**From theoretical possibility
to near practicality.**

BY MICHAEL WALFISH AND ANDREW J. BLUMBERG

Verifying Computations without Reexecuting Them

IN THIS SETUP, a single reliable PC can monitor the operation of a herd of supercomputers working with possibly extremely powerful but unreliable software and untested hardware.

—Babai, Fortnow, Levin, Szegedy, 1991⁴

How *can* a single PC check a herd of supercomputers with unreliable software and untested hardware?

This classic problem is particularly relevant today, as much computation is now outsourced: it is performed by machines that are rented, remote, or both. For example, service providers (SPs) now offer storage, computation, managed desktops, and more. As a result, relatively weak devices (phones, tablets, laptops, and PCs) can run computations (storage, image processing, data

analysis, video encoding, and so on) on banks of machines controlled by someone else.

This arrangement is known as cloud computing, and its promise is enormous. A lone graduate student with an intensive analysis of genome data can now rent a hundred computers for 12 hours for less than \$200. And many companies now run their core computing tasks (websites, application logic, storage) on machines owned by SPs, which automatically replicate applications to meet demand. Without cloud computing, these examples would require buying hundreds of physical machines when demand spikes ... and then selling them back the next day.

But with this promise comes risk. SPs are complex and large-scale, making it unlikely that execution is always correct. Moreover, SPs do not necessarily have strong incentives to ensure correctness. Finally, SPs are black boxes, so faults—which can include misconfigurations, corruption of data in storage or transit, hardware problems, malicious operation, and more³³—are unlikely to be detectable. This raises a central question, which goes beyond cloud computing: *How can we ever trust results computed by a third-party, or the integrity of data stored by such a party?*

A common answer is to replicate computations.^{15,16,34} However, replication assumes that failures are uncorrelated, which may not be a valid assumption: the hardware and software

» key insights

- Researchers have built systems that allow a local computer to efficiently check the correctness of a remote execution.
- This is a potentially radical development; there are many applications, such as defending against an untrusted hardware supply chain, providing confidence in cloud computing, and enabling new kinds of distributed systems.
- Key enablers are PCPs and related constructs, which have long been of intense theoretical interest.
- Bringing this theory to near practicality is the focus of an exciting new interdisciplinary research area.

platforms in cloud computing are often homogeneous. Another answer is auditing—checking the responses in a small sample—but this assumes that incorrect outputs, if they occur, are relatively frequent. Still other solutions involve trusted hardware³⁹ or attestation,³⁷ but these mechanisms require a chain of trust and assumptions that the hardware or a hypervisor works correctly.

But what if the third party returned its results along with a proof that the results were computed correctly? And what if the proof were inexpensive to check, compared to the cost of redoing the computation? Then few assumptions would be needed about the

types of faults that can occur: either the proof would check, or not. We call this vision *proof-based verifiable computation*, and the question now becomes: *Can this vision be realized for a wide class of computations?*

Deep results in complexity theory and cryptography tell us that in principle the answer is “yes.” Probabilistic proof systems^{24,44}—which include interactive proofs (IPs),^{3,26,32} probabilistically checkable proofs (PCPs),^{1,2,44} and argument systems¹³ (PCPs coupled with cryptographic commitments³⁰)—consist of two parties: a *verifier* and a *prover*. In these protocols, the prover can efficiently convince the verifier of a mathemati-

cal assertion. In fact, the acclaimed PCP theorem,^{1,2} together with refinements,²⁷ implies that a verifier only has to check three randomly chosen bits in a suitably encoded proof!

Meanwhile, the claim “*this* program, when executed on *this* input, produces *that* output” can be represented as a mathematical assertion of the necessary form. The only requirement is the verifier knows the program, the input (or at least a digest, or fingerprint, of the input), and the purported output. And this requirement is met in many uses of outsourced computing; examples include Map Reduce-style text processing, scientific computing and simulations, da-

tabase queries, and Web request-response.^a Indeed, although the modern significance of PCPs lies elsewhere, an original motivation was verifying the correctness of remotely executed computations: the paper quoted in our epigraph⁴ was one of the seminal works that led to the PCP theorem.

However, for decades these approaches to verifiable computation were purely theoretical. Interactive protocols were prohibitive (exponential-time) for the prover and did not appear to save the verifier work. The proofs arising from the PCP theorem (despite asymptotic improvements^{10,20}) were so long and complicated that it would have taken thousands of years to generate and check them, and would have needed more storage bits than there are atoms in the universe.

But beginning around 2007, a number of theoretical works achieved results that were especially relevant to the problem of verifying cloud computations. Goldwasser et al., in their influential Muggles paper,²⁵ refocused the theory community's attention on verifying outsourced computations, in the context of an interactive proof system that required only *polynomial* work from the prover, and that ap-

plied to computations expressed as certain kinds of circuits. Ishai et al.²⁹ proposed a novel cryptographic commitment to an entire linear function, and used this primitive to apply *simple* PCP constructions to verifying general-purpose outsourced computations. A couple of years later, Gentry's breakthrough protocol for fully homomorphic encryption (FHE)²³ led to work (GGP) on *non-interactive* protocols for general-purpose computations.^{17,21} These developments were exciting, but, as with the earlier work, implementations were thought to be out of the question. So the theory continued to remain theory—until recently.

The last few years have seen a number of projects overturn the conventional wisdom about the hopeless impracticality of proof-based verifiable computation. These projects aim squarely at building real systems based on the theory mentioned earlier, specifically PCPs and Muggles (FHE-based protocols still seem too expensive). The improvements over naive theoretical protocols are dramatic; it is not uncommon to read about factor-of-a-trillion speedups. The projects take different approaches, but broadly speaking, they apply both refinements of the theory and systems techniques. Some projects include a full pipeline: a programmer specifies a computation in a high-level language, and then a compiler (a) transforms the computation to the for-

malism that the verification machinery uses and (b) outputs executables that implement the verifier and prover. As a result, achieving verifiability is no harder for the programmer than writing the code in the first place.

The goal of this article is to survey this blossoming area of research. This is an exciting time for work on verifiable computation: while none of the works we discuss here is practical enough for its inventors to raise venture capital for a startup, they merit being referred to as “systems.” Moreover, many of the open problems cut across subdisciplines of computer science: programming languages, parallel computing, systems, complexity theory, and cryptography. The pace of progress has been rapid, and we believe real applications of these techniques will appear in the next few years.

