14
Finite Geometries, Codes, Latin Squares, and Other Pretty Creatures

14.1   Small Exotic Worlds

The Fano plane is a really small world (Figure 14.1). It only has 7 points, which form 7 lines. In the figure, 6 of these lines are straight and one is circular; but for the inhabitants of this tiny world (the Fanoans), the straight and curved lines look the same. Also, in our figure it seems that the circular line intersects some of the straight lines twice; but the intersection points that are not marked are not real intersection points, they are there only because we want to draw a picture of this very different world in our Euclidean plane.
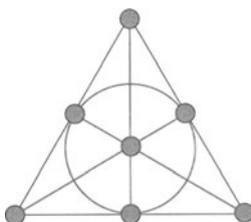

FIGURE 14.1. The Fano plane.

The Fanoans are very proud of their world. They say that it is tiny but perfect in various ways. They point out that

(a) *through any two points their world has a unique line.*

If you check this out and admit it's true, but reply that this also holds in our own world, they go on and boast that

(b) *any two lines have exactly one intersection point.*

This is certainly not true in our Euclidean world (we have parallel lines), so we have to admit that this is nice indeed. But then we can draw a new figure (Figure 14.2) and point out that this rather uninteresting construction also has properties (a) and (b). But the Fanoans are ready for this attack: "It would be enough if we pointed out that in our world,
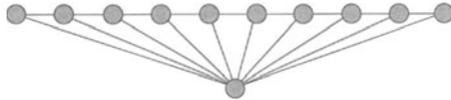
(c) *every line has at least 3 points.*"

FIGURE 14.2. An ugly plane

Our Fanoan friend goes on: "Theoretical physicists have shown that just from (a), (b), and (c), many properties of our world can be derived. For example,

(d) *all our lines have the same number of points.*

"Indeed, let K and L be two lines. By (b), they have an intersection point p; by (c), they contain other points (at least two, but we only need one right now). Let us select a point q from K and a point r from L that are both different from p. By (a), there is a line M through q and r. Let s be a third point on M (which exists by (c); see Figure 14.3).
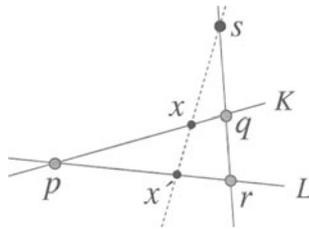

FIGURE 14.3. All lines have the same number of points.

"Consider any point x on K, and connect it to s. This line intersects L at a point x' (we can think of this as projecting K onto L from center s). Conversely, given x', we get x by the same construction. Thus this projection establishes a bijection between K and L, and so they have the same number of points.

"This argument also shows that

(e) *all our points have the same number of lines through them.*

"No doubt you can prove this yourself, if you study the previous argument carefully."

Our intelligent Fanoan friend adds, "These theoretical physicists (obviously having a lot of time on their hands) also determined that the fact that the number of points on each line is 3 does *not* follow from (a), (b), and (c); they say that alternative universes can exist with 4, 5 or 6 points on a line. Imagining this is beyond me, though! But they say that no universe could have 7 points on a line; there could be universes with 8, 9, and 10 points on each line, but 11 points are impossible again. This is, of course, a favorite topic for our science fiction writers."

The Fanoans hate Figure 14.2 for another reason: they are true egalitarians, and the fact that one point is special is intolerable in their society. You may raise here that the Fano plane also has a special point, the one in the middle. But they immediately explain that this is again an artifact of our drawing. "In our world,

(f) *all points and all lines are alike*

in the sense that if we pick any two points (or two lines), we can just rename every point so that one of them becomes the other, and nobody will notice the difference." You may trust them about this, or you may verify this claim by solving exercise 14.1.7.

Let us leave the Fano plane now and visit a larger world, the *Tictactoe plane*. This has 9 points and 12 lines (Figure 14.4). We have learned from our excursion to the Fano plane that we have to be careful with drawing these strange worlds, and so we have drawn it in two ways: In the second figure, the first two columns are repeated, so that the two families of lines (one leaning right, one leaning left) can be seen better.[1]

The Tictacs boast that they have a much more interesting world than the Fanoans. It is still true that any two points determine a single line; but two lines may be intersecting or parallel (which simply means that they don't intersect). One of our Tictac friends explains: "I heard that your mathematicians have been long concerned about the statement that

(g) *for any line and any point not on the line, there is one and only one line that goes through the point and is parallel to the given line.*

They called it the Axiom of Parallels or Euclid's Fifth Postulate. They were trying to prove it from other basic properties of your world, until eventually

---

[1]If you have learned about matrices and determinants, you may recognize the following description of this world: if we think of the points in this plane as the entries of a 3×3 matrix, then the lines are the rows and columns of the matrix and expansion terms of its determinant. The second drawing in the figure corresponds to Sarrus's Rule in the theory of determinants.
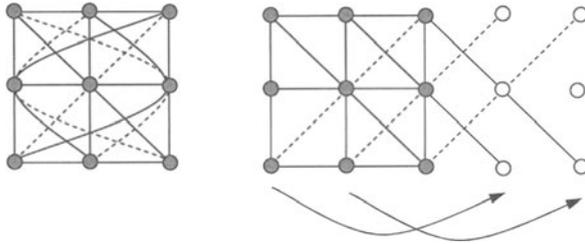
FIGURE 14.4. The Tictactoe plane.

they showed that this cannot be done. Well, this is true in our world, and since our world is finite, it is easy to check that it is true." (We hope our readers will accept the challenge.)

