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 8.5 to 9.5. The review period is 7/1/2016 to 6/30/2019. No advance beyond a single step is possible based on the current packet, as a 1.5 or 2 step advance would go to Above Scale, enlarging the review period and requiring the acquisition of letters.

Preamble. I have always regarded the preparation of these statements as a time to reflect on where I am in the arc of my career. This time around I find such reflection sad. While my work remains strong, my enthusiasm for research has waned, and the teaching often feels a little hopeless—a too-little-too-late effort to help too-few students face the reality of a world on the brink of collapse.

Computer science was always an odd place for me to land. I liked theory, and CS-stuff came easily, so I did it. Yet I always viewed computers with mild disdain. Now, each day, I watch the human-phone hybrids that have replaced our students and think: my discipline has brought the world *this*? I feel ashamed.

Increasingly isolated, I have been shifting more of my energy to teaching, particularly to ECS 188 (ethics and technology). I shifted my service to work on more structural concerns about the cryptographic community. I redirected research to "meta-theory" and, since 2013, to matters connected to privacy and democracy.

Research

I published six papers during the review period (counting one in-press). These are in CRYPTO (the traditional tier-1 conference), ASIACRYPT (nowadays a tier-1 conference), PoPETS (tier-1 conference-then-journal for privacy-enhancing work), *Journal of Cryptography* (the main journal in the field), an invited (but refereed) paper for *IEEE Security and Privacy* (a magazine-like journal), and an invited paper for LATINCRYPT (a minor but pleasant conference). This is a fairly typical rate of output for me. I usually work on papers for 1–3 years and obsess on writing to an extent that others would find absurd. My papers often treat some hitherto unseen problem, often born from a recognition that some aspect of our disciplinary culture has left some blind spot. Let me describe my two most significant efforts during the period.

▶ **Simplifying Game-Based Definitions [85], [86].** The most interesting technical idea I had during the last few years is a concept I call *indistinguishability up to correctness*, or IND|C. You need first to understand that a significant part of cryptography is about crafting *definitions*. Often given as pseudocode-described *games*, they associate a real number to an *adversary* interacting with a *protocol*. This number, the adversary's *advantage*, measures how well it disrupts the protocol's aim. But the games get so subtle or embed so many implicit choices and assumptions that authors make errors. Readers often can't figure out if a definition is right, wrong, or too ambiguous for terms like that to even make sense. In a world like this, cryptographic definitions should be deeply contested territory. They are not. A culture has emerged that mostly ignores definitional complexity, definitional imprecision, and the extent to which definitions align with real-world motivations. Slow, hard, philosophical, and uncertain, worrying about these things is incompatible with the popular ways of working and publishing.

I have long wanted to rationalize and elevate the definition-making enterprise. To get there, I have developed a framework, IND|C, for giving definitions. One begins by writing pseudocode for a pair of *utopian* games, one specifying the behavior of a (protocol-dependent) *real* oracle, the other, an *ideal* oracle. One might "like" to define the adversary's advantage as the difference in probabilities between the adversary outputting 1 when interacting with these two oracles, but one can't do that because the oracles are, well, utopian: there are *trivial* ways available to the adversary to distinguish the oracles. Conventionally, one would identify these ways and tweak the games so as to exclude them. The starting point for IND|C is the realization that one needn't do this by hand: one can automatically and generically outlaw trivial wins by attending to protocol correctness. A correct protocol must accomplish some aim, even in the absence of an adversary. This will make unavoidable certain ways to distinguish the real and ideal utopian games.

One begins with a pair of (distinguishable) utopian games, a real and an ideal one, and a correctness condition, which is just a set of protocols (the ones deemed correct). The real and ideal utopian games are then *silenced*—each oracle is forced to shut up—if the oracle's response, in the real setting, is dictated by the correctness condition alone. That is, you silence an oracle if the answer to a query is determined just by knowing the conversation so far and the fact that the underlying protocol *is* correct. This becomes the operative notion of triviality. In a two-paper series (with one more to come), both with grad student James Zhang, we develop this idea [85] and apply it to increasingly complex examples [85], [86].

This is an abstract and speculative line of work—definitions for writing definitions. But I believe it can ultimately serve as a way to produce more rigorous and comprehensible game-based cryptographic definitions. It might also usurp some of what UC (universal composability) so poorly accomplishes.