A note about scope. We focus on solutions that provide integrity to, and in principle are very efficient for, the verifier.^b Thus, we do not cover exciting work on efficient implementations of secure multiparty protocols.^{28,31} We also exclude FHE-based approaches based on GGP²¹ (as noted earlier, these techniques seem too expensive) and the vast body of domain-specific solutions (surveyed elsewhere^{36,42,47}).

A Problem and Theoretical Solutions

The problem statement and some observations about it. A verifier sends the specification of a computation p (for example, the text of a program) and input x to a *prover*. The prover computes an output y and returns it to the verifier. If $y = p(x)$, then a correct prover should be able to convince the verifier of y 's correctness, either by answering some questions or by providing a certificate of correctness. Otherwise, the verifier should reject y with high probability.

In any protocol that solves this problem, we desire three things. First, the protocol should provide some advantage to the verifier: either the protocol should be cheaper for the verifier than executing $p(x)$ locally, or else the protocol should handle computations p that

a The condition does not hold for “proprietary” computations whose logic is concealed from the verifier. However, the theory can be adapted to this case too, as we discuss near the end of the article.

b Some of the systems (those known as *zero knowledge SNARKs*) also keep the prover's input private.

Figure 1. A framework for solving the problem in theory.

Framework in which a verifier can check that, for a computation p and desired input x , the prover's purported output y is correct. Step 1: The verifier and prover compile p , which is expressed in a high-level language (for example, C) into a Boolean circuit, C . Step 2: the prover executes the computation, obtaining a transcript for the execution of C on x . Step 3: the prover encodes the transcript, to make it suitable for efficient querying by the verifier. Step 4: the verifier probabilistically queries the encoded transcript; the structure of this step varies among the protocols (for example, in some of the works,^{7,36} explicit queries are established before the protocol begins, and this step requires sending only the prover's responses).

the verifier could not execute itself (for example, operations on state private to the prover). Second, we do not want to make any assumptions that the prover follows the protocol. Third, p should be general; later, we will have to make some compromises, but for now, p should be seen as encompassing all C programs whose running time can be statically bounded given the input size.

Some reflections about this setup are in order. To begin with, we are willing to accept some overhead for the prover, as we expect assurance to have a price. Something else to note is that whereas some approaches to computer security attempt to reason about what *incorrect* behavior looks like (think of spam detection, for instance), we will specify *correct* behavior and ensure anything other than this behavior is visible as such; this frees us from having to enumerate, or reason about, the possible failures of the prover.

Finally, one might wonder: How does our problem statement relate to NP-complete problems, which are easy to check but believed to be hard to solve? The answer is that the “check” of an NP solution requires the checking entity to do work polynomial in the length of the solution, whereas our verifier will do far less work than that! (Randomness makes this possible.) Another answer is that many computations (for example, those that run in deterministic polynomial time) do not admit an asymmetric checking structure—unless one invokes the fascinating body of theory that we turn to now.

A framework for solving the problem in theory. A framework for solving the problem in theory is depicted in Figure 1. Because *Boolean circuits* (networks of AND, OR, NOT gates) work naturally with the verification machinery, the first step is for the verifier and prover to transform the computation to such a circuit. This transformation is possible because any of our computations p is naturally modeled by a Turing Machine (TM); meanwhile, a TM can be “unrolled” into a Boolean circuit that is not much larger than the number of steps in the computation.

Thus, from now on, we will talk only about the circuit C that represents our computation p (Figure 1, step 1). Consistent with the problem statement earlier, the verifier supplies the input

Encoding a Circuit's Execution in a Polynomial

This sidebar demonstrates a connection between *program execution* and *polynomials*. As a warmup, consider an AND gate, with two (binary) inputs, z_1, z_2 . One can represent its execution as a *function*:

$$\text{AND}(z_1, z_2) = z_1 \cdot z_2.$$

Here, the function AND behaves exactly as the gate would: it evaluates to 1 if z_1 and z_2 are both 1, and it evaluates to 0 in the other three cases. Now, consider this function of three variables:

$$\begin{aligned} f_{\text{AND}}(z_1, z_2, z_3) &= z_3 - \text{AND}(z_1, z_2) \\ &= z_3 - z_1 \cdot z_2. \end{aligned}$$

Observe that $f_{\text{AND}}(z_1, z_2, z_3)$ evaluates to 0 when, and only when, z_3 equals the AND of z_1 and z_2 . For example, $f_{\text{AND}}(1, 1, 1) = 0$ and $f_{\text{AND}}(0, 1, 0) = 0$ (both of these cases correspond to correct computation by an AND gate), but $f_{\text{AND}}(0, 1, 1) \neq 0$.

We can do the same thing with an OR gate:

$$f_{\text{OR}}(z_1, z_2, z_3) = z_3 - z_1 - z_2 + z_1 \cdot z_2.$$

For example, $f_{\text{OR}}(0, 0, 0) = 0$, $f_{\text{OR}}(1, 1, 1) = 0$, and $f_{\text{OR}}(0, 1, 0) \neq 0$. In all of these cases, the function is determining whether its third argument (z_3) does in fact represent the OR of its first two arguments (z_1 and z_2). Finally, we can do this with a NOT gate:

$$f_{\text{NOT}}(z_1, z_2) = 1 - z_1 + z_2.$$

The intent of this warmup is to communicate that *the correct execution of a gate can be encoded in whether some function evaluates to 0*. Such a function is known as an *arithmetization* of the gate.

Now, we extend the idea to a line $L(t)$ over a dummy variable, t :

$$L(t) = (z_3 - z_1 \cdot z_2) \cdot t.$$

This line is parameterized by z_1, z_2 , and z_3 ; depending on their values, $L(t)$ becomes different lines. A crucial fact is that this line is the 0-line (that is, it covers the horizontal axis, or equivalently, evaluates to 0 for all values of t) if and only if z_3 is the AND of z_1 and z_2 . This is because the y-intercept of $L(t)$ is always 0, and the slope of $L(t)$ is given by the function f_{AND} . Indeed, if $(z_1, z_2, z_3) = (1, 1, 0)$, which corresponds to an incorrect computation of AND, then $L(t) = t$, a line that crosses the horizontal axis only once. On the other hand, if $(z_1, z_2, z_3) = (0, 1, 0)$, which corresponds to a correct computation of AND, then $L(t) = 0 \cdot t$, which is 0 for all values of t .

