

CURRENT EVENTS BULLETIN

Friday, January 18, 2019, 1:00 PM to 4:45 PM

Room 307, Baltimore Convention Center

Joint Mathematics Meeting, Baltimore, MD

1:00 PM | **Bhargav Bhatt**
University of Michigan

Perfectoid Geometry and its Applications


Hear about a great new tool for number theory that allows you to treat a prime number like a variable.



2:00 PM | **Thomas Vidick**
California Institute of Technology

Verifying Quantum Computations at Scale: a Cryptographic Leash on Quantum Devices

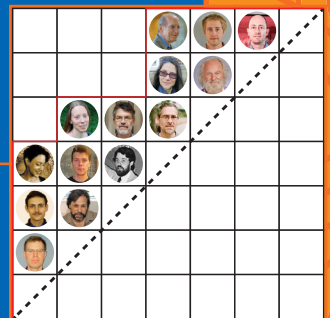
Quantum computers would be amazing—but could you trust one?



3:00 PM | **Stephanie van Willigenburg**
University of British Columbia

The Shuffle Conjecture


Dyck paths, symmetric functions—what's it all about?



4:00 PM | **Robert Lazarsfeld**
Stony Brook University

Tangent Developable Surfaces and the Equations Defining Algebraic Curves

Some lovely old geometry and an algebraic idea inspired by topology have simplified the proof of a central conjecture about generic Riemann surfaces.



Introduction to the Current Events Bulletin

Will the Riemann Hypothesis be proved this week? What is the Geometric Langlands Conjecture about? How could you best exploit a stream of data flowing by too fast to capture? I think we mathematicians are provoked to ask such questions by our sense that underneath the vastness of mathematics is a fundamental unity allowing us to look into many different corners -- though we couldn't possibly work in all of them. I love the idea of having an expert explain such things to me in a brief, accessible way. And I, like most of us, love common-room gossip.

The Current Events Bulletin Session at the Joint Mathematics Meetings, begun in 2003, is an event where the speakers do not report on their own work, but survey some of the most interesting current developments in mathematics, pure and applied. The wonderful tradition of the Bourbaki Seminar is an inspiration, but we aim for more accessible treatments and a wider range of subjects. I've been the organizer of these sessions since they started, but a varying, broadly constituted advisory committee helps select the topics and speakers. Excellence in exposition is a prime consideration.

A written exposition greatly increases the number of people who can enjoy the product of the sessions, so speakers are asked to do the hard work of producing such articles. These are made into a booklet distributed at the meeting. Speakers are then invited to submit papers based on them to the *Bulletin of the AMS*, and this has led to many fine publications.

I hope you'll enjoy the papers produced from these sessions, but there's nothing like being at the talks -- don't miss them!

David Eisenbud, Organizer
Mathematical Sciences Research Institute
de@msri.org

For PDF files of talks given in prior years, see
<http://www.ams.org/ams/current-events-bulletin.html>.
The list of speakers/titles from prior years may be found at the end of this booklet.

PERFECTOID GEOMETRY AND ITS APPLICATIONS

BHARGAV BHATT

ABSTRACT. There is a strong and classical analogy, linking number theory and algebraic geometry, between the field of rational numbers and the field of rational functions in one variable. Perfectoid geometry animates this analogy by providing a context where one can treat a (fixed) prime number like a variable. The resulting notion has helped solve long-standing problems in diverse areas of mathematics—not just in number theory and algebraic geometry, but also in commutative algebra and algebraic topology. In these notes, I will explain the definition of a perfectoid ring and discuss some applications.

1. INTRODUCTION

Fix a prime number p . The field \mathbf{Q}_p of p -adic numbers was introduced by Hensel in 1897 [Hen88] as a receptacle for “ p -adic Taylor expansions” of rational numbers. It has since proven useful across mathematics.

Recall that \mathbf{Q}_p is defined as the completion of the field \mathbf{Q} of rational numbers with respect to the p -adic norm given by $|x|_p := p^{-v_p(x)}$ for $x \in \mathbf{Q}$, where $v_p(x)$ denotes the p -adic valuation of x (i.e., the highest power of p , possibly negative, appearing when x is expressed as a product of distinct prime powers). As large powers of p are small in this norm, a typical element of \mathbf{Q}_p can be written as a power series

$$\sum_{i=n}^{\infty} a_i p^i \tag{1}$$

with $a_i \in \{0, 1, 2, \dots, p-1\}$ and $n \in \mathbf{Z}$. By analogy, note that an element of the field $\mathbf{F}_p((t))$ of Laurent series over \mathbf{F}_p can also be written as a power series

$$\sum_{i=n}^{\infty} a_i t^i \tag{2}$$

with $a_i \in \{0, 1, 2, \dots, p-1\}$ and $n \in \mathbf{Z}$.

Despite the similarity between (1) and (2), the fields \mathbf{Q}_p and $\mathbf{F}_p((t))$ are not isomorphic, even as abelian groups: the former has characteristic 0 while the latter has characteristic p . Nevertheless, comparing the shapes of (1) and (2), it seems natural to speculate about the existence of a more precise relationship between the fields \mathbf{Q}_p and $\mathbf{F}_p((t))$ that switches the roles of p and t . Besides its aesthetic beauty, any precise relationship could be used to transport structure from one side to the other. From the perspective of an algebraic geometer interested in the study of solution sets of polynomial

equations, the prospect of such a relationship is especially tantalizing: one could use the Frobenius map (a resolutely characteristic p phenomenon) to study solution sets over the characteristic 0 field \mathbf{Q}_p !

In the last century, various attempts have been made at addressing the speculation in the previous paragraph. Early examples include the work of Krasner [Kra57] on “approximating” local fields of characteristic p in terms of sequences of local fields of characteristic 0, the work of Deligne [Del84] on the notion of “close” local fields with applications to ramification theory, and the work of Kazhdan [Kaz86] exploiting close local fields to relate representation theory in characteristic 0 and characteristic p . The limiting version of these ideas finds a natural home in a result of Fontaine and Wintenberger [FW79] giving an equivalence between the Galois theory of certain “infinitely ramified” extensions K and K^\flat of \mathbf{Q}_p and $\mathbf{F}_p((t))$, respectively; this result (reviewed in §2) forms one of the basic building blocks of the subject of p -adic Hodge theory.

The main goal of these notes (covered in §3) is to discuss a recent satisfying answer to the speculation addressed above provided by Scholze’s thesis [Sch12]¹: Scholze introduces a category of geometric objects called *perfectoid spaces* over each of the fields K and K^\flat from the previous paragraph, and then proves the resulting categories are equivalent to each other! This equivalence is robust enough to permit transport of structures from one side to the other. In particular, it includes the Fontaine–Wintenberger theorem as a very special case (namely, when the perfectoid space is a point), it provides some geometric meaning to the operation of “replacing p with t ”, and it gives us a tool to study algebraic geometry over the characteristic 0 field K in terms of algebraic geometry over the characteristic p field K^\flat .

Despite its youth, perfectoid geometry has already been useful in the resolution of long standing problems in multiple mathematical disciplines. While it would be impossible to summarize all the applications in a short space, we mention at least two significant ones. In §4, we explain how it led to the resolution of a fundamental problem in commutative algebra. In §5, we discuss the improvements in our understanding of the cohomology of algebraic varieties resulting from perfectoid geometry. In both cases, we do not give any proofs, but focus instead on how one uses perfectoid spaces to prove theorems about objects of classical interest.

2. MOTIVATION: THE FONTAINE–WINTENBERGER THEOREM

In this section, we discuss (a special case of) a theorem of Fontaine and Wintenberger that gives an equivalence between the Galois theory of an extension K of \mathbf{Q}_p with an extension K^\flat of $\mathbf{F}_p((t))$. The fundamental idea here is that the heuristic analogy between \mathbf{Q}_p and $\mathbf{F}_p((t))$ alluded to in §1 can be lifted to an equivalence of Galois theories if we pass up to certain

¹Related ideas were developed contemporaneously in the work of Kedlaya and Liu, and can now be found in [KL15].

infinitely ramified extensions K and K^\flat , respectively. In fact, there are many choices of such pairs, and we stick to the following one to illustrate the idea.

Notation 2.1 (The fields K and K^\flat). Consider the (infinite algebraic) extension $\mathbf{Q}_p(p^{1/p^\infty}) := \cup_n \mathbf{Q}_p(p^{1/p^n})$ of \mathbf{Q}_p obtained by adjoining a compatible system of p -power roots of p . The defining norm on \mathbf{Q}_p passes uniquely up to any finite extension, and hence also up to $\mathbf{Q}_p(p^{1/p^\infty})$. The induced topology on $\mathbf{Q}_p(p^{1/p^\infty})$, however, is not complete as $\mathbf{Q}_p(p^{1/p^\infty})$ has infinite degree over \mathbf{Q}_p . Let $K := \mathbf{Q}_p(p^{1/p^\infty})^\wedge$ denote its completion. Similarly, we let $K^\flat := \mathbf{F}_p((t))(t^{1/p^\infty})^\wedge$ be the analogous construction in characteristic p .

Remark 2.2. Despite the apparent symmetry, the construction of K^\flat from $\mathbf{F}_p((t))$ is slightly more canonical than that of K from \mathbf{Q}_p : the extension $\mathbf{F}_p((t))(t^{1/p^\infty})$ of $\mathbf{F}_p((t))$ can be described as the smallest perfect field containing $\mathbf{F}_p((t))$ and is thus independent of the choice of the variable t , while no such description is available for the extension K of \mathbf{Q}_p .

Remark 2.3 (Relating K to K^\flat algebraically). The characteristic 0 field K and the characteristic p field K^\flat evidently have some formal similarity in their constructions. In fact, they can be related algebraically too on the basis of the following observation: the topology on K provides access to an auxiliary ring whose mod p reduction is related to the analogous construction for K^\flat . We explain this carefully next, as it is the basis of a fundamental operation on perfectoid spaces.

The p -adic metric on K endows K with the structure of a nonarchimedean field, so the unit ball

$$\mathcal{O}_K := \{x \in K \mid |x|_p \leq 1\} \subset K$$

is an open subring of K , and p is not invertible in this subring (since $|1/p|_p = p > 1$). In fact, it is not difficult to explicitly identify \mathcal{O}_K as the p -adic completion of $\mathbf{Z}[x^{1/p^\infty}]/(x - p)$, where $\mathbf{Z}[x^{1/p^\infty}] = \cup_n \mathbf{Z}[x^{1/p^n}]$ as before. Similarly, as the t -adic metric on K^\flat is also nonarchimedean, the unit ball $\mathcal{O}_{K^\flat} \subset K^\flat$ is a ring that can be described explicitly as the t -adic completion of $\mathbf{F}_p[t^{1/p^\infty}]$. The relationship between K and K^\flat is now easy to explain: sending p to t gives an isomorphism

$$\mathcal{O}_K/p \simeq \mathcal{O}_{K^\flat}/t$$

of rings via the chain

$$\mathcal{O}_K/p \simeq \mathbf{Z}[x^{1/p^\infty}]/(x - p, p) \simeq \mathbf{F}_p[x^{1/p^\infty}]/x \xrightarrow{x \mapsto t} \mathbf{F}_p[t^{1/p^\infty}]/t \simeq \mathcal{O}_{K^\flat}/t. \quad (3)$$

Fontaine took this one step further in [Fon94, §1.2] and gave a formula recovering \mathcal{O}_{K^\flat} (and thus $K^\flat \simeq \text{Frac}(\mathcal{O}_{K^\flat})$) from K : the map $\mathcal{O}_{K^\flat} \rightarrow \mathcal{O}_K/p$ constructed above lifts uniquely to an isomorphism

$$\mathcal{O}_{K^\flat} \simeq \varprojlim_{x \mapsto x^p} \mathcal{O}_K/p := \{(x_n \in \mathcal{O}_K/p)_{n \geq 0} \mid x_{n+1}^p = x_n\} \quad (4)$$

of rings. Following Scholze’s terminology from [Sch12], the ring appearing on the right of (4) is nowadays called the *tilt* of \mathcal{O}_K ; the operation of tilting is discussed further in §3.1 below.

We are now in a position to state the promised theorem:

Theorem 2.4 (Fontaine–Wintenberger [FW79]). *The category of finite extensions of K is naturally equivalent to the category of finite extensions of K^\flat . In particular, there is a canonical isomorphism $\mathrm{Gal}(\overline{K}/K) \simeq \mathrm{Gal}(\overline{K}^\flat/K^\flat)$ of the absolute Galois groups of K and K^\flat .*

We refer to the equivalence in Theorem 2.4 as the *tilting equivalence*. Let us first explain how to tilt in a simple example.

Example 2.5. Assume $p \neq 2$. Then $K(\sqrt{p})$ is a non-trivial degree 2 extension of K . The corresponding extension of K^\flat is given by $K^\flat(\sqrt{t})$; this can be checked using the characterization given in Remark 2.6 and the computation in Example 3.11. Thus, in this example, the correspondence between characteristic 0 and characteristic p objects is literally given by replacing p with t . (This need not be the case for more complicated field extension.)

Remark 2.6 (Characterizing the tilting equivalence). A pair of finite extensions L/K and L^\flat/K^\flat match up under the tilting equivalence of Theorem 2.4 exactly when we have an isomorphism $\mathcal{O}_L/p \simeq \mathcal{O}_{L^\flat}/t$ of rings sitting over the isomorphism $\mathcal{O}_K/p \simeq \mathcal{O}_{K^\flat}/t$ mentioned in Remark 2.3; here we implicitly use that a finite extension of a nonarchimedean field is also a nonarchimedean field in an essentially unique fashion.

Remark 2.7 (Variants). As mentioned before, the tilting equivalence from Theorem 2.4 is not specific to the particular choice of K and K^\flat . See Corollary 3.13 for a vast generalization.

Theorem 2.4 has many applications, especially in p -adic Hodge theory, primarily because it is often easier to work in characteristic p than in characteristic 0. We record one such one application next, together with a sketch of a proof to indicate how Theorem 2.4 allows us to use the Frobenius map on K^\flat to prove something about K .

Corollary 2.8 (Cohomological dimension estimate). *Let M be a finite dimensional \mathbf{F}_p -vector space with a continuous action of $\mathrm{Gal}(\overline{K}/K)$. Then the continuous Galois cohomology groups $H^i(\mathrm{Gal}(\overline{K}/K), M)$ vanish for $i \geq 2$.*

Proof sketch. By Theorem 2.4, it suffices to prove the analogous statement for K^\flat . A standard argument allows us to reduce to the case where $M = \mathbf{F}_p$ is the trivial representation. Using the Artin-Schreier short exact sequence

$$0 \rightarrow \mathbf{F}_p \rightarrow \overline{K}^\flat \xrightarrow{x \mapsto x^p - x} \overline{K}^\flat \rightarrow 0,$$

it suffices to show that $H^i(\mathrm{Gal}(\overline{K}^\flat/K^\flat), \overline{K}^\flat) = 0$ for $i > 0$. But this is essentially the additive form of Hilbert’s theorem 90. \square

3. PERFECTOID RINGS AND SPACES

The goal of this section is to define perfectoid rings (§3.2) and give some of their most important properties, such as the almost purity theorem (§3.3). We begin in §3.1 with the tilting functor, which attaches a characteristic p ring to any commutative ring, and is one of the key tools used to study perfectoid rings. Essentially all results discussed in this section can be found in [Sch12] or [KL15]; the only exception is Corollary 3.14, which shall appear in the forthcoming [BS].

3.1. Fontaine’s tilting functor. The tilting functor is the main tool for going from characteristic 0 to characteristic p .

Construction 3.1 (Tilting). Let R be a commutative ring. The *tilt* R^\flat of R is defined as the characteristic p ring coming from the “inverse limit perfection” of R/p , i.e.,

$$R^\flat := \varprojlim_{x \mapsto x^p} R/p := \{(x_n \in R/p)_{n \geq 0} \mid x_{n+1}^p = x_n\}.$$

One easily checks that R^\flat is a *perfect* ring of characteristic p , i.e., the Frobenius endomorphism of R^\flat is bijective. In particular, each $f \in R^\flat$ admits a unique p -th root $f^{1/p} \in R^\flat$; in terms of sequences $f := (x_n)_{n \geq 0}$ as above, the p -th root is given by simply shifting the sequence, i.e., $f^{1/p} := (x_{n+1})_{n \geq 0}$.

Note that R^\flat only depends on R/p and hence only on the p -adic completion of R . Thus, in the sequel, we typically restrict attention to p -adically complete rings.

Remark 3.2 (The \sharp -map). For any p -adically complete commutative ring R , an elementary number theory exercise shows that the natural map $R \rightarrow R/p$ induces a bijection

$$\varprojlim_{x \mapsto x^p} R \simeq \varprojlim_{x \mapsto x^p} R/p =: R^\flat$$

of multiplicative monoids. Explicitly, the inverse sends a p -power compatible sequence $(x_0, x_1, x_2, \dots) \in \varprojlim_{x \mapsto x^p} R/p$ to the sequence $(y_0, y_1, y_2, \dots) \in \varprojlim_{x \mapsto x^p} R$

given by

$$y_n = \lim_{m \rightarrow \infty} \tilde{x}_{n+m}^{p^m},$$

where $\tilde{x}_r \in R$ denotes a lift of $x_r \in R/p$. Consequently, we can view elements of R^\flat as p -power compatible sequences (y_0, y_1, y_2, \dots) of elements of R ; the multiplication rule on sequences is the obvious one (componentwise), while the addition is slightly funnier. A tangible observable of this construction is that remembering only the first term of the sequence gives a multiplicative (but typically not additive) map $\sharp : R^\flat \rightarrow R$ connecting the characteristic p ring R^\flat to the original ring R . The image of this map is exactly those elements of R that admit a compatible system of p -power roots.

Example 3.3 (The tilt of \mathbf{Z}_p). Let $R = \mathbf{Z}_p$ be the p -adic integers, i.e., the p -adic completion of \mathbf{Z} or equivalently the unit ball in the nonarchimedean

field \mathbf{Q}_p . Then $R/p \simeq \mathbf{F}_p$; as this ring has a bijective Frobenius, we have $R^b \simeq \mathbf{F}_p$ as well. In this case, the map $\sharp : \mathbf{F}_p \rightarrow R^b$ sends 0 to 0 and the rest is determined by the inverse to the following bijection: the $(p-1)$ -st roots of 1 in \mathbf{Z}_p map bijectively down to \mathbf{F}_p^* via Hensel's lemma.

More generally, for readers familiar with the Witt vectors $W(S)$ of a perfect ring S of characteristic p , we remark that $W(S)^b \simeq S$, and the resulting map $\sharp : S \rightarrow W(S)$ is the Teichmüller map; the example in the previous paragraph was simply $S = \mathbf{F}_p$.

The tilting operation turns out to be quite uninteresting if we only consider rings familiar from commutative algebra or algebraic geometry:

Example 3.4 (Tilting finite type algebras). Let $R := \mathbf{Z}[x_1, \dots, x_n]$ (or its p -adic completion). Then $R/p \simeq \mathbf{F}_p[x_1, \dots, x_n]$. The Frobenius map on R/p is injective with image $\mathbf{F}_p[x_1^p, \dots, x_n^p]$. Iterating this observation, we find that

$$R^b = \bigcap_m \mathbf{F}_p[x_1^{p^m}, \dots, x_n^{p^m}] = \mathbf{F}_p$$

is simply the set of constant polynomials.

More generally, one can show that if S is a quotient of the ring R from the previous paragraph, then $S^b \simeq \prod_{i=1}^r \mathbf{F}_p$, where r is the number of connected components of $\text{Spec}(S/p)$.

The preceding example shows that the tilting operation is lossy if the Frobenius map on R/p is not surjective. As we shall see soon, this cannot happen for perfectoid rings (by fiat).

3.2. Perfectoid rings. The main definition of these notes is the following.

Definition 3.5 (Perfectoid rings). Let R be a p -torsionfree² and p -adically complete ring. We say that R is *perfectoid* if

- (1) there exists some $\varpi \in R$ such that $\varpi^p = pu$ for a unit $u \in R$, and
- (2) the Frobenius map $R/p \rightarrow R/p$ is surjective with kernel generated by the image of any element $\varpi \in R$ as in (1).

The category of perfectoid rings is simply the full subcategory of all commutative rings spanned by perfectoid rings.

Remark 3.6 (The various meanings of “perfectoid”). The notion introduced in Definition 3.5 is borrowed from [BMS16, §3] and sometimes called *integral perfectoid* in the literature to emphasize its integral nature; in these cases, the term perfectoid is used to describe $R[1/p]$ (viewed as a topological ring with a neighbourhood basis of 0 given by $p^n R$). These notions are almost equivalent: given the topological ring $R[1/p]$, one can recover R up to some very small (“almost zero”) ambiguity. The original notion of a perfectoid algebra from [Sch12] depended on working over a fixed perfectoid field; to

²It is also possible to define perfectoid rings without requiring p -torsionfreeness, and including perfect characteristic p rings as a special case (see [BMS16, §3]). However, in the interest of simplicity, we stick to the p -torsionfree setting.

the best of our knowledge, the first field independent (or “absolute”) notion is due to Fontaine [Fon13, §1.1]

Condition (2) in Definition 3.5 rules out the phenomenon observed in Example 3.4, while condition (1) in Definition 3.5 rules out examples like \mathbf{Z}_p . In fact, nonzero perfectoid rings have to be quite large as they are always non-noetherian; one can even show that the ideal \sqrt{pR} is so large that it is its own square (see Lemma 3.9). Nevertheless, if one is willing to leave the world of noetherian rings (the familiar setting of commutative algebra and algebraic geometry), then there are many interesting examples.

Example 3.7. We record some basic examples of perfectoid rings and their tilts.

- (1) (Extracting p -th roots of p) Let $K := \mathbf{Q}_p(p^{1/p^\infty})^\wedge$ be the field studied in §2. Then the ring $\mathcal{O}_K = \mathbf{Z}[p^{1/p^\infty}]^\wedge \subset K$ from Remark 2.3 is perfectoid. Indeed, taking $\varpi = p^{1/p} \in \mathcal{O}_K$, one checks that both conditions in Definition 3.5 are satisfied using the formula in (3). The tilt $(\mathcal{O}_K)^\flat$ is then identified with the ring $\mathcal{O}_{K^\flat} \subset K^\flat$ also studied §2 via (4). The resulting \sharp -map $\mathcal{O}_{K^\flat} \rightarrow \mathcal{O}_K$ sends t^{1/p^n} to p^{1/p^n} for all n . In this case, we have $\sqrt{p\mathcal{O}_K} = \cup_n p^{1/p^n} \mathcal{O}_K$.
- (2) (The perfectoid polynomial ring) Continuing the notation of the previous example, let R be the p -adic completion of $\mathcal{O}_K[x^{1/p^\infty}]$. Then again taking $\varpi = p^{1/p} \in \mathcal{O}_K \subset R$, one checks that both conditions in Definition 3.5 are satisfied, so R is perfectoid. The tilt R^\flat in this case is identified with the analogous object over \mathcal{O}_{K^\flat} , i.e., R^\flat is the t -adic completion of $\mathcal{O}_{K^\flat}[x^{1/p^\infty}]$. The \sharp -map then sends $x^{1/p^n} \in R^\flat$ to $x^{1/p^n} \in R$. The ideal \sqrt{pR} has the same shape as in (1). There are similar examples with a larger number of variables.
- (3) (The cyclotomic extension) Let $\mathbf{Q}_p^{cyc} := \mathbf{Q}_p(\mu_{p^\infty})^\wedge$ be the completion of the algebraic extension of \mathbf{Q}_p obtained by adjoining a compatible system of p -power roots of 1. Like the field K in (1) above, the field \mathbf{Q}_p^{cyc} is a nonarchimedean field, and its unit ball $\mathbf{Z}_p^{cyc} \subset \mathbf{Q}_p^{cyc}$ is a perfectoid ring. In fact, one can describe \mathbf{Z}_p^{cyc} explicitly: it is isomorphic to the p -adic completion of

$$\mathbf{Z}[q^{1/p^\infty}]/(1 + q + q^2 + \dots + q^{p-1}),$$

with q corresponding to a primitive p -th root of 1. The element $\varpi \in R$ required in Definition 3.5 can be taken to be the image of $1 + q^{1/p} + q^{2/p} + \dots + q^{(p-1)/p}$. In fact, the p -power compatible system $\{1 + q^{1/p^n} + q^{2/p^n} + \dots + q^{(p-1)/p^n} \in \mathbf{Z}_p^{cyc}/p\}_{n \geq 0}$ gives an element $t \in \mathbf{Z}_p^{cyc, \flat}$ inducing an isomorphism of $\mathbf{Z}_p^{cyc, \flat}$ with the t -adic completion of $\mathbf{F}_p[t^{1/p^\infty}]$.

Remark 3.8 (The fibers of tilting: the Fargues-Fontaine curve). Note that the perfectoid rings from (1) and (3) of Example 3.7 are distinct, but have isomorphic tilts. Thus, the tilting operation is “many-to-one”. In fact, given

a perfectoid ring R , the collection of all “untilts” of R^b can be organized into an extremely interesting geometric structure known as the Fargues-Fontaine “curve” X_{FF} . Even though X_{FF} is not a curve in the traditional sense of algebraic geometry, perfectoid geometry has made it possible to take known results and techniques used to study curves in characteristic p and transfer them to the curve X_{FF} . This transfer has already had some stunning applications in the local Langlands program: Fargues observed [Far] that the statement of the geometric Langlands correspondence for a curve over \mathbf{F}_q can be adapted to the setting of X_{FF} , leading to a beautiful conjectural geometrization of the local Langlands correspondence for p -adic fields. Essentially simultaneously, Scholze embarked on a program in [SW, Sch17] to transplant V. Lafforgue’s breakthrough [Laf] on the global Langlands correspondence for arbitrary reductive groups over function fields to the setting of p -adic fields using X_{FF} as the curve. A more thorough expository account of X_{FF} can be found in [Mor18], and the definitive reference is [FF].

However, the tilting operation becomes “one-to-one” over a fixed base, i.e., given a perfectoid ring R , the functor $S \mapsto S^b$ gives a fully faithful embedding of the category of perfectoid R -algebras into the category of R^b -algebras.

For all three perfectoid rings R discussed in Example 3.7, the corresponding tilt R^b had a preferred element $t \in R^b$ playing roughly the same role as $p \in R$; for example, we have $R/p \simeq R^b/t$ in all examples. It turns out that a similar picture holds true for any perfectoid ring.

Lemma 3.9. *Let R be a perfectoid ring. Then there exists an element ϖ as in Definition 3.5 that admits a compatible system $\{\varpi^{1/p^n}\}_{n \geq 0}$ of roots in R . Viewing this system as an element $t^{1/p} \in R^b$ via Remark 3.2, the natural map $R^b \rightarrow R/p$ gives an isomorphism*

$$R/p \simeq R^b/t$$

of rings. We also have $\sqrt{pR} = \cup_n \varpi^{1/p^n} R$, whence $(\sqrt{pR})^2 = \sqrt{pR}$.

Thus, a perfectoid ring R is related to the perfect characteristic p ring R^b via a “correspondence”

$$R \rightarrow R/p \simeq R^b/t \leftarrow R^b \quad (5)$$

of rings, with t as in Lemma 3.9. If one only keeps track of the multiplicative structure, the \sharp -map from Remark 3.2 lifts this correspondence to a map $R^b \rightarrow R$. Through this correspondence, many remarkable properties of perfect characteristic p rings can be transferred to yield similar properties of perfectoid rings. For example, perfectoid rings are reduced and even seminormal; they tend to have finite global dimension in the sense of homological algebra; the category of perfectoid rings is closed under many operations (such as pushouts) in the category of all p -adically complete rings; and differential forms on a perfectoid ring form a p -divisible group.

3.3. The almost purity theorem. Perhaps the most important property of perfectoid rings is the almost purity theorem. This property refines deep previous work of Faltings [Fal88, Fal02, Fal99] that built on Tate’s seminal paper [Tat67]. One of the key notions introduced by Faltings was “almost mathematics”. Roughly speaking, this refers to doing commutative algebra over a perfectoid ring R whilst systematically ignoring modules killed by the (very large) ideal \sqrt{pR} ; since this ideal is its own square (Lemma 3.9), modules killed by \sqrt{pR} are closed under extensions, so ignoring them is a sensible operation. In this language, the almost purity theorem states:

Theorem 3.10 (The almost purity theorem). *Let R be a perfectoid ring. Let S be the integral closure of R in a finite étale extension of $R[1/p]$. Then S is perfectoid and $R \rightarrow S$ is almost finite étale.*

Let us briefly explain the terms appearing above. The property of being finite étale for a map of commutative rings is the algebraic translation (under the “rings = affine varieties” dictionary of algebraic geometry) of the notion of a covering space map of finite degree in topology. Thus, “almost finite étale” means roughly that the algebraic obstructions to being finite étale, such as the functor $\mathrm{Ext}_R^1(S, -)$, are annihilated by \sqrt{pR} ; see [GR03] for an in-depth development of many concepts from algebraic geometry and commutative algebra in the context of almost mathematics. Summarizing geometrically, Theorem 3.10 says that if a finite cover of $\mathrm{Spec}(R)$ only branches over the divisor $\mathrm{Spec}(R/p) \subset \mathrm{Spec}(R)$, then it is almost unbranched. (See Remark 3.12 for why this is at least heuristically reasonable.)

To give a sense of the algebraic information captured by Theorem 3.10, let us explain how to see the promised behaviour directly in Example 2.5.

Example 3.11 (Almost purity for a simple quadratic extension). Consider the perfectoid ring $\mathcal{O}_K = \mathbf{Z}[p^{1/p^\infty}]^\wedge$ from Example 3.7 (1) with fraction field K , so $\sqrt{p\mathcal{O}_K} = \cup_n p^{1/p^n} \mathcal{O}_K$. Assume $p \neq 2$, and consider the quadratic extension $L = K(\sqrt{p})$. We shall explain why \mathcal{O}_L is almost projective as an \mathcal{O}_K -module (as predicted by Theorem 3.10). In other words, we check why the functor $\mathrm{Ext}_{\mathcal{O}_K}^1(\mathcal{O}_L, -)$ is annihilated by p^{1/p^n} for all $n \geq 0$.

