Fermat's Two Squares Theorem Revisited

Bhaskar Bagchi

The Two Squares Theorem

Throughout this article, p is a prime such that $p \equiv 1 \pmod 4$. $\mathbb{N}$ and Z will denote, as usual, the set of all natural numbers (excluding zero) and the set of all integers (positive, negative or zero), respectively. Recall that the celebrated two squares theorem (first stated by Fermat and proved by Euler) says that p can be written as a sum of two perfect squares. Clearly one of these two squares must be even (and the other one is odd). Therefore this theorem may be formulated as saying that there exists $(x, y) \in \mathbb{N} \times \mathbb{N}$ such that $x^2 + 4y^2 = p$. Any such pair (x, y) will be referred to as a representation of p. (Actually, as is well known, the representation is unique. For a proof, see for instance – Niven and Zuckerman[2].)

Bhaskar Bagchi is with the Indian Statistical Institute since 1971, first as a student and then as a member of the faculty. He is interested in diverse areas of mathematics like combinatorics, elementary and analytic number theory, functional analysis, combinatorial topology and statistics.

Permutations

G H Hardy writes that the two squares theorem 'is ranked, very justly, as one of the finest in arithmetic'. So it comes as a surprise to learn that its finest proof was found only in 1990. In that year, D Zagier modified a proof of the two squares theorem due to Heathbrown to create a remarkably short and elegant proof. Although Zagier's proof was presented in detail by Shirali in a *Resonance* article[3], we shall begin with a brief account of this proof. To do so, we need to recall some facts about permutations.

If X is a finite set then by a permutation of X we mean a function from X into itself under which each element of X has a unique pre-image. If π and σ are any two permutations of X then we can form their 'product' $\pi\sigma$ by composition: $\pi\sigma(x) := \pi(\sigma(x))$, x in X. If X is of size n, there are only $n!$ permutations of X and they form a group with this product rule. (Though, strictly speaking, we need no group theory

for this article, familiarity with the elements of this theory will still be useful.) Since we have defined the product of any two permutations, in particular we can form the powers $\pi = \pi^1, \pi^2 \cdot$ of any given permutation π. Since there are only finitely many distinct permutations of X, some two of the powers of π must actually be equal. By cancellation, it follows that there must exist a natural number m such that π^m is the identity permutation id fixing all elements of X. The smallest such number is called the order of π. A permutation of order two is called an *involution.*

Any permutation π of X breaks up ('partitions') X into one or more parts such that two elements x and y of X are in the same part if and only if some power of π takes x to y. These parts are called the *orbits* of π. The singleton orbits are just the fixed points of π. A permutation of X is said to be transitive on X if it has only one orbit (namely the whole of X).

It is easy to convince oneself that the size of any orbit of a permutation divides the order of the permutation. In particular, if the permutation π has prime order q then (as 1 and q are the only divisors of q) each orbit is either a fixed point or has size q. It follows that, in this case, the number of fixed points of π is congruent modulo q to the size n of X. Hence π has a fixed point if n is not a multiple of q. As a special case ($q = 2$) of this observation we see that an involution of X has a fixed point in X if X is an odd set (i.e., the number of elements of X is odd). This is the key fact which makes Zagier's proof (and its constructive versions presented here) work.

Zagier's Proof

Now we come to Zagier's proof. Let S denote the subset of $\mathbb{N} \times \mathbb{N} \times \mathbb{N}$ defined by

$$S = \{(x, y, z) \in \mathbb{N} \times \mathbb{N} \times \mathbb{N} : x^2 + 4yz = p\}.$$

Clearly S is a finite set. Zagier defines two involutions α and β of S by

$$\alpha(x,y,z) = \begin{cases} (x+2z, z, y-x-z) & \text{if } x < y-z, \\ (2y-x, y, x+z-y) & \text{if } y-z < x < 2y, \\ (x-2y, x+z-y, y) & \text{if } x > 2y. \end{cases}$$

$$\beta(x,y,z) = (x,z,y).$$

The involution α of the finite set S has a unique fixed point (namely $(1,1,\frac{p-1}{4})$). It follows that S is an odd set. Therefore, the involution β of the odd set S must have an odd number (hence at least one) of fixed points in S. But $(x,y) \mapsto (x,y,y)$ is a bijection of the set of representations of p onto the set of fixed points of β. Hence p has at least one representation (as a sum of two squares). This completes Zagier's proof of the two squares theorem.