We can generalize this idea to higher order polynomials (a line is just a degree-1 polynomial). Consider the following degree-2 polynomial, or parabola, $Q(t)$ in the variable t :

$$Q(t) = [z_1 \cdot z_2 (1 - z_3) + z_3 (1 - z_1 \cdot z_2)] t^2 + (z_3 - z_1 \cdot z_2) \cdot t.$$

As with $L(t)$, the parabola $Q(t)$ is parameterized by z_1, z_2 , and z_3 ; they determine the coefficients. And as with $L(t)$, this parabola is the 0 parabola (all coefficients are 0, causing the parabola to evaluate to 0 for all values of t) if and only if z_3 is the AND of z_1 and z_2 . For example, if $(z_1, z_2, z_3) = (1, 1, 0)$, which is an incorrect computation of AND, then $Q(t) = t^2 - t$, which crosses the horizontal axis only at $t = 0$ and $t = 1$. On the other hand, if $(z_1, z_2, z_3) = (0, 1, 0)$, which is a *correct* computation of AND, then $Q(t) = 0 \cdot t^2 + 0 \cdot t$, which of course is 0 for all values of t .

Summarizing, $L(t)$ (resp., $Q(t)$) is the 0-line (resp., 0-parabola) when and only when $z_3 = \text{AND}(z_1, z_2)$. This concept is powerful, for if there is an efficient way to check whether a polynomial is 0, then there is now an efficient check of whether a circuit was executed correctly (here, we have generalized to circuit from gate). And there are indeed such checks of polynomials, as described in the sidebar on page 78.

Probabilistically Checking a Transcript's Validity

This sidebar explains the idea behind a fast probabilistic check of a transcript's validity. As noted in the text, computations are expressed as Boolean circuits. As an example, consider the following computation, where x_1 and x_2 are bits:

```
if ( $x_1 \neq x_2$ ) {  $y = 1$  } else {  $y = 0$  }
```

This computation could be represented by a single XOR gate; for illustration, we represent it in terms of AND, OR, NOT:

To establish the correctness of a purported output y given inputs x_1, x_2 , the prover must demonstrate to the verifier that it has a *valid transcript* (see text) for this circuit. A naive way to do this is for the prover to simply send the transcript to the verifier, and for the verifier to check it step-by-step. However, that would take as much time as the computation.

Instead, the two parties encode the computation as a polynomial $Q(t)$ over a dummy variable t . The sidebar on page 77 gives an example of this process for a single gate, but the idea generalizes to a full circuit. The result is a polynomial $Q(t)$ that evaluates to 0 for all t if and only if each gate's output in the transcript follows correctly from its inputs.

Generalizing the single-gate case, the coefficients of $Q(t)$ are given by various combinations of $x_1, x_2, z_1, z_2, z_3, z_4, y$. Variables corresponding to inputs x_1, x_2 and output y are hard-coded, ensuring that the polynomial expresses a computation based on the correct inputs and the purported output.

Now, the verifier wants a probabilistic and efficient check that $Q(t)$ is 0 everywhere (see sidebar, page 77). A key fact is: if a polynomial is *not* the zero polynomial, it has few roots (consider a parabola: it crosses the horizontal axis a maximum of two times). For example, if we take $x_1 = 0, x_2 = 0, y = 1$, which is an *incorrect* execution of the above circuit, then the corresponding polynomial might look like this:

and a polynomial corresponding to a *correct* execution is simply a horizontal line on the axis.

The check, then, is the following. The verifier chooses a random value for t (call it τ) from a pre-existing range (for example, integers between 0 and M , for some M), and evaluates Q at τ . The verifier accepts the computation as correct if $Q(\tau) = 0$ and rejects otherwise. This process occasionally produces errors since even a non-zero polynomial Q is zero sometimes (the idea here is a variant of “a stopped clock is right twice per day”), but this event happens rarely and is independent of the prover's actions.

But how does the verifier actually evaluate $Q(\tau)$? Recall that our setup, for now, is that the prover sends a (possibly long) *encoded transcript* to the verifier. The sidebar on page 81 explains what is in the encoded transcript, and how it allows the verifier to evaluate $Q(\tau)$.

x , and the prover executes the circuit C on input x and claims the output is y .^c In performing this step, the prover is expected to obtain a *valid transcript* for $\{C, x, y\}$ (Figure 1, step 2). A *transcript* is an assignment of values to the circuit wires; in a valid transcript for $\{C, x, y\}$, the values assigned to the input wires are those of x , the intermediate values correspond to the correct operation of each gate in C , and the values assigned to the output wires are y . Notice that if the claimed output is incorrect—that is, if $y \neq p(x)$ —then a valid transcript for $\{C, x, y\}$ simply does not exist.

Therefore, if the prover could establish a valid transcript exists for $\{C, x, y\}$, this would convince the verifier of the correctness of the execution. Of course, there is a simple proof that a valid transcript exists: the transcript itself. However, the verifier can check the transcript only by examining all of it, which would be as much work as having executed p in the first place.

Instead, the prover will *encode* the transcript (Figure 1, step 3) into a longer string. The encoding lets the verifier detect a transcript's validity by inspecting a small number of randomly chosen locations in the encoded string and then applying efficient tests to the contents found therein. The machinery of PCPs, for example, allows exactly this (see the three accompanying sidebars).

However, we still have a problem. The verifier cannot get its hands on the entire encoded transcript; it is longer—astronomically longer, in some cases—than the plain transcript, so reading in the whole thing would again require too much work from the verifier. Furthermore, we do not want the prover to have to write out the whole encoded transcript: that would also be too much work, much of it wasteful, since the verifier looks at only small pieces of the encoding. And unfortunately, we cannot have the verifier simply ask the prover

^c The framework also handles circuits where the prover supplies some of the inputs and receives some of the outputs (enabling computations over remote state inaccessible to the verifier). However, the accompanying techniques are mostly beyond our scope (we will briefly mention them later). For simplicity we are treating p as a pure computation.

what the encoding holds at particular locations, as the protocols depend on the element of surprise. That is, if the verifier's queries are known in advance, then the prover can arrange its answers to fool the verifier.

As a result, the verifier must issue its queries about the encoding carefully (Figure 1, step 4). The literature describes three separate techniques for this purpose. They draw on a richly varied set of tools from complexity theory and cryptography, and are summarized next. Afterward, we discuss their relative merits.