By chasing a long exact sequence, it is enough show the following: for each $n \geq 0$, there exists a projective \mathcal{O}_K -submodule $S_n \subset \mathcal{O}_L$ with cokernel killed by p^{1/p^n} . To construct S_n , note that $p^{1/(2p^n)} \in L$ for all $n \geq 0$ since $p^{1/p^n}, p^{1/2} \in L$. Now $|p^{1/(2p^n)}| = \sqrt{|p^{1/p^n}|} < 1$, so $p^{1/(2p^n)} \in \mathcal{O}_L$. Let $S_n := \mathcal{O}_K[p^{1/(2p^n)}] \subset \mathcal{O}_L$ be the \mathcal{O}_K -subalgebra generated by $p^{1/(2p^n)}$. We then have a presentation $S_n \simeq \mathcal{O}_K[x_n]/(f(x_n))$, where $f(x) := x^2 - p^{1/p^n}$. A classical result of Dedekind (see [HS06, Theorem 12.1.1]) implies that the discriminant $f'(x_n) = 2x_n$ kills the quotient \mathcal{O}_L/S_n . As 2 is invertible, it follows that $x_n^2 = p^{1/p^n}$ also annihilates \mathcal{O}_L/S_n , as promised.

Remark 3.12 (Why is it called a purity theorem?). The name “almost purity theorem” is appropriate in view of the analogy with the classical Zariski–Nagata purity theorem in algebraic geometry. The latter says that if $f : X \rightarrow Y$ is a finite surjective map of complex varieties with X normal

and Y smooth, then the ramification locus of f is a (possibly empty) finite union of pure codimension 1 subvarieties of Y . In particular, if we already know that the ramification locus of f has codimension ≥ 2 , then f must be unramified.

The almost purity theorem is a similar assertion for the map $\mathrm{Spec}(S) \rightarrow \mathrm{Spec}(R)$ (with notation as in Theorem 3.10). Indeed, by hypothesis, this morphism is unramified at all points of $\mathrm{Spec}(R)$ where p is invertible, so the ramification locus is contained in the complementary locus $\{p = 0\} \subset \mathrm{Spec}(R)$, which has codimension 1. But the defining function p of this locus is already “infinitely ramified” in R via the perfectoid condition (see Lemma 3.9), so it cannot ramify further in S . There is thus no ramification in codimension 1, so one might expect that f is unramified by analogy with the previous paragraph; Theorem 3.10 says that this heuristic reasoning is almost true (in the technical sense of almost mathematics).

We also refer to [Kis, Sch] for related remarks.

One reason the almost purity theorem is useful is that it permits us to transfer information from the characteristic 0 ring $R[1/p]$ to the characteristic p ring $R^\flat[1/t]$ by passage through the correspondence (5). For example, this strategy leads to the following vast generalization of Theorem 2.4.

Corollary 3.13. *Let R be a perfectoid ring, and fix $t \in R^\flat$ as in Lemma 3.9. Then the category of finite étale $R[1/p]$ -algebras is equivalent to the category of finite étale $R^\flat[1/t]$ -algebras. In particular, we have an isomorphism*

$$\pi_1(\mathrm{Spec}(R[1/p])) \simeq \pi_1(\mathrm{Spec}(R^\flat[1/t]))$$

of fundamental groups.

Sketch of proof. Let S be the integral closure of R in a finite étale $R[1/p]$ -algebra. Theorem 3.10 implies that S is perfectoid and that $R \rightarrow S$ is almost finite étale. Reducing modulo p , we get that $R^\flat/t \simeq R/p \rightarrow S^\flat/t \simeq S/p$ is also almost finite étale. Using almost mathematical analogs of standard results in algebraic geometry and the tilting correspondence, one then proves that $R^\flat \rightarrow S^\flat$ is also almost finite étale (where “almost” is now meant with respect to the ideal $\sqrt{tR^\flat} = \cup_n t^{1/p^n} R^\flat$). Inverting t shows that $R^\flat[1/t] \rightarrow S^\flat[1/t]$ is finite étale, which gives the functor in one direction. The other direction is similar (and easier, as the analog of Theorem 3.10 is easier in characteristic p). \square

For completeness, we remark that one may also extend Corollary 2.8 to arbitrary perfectoid rings.

Corollary 3.14 (Cohomological dimension estimate). *Let R be a perfectoid ring. Then the \mathbf{F}_p -cohomological dimension of $\mathrm{Spec}(R[1/p])_{\acute{e}t}$ is ≤ 1 .*

Remark 3.15 (A comment on perfectoid spaces). In our exposition above, we have avoided discussing the theory of perfectoid *spaces*, sticking to perfectoid *rings* as much as possible. While this suffices for a first glimpse, it is important to develop the theory of perfectoid spaces for most serious

applications (including the proof of the almost purity theorem, even though the latter is formulated purely algebraically!). Due to the p -adically complete nature of the rings involved, this theory is essentially analytic: to each perfectoid ring R , one attaches a so-called affinoid perfectoid space $\mathrm{Spa}(R[1/p], R)$ using Huber’s theory of adic spaces [Hub93, Hub94, Hub96]. Roughly speaking, points of this space are identified with certain nonarchimedean valuations on $R[1/p]$, the ring of analytic functions on this space is $R[1/p]$, and the basic open subsets are the loci where certain inequalities are satisfied; one then defines a perfectoid space to be an adic space that is locally of the form $\mathrm{Spa}(R[1/p], R)$. See [Sch12] for details.

Note that the perfectoid space $\mathrm{Spa}(R[1/p], R)$ is best viewed as a tool to study the topological ring $R[1/p]$ and not the perfectoid ring R (i.e., the space “lives in” characteristic 0). As mentioned before, the difference is not too significant as the topological ring $R[1/p]$ almost recovers R .

4. APPLICATION: THE DIRECT SUMMAND CONJECTURE

The first major application of perfectoid geometry we discuss comes from commutative algebra: it is the resolution of Hochster’s *direct summand conjecture* (DSC) by André [And18b, And18a].

The DSC asserts a fundamental property of “regular” rings. Recall that regularity is the commutative algebraic formulation of the geometric notion of smoothness; such rings are always products of domains, and typical examples include fields and PIDs (such as \mathbf{Z}), as well as polynomial or power series rings over such rings.

Theorem 4.1 (DSC). *Let R be a regular ring, and let $f : R \rightarrow S$ be an injective ring map that exhibits S as a finitely generated R -module. Then f admits an R -linear splitting.*

Theorem 4.1 is relatively easy when R has characteristic 0 (i.e., contains \mathbf{Q}), and was settled by Hochster himself [Hoc73] when R has positive characteristic (i.e., contains \mathbf{F}_p for some prime p). Prior to André’s work, the best general mixed characteristic result was Heitmann’s [Hei02], settling Theorem 4.1 when $\dim(R) \leq 3$. In particular, the conjecture was wide open for $R = \mathbf{Z}[x, y, z]$ before André’s work.

Remark 4.2 (Historical comments). The DSC was formulated in the late 60s by Hochster. It has inspired significant amounts of research in commutative algebra; in particular, the modern theory of F -singularities traces its origins to ideas stemming from the resolution of a special case of this conjecture. The DSC has also occupied a central spot in the network of interlocked conjectures in commutative algebra that came to be known as the *homological conjectures*; see [Hoc07, Hoc16] for a survey of the state of affairs immediately prior to Theorem 4.1. Thanks to Theorem 4.1 and related ideas, many of these conjectures, with the notable exception of Serre’s positivity conjecture, are now resolved (see Remark 4.3).

We now give a brief outline of the proof of Theorem 4.1, highlighting the critical role played by the theory of perfectoid rings and spaces.

Outline of the proof of Theorem 4.1 in mixed characteristic. The proof can be divided broadly into three steps.

- (1) Reduce to the case where $R = V[[x_1, \dots, x_n]]$ is a formal power series over a p -adic valuation ring V . (This reduction was known for decades, and is due to Hochster himself.)
- (2) Settle the case where $f[1/p]$ is unramified. The key idea here is to pass from R up to the perfectoid ring R_∞ obtained as the p -adic completion of $R[p^{1/p^\infty}, x_i^{1/p^\infty}]$ (see Example 3.7 (2)), to use the almost purity theorem (Theorem 3.10) to get an “almost splitting” over R_∞ , and then to descend back down to R using faithful flatness of $R \rightarrow R_\infty$ (which makes crucial use of the regularity of R). See [Bha14] for a careful exposition of this step.
- (3) Reduce the general case to the one in (2). This reduction involves first passing up to a carefully constructed perfectoid extension R_∞ of R where a certain discriminant for the map f acquires a compatible system of p -power roots); once R_∞ is constructed, the key new idea is to use (a strengthened form of) the perfectoid analog of the Riemann extension theorem [Sch15, §II.3] to push all the ramification into characteristic p , which then permits us to argue as in (2).

A more complete outline can be found in [And18c]. □

Remark 4.3 (Progress inspired by Theorem 4.1). As mentioned before, the ideas informing Theorem 4.1 have already had significant impact in commutative algebra going well beyond the direct summand conjecture. A non-exhaustive list of examples includes: a simplified proof of DSC and the resolution of the “derived” direct summand conjecture [Bha18a], the analog of Boutot’s theorem in mixed characteristic [HM], the development of a theory of test ideals in mixed characteristic [MS18] robust enough to lift [ELS01] to mixed characteristic, and the construction of weakly functorial big CM algebras and the consequent resolution of many other homological conjectures [And18d, Gab18]. We refer the interested reader to a forthcoming survey article by Ma and Schwede.

5. APPLICATION: THE COHOMOLOGY OF ALGEBRAIC VARIETIES

In this section, we discuss how perfectoid geometry has improved our understanding of the cohomology of algebraic varieties. We begin by recalling the classical Hodge decomposition theorem over the complex numbers in §5.1; the p -adic analog of this result is the subject of §5.2. In §5.3, we explain a (consequence of) a recent result in integral p -adic Hodge theory, which gives geometric representatives for torsion cohomology classes.

5.1. Complex Hodge theory. Hodge theory began with results of de Rham and Hodge describing the singular cohomology of compact complex manifolds in terms of differential data. A summary of this description is the following result [Voi07, §II.6]:

Theorem 5.1 (de Rham, Hodge). *Let X be a compact complex manifold. Assume X admits a Kähler metric. For each $n \geq 0$, there is a canonical isomorphism (called the Hodge decomposition)*

$$H^n(X; \mathbf{C}) \simeq \bigoplus_{i+j=n} H^i(X, \Omega_X^j),$$

where Ω_X^j denotes the sheaf of holomorphic differential j -forms on X .

Theorem 5.1 is a remarkable assertion about the relation between the topology and the geometry of X : the singular cohomology groups $H^n(X; \mathbf{C})$ are purely topological invariants of X , while the cohomology groups $H^i(X, \Omega_X^j)$ are defined in terms of the complex structure on X . In particular, this result allows us to interpret data, such as cycles $\gamma \in H_i(X; \mathbf{C})$, geometrically in terms of differential forms on X . We shall see in §5.3 how one can do the same for cycles $\gamma \in H_i(X; \mathbf{F}_p)$ with torsion coefficients.

5.2. Rational p -adic Hodge theory. Fix a prime number p . Our goal in this section is to discuss the analog of Theorem 5.1 in p -adic geometry. Thus, we shall work over a fixed complete and algebraically closed nonarchimedean extension C of \mathbf{Q}_p (playing the role of the complex numbers). The p -adic analog of Theorem 5.1 describes p -adic étale cohomology (which is the analog of singular cohomology in this setting) in terms of differential forms and is summarized in the following result:

Theorem 5.2 (Hodge-Tate theory). *Let X/C be a proper smooth rigid analytic space (for example, a smooth projective variety). Then there exists a degenerate E_2 -spectral sequence (called the Hodge-Tate spectral sequence)*

$$E_2^{i,j} : H^i(X, \Omega_{X/C}^j)(-j) \Rightarrow H_{\text{ét}}^{i+j}(X; C),$$

where the twist $(-j)$ denotes the Tate twist³.

Moreover, if X is defined over a discretely valued subfield $K \subset C$, then the degeneration of the preceding spectral sequence is canonical, i.e., there exists a canonical isomorphism

$$H_{\text{ét}}^n(X; C) \simeq \bigoplus_{i+j=n} H^i(X, \Omega_{X/C}^j)(-j)$$

of C -vector spaces.

³The Tate twist $(-j)$ in p -adic Hodge theory is roughly analogous to the factors of $(2\pi i)^j$ that appear in complex Hodge theory when relating singular and de Rham cohomology, such as $\int_{S^1} \frac{dz}{z} = 2\pi i$. These can be ignored at first pass, especially if one is solely interested in geometric questions, and not arithmetic ones.

Remark 5.3 (A comment on the proofs). The existence of the Hodge-Tate spectral sequence was established in [Sch13b, §3.3] (which builds on [Sch13a]), while its degeneration was proven [BMS16, §13] using a remarkable algebraization technique due to Conrad-Gabber [CG]. However, it is important to note that Theorem 5.2 records the culmination of a long line of successful research beginning with the work of Tate [Tat67] and Fontaine [Fon82, Fon81]. In particular, special cases of Theorem 5.2 were proven earlier using many different methods (including almost mathematics [Fal88, Fal02, Fal99], syntomic cohomology [FM87, Tsu99], K -theory [BK86, Niz98], and derived de Rham cohomology [Bei12]), and the second half of Theorem 5.2 can be proven directly without knowing in advance about the degeneration of the Hodge-Tate spectral sequence.

We now sketch *very informally* the key steps involved in proving Theorem 5.2 using perfectoid spaces.

- (1) Find a “cover” $f : X_\infty \rightarrow X$ where X_∞ is a perfectoid space, and relate $H_{\text{ét}}^*(X_\infty, C)$ to the coherent cohomology $H^*(X_\infty, \mathcal{O}_{X_\infty})$ on X_∞ using the theory of perfectoid spaces (e.g., the Artin-Schreier sequence, as in Corollary 2.8).
- (2) Calculate $H_{\text{ét}}^*(X, C)$ using the descent spectral sequence for the map f from (1). Differential forms on X naturally appear on the E_2 -page of the relevant spectral sequence.

We refer to [Bha17, Lecture II] for a more thorough exposition of this strategy in the case of abelian varieties.

Remark 5.4 (New period maps). One can use Theorem 5.2 to define certain new “period maps” analogous to constructions in classical Hodge theory. Let us recall the latter first, and then say a few words about the former.

The Hodge decomposition in Theorem 5.1 does not vary holomorphically when X varies in a holomorphic family. Instead, it is the *Hodge filtration* that varies holomorphically, i.e.,

$$\text{Fil}^k H^n(X; \mathbf{C}) := \bigoplus_{i+j=n, j \geq k} H^i(X, \Omega_X^j) \subset H^n(X; \mathbf{C})$$

varies holomorphically with X . This means the following: given a family $\mathcal{X} \rightarrow S$ of compact Kähler manifolds over a simply connected base S , then (after choosing parallel transport isomorphisms between the cohomologies of all the fibers) there is a holomorphic “period map” from S to a suitable Grassmannian given by sending $s \in S$ to the subspace $\text{Fil}^k H^n(\mathcal{X}_s; \mathbf{C}) \subset H^n(\mathcal{X}_s; \mathbf{C})$, where \mathcal{X}_s is the fibre over s . This period map plays an important role in understanding moduli spaces of algebraic varieties over \mathbf{C} .

Similarly, in the p -adic setting, the Hodge-Tate filtration (i.e., the filtration induced by the Hodge-Tate spectral sequence) varies p -adically. One may then mimic the previous discussion to define p -adic period maps (called the *Hodge-Tate period maps*) for certain families of proper smooth rigid spaces. These maps were first discovered by Scholze in [Sch15] in his

study of the moduli space of abelian varieties, and formed the key geometric ingredient in his work on the torsion Langlands correspondence. In fact, an early hint of the power of this technique can be found in the work of Scholze–Weinstein [SW13] that uses the Hodge–Tate filtration to give a p -adic analog of Riemann’s classification of abelian varieties. Relatedly, the Hodge–Tate period map also plays a key role in the work of Caraiani–Scholze [CS17] proving torsionfreeness results for the cohomology of Shimura varieties.

The preceding two period maps, though formally analogous, are distinct: the Hodge filtration over \mathbf{C} and the Hodge–Tate filtration over C go in opposite directions. For example, $H^1(X, \mathcal{O}_X)$ is a quotient of $H^1(X; \mathbf{C})$ for a compact Kähler complex manifold X , while $H^1(Y, \mathcal{O}_Y)$ is a subspace of $H_{\text{ét}}^1(Y; C)$ for a proper smooth rigid space Y/C .

5.3. Integral p -adic Hodge theory. Theorem 5.1 provides a relative satisfactory description of singular cohomology with rational coefficients via differential forms. However, it does not answer the following question:

- (*) Given (say) a smooth projective variety X over \mathbf{C} , is there a way to relate torsion classes in $H^i(X, \mathbf{Z})$ to differential forms?

Recently, a partial answer to (*) was provided in [BMS16], fusing ideas from perfectoid geometry with constructions in homological algebra. Roughly speaking, the answer states that p -torsion classes in $H^i(X, \mathbf{Z})$ can be understood as differential forms on the “mod p reduction” of X . A special case can be stated precisely as follows:

Theorem 5.5. *Let $X \subset \mathbf{P}_{\mathbf{Q}}^n$ be a smooth projective variety defined by equations with coefficients in \mathbf{Z} . Assume that p is a prime of good reduction, i.e., reducing the equations defining X modulo p gives a smooth projective variety $X_p \subset \mathbf{P}_{\mathbf{F}_p}^n$. If we write X^{an} for the complex manifold $X(\mathbf{C})$ attached to X , then we have*

$$\dim_{\mathbf{F}_p} H^i(X^{an}; \mathbf{F}_p) \leq \dim_{\mathbf{F}_p} H_{dR}^i(X_p).$$

Moreover, this inequality can be sharp.

As a corollary, one obtains the following purely algebraic criterion for proving the non-existence of torsion in singular cohomology.

Corollary 5.6. *In the notation from Theorem 5.5, if $\dim_{\mathbf{C}} H_{dR}^i(X^{an}) = \dim_{\mathbf{F}_p} H_{dR}^i(X_p)$, then $H^{i+1}(X^{an}, \mathbf{Z})$ has no p -torsion.*

Remark 5.7. Theorem 5.5 more generally holds for a proper smooth formal scheme over a p -adic valuation ring (see [BMS16]). In fact, one can also replace “smooth” with “semistable” [CK], which is a much more ubiquitous hypothesis in practice. Besides [BMS16], we refer to the surveys [Mor16, Sch16, Bha18b] for more leisurely expositions.

REFERENCES

- [And18a] Yves André, *La conjecture du facteur direct*, Publ. Math. Inst. Hautes Études Sci. **127** (2018), 71–93. MR 3814651
- [And18b] ———, *Le lemme d’Abhyankar perfectoïde*, Publ. Math. Inst. Hautes Études Sci. **127** (2018), 1–70. MR 3814650
- [And18c] ———, *Perfectoid spaces and the homological conjectures*, Available at <https://arxiv.org/abs/1801.10006>, for the proceedings of the ICM 2018.
- [And18d] ———, *Weak functoriality of Cohen-Macaulay algebras*, Available at <https://arxiv.org/abs/1801.10010>.
- [Bei12] Alexander Beilinson, *p -adic periods and derived de Rham cohomology*, J. Amer. Math. Soc. **25** (2012), no. 3, 715–738. MR 2904571
- [Bha14] Bhargav Bhatt, *Almost direct summands*, Nagoya Mathematical Journal **214** (2014), 195 – 204.
- [Bha17] ———, *The Hodge-Tate decomposition via perfectoid spaces*, Lectures for a series at the Arizona Winter School in 2017, notes available at <http://swc.math.arizona.edu/aws/2017/index.html>.
- [Bha18a] ———, *On the direct summand conjecture and its derived variant*, Invent. Math. **212** (2018), no. 2, 297–317. MR 3787829
- [Bha18b] ———, *Specializing varieties and their cohomology from characteristic 0 to characteristic p* , 43–88. MR 3821167
- [BK86] Spencer Bloch and Kazuya Kato, *p -adic étale cohomology*, Inst. Hautes Études Sci. Publ. Math. (1986), no. 63, 107–152. MR 849653 (87k:14018)
- [BMS16] Bhargav Bhatt, Matthew Morrow, and Peter Scholze, *Integral p -adic Hodge theory*, Available at <http://arxiv.org/abs/1602.03148>, submitted.
- [BS] Bhargav Bhatt and Peter Scholze, *Prisms and prismatic cohomology*, In preparation.
- [CG] Brian Conrad and Ofer Gabber, *Spreading out of rigid-analytic varieties*, in preparation.
- [CK] Kestutis Cesnavicus and Teruhisa Koshikawa, *The A_{inf} -cohomology in the semistable case*, Available at <https://arxiv.org/abs/1710.06145>.
- [CS17] Ana Caraiani and Peter Scholze, *On the generic part of the cohomology of compact unitary Shimura varieties*, Ann. of Math. (2) **186** (2017), no. 3, 649–766. MR 3702677
- [Del84] P. Deligne, *Les corps locaux de caractéristique p , limites de corps locaux de caractéristique 0*, Representations of reductive groups over a local field, Travaux en Cours, Hermann, Paris, 1984, pp. 119–157. MR 771673
- [ELS01] Lawrence Ein, Robert Lazarsfeld, and Karen E. Smith, *Uniform bounds and symbolic powers on smooth varieties*, Invent. Math. **144** (2001), no. 2, 241–252. MR 1826369
- [Fal88] Gerd Faltings, *p -adic Hodge theory*, J. Amer. Math. Soc. **1** (1988), no. 1, 255–299. MR 924705 (89g:14008)
- [Fal99] ———, *Integral crystalline cohomology over very ramified valuation rings*, J. Amer. Math. Soc. **12** (1999), no. 1, 117–144. MR 1618483 (99e:14022)
- [Fal02] ———, *Almost étale extensions*, Astérisque (2002), no. 279, 185–270, Cohomologies p -adiques et applications arithmétiques, II. MR 1922831 (2003m:14031)
- [Far] Laurent Fargues, *Geometrization of the local Langlands correspondence: an overview*, Available at <https://webusers.imj-prg.fr/~laurent.fargues/Prepublications.html>.
- [FF] Laurent Fargues and Jean-Marc Fontaine, *Courbes et fibrés vectoriels en théorie de Hodge p -adique*, available at http://webusers.imj-prg.fr/~laurent.fargues/Courbe_fichier_principal.pdf.

- [FM87] Jean-Marc Fontaine and William Messing, *p-adic periods and p-adic étale cohomology*, Current trends in arithmetical algebraic geometry (Arcata, Calif., 1985), Contemp. Math., vol. 67, Amer. Math. Soc., Providence, RI, 1987, pp. 179–207. MR 902593 (89g:14009)
- [Fon81] Jean-Marc Fontaine, *Formes différentielles et modules de Tate des variétés abéliennes sur les corps locaux*, Invent. Math. **65** (1981), no. 3, 379–409. MR 643559
- [Fon82] ———, *Sur certains types de représentations p-adiques du groupe de Galois d'un corps local; construction d'un anneau de Barsotti-Tate*, Ann. of Math. (2) **115** (1982), no. 3, 529–577. MR 657238 (84d:14010)
- [Fon94] ———, *Le corps des périodes p-adiques*, Astérisque (1994), no. 223, 59–111, With an appendix by Pierre Colmez, Périodes p-adiques (Bures-sur-Yvette, 1988). MR 1293971
- [Fon13] ———, *Perfectoïdes, presque pureté et monodromie-poids (d'après Peter Scholze)*, Astérisque (2013), no. 352, Exp. No. 1057, x, 509–534, Séminaire Bourbaki. Vol. 2011/2012. Exposés 1043–1058. MR 3087355
- [FW79] Jean-Marc Fontaine and Jean-Pierre Wintenberger, *Extensions algébrique et corps des normes des extensions APF des corps locaux*, C. R. Acad. Sci. Paris Sér. A-B **288** (1979), no. 8, A441–A444. MR 527692
- [Gab18] Ofer Gabber, *Talk titled “Remarks on big Cohen-Macaulay algebras and on ramification over log-regular rings” at MSRI*, Notes available at <https://www.msri.org/workshops/842/schedules/23854>.
- [GR03] Ofer Gabber and Lorenzo Ramero, *Almost ring theory*, Lecture Notes in Mathematics, vol. 1800, Springer-Verlag, Berlin, 2003. MR 2004652 (2004k:13027)
- [Hei02] Raymond C. Heitmann, *The direct summand conjecture in dimension three*, Ann. of Math. (2) **156** (2002), no. 2, 695–712. MR MR1933722 (2003m:13008)
- [Hen88] K. Hensel, *Ueber die Darstellung der Zahlen eines Gattungsbereiches für einen beliebigen Primdivisor*, J. Reine Angew. Math. **103** (1888), 230–237. MR 1580165
- [HM] Raymond Heitmann and Linquan Ma, *Big Cohen-Macaulay algebras and the vanishing conjecture for maps of Tor in mixed characteristic*, Available at <https://arxiv.org/abs/1703.08281>.
- [Hoc73] M. Hochster, *Contracted ideals from integral extensions of regular rings*, Nagoya Math. J. **51** (1973), 25–43. MR 0349656
- [Hoc07] ———, *Homological conjectures, old and new*, Illinois J. Math. **51** (2007), no. 1, 151–169 (electronic). MR MR2346192 (2008j:13034)
- [Hoc16] ———, *Homological conjectures and lim Cohen-Macaulay sequences*, To appear in the *Proceedings of the Conference on Homological and Computational Methods in Commutative Algebra (honoring Winfried Bruns), Cortona Italy, 2016*.
- [HS06] Craig Huneke and Irena Swanson, *Integral closure of ideals, rings, and modules*, London Mathematical Society Lecture Note Series, vol. 336, Cambridge University Press, Cambridge, 2006. MR 2266432
- [Hub93] R. Huber, *Continuous valuations*, Math. Z. **212** (1993), no. 3, 455–477. MR 1207303 (94e:13041)
- [Hub94] ———, *A generalization of formal schemes and rigid analytic varieties*, Math. Z. **217** (1994), no. 4, 513–551. MR 1306024 (95k:14001)
- [Hub96] Roland Huber, *Étale cohomology of rigid analytic varieties and adic spaces*, Aspects of Mathematics, E30, Friedr. Vieweg & Sohn, Braunschweig, 1996. MR 1734903 (2001c:14046)
- [Kaz86] D. Kazhdan, *Representations of groups over close local fields*, J. Analyse Math. **47** (1986), 175–179. MR 874049
- [Kis] Mark Kisin, *Review of Faltings’ paper “Almost étale extensions” on mathscinet*.
- [KL15] Kiran S. Kedlaya and Ruochuan Liu, *Relative p-adic Hodge theory: foundations*, Astérisque (2015), no. 371, 239. MR 3379653

- [Kra57] Marc Krasner, *Approximation des corps valués complets de caractéristique $p \neq 0$ par ceux de caractéristique 0*, Colloque d’algèbre supérieure, tenu à Bruxelles du 19 au 22 décembre 1956, Centre Belge de Recherches Mathématiques, Établissements Ceuterick, Louvain; Librairie Gauthier-Villars, Paris, 1957, pp. 129–206. MR 0106218
- [Laf] Vincent Lafforgue, *Chtoucas pour les groupes réductifs et paramétrisation de Langlands globale*, Available at <https://arxiv.org/abs/1209.5352>.
- [Mor16] Matthew Morrow, *Notes on the A_{inf} -cohomology of integral p -adic Hodge theory*, Available at <https://arxiv.org/abs/1608.00922>.
- [Mor18] ———, *The Fargues-Fontaine curve and diamonds (d’après Fargues, Fontaine and Scholze)*, Séminaire N. Bourbaki, exposé 1150.
- [MS18] Linquan Ma and Karl Schwede, *Perfectoid multiplier/test ideals in regular rings and bounds on symbolic powers*, Available at <https://arxiv.org/abs/1705.02300>, to appear in *Inventiones*.
- [Niz98] Wiesława Nizioł, *Crystalline conjecture via K -theory*, Ann. Sci. École Norm. Sup. (4) **31** (1998), no. 5, 659–681. MR 1643962
- [Sch] Peter Scholze, Answer to a Mathoverflow question at <https://mathoverflow.net/questions/132438/why-is-faltings-almost-purity-theorem-a-purity-theorem>.
- [Sch12] ———, *Perfectoid spaces*, Publ. Math. Inst. Hautes Études Sci. **116** (2012), 245–313. MR 3090258
- [Sch13a] ———, *p -adic Hodge theory for rigid-analytic varieties*, Forum Math. Pi **1** (2013), e1, 77. MR 3090230
- [Sch13b] ———, *Perfectoid spaces: a survey*, Current developments in mathematics 2012, Int. Press, Somerville, MA, 2013, pp. 193–227. MR 3204346
- [Sch15] ———, *On torsion in the cohomology of locally symmetric varieties*, Ann. of Math. (2) **182** (2015), no. 3, 945–1066. MR 3418533
- [Sch16] ———, *Canonical q -deformations in arithmetic geometry*, Available at www.math.uni-bonn.de/people/scholze/.
- [Sch17] ———, *Étale cohomology of diamonds*, Available at www.math.uni-bonn.de/people/scholze/.
- [SW] Peter Scholze and Jared Weinstein, *p -adic geometry*, Lecture notes from course at UC Berkeley in Fall 2014, available at <https://math.berkeley.edu/~jared/Math274/ScholzeLectures.pdf>.
- [SW13] ———, *Moduli of p -divisible groups*, Camb. J. Math. **1** (2013), no. 2, 145–237. MR 3272049
- [Tat67] J. T. Tate, *p -divisible groups.*, Proc. Conf. Local Fields (Driebergen, 1966), Springer, Berlin, 1967, pp. 158–183. MR 0231827 (38 #155)
- [Tsu99] Takeshi Tsuji, *p -adic étale cohomology and crystalline cohomology in the semi-stable reduction case*, Invent. Math. **137** (1999), no. 2, 233–411. MR 1705837 (2000m:14024)
- [Voi07] Claire Voisin, *Hodge theory and complex algebraic geometry. I*, english ed., Cambridge Studies in Advanced Mathematics, vol. 76, Cambridge University Press, Cambridge, 2007, Translated from the French by Leila Schneps. MR 2451566

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN, 530 CHURCH STREET,
ANN ARBOR, MI 48109, USA

VERIFYING QUANTUM COMPUTATIONS AT SCALE: A CRYPTOGRAPHIC LEASH ON QUANTUM DEVICES

THOMAS VIDICK

ABSTRACT. Quantum computing enthusiasts hope to soon reach a stage where engineered devices based on the laws of quantum mechanics are able to implement computations that can no longer be emulated on a classical computer. Once that stage is reached, will it be possible to verify the results of the quantum device?

