achieve it, and the protection is less likely to end up inadequate. **Robust-AE [79]** strengthens conventional AE in useful ways. For example, if a plaintext already has some redundancy within, checking for it on decryption will enhances authenticity to the extent one would hope. Our paper not only motivates and defines robust-AE but shows how to achieve it with startling efficiency: about the cost of AES-CTR (counter mode). Previous methods used more work to achieve less. Our scheme, AEZ, is already in use by a leading security company, Silent Circle, and is a third-round finalist in the CAESAR competition.

That same competition attracted about 20 proposals for schemes that targeted a security notion, OAE1, suggested by colleagues. Authors were claiming that OAE1 was almost as strong as MRAE (misuse-resistant AE) [68] despite allowing encryption to be online. The claim is untrue, and the definition of OAE1 doesn't capture what it should. Our work on **online-AE [80]** exposed this, and explained how online-AE *should* be defined. (Work like this doesn't make you a lot of friends: we're effectively saying that an awful lot of people do pointless stuff and don't understanding the meaning of the definitions they use.)

Our paper on **generic composition [75]** upends a different aspect of AE research: the belief, going back to [Bellare-Namprempre 2000], that there are three generic ways to construct an AE scheme—encrypt-then-MAC, MAC-then-encrypt, and encrypt-and-MAC—with only the first being right. We show that the truth of this claim depends on inessential and outmoded definitional choices. When one switches to more modern definitions, ones better aligned with cryptographic practice, the claim is false. We explain how a misunderstanding of this issue resulted in a completely broken algorithm in international standard ISO 19772.

► **Recognition** ▷ I won the first Levchin Prize for Real-World Cryptography (Jan. 2016). The citation says: "Phil is a giant in the area of symmetric encryption. The award is given for his groundbreaking practice-oriented research, authenticated encryption and his work on format preserving encryption which has had an exceptional impact on real-world cryptography." ▷ Paper [76] won the 2015 PETS prize, as mentioned above. ▷ I gave the IACR Distinguished Lecture of 2015. (The IACR is the professional organization for cryptographers, analogous to the ACM or IEEE, organizations that, for historical and political reasons, play only a small role in publishing cryptographic work.) There is one such lecture per year, the speaker selected by the IACR Board of Directors. ▷ I spent the summer and fall of 2014 as a visiting professor at ETH Zürich (invited position, voted on by the department). ETH is one of the best universities in the world. ▷ I spent a month as a visiting professor at ENS Paris (invited position, 2015). ▷ I was invited to give an NSF WATCH talk (March 2016). These are watched by about 200 people in government. ▷ Distinguished Lecture at UC San Diego. ▷ Distinguished Lecture at MIT. ▷ My current citation count is 25,701 and my h-index is 66 (as per Google scholar; these up 33% and 20% since my last merit).

Service

▶ **Public/community service** I spent much of 2015 working on an essay that I wrote to accompany an invited talk (the IACR Distinguished Lecture), "The Moral Character of Cryptographic Work" [I33]. I hope to be remembered for this as much as for any of my technical papers. The essay is extensively researched, but it is not technical. It is explicitly normative, and unabashedly political. The essay ended up rather long (48 pages, 212 footnotes), but should be understandable by anyone. Its abstract is as follows:

> Cryptography rearranges power: it configures who can do what, from what. This makes cryptography an inherently *political* tool, and it confers on the field an intrinsically *moral* dimension. The Snowden revelations motivate a reassessment of the political and moral positioning of cryptography. They lead one to ask if our inability to effectively address mass surveillance constitutes a failure of our field. I believe that it does. I call for a community-wide effort to develop more effective means to resist mass surveillance. I plead for a reinvention of our disciplinary culture to attend not only to puzzles and math, but, also, to the societal implications of our work.

Reaction to the essay was swift, and highly positive. Even people who disagreed with much of what seemed to like the fact that I was saying it. I was deluged with emails and calls. Some were interview requests, and I did a couple before I realized how uncomfortable they made me. There were lots blog postings and tweets (eg, Cory Doctorow and Bruce Schneier). I got many talk invitations, which continue to this day. I found that the essay was getting assigned to students in all sorts of classes (eg, MIT's 6.033, Computer Systems Engineering; U. of Michigan's EECS 588, Computer and Network Security; UT Austin's Math 343L, Applied Number Theory). While I never anticipated publishing the essay (nowadays, just posting to ePrint gets plenty of notice), I received many invitations from editors to publish the work, or something derivative. Academic journals like *Philosophy and Technology*; popular magazines like *Scientific American*; and things in between like *CACM*. But the offer that sounded most appealing was an invitation from Princeton University Press to publish the work as a short book. I agreed to that, and am working on the revisions now.