► *Use the power of interaction.* One set of protocols proceeds in rounds: the verifier queries the prover about the contents of the encoding at a particular location, the prover responds, the verifier makes another query, the prover responds, and so on. Just as a lawyer's questions to a witness restrict the answers the witness can give to the next question, until a lying witness is caught in a contradiction, the prover's answers in each round about what the encoding holds limit the space of valid answers in the next round. This continues until the last round, at which point a prover that has answered perfidiously at any point—by answering based on an invalid transcript or by giving answers that are untethered to any transcript—simply has no valid answers. This approach relies on interactive proof protocols,^{3,26,32} most notably Muggles,²⁵ which was refined and implemented.^{18,45–47}

► *Extract a commitment.* These protocols proceed in two rounds. The verifier first requires the prover to *commit* to the full contents of the encoded transcript; the commitment relies on standard cryptographic primitives, and we call the committed-to contents a *proof*. In the second round, the verifier generates queries—locations in the proof the verifier is interested in—and then asks the prover what values the proof contains at those locations; the prover is forced to respond consistent with the commitment. To generate queries and validate answers, the verifier uses PCPs (they enable probabilistic checking, as described in the sidebar entitled “Probabilistically Checkable Proofs”). This approach

**This is an exciting
time for work
on verifiable
computation.**

was outlined in theory by Kilian,³⁰ building on the PCP theorem.^{1,2} Later, Ishai et al.²⁹ (IKO) gave a drastic simplification, in which the prover can commit to a proof without materializing the whole thing. IKO led to a series of refinements, and implementation in a system.^{40–43,47}

► *Hide the queries.* Instead of extracting a commitment and then revealing its queries, the verifier pre-encrypts its queries—as with the prior technique, the queries describe locations where the verifier wants to inspect an eventual proof, and these locations are chosen by PCP machinery—and sends this description to the prover prior to their interaction. Then, during the verification phase, powerful cryptography achieves the following: the prover answers the queries without being able to tell which locations in the proof are being queried, and the verifier recovers the prover's answers. The verifier then uses PCP machinery to check the answers, as in the commitment-based protocols. The approach was described in theory by Gennaro et al.²² (see also Bitansky et al.¹²), and refined and implemented in two projects.^{7,9,36}

Progress: Implemented Systems

The three techniques described are elegant and powerful, but as with the prior technique, naive implementations would result in preposterous costs. The research projects that implemented these techniques have applied theoretical innovations and serious systems work to achieve *near* practical performance. Here, we explain the structure of the design space, survey the various efforts, and explore their performance (in doing this, we will illustrate what “near practical” means).

We restrict our attention to *implemented systems* with published experimental results. By “system,” we mean code (preferably publically released) that takes some kind of representation of a computation and produces executables for the verifier and the prover that run on stock hardware. Ideally, this code is a compiler toolchain, and the representation is a program in a high-level language.

The landscape. As depicted in Figure 2, we organize the design space in

terms of a three-way trade-off among cost, expressiveness, and functionality.^d Here, cost mainly refers to *setup costs* for the verifier; as we will see, this cost is the verifier's largest expense, and affects whether a system meets the goal of saving the verifier work. (This setup cost also correlates with the prover's cost for most of the systems discussed.) By expressiveness, we mean the class of computations the system can handle while providing a benefit to the verifier. By functionality, we mean whether the works provide properties like non-interactivity (setup costs amortize indefinitely), *zero knowledge*^{24,26} (the computation transcript is hidden from the verifier, giving the prover some privacy), and public verifiability (anyone, not just a particular verifier, can check a proof, provided that the party who generated the queries is trusted).

CMT, Allspice, and Thaler. One line of work uses “the power of interaction;” it starts from Muggles,²⁵ the interactive proof protocol mentioned earlier. CMT^{18,46} exploits an algebraic insight to save orders of magnitude for the prover, versus a naive implementation of Muggles.

For circuits to which CMT applies, performance is very good, in part because Muggles and CMT do not use cryptographic operations. In fact, refinements by Thaler⁴⁵ provide a prover

that is optimal for certain classes of computations: the costs are only a constant factor (10–100×, depending on choice of baseline) over executing the computation locally. Moreover, CMT applies in (and was originally designed for) a streaming model, in which the verifier processes and discards input as it comes in.

However, CMT's expressiveness is limited. First, it imposes requirements on the circuit's geometry: the circuit must have structurally similar parallel blocks. Of course, not all computations can be expressed in that form. Second, the computation cannot use order comparisons (less-than, and so on).

Allspice⁴⁷ has CMT's low costs but achieves greater expressiveness (under the amortization model described next).

Pepper, Ginger, and Zaatar. Another line of work builds on the “extract a commitment” technique (called an “efficient argument” in the theory literature.^{13,30}). Pepper⁴² and Ginger⁴³ refine the protocol of IKO. To begin with, they represent computations as arithmetic constraints (that is, a set of equations over a finite field); a solution to the constraints corresponds to a valid transcript of the computation. This representation is often far more concise than Boolean circuits (used by IKO and in the proof of the PCP theorem¹) or arithmetic circuits (used by CMT). Pepper and Ginger also strengthen IKO's commitment primitive, explore low-overhead PCP encodings for certain computations, and apply a number of systems techniques (such as parallelization on

distributed systems).

Pepper and Ginger dramatically reduce costs for the verifier and prover, compared to IKO. However, as in IKO, the verifier incurs setup costs. Both systems address this issue via amortization, reusing the setup work over a *batch*: multiple *instances* of the same computation, on different inputs, verified simultaneously.

Pepper requires describing constraints manually. Ginger has a compiler that targets a larger class of computations; also, the constraints can have auxiliary variables set by the prover, allowing for efficient representation of not-equal-to checks and order comparisons. Still, both handle only straight line computations with repeated structure, and both require special-purpose PCP encodings.

Zaatar⁴¹ composes the commitment protocol of Pepper and Ginger with a new linear PCP;^{1,29} this PCP adapts an ingenious algebraic encoding of computations from GGPR²² (which we return to shortly). The PCP applies to all pure computations; as a result, Zaatar achieves Ginger's performance but with far greater generality.

Pinocchio. Pinocchio³⁶ instantiates the technique of hiding the queries. Pinocchio is an implementation of GGPR, which is a *noninteractive* argument. GGPR can be viewed as a probabilistically checkable encoding of computations that is akin to a PCP (this is the piece that Zaatar adapts) plus a layer of sophisticated cryptography.^{12,41} GGPR's encoding is substantially more concise than prior approaches, yielding major reductions in overhead.

