On Incentives and Cryptanalysis

Nadia Heninger

UC San Diego

June 12, 2024

A reductionist view of cryptographic research

Inputs

- Funding
- Human effort
- Coffee

Outputs

- Papers
- Trained students
- A vetted set of
  - cryptographic algorithms,
  - protocols,
  - recommended parameters,
  - and analysis frameworks

  for use by the world

The dual nature of cryptography

Our field has a constructive side and a destructive side.

One side builds cryptography, the other destroyes it.

Tempting thought: Perhaps cryptanalysis research is only necessary because complexity theory doesn't have strong computational lower bounds *yet*.

However: Even if we could prove strong computational lower bounds, attack work is still necessary to establish desired security properties to prove.

Messaging from the two sides of cryptography

Researchers on the constructive side are exploring the underlying truth of the computational universe and building a bright future full of privacy-preserving technologies that allow companies to collect and process all the data they want without any legal liability.

Researchers on the destructive side are destroyeing perfectly good cryptographic algorithms with impressive one-off tricks.

My claim: There is little market for cryptanalysis.

But market incentives influence the work that gets done.

Like all good dualities, there is no light without the dark.

For some areas of cryptography, the two sides of our field are out of balance.

This imbalance leads to insecurity.

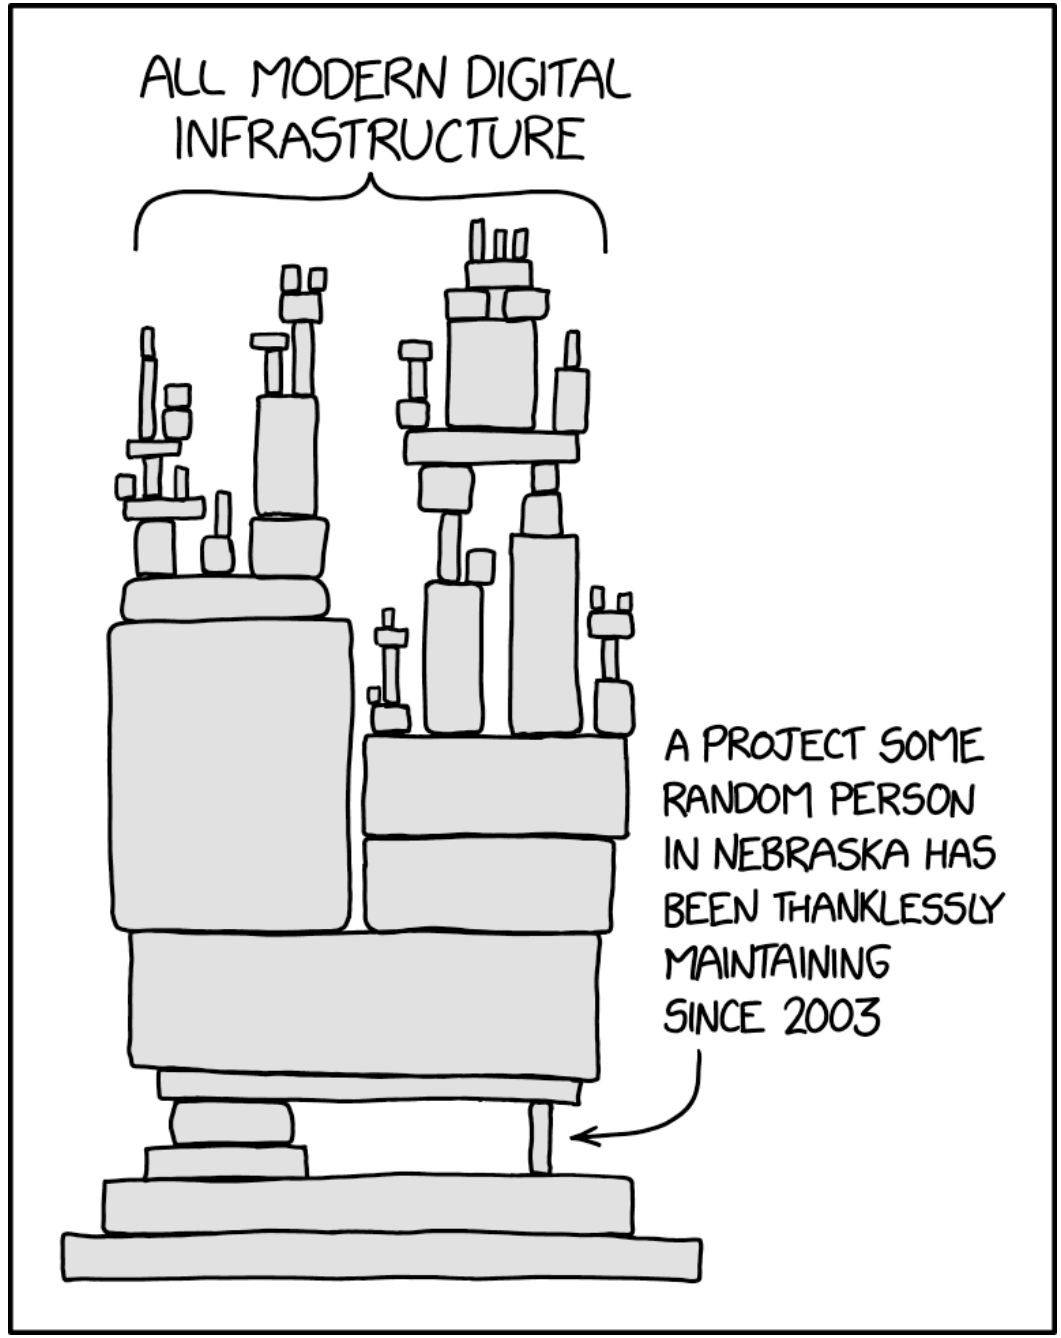
ALL MODERN DIGITAL INFRASTRUCTURE
A PROJECT SOME RANDOM PERSON IN NEBRASKA HAS BEEN THANKLESSLY MAINTAINING SINCE 2003

