PI MU EPSILON: PROBLEMS AND SOLUTIONS: FALL 2018

STEVEN J. MILLER (EDITOR)

1. Problems: Spring 2020

This department welcomes problems believed to be new and at a level appropriate for the readers of this journal. Old problems displaying novel and elegant methods of solution are also invited. Proposals should be accompanied by solutions if available and by any information that will assist the editor. An asterisk (*) preceding a problem number indicates that the proposer did not submit a solution.

Solutions and new problems should be emailed to the Problem Section Editor Steven J. Miller at sjm1@williams.edu; proposers of new problems are strongly encouraged to use LaTeX. Please submit each proposal and solution preferably typed or clearly written on a separate sheet, properly identified with your name, affiliation, email address, and if it is a solution clearly state the problem number. Solutions to open problems from any year are welcome, and will be published or acknowledged in the next available issue; if multiple correct solutions are received the first correct solution will be published (if the solution is not in LaTeX, we are happy to work with you to convert your work). Thus there is no deadline to submit, and anything that arrives before the issue goes to press will be acknowledged. Starting with the Fall 2017 issue the problem session concludes with a discussion on problem solving techniques for the math GRE subject test.

Earlier we introduced changes starting with the Fall 2016 problems to encourage greater participation and collaboration. First, you may notice the number of problems in an issue has increased. Second, any school that submits correct solutions to at least two problems from the current issue will be entered in a lottery to win a pizza party (value up to $100). Each correct solution must have at least one undergraduate participating in solving the problem; if your school solves $N \geq 2$ problems correctly your school will be entered $N \geq 2$ times in the lottery. Solutions for problems in the Spring Issue must be received by October 31, while solutions for the Fall Issue must arrive by March 31 (though slightly later may be possible due to when the final version goes to press, submitting by these dates will ensure full consideration). The winning school from the Fall problem set is **ADD**. Also in the last issue one problem solver was accidentally omitted; Problem #1359 was also solved by Kenneth Davenport (SCI-Dallas, Dallas, PA).

Finally, in this issue we are introducing a new feature. Each year a distinguished mathematician gives the J. Sutherland Frame Pi Mu Epsilon Lecture at MathFest. In 2019 that speaker was Alice Silverberg, Distinguished Professor at the University of California, Irvine; her talk is available at , and Problem #1366 is inspired by her lecture. For 2020 the speaker will be Florian Luca from the University of the Witwatersrand. The abstract for his talk, *Arithmetic and Digits*, is the following: *In our recent paper in the Monthly (October, 2019) with Pante Stănică, we looked at perfect squares which arise when concatenating*


FIGURE 1. Pizza motivation; can you name the theorem that's represented here?

two consecutive positive integers like $183184 = 428^2$ *with the smaller number to the left, or* $98029801 = 9901^2$ *with the larger number to the left. My talk will present variations on this topic with the aim of providing the audience with examples of numbers which are both arithmetically interesting (like perfect squares) while their digital representations obey some regular patterns. The examples will not be limited to perfect squares, but will also include other old friends like Fibonacci numbers and palindromes.*

#1366: *Proposed by Alice Silverberg (University of California, Irvine).*

Solutions to the following problems would have interesting applications to cryptography and computer security.

Problem 1: Find a way for four or more parties to create a shared secret (to use as a secret key in a symmetric key encryption scheme), with one round of broadcasts through an insecure channel.

Diffie-Hellman key agreement [2] solves this problem for two parties (see below). Pairings on elliptic curves or abelian varieties lead to a solution for three parties [3]. The problem is open for four or more parties [1].

The next problem, which was raised in [1], is a group-theoretic approach to Problem 1 that would generalize the solutions that are known in the cases of two and three parties.

Problem 2: For each (or some) fixed integer $n \geq 3$, find a large prime number p, groups G and G_T of order p, a generator g of the cyclic group G, and a function

$$e : G^n \to G_T$$

such that:

(a) **it is easy to compute** $e(g_1, \ldots, g_n)$ **for all** $g_i \in G$,
(b) $e(g_1^{a_1}, \ldots, g_n^{a_n}) = e(g_1, \ldots, g_n)^{a_1 \cdots a_n}$ **for all** $g_i \in G$ **and all integers** $a_1, \ldots, a_n$, **and**
(c) **it is hard to compute**

$$e(g, \ldots, g)^{a_1 \cdots a_{n+1}}$$

for random unknown integers $a_1, \ldots, a_{n+1}$, **even given** $g, g^{a_1}, g^{a_2}, \ldots, g^{a_{n+1}}$.