The cryptography also provides many benefits. It hides the queries, which allows them to be reused. The result is a protocol with minimal interaction (after a per-computation setup phase, the verifier sends only an instance's input to the prover) and, thus, qualitatively better amortization behavior. Specifically, Pinocchio amortizes per-computation setup costs over all future instances of a given computation; by contrast, recall that Zaatar and Allspice amortize their per-computation costs only over a batch. GGPR's and Pinocchio's cryptography also yield zero knowledge and public verifiability.

Compared to Zaatar, Pinocchio brings some additional expense in the

d For space, we omit recent work that is optimized for specific classes of computations; these works can be found with this article in the ACM Digital Library under Supplemental Material.

Figure 2. Design space of implemented systems for proof-based verifiable computation.

There is a three-way trade-off among cost, expressiveness, and functionality. Higher in the figure means lower cost, and rightward generally means better expressiveness. The shaded systems achieve non-interactivity, zero knowledge, and other cryptographic properties. (Pantry, Buffet, and TinyRAM achieve these properties by leveraging Pinocchio.) Here, “regular” means structurally similar parallel blocks; “straight line” means not many conditional statements; and “pure” means computations without side effects.

prover's costs and the verifier's setup costs. Pinocchio's compiler initiated the use of C syntax in this area, and includes some program constructs not present in prior work. The underlying computational model (unrolled executions) is essentially the same as Ginger's and Zaatar's.^{41,43}

Although the systems we have described so far have made tremendous progress, they have done so within a programming model that is not reflective of real-world computations. First, these systems require loop bounds to be known at compile time. Second, they do not support indirect memory references scalably and efficiently, ruling out RAM and thus general-purpose programming. Third, the verifier must handle all inputs and outputs, a limitation that is at odds with common uses of the cloud. For example, it is unreasonable to insist that the verifier materialize the entire (massive) input to a MapReduce job. The projects described next address these issues.

TinyRAM (BCGTV and BCTV). BCGTV⁷ compiles programs in C (not just a subset) to an innovative circuit representation.⁶ Applying prior insights,^{12,22,41} BCGTV combines this circuit with proof machinery (including transcript encoding and queries) from Pinocchio and GGPR.

BCGTV's circuit representation consists of the unrolled execution of a general-purpose MIPS-like CPU, called TinyRAM (and for convenience we sometimes use "TinyRAM" to refer to BCGTV and its successors). The circuit-as-unrolled-processor provides a natural representation for language features like data-dependent looping, control flow, and self-modifying code. BCGTV's circuit also includes a permutation network that elegantly and efficiently enforces the correctness of RAM operations.

BCTV⁹ improves on BCGTV by retrieving program instructions from RAM (instead of hard-coding them in the circuit). As a result, all executions with the same number of steps use the same circuits, yielding the best amortization behavior in the literature: setup costs amortize over all computations of a given length. BCTV also includes an optimized implementation of Pinocchio's protocol that cuts costs by constant factors.

Despite these advantages, the gen-

Probabilistically Checkable Proofs (simplified)

This sidebar answers the following question: how does the prover encode its transcript, and how does the verifier use this encoded transcript to evaluate Q at a randomly chosen point? (The encoded transcript is known as a probabilistically checkable proof, or PCP. For the purposes of this sidebar, we assume that the prover sends the PCP to the verifier; in the main text, we will ultimately avoid this transmission, using commitment and other techniques.)

A naive solution is a protocol in which: the prover claims that it is sending $\{Q(0), \dots, Q(M)\}$ to the verifier, the verifier chooses one of these values at random, and the verifier checks whether the randomly chosen value is 0. However, this protocol does not work: even if there is no valid transcript, the prover could cause the verifier's "check" to always pass, by sending a string of zeroes.

Instead, the prover will encode the transcript, z , and the verifier will impose structure on this encoding; in this way, both parties together form the required polynomial Q . This process is detailed in the rest of this sidebar, which will be somewhat more technical than the previous sidebars. Nevertheless, we will be simplifying heavily; readers who want the full picture are encouraged to consult the tutorials referenced in the supplemental material (online).

As a warmup, observe that we can rewrite the polynomial Q by regarding the " z " variables as unknowns. For example, the polynomial $Q(t)$ in the sidebar on page 77 can be written as:

$$Q(t, z_1, z_2, z_3) = (-2t^2) \cdot z_1 \cdot z_2 \cdot z_3 + (t^2 - t) \cdot z_1 \cdot z_2 + (t^2 + t) \cdot z_3.$$

An important fact is that for any circuit, the polynomial Q that encodes its execution can be represented as a linear combination of the components of the transcript and pairwise products of components of the transcript. We will now state this fact using notation. Assume that there are n circuit wires, labeled $(z_1, \dots, z_n)$, and arranged as a vector $\vec{z}$. Further, let $\langle \vec{a}, \vec{b} \rangle$ denote the dot product between two vectors, and let $\vec{a} \otimes \vec{b}$ denote a vector whose components are all pairs $a_i b_j$. Then we can write Q as

$$Q(t, \vec{z}) = g_0(t) + \langle \vec{g}_1(t), \vec{z} \rangle + \langle \vec{g}_2(t), \vec{z} \otimes \vec{z} \rangle,$$

where g_0 is a function from t to scalars, and $\vec{g}_1$ and $\vec{g}_2$ are functions from t to vectors.

Now, what if the verifier had a table that contained $\langle \vec{g}_1, \vec{z} \rangle$ for all vectors $\vec{z}$ (in a finite vector space), and likewise a table that contained $\langle \vec{g}_2, \vec{z} \otimes \vec{z} \rangle$ for all vectors $\vec{z}$? Then, the verifier could evaluate $Q(\tau)$ by inspecting the tables in only one location each. Specifically, the verifier would randomly choose τ , then compute $g_0(\tau)$, $\vec{g}_1(\tau)$, and $\vec{g}_2(\tau)$; then use the two tables to look up $\langle \vec{g}_1(\tau), \vec{z} \rangle$ and $\langle \vec{g}_2(\tau), \vec{z} \otimes \vec{z} \rangle$; and add these values to $g_0(\tau)$. If the tables were produced correctly, this final sum (of scalars) will yield $Q(\tau, \vec{z})$.

However, a few issues remain. The verifier cannot know that it actually received tables of the correct form, or that the tables are consistent with each other. So the verifier performs additional spot checks; the rough idea is that if the tables deviate too heavily from the correct form, then the spot checks will pick up the divergence with high probability (and if the tables deviate from the correct form but only mildly, the verifier still recovers $Q(\tau)$).