Shirali's Conjecture

Zagier notes in his paper that his proof 'is not constructive: it does not give a method to actually find the representation of p as a sum of two squares'. Perhaps provoked by this statement, S A Shirali gave a conjectural way to 'constructivize' this proof. Shirali's conjecture may be phrased as follows. Define a finite subset $\widehat{S}$ of $Z \times \mathbb{N} \times \mathbb{N}$ by

$$\widehat{S} = \{(x,y,z) \in Z \times \mathbb{N} \times \mathbb{N} : \ x+y > z \text{ and } x^2+4yz = p\}.$$

Define a function $\hat{\gamma} : \widehat{S} \to \widehat{S}$ by

$$\hat{\gamma}(x,y,z) = \begin{cases} (x+2z, y-x-z, z) & \text{if } x+z < y, \\ (2y-x, x+z-y, y) & \text{if } x+z > y. \end{cases}$$

Then Shirali conjectures that the orbit of the point $(1, \frac{p-1}{4}, 1)$ under $\hat{\gamma}$ contains a point of the form (x,y,y). That is, to obtain a point $(x,y,y) \in \widehat{S}$ (and hence a square plus square representation of p), begin with the point $(1, \frac{p-1}{4}, 1)$ and look at the successive iterates (powers) of $\hat{\gamma}$ on this point until a point (x,y,y) is obtained.

(Actually, Shirali defines his function on the (infinite) set of all points (x,y,z) in $Z \times Z \times Z$ satisfying $x^2+4yz = p$, and

proposes to begin with the α-fixed point $(1, 1, \frac{p-1}{4})$. However, we observed that this function fixes the finite subset $\widehat{S}$ introduced above and on this subset it restricts to $\hat{\gamma}$ as defined. Though the α-fixed point itself does not belong to this subset, its image under Shirali's original function is $(1, \frac{p-1}{4}, 1)$, which does belong. Therefore our formulation of the conjecture is entirely equivalent to Shirali's original formulation.)

A Constructive Version of Zagier's Proof

Notice that the function $\hat{\gamma}$ is a 'perturbation' of the permutation $\gamma := \alpha\beta$ of S obtained by composing Zagier's involutions α and β. So it is natural to ask if Shirali's conjecture is valid with $\hat{\gamma}$ replaced by γ. In the following theorem, we show that this modified conjecture is indeed correct. Note that we now stay within the set S, and this is closer to Zagier's original proof.

Theorem. *Let k denote the size of the orbit T under $\gamma := \alpha\beta$ which contains the α-fixed point a. Then k is odd; T contains a unique β-fixed point b and it is given by the formula $b = \gamma^{(k-1)/2}(a)$. In fact, the orbit T satisfies the symmetry relation $\gamma^{k-1-n}(a) = \beta(\gamma^n(a))$ for $0 \leq n \leq k-1$.*

Thus, to obtain a β-fixed point (x, y, y) (and hence a representation $p = x^2 + (2y)^2$), begin with the α-fixed point and iterate $\alpha\beta$ on it; in a finite number of steps you will reach a β-fixed point. This theorem shows that exactly half of the orbit has to be traversed before this point is reached; and the remaining half of the orbit may be found (in reverse order) simply by applying β to the first half.