xkcd/2347

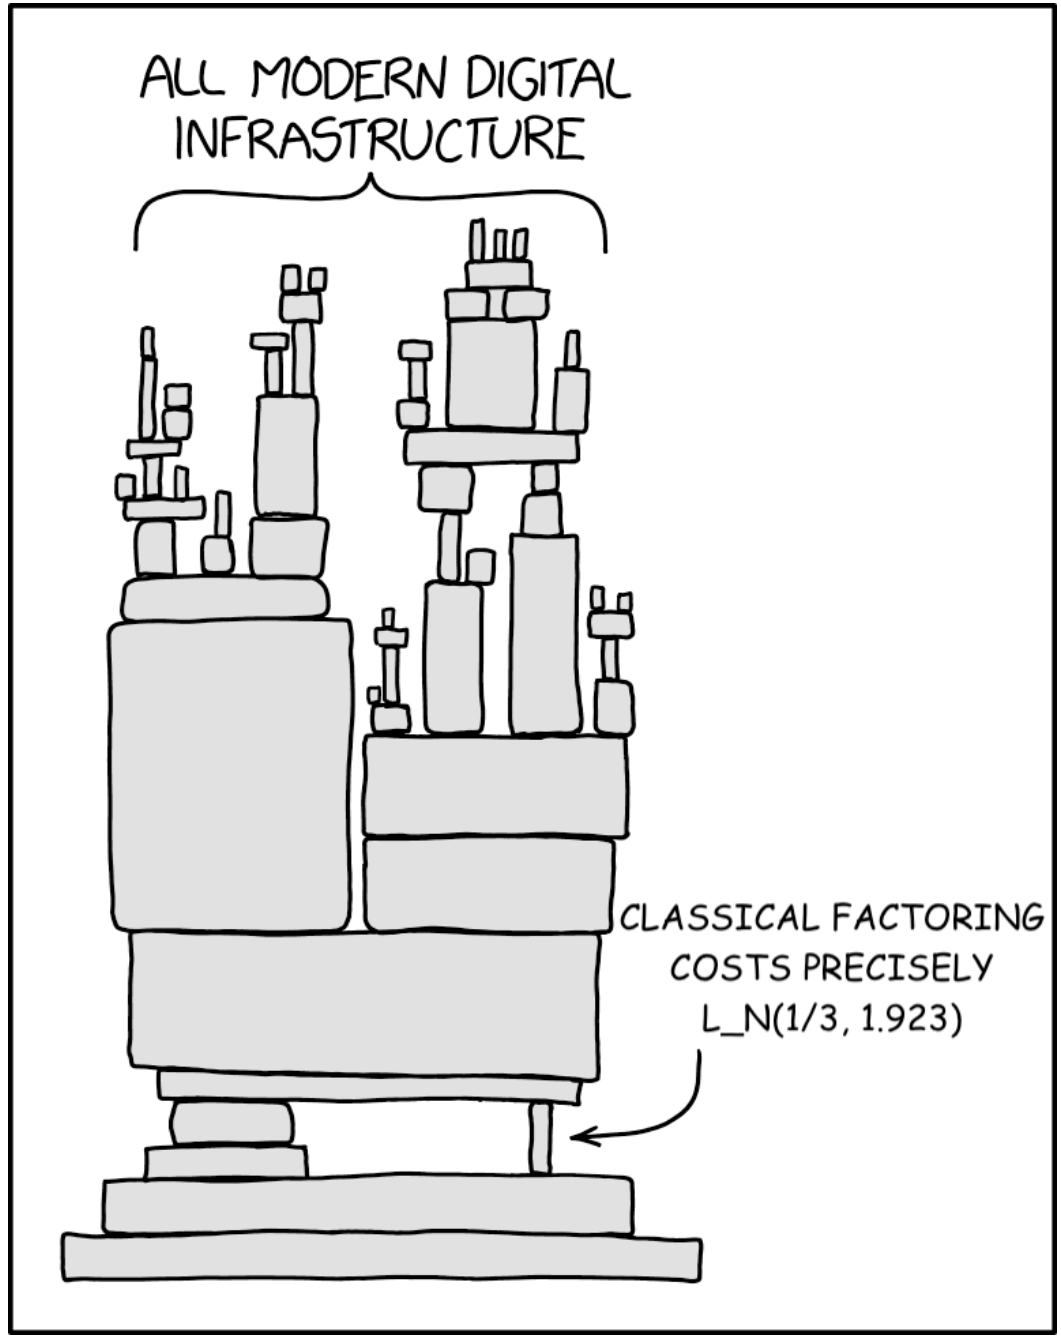

Apologies to xkcd/2347

I wanted to explore the position of cryptanalysis in different fields of cryptography.

I did a collection of totally unscientific interviews with a biased selection of researchers via chat, email, and Zoom.

- Scott Aaronson
- Paul Kocher
- Emmanuel Thomé
- Wouter Castryck and Thomas Decru
- Thomas Espitau
- Martin Albrecht
- Plus helpful feedback from Matt Green

Martin: “You’re asking all these weird introspection questions that are kind of hard to answer.”

Disclaimer: I live in the open research world

Paul Kocher: “The balance is really different in the government/classified world, where offense/attack get the lion’s share of the budget and glory. They may have the opposite imbalance from the academic world.”

Prioritizing offense over defense has led to problems:

- DES key strength
- Dual EC DRBG
- Continuing mistrust around standardization and algorithm recommendations

Area 1: Quantum computing

The looming threat of quantum computers is why you're all here today.

For cryptographers, quantum computers means Shor's algorithm.

But people building quantum computers mostly don't want to talk about Shor's algorithm.

They want to talk about quantum chemistry simulations:

- Solving world hunger through better fertilizer.
- Solving disease through drug discovery.
- Solving energy problems with batteries and solar cells.

Quantum Computing: Progress and Prospects

2019 National Academies study

"For quantum computing to be similarly successful, it must either create a virtuous cycle to fund the development of increasingly useful quantum computers (with government funding required to support this effort until this stage is reached) or be pursued by an organization committed to providing the necessary investment in order to achieve a practically useful machine even in the absence of intermediate returns or utility (although the total investment is likely to be prohibitively large)."

Scott Aaronson estimates the current investment in quantum computing is $O(1) billion per year.

I asked what this is for.