At this point, we have answered the question at the beginning of this sidebar: the correct encoding of a valid transcript z is the two tables of values. In other words, these two tables form the probabilistically checkable proof, or PCP.

Notice that the two tables are exponentially larger than the transcript. Therefore, the prover cannot send them to the verifier or even materialize them. The purpose of the three techniques discussed in the text—interactivity, commitment, hide the queries—is, roughly speaking, to allow the verifier to query the prover about the tables without either party having to materialize or handle them.

eral approach brings a steep price: BCTV's circuits are orders of magnitude larger than Pinocchio's and Zaatar's for the same high-level programs. As a result, the verifier's setup work and the prover's costs are orders of magnitude higher, and BCTV is restricted to very short executions. Nevertheless, BCGTV and BCTV intro-

duce important tools.

Pantry and Buffet. Pantry¹⁴ extends the computational model of Zaatar and Pinocchio, and works with both systems. Pantry provides a general-purpose approach to state, yielding a RAM abstraction, verifiable Map Reduce, verifiable queries on remote databases, and—using Pinocchio's

zero knowledge variant—computations that keep the prover’s state private. To date, Pantry is the only system to extensively use the capability of argument protocols (the “extract a commitment” and “hide the queries” techniques) to handle computations for which the verifier does not have all the input. In Pantry’s approach—which instantiates folklore techniques—the verifier’s explicit input includes a digest of the full input or state, and the prover is obliged to work over state that matches this digest.

Under Pantry, every operation against state compiles into the evaluation of a cryptographic hash function. As a result, a memory access is tens of thousands of times more costly than a basic arithmetic operation.

Buffet⁴⁸ combines the best features of Pantry and TinyRAM. It slashes the cost of memory operations in Pantry by adapting TinyRAM’s RAM abstraction. Buffet also brings data-dependent looping and control flow to Pantry (without TinyRAM’s expense), using a loop flattening technique inspired by the compilers literature. As a result, Buffet supports an expansive subset of C (disallowing only function pointers and goto) at costs orders of magnitude lower than both Pantry and TinyRAM. As of this writing, Buffet appears to achieve the best mix of performance and generality in the literature.

A brief look at performance. We will

answer three questions:

1. *How do the verifier’s variable (per-instance) costs compare to the baseline of local, native execution?* For some computations, this baseline is an alternative to verifiable outsourcing.

2. *What are the verifier’s setup costs, and how do they amortize?* In many of the systems, setup costs are significant and are paid for only over multiple instances of the same circuit.

3. *What is the prover’s overhead?*

We focus only on CPU costs. On the one hand, this focus is conservative: verifiable outsourcing is motivated by more than CPU savings for the verifier. For example, if inputs are large or inaccessible, verifiable outsourcing saves network costs (the naive alternative is to download the inputs and locally execute); in this case, the CPU cost of local execution is irrelevant. On the other hand, CPU costs provide a good sense of the overall expense of the protocols. (For evaluations that take additional resources into account, see Braun et al.¹⁴)

The data we present is from re-implementations of the various systems by members of our research group, and essentially match the published results. All experiments are run on the same hardware (Intel Xeon E5-2680 processor, 2.7Ghz, 32GB RAM), with the prover on one machine and the verifier on another. We perform three runs per experiment; experimental variation is minor, so we just report

the average. Our benchmarks are 128×128 matrix multiplication (of 64-bit quantities, with full precision arithmetic) and PAM clustering of 20 vectors, each of dimension 128. We do not include data for Pantry and Buffet since their innovations do not apply to these benchmarks (their performance would be the same as Zaatar or Pinocchio, depending on which machinery they were configured to run with). For TinyRAM, we report extrapolated results since, as noted earlier, TinyRAM on current hardware is restricted to executions much smaller than these benchmarks.

Figure 3 depicts per-instance verification costs, for matrix multiplication, compared to two baselines. The first is a native execution of the standard algorithm, implemented with floating-point operations; it costs 3.5ms, and beats all of the systems at the given input size.^e (At larger input sizes, the verifier would do better than native execution: the verifier’s costs grow linearly in the input size, which is only $O(m^2)$; local execution is $O(m^3)$.) The second is an implementation of the algorithm using a multi-precision library; this baseline models a situation in which complete precision is required.

We evaluate setup costs by asking about the *cross-over point*: how many instances of a computation are required to amortize the setup cost in the sense the verifier spends fewer CPU cycles on outsourcing versus executing locally? Figure 4 plots total cost lines and cross-over points versus the second baseline.

To evaluate prover overhead, Figure 5 normalizes the prover’s cost to the floating-point baseline.

Summary and discussion. Performance differences among the systems are overshadowed by the general nature of costs in this area. The verifier is practical if its computation is amenable to one of the less expensive (but more restricted) protocols, or if there are a large number of instances that will be run (on different inputs). And when state is remote, the verifier does not need to be faster than local computation because it would be difficult—or impossible, if the remote state is private—for the verifier to perform the computation itself (such

Figure 3. Per-instance verification costs applied to 128×128 matrix multiplication of 64-bit numbers.

^e Systems that report verification costs beating local execution choose very expensive baselines for local computation.^{36,41–43}

applications are evaluated elsewhere¹⁴).

The prover, of course, has terrible overhead: several orders of magnitude (though as noted previously, this still represents tremendous progress versus the prior costs). The prover's practicality thus depends on your ability to construct appropriate scenarios. Maybe you're sending Will Smith and Jeff Goldblum into space to save Earth; then you care a lot more about correctness than costs (a calculus that applies to ordinary satellites, too). More prosaically, there is a scenario with an abundance of server CPU cycles, many instances of the same computation to verify, and remotely stored inputs: data-parallel cloud computing. Verifiable MapReduce¹⁴ is therefore an encouraging application.

Open Questions and Next Steps

The main issue in this area is performance, and the biggest problem is the prover's overhead. The verifier's per-instance costs are also too high. And the verifier's setup costs would ideally be controlled while retaining expressivity. (This is theoretically possible,^{8,11} but overhead is very high: in a recent implementation,⁸ the prover's computational costs are orders of magnitude higher than in TinyRAM.)

The computational model is a critical area of focus. Can we identify or develop programming languages that are expressive yet compile efficiently to the circuit or constraint formalism? More generally, can we move beyond this intrinsically costly formalism?