A solution to Problem 2 would solve Problem 1, as follows. Suppose we have $n+1$ parties. The parties publicly agree on a solution to Problem 2, namely, a prime p, groups G and G_T, a map e, and a generator g of G that satisfy (a), (b), and (c). For $i = 1, \ldots, n+1$, party

i chooses a random secret integer a_i between 1 and p, computes g^{a_i}, and broadcasts that element of the group G. The shared secret is

$$e(g, \ldots, g)^{a_1 \cdots a_{n+1}} \in G_T.$$

All of the $n+1$ parties involved can compute it; the i^{th} party computes it by computing

$$e(g^{a_1}, \ldots, g^{a_{i-1}}, g^{a_{i+1}}, \ldots, g^{a_{n+1}})^{a_i},$$

which gives the desired result by (b). By (c), it is hard for anyone else to learn this shared element of the target group G_T.

For two parties, let $n = 1$, take G to be a (large) prime order subgroup of the multiplicative group of a finite field, let $G_T = G$, and let e be the identity map on G. The above protocol recovers Diffie-Hellman key agreement. Namely, Alice and Bob publicly share the prime p, the group G, and a generator g. Alice (respectively, Bob) chooses a random secret integer a (respectively, b) between 1 and p, computes g^a (respectively, g^b), and broadcasts it. The shared secret is g^{ab}, which Alice obtains by computing $(g^b)^a$ and Bob obtains by computing $(g^a)^b$. The security is based on the assumption that G was chosen so that no one else can compute g^{ab}, for random unknown a and b, even if they know g, g^a, and g^b.

Taking G to be a prime order subgroup of the group of points on a suitable elliptic curve over a finite field gives Elliptic Curve Diffie-Hellman.

For three parties, let $n = 2$. Problem 2 can be solved by taking G to be a (large) prime order subgroup of the group of points on a suitable elliptic curve E over a finite field F, and letting e be a certain (suitably modified Weil or Tate) pairing associated to E. The group G_T will be a subgroup of the multiplicative group of a certain extension field of the finite field F.

Beyond that, these problems are an area of current research.

Applications of Problem 2 to key agreement, broadcast encryption, and digital signatures were given in [1]. That paper also gave evidence that when $n \geq 3$ it might be difficult to use algebraic geometry to find a solution to Problem 2, and therefore it might be productive to look more broadly than number theory and algebraic geometry for solutions to Problems 1 and 2.

REFERENCES

[1] Dan Boneh and Alice Silverberg, *Applications of multilinear forms to cryptography*, in Topics in Algebraic and Noncommutative Geometry: Proceedings in Memory of Ruth Michler, Contemporary Mathematics **324**, American Mathematical Society, Providence (2003), 71–90.
[2] Whitfield Diffie and Martin E. Hellman, *New Directions in Cryptography*, IEEE Transactions on Information Theory **22** (1976), 644–654.
[3] Antoine Joux, *A one round protocol for tripartite Diffie-Hellman*, Proc. ANTS IV, Lecture Notes in Computer Science **1838** (2000), 385–394.

#1367: *Proposed by Thomas Garrity, Steven J. Miller and Chenyang Sun (Williams College).*

We say a positive integer is a side-into-hypothesis number if it is both a side of a right triangle with integer lengths and the hypotenuse of a triangle with integer sides. Thus 5 is such a number, from the 3–4–5 triangle and the 5–12–13.

(a) Prove or disprove: there are infinitely many side-into-hypothesis numbers.

(b) Prove or disprove: there exists an infinite sequence of positive integers $\{x_n\}$ such that x_n is a side of a right triangle with x_{n+1} as a hypotenuse.