“There are the grounded people who correctly expect quantum simulation as the first big killer app, and who knows what else could come later?

There are the ones who talk about speedups for optimization or finance or classical ML in the near future. I think these people are mostly either fooling themselves, fooling others, or fooled by others.

Eventually, sure, Grover-like speedups could come into play for all these areas; the issue is that probably won’t beat classical for a VERY long time.”

Baseline setting/threat modeling

Realistically, if your attackers are breaking cryptography they are probably:

- pirating copy-protected content
- or a government.

This is good news! Cryptography is almost certainly not the weakest link in computer security.

All of the phishing and social engineering that actually compromise systems is mostly someone else's problem.

There is *no mass market* for Shor's algorithm.

The only customers are a handful of governments.

The people building quantum computers think of Shor's algorithm as a proof of concept and not a product.

Conclusions:

If Shor's algorithm becomes feasible, it is a by-product of other things the market cares more about.

Scale of engineering challenges implies classified progress probably not too far ahead of open research.

2048-bit RSA factoring: classical or quantum first?

Wouter Castryck: quantum

Emmanuel Thomé: classical

Thomas Espitau: Original Shor's algorithm infeasible (millions of physical qubits), but new opinion after recent algorithmic improvements:

	gates	qubits
Shor 1994	$O(n^2 \log n)$	$O(n)$
Regev 2023	$O(n^{3/2} \log n)$	$O(n^{3/2})$
Ragavan Vaikuntanathan 2024	$O(n^{3/2} \log n)$	$O(n \log n)$

Algorithmic advances and quantum engineering may meet in the middle.

Area 2: Classical factoring and discrete log

Yes, people still use RSA and mod p Diffie-Hellman

- $\approx$ 30% of SSH connections use RSA host keys in active and passive data collected from UCSD
- $\approx$ 90% of TLS 1.2 connections use RSA signatures in passive data from UCSD
- TLS 1.3 supports a list of “named” finite field Diffie-Hellman groups listed in RFC 7919