There are also questions in systems. For example, can we develop a realistic database application, including concurrency, relational structures, and so on? More generally, an important test for this area—so far unmet—is to run experiments at realistic scale.

Another interesting area of investigation concerns privacy. By leveraging Pinocchio, Pantry has experimented with simple applications that hide the prover's state from the verifier, but there is more work to be done here and other notions of privacy that are worth providing. For example, we can provide verifiability while concealing the program that is executed (by composing techniques from Pantry, Pinocchio, and TinyRAM). A speculative application is to produce versions of

Bitcoin in which transactions can be conducted anonymously, in contrast to the status quo.^{5,19}

Reflections and Predictions

It is worth recalling that the intellectual foundations of this research area really had nothing to do with practice. For example, the PCP theorem is a landmark achievement of complexity theory, but if we were to implement the theory as pro-

posed, generating the proofs, even for simple computations, would have taken longer than the age of the universe. In contrast, the projects described in this article have not only built systems from this theory but also performed experimental evaluations that terminate before publication deadlines.

So that is the encouraging news. The sobering news, of course, is these systems are basically toys. Part of the rea-

Figure 4. Total costs and cross-over points (extrapolated), for 128×128 matrix multiplication.

The slope of each line is the per-instance cost (depicted in Figure 3); the y-intercepts are the setup costs and equal 0 for local, CMT, and Thaler. The cross-over point is the x-axis point at which a system's total cost line crosses its "local" line. The cross-over points for Zaatar and Pinocchio are in the thousands; the special-purpose approaches do far better but do not apply to all computations. Pinocchio's cross-over point could be improved by constant factors, using TinyRAM's optimized implementation.⁹ Although it has the worst cross-over point, TinyRAM has the best amortization regime, followed by Pinocchio and Zaatar (see text).

Figure 5. Prover overhead normalized to native execution cost for two computations. Prover overheads are generally enormous.

son we are willing to label them near-practical is painful experience with what the theory *used* to cost. (As a rough analogy, imagine a graduate student's delight in discovering hexadecimal machine code after years spent programming one-tape Turing machines.)

Still, these systems are arguably useful in some scenarios. In high-assurance contexts, we might be willing to pay a lot to know that a remotely deployed machine is executing correctly. In the streaming context, the verifier may not have space to compute locally, so we could use CMT¹⁸ to check the outputs are correct, in concert with Thaler's refinements⁴⁵ to make the prover truly inexpensive. Finally, data-parallel cloud computations (like MapReduce jobs) perfectly match the regimes in which the general-purpose schemes perform well: abundant CPU cycles for the prover and many instances of the same computation with different inputs.

Moreover, the gap separating the performance of the current research prototypes and plausible deployment in the cloud is a few orders of magnitude—which is certainly daunting, but, given the current pace of improvement, might be bridged in a few years.

More speculatively, if the machinery becomes truly low overhead, the effects will go far beyond verifying cloud computations: we will have new ways of building systems. In any situation in which one module performs a task for another, the delegating module will be able to check the answers. This could apply at the micro level (if the CPU could check the results of the GPU, this would expose hardware errors) and the macro level (distributed systems could be built under very different trust assumptions).

But even if none of this comes to pass, there are exciting intellectual currents here. Across computer systems, we are starting to see a new style of work: reducing sophisticated cryptography and other achievements of theoretical computer science to practice.^{28,35,38,49}

These developments are likely a product of our times: the preoccupation with strong security of various kinds, and the computers powerful enough to run previously “paper-only” algorithms. Whatever the cause, proof-based verifiable computation is an excellent example of this tendency: not only does it compose theoretical refinements with systems

techniques, it also raises research questions in other sub-disciplines of computer science. This cross-pollination is the best news of all.

Acknowledgments. We thank Sri-nath Setty, Justin Thaler, Riad Wahby, Alexis Gallagher, the anonymous *Communications* reviewers, Boaz Barak, William Blumberg, Oded Goldreich, Yuval Ishai and Guy Rothblum. C