#1368: *Proposed by Steven J. Miller (Williams College).*
Consider an $n \times n$ chessboard. It is a famous (and difficult!) problem to place n queens and maximize the number of squares they do not attack; however, if we only care about the percentage in the limit then it is significantly easier. Prove that as $n \to \infty$ we can place the queens in such a way so that 100% of the squares are not attacked. Note this does not mean every square is safe; it means if $s(n)$ is the maximum number of safe squares, then $s(n)/n^2 \to 1$. Can you obtain good upper and lower bounds for $n^2 - s(n)$?

#1369: *Proposed by Steven J. Miller and Chenyang Sun (Williams College).*
Consider an $n \times n$ chessboard. The previous problem is greatly simplified if instead of queens we place rooks. Determine an optimal placement of n rooks to maximize the number of safe squares.

#1370: *Proposed by Eugen J. Ionascu (Columbus State University).*
Consider three points A, B and C chosen uniformly at random inside of the region (see Figure 1 below, the yellow region)

$$\mathcal{R}_{r,\alpha} = \{(x,y) \in \mathbb{R}^2 | x = t\cos\theta, y = t\sin\theta, \ \pi \geq |\theta| \geq \alpha, t \in [0,r]\},$$

$$\alpha = \pi a, \quad a \in [0, \tfrac{1}{2}], \quad r > 0;$$

(thus we are choosing from the uniform distribution with respect to the area). Show that the probability that the resulting triangle $\triangle ABC$ contains the origin $O(0,0)$ is equal to

$$\mathcal{P}_a = \frac{(1+a)(1-2a)^2}{4(1-a)^3}.$$

Figure 1

GRE Practice #6: The following is essentially the famous German Tank Problem. There are N tanks, labeled $1, 2, \ldots, N$, and k distinct serial numbers $s_1, s_2, \ldots, s_k$ are observed,

4

with $m = \max s_i$ the largest value. If every subset of size k is equally likely, what is the best predictor for N as a function of m and k?
(a) mk (b) $m(1 + 1/k)$ (c) $m(1 - 1/k)$ (d) $m(1 + 1/k) - 1$ (e) $m(1 - 1/k) - 1$.

2. SOLUTIONS

#1364: *Proposed by Ron Evans (UCSD) and Steven J. Miller (Williams College).*
(a) Two people are playing pool, which has 15 balls labeled from 1 to 15. On the first turn someone sinks three of these. What is the probability that none of the three numbers of the sunk balls are adjacent? Thus if they sink the 4, 8 and 12 that would work, but if they sank the 4, 8 and 9 it would not.
(b) Assume now they play super-pool, where there are 2020 balls on the table, and now someone sinks 314 balls on their first turn. What is the probability none of the 314 sunk have adjacent numbers?
(Bonus: More generally, what would the answer be if there were N balls and they sink k on their first turn?)

*Solution by **Cal Poly Pomona Problem Solving Group**.*

We will use the inclusion-exclusion principle of counting. There are 14 pairs of adjacent numbers in the sequence

$$1\ 2\ 3\ 4\ 5\ 6\ 7\ 8\ 9\ 10\ 11\ 12\ 13\ 14\ 15.$$

and the number of choices for the third ball is given by $\binom{13}{1}$. So there are $14 \cdot \binom{13}{1}$ ways to choose three balls containing at least one adjacent pair. However, this over-counts the combinations which contain a sequence of three numbers (e.g. 1,2,3). To account for this, we subtract the number of ways to get three in a row. In this sequence, there are 13 ways to get three in a row. So the total number of ways to get at least one adjacent pair is

$$14 \cdot \binom{13}{1} - 13$$

.

The probability that none of the three numbers of the sunk balls are adjacent is

$$1 - \frac{14 \cdot \binom{13}{1} - 13}{\binom{15}{3}} = \frac{22}{35}$$

(b) Assume now they play super-pool, where there are 2020 balls on the table, and now someone sinks 314 balls on their first turn. What is the probability none of the 314 sunk balls have adjacent numbers? (Bonus: more generally, what would the answer be if there were N balls and they sink k on their first turn?)