Proof of the Theorem

Since α and β are involutions, α 'normalises' γ: $\alpha\gamma\alpha^{-1} = \beta\alpha = \gamma^{-1}$. Therefore α maps the orbits of γ to orbits of γ. (To see this, let s_1 and s_2 be two points from a common γ-orbit. By definition, this means that there is an integer ℓ such that $\gamma^\ell(s_1) = s_2$. Then $\alpha(s_2) = \alpha\gamma^\ell(s_1) = \alpha\gamma^\ell\alpha^{-1}(\alpha(s_1)) = \gamma^{-\ell}(\alpha(s_1))$. Thus, whenever s_1, s_2 in S are from a common γ-orbit, $\alpha(s_1)$ and $\alpha(s_2)$ are also in a

common γ-orbit. So the image under α of any γ-orbit is again a γ-orbit.) In particular, if T is the orbit under γ which contains the fixed point a of α, then $\alpha(T)$ is an orbit which meets the orbit T in this fixed point, hence we must have $\alpha(T) = T$ Since the restriction of α to T is an involution of T with a unique fixed point, it follows as before that T is an odd set. Since both α and γ fix T, so does $\beta = \alpha\gamma$. Thus (the restriction to T of) β is an involution of the odd set T and hence β must have a fixed point b in T So there is an ℓ, $0 \leq \ell \leq k-1$, such that $b = \gamma^{\ell}(a)$ is fixed by β. To prove the uniqueness of this fixed point, it suffices to show that $k = 2\ell + 1$ is forced on us.

For $m \in Z$, we have $\beta(\gamma^m(b)) = \beta\gamma^m\beta^{-1}(\beta(b)) = \gamma^{-m}(b)$. Substituting $\gamma^{\ell}(a)$ for b, we find that the orbit T has a two fold symmetry around its ℓth term:

$$\gamma^{\ell+m}(a) = \beta(\gamma^{\ell-m}(a)) \quad \forall m \in Z.$$

In particular, taking $m = \ell + 1$ in this identity, we get $\gamma^{2\ell+1}(a) = \beta\gamma^{-1}(a) = \beta^2\alpha(a) = \alpha(a) = a$. From the definition of k one sees that an integer h satisfies $\gamma^h(a) = a$ iff h is an integral multiple of k. Since $h = 2\ell + 1$ satisfies this condition, $2\ell + 1$ is a multiple of k. Since $1 \leq 2\ell + 1 < 2k$, this forces $2\ell + 1 = k$. Finally, substituting $\ell = \frac{k-1}{2}$, $m = \frac{k-1}{2} - n$ in the displayed identity, we get the last assertion of the theorem. □

Shirali's Conjecture Vindicated

Define the involutions $\hat{\alpha}$ and $\hat{\beta}$ of the finite set $\widehat{S}$ as follows.

$$\hat{\alpha}(x, y, z) = (2z - x, x + y - z, z),$$

$$\hat{\beta}(x, y, z) = \begin{cases} (-x, y, z) & \text{if } x + z < y, \\ (x, z, y) & \text{if } x + z > y. \end{cases}$$

One readily verifies that (i) these are indeed involutions of $\widehat{S}$, (ii) $\hat{\alpha}$ has a unique fixed point, namely $\hat{a} := (1, \frac{p-1}{4}, 1)$, and $(x, y) \mapsto (x, y, y)$ is a bijection from the representations of p onto the fixed points of $\hat{\beta}$. Thus, in Zagier's proof, one may replace α, β, S by $\hat{\alpha}$, $\hat{\beta}$, $\widehat{S}$, respectively. Finally,

Shirali's function $\hat{\gamma}$ is related to these involutions by $\hat{\gamma} = \hat{\alpha}\hat{\beta}$. Therefore the indicated substitutions in the proof of the above theorem yields a 'hatted' version of the theorem. In particular, this proves Shirali's conjecture.

Uniqueness of the Square Plus Square Representation of p

Aside from being non-constructive, Zagier's proof has another shortcoming. As already mentioned, the prime p has a unique representation as a sum of two squares. Or, what amounts to the same thing, β also has a unique fixed point in S. But this does not emerge from Zagier's proof (or from its constructive variations given above). We are unable to remedy this defect. Notice, however, that in view of the uniqueness assertion in the above theorem, it would suffice to show that γ acts transitively on S. (For, this would mean that $T = S$, and we know that β has a unique fixed point in T.) Computations by hand show that this is indeed correct for the primes below hundred. One might therefore be tempted to conjecture that, generally, γ acts transitively on S. If correct, this would provide a neat explanation for the uniqueness of the β-fixed point. Unfortunately this conjecture is incorrect. Its validity for small primes turns out to be yet another instance of the 'strong law of small numbers'. (If you have never heard of this law then you are urged to take a look at the beautiful article by Guy[1]).