References

- Arora, S., Lund, C., Motwani, R., Sudan, M. and Szegedy, M. Proof verification and the hardness of approximation problems. *JACM* 45, 3 (May 1998), 501–555. (Prelim. version FOCS 1992).
- Arora, S. and Safra, S. Probabilistic checking of proofs: A new characterization of NP. *JACM* 45, 1 (Jan. 1998), 70–122. (Prelim. version FOCS 1992).
- Babai, L. Trading group theory for randomness. In *Proceedings of STOC*, 1985.
- Babai, L., Fortnow, L., Levin, A. and Szegedy, M. Checking computations in polylogarithmic time. In *Proceedings of STOC*, 1991.
- Ben-Sasson, E. et al. Decentralized anonymous payments from Bitcoin. *IEEE Symposium on Security and Privacy*, 2014.
- Ben-Sasson, E., Chiesa, A., Genkin, D. and Tromer, E. Fast reductions from RAMs to delegatable succinct constraint satisfaction problems. In *Proceedings of ITCS*, Jan. 2013.
- Ben-Sasson, E., Chiesa, A., Genkin, D. and Tromer, E. SNARKs for C: Verifying program executions succinctly and in zero knowledge. In *Proceedings of CRYPTO*, Aug. 2013.
- Ben-Sasson, E., Chiesa, A., Tromer, E. and Virza, M. Scalable zero knowledge via cycles of elliptic curves. In *Proceedings of CRYPTO*, Aug. 2014.
- Ben-Sasson, E., Chiesa, A., Tromer, E. and Virza, M. Succinct non-interactive zero knowledge for a von Neumann architecture. *USENIX Security*, (Aug. 2014).
- Ben-Sasson, E., Goldreich, O., Harsha, P., Sudan, M. and S. Vadhan, S. Robust PCPs of proximity, shorter PCPs and applications to coding. *SIAM J. on Comp.* 36, 4 (Dec. 2006), 889–974.
- Bitansky, N., Canetti, R., Chiesa, A. and Tromer, E. Recursive composition and bootstrapping for SNARKs and proof-carrying data. In *Proceedings of STOC*, June 2013.
- Bitansky, N., Chiesa, A., Ishai, Y., Ostrovsky, R. and Paneth, O. Succinct non-interactive arguments via linear interactive proofs. In *Proceedings of IACR TCC*, Mar. 2013.
- Brassard, G., Chaum, D. and Crépeau, C. Minimum disclosure proofs of knowledge. *J. Comp. and Sys. Sciences* 37, 2 (Oct. 1988), 156–189.
- Braun, B., Feldman, A.J., Ren, Z., Setty, S., Blumberg, A.J., and Walfish, M. Verifying computations with state. In *Proceedings of SOSP*, Nov. 2013.
- Canetti, R., Riva, B. and Rothblum, G. Practical delegation of computation using multiple servers. *ACM CCS*, 2011.
- Castro, M. and Liskov, B. Practical Byzantine fault tolerance and proactive recovery. *ACM Trans. on Comp. Sys.* 20, 4 (Nov. 2002), 398–461.
- Chung, K.M., Kalai, Y. and Vadhan, S. Improved delegation of computation using fully homomorphic encryption. In *Proceedings of CRYPTO*, 2010.
- Cormode, G., Mitzenmacher, M. and Thaler, J. Practical verifiable computation with streaming interactive proofs. In *Proceedings of ITCS*, 2012.
- Danezis, G., Fournet, C., Kohlweiss, M. and Parno, B. Pinocchio coin: Building zerocoin from a succinct pairing-based proof system. In *Proceedings of the Workshop on Language Support for Privacy-enhancing Technologies*, Nov. 2013.
- Dinur, I. The PCP theorem by gap amplification. *JACM* 54, 3 (June 2007), 2:1–12:44.
- Gennaro, R., Gentry, C. and Parno, B. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Proceedings of CRYPTO*, 2010.
- Gennaro, R., Gentry, C. and Parno, B. and Raykova, M. Quadratic span programs and succinct NIZKs without PCPs. In *Proceedings of EUROCRYPT*, May 2013.
- Gentry, C. A fully homomorphic encryption scheme. PhD thesis, Stanford University, 2009.
- Goldreich, O. Probabilistic proof systems—A primer. *Foundations and Trends in Theoretical Computer Science* 3, 1 (2007), 1–91.
- Goldwasser, S., Kalai, Y.T. and Rothblum, G.N. Delegating computation: Interactive proofs for muggles. In *Proceedings of STOC*, May 2008.
- Goldwasser, S., Micali, S. and Rackoff, C. The knowledge complexity of interactive proof systems. *SIAM J. on Comp.* 18, 1 (1989), 186–208.
- Håstad, J. Some optimal inapproximability results. *JACM* 48, 4 (July 2001), 798–859. (Prelim. version STOC 1997).
- Huang, Y., Evans, D., Katz, J. and Malka, L. Faster secure two-party computation using garbled circuits. In *USENIX Security*, 2011.
- Ishai, Y., Kushilevitz, E., and Ostrovsky, R. Efficient arguments without short PCPs. In *Proceedings of the Conference on Computational Complexity (CCC)*, 2007.
- Kilian, J. A note on efficient zero-knowledge proofs and arguments (extended abstract). In *Proceedings of STOC*, 1992.
- Kreuter, B., Shelat, A. and Shen, C.H. Billion-gate secure computation with malicious adversaries. *USENIX Security* (Aug. 2012).
- Lund, C., Fortnow, L., Karloff, H.J., and Nisan, N. Algebraic methods for interactive proof systems. *JACM* 39, 4 (1992), 859–888.
- Mahajan, P. et al. Depot: Cloud storage with minimal trust. *ACM Trans. on Comp. Sys.* 29, 4 (Dec. 2011).
- Malkhi, D. and Reiter, M. Byzantine quorum systems. *Distributed Computing* 11, 4 (Oct. 1998), 203–213. (Prelim. version *Proceedings of STOC* 1997).
- Narayan, A. and Haebleren, A. DJoin: Differentially private join queries over distributed databases. In *Proceedings of OSDI*, 2012.
- Parno, B., Gentry, C., Howell, J. and Raykova, M. Pinocchio: Nearly practical verifiable computation. *IEEE Symposium on Security and Privacy*, (May 2013).
- Parno, B., McCune, J.M. and Perrig, A. *Bootstrapping Trust in Modern Computers*. Springer, 2011.
- Popa, R.A., Redfield, C.M.S., Zeldovich, N. and Balakrishnan, H. CryptDB: Protecting confidentiality with encrypted query processing. In *Proceedings of SOSP*, 2011.
- Sadeghi, A.R., Schneider, T., and Winandy, M. Token-based cloud computing: Secure outsourcing of data and arbitrary computations with lower latency. In *Proceedings of TRUST*, 2010.
- Setty, S., Blumberg, A.J. and Walfish, M. Toward practical and unconditional verification of remote computations. In *Proceedings of HotOS*, May 2011.
- Setty, S., Braun, B., Vu, V., Blumberg, A.J., Parno, B. and Walfish, M. Resolving the conflict between generality and plausibility in verified computation. In *Proceedings of EuroSys*, Apr. 2013.
- Setty, S., McPherson, R., Blumberg, A.J., and Walfish, M. Making argument systems for outsourced computation practical (sometimes). In *Proceedings of NDSS*, 2012.
- Setty, S., Vu, V., Panpalia, N., Braun, B., Blumberg, A.J. and Walfish, M. Taking proof-based verified computation a few steps closer to practicality. *USENIX Security*, Aug. 2012.
- Sudan, M. Probabilistically checkable proofs. *Commun. ACM* 52, 3 (Mar. 2009), 76–84.
- Thaler, J. Time-optimal interactive proofs for circuit evaluation. In *Proceedings of CRYPTO*, Aug. 2013.
- Thaler, J., Roberts, M., Mitzenmacher, M. and Pfister, H. Verifiable computation with massively parallel interactive proofs. In *Proceedings of USENIX HotCloud Workshop*, (June 2012).
- Vu, V., Setty, S., Blumberg, A.J. and Walfish, M. A hybrid architecture for interactive verifiable computation. *IEEE Symposium on Security and Privacy*, (May 2013).
- Wahby, R.S., Setty, S., Ren, Z., Blumberg, A.J. and Walfish, M. Efficient RAM and control flow in verifiable outsourced computation. In *Proceedings of NDSS*, Feb. 2015.
- Wolinsky, D.I., Corrigan-Gibbs, H., Ford, B. and Johnson, A. Dissent in numbers: Making strong anonymity scale. In *Proceedings of OSDI*, 2012.

Michael Walfish (mwalfish@cs.nyu.edu) is an associate professor in the computer science department at New York University, New York City.

Andrew J. Blumberg (blumberg@math.utexas.edu) is an associate professor of mathematics at the University of Texas at Austin.

Copyright held by authors.