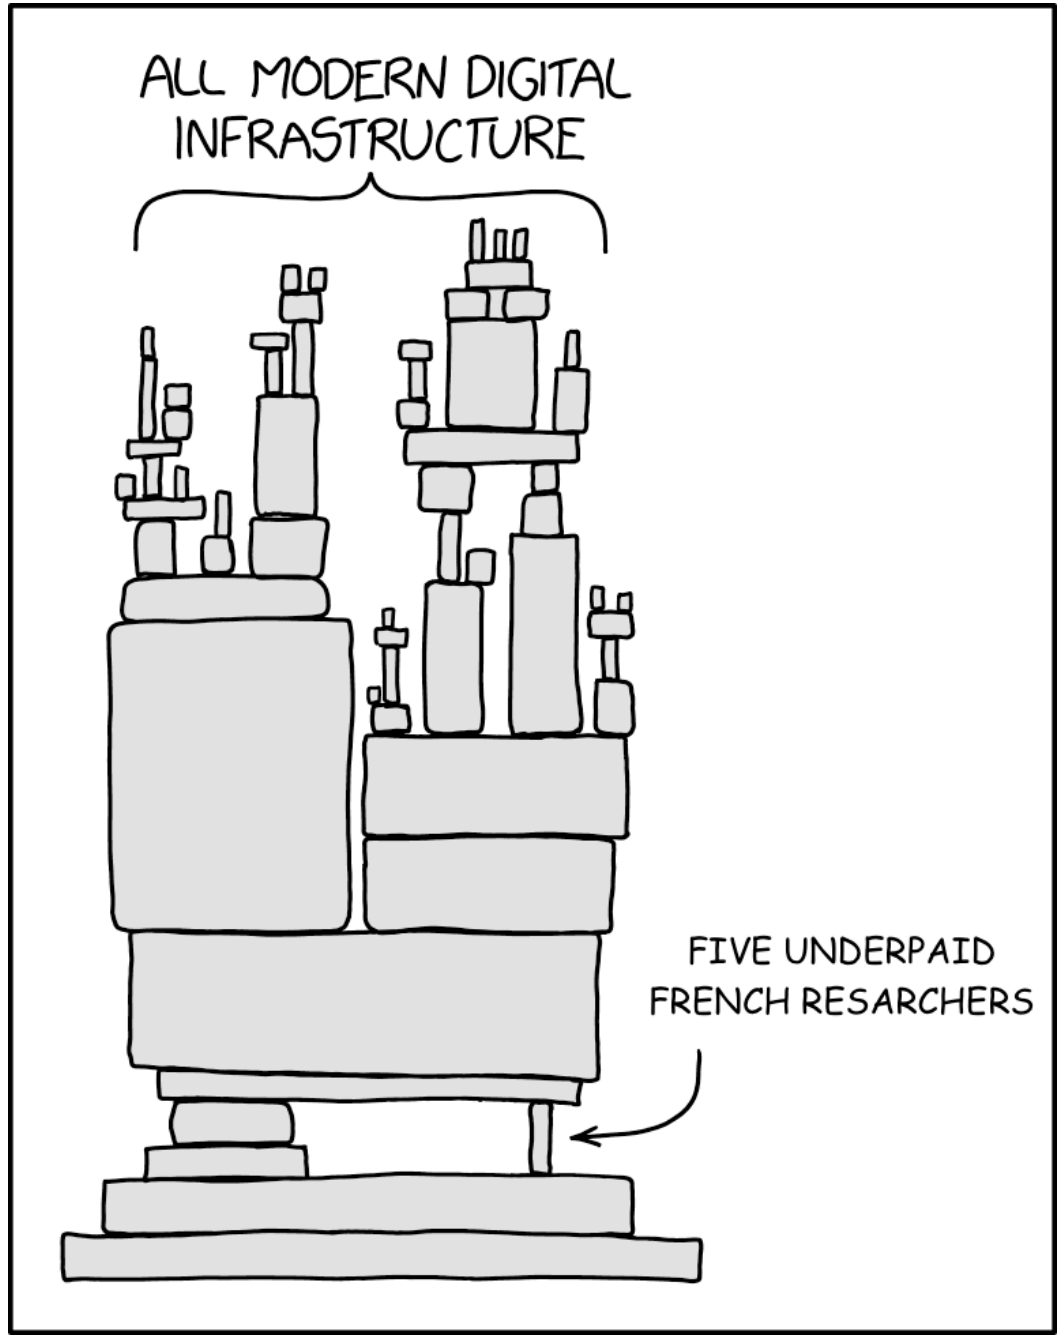

Apologies to xkcd/2347

The number field sieve algorithm for factoring:

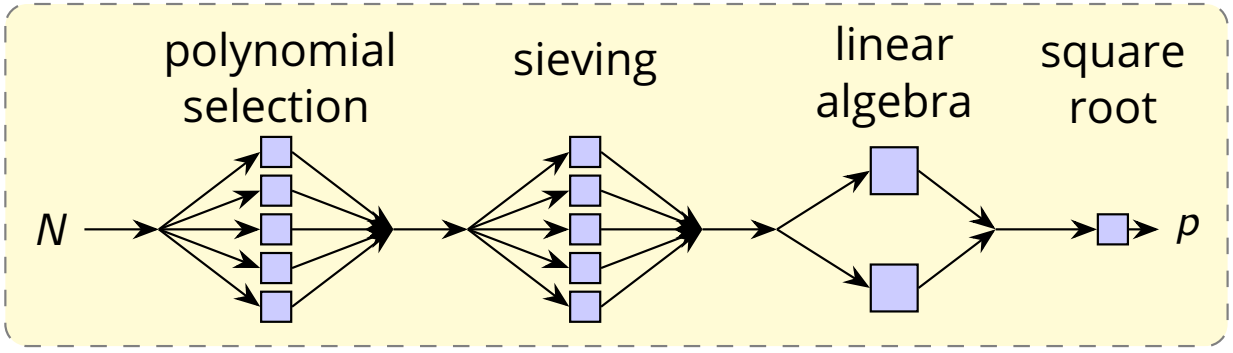

The number field sieve algorithm for discrete log:

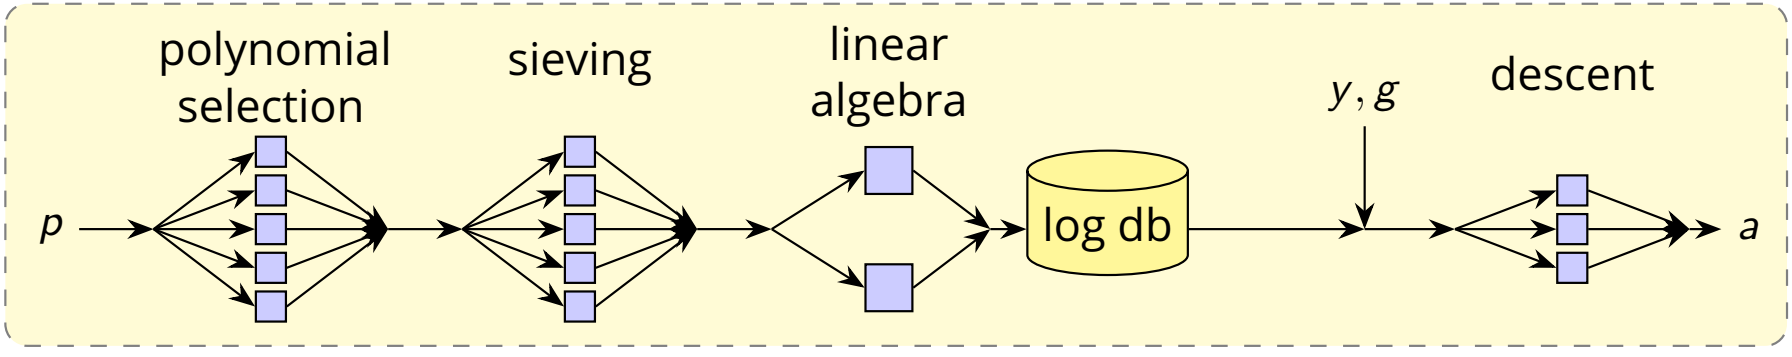

$$L_N(1/3, \sqrt[3]{64/9}) = e^{(1.923+o(1))(\ln N)^{1/3}(\ln \ln N)^{2/3}}$$

Factoring and discrete log records

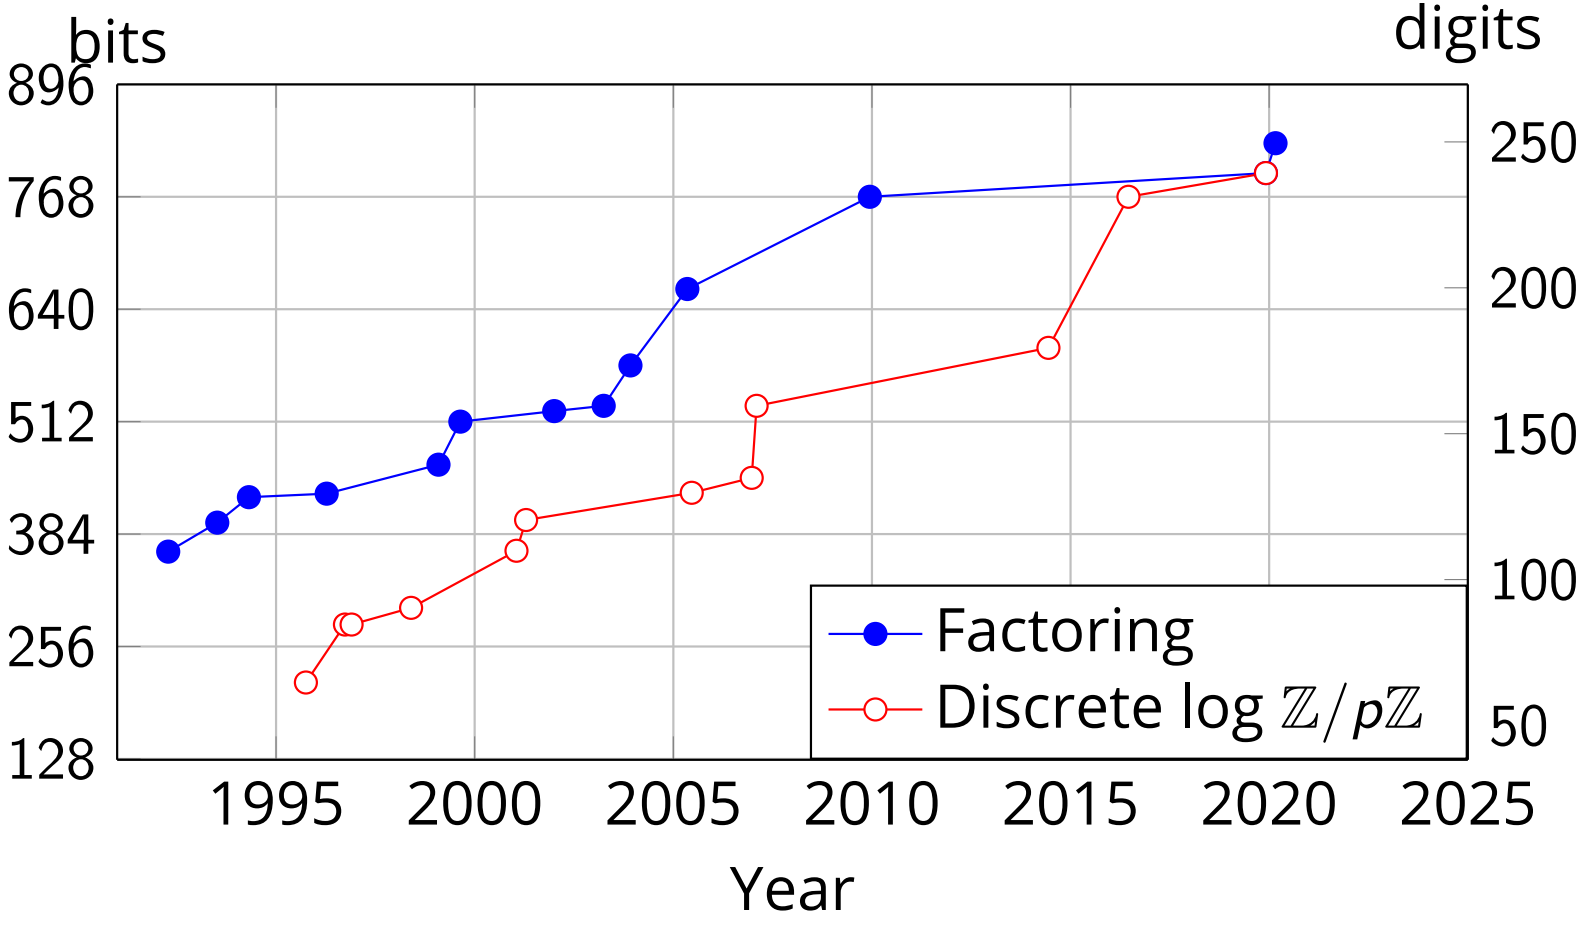

CADO-NFS

Crible Algébrique: Distribution, Optimisation - Number Field Sieve

Home
Prerequisites
Download
Development
Bugs / Support
User Reports

Introduction

CADO-NFS is a complete implementation in C/C++ of the Number Field Sieve (NFS) algorithm for factoring integers and computing discrete logarithms in finite fields. It consists in various programs corresponding to all the phases of the algorithm, and a general script that runs them, possibly in parallel over a network of computers. CADO-NFS is distributed under the Gnu Lesser General Public License (LGPL) version 2.1 (or any later version).

CADO-NFS is the result of a collaborative effort involving many persons, over various periods of time. The current list of active contributors can be extracted from the git repository or from the openhub.net page. A tentative list of CADO-NFS authors is (alphabetical order):

- Shi Bai
- Razvan Barbulescu
- Cyril Bouvier
- Richard Brent
- Christophe Clavier
- Jérémie Detrey
- Andreas Enge
- Alain Filbois
- Nuno Franco
- Pierrick Gaudry