▶ **Anonymous authenticated-encryption [88].** This paper points out that authenticated-encryption (AE) definitions—indeed all encryption definitions—implicitly violate privacy. (Privacy in the sense of anonymity: keeping the identity of the encrypting party secret.) The problem can be seen from a typical decryption API. What has become the standard interface (due to my own work) looks like $\mathrm{Decrypt}_K^{N,A}(C)$ where K is the shared key, N is the nonce (chosen by the encrypting party and used at most once per session), A is the associated data (a value authenticated but not made confidential), and C is the the ciphertext. Consider even the first of these arguments, K. How can the decrypting party know *which* key K to use? In practice, it will usually know this because there will be something like a *session-ID* (SID) that flows alongside the ciphertext C. But providing such an SID is privacy-violating. Arguments N and A have similar problems. In short, cryptographer's basic formulation of encryption—all the way down to the fundamental conception of what transformations are performed on what data—is privacy-violating in it most natural use.

My paper on *anonymous AE* (anAE), joint with grad student John Chan, reformulates AE in a privacy-friendly way. We demand that one be able to augment the customary Decrypt operation with a collection of alternative calls whose use will not destroy privacy. We execute this idea in a way that permits simple schemes built from any conventional AE scheme.

Other work. Papers [83] and [84] are retrospective, introspective works. The first focuses on the evolution of practice-oriented provable security; the second, on the role of definitions in modern cryptography. Both papers are explicitly philosophical and sociological.

Many people see expository, survey, or position papers as nearly valueless. They are wrong. Papers [83] and [84] offer perspectives on my field that, realistically, could only be provided by me or Mihir Bellare. Numerous graduate students have told me that my handful of expository pieces shaped their understanding of our field. My non-technical paper *The Moral Character of Cryptographic Work*, written during my last review period, is the work I am most proud of over my entire career.

Paper [87] describes a bizarre way to create a blockcipher, based on a particular card shuffle. There is something mildly shocking in the idea that there are ways to make blockciphers that look so different from conventional ways (Feistel networks, SP networks, ARX constructions); and even more shocking that alternatives can have far better concrete security bounds. I have followed this up with even stranger and quantitatively better ways to encipher—the "swap-or-not" shuffle and the "sometimes-recurse" shuffle.:w By now I have demonstrated a productive connection between blockcipher construction and card shuffling.

Standing. At the time of this writing, 32650 papers cite my work, and my h-index is 78. (All data in this paragraph from Google Scholar on 10/12/2019.) These are up 27% and 18% from my last action. I am currently the 1st, 6th, and 14th most cited researcher in ethics-and-technology, privacy, and cryptography, respectively—the three areas that Google Scholar (correctly) slots me into.

Teaching While I have always cared deeply about teaching, I once cared more about research. I saw myself as a researcher who also taught, not as a teacher who also did research. This has changed. I now see my teaching—particularly my teaching of ethics—as the most valuable thing that I am doing in life. ECS 188 often upends students' worldview and their life trajectories; how is that less important than my research? I have come to regard my previous valuation of research over teaching as implicitly arrogant ("anyone can teach—being a top researcher is more special") and conformist ("UCD values research over teaching—so I should, too").

I taught eight podium classes during the three-year review period: seven classes of ECS 188 (Ethics in an Age of Technology) and one class of ECS 127 (Modern Cryptography). The teaching load was relatively low because I had accumulated excess teaching credits that the department wanted to "pay off" (resulting in a one-year period with no teaching at all). Summary data is as follows.

Class (Term)	Mean	Median	Response Rate	1	2	3	4	5
ECS 188–1 (Winter 17)	4.8	5	22/22 = 100%				5	17
ECS 188–2 (Winter 17)	4.4	5	22/23 = 96%		1	2	7	12
ECS 188–1 (Fall 17)	4.6	5	25/25 = 100%			1	7	17
ECS 188–2 (Fall 17)	4.7	5	25/26 = 96%			1	5	19
ECS 188–3 (Fall 17)	4.6	5	22/24 = 92%			2	5	15
ECS 127 (Winter 19)	4.3	5	70/97 = 72%		5	6	21	37
ECS 188–1 (Spring 19)	4.7	5	23/27 = 85%			2	3	18
ECS 188–2 (Spring 19)	4.5	5	17/26 = 65%			3	2	12

In the table above, evaluations prior to Winter 2019 were done on paper—in class, after students' final exams were turned in. This is why my response rate are nearly 100%. Beginning in Winter 2019, the department stopped allowing me to use paper evaluations. A faculty member had complained that it was impractical for *everyone* to do paper evaluations and that my doing evaluations this might give me (and the one other holdout) an unfair "advantage." While I seriously doubt that doing student evaluations immediately after a difficult final exam *increases* scores, it definitely increased response rates. And given the no-tech-in-class rule I usually adopt for ECS 188, online evaluations continue to feel inappropriate.