Recently Mahadev introduced a solution to the following problem: given black-box access to a quantum device, i.e. given only the ability to generate classical instructions and obtain classical readout information in return, is it possible to delegate a quantum computation to the device in a way that the outcome obtained can be verified on a classical computer — even when the quantum device may be faulty or even adversarially designed to fool the verification procedure?

Mahadev’s solution combines the framework of interactive proof systems from complexity theory with quantum information and an ingenious use of classical cryptographic techniques to tie a “cryptographic leash” around the quantum device. In these notes I give a self-contained introduction to her elegant solution, explaining the required concepts from complexity, quantum computing and cryptography and how they are brought together in Mahadev’s protocol for the verification of quantum computations.

Quantum mechanics has been a source of endless fascination throughout the 20th century — and continues to be in the 21st. Two of the most thought-provoking aspects of the theory are the *exponential scaling* of parameter space (a pure state of n qubits requires $2^n - 1$ complex parameters to be fully specified), and the *uncertainty principle* (measurements represented by non-commuting observables cannot be performed simultaneously without perturbing the state). The conceptual difficulty of the problem of verification of quantum computations stems from both aspects. Suppose given the description of an experiment that can be modeled in quantum mechanics — say, a number N of individual photons are emitted by lasers in a certain configuration, then made to interact according to optical equipment such as mirrors and beam-splitters, and finally some observation is made, for example counting the number of photons that hit a strategically located detector within a certain time period. Quantum mechanics provides a set of formal rules that, *in principle*, allow for the computation of the distribution of possible outcomes obtained in this experiment — what is the probability that any number of photons hit the detector within the prescribed time frame. These rules yield extremely precise predictions that have been verified in countless experiments. In general however, computing the prediction requires a number of operations that scales exponentially with N , the total number of photons in the experiment. What this means in practice is that as soon as N exceeds, say, 80, it becomes all but infeasible, using even

the most powerful supercomputers available today, to predict the outcome of any nontrivial quantum experiment.

This should come as no surprise. Indeed it is the same difficulty, that of classical simulation of quantum evolutions, that prompted Feynman to bring forward the idea of a quantum computer in the first place: a computer that by its very nature would have the ability to efficiently simulate any quantum process. While such a “universal quantum simulator” remains a distant technological challenge, smaller-scale quantum devices have begun to appear that will soon have the capacity to simulate the evolution of specific quantum-mechanical systems, enabling physicists to e.g. make predictions regarding properties of new materials. Such simulators will become interesting the day when they are able to generate predictions that could not have been obtained on a classical computer. Assuming such a “classically unsimulatable quantum simulator”, can we check that the simulator is accomplishing the task it was asked to perform — given that the task could not have been accomplished on a classical computer? If the simulator makes a wrong prediction, be it due to a default in the implementation, or even by malice, is there any way that the error can be detected without having to rely on yet another quantum simulator to duplicate the first simulator’s results?

In addition to the exponential scaling of quantum mechanics, that provides a barrier to the direct classical simulation of quantum experiments, other quantum phenomena, such as the uncertainty principle, place fundamental limits on our ability to verify that a quantum mechanical evolution proceeds as expected. Any attempt by the experimentalist at making intermediate observations on the state of a subset of the elementary particles involved in the operation of her quantum simulator risks altering the outcome of the computation, so that it is not clear at all if the results of low-level benchmarks can be meaningfully pieced together to certify the simulator’s final result.

Not all is lost. There obviously are *some* quantum computations whose outcome can be easily checked. A famous example is factoring: having run Shor’s quantum algorithm for finding a prime factor p of an integer n , it is easy to execute Euclid’s algorithm to check that $p|n$. If every quantum computation had this property — that the correct outcome of the computation can be efficiently verified on a classical computer — then our question would be moot. However, there are very good reasons to think that this is not the case. In the language of complexity theory, the set of languages (informally, problems) that can be decided efficiently with the help of a classical (probabilistic) computer is denoted BPP (for “bounded-error probabilistic polynomial-time”), while with a quantum computer one gets BQP (for “bounded-error quantum polynomial-time”). Clearly BPP is a subset of BQP, and it is generally believed that the latter is strictly larger, with factoring being a possible problem that would separate the two classes. While factoring lies in NP (for “non-deterministic polynomial-time”), the class of languages that can be efficiently *verified* on a classical computer, given the right witness, or proof, complexity theorists strongly believe that there are problems in BQP that do not even lie in NP (they also believe that there are problems in NP, such as the traveling salesman problem or any NP-complete problem, that are not in BQP, so that the two classes are incomparable).¹ Can all efficiently describable experiments based on

¹In fact, a much stronger claim can be proven in the restricted model of “oracle separations”, in which one places restrictions on the ways through which the computer can access its input [14].

quantum mechanics be executed in a way that allows the outcome of the experiment to be certified correct (i.e. matches the predictions of quantum mechanics) using a verification procedure that does not itself rely on the manipulation of quantum information?




FIGURE 1. Known inclusions between complexity classes. Note that NP is not known to contain BPP because in NP, the verification procedure is assumed deterministic. With a randomized verifier, one obtains the class MA (for “Merlin-Arthur”) that does contain BPP.

In these notes I present a solution to the problem of verification of quantum computation due to Mahadev [12]. Her result makes verification possible by injecting an additional ingredient, the use of a post-quantum cryptographically secure scheme, to restore symmetry between the quantum device and the classical experimentalist attempting to verify it. Informally, the experimentalist will use the cryptography to set things up in a way such that the device has no choice but to implement the required computation, or be detected — or break the cryptographic assumption, which is why it is important that we select a scheme that is post-quantum secure, i.e. believed secure even against full-fledged quantum computers. The notes are written in a way that aims to make the most important insights of Mahadev’s work accessible to any mathematician, with or without background in complexity theory or quantum information. As a result, the presentation will remain rather informal, and we alert the reader whenever the discussion makes an important shortcut.

The outline for the remainder of the notes is as follows. First we introduce the formal setup for the verification problem. This will require us to delve in the wonderful framework of *interactive protocols* from complexity theory, that we illustrate with a toy scheme for the verification of *classical* computation. Then we describe a specific problem about quantum states, which the verification question reduces to. Finally, we describe the protocol of Mahadev and give the key ideas that go in its analysis.

To keep the presentation focused we do not survey prior works and other approaches to verification in any depth. The question has a long history, that prior to Mahadev’s work had resulted in partial answers for different models of verification. Some of the most important results include the concurrent works of Aharonov et al. [1] and Broadbent et al. [5] showing how to achieve verification in a model

where the verification procedure itself has access to a small, trusted quantum computer, and the work of Reichardt et al. [17] in a model where the verification procedure is entirely classical, but has access to two spatially isolated quantum computers, sharing entanglement, whose implementation of a quantum computation it aims to verify. In contrast to Mahadev’s result presented here, these works achieve information-theoretic (instead of computational) soundness guarantees in their respective models. For an in-depth discussion of these and related works, I recommend the recent survey by Gheorghiu et al. [9].

1. INTERACTIVE PROTOCOLS

The concept of an *interactive proof system* can be difficult to digest for the mathematician, in part because it involves some amount of personification — there is the *verifier*, the *prover* — and even worse, these imaginary beings are ascribed *intentions* — the prover is *trying* to demonstrate something to the verifier, while the verifier *attempts* to catch any *cheating behavior* from the prover — not the kind of language that is frequently used in, say, the theory of differential equations or operator spaces. Please bear with me — interactive proofs are the single most powerful idea to have emerged out of complexity theory since the 1990s, and they are a key element of Mahadev’s solution.

It all starts with the complexity class NP, whose study originates in the works of Cook, Karp, and Levin in the 1970s. A complexity class is a collection of *languages*. A (promise) language L is a pair of subsets $L_{yes}, L_{no} \subseteq \{0, 1\}^*$, the set of sequences over the alphabet $\{0, 1\}$ of arbitrary but finite length (a.k.a. “bit strings”). For example, L_{yes} could be the set of (suitably encoded) 3-SAT formulas that admit a satisfying assignment,² and L_{no} the set of formulas that are not satisfiable. The language $L = (L_{yes}, L_{no})$ is called 3-SAT. A language L is in the class NP if there exists a real polynomial p and a Turing machine V (for our purposes, the reader may replace the intimidating notion of “Turing machine” by any intuitive notion of efficient computation, such as an “algorithm”) such that V has two input tapes, and is such that for all $x \in L_{yes} \cup L_{no}$, (i) if $x \in L_{yes}$ then there exists a w , the *witness*, or *proof*, such that $V(x, w)$ halts in at most $p(|x|)$ steps, where $|x|$ denotes the length of x , and returns 1 (for “accept”), and (ii) if $x \in L_{no}$ then for any w , $V(x, w)$ halts in at most $p(|x|)$ steps and returns 0 (for “reject”). Property (i) is usually referred to as the *completeness* condition (valid statements have an accepted proof), and (ii) as the *soundness* condition (invalid statements have no accepted proof). The fact that 3-SAT is in NP follows since given as input a 3-SAT formula φ , if the formula is satisfiable then there exists a witness w that proves this (e.g. w is a satisfying assignment for φ), whereas if the formula is not satisfiable, it is easy for V to check that any given purported assignment w is indeed invalid.

Informally, NP captures the collection of problems that have efficiently verifiable solutions, such as factoring, 3-SAT, the traveling salesman problem, or mathematical theorems (that have reasonably short proofs within a prespecified axiomatic system). A key insight from research in complexity theory in the 1990s is that the collection of languages that admit “efficient verification” can be substantially extended by allowing *interactive protocols* for verification. Consider a verification

²A 3-SAT formula is an AND of 3-variable ORs, e.g. a Boolean formula of the form $\varphi = (x_1 \vee x_2 \vee \bar{x}_3) \wedge (\bar{x}_2 \vee x_5 \vee \bar{x}_6) \wedge \dots$, where the variables $x_i \in \{0, 1\}$ and \bar{x}_i denotes variable negation.

procedure V , referred to as the *verifier* (personification, here it comes...), that is allowed to “ask questions” about a claimed proof to another entity, the *prover*, as illustrated in Figure 2. Can this provide an advantage, in the sense of allowing verification of languages that lie beyond those languages in NP, which admit “static” proofs?




FIGURE 2. An example of a 4-message interactive proof between a verifier V , assumed to be computationally bounded, and an all-powerful prover P .

Note that in the context of interactive proofs it is always assumed that both verifier and prover are given access to an input x (e.g. a 3-SAT formula), and that the prover aims to convince the verifier that the input is in the language (e.g. the formula is satisfiable). The verifier, in turn, aims to make the right decision, by taking information from the prover but “verifying” it before making a decision. This reflects an asymmetry in the definition of NP: for valid statements, there should exist a proof, while for invalid statements, no purported proof should be accepted. In contrast to NP, for general interactive proof systems we allow the verifier to be randomized and sometimes make an erroneous decision, as long as for every input the probability of making an error in deciding that input is small: this ensures that repeating the verification procedure sufficiently many times and taking the majority decision yields an outcome that is erroneous with arbitrarily small probability.³ Formally, we have the following two requirements:

- *Completeness*: For any $x \in L_{yes}$, there exists a prover that is accepted by the verifier with probability at least $\frac{2}{3}$.
- *Soundness*: For any $x \in L_{no}$ and any prover, the verifier accepts with probability at most $\frac{1}{3}$.

The set of languages for which there exists an interactive proof system satisfying these two requirements is denoted IP. To gain intuition as to how interaction may help, consider the following toy problem. The input x is interpreted as a bivariate polynomial $P(y, z)$ defined over a prime field \mathbb{F}_p , such that P has degree at most d in each of its variables. Think of p as much larger than d : the input size, measured in number of bits, is roughly $(d + 1)^2 \log p$ (a list of coefficients), and p could be exponentially larger. Let L_{yes} be the set of all such polynomials such that $S = \sum_{y,z} P(y, z) = 0$, and L_{no} the set of polynomials such that $S \neq 0$. Computing S naïvely takes time $O(p^2)$, where the $O(\cdot)$ hides a multiplicative factor of order the time it takes to evaluate P at any given point, i.e. polynomial in d and $\log p$.

³It is not hard to see that interactive proof systems with deterministic verifiers cannot decide languages beyond NP, so this is an essential modification.

Consider the following protocol, using which the verifier can be led to making the right decision while having to invest a computational effort of $O(p)$ only. The first step in the protocol is a message from the prover to the verifier, which consists in the claimed value of S , together with a “supporting statement” in the form of a degree- d polynomial $Q(z)$ such that $S = \sum_y Q(z)$. (The honest prover can find such a Q without difficulty by setting $Q(z) = \sum_y P(y, z)$.) Upon receipt of an arbitrary (S, Q) , the verifier first checks the equality $S = \sum_z Q(z)$, which takes $O(p)$ time. Next, she selects a uniformly random $z^* \in \mathbb{F}_p$ and checks that $Q(z^*) = \sum_y P(y, z^*)$, which again requires time $O(p)$. Note that, if it was the case that $Q(\cdot) \neq \sum_y P(y, \cdot)$, by the Schwartz-Zippel lemma the probability that $Q(z^*) = \sum_y P(y, z^*)$ would be at most d/p , which is small provided p is much larger than d , as we assumed. Thus the verifier makes the right decision with high probability, on any input P .

This “protocol” reduces the verifier’s effort from order p^2 to order p . The attentive reader will have realized that the protocol is in fact non-interactive — there is a single message from the prover to the verifier. Applying the same idea to polynomials with more variables yields an interactive protocol, in which variables are randomly fixed by the verifier one at a time, with exponential savings in the amount of time required for verification, from order p^m to order mp , with m representing the number of variables. This idea, of efficiently verifying claims about the sum of the values taken by a multivariate polynomial, forms the basis of a celebrated result in complexity theory, the inclusion of PSPACE (the class of languages that can be decided using polynomial space, but arbitrary time) in IP [11, 18]. While the state of the art in complexity theory does not allow one to prove that PSPACE is a strictly larger class than NP, this is generally believed to be the case, so that interaction seems to significantly broaden the class of languages that have efficient verifiers.

In fact, it sufficiently broadens it so as to encompass the class BQP of languages that can be efficiently decided on a *quantum* computer! Using the idea of Feynman path integrals the probability that a quantum circuit returns the outcome “0” can be expressed as the sum of exponentially many complex numbers, each of which can be directly computed by taking a product of entries of the matrices that specify the quantum gates of the circuit; this exponential sum can be exactly computed (modulo finite-precision issues) in polynomial space. Given that PSPACE is in IP, it follows that there exists an interactive protocol, of the form described above, that allows a classical polynomial-time verifier to verify the outcome of a quantum computation by asking questions to an untrusted prover. But there is a major hitch. The model of interactive proofs does not place limitations on the computational effort required of the prover, even in the case when it needs to come up with the “right” argument to convince the verifier (i.e. for an input $x \in L_{yes}$). In the protocol for computing sums of polynomials described earlier, the prover has to compute multiple exponential-sized intermediate sums, that in general will take time p^m . Unfortunately, following the proofs that BQP is in PSPACE is in IP leads to a very similar protocol, in which the prover has to compute exponentially large sums for which there is no reason to think there would be an efficient quantum algorithm.

There has been very little progress in tweaking the protocols obtained in this way to decrease the computational effort required for the honest prover (see e.g. the recent [2]). Mahadev’s work takes a different path by explicitly introducing a

computational limitation on the cheating prover: that it does not have the ability to break a cryptographic scheme that is believed secure against quantum polynomial-time attackers. To explain her approach we first introduce the model of quantum computation and the specific problem that is verified by her protocol; this is done in the next section.

2. QUANTUM COMPUTATIONS

In this section we give a light introduction to the formalism of quantum computing. Our goal in doing so is to provide the minimal background required to describe a reduction from the problem of deciding the outcome of a quantum circuit to another problem, of deciding the existence of a quantum state that satisfies certain simple constraints, that is the starting point for Mahadev’s protocol. The latter decision problem is formulated in Theorem 2.3, so that the impatient reader may skip ahead, read the statement of the theorem, and proceed to the next section, in which we describe Mahadev’s protocol. (Notation used in the theorem statement is introduced progressively throughout this section.)




FIGURE 3. Schematic representation of an instance of the Ising spin problem. Here each σ_i is a variable in $\{0, 1\}$, and each J_{ij} a fixed coupling constant in $[-1, 1]$ (not all J_{ij} are represented on the figure). The goal is to find an assignment to the variables that minimizes the expression (2.1).

2.1. Warmup: the Ising spin problem. As a warmup, let’s consider a problem from classical statistical physics, and see how we can reduce the verification of a *classical* computation to it. Consider a graph with n vertices, such that each vertex $i \in \{1, \dots, n\}$ is associated a value (or, “state”) $\sigma_i \in \{0, 1\}$, and each edge (i, j) is associated a real weight (or, “coupling constant”) J_{ij} such that $|J_{ij}| \leq 1$. (See Figure 3.) Vertices represent particles that can be in one of two states, $\sigma_i = 0$ or $\sigma_i = 1$, and edges represent interactions between particles, where the interaction can be attractive ($J_{ij} < 0$) or repulsive ($J_{ij} > 0$). The *energy* of a configuration $\sigma \in \{0, 1\}^n$ is defined as⁴

$$(2.1) \quad H_{\text{ising}}(\sigma) = - \sum_{(i,j)} J_{ij} (-1)^{\sigma_i + \sigma_j} .$$

Informally, the energy functional (a.k.a. *Hamiltonian*) H_{ising} measures the number of edge constraints that are not satisfied by an assignment σ , with each violation

⁴Technically the expression in (2.1) can take negative values, which may not seem appropriate for an “energy”. Correcting this is just a matter of introducing an additive shift.

giving a penalty of $|J_{ij}|$. It is well-known that the problem of deciding, given as input n , the coefficients J_{ij} , and two thresholds a, b such that $b - a$ is at least a constant independent of n , whether the minimum of H_{ising} over all $\sigma \in \{0, 1\}^n$ is less than a , or larger than b , is an NP-complete problem (i.e. any problem in NP reduces to it); moreover, this holds even for the case where $J_{ij} \in \{-1, 0, 1\}$ for all (i, j) .




FIGURE 4. The tableau of a classical circuit. Here $x = 10$, $w = 10$, and there is one ancilla qubit, initialized to 0. The circuit has 5 reversible NAND gates $(a, b) \mapsto (a, 1 - ab)$. The tableau is given by $\sigma \in \{0, 1\}^{25}$, that represents the state of each of the 5 circuit wires at successive stages of an execution of the circuit.

To understand why the Ising problem is as hard as any problem in NP, let's see how we can reduce the problem of deciding whether, given a classical circuit \mathcal{C} acting on $n + m + r$ bits and an input string $x \in \{0, 1\}^n$, there exists a string $w \in \{0, 1\}^m$ such that the circuit accepts $(x, w, 0^r)$. By definition any problem in NP can be expressed in this form, with x the input, w the witness, and \mathcal{C} the verifier's circuit. In general \mathcal{C} is specified by a sequence of gates (C_1, \dots, C_ℓ) taken from some fixed gate set, that we may without loss of generality restrict to the sole (reversible) NAND gate, that maps (a, b) to $(a, 1 - ab)$. Next consider the *tableau* of the computation performed by the circuit. This is simply a list of values for all wires in the circuit, starting from the input wires (initialized to $(x, w, 0^r)$) to the output wire (that should equal 1, which stands for "accept"). (See also Figure 4.) Given a tableau, it is possible to verify that the tableau is correct by checking the propagation of each gate, one at a time: if the inputs to a NAND gate are $\sigma_{i_1}, \sigma_{i_2} \in \{0, 1\}$, the output should be $(\sigma_{i_3} = \sigma_{i_1}, \sigma_{i_4} = 1 - \sigma_{i_1}\sigma_{i_2})$. Wires corresponding to the input string x should be initialized to the right value, whereas wires associated with the witness string w can take arbitrary values. We can express this set of constraints as a Hamiltonian $H_{\mathcal{C}} : \{0, 1\}^T \rightarrow \mathbb{R}$, where T is the total number of wires in the circuit:

$$(2.2) \quad H_{\mathcal{C}} = H_{in} + H_{prop} + H_{out},$$

where H_{in} is a summation of energy penalties⁵ $1_{\sigma_{i_k} \neq x_k}$ with i_k the input wire associated with the k -th bit of x and $1_{\sigma_{i_j} \neq 0}$ with i_j the input wire associated with the j -th ancilla bit, H_{prop} a summation of penalties of the form $1_{\sigma_{i_k} \neq \sigma_{i_j}} + 1_{\sigma_{i_{k+1}} \neq 1 - \sigma_{i_j}\sigma_{i_\ell}}$ for each NAND gate mapping (i_j, i_ℓ) to (i_k, i_{k+1}) , and H_{out} consists of a single penalty term $1_{\sigma_{i_T} \neq 1}$, for i_T the output wire. Then, $H_{\mathcal{C}}$ is such that there

⁵We use the notation 1_E for the indicator that event E occurs.

exists σ such that $H_{\mathcal{C}}(\sigma) = 0$ if and only if there is a witness w such that \mathcal{C} accepts $(x, w, 0^r)$; otherwise, $H_{\mathcal{C}}(\sigma) \geq 1$ for all σ .

Note that $H_{\mathcal{C}}$ doesn't quite have the form (2.1) of an Ising spin Hamiltonian yet: some of its terms involve three variables at a time, and moreover not all terms are directly expressed as a function of the parity of the sum of two variables. With a little more work, using so-called "gadgets" it is possible to complete the reduction and find an $H'_{\mathcal{C}}$ that is equivalent to $H_{\mathcal{C}}$ (in terms of capturing the same decision problem) but is of the form (2.1).

Remark 2.1. By adding penalty terms $1_{\sigma_{i_k} \neq 0}$ for all input wires associated to w we can force the witness to be 0, in which case the Hamiltonian directly captures the problem of deciding if the circuit accepts its input x , or not. The above reduction thus maps the problem of deciding whether a given classical circuit returns the outcome 1, or not, to the problem of deciding whether an Ising Hamiltonian has a configuration with energy 0, or not.

Our next step is to devise a quantum analogue of this reduction. For this we first introduce some of the basics of quantum computation: quantum states, operations, and measurements.

2.2. Quantum states and observables.

States. An n -qubit quantum state is specified by a *density matrix*, a positive semi-definite matrix ρ on the 2^n -dimensional Hilbert space $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ such that ρ has trace 1. Density matrices generalize classical probability distributions over n bits, as the latter can be represented by a probability vector $p : \{0, 1\}^n \rightarrow [0, 1]$ that we can embed on the diagonal of a density matrix.

Even though in general we may allow arbitrary Hilbert spaces, for convenience we generally assume that the space comes endowed with a canonical decomposition as a tensor product of n copies of \mathbb{C}^2 , for some finite integer n , and that moreover each copy of \mathbb{C}^2 has a canonical basis (e_1, e_2) . We generally use the "ket" notation to write the canonical basis as $|0\rangle = e_1$, $|1\rangle = e_2$, and we refer to this basis as the "computational basis". A quantum state is called *pure* if its density matrix has rank 1; in this case we can also represent the state as a unit vector expressed in the canonical basis $\{e_{i_1} \otimes \cdots \otimes e_{i_n}\}$, or $\{|e_{i_1} \cdots e_{i_n}\rangle\}$ in the more compact ket notation. An arbitrary pure state thus has an expansion $|\psi\rangle = \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle$, where the $\{\alpha_x\}$ are complex coefficients such that $\sum_x |\alpha_x|^2 = 1$. The associated density matrix is the rank-1 projection $\rho = |\psi\rangle\langle\psi| = |\psi\rangle\langle\psi|$, where the "bra" notation $\langle\psi| = (|\psi\rangle)^\dagger$ is used for the conjugate-transpose.

Evolution. Evolutions in quantum mechanics are unitary. For a unitary U , a pure state $|\psi\rangle$ evolve as $|\psi\rangle \mapsto U|\psi\rangle$, and a density matrix ρ evolves as $\rho \mapsto U\rho U^\dagger$. For U to be implementable on a quantum computer we require that it decomposes as a product $U = U_T \cdots U_1$, where each U_i is a unitary that acts non-trivially on at most two qubits, i.e. it can be written as a tensor product of a unitary on $\mathbb{C}^2 \otimes \mathbb{C}^2$ with the identity on the remaining space. Moreover, each U_i should be taken from a finite "gate set" of allowed operations on the computer. A sample gate set contains the unitaries $\{H, T, CNOT\}$, where $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ is the Hadamard gate, $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$ the T (sometimes also called $\pi/8$) gate, and CNOT the two-qubit unitary that sends $|a\rangle|b\rangle \mapsto |a\rangle|a \oplus b\rangle$ for any $a, b \in \{0, 1\}$.

A fundamental theorem in quantum computing, the Solovay-Kitaev theorem, states that for the purpose of efficient circuit representations any finite set of 1 and 2-qubit gates is as good as any other, as long as it generates a dense subgroup in $SU(2)$ (which is the case for the above-defined set). More formally,

Theorem 2.2 (Solovay-Kitaev '97). *There is a constant c such that for any finite gate set $G \subseteq SU(2)$ such that the group $\langle G \rangle$ generated by G is dense in $SU(2)$ and G is closed under inverse, for any $\varepsilon > 0$ there is an $\ell = O(\log^c(1/\varepsilon))$ such that G^ℓ is an ε -net in $SU(2)$.⁶*

Given a function $f : \{0,1\}^n \rightarrow \{0,1\}^m$ specified by a classical circuit, it is possible to devise a quantum circuit of comparable size for the unitary U_f that maps $|x\rangle|b\rangle$ to $|x\rangle|f(x) \oplus b\rangle$ for $x \in \{0,1\}^n$, $b \in \{0,1\}^m$. (This observation is not completely trivial, as the classical circuit may use non-reversible gates, for which there is no direct quantum analogue; nevertheless, it is possible to show that any classical circuit over, say, $\{AND, OR, NOT\}$ can be efficiently simulated by a reversible circuit.) We often consider the application of the unitary U_f “in superposition”: by linearity,

$$U_f : \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle |0^m\rangle \mapsto \sum_{x \in \{0,1\}^n} \alpha_x |x\rangle |f(x)\rangle .$$

Observables. It remains to discuss measurements. A measurement of a set of qubits is specified by an orthonormal basis of the Hilbert space associated with the qubits. The outcome of the measurement is the label of one of the basis vectors, and the probability with which each basis vector is obtained equals the squared norm of the component of the state that is in the direction of that basis vector. Formally, suppose $|\psi\rangle$ is a state in $(\mathbb{C}^2)^{\otimes n} \otimes (\mathbb{C}^2)^{\otimes m}$, and that the first n qubits of $|\psi\rangle$ are measured in the orthonormal basis $\{|\phi_i\rangle\}_{i \in 1, \dots, 2^n}$ of $(\mathbb{C}^2)^{\otimes n}$. To compute the probability of the i -th outcome being obtained, we expand $|\psi\rangle$ in the basis $\{|\phi_i\rangle\}$ as

$$|\psi\rangle = \sum_{i=1}^{2^n} |\phi_i\rangle |\phi'_i\rangle ,$$

where the $|\phi'_i\rangle$ are arbitrary vectors in $(\mathbb{C}^2)^{\otimes m}$ (not necessarily normalized or orthogonal). The probability of the i -th outcome is then given by $\| |\phi'_i\rangle \|^2$. It will later be important to remember that a measurement *collapses* the state: once the outcome i has been obtained and recorded, the state undergoes a non-unitary evolution $|\psi\rangle \mapsto |\psi_i\rangle = \frac{|\phi_i\rangle |\phi'_i\rangle}{\| |\phi_i\rangle \|}$.⁷

A measurement in the basis $\{|\phi_i\rangle\}$, together with a choice of a real number λ_i associated with each outcome i , can be succinctly represented as an “observable” $O = \sum_i \lambda_i |\phi_i\rangle \langle \phi_i|$. For a quantum state ρ , the real number $\text{Tr}(O\rho)$ is precisely the expectation of λ_i , under the distribution on i obtained by measuring the state ρ in the basis $\{|\phi_i\rangle\}$. An example is the observable associated with a measurement of a

⁶An ε -net is a set of points $S \subseteq SU(2)$ such that for all $U \in SU(2)$, there is $V \in S$ such that $\|U - V\| \leq \varepsilon$. Here the norm is the operator norm, but any norm would give essentially the same result, since the space has small dimension.

⁷If the conflict between the statements that “quantum mechanics requires all evolutions to be unitary” and “a measurement is an irreversible process” puts you ill at ease, you are not alone.

qubit in the computational basis $\{|0\rangle, |1\rangle\}$, labeling the first outcome “1” and the second “−1”, whose associated observable is the Pauli σ_Z matrix,

$$\sigma_Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Similarly, a measurement in the Hadamard basis $\{H|0\rangle, H|1\rangle\}$ is represented by the Pauli σ_X observable,

$$\sigma_X = H\sigma_ZH^\dagger = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

2.3. Quantum spin problems. With the basics of quantum computing in place, we return to the Ising spin problem introduced in Section 2.1 and discuss its “quantization”. Recall the Hamiltonian H_{ising} introduced in (2.1). We can interpret this Hamiltonian as an “energy functional”, that associates an energy to any configuration σ . In quantum mechanics, a Hamiltonian is any linear operator on Hilbert space, with the restriction that the operator should be Hermitian (and bounded; for convenience here we only consider finite-dimensional spaces, so that the latter condition is automatic). The interpretation is that the Hamiltonian associates a definite energy $\lambda_i \in \mathbb{R}$ to any quantum state that happens to be in one of its eigenstates $|\phi_i\rangle$. The energy of an arbitrary state ρ is then computed linearly as $\text{Tr}(H\rho)$.⁸ Often it is also required that the Hamiltonian be *local*, meaning that H can be expressed as a sum of a polynomial number of terms h_i , each of them the tensor product of the identity on $(n - k)$ qubits, and an arbitrary Hamiltonian on the remaining k qubits, for some constant k .