- Laurent Grémy
- Aurore Guillevic
- Nadia Heninger
- Laurent Imbert
- Alexander Kruppa
- Jérome Milan
- François Morain
- Lionel Muller
- Thomas Prest
- Thomas Richard
- Emmanuel Thomé
- Marion Videau
- Paul Zimmermann

Citing CADO-NFS

The recommended way to cite CADO-NFS in a scientific publication is (bibtex entry):

- `The CADO-NFS Development Team. CADO-NFS, An Implementation of the Number Field Sieve Algorithm, Release 2.3.0, 2017, http://cado-nfs.inria.fr/`

where the release number and date should be changed to correspond to the version you actually used. If you used the development version, it is a good idea to give the git revision number. You can, at the very least, say "development version" and give the current date. You can base your citation entry on the following template (bibtex entry for the development version, to be completed):

- `The CADO-NFS Development Team. CADO-NFS, An Implementation of the Number Field Sieve Algorithm, development version, 20XX, http://cado-nfs.inria.fr/`

“Recent” computational advances

Asymptotic general number field sieve running time is:

$$L_p(1/3, \sqrt[3]{64/9}) = e^{(1.923+o(1))(\ln p)^{1/3}(\ln \ln p)^{2/3}}$$

This has been the same since the early 1990s.