Sample student comments on 188: "This is the class where I learned the most in my educational journey" [188-1.W17]. "This class changed the way I see many things, and I believe might change the way I live my life in a significant way." [188-2.W17]. "This class has been an eye-opener & is seriously value-changing. I will look back upon this class as among the most, if not the most, influential in my 5 year college career" [188-1.F17]. "made me care more about humanity" [188-1.F17]. "He was extremely inspirational and impactful" [188-1.F17]. "His approachability and lack of fear in pointing out why we are failing as a civilization" [188-1.F17]. "extremely profound and wise. A joy to listen to" [188-1.F17]. "I like best that the professor, for lack of a better phrase, gives a shit" [188-1.F17]. "He changed my life, thought outside the norm and made us do the same" [188-2.F17]. "Rogaway combines strong ethical beliefs with humbleness and humor to not just convince you that injustice exists in the world, but to want to do something about it." [188-2.F17]. "Amazingly relevant and well-curated source material" [188-2.F17]. "Professor Rogaway is the most inspirational professor I have ever had" [188-2.F17]. "No course has taught me more about the world" [188-2.F17]. "he show me it's possible to have conviction" [188-3.F17]. "Most educational course I've taken at Davis" [188-1.S19]. "My favorite class ever" [188-1.S19]. "This class really made me re-evaluate what I've been taught in school for years and taught me to be more critical" [188-1.S19]. "This class really changed my perspective on almost everything. The material can be quite depressing, but I see why its necessary to present it in that way" [188-2.S19]. "I honestly believe this course should be required to be taken by all students at UC Davis, but if anything at least also make it required to be taken by Computer Science Majors as well" [188-2.S19]. "Phil produced a breathtaking amount of reading materials that were insightful, eye opening, as well as intellectually and emotionally challenging. . . . [It] has changed my worldview on so many issues already: urgency of climate destruction, exploitative nature of technology companies, deep analysis of our social structure. For the first time in my college career, I felt intellectually stimulated via multiple disciplines, rather than just the engineering skills" [188-2.S19]. **And some negative comments:** "Very politically biased" [188-1.W17]. "The discussion was OK, but so many people didn't contribute" [188-2.W17]. "So much work, Too much!!!" [188-3.F17]. "I feel like this class was designed to make students dislike their career path/majors" [188-3.F17].

Sample student comments on 127: "Best class ever." "This was by far one of the most interesting and eye opening classes I have ever taken." "This course is by far one of the more useful, interesting, and academically challenging course I've taken as an undergraduate. Professor Rogaway truly communicated the depth of the subject while also making it accessible. It's very empowering to have access to a professor with such expertise in the field". "My favorite course here at UC Davis. The material was fun and challenging and the professor was by far the best professor I've ever had." "Awesome. Best course I've had. Don't change a thing." "I love how I honestly feel like I can read cryptography research papers now and understand them because you didn't dumb it down. Also you are very funny and made class extremely interesting and enjoyable." "very thoughtful about how he chooses his words, which makes him a great teacher. Hands down, he is the best instructor I have had at Davis." "His explanations during office hours were amazing." "The professor knows the material incredibly well and makes lectures interesting and useful." "To be honest you're one of the best professors I've taken in Davis and this is in a large part because it is very apparent that you have passion both for what you are teaching and the students you are teaching." **And some negative comments:** "I gave up on this class due to how much effort was required for how little reward you'd get." "I was not able to learn anything since I was not able to even understand the homeworks, let alone complete them". "This was a ridiculously stressful class." "not having a textbook was brutal."

Since 2004, when I reinvented and first taught ECS 188, the course has come to play a central role in my life. It has required much of me personally: becoming vegetarian, changing what and how much I read, shifting my research focus, increasing my political activism. As our species moves towards collapse, destroying the earth's climate and biome, we must grieve, resist, and try to understand what has been our story. As technologists and human beings, we need to see how our understanding and manipulation of nature became the instrument of our demise. This is not an easy thing to do—intellectually or emotionally. We have been spoon-fed a constant diet of propaganda and techno-optimism. I try to get students to push it aside.

In terms of graduate advising, I advised 1–3 Ph.D. students during this review period. One student, James Zhang, graduated last Spring and took a job at Google. This has made me agitated. All my students took academic positions except the two weakest, both of whom went to Google. I have long felt that if a student of mine went to the NSA then I should stop training graduate students. Should I be thinking of Google differently?

