Mathematicians Discuss the Snowden Revelations

In the first part of 2013, Edward Snowden, a former contractor for the National Security Agency (NSA), handed over to journalists a trove of secret NSA documents. First described in the media in June 2013, these documents revealed extensive spying programs of the NSA and other governmental organizations, such as the United Kingdom's GCHQ (Government Communications Headquarters). The disclosures reverberated around the world, influencing the bottom lines of big businesses, the upper echelons of international relations, and the everyday activities of ordinary people whose lives are increasingly mirrored in the Internet and on cell phone networks.

The revelations also hit home in the mathematical sciences community. The NSA is often said to be the world's largest employer of mathematicians; it's where many academic mathematicians in the US see their students get jobs. The same is true for GCHQ in the UK. Many academic mathematicians in the US and the UK have done work for these organizations, sometimes during summers or sabbaticals. Some US mathematicians decided to take on NSA work after the 9/11 attacks as a contribution to national defense.

Another tie to the mathematical sciences community comes through the Mathematical Sciences Program (MSP), which the NSA launched in the mid-1980s (see `http://www.nsa.gov/research/math_research/`). The MSP provides grants for unclassified research by individuals, conferences, research experiences for undergraduates, and a few other infrastructure projects, as well as for sabbaticals at the NSA. While the program is quite small—due to recent cuts, its budget is expected to be US$4 million in 2015—it is a significant source of support for some areas of mathematics. Since the early 1990s, the AMS has assisted with administration of the program by convening panels to review proposals for individual grants and for conferences, and to make recommendations to the NSA about which ones to fund.

On the suggestion of one of us (Harris), the *Notices* decided to host a discussion of the NSA. (The controversy over GCHQ is in many ways similar, but the *Notices*, being the journal of record of the AMS, is focusing on NSA.) Three unsolicited pieces arrived to open the discussion even before we had finalized plans for its format (all of the following articles are available at `http://www.ams.org/notices`):

> *Letter to the Editor:* "AMS Should Sever Ties with the NSA", by Alexander Beilinson (December 2013)
>
> *Opinion:* "Dear NSA: Long-Term Security Depends on Freedom", by Stefan Forcey (January 2014)
>
> *Communication:* "The NSA Back Door to NIST", by Thomas C. Hales (February 2014)

Other discussion of the issue includes an article by Edward Frenkel, "The perils of hacking math", which appeared in the online magazine *Slate* on September 30, 2013. Mathematicians in the US are not alone in feeling an urgent need for a public discussion of the implications of their institutional relations with surveillance agencies. The April 2014 issue of the *London Mathematical Society Newsletter* carried an opinion piece, "Should mathematicians cooperate with GCHQ?", by Tom Leinster of the University of Edinburgh (see `http://newsletter.lms.ac.uk`). Soon thereafter, Leinster wrote "Maths spying: The quandary of working for the spooks", which appeared in the April 23, 2014, issue of *New Scientist* magazine (see `http://www.newscientist.com`). That article was syndicated in *Slate* and sparked international media coverage, including articles on the French website Mediapart and in the German magazines *Der Spiegel* and *Die Zeit* online.

Over the past several months we have solicited articles from mathematicians whom we believed would have useful and informative views on this subject. Two of the resulting articles appear here. Both articles, as well as the other pieces mentioned above, are critical of the NSA. In aiming to present a balanced discussion representing a variety of views, we made many efforts to seek out authors

DOI: http://dx.doi.org/10.1090/noti1138

whom we thought might write in defense of the NSA. However, this proved difficult; some of those who turned us down might be under legal restrictions that greatly limit what they can say in public. We are continuing our efforts and intend in future issues to publish additional articles representing other viewpoints.