Using the notation introduced in the previous section, the Ising spin Hamiltonian can be recast as a quantum Hamiltonian, $H'_{ising} = -\sum_{(i,j)} J_{ij} \sigma_Z^i \sigma_Z^j$, where σ_Z^i is shorthand for the observable that is σ_Z on the i -th qubit and the identity on the others, $\sigma_Z^i = \text{Id}^{\otimes(i-1)} \otimes \sigma_Z \otimes \text{Id}^{\otimes(n-i)}$. Since this Hamiltonian is diagonal in the computational basis, its eigenstate with smallest eigenvalue, also called its “ground state” or minimal energy state, is always attained at a pure computational basis state $|\sigma\rangle$, for some $\sigma \in \{0, 1\}^n$.

Things get more interesting when we consider Hamiltonians made of a combination of non-commuting observables. Consider for example the 2-qubit Hamiltonian

$$(2.3) \quad H_{EPR} = -\frac{1}{2}(\sigma_X^1 \sigma_X^2 + \sigma_Z^1 \sigma_Z^2).$$

As a matrix, this can be written as

$$H_{EPR} = \begin{pmatrix} -1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & -1 & 1 & 0 \\ -1 & 0 & 0 & -1 \end{pmatrix}.$$

⁸The reader may have noticed that the syntactic requirements for “Hamiltonians” and “observables” are identical. Physically, a Hamiltonian is meant to represent a specific observable, that corresponds to the energy of a system; mathematically, the two notions are interchangeable.

What is remarkable about this Hamiltonian is that its smallest-eigenvalue eigenstate is a state

$$|\phi^+\rangle = \frac{1}{\sqrt{2}}(|0\rangle|0\rangle + |1\rangle|1\rangle) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 0 \\ 0 \\ 1 \end{pmatrix},$$

also known as “EPR pair”, that has the property of being *entangled*: it cannot be expressed in the form $|\phi_1\rangle \otimes |\phi_2\rangle$, for any $|\phi_1\rangle$ and $|\phi_2\rangle$.

The possibility for quantum Hamiltonians to force entanglement in their ground state distinguishes them from classical Hamiltonians, whose eigenstates are computational basis states, and in particular product states. As a result, while a classical Hamiltonian always has a minimal energy configuration that can be described using $O(n)$ bits (hence, as already observed, the problem of deciding its minimum energy is in NP), for quantum Hamiltonians this need not be the case. The complexity class QMA (for “Quantum Merlin-Arthur”) is the quantum analogue of NP: QMA is the collection of all (promise) languages $L = L_{yes} \cup L_{no}$ such that it is possible for a quantum polynomial-time verifier to correctly decide whether an input $x \in L_{yes}$ or $x \in L_{no}$, with error at most $\frac{1}{3}$, with the help of a “quantum proof” $|\phi\rangle$ provided by an all-powerful, but untrusted, quantum prover. The problem of deciding the minimal energy of a local Hamiltonian, to within some inverse polynomial precision, is an example of a problem that is in QMA: informally, given a claimed minimum-eigenvalue eigenstate presented as a quantum state, it is possible to estimate the associated eigenvalue by making the appropriate energy measurement. Moreover, Kitaev established a quantum analogue of NP-completeness of 3-SAT by showing that the local Hamiltonian problem is QMA-complete, i.e. the constraints expressed by any polynomial-time quantum verification procedure can be reduced to constraints of the form checked by a local Hamiltonian.

2.4. Certificates for quantum computations. Back to our quest for a certificate for quantum computations. We have seen that the computation carried out by a classical circuit could be represented by a “tableau”, such that the property of being a valid tableau can be encoded in a classical Hamiltonian, thereby reducing the task of deciding whether a classical circuit accepts its input to the task of deciding whether the associated Hamiltonian (that can be efficiently computed from the circuit) has a small enough eigenvalue.

What is the correct notion of a tableau for quantum circuits? The first idea is to consider the juxtaposition of the quantum state of an ℓ -gate circuit at each step of the computation, i.e. the tensor product $|\psi_0\rangle \otimes \cdots \otimes |\psi_\ell\rangle$ of the states $|\psi_i\rangle$ obtained by executing the circuit from scratch and stopping after i gates have been applied. While this is a well-defined $n(\ell+1)$ -qubit quantum state (see Figure 5) the property of being a valid “quantum tableau” cannot be enforced using a *local* Hamiltonian! The reason is subtle, and has to do with the possible presence of entanglement at intermediate steps of the computation. Indeed, there are quantum states that are very different, in the sense that they are perfectly distinguishable by some *global* observable, yet cannot be distinguished at all by any *local* observable, that would act on at most, say, half the qubits. An example is given by the two n -qubit “cat” (named after the homonymous animal) states

$$|\psi_\pm\rangle = \frac{1}{\sqrt{2}}(|0 \cdots 0\rangle \pm |1 \cdots 1\rangle).$$




FIGURE 5. Two different ways to create a tableau from a quantum circuit. The state $|\psi_{\text{tableau}}\rangle$ is the tensor product of the state of the circuit at each time step. The state $|\psi_{\text{hist}}\rangle$ is their superposition, indexed by a clock register that goes from $|0\rangle$ to $|T\rangle$.

The two states $|\psi_+\rangle$ and $|\psi_-\rangle$ are easily seen to be orthogonal, so that they can be perfectly distinguished by a measurement (using any orthonormal basis that contains both states). But it is an exercise to verify that for any observable that acts on at most $(n-1)$ of the n qubits, both states give exactly the same expectation value. (Informally, this is because any measurement on a strict subset of the qubits of the state necessarily destroys the coherence — the only relevant information, the \pm sign, is encoded “globally” and cannot be accessed locally.) Note that this is a uniquely quantum phenomenon: if two classical strings of bits have each of their bits equal, one pair at a time, then the strings are “globally” identical. Not so for quantum states.

So naïve tableaus will not do. In the late 1990s the physicist Alexei Kitaev introduced a very powerful idea that provides a solution. Kitaev’s idea is to replace the juxtaposition of snapshot states by their *superposition* (see Figure 5). A special ancilla system, called the “clock”, is introduced to index different elements of the superposition. Thus, instead of defining a tableau as $|\psi_0\rangle \cdots |\psi_\ell\rangle$, Kitaev considers the state

$$(2.4) \quad |\psi_{\text{hist}}\rangle = \frac{1}{\sqrt{\ell+1}} \sum_{i=0}^{\ell} |i\rangle |\psi_i\rangle.$$

Kitaev showed that, assuming the clock register is encoded in unary, it is possible to check the correct propagation of every step of the circuit directly on this superposition by only applying local observables, in a manner very similar to what we did for classical tableaus: there is a set of observables H_{in} that checks that $|\psi_0\rangle$ has the right format; a set of observables H_{prop} that checks propagation of the circuit, and an observable H_{out} that checks that the “output qubit” of the circuit is in the right state. The key point that makes this possible is that, while equality of quantum states cannot be decided locally when the states are juxtaposed, it becomes possible when they are given in superposition: as an exercise, verify that a measurement of the first qubit of the state

$$|\psi_{\text{SWAP}}\rangle = \frac{1}{\sqrt{2}} (|0\rangle |\psi_0\rangle + |1\rangle |\psi_1\rangle)$$

in the Hadamard basis $\{H|0\rangle, H|1\rangle\}$ returns the first outcome with probability exactly $\frac{1}{2}(1 + |\langle\psi_0|\psi_1\rangle|^2)$. With more work, replacing the use of gadgets for the

classical case by techniques from perturbation theory, it is possible to write the resulting Hamiltonian as a linear combination of local terms that all take the form of the EPR Hamiltonian (2.3). (Such a Hamiltonian is called a Hamiltonian “in XZ form”, for obvious reasons.) The result is the following theorem from [6].

Theorem 2.3. *For any integer $n \geq 1$ there are $n' = \text{poly}(n)$, $a = a(n)$ and $\delta \geq 1/\text{poly}(n)$ such that the following holds. Given an ℓ -gate quantum circuit $\mathcal{C} = C_1 \cdots C_\ell$ acting on n qubits, such that $\ell = \text{poly}(n)$, and an input x for the circuit, there exist efficiently computable real weights $\{J_{ij}, i, j \in \{1, \dots, n'\}\}$ such that $|J_{ij}| \leq 1$ and if*

$$(2.5) \quad H_{\mathcal{C}} = - \sum_{i,j} \frac{J_{ij}}{2} (\sigma_X^i \sigma_X^j + \sigma_Z^i \sigma_Z^j),$$

then:

- (Completeness) *If the circuit \mathcal{C} accepts its input x with probability at least $2/3$,⁹ then the smallest eigenvalue of $H_{\mathcal{C}}$ is at most a ;*
- (Soundness) *If the circuit \mathcal{C} accepts its input x with probability at most $1/3$, then the smallest eigenvalue of $H_{\mathcal{C}}$ is at least $a + \delta$.*

Remark 2.4. Analogously to Remark 2.1 about the Ising spin problem being NP-hard, it is possible to modify Theorem 2.3 so that the completeness and soundness statements specify that “if there exists a state $|\phi\rangle$ such that \mathcal{C} accepts on input $(x, |\phi\rangle)$ with probability at least $2/3$...” and “if there does not exist a state $|\phi\rangle$ such that \mathcal{C} accepts on input $(x, |\phi\rangle)$ with probability greater than $1/3$...” respectively. Thus, Theorem 2.3 can be adapted to show that the problem of estimating the minimal energy of a Hamiltonian of the form (2.5) is a QMA-complete problem.

Theorem 2.3 provides us with a roadmap for the verification of quantum circuits: it is sufficient to verify the *existence* of a quantum state that yields certain statistics, when some of its qubits are measured in the computational (σ_Z observable) or Hadamard (σ_X observable) basis. The reason this can be considered progress is that we no longer need to check any dynamics; it is sufficient to collect measurement statistics and estimate the energy. In particular, the theorem readily leads to a verification protocol in a model where the prover has a full quantum computer, and the verifier only has a limited quantum device — namely, a one-qubit memory, together with the ability to measure the qubit using either the σ_X or σ_Z observables.

Such a verification protocol was introduced in [7], and can be summarized as follows. First, given as input a description of the circuit \mathcal{C} the prover is asked to prepare a state $|\psi\rangle$ such that $|\psi\rangle$ is a minimum-eigenvalue eigenstate of the Hamiltonian $H_{\mathcal{C}}$ given in (2.5). While it may not be immediately obvious at the level of our description, it is possible to prepare such a “history state” (2.4) by executing a quantum circuit that is only mildly more complex than the original circuit \mathcal{C} . The prover is then asked to send the qubits of $|\psi\rangle$ to the verifier one at a time. The verifier secretly selects a random pair (i, j) such that $J_{ij} \neq 0$ ahead of time (without telling the prover), and measures qubits i and j , when she receives them, either both using σ_X , or both using σ_Z . The verifier ignores all other qubits received from the prover. For simplicity, assume all $J_{ij} \in \{-1, 0, 1\}$; it is straightforward to adapt the outline to the general case where $J_{ij} \in [-1, 1]$. Then,

⁹The constants $1/3$ and $2/3$ are a matter of convention, and can be replaced by any constants $0 < s < c < 1$.

if the product of the two outcomes obtained matches the sign of J_{ij} , the verifier accepts; otherwise, she rejects. It is straightforward to check that this procedure accepts any state $|\psi\rangle$ with probability $\frac{1}{2} - \frac{1}{2s} \langle \psi | H_C | \psi \rangle$, where $s = |\{(i, j) : J_{ij} \neq 0\}|$. Repeating the protocol a number of times that scales quadratically with s/δ and accepting if and only if a fraction at least $\frac{1}{2} - \frac{a+\delta/2}{2s}$ accept results in a protocol that accepts valid claims, and rejects erroneous claims, with probability close to 1.

Even though the verifier’s “quantumness” in this protocol is limited — she only needs to hold one qubit at a time — this capability is crucial for the analysis, as it is used to guarantee the “existence” of the state that is being measured: it allows us to meaningfully talk about “the state $|\psi\rangle$ whose first qubit is the first qubit received by the verifier; whose second qubit is the second qubit received by the verifier; etc.”. These qubits are distinct, because the verifier has seen and then discarded them (it would be a different matter if they were returned to the prover). In particular, the fact that a one-qubit computer can be trivially simulated on a classical piece of paper is immaterial to the argument.

With a classical verifier things become substantially more delicate. How can we verify the existence of an n -qubit state with certain properties, while having only access to classical data about the state, data that, for all we know a priori, could have been generated by a simple — classical — laptop? To achieve this we need to find a way for the verifier to establish that the prover holds an n -qubit state, without ever having the ability to directly probe even a single qubit of that state. The major achievement in Mahadev’s work is a method to do just this; it is the topic of the next section.

3. VERIFYING QUANTUM COMPUTATIONS

In the previous section we reduced the problem of verifying the outcome of an arbitrary quantum computation to the following decision problem.

Input: An integer n , the description of an n -qubit Hamiltonian H in XZ form,

$$(3.1) \quad H = - \sum_{1 \leq i < j \leq n} \frac{J_{ij}}{2} (\sigma_X^i \sigma_X^j + \sigma_Z^i \sigma_Z^j),$$

a real number a , and $\delta > 0$.

Promise: The smallest eigenvalue of H is either less than a , or at least $a + \delta$.

Decision: Accept if and only if the smallest eigenvalue of H is less than a . (We refer to such H as “YES” instances.)

The reduction guarantees that the “promise gap” δ can be taken to be at least some inverse polynomial in n . For ease of exposition we will further assume that all coefficients J_{ij} lie in $\{-1, 0, 1\}$, and that $a = -\sum |J_{ij}|$. In physical language this corresponds to a “frustration-free” Hamiltonian, meaning that in the case of a YES instance an eigenstate with smallest eigenvalue a is also an eigenstate with eigenvalue 1 of each of the local terms $h_{ij} = \frac{1}{2}(\sigma_X^i \sigma_X^j + \sigma_Z^i \sigma_Z^j)$. (This last assumption is with loss of generality, as it is not hard to see that the resulting problem lies in NP; nevertheless, we make it because it helps simplify the presentation without hiding any interesting steps.)

Our starting point for deciding this problem is the protocol described at the end of Section 2.4. To remove any quantum computation at the verifier’s side, we

“delegate” the verifier’s single-qubit measurements in that protocol to the prover, who is asked to report the classical measurement outcome. If this procedure were implemented naïvely the prover could easily claim that any Hamiltonian is a YES instance: the first time it is asked to measure a qubit, the prover returns an arbitrary outcome in $\{\pm 1\}$; the second time, it has learned the pair (i, j) the verifier is interested in, and can return an outcome that leads to acceptance (i.e. one that is identical to the first outcome in case $J_{ij} > 0$, and opposite otherwise).

In order to obtain a sound protocol, the verifier faces what may seem like an insurmountable task: ensure that the prover reports measurement outcomes that are consistent with outcomes obtained by measuring the *right* qubits of a *fixed* n -qubit state in the *right basis*. Here “fixed” means that the prover should not change the state it measures depending on the verifier’s question! The main idea we will deploy to achieve this originates in the construction of *commitment protocols* in cryptography. The key ingredient in Mahadev’s verification protocol for quantum computations is a commitment protocol that allows a quantum prover to “commit” to a quantum state using only classical information. Before introducing this, we review the classical notion of commitment, and how it can be achieved using collision-resistant hash functions.

3.1. Classical commitments. Consider the following toy task of “coin-flipping over the phone”: Alice is at work; Bob is at home; they would like to decide over the phone who will cook dinner tonight. Neither volunteers: they need to flip a coin. Clearly neither of them trusts the other to do this properly, so they need a protocol that makes it infeasible for either party to bias the outcome in their favor. Here is a way to achieve this using “commitments”. Bob chooses a value $b \in \{0, 1\}$ — ideally, he chooses it uniformly at random, but this is up to him. He then “commits” to b by sending Alice some information c — think of Bob as inserting a piece of paper with b written on it in a large safe, handing the safe to Alice, but keeping the key to himself. Then, Alice herself chooses a bit $a \in \{0, 1\}$, and announces it directly to Bob. Finally, Bob reveals his bit b by giving Alice the “key” r to the safe. Alice uses r to open the safe and check Bob’s claimed value for b . If the check goes through, they jointly agree that the bit $d = a \oplus b$ is likely to be unbiased. Finally, they use d to designate the night’s cook-in-chief.

Now that we’re convinced of the fundamental need for commitments, let’s discuss how to construct them. The coin-flipping protocol described in the previous paragraph implicitly relies on two requirements: the commitment should be *hiding* (Alice does not learn b , unless Bob explicitly reveals it to her) and *binding* (once he has committed, Bob cannot “reveal” any value other than the one he committed to).

Observe that there does not exist a commitment scheme such that both properties hold with perfect, information-theoretic security. Clearly, if Bob’s commitment c contains no information about b , then the distributions of c conditioned on $b = 0$ and $b = 1$ must be identical; in particular, the scheme couldn’t be binding. Thus achieving commitments requires a limitation on either Alice or Bob’s abilities, i.e. to relax either the hiding or binding property to hold *provided the malicious party, Alice or Bob, is computationally bounded*. It is possible to achieve schemes that are computationally binding and information-theoretically hiding, or vice-versa; motivated by our target application we discuss an example of the former kind of scheme. The scheme relies on the existence of a family of *collision-resistant hash functions*

(CRHF). A CRHF is a family of functions $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^{m(n)}\}$, for $n \geq 1$ and $m(n) < n$, such that for any x , $f_n(x)$ can be evaluated efficiently, but it is computationally infeasible to find a pair of inputs $x \neq x'$ such that $f_n(x) = f_n(x')$.¹⁰ Collision-resistant hash functions are widely used in cryptography, and many constructions are known based on assumptions such as the Diffie-Hellman assumption about hardness of the discrete logarithm problem or the Learning with Errors problem about hardness of solving noisy linear equations. (We sketch a construction based on the latter in Section 4. *Unconditionally* proving the existence of CRHF would imply that $P \neq NP$,¹¹ so we have to rely on computational assumptions.)




FIGURE 6. A computationally binding commitment protocol based on the use of a CRHF family $\{f_n\}$. The symbol $x \leftarrow_R S$ means that x is selected uniformly at random from the finite set S .

Let $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^{n/2}\}$ be a CRHF family. Here is a commitment scheme based on f_n (see also Figure 6). Both parties agree on a “security parameter” $n \geq 1$. To commit to a bit $b \in \{0, 1\}$, Bob selects a uniformly random $r \in \{0, 1\}^{n-1}$ and sends $c = f_n(b||r)$ to Alice, where the symbol $||$ is used to denote string concatenation. To reveal b , Bob sends both b and r to Alice, who checks that $c = f_n(b||r)$. This scheme is computationally binding, because to “change his mind” Bob needs to identify r_0 and r_1 such that $f_n(0||r_0) = f_n(1||r_1)$, which is a collision. Note that the scheme as described is not necessarily hiding, as it is not part of the definition of a CRHF that f_n should hide information about its input. For our purposes this will not matter, as we will be looking for a different property than hiding (in our context we mostly care about protecting the verifier (Alice) against the prover (Bob), not the converse). Nevertheless, we mention that it is possible to achieve a scheme that is information-theoretically hiding using the following small variant. To commit to $b \in \{0, 1\}$, Bob selects uniformly random $(u, x) \in \{0, 1\}^n$ such that $u \cdot x = b$, and sends $c = (c_1, c_2) = (u, f_n(x))$ to Alice. When Bob is asked to reveal b , he sends b and x to Alice, who checks that $c_1 \cdot x = b$ and $c_2 = f_n(x)$. Informally, this variant is information-theoretically hiding because $f_n(x)$ only reveals about $n/2$ bits of information about x , so that to a party that does not know x , the distribution of u , even conditioned on $u \cdot x = b$ for some b , is statistically close to uniform. In other words, the distributions $(u, f_n(x))$ conditioned on $b = 0$ or $b = 1$ are within total variation distance of each other that is exponentially small in n .

¹⁰Formally, given 1^n as input, no randomized polynomial-time procedure can produce a colliding pair (x, x') with more than negligible in n probability. A negligible function $\varepsilon(n)$ is one such that $\varepsilon(n)p(n) \rightarrow_{n \rightarrow \infty} 0$ for any polynomial $p(n)$.

¹¹The converse implication is not known to hold.

3.2. Committing to a qubit. Recall that our goal is to constrain the quantum prover to correctly report the outcomes of measurements on a pair of qubits of an n -qubit state $|\psi\rangle$ that it holds, without being able to change the state or the measurement as a function of the specific qubit(s) that it is asked to measure. Note the analogy to the situation discussed in the previous section: ideally, we would like the prover to “commit” to each of the n qubits of its state $|\psi\rangle$ ahead of time, in a way that any subsequent request to “reveal” the outcome of a measurement of one or more of the qubits in a basis of the verifier’s choice can only be answered by the prover with the correct outcome of the requested measurement on the qubits, lest the prover be caught cheating. This guarantee should hold as long as the prover does not break the computational assumption that breaks the commitment scheme. (The assumption that the scheme described here ultimately rests on is the hardness of the “Learning with Errors” problem, introduced in Section 4.)

Commitments to qubits, using qubits, can be devised using similar ideas as in the classical case. The key innovation that underlies Mahadev’s scheme is a method to “commit” to a qubit using only classical information. Let’s start with the most naïve adaptation of the commitment scheme described in Section 3.1, that would apply it directly, in superposition, to a quantum state. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be an arbitrary function (we will progressively formulate four specific requirements on f , that are stronger than those of a CRHF, so we start with the general case). Starting from an arbitrary single-qubit state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, the prover can compute the classical commitment by first creating an additional register that contains a uniform superposition over all $(n - 1)$ -bit strings and then computing the commitment:

$$(3.2) \quad \begin{aligned} |\psi\rangle = \alpha|0\rangle + \beta|1\rangle &\mapsto (\alpha|0\rangle + \beta|1\rangle) \left(\frac{1}{\sqrt{2^{n-1}}} \sum_{r \in \{0,1\}^{n-1}} |r\rangle \right) |0^n\rangle \\ &\mapsto |\psi'\rangle = \frac{1}{\sqrt{2^{n-1}}} \sum_{r \in \{0,1\}^{n-1}} \alpha|0\rangle|r\rangle|f(0||r)\rangle + \beta|1\rangle|r\rangle|f(1||r)\rangle. \end{aligned}$$

Now suppose that the prover measures the last register in the computational basis to obtain a classical “commitment string” $c \in \{0, 1\}^m$, that it returns to the verifier. Recall from Section 2.2 that any measurement induces a collapse of the state. This is not something that we want to happen in an uncontrolled fashion, because we don’t want the commitment procedure to destroy the committed qubit, that may itself be in a superposition. A collapse can be avoided by ensuring that any measured string c has exactly one preimage of the form $0||r_0$, and one preimage of the form $1||r_1$. So let’s make the assumption that

Assumption (2T01): Both functions $f_0 : r \mapsto f(0||r)$ and $f_1 : r \mapsto f(1||r)$ are injective, and they have the same range.

If this holds then upon having obtained a measurement outcome $c = f(0||r_0) = f(1||r_1)$ for some $r_0, r_1 \in \{0, 1\}^{n-1}$ the state $|\psi'\rangle$ collapses to the post-measurement state consistent with the outcome obtained,

$$(3.3) \quad |\psi''\rangle = (\alpha|0\rangle|r_0\rangle + \beta|1\rangle|r_1\rangle)|c\rangle.$$

This state is interesting. Observe that the initial superposition that defined $|\psi\rangle$ got slightly muddled by the inclusion of r_0 and r_1 . Nevertheless, morally the superposition is preserved: most importantly, the coefficients α, β that define it are

unchanged. In fact, the state $|\psi''\rangle$ is unitarily related to $|\psi\rangle$, by the simple unitary

$$(3.4) \quad U_c : |0\rangle|r_0\rangle \mapsto |0\rangle|0\rangle, \quad |1\rangle|r_1\rangle \mapsto |1\rangle|0\rangle$$

(extended in an arbitrary way to the whole space). However, although this unitary exists, it may not be easy for the prover to implement it! This is because doing so seems to require the identification of both r_0 and r_1 from c , which would require the prover to find a collision for f . Even though we only need it later, let's already make an assumption about f that is known as *collapsing*, a stronger property than being collision resistant.

Assumption (C): Consider the following abstract game between the prover and a trusted (quantum) “challenger”.¹² First, the prover is required to prepare an arbitrary state of the form $|\phi\rangle = \sum_x \alpha_x |x\rangle$, where x ranges over the domain of f . The prover hands the state $|\phi\rangle$ over to the challenger, who evaluates f in superposition on $|\phi\rangle$ and measures the image register, obtaining a c in the range of f and the (suitably normalized) post-measurement state $|\phi'\rangle = \sum_{x:f(x)=c} \alpha_x |x\rangle$. The challenger returns to the prover the string c together with either the state $|\phi'\rangle$ or the probabilistic mixture $\sum_{x:f(x)=c} |\alpha_x|^2 |x\rangle\langle x|$ obtained by measuring the same state $|\phi'\rangle$ in the computational basis (and throwing away the outcome). The prover wins if it correctly guesses which is the case. Assumption (C) on the function f states that no quantum polynomial-time prover can succeed in this game with probability non-negligibly larger than $\frac{1}{2}$.

The reason that Assumption (C) implies collision resistance is that, if the function were not collision resistant, the prover could identify a colliding pair (x_0, x_1) and submit $|\phi\rangle = \frac{1}{\sqrt{2}}(|x_0\rangle + |x_1\rangle)$ to the challenger. It could then measure the challenger's response in a basis containing the two states $\frac{1}{\sqrt{2}}(|x_0\rangle \pm |x_1\rangle)$ and guess that, in case the “-” outcome is obtained, the challenger must have measured; in the other case, the prover guesses at random.

At this point we have completed the description of the commitment phase of the protocol; this phase is summarized in the top part of Figure 7. Before proceeding to a description of the reveal phase, we give a formal definition for a single-qubit state σ to which the prover is “committed” to at the end of the commitment phase. Note that we have to be careful with the meaning that we ascribe to any such definition: for all we know, at this point it could be that the prover has not yet performed any “quantum” actions; in fact, it could even have selected c without performing a single evaluation of the function f . It would in any case be hopeless to make claims about the “true” state of the prover, as the latter may include information that is never made accessible to the verifier.

What matters for our ultimate verification goal is that there *exists* a state that underlies the measurement outcomes revealed by the prover, and that this state is independent of the basis with respect to which the prover is asked to “reveal”. To define such a state we look ahead to the structure of the reveal phase. In the classical protocol described in Figure 6 the phase consists of the prover returning the

¹²This game is not meant to be executed in the protocol; rather, it is meant to indicate a task that the prover should *not* be able to complete.

bit $b \in \{0, 1\}$ and the string $r \in \{0, 1\}^{n-1}$ that allows the verifier to check that the commitment equals to the claimed value b . In the quantum case, the reveal phase comes in two flavors: the “ Z -reveal phase” and the “ X -reveal phase”. Informally, in either phase the prover reveals to the verifier information that allows it to obtain the outcome of a measurement of the committed qubit in the corresponding basis. Since measurements in the computational and Hadamard basis cannot be performed simultaneously, only one of the reveal phases can be executed, at the verifier’s choice.

As we will see, the information requested from the prover in either of the two possible reveal phases is, analogously to the classical case, a pair formed by a single bit followed by an $(n - 1)$ -bit string: n bits in total. We refer to these n bits as the prover’s “reveal string”. In general the prover’s actions can be modeled by two measurements, one for each possible reveal phase. Each of the measurements has outcomes that range in the set of n -bit strings. Since the state of the prover after the commitment phase is only defined up to an arbitrary choice of basis, we may without loss of generality assume that in one of the phases, say the Z -reveal phase, the associated measurement is a direct measurement of n of the prover’s qubits in the computational basis. In other words, we may fix a basis in which the prover’s post-commitment state can be expressed as

$$(3.5) \quad |\tilde{\psi}\rangle = \sum_{b \in \{0,1\}, r \in \{0,1\}^{n-1}} \tilde{\alpha}_{b,r} |b\rangle |r\rangle |\phi_{b,r}\rangle,$$

for arbitrary coefficients $\tilde{\alpha}_{b,r}$ and normalized states $|\phi_{b,r}\rangle$, and such that moreover in the Z -reveal phase the prover directly measures the first n qubits of $|\tilde{\psi}\rangle$ in the computational basis and uses the n -bit outcome (b, r) as its reveal string.

Having fixed a convenient basis for the prover’s measurement in the Z -reveal phase, the measurement that the prover performs to produce its n -bit X -reveal string is expressed as the composition of an arbitrary unitary V acting on the entirety of the prover’s space, followed by a measurement of the first n qubits in the Hadamard basis. Using that $\{\text{Id}, \sigma_X, \sigma_Z, \sigma_X \sigma_Z\}$ form a basis for the space of linear operators on \mathbb{C}^2 , any such unitary has an expansion as

$$V = \text{Id} \otimes V_I + \sigma_X \otimes V_X + \sigma_Z \otimes V_Z + \sigma_X \sigma_Z \otimes V_{XZ}.$$

It is not hard to verify that the linear map defined by

$$(3.6) \quad \tilde{V} : |b\rangle |\phi\rangle \mapsto (|b\rangle V_I |\phi\rangle + (-1)^b |b\rangle V_Z |\phi\rangle) |0\rangle + (|b\rangle V_X |\phi\rangle + (-1)^b |b\rangle V_{XZ} |\phi\rangle) |1\rangle$$

where $|\phi\rangle$ is arbitrary, is an isometry, hence is an admissible operation in quantum mechanics.¹³ (Note that \tilde{V} increases dimension by a factor 2. The new qubit register, in third position, is called a “purifying system”.) We are ready to make a crucial definition.