A second student, Zane Rubaii, did a year with me but decided that his sociopolitical orientation was inconsistent with a career in CS. I helped him to transfer to the Ph.D. program in economics.

A third student, John Chan, remains with me. He has progressed slowly.

I have no postdocs at the moment. My last post-doc, Atul Luykx, moved on near the beginning of this review period. He is now head of cryptography for Visa Research.

I received one NSF award, "Crypto-for-Privacy," for 500K. While technical, the proposal was also political, so I was pleased that the NSF program managers didn't kill it. I don't need a ton of money and usually get requested awards, so I must ask for NSF for funding about once ever three years.

Service

Departmental service. I have cut back some on departmental service. Two or three years ago I withdrew participation from faculty recruiting (apart from LSOE/LPSOE hires). The values that had come to shape our recruiting were just not ones I shared. I had tried for years to speak out and to reorient our efforts, but had consistently failed. Faculty recruiting repeatedly brought out the worst in people. It was depressing. I consider it my worst professional failing at UCD that I was unable to move our department's evolution in a direction responsive to the era in which we live: the era when computer science became central to mass surveillance, political manipulation, and social alienation.

I continue to serve on CSUGA, but no longer chair it. Probably the most significant thing I worked on for that was the default academic-misconduct policy for our department (trying to institutionalize what I have long done for my own teaching). I often cover events like visit-day. I often cover UGEP meetings, but am no longer our regular representative.

Community service. I have expanded my community service, and met with some success. From Jan 2016 to Dec 2018 I served as a member of the IACR Board of Directors. (For political and historical reasons, all major cryptographic venues fall under the International Association of Cryptologic Research (IACR), not, say, the ACM or IEEE). Being a Board Member is an elected position. I won election by a large margin.

As a Board Member I soon found the topic I felt most useful to work on: COIs. A failure to enforce COIs when reviewing conference-papers had become a major problem in the community, distorting programs and angering authors. While submissions to IACR conferences are anonymous (apart from the one theory conference), reviewers routinely know who are the authors anyway. PC members would push for the papers of friends or close colleagues. COIs lacked any agreed-upon definition, and it was the whim of each conference chair, and each individual, when to declare one. For more than a decade authors and PC members had been telling me confidential stories of extraordinary abuse. Some senior, highly-respected people were the worst offenders. For various sociological reasons, the crypto-theory subcommunity, the subcommunity in which I had grown up, seemed to be the most nepotistic of all.

Over a period of about 1.5 years, I managed to formalize a COI policy and get it adopted by the IACR. It applies to all IACR-sponsored publication venues (so all major conferences and venues). When you submit a paper you now fill out a form in which you must disclose all COIs. Automatic COIs are formally defined by the IACR. Reviewers may not review papers or witness discussions when a defined-COI exists. It is long-overdue change for how my field works. Pushing this change earned me some enemies. One theory-focused board member claimed that the entire crypto-theory community would split off from the IACR if we went through with it. But we did go ahead, and the theory people did not leave.

I also served on the IACR Ethics Committee. This became an important job. We dealt with cases plagiarism, sexual harassment, and more. I am proud of how we handled the complaints brought to our attention. One of our outputs is a formal Code of Conduct that gets adopted for each IACR event. There are now clear channels for reporting problems and getting help.

A final piece of IACR service I am happy with was my push to greatly increase funding for schools (mostly 1-week summer schools for grad students and young researchers). I helped double the IACR's default funding for these schools. We do several per year.

In a different direction, beginning in 2014, a competition named CAESAR was held for identifying next-generation authenticated-encryption (AE) schemes. A few months ago, my own submission, OCB, was selected as one of six winners (from the 53 submissions); it and AEGIS-128 are the chosen mechanisms for high-performance AE applications. Beyond this, I have been at the center of AE since its inception (the father of the area, one could fairly say). Nowadays, real-world cryptographic protocols invariably employ AE when they need symmetric encryption. The recently approved TLS 1.2 standard allows nothing else. The phrase "authenticated encryption" now appears on 187,000 webpages and 10,900 academic papers.

I have continued to give talks about the moral character of cryptography, and continue to get emails and inquiries about this. My book on this has stalled, however, with all but one chapter complete. I have been feeling too hopeless on the topic of this book to finish it. I do hope this will change. I am in desperate need of sabbatical.

Sincerely,

Phillip Rogaway
Davis, California, USA
October 14, 2019