In his *New Scientist* piece, Tom Leinster writes, "Mathematicians must decide: do we cooperate with the intelligence services or not?… we mathematicians should talk about this." Frenkel, drawing a parallel with the ethical questions physicists faced with the invention of nuclear weapons, writes, "Members of my community must initiate a serious discussion about our role in this brave new world." What do you think? We look forward to hearing your opinions on these and similar questions. We also welcome all suggestions about how to make this discussion a thoughtful and informative one. Unsolicited submissions are welcome. Inquiries and submissions may be sent to `notices-snowden@ams.org`. Articles of 800 words or less are preferred. Those that are of 400 words or less can be considered as Letters to the Editor and should be sent to `notices-letters@ams.org`.

— *Michael Harris*
Institut de Mathématiques de Jussieu
`harris@math.jusseiu.fr`
and Columbia University
`harris@math.columbia.edu`

— *Allyn Jackson*
Notices *Deputy Editor*
`axj@ams.org`

The NSA: A Betrayal of Trust

Keith Devlin

Both as an American citizen and as a citizen in what is a highly integrated global society, I have opinions on many aspects of NSA surveillance. As someone who became a US citizen by choice, I have spent much time reflecting on what it means (or should mean) to be a citizen in a nation having so much power.

Since the Snowden revelations first broke, I have expressed those opinions publicly on social media and in a small number of published interviews that, like a great deal of information these days, can easily be accessed with a few keystrokes. In this article, written in response to an invitation from the *Notices of the American Mathematical Society*, I will focus on the one area where my opinion is informed by my mathematical expertise and five years of in-depth, Department-of-Defense-funded research in the area of extracting actionable information from vast amounts of data.

I concentrate on whether indiscriminate "vacuuming up" of personal information that, according to the documents Edward Snowden has released, the NSA has routinely engaged in for several years can effectively predict terrorist attacks. I'll say up front that, based on everything I learned in those five years, blanket surveillance is *highly* unlikely to prevent a terrorist attack and is a dangerous misuse of resources that, *if used in other ways, possibly could prevent attacks* (such as the 2013 Boston Marathon bombing). Anyone with a reasonable sense of large numbers could surmise a similar conclusion. When the goal is to identify a very small number of key signals in a large ocean of noise, indiscriminately increasing the size of the ocean is self-evidently not the way to go.

I reach my conclusion having spent five years looking at this problem in depth. From early 2002 until the middle of 2006, I worked on a Defense Department research project called NIMD (Novel Intelligence from Massive Data, `http://www.sourcewatch.org/index.php?title=Novel_Intelligence_from_Massive_Data`), funded by ARDA, the Advanced Research and Development Agency (`http://www.sourcewatch.org/index.php/Advanced_Research_and_Development Activity`). I did so under contract to Veridian Inc. It was a nonclassified project. I never sought nor had security clearance. Those of us involved were free to publish our results, but we were asked not to make public statements about the project or our involvement. I was happy to go along with that request. In particular, I never mentioned this work in any of my "Math Guy" appearances on National Public Radio nor in any of my regular columns for the Mathematical Association of America.

The only reason I am putting these words down now is the feeling of intense betrayal I suffered when I learned how my government and the leadership of my intelligence community took the work I and many others did over many years, with a genuine desire to prevent another 9/11 attack, and subverted it in ways that run totally counter to the founding principles of the United States, that cause huge harm to the US economy, and that moreover almost certainly *weaken* our ability to defend ourselves. During the project, I interacted with many other individuals, including other academic researchers, intelligence workers, and a few government and military personnel. Nevertheless, what my words express below is my considered and informed *opinion*. I never had, nor do I now have, access to any information beyond what is publicly available.

Keith Devlin is a mathematician at Stanford University. His email address is `devlin@stanford.edu`.

Over the course of my work on NIMD, I saw systems demonstrated under nonclassified circumstances that, in a few seconds, could produce incredibly detailed and deeply personal profiles of individuals based on an Internet search that pulled in many isolated *publicly available* facts. So when I hear officials from President Obama down say, "It's just metadata," I smell a deliberate attempt to mislead the population they are supposed to serve.