Define $a_{N,k}$ as the number of ways that k out of N balls can be sunk with no adjacencies. The following table shows $a_{N,k}$ for $N, k \in \{1, 2, ..., 8\}$. Notice that the diagonals of the table give the coefficients of the binomial expansion $(1 + x)^n$, $n \in \mathbb{N}$.

k \ N	1	2	3	4	5	6	7	8
0	1	1	1	1	1	1	1	1
1	1	2	3	4	5	6	7	8
2	0	0	1	3	6	10	15	21
3	0	0	0	0	1	4	10	20
4	0	0	0	0	0	0	1	5
5	0	0	0	0	0	0	0	0
6	0	0	0	0	0	0	0	0
7	0	0	0	0	0	0	0	0
8	0	0	0	0	0	0	0	0

We conjecture that $a_{N,k} = \binom{N-k+1}{k}$ for $k \leq \lceil \frac{N}{2} \rceil$, and $a_{N,k} = 0$, for $k > \lceil \frac{N}{2} \rceil$. Notice that the number of subsets of k elements of $\{1, 2, ..., N\}$ with no consecutive numbers satisfies the recursion

$$a_{N,k} = a_{N-1,k} + a_{N-2,k-1}$$

with the conditions that $a_{N,0} = 1$ for every N,

$$a_{1,k} = \begin{cases} 1 & \text{if } k = 0 \text{ or } 1 \\ 0 & \text{otherwise} \end{cases},$$

and

$$a_{2,k} = \begin{cases} k+1 & \text{if } k = 0 \text{ or } 1 \\ 0 & \text{otherwise}. \end{cases}$$

To see this, consider such a subset. If the subset does not include the number 1, then its k elements can be chosen in $a_{N-1,k}$ ways from $\{2, ..., N\}$. If the subset contains 1, then 2 cannot be in it, so the remaining $k - 1$ elements must be chosen from $\{3, ..., N\}$, which can be done in $a_{N-2,k-1}$ ways.

To prove our conjecture, we will proceed by induction using the recursion above. The fact that $a_{N,k} = 0$ for $k > \lceil \frac{N}{2} \rceil$, $\forall N, k$ is obvious. The initial conditions provide our basis step. Assume that $a_{M,l} = \binom{M-l+1}{l}$ is satisfied for $M \in \{1, 2, ..., N-1\}$, $\forall l \leq \lceil \frac{M}{2} \rceil$. Assume $1 < k < \lceil \frac{N}{2} \rceil$. Then

$$
\begin{aligned}
a_{N-1,k} + a_{N-2,k-2} &= \binom{N-k}{k} + \binom{N-k}{k-1} \\
&= \binom{N-k+1}{k}, \text{ since } k \leq \left\lceil \frac{N-1}{2} \right\rceil \text{ and } k - 1 \leq \left\lceil \frac{N-2}{2} \right\rceil \\
&= a_{N,k}
\end{aligned}
$$

Consider the case where $k = \lceil \frac{N}{2} \rceil$.
If N is even, we let $N = 2s$, $s \in \mathbb{N}$ and $k = \lceil \frac{N}{2} \rceil = s$. Then

$$
\begin{aligned}
a_{N-1,k} + a_{N-2,k-2} &= 1 + \binom{s}{s-1} \\
&= 1 + \binom{k}{k-1} \\
&= \binom{k+1}{k} \\
&= \binom{N-k+1}{k} \\
&= a_{N,k}
\end{aligned}
$$

7

If N is odd, we let $N = 2s + 1$, $s \in \mathbb{N}$ and $k = \lceil \frac{N}{2} \rceil = s + 1$. Then

$$a_{N-1,k} + a_{N-2,k-2} = 0 + \binom{s}{s}$$
$$= 1$$
$$= \binom{N-k+1}{k}$$
$$= a_{N,k}$$

Thus the probability that none of the k sunk balls out of N total balls are adjacent is

$$\frac{a_{N,k}}{\binom{N}{k}}.$$

Finally, the probability that none of the 314 sunk balls out of 2020 balls are adjacent is

$$\frac{a_{2020,314}}{\binom{2020}{314}} = \frac{\binom{1707}{314}}{\binom{2020}{314}}$$
$$= 7.1223 \cdot 10^{-26}$$

#1365: *Proposed by Ron Evans (UCSD) and Steven J. Miller (Williams College).*
We generalize the previous problem. We have N balls, numbered 1 to N.
(a) Assume exactly k balls are sunk on the first turn. What is the probability that there are no three consecutive numbers among them?
(b) Assume exactly k balls are sunk on the first turn, so there are $N - k$ that are not sunk. What is the probability that there are no three consecutive numbers among the k sunk and also there are no three consecutive numbers among the $N - k$ not sunk?