2013: Discrete log descent phase improved to $L_p(1/3, 1.232)$ [Barbulescu]

2013: A quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic [Barbulescu, Gaudry, Joux, Thomé] (Function Field Sieve)

Any (even small) improvement in NFS running time would have a large impact on RSA/DH key sizes:

- $L_p(1/3, \sqrt[3]{32/9})$ is academically easy for 1024 bits
- $L_p(1/4, c)$ could take out 2048 bits.

Emmanuel Thomé: ≈1M core-years for 1024-bit RSA.
Probably no less than 250K, no more than 2M.

	core-years	year done	EC2 cost
RSA-512	1	1999	$90
RSA-768	1,000	2009	$90,000
RSA-829	2,700	2020	$240,000
RSA-1024	1,000,000	????	$90,000,000

	core-years	year done	EC2 cost
DH-512	10	≈2007	$1,000
DH-768	5,000	2016	$450,000
DH-795	3,000	2019	$270,000
DH-1024	3,000,000	????	$270,000,000

1vCPU: $0.01/hour

Lack of incentives for factoring/discrete log research

- Separation between constructions and cryptanalysis.
  - Clean hardness assumptions allow for clean constructive proofs but insulate users from mathematics.
  - Emmanuel: “Separation lets people think that the math assumption is something you can store for good in the area of knowledge that it’s known to be hard.”
- Funding situation is okay for France.
  - INRIA supports software development.
  - Emmanuel’s group funded mostly by ANR.
  - No industry interest.
  - Doesn’t fit in EU cybersecurity framework.
- Not a good topic for grad students.
  - Requires deep background in both number theory and computer implementation.
  - Not trendy or sexy.
  - Unclear reward.

What is the correct running time for factoring/dlog?

Emmanuel Thomé and Antoine Joux: Quasi-polynomial.

But: Community interest and funding are minimal.

- Compare funding level of at most low hundreds of thousands of dollars/year to $O(1)B for quantum.

Increased investment would help, but unclear payoff:

- Emmanuel: Small-characteristic dlog improvement could have been found 5–10 years earlier.
- Sometimes breaks in cryptography have constructive implications; what constructive applications could factoring enable?

Unique INRIA structure facilitates this type of research in a way that US academia does not.

Opportunity!

If you want to speed up the adoption of post-quantum cryptography, you may not need to wait for the scientists to build an actual quantum computer.

You could pro-actively take out the competition classically!

Area 3: Isogenies

An efficient key recovery attack on SIDH

Wouter Castryck[1,2] and Thomas Decru[1]

[1] imec-COSIC, KU Leuven, Belgium
[2] Vakgroep Wiskunde: Algebra en Meetkunde, Universiteit Gent, Belgium

"A run on the SIKEp434 parameters, previously believed to meet NIST?s quantum security level 1, took roughly 10m, again on a single core. We also ran the code on random instances of SIKEp503 (level 2), SIKEp610 (level 3) and SIKEp751 (level 5), which on average took about 20m, 55m and 3h15m, respectively."

What is up with isogenies?

Wouter Castryck and Thomas Decru:

- Emphasize that pure isogeny problems are definitely not broken, only those with auxiliary information.
- They are optimistic about isogenies.
- This kind of cryptanalysis is like going back to the good old days of the 1980s and 1990s; it is not worrying.
- Attack follows from techniques that were known in the 1990s.
- If more mathematicians had looked at cryptography, SIKE could have been broken earlier.

Incentives for isogeny research

- It's a small field; little division between constructive cryptography and cryptanalysis.
- Castryck and Decru were not trying to find the SIKE break but discovered relevant math theorems.
- Funding:
  - Both mathematicians at heart, but funding situation better for cryptography.
  - Plenty of national and EU government funding.
  - NIST competition has been a great catalyst for research and funding.
- A lot of activity, enthusiasm, new mathematical ideas in isogeny club.

Exciting times for isogenies

The fundamental problems in isogeny-based cryptography are very appealing to computational number theorists.

Wouter Castryck: Specific analogy to story with pairings.

- Initially used to break schemes, then constructive use.
- Higher-dimensional isogenies are now being used to build schemes.

Hopes for fancy crypto constructions from isogenies.

What about elliptic curve discrete log?

Wouter trusts elliptic curve discrete log less than factoring.

Dictionary between number theory and algebraic geometry:

number field	function field
$\mathbb{Z}$	$\mathbb{F}[x]$
factoring hard	factoring easy
SVP hard	SVP easy
Riemann Hypothesis	Weil conjectures

Elliptic curve discrete log lives on the right side of this table.