Metadata tells you practically everything you need to know! In fact, much of the focus of my NIMD work was on the degree to which contextual features of signals (information sources) play a role in the knowledge that can be acquired from that signal. I was asked to join Veridian's project in NIMD precisely to look at that issue.

The invitation to join the Veridian team that successfully bid for one of the thirteen NIMD projects that were eventually funded came as a result of research I had carried out since the late 1980s, much of which was summarized in a series of books [1]–[4]. That research focused on analyzing the role played by different kinds of contexts in the acquisition and transmission of information. Having pursued that research for many years in a purely academic fashion, I was at first surprised to find that in the early post-9/11 world, it suddenly occupied a central position.

Well, not exactly central. My work occupied one edge of the central focus of NIMD. While a lot of the program's research was focused on developing technologies that would (they hoped) in future help the US intelligence community "connect the dots" in order to prevent another terrorist attack, the Veridian project was from the start focused on trained human analysts. The mission was to find ways to make them better. In our project, cognitive science and psychology played a far greater role than writing code. So we spent a lot of time thinking about what happened to any results that the ever bigger and more powerful computer systems spewed out. How could we take an impossibly large amount of data and produce a human-sized output that a trained analyst could make effective use of? It would involve filtering, condensing, fusing, and processing information to a truly gigantic degree to provide that analyst (actually a team of analysts) with something manageable. And that was just the first step. That analyst would have to take his or her conclusions and start a cascade of persuasion and decision-making running up through the command chain until it landed on the desk of a person who could initiate an action—an action having huge ramifications for public safety, the pursuit of which would carry the risks of danger to many people and of possible massive political fallout.

That highly significant, human part of the decision chain tends to be totally overlooked when intelligence leaders and politicians talk in glowing terms about the safety yielded by massive data processing of huge trawls of information. But it should not be ignored. It is a crucial factor. It's also the factor I spent four years trying to address and hence the one thing I want to add to the debate.

Data mining systems don't identify and take out terrorist groups; people do. And those people—and those who send them into harm's way on our behalf—require not only accurate information but sufficient meta-information (information about the origins and reliability of that information) to have confidence in any decision they make. Veridian asked me to investigate whether the largely theoretical ideas I had been pursuing in my research since the mid-1980s could be brought to bear on this problem.

I think I am not being unfair to any of the many really talented teams that worked on NIMD when I say that we did not find a way to analyze vast amounts of wide-focused (i.e., not focused) intelligence data and provide intelligence analysts with the kind of information they needed to take preventive action, given all that would entail. On the contrary, everything we learned made it even clearer that such was an impossible goal.

I was, for example, not at all surprised to learn that the Boston Marathon bombers were on counterterrorist watch lists all the time they were planning and then carrying out their heinous act. That does not indicate a "failure" of counterterrorism. I guarantee that the massive computer data searches were turning up hundreds (maybe thousands; I have no way of knowing) of cases that had similar profiles. Even if an intelligence agent were to "have a hunch" that one of those many cases was about to blow—and it really would have to be a hunch—what are the chances that it would make its way successfully up the command chain to produce effective countermeasures?

And the bigger you make the dataset, the wider the information trawl, the more unlikely that it will lead to an effective countermeasure. Thus, not only did NIMD fail to meet its goal, but as the data collection grew (we did not know about the pending degree of growth at the time, of course, nor its scope), the more inaccessible that goal became.

It is reasonable to assume that the number of genuine potential terrorists is small and not growing (at least not dramatically). Consequently, the bigger the data trawl, the harder it is to spot the bad guys, no matter how much computing power you bring to the problem.

What we did learn from NIMD—at least, what I took away from the experience (I cannot speak for everyone else, though others agreed with me)—is that the methods and tools we developed could be of real benefit *if they were utilized in a highly targeted way*.

That's the real NIMD message. Use of the search and analysis methods should be narrow and deep, not wide and shallow. Focus all those tools and all that computing power on *deep* investigations of high-likelihood terrorists so that an intelligence analyst *can* be provided with all the information required to initiate a command chain call that *will* result in decisive action.