Definition 3.1 (Committed qubit). Given a commitment string $c \in \{0, 1\}^m$ and an arbitrary post-commitment state for the prover of the form (3.5), let σ be the single-qubit state obtained from $|\tilde{\psi}\rangle$ by applying the isometry \tilde{V} defined in (3.6) and returning the first qubit of the resulting state. We refer to σ as the *committed qubit*.

¹³For the expert, \tilde{V} is obtained from V by a σ_Z -twirl, followed by a conditional σ_X bit-flip. The motivation for this definition will become clear in the analysis of the X -reveal phase.

Note that the verifier does not know the state σ ; in fact, strictly speaking σ is not present on the prover's space at any time. The point is that σ exists, and is well-defined (mathematically) as a function only of the prover's post-commitment state $|\tilde{\psi}\rangle$ and the unitary V .

Remark 3.2. The state σ introduced in Definition 3.1 is a single-qubit state. Eventually we need the prover to commit to an n' -qubit state. This is done by requiring n' commitment strings $c_1, \dots, c_{n'}$. There are still only two reveal phases; in each phase, the prover reports n' n -bit strings that are meant to reveal to the verifier the outcomes obtained by measuring all n' qubits in the same basis, computational or Hadamard. In her paper [12] Mahadev gives a more complicated construction that allows mixing of the two bases (so, there are $2^{n'}$ possible reveal phases). This is not needed here due to the specific form (3.1) of the Hamiltonian we aim to verify.

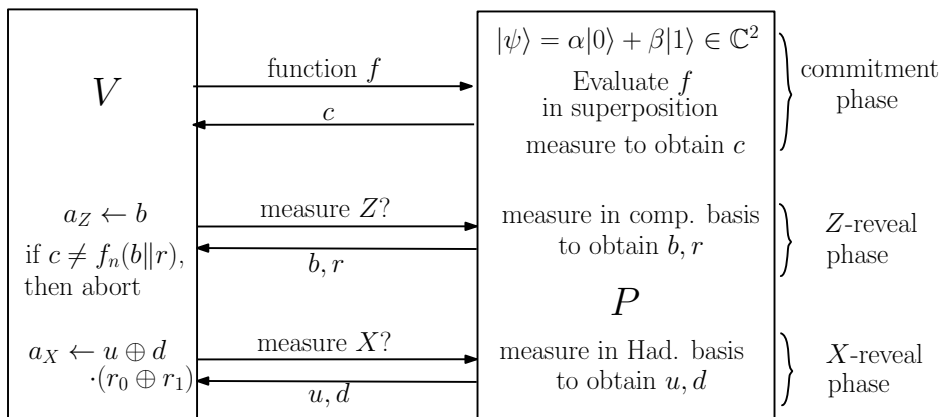


FIGURE 7. Committing to a qubit.

The preceding discussion establishes a definition of a state σ , that may not exist directly on the provers' space at any time in the protocol, but that is explicitly defined from states and operators that are a function of the provers'. To conclude it remains to describe each of the X -reveal and Z -reveal phases, and show how the verifier can extract truthful computational or Hadamard basis measurement outcomes on σ from the prover's reveal string in each phase.

3.2.1. Measuring in the computational basis. We start with a description of the Z -reveal phase. Recall that we made a choice of basis for the prover's space such that its post-commitment state is of the form $|\tilde{\psi}\rangle$ in (3.5), and moreover in the Z -reveal phase the prover directly returns the outcome (b, r) of a measurement of the first n qubits of $|\tilde{\psi}\rangle$ in the computational basis. Having received (b, r) , the verifier records the bit $a_Z = b \in \{0, 1\}$. We call a_Z the verifier's "decoded bit" for the computational basis.

Our goal is to show that the distribution of a_Z is identical to the distribution obtained by measuring the committed qubit σ in the computational basis. According to Definition 3.1, the committed qubit is defined from $|\tilde{\psi}\rangle$ by applying the isometry \tilde{V} defined in (3.6) and returning the first qubit. Observe that \tilde{V} has a block-diagonal form: it stabilizes the spaces $|0\rangle \otimes \mathcal{H}'$ and $|1\rangle \otimes \mathcal{H}'$, where \mathcal{H}' is

the Hilbert space associated with all but the prover's first qubit. As a result the outcome of a measurement of the first qubit of $|\tilde{\psi}\rangle$, or of $\tilde{V}|\tilde{\psi}\rangle$, in the computational basis, are identically distributed: the verifier's decoded bit $a_Z = b$ has exactly the right distribution.

Of course we set things up in this way, essentially defining the committed qubit so that the property holds. The analysis of the X -reveal phase is more challenging. Before turning to it, we add a small test to the Z -reveal phase, whose purpose will become clear later. We call the test the “preimage test”: in this test, the verifier checks that $b||r$ is a preimage of c under f , i.e. that $f(b||r) = c$. In case the test fails, the verifier aborts (irrespective of the value taken by the decoded bit a_Z). Note that this test is identical to the test performed by the verifier in the reveal phase of the classical commitment protocol described in Section 3.1.

An honest prover, whose state $|\tilde{\psi}\rangle$ is the state $|\psi''\rangle$ in (3.3), always passes the preimage test. For the purpose of the analysis of the X -reveal phase given below we make the simplifying assumption that, in case the Z -reveal phase is executed, the prover *always* returns a pair (b, x) that passes the verifier's preimage test.¹⁴ As a consequence, the expression for the state $|\tilde{\psi}\rangle$ simplifies to

$$(3.7) \quad |\tilde{\psi}\rangle = \sum_{b \in \{0,1\}} \tilde{\alpha}_b |b\rangle |r_b\rangle |\phi_b\rangle,$$

where r_0 and r_1 are such that $f(0||r_0) = f(1||r_1) = c$, since using Assumption (2TO1) all other (b, r) would be rejected by the preimage test. This additional assumption requires that in any execution of the commitment protocol, there is a positive probability that the verifier executes the Z -reveal phase, as otherwise the preimage test would never be executed. In the case of the verification protocol, each of the two reveal phases is chosen with probability $1/2$, so that this is not an issue.

3.2.2. Measuring in the Hadamard basis. Similarly to the Z -reveal phase, in the X -reveal phase the verifier expects an X -reveal string $(u, d) \in \{0, 1\} \times \{0, 1\}^{n-1}$ from the prover, supposedly obtained as the outcome of a measurement of the first n qubits of the post-commitment state $|\tilde{\psi}\rangle$ in (3.7) in the Hadamard basis. Upon receiving such a (u, d) from the prover, the verifier records the decoded bit

$$(3.8) \quad a_X = u \oplus d \cdot (r_0 \oplus r_1) \in \{0, 1\}.$$

In addition, in case $d = 0^{n-1}$ the verifier aborts (the purpose for this additional condition will become clear later). To understand the expression (3.8), consider the outcome of a Hadamard measurement on the first n qubits of a state of the form $|\psi''\rangle$ in (3.3), when for example $\alpha = \beta = \frac{1}{\sqrt{2}}$. Applying the H gate on the first n qubits yields (omitting the last register, that contains $|c\rangle$)

$$\begin{aligned} H^{\otimes n} |\psi''\rangle &= \frac{1}{\sqrt{2^{n+1}}} \sum_{b \in \{0,1\}} \left(\sum_{u \in \{0,1\}} (-1)^{ub} |u\rangle \right) \otimes \left(\sum_{d \in \{0,1\}^{n-1}} ((-1)^{d \cdot r_0} + (-1)^{d \cdot r_1}) |d\rangle \right) \\ &= \frac{1}{\sqrt{2^{n-1}}} \sum_{u \in \{0,1\}, d \in \{0,1\}^{n-1}} 1_{d \cdot r_0 = u \oplus (d \cdot r_1)} |u\rangle |d\rangle, \end{aligned}$$

¹⁴More generally, the prover may have a small probability of failure, leading to an error term that needs to be accounted for. This is a technical issue that can be accommodated using standard arguments, so we ignore it here.

so that in this case, for any outcome (u, d) that can be obtained with non-zero probability by the prover, the verifier’s decoded bit is $a_X = u \oplus d \cdot (r_0 \oplus r_1) = 0$, which agrees with the outcome of a measurement of the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ in the Hadamard basis. (Moreover, the probability of obtaining $d = 0^{n-1}$ is exponentially small, so that the honest prover has only a tiny chance of leading to an abort.) Informally, the addition of $d \cdot (r_0 \oplus r_1)$ to the bit u acts as a “decoding” operation that accounts for interference created by the strings r_0, r_1 that have been appended to the prover’s qubit in the commitment phase. Note that this decoding procedure requires the verifier to have the ability to recover the preimages r_0 and r_1 from the prover’s commitment string c . This motivates an additional assumption on the function f :

Assumption (T): There is “trapdoor information” such that, given the trapdoor, it is possible to efficiently invert f , i.e. given c , recover r_0, r_1 such that $f(0||r_0) = f(1||r_1) = c$.

The above argues for “completeness” of the phase, meaning that the verifier’s decoded bit is correct in case the prover performs the intended actions. In general, as discussed earlier the actions of an arbitrary prover can be modeled as the application of a unitary V acting on the entirety of the prover’s space, followed by a measurement of the first n qubits in the Hadamard basis. Our task is then to argue that the verifier’s decoded bit a_X , as defined in (3.8), has the same distribution as the result of a Hadamard measurement of the committed qubit (Definition 3.1). We will show a weaker statement, which is that the two distributions are *computationally indistinguishable*. For distributions on a single bit, the two notions are essentially equivalent, but this would not be the case if we considered measurements on more qubits.¹⁵

We start from the distribution of the verifier’s decoded bit a_X . Recall from Section 2.2 that given a quantum state $\tilde{\rho} = |\tilde{\psi}\rangle\langle\tilde{\psi}|$ the expectation value of an observable O is given by $\text{Tr}(O\tilde{\rho}) = \langle\tilde{\psi}|O|\tilde{\psi}\rangle$. Here the observable O is obtained by first applying V , followed by a Hadamard $H^{\otimes n}$ on the first n qubits, then a measurement in the computational basis of all qubits but the first to obtain the string d (this corresponds to applying the projection $|d\rangle\langle d|$), then a σ_X bit-flip on the first qubit, as a function of the outcome d obtained (this corresponds to applying the unitary $\sigma_X^{d \cdot (r_0 \oplus r_1)}$), and finally measuring the first bit in the computational basis (this corresponds to the observable σ_Z). As a result, the expectation value of $(-1)^{a_X}$ can be expressed as

$$\mathbb{E}[(-1)^{a_X}] = \sum_{d \in \{0,1\}^{n-1}} \langle\tilde{\psi}|V^\dagger H^{\otimes n} ((\sigma_X^{d \cdot (r_0 \oplus r_1)}) \sigma_Z \sigma_X^{d \cdot (r_0 \oplus r_1)}) \otimes |d\rangle\langle d| H^{\otimes n} V |\tilde{\psi}\rangle.$$

The expectation value of $(-1)^b$, where b is the outcome of a measurement of the committed qubit σ in the Hadamard basis, can be expressed similarly except that due to the use of the isometry \hat{V} in the definition of the committed qubit, the

¹⁵This is a limitation of Mahadev’s approach, that is unimportant for the application at hand — recall that the verifier only needs to verify measurements on pairs of qubits at a time — but would be interesting to lift in future work.

unitary V has been conjugated by a random σ_Z operator:¹⁶

$$\begin{aligned} \mathbb{E}[(-1)^b] &= \frac{1}{2} \mathbb{E}[(-1)^{a_X}] \\ &+ \frac{1}{2} \sum_{d \in \{0,1\}^{n-1}} \langle \tilde{\psi} | \sigma_Z V^\dagger \sigma_Z H^{\otimes n} ((\sigma_X^{d \cdot (r_0 \oplus r_1)}) \sigma_Z \sigma_X^{d \cdot (r_0 \oplus r_1)}) \otimes |d\rangle\langle d| H^{\otimes n} \sigma_Z V \sigma_Z | \tilde{\psi} \rangle \\ &= \frac{1}{2} \mathbb{E}[(-1)^{a_X}] \\ &- \frac{1}{2} \sum_{d \in \{0,1\}^{n-1}} \langle \tilde{\psi} | \sigma_Z V^\dagger H^{\otimes n} ((\sigma_X^{d \cdot (r_0 \oplus r_1)}) \sigma_Z \sigma_X^{d \cdot (r_0 \oplus r_1)}) \otimes |d\rangle\langle d| H^{\otimes n} V \sigma_Z | \tilde{\psi} \rangle, \end{aligned}$$

where all σ_Z operators act on the first qubit, and for the second line we used $\sigma_Z H = H \sigma_X$ to commute the innermost σ_Z all the way to the middle, where we simplified $\sigma_X \sigma_Z \sigma_X = -\sigma_Z$. Taking the difference between the two terms, simplifying the middle expression using anti-commutation, and cancelling out cross-terms gives

$$\begin{aligned} &\left| \mathbb{E}[(-1)^{a_X}] - \mathbb{E}[(-1)^b] \right| \\ &= \frac{1}{2} \left| \sum_{d \in \{0,1\}^{n-1}} (-1)^{d \cdot (r_0 \oplus r_1)} (\langle 0, r_0, \phi_0 | V^\dagger H^{\otimes n} (\sigma_Z \otimes |d\rangle\langle d|) H^{\otimes n} V | 0, r_0, \phi_0 \rangle \right. \\ (3.9) \quad &\left. + \langle 1, r_1, \phi_1 | V^\dagger H^{\otimes n} (\sigma_Z \otimes |d\rangle\langle d|) H^{\otimes n} V | 1, r_1, \phi_1 \rangle) \right|, \end{aligned}$$

where to write this last expression we used the assumption that the state $|\tilde{\psi}\rangle$ can be expressed as in (3.7), i.e. the prover succeeds in the verifier's preimage test in the Z -reveal phase with probability 1. To argue that the right-hand side of (3.9) cannot be large, we make the following final assumption.

Assumption (HC): No quantum polynomial-time procedure can, given as input a description of f , return a quadruple (c, r, u, d) such that $f(b||r) = c$ for some $b \in \{0,1\}$, $d \neq 0^{n-1}$, and $u = d \cdot (r_0 \oplus r_1)$, where r_0, r_1 are the two preimages of c , with probability non-negligibly larger than $\frac{1}{2}$.

If the expression on the right-hand side of (3.9) were non-negligible, there would be a violation of Assumption (HC): a quantum polynomial-time ‘‘adversary’’ (to the assumption) could simulate the prover to prepare $|\tilde{\psi}\rangle$, then measure the first n qubits register to obtain (b, r_b) . It would then apply the unitary V and measure the first n qubits in the Hadamard basis to obtain a string (u, d) . Finally, the adversary would return the tuple (c, r, u, d) ; (3.9) exactly measures the correlation of the bit u with the correct value $d \cdot (r_0 \oplus r_1)$.

Thus Assumption (HC) guarantees that the expression in (3.9) is negligibly small, completing the soundness analysis of the Hadamard basis submeasurement protocol: the verifier's decoded bit a_X is negligibly close in distribution to the outcome of a measurement of the committed qubit in the Hadamard basis. (At this point we warn the reader that our argument only considers the case of a single qubit. Extending

¹⁶For the second part of the state on the right-hand side of (3.6), corresponding to the last qubit being in state $|1\rangle$, there is also a missing ‘‘ σ_X ’’ operator on the first qubit, labeled $|b\rangle$. A σ_X has no effect on the outcome of a measurement in the Hadamard basis, so it can be ignored here.

it to the case of n qubits can be done, but requires the use of assumption (C) to argue that the σ_Z operators applied on committed qubits that are eventually not measured by the verifier do not have a noticeable effect on the marginal distribution of outcomes for the two qubits that are considered.)

3.3. Summary. To conclude we summarize the verification protocol introduced progressively in the previous sections. (See also Figure 7.) The input to the protocol is the classical description of an n -qubit Hamiltonian H of the form (3.1). The goal of the verifier is to determine if the smallest eigenvalue of H is less than a , or larger than $a + \delta$. For this she can interact with a prover, that claims to have the ability to prepare an eigenstate $|\psi\rangle$ of H with associated eigenvalue no larger than a . In order to check this claim, the verifier proceeds as follows.

In a first phase, she selects n functions $f_1, \dots, f_n : \{0, 1\}^{n(\lambda)} \rightarrow \{0, 1\}^{m(\lambda)}$ satisfying assumptions (2TO1), (C), (T), and (HC). Here λ is an integer that plays the role of “security parameter”; the larger the λ the more secure the scheme. Generally one may think of λ as being of the same order as n . The verifier sends the public information allowing evaluation of the functions f_i to the prover. The prover replies with “commitment strings” c_1, \dots, c_n . This ends the commitment phase.

In the second phase, the verifier runs either the Z -reveal phase, or the X -reveal phase, each chosen with probability $\frac{1}{2}$. In the case of a Z -reveal phase, the verifier asks the prover for the outcome of a measurement of each of its committed qubits in the computational basis. The prover sends back n pairs (b_i, r_i) , for $i \in \{1, \dots, n\}$. The verifier checks that for each i , $f_i(b_i \| r_i) = c_i$; if not she aborts. Then, the verifier records outcomes $a_{Z,i} = b_i$. In the case of an X -reveal phase, the prover sends back n pairs (u_i, d_i) and the verifier records outcomes $a_{X,i} = u_i \oplus d_i \cdot (r_{0,i} \oplus r_{1,i})$, where $0 \| r_{0,i}$ and $1 \| r_{1,i}$ are the two preimages of c_i under f_i (that the verifier computes using the trapdoor information for f_i).

Finally, the verifier selects a random (i, j) such that $J_{ij} \neq 0$. If $J_{ij} > 0$ and $a_{X,i} = a_{X,j}$ or $a_{Z,i} = a_{Z,j}$ (depending on the basis subprotocol performed), or if $J_{ij} < 0$ and $a_{X,i} \neq a_{X,j}$ or $a_{Z,i} \neq a_{Z,j}$, the verifier accepts. Otherwise, she rejects.

To show that this protocol is sound, we argue that, assuming the prover’s actions lead to a small probability of the verifier aborting, there must exist an n -qubit “committed state” ρ , that can be defined from the provers’ actions, and has the following property: the distribution of any pair of decoded bits recorded by the verifier for measurements in the computational, or Hadamard, basis is negligibly close to the distribution of outcomes obtained by directly measuring the committed state in the appropriate basis. The proof of this follows the outline given in the preceding sections by extending the definition of a committed qubit in Definition 3.1 to n committed qubits in the natural way, and extending the arguments in Section 3.2.1 and Section 3.2.2 to apply to two qubits at a time. A detailed proof would require a more formal treatment than we aim for here; we refer the interested reader to the paper [12].

Our description would not be satisfying if we did not discuss assumptions (2TO1), (C), (T), and (HC). Are these assumptions reasonable? While (2TO1) and (T) are fairly standard assumptions in classical cryptography, for which it is possible to find multiple constructions, Assumption (C) is less common (though it has been used in different contexts in quantum cryptography), and Assumption (HC) is even less usual (though it can be seen as a strengthening of a more standard “(non-adaptive) hardcore bit property”). In the next section we sketch a construction of a function

family satisfying all four assumption simultaneously, based on the computational hardness of the “Learning with Errors” problem in cryptography.

4. A CONSTRUCTION BASED ON THE LEARNING WITH ERRORS PROBLEM

In the previous section we have identified four assumptions on a family of functions $\{f_n : \{0, 1\}^n \rightarrow \{0, 1\}^m\}$ that are sufficient for the resulting verification protocol to be computationally sound. Can the four assumptions be simultaneously satisfied? Strictly speaking, we do not know the answer. In this section we sketch a construction that *nearly* satisfies the assumptions. The construction appears in [4], and a mild modification of it is used in Mahadev’s scheme. Even though the assumptions introduced in the previous section will not all be strictly satisfied by the construction, it is possible to verify that the protocol remains sound.

4.1. The LWE problem. Our starting point is the *Learning with Errors* (LWE) problem, introduced by Regev [15]. The hardness of this problem has become a widely used computational assumption in cryptography, for at least three reasons. The first is that it is very versatile, allowing the implementation of advanced primitives such as fully homomorphic encryption [8], program obfuscation [20], traitor tracing [10], and many others. The second is that the assumption can be reduced to the hardness of *worst-case* computational problems on lattices: an efficient procedure that breaks the LWE assumption *on average* can be used to solve the closest vector problem in (almost) *any* lattice. The third reason, that is most relevant to the use of the LWE assumption for the verification protocol presented here, is that in contrast to the RSA assumption or the discrete logarithm problem so far it is believed that the LWE problem may be hard for quantum computers, so that cryptographic schemes based on it remain (to the best of published knowledge) secure against quantum attacks.

The LWE assumption comes in multiple flavors, all roughly equivalent. Here we formulate the *decisional LWE* assumption on the difficulty of distinguishing samples from two distributions. To state the problem, fix a size parameter $n \geq 1$, an integer modulus $q \geq 2$, a number of equations $m \geq n \log q$, and an error distribution χ over \mathbb{Z}_q . Given χ , write χ^m for the distribution over \mathbb{Z}_q^m that is obtained by sampling each entry of a vector independently according to χ . The decisional LWE assumption is the following.

(Decisional LWE) Let A be a uniformly random matrix in $\mathbb{Z}_q^{m \times n}$, s a uniformly random vector in $\{0, 1\}^n$, e a random vector in \mathbb{Z}_q^m drawn from χ^m , and r a uniformly random vector in \mathbb{Z}_q^m . Then no classical or quantum probabilistic polynomial-time procedure can distinguish $(A, As + e)$ from (A, r) .

We include a few words of explanation for the reader unaccustomed with the notion of computational indistinguishability between ensembles of distributions. Note that the distribution of $(A, As + e)$ and the distribution of (A, r) are in general very far from each other: provided m is sufficiently larger than n a random vector r will not lie in the column span of A , nor even be close to it. What the (decisional) LWE assumption asserts is that, even though in principle these distributions are far from each other, it is computationally difficult, given a sample from the one or the other,

to tell which is the case.¹⁷ Note that without the error vector e the task would be easy: given (A, y) , solve for $As = by$ and check whether the solution has coefficients in $\{0, 1\}$. The LWE assumption is that the inclusion of e makes the task substantially more arduous. In particular, it is well-known that Gaussian elimination is very sensitive to errors, which rules out the most natural approach. To the reader with a geometric mind, it might help to picture a discrete lattice (all integer linear combinations of the columns of A , as a subset of \mathbb{R}^m) such that to each lattice point is added a little noise, in the form of a discrete Gaussian distribution with small variance centered at the lattice point. Even though all the Gaussian “blobs” thus obtained are well separated, given a point in any one of them, it is (conjecturally) hard to recover the center of the blob, i.e. the closest lattice vector.

We comment briefly on the choice of parameters. The integer n should generally be thought of as the security parameter; the larger the more secure (in particular it is always possible to guess s and check validity of the equations, giving an attack in time roughly 2^n). The modulus q should be at least polynomial in n , but can be as large as exponential; this will be the case in our construction. The error distribution χ can be chosen in multiple ways. A common choice is to set χ a discretized centered Gaussian distribution with variance αq , for some small parameter α (typically chosen as an inverse polynomial function of n); this is generally denoted $D_{\mathbb{Z}_q, \alpha q}$. For more details on LWE and its applications, we refer to the survey [16].

4.2. Construction. To specify the function f we describe how public and private parameters for the function are chosen. Let λ be an integer that plays the role of security parameter (i.e. the number 2^λ is thought of as an estimate of the time required to break assumptions such as (HC)).

First, integers n, m and a modulus q are chosen such that $n = \Omega(\lambda)$, $q \geq 2$ is a prime, and $m = \Omega(n \log q)$. Then, a matrix $A \in \mathbb{Z}_q^{m \times n}$ is sampled at random, together with a “trapdoor” in the form of a matrix $R \in \mathbb{Z}_q^{\ell \times m}$, where $n \leq \ell \leq m$ is a parameter. The sampling procedure has the property that the distribution of A is statistically close to uniform, and R is such that $G = RA \in \mathbb{Z}_q^{\ell \times n}$ is a “nice” matrix, in the sense that given $b = Gs + e$, for any $s \in \mathbb{Z}_q^n$ and e small enough, it is computationally easy to recover s .¹⁸ That such a sampling procedure would exist and be efficiently implementable is non-trivial, and relies on the underlying lattice structure given by the columns of A ; see [13]. Finally, a uniformly random $s \in \{0, 1\}^n$, and a random $e \in \mathbb{Z}_q^m$ distributed according to $D_{\mathbb{Z}_q, \alpha q}$ with α of order $1/(\sqrt{mn \log q})$,¹⁹ are sampled. The public information is $(A, y = As + e)$. The private information is the pair (R, s) .

Next, we discuss how the function can be evaluated, given the public parameters (A, y) . We define two functions f_0, f_1 that should be understood as $f(0||\cdot)$ and $f(1||\cdot)$ respectively. For $b \in \{0, 1\}$ the function f_b takes as input an $x \in \mathbb{Z}_q^n$ (that can be seen as an element of \mathbb{Z}_2^{wn} for $w = \lceil \log q \rceil$) and returns $Ax + e' + by$, which is an element of $\mathbb{Z}_q^m \subseteq \mathbb{Z}_2^{wm}$. Here, e' is a vector sampled at random from

¹⁷Computational hardness only makes sense as the input size goes to infinity, which is why to be precise we should consider a family of distributions, parametrized by the integer n , and argue that the samples become harder and harder to distinguish as $n \rightarrow \infty$.

¹⁸One can think of G as a matrix whose rows are almost orthonormal, so that Gaussian elimination on G induces only small propagation of the errors.

¹⁹The precise choice of α is delicate, and the parameters given here should only be treated as indicative; we refer to [4, Section 8] for the right setting of parameters.

a distribution $D_{\mathbb{Z}_q, \alpha'q}$ such that α' is “much larger” than α . The inclusion of e' makes f a “randomized” function, which is the main way in which the construction differs from the requirements expressed in Section 3. A formal way around this is to think of f_b as the function that returns not $Ax + e' + by$, but the *distribution* of $Ax + e' + by$, when $e' \sim D_{\mathbb{Z}_q, \alpha'q}$ and all other variables are fixed. In practice, the evaluation of f on a quantum computer (as required of the “honest” prover in the verification protocol) involves preparing a weighted superposition over all error vectors, and computing the function in superposition.

We would, of course, rather do away with this complication. Why is the error vector necessary? It is there to satisfy the important requirement that the functions f_0 and f_1 are injective with overlapping ranges, i.e. Assumption (2TO1). Injectivity follows from the existence of the trapdoor for A and an appropriate setting of the standard deviation of the error distribution, which guarantee that (given the trapdoor) x can be recovered from $Ax + e' + by$ (with high probability over the choice of e'). To make the function ranges overlap, we need the distribution of $Ax + e'$ to have the same support as the distribution of $Ax' + e' + y = A(x' + s) + (e' + e)$. The first distribution considers an arbitrary vector in the column span of A , shifted by e ; the second considers the same, except that the shift is by $(e' + e)$. For the two distributions to (almost) match, we need the distribution of e' to (almost) match the distribution of $e + e'$. This is possible as long as the standard deviation $\sigma' = \alpha'q$ is substantially larger than the standard deviation $\sigma = \alpha q$; provided this holds it is an exercise to compute the statistical distance between the two Gaussians and verify that it can be made very close to 1.

With this important caveat in place, we have specified the function f , and verified property (2TO1). Property (T) follows from the existence of the secret information (R, s) . Given a $b \in \{0, 1\}$ and an element $c = Ax + e' + by = A(x + bs) + (e' + be)$ in the range of f_b it is possible to use the trapdoor matrix R to recover $x + bs$ and subtract bs to deduce the preimage x of c under f_b .

The two remaining properties, the collapsing property (C) and the hardcore bit property (HC), require more work, and we refer to [4] for a detailed exposition. We remark that the two properties are not entirely new. Property (C) was been introduced by Unruh as a strengthening of the classical property of collision resistance required for his work on the security of commitment protocols that are computationally binding against quantum adversaries [19]. Similar “hardcore bit” properties to (HC) have been shown for many LWE-based cryptographic schemes (see e.g. [3]). Usually the property states that “for any vector $d \in \mathbb{Z}_q^n \setminus \{0\}$, the value $d \cdot s \in \mathbb{Z}_q$ is indistinguishable from uniform, even given a sample $(A, As + e)$ ”. Our property (HC) is subtly stronger, in that the adversary may choose the vector d itself, possibly as a function of the sample $(A, As + e)$. An additional difficulty stems from the specific “bit” that the adversary predicts in our setting. In the definition of Assumption (HC) this bit is the value $u = d \cdot (r_0 \oplus r_1)$, where r_0, r_1 are the *binary representation* of the two preimages in \mathbb{Z}_q^n , x_0 and $x_1 = x_0 - s$, of the prover’s commitment string $c \in \mathbb{Z}_q^m$. (Recall that the use of the binary representation came from the requirements on the honest prover, that is asked to perform a measurement in the Hadamard basis, yielding a binary string of outcomes.) It is in order to complete the argument showing that a procedure that returns the information asked for in Assumption (HC), i.e. the quadruple (c, r, u, d) , can be turned into a procedure that breaks the decisional LWE assumption, that we need

to assume that the secret vector s is a binary vector. The result is a somewhat roundabout construction that we may hope will be simplified in future work.

Acknowledgments. I am indebted to Urmila Mahadev for numerous conversations that helped clarify her work. I thank Alexandru Georghiu, Urmila Mahadev and Oded Regev for comments on earlier versions of these notes.