The following, by Ivan Ventura of Cal Poly Pomona, is related to Problems #1364 and #1365.

Claim 2.1. *Let $A_N = \{1, 2, \ldots, N - 1, N\}$. The number of subsets of size k of A_N with no m consecutive elements is given by the coefficient of x^k in $(1 + x + \ldots + x^{m-1})^{N-k+1}$, which is*

$$\sum_{j=0}^{N-k+1} (-1)^j \binom{N-k+1}{j} \binom{N-jm}{N-k}.$$

Proof. Let $\mathcal{S} = \{A \subset A_N \mid A \text{ does not contain } m \text{ consecutive elements}, |A| = k\}$.
Consider the subsets of A_N, defined by $B_i = \{i, i+1, \ldots, i+k-1\}$ for $1 \leq i \leq N - k + 1$ and define $\mathcal{B}$ to be the set of all selections of k elements of $\{B_1, B_2, \ldots, B_{N-k+1}\}$, where no B_i is chosen more than $m - 1$ times.

Define the function $\varphi : \mathcal{B} \to \mathcal{S}$ as follows. Let $\beta \in \mathcal{B}$, where

$$\beta = B_{i_1}, B_{i_2}, B_{i_3}, \ldots, B_{i_k}$$

where $i_1 \leq i_2 \leq \ldots \leq i_k$. Define

$$\varphi(\beta) = \{i_j + j - 1 \mid 1 \leq j \leq k\} = \{i_1, i_2 + 1, i_3 + 2, \ldots, i_k + k - 1\}.$$

We claim φ is a bijective function. First, note that since $i_j \leq i_{j+1}$, we have $i_j + j - 1 < i_{j+1} + j$, so $\varphi(\beta)$ has k elements. Furthermore, $i_j + j - 1$ and $i_{j+1} + j$ are consecutive if and only if $i_j = i_{j+1}$. Since each B_i can appear at most $m - 1$ times, we see that $\varphi(\beta)$ is in $\mathcal{S}$.

Take $\{a_1, a_2, \ldots, a_k\} \in \mathcal{S}$, where the a_j are listed in increasing order. Let $i_j = a_j - j + 1$ and define

$$\psi(\{a_1, a_2, \ldots, a_k\}) = B_{i_1}, B_{i_2}, B_{i_3}, \ldots, B_{i_k}.$$

It is easy to see that $\psi = \varphi^{-1}$, so φ is bijective. Thus it suffices to find $|\mathcal{B}|$.

$|\mathcal{B}|$ is the coeffcient of x^k in the function $(1 + x + \ldots + x^{m-1})^{N-k+1}$, which is the generating function for choosing k elements from a set with $N - k + 1$ elements, where each element can be chosen at most $m - 1$ times.

Example. To illustrate the proof above, we will evaluate φ when $N = 6$, $k = 3$. In this case we have that $A_6 = \{1, 2, 3, 4, 5, 6\}$, $B_1 = \{1, 2, 3\}$, $B_2 = \{2, 3, 4\}$, $B_3 = \{3, 4, 5\}$, and $B_4 = \{4, 5, 6\}$.
We will start in the case that $m = 2$. In this case, the elements of $\mathcal{B}$ are three element selections of $\{B_1, B_2, B_3, B_4\}$ that contain at most one copy of each B_i. $\mathcal{S}$ contains all three element subsets of A_6 that do not have two consecutive numbers. To find $\varphi(B_1, B_3, B_4)$ and $\varphi(B_1, B_2, B_3)$ we make the following arrays. The diagonal elements of the arrays are the elements of the output.