How do you identify those (relatively few) high-likelihood targets? The way intelligence communities always have: HUMINT (human intelligence). Not only is that the only effective way known, it does not require breaking laws and trampling the US Constitution. You get a court order and proceed lawfully. It's supposed to be the American way.

At the end of my NIMD work, I summarized some of what I had done in an unpublished paper that I posted on my Stanford homepage. It remains there to this day, dated July 15, 2005 (`http://www.stanford.edu/~kdevlin/Papers/Context_in_Reasoning.pdf`), bearing the annotation that it is an unfinished draft. In many ways, I wrote it as a road map of what to try next.

When I look back on that document now, it does not seem to represent much progress. (It also indicates that I was a very tiny cog in a very large engine. I made no major breakthroughs. I was just one among many mathematicians and others making small incremental steps in a very complex and messy domain.) On the other hand, when NIMD started in 2002, there wasn't even a sketch, let alone a road map.

After NIMD came to an end, I continued to pursue similar ideas in two subsequent Defense Department projects, first for a US naval contractor developing systems to process videos from surveillance drones, then a division of the US Army tasked with protecting US troops. When the army project came to an end in 2011, I assumed I would continue the work one way or another. I have, after all, learned a lot about this domain over the past twelve years.

But my purpose throughout has been to defend democratic freedom, not trample it. Personally, I would not trade freedom in order to prevent terrorist attacks, even if they were more frequent than the current de facto frequency of every ten years or so. If you do that, the terrorists have won. To give up those freedoms to run an Orwellian surveillance program that, based on the intelligence community's own research, is known to not only not work but to divert resources that if properly targeted (i.e., narrow and deep) could work, is completely wrong.

As things currently stand, I would not collaborate further with any of the US intelligence services. They have betrayed all of us who were glad to do what we could for the benefit of the free world and have used our work to trample over the Fourth Amendment, to do immense harm to US economic competitiveness, to weaken the Internet on which modern society depends, and to expose us to increased danger from our enemies (the latter two are "own-goals" that result from deliberately weakening the mathematical cryptosystems used in the Internet). I urge all my fellow mathematicians to take a similar stand.

References

[1] Keith Devlin, *Logic and Information*, Cambridge University Press, 1991.

[2] Keith Devlin and Duska Rosenberg, *Language at Work: Analyzing Communication Breakdown in the Workplace to Inform Systems Design*, Center for the Study of Language and Information, Stanford University and Cambridge University Press, 1996.

[3] Keith Devlin, *Goodbye Descartes: The End of Logic and the Search for a New Cosmology of the Mind*, Wiley, 1997.

[4] ______, *InfoSense: Turning Information into Knowledge*, W. H. Freeman, 1999.

The Mathematical Community and the National Security Agency

Andrew Odlyzko

The recent revelations about the NSA's spying programs are both dismaying and encouraging. What is encouraging is that they might lead not just to a reform of the intelligence agencies but also to a more serious look at what the ongoing and inevitable erosion of privacy is doing to our society. What is dismaying is less the intrusive data collection itself and more what it reveals about the decision-making processes inside the government.

These are all my personal opinions, but they are opinions based on over three decades of working on cryptography and security. Most of this time was in an industrial research lab. Currently, as a professor in a mathematics department, I regularly teach a course on cryptography. In addition, I am involved in a master's program on security technologies, where I lecture primarily on economics and psychology of security. I should also add that I have never had any kind of security clearance. Therefore I am not privy to any official secrets but at the same time am not restrained in expressing my opinions by any institutional ties.