- ECDL has more structure than factoring.
- It's an isogeny-finding problem.
- Existing isogeny attacks fail because degree is secret.
- But geometry is easier over finite fields.

A Riddle Wrapped in an Enigma

Koblitz and Menezes 2015

Flash back to the NSA's 2015 post quantum announcement:

"For those partners and vendors that have not yet made the transition to Suite B algorithms, we recommend not making a significant expenditure to do so at this point but instead to prepare for the upcoming quantum resistant algorithm transition"

Koblitz and Menezes: "Does the NSA have an $n^{1/3}$-algorithm for finding elliptic curve discrete logs? The reason for wondering about this is that in the latest revision of Suite B the NSA has dropped P-256, leaving only P-384."

Now: NSA not a fan of hybrid post-quantum schemes.

Area 4: Lattices

The industry of lattices

My lattice interviews were with Thomas Espitau (PQShield) and Martin Albrecht (Sandbox AQ).

Both do lattice cryptanalysis for post-quantum startups.

For both companies, having cryptanalysts on staff establishes expertise.

Both researchers discussed balancing industry demands:

- Thomas partners with academia for more far out or large scale cryptanalytic work.
- Martin maintains academic affiliation in order to do applied cryptography and vulnerability finding.

Community and incentives

Lattices are between the situation of factoring and isogenies.

Bigger field, so researchers are more specialized: mostly construction or mostly analysis.

NIST competition has catalyzed plenty of EU and UK funding.

Either positive or negative results from lattice algorithm analysis are a contribution. (In contrast to a failure to improve factoring.)

Now that NIST competition is wrapping up, research may slow down.

Confidence in lattice assumptions

Both Thomas and Martin fairly confident in LWE and SIS; not much do to cryptanalytically unless a scheme makes really bad design choices.

However, fancy cryptography like threshold signatures, functional encryption, FHE all have heavy machinery that doesn't rely on standard assumptions.

They tweak the assumptions in ways that obscure vulnerabilities.

Constructive researchers are producing more lattice assumptions than there is capacity to analyze.

Lattices in cryptography vs. number theory

Thomas Espitau: Lattice questions in cryptography different from questions that interest number theorists.

Fundamental geometric question in cryptography is to reduce the largest possible lattice in the fastest way possible.

In geometry of numbers, often end up with huge number of very small lattices to reduce quickly.

Optimizations in single hard instance vs. batch of easy problems are not the same.

Algebraic vs. unstructured lattices

Martin: If there are surprises, it will be in intersection of algebraic lattices and quantum.

- Existing sharp drop for ideal SVP on a quantum computer at $2^{\sqrt{n}}$ in poly time. [Cramer Ducas Wesolowski 2016]
- Maybe ideal SVP is also easier on a classical computer?

Thomas:

- Fairly confident that actual security of structured variants like NTRU and module lattices is less than current estimates.
- More likely to be a few bits of reduction rather than a complete break.

On the difficulty of partial or negative results

A common complaint is that it's harder to publish partial results in cryptanalysis than in construction.

Martin: In symmetric cryptography, have attacks that scale with number of rounds.

Easier to publish partial results of this kind.

NIST process caused a lot more interest in small improvements for lattice cryptanalysis.

We need to publish more results like "I tried a nontrivial attack for algebraic lattices and it didn't work" in places like CFAIL.

The case of symmetric cryptography and protocols

Emmanuel: Symmetric cryptography has less separation between constructive and analysis side.

Martin and Paul: Modern protocol analysis is clean intersection of cryptanalysis and construction.

When you construct a protocol and a security proof:

- If you succeed you have a security proof.
- If you fail you have an attack.

Conclusions and discussion

Paul Kocher: “Cryptanalysis used to be a lot more fun.

Differential power analysis came about because of the shift toward better algorithms and protocols—my first attacks against smart cards were cryptanalytic breaks, but it was obvious I’d need more effective tools.”

Emmanuel Thomé: “There’s been a shifting balance between cryptography and cryptanalysis in the community. If you look back 25 or 30 years, there were a lot more papers related to cryptanalysis.”

- Post-quantum cryptography might be closer to what classical cryptanalysis was like in the 1990s.
- But this era may be ending, and there’s a risk of neglecting the analysis side going forward.

- NIST process has catalyzed a lot of research on constructions and analysis; excitement may move elsewhere once it’s over.

- Need to ensure ongoing analysis and translation to/from mathematics advances.

- Imbalance between incentives for constructive cryptography vs. cryptanalysis leads to late surprises.

- How many people in the world deeply understand quantum algorithms and the state of the art in post-quantum constructions?
  (e.g. exactly which lattice parameters are hard)

- We need more US number theorists in the open cryptography research community.

- Situation is healthier in Europe.

- Need more research on fancy lattice assumptions.
  - Good project for grad students!
  - (But a focus on constructions is probably best for employment.)

- Classical factoring/discrete log/ECDL are not solved yet.
  - But I don’t know how to fix incentive gap for research.
  - Not a good project for grad students.