$$
\begin{array}{llll}
B_1: & ① & 2 & 3 \\
B_3: & 3 & ④ & 5 \\
B_4: & 4 & 5 & ⑥
\end{array}
\qquad\qquad
\begin{array}{llll}
B_1: & ① & 2 & 3 \\
B_2: & 2 & ③ & 4 \\
B_3: & 3 & 4 & ⑤
\end{array}
$$

Thus we have $\varphi(B_1, B_3, B_4) = \{1, 4, 6\}$ and $\varphi(B_1, B_2, B_3) = \{1, 3, 5\}$. Here is the complete list of values of φ is given below in this case.

$$
\begin{aligned}
\varphi(B_1, B_2, B_3) &= \{1, 3, 5\}, & \varphi(B_1, B_2, B_4) &= \{1, 3, 6\} \\
\varphi(B_1, B_3, B_4) &= \{1, 4, 6\}, & \varphi(B_2, B_3, B_4) &= \{2, 4, 6\}
\end{aligned}
$$

When $m = 3$, $\mathcal{B}$ also includes three element selections of $\{B_1, B_2, B_3, B_4\}$ that contain two copies of one of the B_i and $\mathcal{S}$ contains subsets of A_6 that have two but not three consecutive numbers. The process for evaluating φ in this case is no different. Here we compute $\varphi(B_1, B_1, B_2)$ and $\varphi(B_2, B_4, B_4)$ by making the same arrays.

$$
\begin{array}{llll}
B_1: & ① & 2 & 3 \\
B_1: & 1 & ② & 3 \\
B_2: & 2 & 3 & ④
\end{array}
\qquad\qquad
\begin{array}{llll}
B_2: & ② & 3 & 4 \\
B_4: & 4 & ⑤ & 6 \\
B_4: & 4 & 5 & ⑥
\end{array}
$$

Thus we have $\varphi(B_1, B_1, B_2) = \{1, 2, 4\}$ and $\varphi(B_2, B_4, B_4) = \{2, 5, 6\}$. In this case both $\mathcal{B}$ and $\mathcal{S}$ have 16 elements each so we omit the complete list.

GRE Practice #6: The following is essentially the famous German Tank Problem. There

are N tanks, labeled $1, 2, \ldots, N$, and k distinct serial numbers $s_1, s_2, \ldots, s_k$ are observed, with $m = \max s_i$ the largest value. If every subset of size k is equally likely, what is the best predictor for N as a function of m and k?
(a) mk (b) $m(1 + 1/k)$ (c) $m(1 - 1/k)$ (d) $m(1 + 1/k) - 1$ (e) $m(1 - 1/k) - 1$.

While this problem can be solved through combinatorial manipulations, that is time consuming. Instead, let us look at extreme cases and eliminate four of the five answers.

First, (a) is clearly out as if k is large than m will be large (if $k \geq N/2$ then $m \geq N/2$), and thus mk will exceed N!

We now come to four answers that are very similar. If we take the special case $k = N$ then $m = N$ (every tank is observed) and thus our prediction should be N. The answers we get in this special case are $N + 1$ for (b), $N - 1$ for (c), N for (d) and $N - 2$ for (e), which suggests (d) as the answer.

We could also look at $k = 1$ to gain some reasonableness. Note that we will never guess less than m for the number of tanks, so it makes sense to look at a guess of the form $m + g(m, k)$ for some function g. As m increases we should predict a higher value of N, so it makes sense that g is increasing with m. Similarly, if k increases then we've observed more tanks and we're more confident that the largest observed serial number is close to N, and thus the boost should be decreasing with k. This leads to guessing $g(m, k) = \alpha m/k + \beta$; the answer turns out to be $\alpha = -\beta = 1$. Returning to taking $k = 1$, it's likely that our observed m is around $N/2$, so we would want to double it. While this eliminates (c) and (e), both (b) and (d) survive (which is not surprising, as they differ by a constant).

This problem is an excellent example of the power of spending some time thinking about a problem and getting a feel for the answer.

Email address: sjm1@williams.edu

PROFESSOR OF MATHEMATICS, DEPARTMENT OF MATHEMATICS AND STATISTICS, WILLIAMS COLLEGE, WILLIAMSTOWN, MA 01267