My carefully considered view is that our society has become preoccupied with terrorism to an absurd and harmful degree. That is what has driven the intelligence agencies to the extreme measures

Andrew Odlyzko is professor of mathematics at the University of Minnesota in Minneapolis. His email address is `odlyzko@umn.edu`.

they have taken. Are those measures illegal? Given the enthusiastic support they have generally received from the executive, legislative, and judicial branches of the government, this is debatable. However, as the famous saying goes, much of this activity is worse than a crime; it's stupid. Terrorism is a threat to our society, but it is simply not an existential threat that justifies extraordinary measures. We face a variety of threats—from car accidents, which take about as many lives each month as the 9/11 tragedy, to weather (ranging from sudden disasters, such as hurricanes Katrina and Sandy, to the dangers from climate change), to global avian flu pandemics. The moves taken in the name of fighting terrorism, including the intrusive NSA data collection that has recently come to light and more generally the militarization of our society, are not justified by the dangers we currently face from terrorism. In fact, these moves will likely inhibit our ability to deal with many of the other threats and probably will even inhibit the antiterrorism campaign.

Still, the antiterrorism mantra is driving public policy, and it is corroding the already weakened trust in democratic governance. When high-level officials feel free to give the "least untruthful" answers or provide assurances of careful oversight and of intelligence successes that are then shown to be false, much is lost. For democracy to thrive, people have to be able to rely on both the competence and honesty of officials. The recent events have demonstrated major failings on both counts.

The official reactions to the recent revelations about the NSA's programs reveal a striking persistence of delusions about data security in Washington. The Snowden data breach is regarded as a one-time event. Instead, as the Manning leak (or should I say torrent?) already showed earlier, it should be seen as inevitable. Such disclosures arise from the growth in volume of data, our demonstrated inability to build truly secure systems, and the need for wide information sharing inside the intelligence agencies if those agencies are to be effective. (Let's not forget that one of the key findings of the investigations of 9/11 was that extensive relevant information about the terrorists was available inside the US government but was not shared properly.) A reasonable working assumption should be that several foreign intelligence agencies have extracted similar troves of secrets from US collections (in addition to what they collect on their own) and that this will continue.

The likelihood of a continuing leakage of official secrets is just one consequence of the rapid growth of data. The NSA projects revealed by Snowden are just a forerunner of more serious issues. Most of the data that the NSA has been using came from private organizations, and those are building their business cases on ever more intrusive data collection and exploitation. One report from the latest Consumer Electronics Show said that the "unsettling message" of that event was that "everything will be tracked." What the NSA has been amassing is tiny compared to what will be available soon. Further, most of that will be held in databases much more poorly protected than those of the NSA. Therefore we will have to worry about more than government officials misusing the data for political or other purposes (as J. Edgar Hoover infamously did, but, as far as we know, the NSA has not done recently), or NSA employees tracking their romantic interests (as they apparently have done in many instances). We will also have to watch out for what might be done by even less trustworthy employees of the private organizations controlling that data and by all those who manage to break into those (inevitably insecure) databases.

We will need to figure out how to live in a world where practically everything we not just say or write but even feel (at the physical level, as measured by a variety of sensors that are coming and are sure to be embraced for their health benefits) will be recorded. Therefore, it will potentially be available not just to the NSA but to all those who gain legitimate or illegitimate access to it. Just what laws, regulations, and other measures we as a society adopt to deal with these problems is a very thorny issue to which far more attention should be paid. I hope that the Snowden revelations will stimulate more serious consideration of these issues.

Given the ongoing erosion of privacy, the NSA programs we have learned about do not seem too serious. It's not that I approve of them. I do regard them as largely unnecessary and harmful and, in some cases, such as the deliberate weakening of security standards, inexcusable. I am in favor of curtailing those programs, bringing them under more rigorous oversight, and making them more open. However, I do not see the NSA as a rogue organization engaging in amoral activities. What it has been doing has been done with wide support of almost all responsible officials (even though this support was often gained with the help of large doses of obfuscation, fear, uncertainty, and doubt) and is not that far beyond what various private organizations have been doing. The NSA fills an important role both in spying on numerous hostile actors and setting security standards, and in protecting our information infrastructure. And mathematics plays a key part in enabling those functions. Hence, while I do favor reforms, I do not support the argument for the mathematical community to sever its ties to the NSA, and I do not discourage my students from applying there for jobs.