REFERENCES

1. Dorit Aharonov, Micahel Ben-Or, and Elad Eban, *Interactive Proofs For Quantum Computations*, Arxiv preprint arXiv:0810.5375 (2008).
2. Dorit Aharonov and Ayal Green, *A quantum inspired proof of $P^{\#P} \subseteq IP$* , arXiv preprint arXiv:1710.09078 (2017).
3. Adi Akavia, Shafi Goldwasser, and Vinod Vaikuntanathan, *Simultaneous hardcore bits and cryptography against memory attacks*, Theory of Cryptography Conference, Springer, 2009, pp. 474–495.
4. Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick, *Certifiable randomness from a single quantum device*, arXiv preprint arXiv:1804.00640 (2018).
5. Anne Broadbent, Joseph F. Fitzsimons, and Elham Kashefi, *Universal blind quantum computation*, Arxiv preprint arXiv:0807.4154 (2008).
6. Toby Cubitt and Ashley Montanaro, *Complexity classification of local hamiltonian problems*, SIAM Journal on Computing **45** (2016), no. 2, 268–316.
7. Joseph F Fitzsimons, Michal Hajdušek, and Tomoyuki Morimae, *Post hoc verification of quantum computation*, Physical review letters **120** (2018), no. 4, 040501.
8. Craig Gentry, *A fully homomorphic encryption scheme*, Ph.D. thesis, Stanford University, 2009.
9. Alexandru Gheorghiu, Theodoros Kapourniotis, and Elham Kashefi, *Verification of quantum computation: An overview of existing approaches*, arXiv preprint arXiv:1709.06984 (2017).
10. Rishab Goyal, Venkata Koppula, and Brent Waters, *Collusion resistant traitor tracing from learning with errors*, Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, ACM, 2018, pp. 660–670.
11. Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan, *Algebraic methods for interactive proof systems*, Journal of the ACM (JACM) **39** (1992), no. 4, 859–868.
12. Urmila Mahadev, *Classical verification of quantum computations*, arXiv preprint arXiv:1804.01082 (2018).
13. Daniele Micciancio and Chris Peikert, *Trapdoors for lattices: Simpler, tighter, faster, smaller*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2012, pp. 700–718.
14. Ran Raz and Avishay Tal, *Oracle separation of BQP and PH*, Electronic Colloquium on Computational Complexity (ECCC), vol. 25, 2018, p. 107.
15. Oded Regev, *On lattices, learning with errors, random linear codes, and cryptography*, Journal of the ACM (JACM) **56** (2009), no. 6, 34.
16. ———, *The learning with errors problem*, Invited survey in CCC **7** (2010).
17. Ben W Reichardt, Falk Unger, and Umesh Vazirani, *Classical command of quantum systems*, Nature **496** (2013), no. 7446, 456.
18. A. Shamir, *IP= PSPACE*, Journal of the ACM (JACM) **39** (1992), no. 4, 869–877.
19. Dominique Unruh, *Computationally binding quantum commitments*, Annual International Conference on the Theory and Applications of Cryptographic Techniques, Springer, 2016, pp. 497–527.
20. Daniel Wichs and Giorgos Zirdelis, *Obfuscating compute-and-compare programs under LWE*, 2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2017, pp. 600–611.

CALIFORNIA INSTITUTE OF TECHNOLOGY, PASADENA CA 91106, USA

E-mail address: vidick@cms.caltech.edu

THE SHUFFLE CONJECTURE

STEPHANIE VAN WILLIGENBURG

On the occasion of Adriano Garsia's 90th birthday

ABSTRACT. Walks in the plane taking unit-length steps north and east from $(0, 0)$ to (n, n) never dropping below $y = x$, and parking cars subject to preferences, are two intriguing ingredients in a formula conjectured in 2005, now famously known as the shuffle conjecture.

Here we describe the combinatorial tools needed to state the conjecture. We also give key parts and people in its history, including its eventual algebraic solution by Carlsson and Mellit, which was published in the Journal of the American Mathematical Society in 2018. Finally, we conclude with some remaining open problems.

They can see the topography ...
the treetops, but we can see the
parakeets.

Adriano Garsia

Often, in order to delve deep into the structure of an abstract mathematical construct, the treetops, we need to interpret it concretely with a combinatorial visualization, the parakeets. The shuffle conjecture, as we will see, is one such story. In this article we will integrate the motivation, history and mathematics of the shuffle conjecture as we proceed. Hence we will begin by recalling necessary concepts from combinatorics in Section 1, and from algebra in Section 2, in order to state the shuffle conjecture in Theorem 3.2. This recently proved conjecture is, in essence, a formula for encoding the graded dimensions of the symmetric group representation, in the character of a particular vector space on which the symmetric group \mathfrak{S}_n acts. In Section 3 we also discuss some of the motivation and history of the shuffle conjecture, including its refinement known as the compositional shuffle conjecture whose algebraic resolution by Carlsson and Mellit, announced in 2015 [5] and published in 2018 [6], rocked the combinatorial community. We mention some of their proof ingredients in Section 4, where we also conclude with some future avenues.

1. THE COMBINATORICS OF DYCK PATHS AND PARKING FUNCTIONS

A crucial concept for the statement of the shuffle conjecture is that of parking functions. Although originally studied by Pyke [25], they were introduced as a model for parking n cars subject to preferences by Konheim and Weiss who were


2010 *Mathematics Subject Classification.* Primary 05E05, 05E10, 20C30.

The author was supported in part by the National Sciences and Engineering Research Council of Canada, and in part by funding from the Simons Foundation and the Centre de Recherches Mathématiques, through the Simons-CRM scholar-in-residence program.

studying data storage [21, Section 6: A parking problem – the case of the capricious wives]. Konheim and Weiss also proved that the number of parking functions involving n cars is $(n+1)^{(n-1)}$. Since then these functions have arisen in a plethora of places from hyperplane arrangements [31] to chip-firing [8]. More details on parking functions can be found in, for example, the survey by Yan [34]. Rather than using the original definition, in terms of drivers parking cars, we will instead use an equivalent definition introduced by Garsia, for example in his paper with Haiman [10, p 227]. However, before we do this, we need to define a Dyck path.

Definition 1.1 (Dyck path). A *Dyck path* of order n is a path in the $n \times n$ lattice from $(0,0)$ to (n,n) that consists of n unit-length north steps and n unit-length east steps, which stays weakly above the line $y = x$.

Example 1.2. If we let N denote a unit-length north step, and E denote a unit-length east step, then the following path $NNNEEENNENEENNEE$ from $(0,0)$ in the bottom-left corner to $(8,8)$ in the top-right corner is a Dyck path of order 8.



Definition 1.3 (parking function). A *parking function* of order n is a Dyck path of order n such that each north step has a label, called a *car*, written in the square to its immediate right. The cars are $1, 2, \dots, n$, each occurring exactly once, and cars in the same column increase from bottom to top. We denote the set of all parking functions of order n by PF_n .

Example 1.4. An example of a parking function, which we will use throughout this article, is given in Figure 1.

We now define three statistics on parking functions that will be useful later, the first of which is the area of a parking function.

Definition 1.5 (area). If π is a parking function, then its *area* is the number of complete squares between the Dyck path of π and $y = x$, denoted by $\text{area}(\pi)$.

Example 1.6. If π is the parking function from Figure 1, then by counting the number of complete squares in each row contributing to the area, from bottom to top, we get

$$\text{area}(\pi) = 0 + 1 + 2 + 0 + 1 + 1 + 0 + 1 = 6.$$

The second statistic is slightly more intricate than the area.

Definition 1.7 (dinv). Consider a parking function π , and a pair of cars $\{c_1, c_2\}$ in it.




FIGURE 1. A parking function of order 8

- If the cars c_1, c_2 are in the **same** diagonal (that is, their squares are the same distance from $y = x$) with the larger car occurring further right, then $\{c_1, c_2\}$ is a *primary diagonal inversion*. Let $\text{Dinv}^{pri}(\pi)$ be the set of all such pairs.
- If the cars c_1, c_2 are in **adjacent** diagonals with the larger car occurring in the higher diagonal (that is, its square is distance 1 further from $y = x$ than that of the smaller car) and further left, then $\{c_1, c_2\}$ is a *secondary diagonal inversion*. Let $\text{Dinv}^{sec}(\pi)$ be the set of all such pairs.

Then

$$\text{dinv}(\pi) = |\text{Dinv}^{pri}(\pi)| + |\text{Dinv}^{sec}(\pi)|.$$

Example 1.8. If π is from Figure 1, then $\{3, 7\}$ is a primary diagonal inversion, but $\{5, 7\}$ is not, since the smaller car 5 occurs further right. Likewise $\{5, 8\}$ is a secondary diagonal inversion, but $\{3, 4\}$ is not, since the smaller car 3 occurs in the higher diagonal and further left. Note that $\{4, 8\}$ is neither type of diagonal inversion since the cars are not in the same or adjacent diagonals.

Hence,

$$\begin{aligned} \text{Dinv}^{pri}(\pi) &= \{\{2, 4\}, \{3, 6\}, \{3, 7\}, \{3, 5\}, \{6, 7\}\} \\ \text{Dinv}^{sec}(\pi) &= \{\{1, 3\}, \{1, 6\}, \{1, 7\}, \{6, 8\}, \{7, 8\}, \{5, 8\}\} \end{aligned}$$

so

$$\text{dinv}(\pi) = |\text{Dinv}^{pri}(\pi)| + |\text{Dinv}^{sec}(\pi)| = 5 + 6 = 11.$$

Our third statistic is a permutation associated to a parking function.

Definition 1.9 (word). If π is a parking function, then its *word* is the permutation in one-line notation obtained by reading cars from the diagonal furthest from $y = x$ to the diagonal $y = x$, and within a diagonal reading from right to left. We denote this by $\text{word}(\pi)$.

Example 1.10. If π is from Figure 1, then

$$\text{word}(\pi) = 85763142.$$

With our three statistics now defined, we end this section by recalling the i -descent set of a permutation, in our case specialized to the word of a parking function.

Definition 1.11 (*ides*). If π is a parking function, then its i -descent set is

$$\text{ides}(\pi) = \{i \mid i + 1 \text{ is left of } i \text{ in } \text{word}(\pi)\}.$$

Example 1.12. If π is from Figure 1 with $\text{word}(\pi) = 85763142$ from Example 1.10, then

$$\text{ides}(\pi) = \{2, 4, 6, 7\}.$$

2. THE ALGEBRAS OF QUASISYMMETRIC AND SYMMETRIC FUNCTIONS

We now start to turn our attention to the algebraic ingredients needed to state the shuffle conjecture after first recalling the notions of compositions and partitions.

A *composition* α of n , denoted by $\alpha \vDash n$, is a list of positive integers $\alpha = \alpha_1 \alpha_2 \cdots \alpha_{\ell(\alpha)}$ such that $\sum_{i=1}^{\ell(\alpha)} \alpha_i = n$. We call the α_i the *parts* of α , call n the *size* of α and call $\ell(\alpha)$ the *length* of α . If, furthermore, $\alpha_1 \geq \alpha_2 \geq \cdots \geq \alpha_{\ell(\alpha)}$, then we say that α is a *partition* of n , and denote this by $\alpha \vdash n$. For example, 332 is both a composition and partition, with size 8 and length 3.

Now we focus on defining the algebra of quasisymmetric functions, before using them to define the algebra of symmetric functions.

The algebra of quasisymmetric functions, QSym , is a subalgebra of $\mathbb{C}[[z_1, z_2, \dots]]$, meaning that QSym is a vector space, over \mathbb{C} , of formal power series in the variables z_1, z_2, \dots , in which we can also multiply the elements together. A basis for it is given by the set of all fundamental quasisymmetric functions that we now define in the variables $Z = \{z_1, z_2, \dots\}$, indexed by subsets of $[n - 1] = \{1, 2, \dots, n - 1\}$.

Definition 2.1 (*fundamental quasisymmetric function*). Let $S = \{s_1, s_2, \dots, s_{|S|}\} \subseteq [n - 1]$. Then the *fundamental quasisymmetric function* $F_{n,S}$ is defined to be

$$F_{n,S} = \sum z_{i_1} z_{i_2} \cdots z_{i_n}$$

where the sum is over all n -tuples (i_1, i_2, \dots, i_n) satisfying

$$i_1 \leq i_2 \leq \cdots \leq i_n \text{ and } i_j < i_{j+1} \text{ if } j \in S.$$

Example 2.2. We have that

$$F_{3,\{1\}} = z_1 z_2^2 + z_1 z_3^2 + z_2 z_3^2 + \cdots + z_1 z_2 z_3 + z_1 z_2 z_4 + \cdots$$

whereas

$$F_{3,\{2\}} = z_1^2 z_2 + z_1^2 z_3 + z_2^2 z_3 + \cdots + z_1 z_2 z_3 + z_1 z_2 z_4 + \cdots.$$

Quasisymmetric functions were first mentioned implicitly in Stanley's thesis, with regard to P -partitions, published in 1972 [30], and then Gessel developed and published much of the classical theory explicitly in 1984 [11]. Since then they have arisen in a variety of areas, for example, from probability [18] to category theory [1]. However, our interest lies in a special case of a result from Gessel's original paper [11, Theorem 3]. For this we first need to define Young diagrams and Young tableaux.

Given a partition $\lambda = \lambda_1 \lambda_2 \cdots \lambda_{\ell(\lambda)} \vdash n$, we define its *Young diagram*, also denoted by λ , to be the array of left-justified boxes with λ_i boxes in row i from the top. Given the Young diagram, λ , a *standard Young tableau (SYT)* of *shape* λ ,

T , is a filling of the n boxes of λ with $1, 2, \dots, n$ each appearing exactly once such that the entries in the rows increase when read from left to right, and the entries in each column increase when read from top to bottom. We denote the set of all SYTs of shape λ by $SYT(\lambda)$.

Example 2.3. We have that $T = \begin{array}{|c|c|c|c|} \hline 1 & 3 & 4 & 5 \\ \hline 2 & 6 & 8 & \\ \hline 7 & & & \\ \hline \end{array}$ is an SYT of shape $431 \vdash 8$.

Given an SYT, T , of shape $\lambda \vdash n$, we define its *descent set* to be

$$\text{Des}(T) = \{i \mid i + 1 \text{ is in the same column or left of } i\} \subseteq [n - 1].$$

Example 2.4. If T is from Example 2.3, then

$$\text{Des}(T) = \{1, 5, 6\} \subseteq [7].$$

We can now define the algebra of symmetric functions, Sym , which is a subalgebra of QSym . This algebra is so called because its elements are invariant under any permutation of its variables, and a basis for it is the set of all Schur functions that we now define as a special case of [11, Theorem 3].

Definition 2.5. Let $\lambda \vdash n$. Then the *Schur function* s_λ is defined to be

$$s_\lambda = \sum_{T \in SYT(\lambda)} F_{n, \text{Des}(T)}.$$

Example 2.6. We have that $s_{21} = F_{3, \{1\}} + F_{3, \{2\}}$ from the SYTs below.

$$\begin{array}{|c|c|} \hline 1 & 3 \\ \hline 2 & \\ \hline \end{array} \quad \begin{array}{|c|c|} \hline 1 & 2 \\ \hline 3 & \\ \hline \end{array}$$

The Schur functions are not the only basis of Sym . Another basis that will be vital to our story is the basis consisting of all elementary symmetric functions: We define the *i -th elementary symmetric function* e_i to be

$$e_i = s_{1^i}$$

where 1^i is the partition consisting of i parts equal to 1. Then if $\lambda = \lambda_1 \lambda_2 \cdots \lambda_{\ell(\lambda)} \vdash n$ we define the *elementary symmetric function* e_λ to be

$$e_\lambda = e_{\lambda_1} e_{\lambda_2} \cdots e_{\lambda_{\ell(\lambda)}}.$$

Similarly, there exists the basis consisting of all complete homogeneous symmetric functions: We define the *i -th complete homogeneous symmetric function* h_i to be

$$h_i = s_i.$$

Then if $\lambda = \lambda_1 \lambda_2 \cdots \lambda_{\ell(\lambda)} \vdash n$ we define the *complete homogeneous symmetric function* h_λ to be

$$h_\lambda = h_{\lambda_1} h_{\lambda_2} \cdots h_{\lambda_{\ell(\lambda)}}.$$

Symmetric functions date back to Girard [12] in 1629, although Schur functions are much younger, dating to an 1815 paper of Cauchy [7]. The Schur functions were named after Schur who in 1901 proved that they were characters of the irreducible polynomial representations of the general linear group [27], while standard Young tableaux were defined by Young in his 1928 publication [35, p 258]. Substantial historical notes on this subject can be found in Stanley's second volume

on enumerative combinatorics [32, Chapter 7], which is also an excellent resource for symmetric functions and some related representation theory, as is the book by Sagan [26].

3. THE SPACE OF DIAGONAL HARMONICS AND THE SHUFFLE CONJECTURE

With our essential combinatorial and algebraic notations now defined, we can begin to work towards our statement of the shuffle conjecture, which is about the vector space of diagonal harmonics. However, before we do that, let us define our desired space.

Let $X_n = \{x_1, x_2, \dots, x_n\}$ and $Y_n = \{y_1, y_2, \dots, y_n\}$. Then the *space of diagonal harmonics*, DH_n , is the vector space of polynomials in these variables, $f(X_n, Y_n)$, which satisfy

$$(3.1) \quad \partial_{x_1}^a \partial_{y_1}^b f(X_n, Y_n) + \partial_{x_2}^a \partial_{y_2}^b f(X_n, Y_n) + \dots + \partial_{x_n}^a \partial_{y_n}^b f(X_n, Y_n) = 0$$

for all $a, b \geq 0$ and $a + b > 0$. That is,

$$DH_n = \{f(X_n, Y_n) \in \mathbb{C}[X_n, Y_n] \mid \sum_{i=1}^n \partial_{x_i}^a \partial_{y_i}^b f(X_n, Y_n) = 0, \forall a, b \geq 0, a + b > 0\}.$$

Example 3.1. DH_2 consists of all polynomials $f(X_2, Y_2) = f(x_1, x_2, y_1, y_2)$ such that

$$\begin{array}{llll} a + b = 1 & \text{gives} & \partial_{x_1} f(X_2, Y_2) + \partial_{x_2} f(X_2, Y_2) = 0 & \text{when } a = 1 \quad b = 0 \\ & & \partial_{y_1} f(X_2, Y_2) + \partial_{y_2} f(X_2, Y_2) = 0 & a = 0 \quad b = 1 \\ a + b = 2 & \text{gives} & \partial_{x_1}^2 f(X_2, Y_2) + \partial_{x_2}^2 f(X_2, Y_2) = 0 & \text{when } a = 2 \quad b = 0 \\ & & \partial_{x_1} \partial_{y_1} f(X_2, Y_2) + \partial_{x_2} \partial_{y_2} f(X_2, Y_2) = 0 & a = 1 \quad b = 1 \\ & & \partial_{y_1}^2 f(X_2, Y_2) + \partial_{y_2}^2 f(X_2, Y_2) = 0 & a = 0 \quad b = 2 \end{array}$$

etc, and we can check that the solution set has basis $\{1, x_1 - x_2, y_1 - y_2\}$.

The symmetric group, \mathfrak{S}_n , acts naturally on DH_n by the *diagonal action* that permutes the X_n and Y_n variables simultaneously. Namely, given $\sigma \in \mathfrak{S}_n$ and $f(X_n, Y_n) = f(x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n)$ we have that

$$\sigma f(x_1, x_2, \dots, x_n, y_1, y_2, \dots, y_n) = f(x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(n)}, y_{\sigma(1)}, y_{\sigma(2)}, \dots, y_{\sigma(n)}).$$

By Equation (3.1) we see that if $f(X_n, Y_n) \in DH_n$, then $\sigma f(X_n, Y_n) \in DH_n$. Furthermore, if we let $DH_n^{c,d}$ be the subspace of DH_n whose elements have total degree c in the variables x_1, x_2, \dots, x_n , and total degree d in the variables y_1, y_2, \dots, y_n , then if $f(X_n, Y_n) \in DH_n^{c,d}$, then $\sigma f(X_n, Y_n) \in DH_n^{c,d}$. This enables us to define the *bigraded Frobenius characteristic* of DH_n

$$(3.2) \quad DH_n[Z; q, t] = \sum_{c,d \geq 0} t^c q^d \sum_{\lambda \vdash n} s_\lambda \text{Mult}(\chi^\lambda, DH_n^{c,d})$$

where, as before s_λ is a Schur function in the variables $Z = \{z_1, z_2, \dots\}$ and $\text{Mult}(\chi^\lambda, DH_n^{c,d})$ is the multiplicity of the irreducible character of \mathfrak{S}_n , χ^λ , in the character of $DH_n^{c,d}$ under the diagonal action of \mathfrak{S}_n .

3.1. The shuffle conjecture. The shuffle conjecture is a combinatorial formula for computing $DH_n[Z; q, t]$ in Equation (3.2), but before we give it and do an example we will briefly recount a skeletal history that motivated it. More details on this fascinating story can be found in the excellent state-of-the-art survey article by Hicks [19], and the illuminating texts by Bergeron [3] and Haglund [13].

In 1988 Kadell looked for [20] and then Macdonald found [24] a generalization of Schur functions, with additional parameters q, t , $P_\lambda[Z; q, t]$ where $\lambda \vdash n$. This generalization specialized to the Schur functions at $q = t$, and to other well-known functions such as the elementary symmetric functions, Hall-Littlewood functions, and Jack symmetric functions, which were likewise recovered by setting q and t to various values. These functions were then transformed by Garsia and Haiman [10, p 194], thus creating modified Macdonald polynomials $\tilde{H}_\lambda[Z; q, t]$. At the same time they were studying $DH_n[Z; q, t]$ and conjectured a formula for it in terms of the $\tilde{H}_\lambda[Z; q, t]$. Bergeron and Garsia noted that this formula was almost identical to the formula for the elementary symmetric functions e_n in terms of $\tilde{H}_\lambda[Z; q, t]$. More precisely, if the coefficient of $\tilde{H}_\lambda[Z; q, t]$ in e_n was \mathcal{C}_λ , then its conjectured coefficient in $DH_n[Z; q, t]$ was

$$t^{n(\lambda)} q^{n(\lambda')} \mathcal{C}_\lambda$$

where if $\lambda = \lambda_1 \lambda_2 \cdots \lambda_{\ell(\lambda)}$, then $n(\lambda) = \sum_{i=1}^{\ell(\lambda)} \lambda_i(i-1)$ and $\lambda' = \lambda'_1 \lambda'_2 \cdots \lambda'_{\ell(\lambda')}$ is the *transpose* of λ , which is the partition created from λ by setting

$$\lambda'_i = \text{number of parts of } \lambda \text{ that are } \geq i.$$

For example, if $\lambda = 211$, then $\lambda' = 31$. This inspired Bergeron to officially define the *nabla operator* in the paper with Garsia [4, Equation (4.11)] as follows.

$$\nabla \tilde{H}_\lambda[Z; q, t] = t^{n(\lambda)} q^{n(\lambda')} \tilde{H}_\lambda[Z; q, t]$$

Hence when Haiman, using algebraic geometry, proved the conjectured formula for $DH_n[Z; q, t]$ [17, Theorem 3.2] this automatically yielded that [17, Proposition 3.5]

$$(3.3) \quad DH_n[Z; q, t] = \nabla e_n$$

since from above

$$e_n = \sum_{\lambda \vdash n} \mathcal{C}_\lambda \tilde{H}_\lambda[Z; q, t]$$

and now it was proved that

$$DH_n[Z; q, t] = \sum_{\lambda \vdash n} t^{n(\lambda)} q^{n(\lambda')} \mathcal{C}_\lambda \tilde{H}_\lambda[Z; q, t].$$

Haiman had also proved [17, Proposition 3.6] that


$$\dim(DH_n) = (n+1)^{(n-1)}.$$

This supported the search for a collection of $(n+1)^{(n-1)}$ objects, such as all parking functions of order n , along with statistics on them, in order to find a formula to compute ∇e_n more easily. The shuffle conjecture of Haglund, Haiman, Loehr, Remmel and Ulyanov [14, Conjecture 3.1.2] conjectured such a formula, which we give now. This conjecture was proved recently, as a consequence of proving a refinement of it called the compositional shuffle conjecture, by Carlsson and Mellit [6, Theorem 7.5]. However, many still refer to it as the shuffle conjecture, and hence we will too.

Theorem 3.2 (the shuffle conjecture).

$$\nabla e_n = \sum_{\pi \in PF_n} t^{\text{area}(\pi)} q^{\text{dinv}(\pi)} F_{n, \text{ides}(\pi)}$$

Example 3.3. Let us compute $n = 2$. In order to compute ∇e_2 , we first need to calculate the elements of PF_2 that are as follows.



They have

$\text{area}(\pi^{(1)}) = 1$	$\text{dinv}(\pi^{(1)}) = 0$
$\text{area}(\pi^{(2)}) = 0$	$\text{dinv}(\pi^{(2)}) = 0$
$\text{area}(\pi^{(3)}) = 0$	$\text{dinv}(\pi^{(3)}) = 1$
$\text{word}(\pi^{(1)}) = 21$	$\text{ides}(\pi^{(1)}) = \{1\}$
$\text{word}(\pi^{(2)}) = 12$	$\text{ides}(\pi^{(2)}) = \emptyset$
$\text{word}(\pi^{(3)}) = 21$	$\text{ides}(\pi^{(3)}) = \{1\}$


and hence

$$\nabla e_2 = tF_{2, \{1\}} + F_{2, \emptyset} + qF_{2, \{1\}} = F_{2, \emptyset} + (q+t)F_{2, \{1\}}.$$

By Equation (3.3) and the definition of $DH_n[Z; q, t]$ in Equation (3.2) we know that ∇e_2 can be written as a positive linear combination of Schur functions. Using Definition 2.5 we have that

$$s_2 = F_{2, \emptyset} \quad \text{and} \quad s_{11} = F_{2, \{1\}}$$

from the respective SYTs



and hence

$$\nabla e_2 = s_2 + (q+t)s_{11}.$$

It is still an open problem to find a formula for ∇e_n that is a manifestly positive linear combination of Schur functions.

We conclude this subsection with an indication of why the shuffle conjecture was so called. The name arose because the coefficient of the monomial $z_1^{\lambda_1} z_2^{\lambda_2} \cdots z_{\ell(\lambda)}^{\lambda_{\ell(\lambda)}}$ in ∇e_n is equal to [14, Corollary 3.3.1]

$$\sum t^{\text{area}(\pi)} q^{\text{dinv}(\pi)}$$

where the sum is over all $\pi \in PF_n$ such that $\text{word}(\pi)$ is a *shuffle* of the lists

$$[1, 2, \dots, \lambda_1], [\lambda_1 + 1, \lambda_1 + 2, \dots, \lambda_1 + \lambda_2], \dots, [m + 1, m + 2, \dots, n]$$

where $m = \sum_{i=1}^{\ell(\lambda)-1} \lambda_i$, that is, within $\text{word}(\pi)$ the numbers within each list appear in order when $\text{word}(\pi)$ is read from left to right.

Example 3.4. Given the lists $[1, 2]$ and $[3, 4]$ note that 1342 is a shuffle of the lists, but 1432 is not since **3** and **4** are not in order.

3.2. The compositional shuffle conjecture. The conjecture that Carlsson and Mellit proved was not the shuffle conjecture from the previous subsection, but rather a refinement of it known as the compositional shuffle conjecture. This refinement by Haglund, Morse and Zabrocki [15, Conjecture 4.5] centred around further symmetric functions C_α , where $\alpha \vDash n$, that satisfy

$$e_n = \sum_{\alpha \vDash n} C_\alpha$$

so that

$$(3.4) \quad \nabla e_n = \sum_{\alpha \vDash n} \nabla C_\alpha$$

and involved a fourth statistic on parking functions, that of a touch composition.

Definition 3.5 (touch). If π is a parking function of order n , then note the set of row numbers from bottom to top where there is a car in a square on the diagonal $y = x$

$$\{i_1 = 1, i_2, \dots, i_k\}.$$

Then the *touch composition* is

$$\text{touch}(\pi) = (i_2 - i_1)(i_3 - i_2) \cdots (n + 1 - i_k).$$

Example 3.6. If π is from Figure 1, then the set of row numbers where there is a car in a square on $y = x$ is $\{1, 4, 7\}$ and hence

$$\text{touch}(\pi) = 332.$$


We can now state the compositional shuffle conjecture of Haglund, Morse and Zabrocki [15, Conjecture 4.5], which was proved by Carlsson and Mellit [6, Theorem 7.5].

Theorem 3.7 (the compositional shuffle conjecture). *Let $\alpha \vDash n$.*

$$\nabla C_\alpha = \sum_{\substack{\pi \in PF_n \\ \text{touch}(\pi) = \alpha}} t^{\text{area}(\pi)} q^{\text{dinv}(\pi)} F_{n, \text{idess}(\pi)}$$

Observe that proving this would immediately prove the shuffle conjecture since if we sum over all $\alpha \vDash n$, then the left-hand side would yield ∇e_n by Equation (3.4) and the right-hand side would lose its touch composition restriction.

Example 3.8. Let us compute $n = 2$. From Example 3.3 we have that the elements of PF_2 are again as follows.



They have

$$\begin{aligned} \text{touch}(\pi^{(1)}) &= 2 \\ \text{touch}(\pi^{(2)}) &= 11 \\ \text{touch}(\pi^{(3)}) &= 11 \end{aligned}$$

hence

$$\begin{aligned}\nabla C_2 &= tF_{2,\{1\}} \\ \nabla C_{11} &= F_{2,\emptyset} + qF_{2,\{1\}}\end{aligned}$$

and from Example 3.3

$$\nabla e_2 = tF_{2,\{1\}} + F_{2,\emptyset} + qF_{2,\{1\}} = \nabla C_2 + \nabla C_{11}.$$

4. THE PROOF AND FURTHER DIRECTIONS

On 25 August 2015 Carlsson and Mellit posted an article on the arXiv [5] titled simply “A proof of the shuffle conjecture”, in which they proved the compositional shuffle conjecture, which in turn proved the shuffle conjecture. In their proof they refined the compositional shuffle conjecture yet further and proved this refinement.

They worked with the right-hand side of the compositional shuffle conjecture under what is known as the ζ map, which takes a parking function π to a new Dyck path with cars placed in the squares along $y = x$ such that when the cars are read from right to left we obtain $\text{word}(\pi)$. This required them to develop an analogue of touch that they called touch' . They also worked with the reverse ordering of cars, so that, for example, in a parking function the cars in the same column *decrease* when read from bottom to top. The list of other ingredients that they were required to create is impressive and included a generalization of the double affine Hecke algebra; partial Dyck paths; numerous operators including raising and lowering operators involving Hecke algebra operators and plethysm, and a modification of Demazure-Lusztig operators; and a recurrence that their refinement satisfied.

To give a further idea of the complexity of the proof, this proof was almost 30 pages in length and took Haglund a full semester to check. In order to make it more accessible to combinatorialists, at the request of Garsia, Haglund and Xin expanded the proof, and their resulting article [16] is 60 pages in length.