We see this as follows.

For each fixed x, the number of points in S with the given first co-ordinate equals $d(\frac{p-x^2}{4})$. Therefore we have the formula

$$\#(S) = \sum_x d(\frac{p-x^2}{4}),$$

where the sum is over all odd numbers x in the range $1 \leq x < \sqrt{p}$. (Here $d(\cdot)$ is the usual divisor function : for $n \in \mathbb{N}$, $d(n)$ is the number of divisors of n including 1 and n.)

Let p be of the form $k^2 + 4$ (for an odd number k). Then in the iterates under γ of the point $a = (1, 1, \frac{p-1}{4})$, the first

co-ordinate increases in steps of 2 until the point $b = (k, 1, 1)$ is reached, then it decreases in steps of 2 until we reach the end point $(1, \frac{p-1}{4}, 1)$ of the orbit. This shows that in this case , the size k of the orbit T is related to the prime p by $p = k^2 + 4$. Also, the sum in the formula for $\#(S)$ given above has $(k+1)/2$ terms of which one term equals 1 while the remaining $(k-1)/2$ terms are ≥ 2. Since $d(n) = 2$ iff n is a prime, it follows that *for a prime of the form* $p = k^2+4$, γ *is transitive on* S *(i.e.,* $k = \#(S)$*) iff* $(p-x^2)/4$ *is a prime for all odd numbers* x *in the range* $1 \leq x < k$. This shows, for instance, that we do not have transitivity for $p = 229$.

Inefficiency of the Algorithm

Clearly, the α-β algorithm needs at most $\frac{1}{2}\#(S)$ steps. Since $d(n) = O(n^\epsilon)$ and the formula for $\#(S)$ has $O(p^{\frac{1}{2}})$ terms in it, the number of necessary iterations is $O(p^{\frac{1}{2}+\epsilon})$. The example of primes of the form square plus four (presumably there are infinitely many such primes) shows that this estimate is close to best possible. In [4], Wagon describes known algorithms whose complexity is polynomial in $\log p$, and the α-β algorithm compares very unfavourably. But it may be that we have looked at the worst case, and for some large class of primes its performance is much better. More over, it may be possible to significantly improve on the performance of the algorithm as follows. The set S can be partitioned into three parts on each of which γ is linear (the permutation $\hat{\gamma}$ is even better in this respect: we have a partition of $\hat{S}$ into two parts on each of which $\hat{\gamma}$ is linear.). The runs of iteration during which the iterates stay in the same piece of S may easily be combined into a single step.

A Combinatorial Lemma

The perceptive reader may have suspected by now that the theorem presented above does not have much to do with primes or their representations by squares. This is indeed correct, and the theorem is a manifestation of a combinatorial phenomenon. Namely, we have:

Lemma. *For any two involutions* α *and* β *of a finite set* S,

For more about Fermat, and the celebrated Fermat's last theorem, see *Resonance*, January 1996 and the Book Reviews section of *Resonance*, March 1999.

there are only three possibilities for any $\alpha\beta$-orbit: (i) neither involution has a fixed point in the orbit, or (ii) each of them has a unique fixed point in the orbit, or (iii) one of them has two fixed points in the orbit while the other has none.

At first glance, this statement may look very strange. (For readers with a reasonable amount of familiarity with groups and group actions, here is a hint for a group theoretic proof of this lemma: think of the group of isometries of a regular polygon.) But here is an elementary ('graph theoretic') proof.

Let $\gamma = \alpha\beta$. Fix a γ-orbit T If neither α nor β has a fixed point in T then there is nothing to prove: we are in case (i) of the lemma. So assume that one of these two involutions has at least one fixed point. Then, arguing as in the proof of the above theorem, one sees that T is fixed by both α and β. Thus T is a union of α-orbits as well as of β-orbits. If T is a singleton then we are in case (ii) and again there is nothing to prove. So we may assume that T has at least two elements. Hence no element of T is fixed by γ.