The essay and talk were not my first attempts to reach beyond the technical in reacting to the Snowden disclosures. ▷ I started in Aug. 2013 with a personal statement on the web condemning mass surveillance. ▷ In Dec. 2013 I gave my first invited talk on the subject. I have given many since. ▷ In Jan. 2014 Joan Feigenbaum and I put out an open letter signed by 53 top researchers in cryptography and security. This letter was very hard to make happen (hundreds of hours, 900+ emails); computer scientists are extremely reluctant to take a public position on such a matter. ▷ In May 2014 I wrote and introduced what is now known as the Copenhagen Resolution. It was unanimously adopted by the IACR membership at Eurocrypt 2014. ▷ I co-organized a Dagstuhl meeting on mass surveillance that took place in Oct. 2014.

▶ **Further pubic/community service** ▷ NIST standardized the format-preserving encryption scheme I proposed, FF1, and the overarching framework, FFX. This is NIST Special Publication 800-38G (Mar. 2016). ▷ My attack on ISO 19772 resulted in a correction to the standard, ISO 19772:2009/Cor 1:2014. I coordinated with ISO to make that happen. ▷ A particularly consequential technological change we are seeing is that TLS 1.3 will permit *only* AE. (I have been advocating AE-only encryption for years.) Attacks on the SSL/TLS handshake should finally subside. ▷ In 2016 I began a three-year term on the IACR Board of Directors. (I got the most votes of any candidate. The top three win.) ▷ I continued to serve as an associate editor for the *J. of Cryptology*. But I am a terrible editor, and finally stepped down in May 2016. ▷ I chaired the IACR Fellows Selection Committee (2015), and was an ordinary member of this committee before that (2013, 2014). ▷ I served on the PETS selection committee (2016). ▷ Like all of us, I served on various program committees (but somewhat fewer than before; nowadays I most often decline) and NSF panels.

▶ **Departmental service** ▷ Throughout the review period I chaired our department's undergraduate committee, CSUGA, and represented us on undergraduate matters on the college committee, UGEP. ▷ I also chaired the CS Graduate Group's advisor's committee. This is one of those committees where the Chair does almost all the work, reviewing endless petitions. I did this conscientiously. ▷ Because of these two committees, I spent lots of time advising students one-on-one, both undergraduates and graduate students. That is perhaps the most pleasant sort of "work" there is to do—I like talking to our students.

Teaching I taught eight classes during the three-year review period: ECS 20 (F13), ECS 120 (S14, S15), ECS 127 (S16), ECS 188 (W15, W15, W16), and ECS 227 (W14). The mean of my mean teaching ratings was 4.48/5. Four of these eight classes had 126–151 students and four had 13–27 students. One quarter (F14) I was on sabbatical. Given how our department's counts things, this adds up to a teaching load just a bit beyond normal.

Despite its obvious problems, I prefer ratemyprofessor evaluations to our own student evaluations, as I find it unconscionable that students have no access to the scores that we base on *their* data. I'm 4.2/5 (41 ratings) on that website. Tags of high multiplicity being "tough grader," "skip class, you won't pass," "hilarious," and "inspirational."

I teach difficult classes, targeting the top students and bringing along the rest of the class along to the extent that I can. Student evaluations reflect this: the statement "I have had to work hard in this class" got a mean score of 9.4/10 for all online evaluations (F13–W15) during the period.

ECS 127 ("Modern Cryptography") is a new class of my own design, first taught in S16. (I tried things out as a 189A in Sprint 2011, but changed much since). I was surprised by the number of students who stuck with this class (about 130), an elective with plenty of theory. I believe the course to be pedagogically unique, embodying a modern viewpoint on cryptography unavailable from any book, but still accessible to an undergraduate. See http://tinyurl.com/127in2016.

In terms of graduate advising, I advised 1–2 Ph.D. students throughout during this review period, plus 0–1 postdocs. I wasn't satisfied with the performance of one of my two grad students and gently nudged him to move on to a more security-oriented colleague at another university. I think that will work better for him. The other grad student is sharper, but he's been slow to produce good work, too. Overall, my graduate students during this three-year period have not been as good as my prior students.

At present, I have two Ph.D. students, plus a postdoc from Belgium.

I have continued to informally mentor former grad student Viet Tung Hoang (graduated June 2013). Tung did a succession of three one-year postdocs, but finally landed a faculty position this Fall at Florida State University. My other former students, apart from one at Google, are all tenured faculty: Cornell Tech, Sacramento State University, Tsinghua University, University of Colorado, and the University of Florida.

Sincerely,

Phillip Rogaway
Davis, California, USA
October 26, 2016