4.1. Further directions. Carlsson and Mellit’s proof of the shuffle conjecture was published in the Journal of the American Mathematical Society in 2018 [6], but there remain many related open problems, some of which we now conclude with.

- (1) **A Schur-positive formula for ∇e_n** By Equation (3.3) and Equation (3.2) we know that when we express ∇e_n as a linear combination of Schur functions

$$\nabla e_n = \sum_{c,d \geq 0} t^c q^d \sum_{\lambda \vdash n} \mathcal{D}_\lambda s_\lambda$$

we have that the coefficients \mathcal{D}_λ must be nonnegative integers since they are counting multiplicities. It remains an open problem to find a *combinatorial* formula for the \mathcal{D}_λ , namely a formula that would compute them directly as nonnegative integers by counting a set of objects.

- (2) **Nabla on other symmetric functions** While the search for a combinatorial formula for ∇e_n has now been concluded with the proof of the shuffle conjecture, it remains to prove the formula of Loehr and Warrington [23, Conjecture 2.1] for

$$\nabla s_\lambda$$

as the formula would generalize the result for ∇e_n since $e_n = s_{1^n}$. However, a conjecture of Loehr and Warrington [22, p 667] for

$$\nabla p_n$$

where p_n is the n -th power sum symmetric function

$$p_n = z_1^n + z_2^n + \cdots$$

was recently proved by Sergel [28, Theorem 4.11] who has also conjectured the existence of a formula [29, Conjecture 3.1] for

$$\nabla m_\lambda$$

where m_λ is the monomial symmetric function

$$m_\lambda = \sum z_{i_1}^{\lambda_1} z_{i_2}^{\lambda_2} \cdots z_{i_{\ell(\lambda)}}^{\lambda_{\ell(\lambda)}}$$

for $\lambda = \lambda_1 \lambda_2 \cdots \lambda_{\ell(\lambda)}$ and the monomials are distinct.

- (3) **A formula for q, t -Kostka polynomials** The modified Macdonald polynomials \tilde{H}_λ , $\lambda \vdash n$, can be expanded as a linear combination of Schur functions

$$\tilde{H}_\lambda = \sum_{\mu \vdash n} \tilde{K}_{\mu\lambda}(q, t) s_\mu$$

where the $\tilde{K}_{\mu\lambda}(q, t)$ are known as q, t -Kostka polynomials. It is still an open problem to find a combinatorial formula for them, although such formulas have been found for $\lambda = m1^{n-m}$ by Stembridge [33, Theorem 2.1], and also for $\lambda = 2^m 1^{n-2m}$ by Fishel [9, Theorem 1.1], and others. Assaf, furthermore, has a theorem that enables the unification of these two cases [2, Theorem 18].

5. ACKNOWLEDGMENTS

The author would like to thank H el ene Barcelo, David Eisenbud, Adriano Garsia, Jim Haglund, Angela Hicks, Richard Stanley and Mike Zabrocki for many fascinating conversations, and the Centre de Recherches Math ematiques and the Laboratoire de Combinatoire et d'Informatique Math ematique where many of the conversations and much of the subsequent writing took place thanks to a Simons CRM Professorship. She is also grateful to Niall Christie, Samantha Dahlberg, Jim Haglund, Angela Hicks, Franco Saliola and Mike Zabrocki for feedback on an earlier draft of the manuscript, and especially to Angela Hicks who additionally generously shared her draft of the history of the shuffle conjecture, and her code that created the diagrams.

REFERENCES

- [1] M. Aguiar, N. Bergeron and F. Sottile, *Combinatorial Hopf algebras and generalized Dehn-Sommerville relations*, Compos. Math. 142 (2006) 1–30.
- [2] S. Assaf, *Toward the Schur expansion of Macdonald polynomials*, Electron. J. Combin. 25 (2018).
- [3] F. Bergeron, *Algebraic combinatorics and coinvariant spaces*, CRC Press (2009).
- [4] F. Bergeron and A. Garsia, *Science fiction and Macdonald's polynomials*, In Algebraic Methods and q -Special Functions, CRM Proceedings and Lecture Notes 22 (1999) 1–52.
- [5] E. Carlsson and A. Mellit, *A proof of the shuffle conjecture*, arXiv:1508.06239.
- [6] E. Carlsson and A. Mellit, *A proof of the shuffle conjecture*, J. Amer. Math. Soc. 31 (2018) 661–697.
- [7] A. Cauchy, *M emoire sur les fonctions qui ne peuvent obtenir que deux valeurs  egales et de signes contraires par suite des transpositions op er ees entre les variables qu'elles renferment*, J.  Ec. Polytech. Oeuvres Series 2 Volume 1 (1815) 91–169.
- [8] R. Cori and D. Rossin, *On the sandpile group of dual graphs*, European J. Combin. 21 (2000) 447–459.

- [9] S. Fishel, *Statistics for special q, t -Kostka polynomials*, Proc. Amer. Math. Soc. 123 (1995) 2961–2969.
- [10] A. Garsia and M. Haiman, *A remarkable q, t -Catalan sequence and q -Lagrange inversion*, J. Algebraic Combin. 5 (1996) 191–244.
- [11] I. Gessel, *Multipartite P -partitions and inner products of skew Schur functions*, In Combinatorics and Algebra, Contemp. Math. 34 (1984) 289–301.
- [12] A. Girard, *Invention nouvelle en l’algèbre*, Amsterdam (1629).
- [13] J. Haglund, *The q, t -Catalan numbers and the space of diagonal harmonics: with an appendix on the combinatorics of Macdonald polynomials*, University Lecture Series, American Mathematical Society (2008).
- [14] J. Haglund, M. Haiman, N. Loehr, J. Remmel and A. Ulyanov, *A combinatorial formula for the character of the diagonal coinvariants*, Duke Math. J. 126 (2005) 195–232.
- [15] J. Haglund, J. Morse and M. Zabrocki, *A compositional shuffle conjecture specifying touch points of the Dyck path*, Canad. J. Math. 64 (2012) 822–844.
- [16] J. Haglund and G. Xin, *Lecture notes on the Carlsson-Mellit proof of the shuffle conjecture*, arXiv:1705.11064.
- [17] M. Haiman, *Vanishing theorems and character formulas for the Hilbert scheme of points in the plane*, Invent. Math. 149 (2002) 371–407.
- [18] P. Hersh and S. Hsiao, *Random walks on quasisymmetric functions*, Adv. Math. 222 (2009) 782–808.
- [19] A. Hicks, *Combinatorics of the diagonal harmonics*, to appear.
- [20] K. Kadell, *A proof of some q -analogues of Selberg’s integral for $k = 1$* , SIAM J. Math. Anal. 19 (1988) 944–968.
- [21] A. Konheim and B. Weiss, *An occupancy discipline and applications*, SIAM J. Appl. Math. 14 (1966) 1266–1274.
- [22] N. Loehr and G. Warrington, *Square q, t -lattice paths and $\nabla(p_n)$* , Trans. Amer. Math. Soc. 359 (2007) 649–669.
- [23] N. Loehr and G. Warrington, *Nested quantum Dyck paths and $\nabla(s_\lambda)$* , Int. Math. Res. Not. IMRN (2008).
- [24] I. Macdonald, *A new class of symmetric functions*, Sémin. Lothar. Combin. 20 (1988) 131–171.
- [25] R. Pyke, *The supremum and infimum of the Poisson process*, Ann. Math. Statist. 30 (1959) 568–576.
- [26] B. Sagan, *The symmetric group. Representations, combinatorial algorithms, and symmetric functions. Second edition*, Springer-Verlag (2001).
- [27] I. Schur, *Über eine Klasse von Matrizen, die sich einer gegebenen Matrix zuordnen lassen*, Dissertation, Berlin (1901).
- [28] E. Sergel, *A proof of the square paths conjecture*, J. Combin. Theory Ser. A 152 (2017) 363–379.
- [29] E. Sergel, *A combinatorial model for ∇m_μ* , arXiv:1804.06037.
- [30] R. Stanley, *Ordered structures and partitions*, Mem. Amer. Math. Soc. 119 (1972).
- [31] R. Stanley, *Hyperplane arrangements, parking functions and tree inversions*, In Mathematical Essays in honor of Gian-Carlo Rota, Birkhäuser-Verlag (1998) 359–375.
- [32] R. Stanley, *Enumerative combinatorics. Volume 2*, Cambridge University Press (1999).
- [33] J. Stembridge, *Some particular entries of the two-parameter Kostka matrix*, Proc. Amer. Math. Soc. 121 (1994) 367–373.
- [34] C. Yan, *Parking functions*, In Handbook of Enumerative Combinatorics, Discrete Math. Appl. (2015) 835–894.
- [35] A. Young, *On quantitative substitutional analysis (third paper)*, Proc. Lond. Math. Soc. (1928) 255–292.

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF BRITISH COLUMBIA, VANCOUVER, BC V6T 1Z2, CANADA

E-mail address: `steph@math.ubc.ca`

TANGENT DEVELOPABLE SURFACES AND THE EQUATIONS DEFINING ALGEBRAIC CURVES

ROBERT LAZARSFELD

INTRODUCTION

Let X be a smooth complex projective curve – or equivalently a compact Riemann surface – of genus $g \geq 2$, and denote by

$$H^{1,0}(X) = \Gamma(X, \Omega_X^1)$$

the \mathbf{C} -vector space of holomorphic 1-forms on X . Recalling that $\dim H^{1,0}(X) = g$, choose a basis

$$\omega_1, \dots, \omega_g \in H^{1,0}(X).$$

It is classical that the ω_i do not simultaneously vanish at any point $x \in X$, so one can define a holomorphic map

$$\phi_X : X \longrightarrow \mathbf{P}^{g-1}, \quad x \mapsto [\omega_1(x), \dots, \omega_g(x)]$$

from X to a projective space of dimension $g - 1$, called the *canonical mapping* of X . With one well-understood class of exceptions, ϕ_X is an embedding, realizing X as an algebraic curve

$$X \subseteq \mathbf{P}^{g-1}$$

of degree $2g - 2$. Any compact Riemann surface admits many projective embeddings, but the realization just constructed has the big advantage of being canonically defined up to a linear change of coordinates on \mathbf{P}^{g-1} . Therefore the extrinsic projective geometry of a canonically embedded curve must reflect its intrinsic geometry, and working this principle out is an important theme in the theory of algebraic curves.

Given any projective variety, one can consider the degrees of its defining equations. An important theorem of Petri from 1922 states that with a slightly wider range of exceptions, a canonical curve $X \subseteq \mathbf{P}^{g-1}$ is cut out by quadrics, i.e. polynomials of degree two. Classically that seemed to be the end of the story, but in the early 1980s Mark Green realized that Petri's result should be the first case of a much more general statement involving higher syzygies. In other words, one should consider not only the defining equations themselves, but the relations among them, the relations among the relations, and so on. The resulting conjecture has attracted a huge amount of attention over the past thirty-five years.

As of this writing, Green's conjecture remains open. However Voisin made a major breakthrough in 2002 by proving that it holds for *general* curves, where one rules out

Research partially supported by NSF grant DMS-1739285.

for instance all the sorts of exceptional cases alluded to above. Her proof introduced a number of very interesting new ideas, but at the end of the day it relied on some difficult and lengthy cohomological calculations. Prior to Voisin’s work, O’Grady and Buchweitz–Schreyer had observed that one might be able to attack the syzygies of generic canonical curves by studying a very concrete and classical object, namely the developable surface of tangent lines to a rational normal curve. A substantial body of experimental evidence supported this proposal, but in spite of considerable effort nobody was able to push through the required computations. In a recent preprint [2], however, Aprodu, Farkas, Papadima, Raicu and Weyman have succeeded in doing so. Their work is the subject of the present report.

This note is organized as follows. Section 1 is devoted to the geometry of canonical curves, the basic ideas around syzygies, and the statement of Green’s conjecture. The case of general curves, and its relation to the tangent surface of rational normal curves occupies §2. Finally, in §3 we explain the main geometric ideas underlying the work of AFPRW. My understanding of this picture was worked out with Lawrence Ein.

We work throughout over the complex numbers. In particular, we completely ignore contributions of [2] to understanding what parts of Green’s conjecture work in positive characteristics.

I thank the authors of [2] for sharing an early draft of their paper. I profited from correspondence with David Eisenbud, Gabi Farkas, Claudiu Raicu, Frank Schreyer and Claire Voisin. I am particularly grateful to Lawrence Ein, with whom I worked to understand the paper of AFPRW from an algebro-geometric perspective.

1. CANONICAL CURVES, SYZYGIES AND GREEN’S CONJECTURE

Canonical curves and Petri’s theorem. Denote by X a smooth complex projective curve of genus $g \geq 2$, and as in the Introduction consider the canonical mapping

$$\phi_X : X \longrightarrow \mathbf{P}^{g-1}, \quad x \mapsto [\omega_1(x), \dots, \omega_g(x)],$$

arising from a basis $\omega_1, \dots, \omega_g \in H^{1,0}(X)$ of holomorphic 1-forms on X . By construction the inverse image of a hyperplane cuts out the zero-locus of a such a 1-form, and therefore consists of $2g - 2$ points (counting multiplicities). It is instructive to consider concretely the first few cases.

Example 1.1. (Genus 2 and hyperelliptic curves). Suppose $g(X) = 2$. Then the canonical mapping is a degree two branched covering

$$\phi_X : X \longrightarrow \mathbf{P}^1.$$

In general, a curve of genus g admitting a degree two covering $X \longrightarrow \mathbf{P}^1$ is called *hyperelliptic*. Thus every curve of genus 2 is hyperelliptic, but when $g \geq 3$ these are in many respects the “most special” curves of genus g . The canonical mapping of a hyperelliptic curve factors as the composition

$$X \longrightarrow \mathbf{P}^1 \subseteq \mathbf{P}^{g-1}$$

of the hyperelliptic involution with an embedding of \mathbf{P}^1 into \mathbf{P}^{g-1} . It follows from the Riemann–Roch theorem that these are the only curves for which ϕ_X is not an embedding:

FACT. If X is non-hyperelliptic, then the canonical mapping

$$\phi_X : X \subseteq \mathbf{P}^{g-1}$$

is an embedding.

Example 1.2. (Genus 3 and 4). Assume that X is not hyperelliptic. When $g(X) = 3$, the canonical mapping realizes X as a smooth curve of degree 4 in \mathbf{P}^2 , and any such curve is canonically embedded. When $g = 4$, ϕ_X defines an embedding

$$X \subseteq \mathbf{P}^3$$

in which X is the complete intersection of a surface of degree 2 and degree 3.

Example 1.3. (Genus 5). This is the first case where one sees the interesting behavior of quadrics through a canonical curve. Consider a canonically embedded non-hyperelliptic curve of genus 5

$$X \subseteq \mathbf{P}^4, \quad \deg X = 8.$$

One can show that there is a three-dimensional vector space of quadrics through X , say with basis Q_1, Q_2, Q_3 . There are now two possibilities:

- (a). X is *trigonal*, i.e. there exists a degree three branched covering

$$\pi : X \longrightarrow \mathbf{P}^1.$$

In this case each of the fibres of π spans a line in \mathbf{P}^4 , and hence any quadric containing X must also contain each of these lines. They sweep out a ruled surface $S \subseteq \mathbf{P}^4$ containing X , and three quadrics through X meet precisely along S :

$$Q_1 \cap Q_2 \cap Q_3 = S.$$

The canonical curve X is cut out in S by some cubic forms.

- (b). X is not trigonal, i.e. cannot be expressed as a 3-sheeted branched covering of \mathbf{P}^1 . Then X is the complete intersection of the the three quadrics containing it:

$$X = Q_1 \cap Q_2 \cap Q_3.$$

This is the general case.

Example 1.4. (Genus 6). Consider finally a non-hyperelliptic canonical curve $X \subseteq \mathbf{P}^5$ of genus 6. Now the polynomials of degree two vanishing on X form a vector space of dimension 6, and there are three cases:

- (a). If X is trigonal, then as above the quadrics through X intersect along the ruled surface S swept out by the trigonal divisors.
- (b). Suppose that X is a smooth curve of degree 5 in \mathbf{P}^2 . In this case the canonical image of X lies on the *Veronese surface*

$$X \subseteq V \subseteq \mathbf{P}^5,$$

a surface of degree 4 abstractly isomorphic to \mathbf{P}^2 , and V is the intersection of the quadrics through X in canonical space.

- (c). The general situation is that X is neither trigonal nor a plane quintic, and then $X \subseteq \mathbf{P}^5$ is cut out by the quadrics passing through it. Note however that X is not the complete intersection of these quadrics, since they span a vector space of dimension strictly greater than the codimension of X .

We conclude this subsection by stating Petri's theorem. Consider a non-hyperelliptic canonical curve $X \subseteq \mathbf{P}^{g-1}$. Let $S = \mathbf{C}[Z_0, \dots, Z_{g-1}]$ be the homogeneous coordinate ring of canonical space \mathbf{P}^{g-1} , and denote by

$$I_X \subseteq S$$

the homogeneous ideal of all forms vanishing on X . We ask when I_X is generated by forms of degree 2: this is the strongest sense in which X might be cut out by quadrics.

Theorem 1.5 (Petri). *The homogeneous ideal I_X fails to be generated by quadrics if and only if X is either trigonal or a smooth plane quintic.*

Note that the Petri-exceptional curves fall into two classes: there is one family (trigonal curves) that appears in all genera, and in addition one "sporadic" case.

In retrospect, Petri's statement suggests some natural questions. For example, how does one detect algebraically curves X that can be expressed as a degree 4 branched covering $X \rightarrow \mathbf{P}^1$, or that arise as smooth plane sextics? Or again, what happens in the generic case, when X does not admit any unusually low degree mappings to projective space? Green's beautiful insight is that one should consider for this not just the generators of I_X but also its higher syzygies.

Syzygies. The idea to study the relations – or syzygies – among the generators of an ideal goes back to Hilbert. Making this precise inevitably involves a certain amount of notation, so perhaps it's best to start concretely with the simplest example.¹

The rational normal curve $C \subseteq \mathbf{P}^3$ of degree 3 is the image of the embedding

$$\nu : \mathbf{P}^1 \longrightarrow \mathbf{P}^3, \quad [u, v] \mapsto [u^3, u^2v, uv^2, v^3].$$

Writing $[Z_0, \dots, Z_3]$ for homogeneous coordinates on \mathbf{P}^3 , it is a pleasant exercise to show that C can be described as the locus where a catalecticant matrix drops rank:

$$C = \left\{ \text{rank} \begin{bmatrix} Z_0 & Z_1 & Z_2 \\ Z_1 & Z_2 & Z_3 \end{bmatrix} \leq 1 \right\}.$$

Therefore C lies on the three quadrics

$$Q_{02} = Z_0Z_2 - Z_1^2, \quad Q_{03} = Z_0Z_3 - Z_1Z_2, \quad Q_{13} = Z_1Z_3 - Z_2^2$$

¹We refer the reader to [6] for an systematic introduction to the theory from an algebraic perspective, and [1] for a more geometric viewpoint.

given by the 2×2 minors of this matrix, and in fact these generate the homogeneous ideal I_C of C . While the Q_{ij} are linearly independent over \mathbf{C} , they satisfy two relations with polynomial coefficients, namely

$$(*) \quad \begin{aligned} Z_0 \cdot Q_{13} - Z_1 \cdot Q_{03} + Z_2 \cdot Q_{02} &= 0 \\ Z_1 \cdot Q_{13} - Z_2 \cdot Q_{03} + Z_3 \cdot Q_{02} &= 0. \end{aligned}$$

One can derive these by repeating a row of the matrix defining C and expanding the resulting determinant along the duplicate row. Moreover it turns out that any relation among the Q_{ij} is a consequence of these.

We recast this discussion somewhat more formally. Write $S = \mathbf{C}[Z_0, \dots, Z_3]$ for the homogeneous coordinate ring of \mathbf{P}^3 . The three quadric generators of I_C determine a surjective map

$$S(-2)^{\oplus 3} \longrightarrow I_C,$$

where $S(-2)$ denotes a copy of S re-graded so that multiplication by the Q_{ij} is degree preserving. The relations in $(*)$ come from choosing generators for the kernel of this map. So the upshot of the previous paragraph is that one has an exact sequence

$$0 \longrightarrow S(-3)^{\oplus 2} \xrightarrow{\begin{pmatrix} Z_0 & Z_1 \\ -Z_1 & -Z_2 \\ Z_2 & Z_3 \end{pmatrix}} S(-2)^{\oplus 3} \xrightarrow{(Q_{13} \quad Q_{03} \quad Q_{02})} I_C \longrightarrow 0$$

of S -modules. This is the *minimal graded free resolution* of I_C .

The general situation is similar. Sticking for simplicity to the one-dimensional case, consider a non-degenerate curve

$$C \subseteq \mathbf{P}^r$$

i.e. one not lying on any hyperplanes. We suppose in addition that C is projectively normal, a technical condition that holds for any embedding of sufficiently large degree (and for non-hyperelliptic canonical curves thanks to a theorem of Noether.) Put

$$S = \mathbf{C}[Z_0, \dots, Z_r],$$

and denote by $I_C \subseteq S$ the homogeneous ideal of C . Then I_C has a minimal resolution E_\bullet of length $r - 1$:

$$(1.1) \quad 0 \longrightarrow E_{r-1} \longrightarrow \dots \longrightarrow E_2 \longrightarrow E_1 \longrightarrow I_C \longrightarrow 0,$$

where $E_i = \oplus S(-a_{i,j})$. It is elementary that $a_{i,j} \geq i + 1$ for every j .

Green realized that the way to generalize classical statements about quadratic generation of I_C is to ask when the first p steps of this resolution are as simple as possible.

Definition 1.6. One says that C satisfies Property (N_p) if

$$E_i = \oplus S(-i - 1) \quad \text{for every } 1 \leq i \leq p.$$

Thus (N_1) holds if and only if I_C is generated in degree 2. The first non-classical condition is (N_2) , which asks in addition that if one chooses quadratic generators $Q_\alpha \in I_C$, then the module of syzygies among the Q_α should be spanned by relations of the form

$$(*) \quad \sum L_\alpha \cdot Q_\alpha = 0,$$

where the L_α are *linear* polynomials. Condition (N_3) would ask that the syzygies among the coefficient vectors describing the relations $(*)$ are themselves generated by polynomials of degree one.

Example 1.7. The twisted cubic $C \subseteq \mathbf{P}^3$ discussed above satisfies (N_2) . On the other hand, an elliptic curve $E \subseteq \mathbf{P}^3$ of degree 4 is the complete intersection of two quadrics, whose homogeneous ideal has a Koszul resolution:

$$0 \longrightarrow S(-4) \longrightarrow S(-2)^{\oplus 2} \longrightarrow I_E \longrightarrow 0.$$

Thus E satisfies (N_1) but not (N_2) .

Return now to a non-hyperelliptic canonical curve $X \subseteq \mathbf{P}^{g-1}$ of genus g . Petri's theorem states that X satisfies (N_1) unless it is trigonal or a smooth plane quintic. Green's conjecture vastly extends this by predicting when X satisfies condition (N_p) .

Green's conjecture. In order to state Green's conjecture, it remains to understand the pattern behind the exceptional cases in Petri's theorem.

Let X be a curve of genus $g \geq 2$, and suppose given a non-constant holomorphic mapping

$$\varphi : X \longrightarrow \mathbf{P}^r.$$

We assume that X does not map into any hyperplanes, in which case we write $r(\varphi) = r$: this is often called the *dimension* or *rank* of φ . If φ has degree d in the sense that a general hyperplane pulls back to d points on X , we set $d(\varphi) = d$. The *Clifford index* of φ is then defined to be

$$\text{Cliff}(\varphi) = d(\varphi) - 2 \cdot r(\varphi).$$

A classical theorem of Clifford states that if $d(\varphi) \leq g - 1$, then

$$\text{Cliff}(\varphi) \geq 0,$$

and equality holds if and only if X is hyperelliptic and $\varphi : X \longrightarrow \mathbf{P}^1$ is the hyperelliptic involution (or a mapping derived from it by a Veronese-type construction).

We now attach an invariant to X by considering the minimum of the Clifford indices of all "interesting" mappings:

Definition 1.8. The Clifford index of X is

$$\text{Cliff}(X) = \min \{ \text{Cliff}(\varphi) \mid d(\varphi) \leq g - 1 \}.$$

One has

$$0 \leq \text{Cliff}(X) \leq \left\lfloor \frac{g-1}{2} \right\rfloor,$$

for every X , the first inequality coming from Clifford's theorem, and the second (as we explain in the next section) from Brill-Noether theory. Moreover

$$\text{Cliff}(X) = 0 \iff X \text{ is hyperelliptic,}$$

and similarly one can show that when X is non-hyperelliptic,

$$\text{Cliff}(X) = 1 \iff X \text{ is either trigonal or a smooth plane quintic.}$$

It is now clear what to expect for higher syzygies:

Conjecture 1.9 (Green, [9]). *Let $X \subseteq \mathbf{P}^{g-1}$ be a non-hyperelliptic canonical curve. Then the Clifford index of X is equal to the least integer p for which Property (N_p) fails for X .*

The case $p = 1$ is exactly Petri's theorem, and the first non-classical case $p = 2$ was established by Schreyer [11] and Voisin [12]. There is a symmetry among the syzygies of canonical curves, and knowing the smallest value of p for which (N_p) fails turns out to determine the grading of the whole resolution of I_X .

One implication in Green's statement is elementary: it was established in the appendix to [9] that if $\text{Cliff}(X) = p$, then (N_p) fails for X . What remains mysterious as of this writing is how to show conversely that unexpected syzygies actually have a geometric origin.

2. GENERAL CURVES OF LARGE GENUS AND THE TANGENT DEVELOPABLE TO RATIONAL NORMAL CURVES

General curves. The most important instance of Green's conjecture – which is the actual subject of the present report – is the case of “general” curves. We start by explaining a little more precisely what one means by this.

In the 1960s, Mumford and others constructed an algebraic variety \mathfrak{M}_g whose points parameterize in a natural way isomorphism classes of smooth projective curves of genus $g \geq 2$. This is the *moduli space* of curves of genus g . One has

$$\dim \mathfrak{M}_g = 3g - 3,$$

formalizing a computation going back to Riemann that compact Riemann surfaces of genus $g \geq 2$ depend on $3g - 3$ parameters. Special classes of curves correspond to (locally closed) proper subvarieties of \mathfrak{M}_g : for example, hyperelliptic curves are parameterized by a subvariety $\mathfrak{H}_g \subseteq \mathfrak{M}_g$ of dimension $2g - 1$, showing again that hyperelliptic curves are special when $g \geq 3$. One says that a statement holds for a *general curve* of genus g if it holds for all curves whose moduli points lie outside a finite union of proper subvarieties of \mathfrak{M}_g .

The question of what mappings $\varphi : X \rightarrow \mathbf{P}^r$ exist for a general curve X was studied classically, and the theory was put on a firm modern footing in the 1970s by Kempf, Kleiman-Laksov, Griffiths-Harris and Gieseker, among others: see [4]. For our purposes, the basic fact is the following:

Theorem 2.1 (Weak form of Brill–Noether theorem). *Let X be a general curve of genus $g \geq 2$. Then there exists a map $\varphi : X \rightarrow \mathbf{P}^r$ of degree d and dimension r if and only if*

$$g \geq (r + 1)(g - d + r).$$

In particular, $\text{Cliff}(X) = \lfloor \frac{g-1}{2} \rfloor$.

Green's conjecture then predicts the shape of the minimal resolution of the ideal of a general canonical curve of genus g . This is the stunning result established by Voisin [13], [14].

Theorem 2.2 (Voisin's Theorem). *Put $c = \lfloor \frac{g-3}{2} \rfloor$. Then a general canonical curve $X \subseteq \mathbf{P}^{g-1}$ of genus g satisfies Property (N_c) .*

The symmetry mentioned following the statement of Green's conjecture imposes limits on how far (N_p) could be satisfied, and one can view Voisin's theorem as asserting that the syzygies of a general canonical curve are "as linear as possible" given this constraint.

General principles imply that the set of curves for which the conclusions of Theorems 2.1 or 2.2 hold are parameterized by Zariski-open subsets of \mathfrak{M}_g . So to prove the results it would suffice to exhibit one curve of each genus g for which the assertions are satisfied. However it has long been understood that this is not a practical approach. Instead, two different strategies have emerged for establishing statements concerning general canonical curves.

The first is to consider singular rational curves. For example, a rational curve $\Gamma \subseteq \mathbf{P}^{g-1}$ of degree $2g-2$ with g nodes can be realized as a limit of canonical curves. The first proof of Theorem 2.1, by Griffiths and Harris, went by establishing that Γ satisfies an appropriate analogue of the statement, and then deducing that 2.1 must hold for a general smooth curve of genus g . A difficulty here is that the nodes themselves have to be in general position, requiring a further degeneration. Eisenbud and Harris subsequently found that it is much better to work with cuspidal curves: we will return to this shortly. More recently, tropical methods have entered the picture to give new proofs of Theorem 2.1.

A different approach, initiated in [10], involves K3 surfaces. These are surfaces

$$S = S_{2g-2} \subseteq \mathbf{P}^g$$

of degree $2g-2$ whose hyperplane sections are canonical curves. It turns out to be quite quick to show that these curves are Brill-Noether general provided that S itself is generic. While it is not easy to exhibit explicitly a suitable K3, it is known by Hodge theory that they exist in all genera. This re-establishes the existence of curves that behave generically from the perspective of Brill-Noether theory.

This was the starting point of Voisin's proof of Theorem 2.2. Under favorable circumstances the resolution of a surface restricts to that of its hyperplane section, so it suffices to show that a general K3 surface of genus g satisfies the conclusion of 2.2. However so far this doesn't really simplify the picture. Voisin's remarkable new idea was to pass to a larger space, namely the Hilbert scheme

$$S^{[c+1]} = \text{Hilb}^{c+1}(S)$$

parameterizing finite subschemes of length $(c+1)$ on S . Voisin showed that the syzygies of S are encoded in a quite simple-looking geometric statement on $S^{[c+1]}$. The required computations turn out to be rather involved, but in a real tour de force Voisin succeeded in pushing them through. Interestingly, it later emerged that her computations could be used

to establish many other cases of Green's conjecture, eg that it holds for a general curve of each gonality, or for every curve appearing on a K3 surface. See [1] for some examples and references.