Now consider the graph G defined as follows. The vertices of G are the elements of T Two distinct elements x, y of T are joined by an edge in G if (and only if) $y = \alpha(x)$ or $y = \beta(x)$ (i.e., if $\{x, y\}$ is an orbit of one of the involutions). Clearly this is an undirected graph. Note that, for each x in T, $\alpha(x)$ and $\beta(x)$ are distinct elements of T – or else x would be fixed by γ, contrary to our assumption. It follows that each vertex x is of degree 1 or 2 in G (i.e., x is joined to one or two vertices) – according as x is or is not fixed by one (and only one) of the two involutions. Since we have assumed that at least one of them has a fixed point in T, it follows that G has at least one vertex of degree one. Also, since $\gamma = \alpha\beta$ is transitive on T (T is a γ-orbit!), it follows that G is connected. Now, here is the punch line: the only connected graphs with all vertices of degree ≤ 2 and at least one vertex of degree 1 are the paths. Hence G is a path. So G has exactly two vertices of degree 1 (the two ends of the path) and hence we are in case (ii) or (iii). This proves the lemma. (Exercise: Continue this argument to see that

if the elements of T are arranged on a circle according to the action of γ, then the two ends of G are placed opposite to each other. This explains the symmetry observed in the theorem.)

A Prime Testing Algorithm?

If $n \equiv 1 \pmod 4$ is a number (not necessarily a prime) which is not a perfect square, then S, α, β may be defined as before with n replacing p. What happens if one runs the α-β algorithm in this case ? Our combinatorial lemma shows that if we look inside the orbit T containing the fixed point $(1, 1, \frac{n-1}{4})$ of α, either we may find a fixed point of β and hence a representation of n as a sum of two squares or we find a second fixed point (x, x, z) of α and hence a nontrivial factorisation $n = x(x + 4z)$ of n. The second case is bound to occur if the square free part of n has a 3 (mod 4) factor (since in this case n has no representation as a sum of two squares). In the former case, of course, we are unable to decide whether n is a prime or not (for instance, this case occurs if n is a number of the form $k^2 + 4$, even when n is composite). If, however, we happen to know a two squares representation of n and the algorithm is lucky enough to produce a second representation, then we can still conclude that n is composite (because a prime has at most one such representation). Perhaps it will be interesting to characterise those numbers n for which the first case occurs.

Suggested Reading

[1] R K Guy, The strong law of small numbers, *Amer. Math. Monthly*, 95, 8, 697–711, 1988.

[2] I Niven and H S Zuckerman, *An Introduction to the Theory of Numbers*, third edition, Wiley, 1972.

[3] S A Shirali, On Fermat's two-squares theorem, *Resonance*, Vol.2, No. 3, 69–73, 1997.

[4] S Wagon, The euclidean algorithm strikes again, *Amer. Math. Monthly*, 97, 2, 125–126, 1990.

[5] D Zagier, A one-sentence proof that every prime $p \equiv 1$ (mod 4) is a sum of two squares, *Amer. Math. Monthly*, 97, 2, 144, 1990.

Address for Correspondence
Bhaskar Bagchi
Math. Stat. Unit
Indian Statistical Institute
Bangalore 560 059, India.
Email:bbagchi@isibang.ac.in

"Any biographical sketch of Pierre de Fermat (1601–1665) is bound to be short. His life spanned the first two-thirds of the seventeenth century but was, truth to tell, rather dull. He never held an appointment at a university nor occupied a chair at a royal academy. By training a lawyer, by profession a magistrate, Fermat published almost nothing during his lifetime, instead conveying his ideas through correspondence and unpublished manuscripts. Because he was not a professional mathematician, Fermat has been dubbed the 'prince of amateurs'. But, if by 'amateur' we mean 'marginally talented novice', the moniker is totally inaccurate."

William Dunham
The Mathematical Universe