At about the same time that Green formulated his conjecture in the early 1980s, Eisenbud and Harris [7] realized that many of the difficulties involved in degenerating to nodal rational curves disappeared if one worked instead with rational curves with g cusps. The advantage of these curves is that they behave Brill-Noether generally without any conditions on the location of the singular points. This raised the possibility that one might use g -cuspidal curves to study syzygies of general canonical curves. It was at this point that Kieran O'Grady, and independently Buchweitz and Schreyer, remarked that it should suffice to understand the syzygies of a very classical object, namely the tangent surface to a rational normal curve.

The tangent developable of a rational normal curve. Let $C \subseteq \mathbf{P}^g$ be a rational normal curve of genus g . By definition this is the image of the embedding

$$\mathbf{P}^1 \hookrightarrow \mathbf{P}^g \quad \text{given by} \quad [s, t] \mapsto [s^g, s^{g-1}t, \dots, st^{g-1}, t^g].$$

One can associate to C (as to any smooth curve) its *tangent surface*

$$T = \text{Tan}(C) \subseteq \mathbf{P}^g,$$

defined to be the union of all the embedded projective tangent lines to C . In the case at hand, one can describe T very concretely. Specifically it is the image of the map

$$(2.1) \quad \nu : \mathbf{P}^1 \times \mathbf{P}^1 \longrightarrow \mathbf{P}^g$$

given matricially by

$$(2.2) \quad \nu([s, t] \times [u, v]) = [u \ v] \cdot \text{Jac}(\mu),$$

where $\text{Jac}(\mu)$ is the $2 \times (g+1)$ matrix of partials of $\mu = [s^g, s^{g-1}t, \dots, st^{g-1}, t^g]$. In other words,

$$\nu([s, t] \times [u, v]) = \left[g \cdot s^{g-1}u, (g-1) \cdot s^{g-2}tu + s^{g-1}v, \dots, g \cdot st^{g-1}v \right].$$

Note that ν is one-to-one, and maps the diagonal $\Delta \subseteq \mathbf{P}^1 \times \mathbf{P}^1$ isomorphically to C . However ν is not an embedding: it ramifies along the diagonal, and T has cuspidal singularities along C . The tangent surface T is a complex-geometric analogue of one of the classes of developable surfaces studied in differential geometry. A pleasant computation shows that $\deg(T) = 2g - 2$.²

The upshot of this discussion is that the hyperplane sections of T are rational curves $\Gamma \subseteq \mathbf{P}^{g-1}$ of degree $2g - 2$ with g cusps – in other words, the degenerations of canonical curves with which one hopes to be able to prove the generic case of Green's conjecture. This led to the

²Either observe that ν is given by an (incomplete) linear series of type $(g-1, 1)$ on $\mathbf{P}^1 \times \mathbf{P}^1$, or use Riemann-Hurwitz for a degree g mapping $\mathbf{P}^1 \rightarrow \mathbf{P}^1$ to see that there are $2g-2$ tangent lines to C meeting a general linear space $\Lambda \subseteq \mathbf{P}^g$ of codimension 2.

Folk-Conjecture 2.3. *The tangent developable surface*

$$T = \text{Tan}(C) \subseteq \mathbf{P}^g$$

satisfies Property (N_p) for $p = \lfloor \frac{g-3}{2} \rfloor$.

With a small argument showing that T indeed has the same syzygies as its hyperplane sections, it has been well understood since the mid 1980s that this would imply the result (Theorem 2.2) that Voisin later proved by completely different methods.

The important thing to observe about 2.3 is that it is a completely concrete statement. Via the parameterization (2.2), the conjecture was quickly verified for a large range of genera using early versions of the computer algebra system Macaulay. That such an utterly down-to-earth assertion could resist proof for thirty-five years has been something of a scandal. Happily, the work of Aprodu, Farkas, Papadima, Raicu and Weyman has remedied this situation.

3. SKETCH OF THE PROOF OF CONJECTURE 2.3

In this section, we outline the main ideas of the work of AFPRW proving Folk Conjecture 2.3.

The actual write-up in [2] is a bit long and complicated, in part because the authors work to extend their results as far as possible to positive characteristics, and in part because they are fastidious in checking that the maps that come up are the expected ones. Here I focus on the essential geometric ideas that seem to underlie their computations. This understanding of the argument was worked out together with Lawrence Ein.

Computing the syzygies of T . The first step in the argument of [2] is to understand the tangent developable $T = \text{Tan}(C)$ and its syzygies in terms of more familiar and computable objects. This culminates in Theorem 3.3 below, which describes the relevant syzygies linear algebraically. Some of the computations of AFPRW apparently elaborate on earlier (unpublished) work of Weyman, as outlined in Eisenbud's notes [5].

A basic principle guiding algebraic geometry holds that spaces are determined by the polynomial functions on them, so we will need to understand those on T . It is in turn natural to expect that functions on the tangent developable should be described using the mapping $\nu : \mathbf{P}^1 \times \mathbf{P}^1 \rightarrow T$ from (2.1), which realizes T as the homeomorphic image of $\mathbf{P}^1 \times \mathbf{P}^1$ cusped along the diagonal. In order to get a sense of how this should go, let us start with a one-dimensional toy example.

Consider then the mapping

$$\nu_0 : \mathbf{A}^1 = \mathbf{C} \rightarrow \mathbf{A}^2 = \mathbf{C}^2, \quad \nu_0(t) = (t^2, t^3).$$

This maps \mathbf{A}^1 homeomorphically onto the cuspidal curve

$$T_0 = \{y^2 = x^3\} \subseteq \mathbf{A}^2$$

in the plane, and the polynomial functions on T_0 are realized as the subring

$$\mathbf{C}[T_0] = \mathbf{C}[t^2, t^3] \subseteq \mathbf{C}[t] = \mathbf{C}[\mathbf{A}^1]$$

of the regular functions on the affine line. The point to note is that we can describe $\mathbf{C}[T_0]$ intrinsically, without using the map ν_0 . Specifically, there is a \mathbf{C} -linear derivation

$$\delta_0 : \mathbf{C}[t] \longrightarrow \mathbf{C} \quad , \quad \delta_0(f) = f'(0),$$

and $\mathbf{C}[T_0] = \ker(\delta_0)$. Moreover while δ_0 is not $\mathbf{C}[t]$ -linear, it is linear over $\mathbf{C}[t^2, t^3]$, giving a short exact sequence $0 \longrightarrow \mathbf{C}[T_0] \longrightarrow \mathbf{C}[\mathbf{A}^1] \longrightarrow \mathbf{C} \longrightarrow 0$ of $\mathbf{C}[T_0]$ -modules.

This model generalizes. Writing \mathcal{O}_X to denote the (sheaf of locally) polynomial functions on a variety X , one has

Proposition 3.1. *There is a \mathbf{C} -linear derivation*

$$\delta : \mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1} \longrightarrow \Omega_{\Delta}^1 \quad , \quad \delta(f) = df|_{\Delta}$$

with $\ker \delta = \mathcal{O}_T$. Moreover, this gives rise to a short exact sequence

$$(3.1) \quad 0 \longrightarrow \mathcal{O}_T \longrightarrow \mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1} \xrightarrow{\delta} \Omega_C^1 \longrightarrow 0$$

(of sheaves) on \mathbf{P}^g .³

It is easy to describe the syzygies of $\mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1}$ and Ω_C^1 , and then the plan is to use (3.1) to analyze the syzygies of T .

At this point we require some additional syzygetic notation. As above denote by $S = \mathbf{C}[Z_0, \dots, Z_g]$ the homogeneous coordinate ring of \mathbf{P}^g , and consider a finitely generated graded S -module M . As in (1.1), M has a minimal graded free resolution E_{\bullet}

$$\dots \longrightarrow E_2 \longrightarrow E_1 \longrightarrow E_0 \longrightarrow M \longrightarrow 0,$$

where $E_i = E_i(M) = \bigoplus S(-a_{i,j})$.⁴ Write

$$K_{i,1}(M) = \{ \text{minimal generators of } E_i(M) \text{ of degree } i+1 \}.$$

This is a finite dimensional vector space whose elements we call i^{th} syzygies of weight 1. (The space $K_{i,q}$ of syzygies of weight q are defined analogously.) For instance the ideal I_C of the twisted cubic $C \subseteq \mathbf{P}^3$ discussed in §1 satisfies

$$\dim K_{2,1}(I_C) = 2 \quad , \quad \dim K_{1,1}(I_C) = 3.$$

When M is the S -module associated to a coherent sheaf \mathcal{F} on \mathbf{P}^g , we write simply $K_{i,1}(\mathcal{F})$. In particular, the weight one syzygies of the tangent developable T – which, as it turns out, govern Conjecture 2.3 – are given by $K_{i,1}(\mathcal{O}_T)$.

Proposition 3.1 then yields

³Strictly speaking, the middle term in (*) is the direct image $\nu_*(\mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1})$, but we wish to minimize sheaf-theoretic notation.

⁴We are purposely introducing a shift in indexing, so that here our resolutions start in homological degree zero rather than one. The reason for this is that we henceforth wish to view the resolution (1.1) of an ideal I as coming from one of S/I with $E_0 = S$.

Corollary 3.2. *For every $i \geq 1$ one has an exact sequence*

$$(3.2) \quad 0 \longrightarrow K_{i,1}(\mathcal{O}_T) \longrightarrow K_{i,1}(\mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1}) \longrightarrow K_{i,1}(\Omega_C^1).$$

Happily, it is quite easy to work out the two right-hand terms in the exact sequence (3.2).

Let U denote the two-dimensional complex vector space of linear functions on \mathbf{P}^1 , so that $\mathbf{P}^1 = \mathbf{P}(U)$ is the projective space of one-dimensional quotients of V . The group $\mathrm{SL}_2(\mathbf{C})$ acts on everything in sight, and in particular the Koszul groups $K_{i,1}$ will be representations of $\mathrm{SL}_2(\mathbf{C})$. After choosing an identification $\Lambda^2 U = \mathbf{C}$, a standard calculation shows that there is a canonical SL_2 -equivariant isomorphism

$$(3.3) \quad K_{i,1}(\mathcal{O}_{\mathbf{P}^1 \times \mathbf{P}^1}) = \Lambda^{i+1} S^{g-2} U \otimes S^{2i} U,$$

as well as a natural inclusion

$$(3.4) \quad K_{i,1}(\Omega_C^1) \subseteq \Lambda^{i+1} S^{g-1} U \otimes S^{i+1} U. \quad ^5$$

In view of Corollary 3.2, one then anticipates a mapping

$$(3.5) \quad \gamma : \Lambda^{i+1} S^{g-2} U \otimes S^{2i} U \longrightarrow \Lambda^{i+1} S^{g-1} U \otimes S^{i+1} U$$

whose kernel is $K_{i,1}(\mathcal{O}_T)$. AFPRW in effect devote very substantial effort to elucidating this map, but the upshot is that it is built from several off-the-shelf pieces. To begin with, there is a natural inclusion

$$(3.6) \quad S^{2i} U \longrightarrow \Lambda^2 S^{i+1} U$$

which is dual to the so-called Wahl map $\Lambda^2 S^{i+1} U^* \longrightarrow S^{2i} U^*$.⁶ Recalling that $\Lambda^{i+1}(A \otimes B)$ contains $\Lambda^{i+1} A \otimes S^{i+1} B$ as a summand for any vector spaces A and B , γ then arises as the composition

$$(3.7) \quad \begin{array}{c} \Lambda^{i+1} S^{g-2} U \otimes S^{2i} U \longrightarrow \Lambda^{i+1} S^{g-2} U \otimes S^{i+1} U \otimes S^{i+1} U \\ \searrow \hspace{10em} \nearrow \\ \Lambda^{i+1}(S^{g-2} U \otimes U) \otimes S^{i+1} U \longrightarrow \Lambda^{i+1} S^{g-1} U \otimes S^{i+1} U. \end{array}$$

We summarize this discussion as

Theorem 3.3. *With γ as just specified, $K_{i,1}(\mathcal{O}_T)$ sits in the exact sequence*

$$0 \longrightarrow K_{i,1}(\mathcal{O}_T) \longrightarrow \Lambda^{i+1} S^{g-2} U \otimes S^{2i} U \xrightarrow{\gamma} \Lambda^{i+1} S^{g-1} U \otimes S^{i+1} U.$$

⁵In arbitrary characteristic, which is the setting considered in [2], the computations are more delicate because one has to distinguish between divided and symmetric powers. Working as we are over \mathbf{C} , we can ignore this.

⁶If W is any two-dimensional \mathbf{C} -vector space with coordinates x, y , the Wahl or Gaussian mapping $\Lambda^2 S^{i+1} W \longrightarrow S^{2i} W$ is given (up to scaling) by $f \wedge g \mapsto \det \begin{pmatrix} f_x & f_y \\ g_x & g_y \end{pmatrix}$.

Hermite reciprocity and Koszul modules. Computations such as (3.3) and (3.4) are made by studying the cohomology of certain Koszul-type complexes. These can be difficult to deal with because they involve high wedge powers of a vector space or vector bundle. One of Voisin’s key insights was that upon passing to a Hilbert scheme, complicated multilinear data are encoded into more geometric questions about line bundles. The next step in the proof of AFPRW is an algebraic analogue of this strategy: one uses a classical theorem of Hermite to reinterpret Theorem 3.3 in a more tractable form involving only symmetric products. (In fact the analogy goes farther: a quick proof of Hermite reciprocity proceeds by interpreting $\Lambda^a S^b U$ as the space of global sections of a line bundle on the projective space $\mathbf{P}^a = \text{Hilb}^a(\mathbf{P}^1)$.)

As above, let U denote a complex vector space of dimension 2. The result in question is the following.

Hermite Reciprocity. *For any $a, b > 0$ there is a canonical $\text{SL}_2(\mathbf{C})$ –linear isomorphism*

$$(3.8) \quad \Lambda^a S^b U = S^{b+1-a} S^a U.$$

(See for example [8, Exercise 11.35].) In positive characteristics this is no longer true, and one of the contributions of [2] is to give a characteristic-free variant.

Plugging this into Theorem 3.3, one arrives at:

Corollary 3.4. *The Koszul group $K_{i,1}(\mathcal{O}_T)$ is the kernel of the map*

$$(3.9) \quad \gamma' : S^{2i} U \otimes S^{g-i-2} S^{i+1} U \longrightarrow S^{i+1} U \otimes S^{g-i-1} S^{i+1} U$$

obtained by pulling back the Koszul differential⁷

$$\Lambda^2 S^{i+1} U \otimes S^{g-i-2} S^{i+1} U \longrightarrow S^{i+1} U \otimes S^{g-i-1} S^{i+1} U$$

along the “co-Wahl” mapping $S^{2i} U \longrightarrow \Lambda^2 S^{i+1} U$ appearing in (3.6).

We now come to one of the main new ideas of [2], namely the introduction of Koszul (or Weyman) modules to study (3.9). To understand the motivation, set $V = S^{i+1} U$, $A = S^{2i} U$, and put $q = g - i - 2$. On the one hand we have from (3.6) an inclusion $A \subseteq \Lambda^2 V$, while for $q \geq 0$ there is a Koszul complex

$$\Lambda^2 V \otimes S^q V \longrightarrow V \otimes S^{q+1} V \longrightarrow S^{q+2} V.$$

The construction of γ' involved splicing these together, giving a three-term complex

$$(3.11) \quad A \otimes S^q V \xrightarrow{\gamma'} V \otimes S^{q+1} V \longrightarrow S^{q+2} V$$

⁷For any vector space V and integer $a > 0$, there is a natural map

$$\Lambda^2 V \otimes S^a V \longrightarrow V \otimes S^{a+1} V$$

which fits into the longer Koszul-type complex

$$(3.10) \quad \Lambda^2 V \otimes S^a V \longrightarrow V \otimes S^{a+1} V \longrightarrow S^{a+2} V \longrightarrow 0.$$

whose left-hand kernel $K = \ker \gamma'$ we would like to understand. Now suppose we knew that (3.11) is exact. Since in any event the map on the right is surjective, this would yield an exact sequence

$$0 \longrightarrow K \longrightarrow A \otimes S^q V \xrightarrow{\gamma'} V \otimes S^{q+1} V \longrightarrow S^{q+2} V \longrightarrow 0,$$

and we could immediately compute $\dim K$. The very nice observation of AFPRW is that the exactness of (3.11) is essentially automatic provided only that $q \geq \dim V - 3$.

Turning to details, let V be any complex vector space of dimension n , and suppose given a subspace $A \subseteq \Lambda^2 V$. As above, this determines for $q \geq 1$ a three-term complex

$$A \otimes S^q V \longrightarrow V \otimes S^{q+1} V \longrightarrow S^{q+2} V$$

whose homology $W_q(V, A)$ is called (the degree q piece of) the Koszul module associated to A and V . The essential result is:

Theorem 3.5. *Assume that no decomposable forms 2-forms $\eta \in \Lambda^2 V^*$ vanish on A . Then*

$$(3.12) \quad W_q(V, A) = 0 \quad \text{for } q \geq \dim V - 3.$$

This was originally proved in characteristic zero in [3] by a relatively painless application of Bott vanishing. An alternative (but largely equivalent) proof in characteristic zero uses vector bundles on projective space and considerations of Castelnuovo–Mumford regularity. In [2] the argument is extended to positive characteristics.

Remark 3.6. (Topological applications of Theorem 3.5) Before [2], the same authors had used Koszul modules in [3] to study some interesting topological questions, involving for example Kähler groups.

Completion of the proof. It is now immediate to complete the proof of Folk Conjecture 2.3. To begin with, using the symmetry in the resolution of T mentioned following the statement of Green’s Conjecture 1.9, one sees that 2.3 is equivalent to the assertion that

$$(3.13) \quad K_{\left[\frac{g}{2}\right], 1}(\mathcal{O}_T) = 0.$$

AFPRW treat separately the case of even and odd genus, so suppose that $g = 2n - 3$ is odd. Put

$$i = \left\lfloor \frac{g}{2} \right\rfloor = n - 2,$$

set $V = S^{i+1} U$ – so that $\dim V = n$ – and let $q = g - i - 2 = n - 3$. Corollary 3.4 shows that $K_{n-2, 1}(\mathcal{O}_T)$ is governed by the complex

$$S^{2(n-2)} U \otimes S^{n-3} V \longrightarrow V \otimes S^{n-2} V \longrightarrow S^{n-1} V$$

computing the Weyman module $W_{n-3}(V, S^{2(n-2)} U)$. The hypotheses of Theorem 3.5 are satisfied, and so $W_{n-3}(V, S^{2(n-2)} U) = 0$. Therefore we get an exact sequence

$$0 \longrightarrow K_{n-2, 1}(\mathcal{O}_T) \longrightarrow S^{2(n-2)} U \otimes S^{n-3} V \xrightarrow{\gamma'} V \otimes S^{n-2} V \longrightarrow S^{n-1} V \longrightarrow 0.$$

A computation of dimensions then shows that $\dim K_{n-2, 1}(\mathcal{O}_T) = 0$, and we are done!

REFERENCES

- [1] Marian Aprodu and J. Nagel, Koszul cohomology and algebraic geometry, AMS University Lecture Series **52** (2010)
- [2] Marian Aprodu, Gavril Farkas, Stefan Papadimu, Claudiu Raicu and Jerzy Weyman, Koszul modules and Green's conjecture, preprint, [arXiv:1810.11635](https://arxiv.org/abs/1810.11635)
- [3] Marian Aprodu, Gavril Farkas, Stefan Papadimu, Claudiu Raicu and Jerzy Weyman, Topological invariants of groups and Koszul modules, preprint. [arxiv:1806.01702](https://arxiv.org/abs/1806.01702)
- [4] Enrico Arbarello, Maurizio Cornalba, Phillip Griffiths and Joe Harris, *The Geometry of Algebraic Curves*, Springer Verlag, 1985.
- [5] David Eisenbud, Green's conjecture: An orientation for algebraists, in Free Resolutions in Commutative Algebra and Algebraic Geometry (Sundance 1990), Res. Notes Math. **2**, Jones and Bartlett, Boston (1992), pp. 51–78
- [6] David Eisenbud, *The Geometry of Syzygies*, Graduate Texts in Math **229**, Springer-Verlag New York, 2005.
- [7] David Eisenbud and Joe Harris, Divisors on general curves and cuspidal rational curves, Invent. Math. **74** (1983), pp. 371–418
- [8] William Fulton and Joe Harris, *Representation Theory*, Graduate Texts in Math **129**, Springer Verlag New York, 1991.
- [9] Mark Green, Koszul cohomology and the geometry of projective varieties, J. Diff. Geom. **19** (1984), pp. 125–171
- [10] Robert Lazarsfeld, Brill-Noether-Petri without degenerations, J. Diff. Geom. **23** (1986), pp 299 — 307.
- [11] Frank-Olaf Schreyer, Syzygies of canonical curves and special linear series, Math. Ann. **275** (1986), pp. 105-137.
- [12] Claire Voisin, Courbes tétragonales et cohomologie de Koszul, J. Reine Angew. Math. **387** (1988), 111–121.
- [13] Claire Voisin, Green's generic syzygy conjecture for curves of even genus lying on a K3 surface, Journal of the European Math Society **4** (2002), pp. 363-404.
- [14] Claire Voisin, Green's canonical syzygy conjecture for generic curves of odd genus, Compositio Math. **141** (2005), pp, 116301190.

DEPARTMENT OF MATHEMATICS, STONY BROOK UNIVERSITY, STONY BROOK, NEW YORK 11794

Email address: `robert.lazarsfeld@stonybrook.edu`

CURRENT EVENTS BULLETIN

Previous speakers and titles

For PDF files of talks, and links to Bulletin of the AMS articles,
see <http://www.ams.org/ams/current-events-bulletin.html>.

January 12, 2018 (San Diego, CA)

Richard D. James, University of Minnesota
Materials from mathematics

Craig L. Huneke, University of Virginia
How complicated are polynomials in many variables?

Isabelle Gallagher, Université Paris Diderot
*From Newton to Navier-Stokes, or how to connect fluid mechanics equations
from microscopic to macroscopic scales*

Joshua A. Grochow, University of Colorado, Boulder
*The Cap Set Conjecture, the polynomial method, and applications
(after Croot-Lev-Pach, Ellenberg-Gijswijt, and others)*

January 6, 2017 (Atlanta, GA)

Lydia Bieri, University of Michigan
Black hole formation and stability: a mathematical investigation.

Matt Baker, Georgia Tech
Hodge Theory in Combinatorics.

Kannan Soundararajan, Stanford University
Tao's work on the Erdos Discrepancy Problem.

Susan Holmes, Stanford University
Statistical proof and the problem of irreproducibility.

January 8, 2016 (Seattle, WA)

Carina Curto, Pennsylvania State University
What can topology tell us about the neural code?

Lionel Levine, Cornell University and *Yuval Peres, Microsoft Research
and University of California, Berkeley
Laplacian growth, sandpiles and scaling limits.

Timothy Gowers, Cambridge University
Probabilistic combinatorics and the recent work of Peter Keevash.

Amie Wilkinson, University of Chicago
What are Lyapunov exponents, and why are they interesting?

January 12, 2015 (San Antonio, TX)

Jared S. Weinstein, Boston University
Exploring the Galois group of the rational numbers: Recent breakthroughs.

Andrea R. Nahmod, University of Massachusetts, Amherst
*The nonlinear Schrödinger equation on tori: Integrating harmonic analysis,
geometry, and probability.*

Mina Aganagic, University of California, Berkeley
String theory and math: Why this marriage may last.

Alex Wright, Stanford University
From rational billiards to dynamics on moduli spaces.

January 17, 2014 (Baltimore, MD)

Daniel Rothman, Massachusetts Institute of Technology
Earth's Carbon Cycle: A Mathematical Perspective

Karen Vogtmann, Cornell University
The geometry of Outer space

Yakov Eliashberg, Stanford University
Recent advances in symplectic flexibility

Andrew Granville, Université de Montréal
*Infinitely many pairs of primes differ by no more than 70 million
(and the bound's getting smaller every day)*

January 11, 2013 (San Diego, CA)

Wei Ho, Columbia University

How many rational points does a random curve have?

Sam Payne, Yale University

Topology of nonarchimedean analytic spaces

Mladen Bestvina, University of Utah

Geometric group theory and 3-manifolds hand in hand: the fulfillment of Thurston's vision for three-manifolds

Lauren Williams, University of California, Berkeley

Cluster algebras

January 6, 2012 (Boston, MA)

Jeffrey Brock, Brown University

Assembling surfaces from random pants: the surface-subgroup and Ehrenpreis conjectures

Daniel Freed, University of Texas at Austin

The cobordism hypothesis: quantum field theory + homotopy invariance = higher algebra

Gigliola Staffilani, Massachusetts Institute of Technology

Dispersive equations and their role beyond PDE

Umesh Vazirani, University of California, Berkeley

How does quantum mechanics scale?

January 6, 2011 (New Orleans, LA)

Luca Trevisan, Stanford University

Khot's unique games conjecture: its consequences and the evidence for and against it

Thomas Scanlon, University of California, Berkeley

Counting special points: logic, Diophantine geometry and transcendence theory

Ulrike Tillmann, Oxford University

Spaces of graphs and surfaces

David Nadler, Northwestern University

The geometric nature of the Fundamental Lemma

January 15, 2010 (San Francisco, CA)

Ben Green, University of Cambridge

Approximate groups and their applications: work of Bourgain, Gamburd, Helfgott and Sarnak

David Wagner, University of Waterloo

Multivariate stable polynomials: theory and applications

Laura DeMarco, University of Illinois at Chicago

The conformal geometry of billiards

Michael Hopkins, Harvard University

On the Kervaire Invariant Problem

January 7, 2009 (Washington, DC)

Matthew James Emerton, Northwestern University

Topology, representation theory and arithmetic: Three-manifolds and the Langlands program

Olga Holtz, University of California, Berkeley

Compressive sensing: A paradigm shift in signal processing

Michael Hutchings, University of California, Berkeley

*From Seiberg-Witten theory to closed orbits of vector fields:
Taubes's proof of the Weinstein conjecture*

Frank Sottile, Texas A & M University

Frontiers of reality in Schubert calculus

January 8, 2008 (San Diego, California)

Günther Uhlmann, University of Washington

Invisibility

Antonella Grassi, University of Pennsylvania

Birational Geometry: Old and New

Gregory F. Lawler, University of Chicago

Conformal Invariance and 2-d Statistical Physics

Terence C. Tao, University of California, Los Angeles

Why are Solitons Stable?

January 7, 2007 (New Orleans, Louisiana)

Robert Ghrist, University of Illinois, Urbana-Champaign
Barcodes: The persistent topology of data

Akshay Venkatesh, Courant Institute, New York University
*Flows on the space of lattices: work of Einsiedler, Katok
and Lindenstrauss*

Izabella Laba, University of British Columbia
From harmonic analysis to arithmetic combinatorics

Barry Mazur, Harvard University
*The structure of error terms in number theory and an introduction
to the Sato-Tate Conjecture*

January 14, 2006 (San Antonio, Texas)

Lauren Ancel Myers, University of Texas at Austin
*Contact network epidemiology: Bond percolation applied
to infectious disease prediction and control*

Kannan Soundararajan, University of Michigan, Ann Arbor
Small gaps between prime numbers

Madhu Sudan, MIT
Probabilistically checkable proofs

Martin Golubitsky, University of Houston
Symmetry in neuroscience

January 7, 2005 (Atlanta, Georgia)

Bryna Kra, Northwestern University
*The Green-Tao Theorem on primes in arithmetic progression:
A dynamical point of view*

Robert McEliece, California Institute of Technology
Achieving the Shannon Limit: A progress report

Dusa McDuff, SUNY at Stony Brook
Floer theory and low dimensional topology

Jerrold Marsden, Shane Ross, California Institute of Technology
New methods in celestial mechanics and mission design

László Lovász, Microsoft Corporation
Graph minors and the proof of Wagner's Conjecture

January 9, 2004 (Phoenix, Arizona)

Margaret H. Wright, Courant Institute of Mathematical Sciences,
New York University
*The interior-point revolution in optimization: History, recent
developments and lasting consequences*

Thomas C. Hales, University of Pittsburgh
What is motivic integration?

Andrew Granville, Université de Montréal
It is easy to determine whether or not a given integer is prime

John W. Morgan, Columbia University
Perelman's recent work on the classification of 3-manifolds

January 17, 2003 (Baltimore, Maryland)

Michael J. Hopkins, MIT
Homotopy theory of schemes

Ingrid Daubechies, Princeton University
Sublinear algorithms for sparse approximations with excellent odds

Edward Frenkel, University of California, Berkeley
Recent advances in the Langlands Program

Daniel Tataru, University of California, Berkeley
The wave maps equation

2019 CURRENT EVENTS BULLETIN

Committee

Matt Baker, *Georgia Tech*

Hélène Barcelo, *Mathematical Sciences Research Institute*

Sandra Cerrai, *University of Maryland*

Henry Cohn, *Microsoft Research*

David Eisenbud, *Mathematical Sciences Research Institute,*
Committee Chair

Susan Friedlander, *University of Southern California*

Joshua Grochow, *University of Colorado, Boulder*

Gregory F. Lawler, *University of Chicago*

Ryan O'Donnell, *Carnegie Mellon University*

Lillian Pierce, *Duke University*


Alice Silverberg, *University of California, Irvine*

Michael Singer, *North Carolina State University*

Kannan Soundararajan, *Stanford University*

Gigliola Staffilani, *Massachusetts Institute of Technology*

Vanessa Goncalves, *AMS Administrative Support*



The back cover graphic is reprinted courtesy of Andrei Okounkov.

Cover graphic associated with Bhargav Bhatt's talk courtesy of the AMS.

Cover graphic associated with Thomas Vidick's talk courtesy of the AMS.

Cover graphic associated with Stephanie van Willigenburg's talk courtesy of the AMS.

Cover graphic associated with Robert Lazarsfeld's talk courtesy of Wikimedia Commons, **Attribution-ShareAlike 3.0 Unported (CC BY-SA 3.0)** (<https://creativecommons.org/licenses/by-sa/3.0/>).