"We have 3 points on each line just as in the Fano plane, but we have 4 lines through each point—more than those Fanoans. All of our points are alike, and so are all of our lines (even though the way you draw them in your own geometry seems to differentiate between straight and curved lines)."
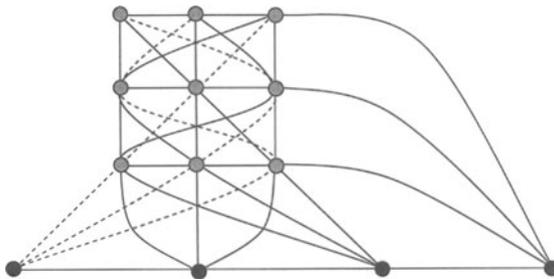
FIGURE 14.5. Extending the Tictactoe plane by 1 line and 4 points at infinity.

When we point out how the Fanoans love their property (b), our Tictac guide replies, "We could easily achieve this ourselves. All we'd have to do is add 4 new points to our plane. Each line should go through exactly one of these new points; parallel lines should go through the same new point, nonparallel lines through different new points. And we could even things out by declaring that the 4 new points also form a line. We could call the new points 'points at infinity' and the new line, the 'line at infinity' (Figure

14.5). Then we would have properties (a) through (g) ourselves.[2] But we prefer to distinguish between finite and infinite points, which makes our world more interesting."
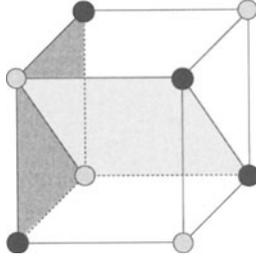

FIGURE 14.6. The Cube space.

Finally, we visit a third tiny world called the *Cube space* (Figure 14.6). While this one has only 8 points, it is much richer than the Tictactoe plane in one sense: It is 3-dimensional! Its lines are uninteresting: Every line has just 2 points, and any two points form a line. But it has planes! In our (deficient) Euclidean picture, the points are arranged as the vertices of a cube. The planes are (1) 4-tuples of points forming a face of the cube (there are six of these), (2) 4-tuples of points on two opposite edges of the cube (6 of these again), (3) the four black points, and (4) the four light points.

The Cube space has the following very nice properties (whose verification is left to the reader):

(A) *Any three points determine a unique plane.*

(The Cubics remark at this point, "In your world, this is only true if the three points are not on a line. Luckily, we never have three points on a line!")

(B) *Any two planes are either parallel (nonintersecting), or their intersection is a line.*

(C) *For any plane and any point outside it, there is exactly one plane through the given point parallel to the given plane.*

(D) *Any two points are alike.*

(E) *Any two planes are alike.*

This last claim looks so unlikely, considering that we have such different kinds of planes, that a proof is in order. Let us label the points of the cube by $A, \ldots, H$ as in Figure 14.7. It is clear that the faces of the cube are alike

---

[2]This construction appending new points at infinity can be carried out in our own Euclidean plane, leading to an interesting kind of geometry, called *projective geometry*.

(one can be moved onto any other by appropriately rotating the cube, and this rotation maps all the other planes onto planes). It is also clear that the planes formed by two opposite edges are alike, and the all-black plane is like the all-light plane (reflecting the cube in its center will interchange black and light vertices).

This was the easy part. Now we do a trickier transformation: We interchange E with F and G with H (Figure 14.7). What happens to the planes?
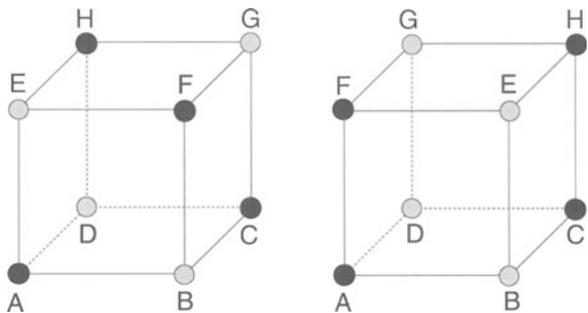

FIGURE 14.7. Why different-looking planes are alike.

Some of them are not changed (even though their points will change places): The top, bottom, front, and back faces and the planes $ABGH$ and $CDEF$ are mapped onto themselves. The left face $ADEH$ is mapped onto the plane $ADFG$, and vice versa. Similarly, the right face $BCFG$ is mapped onto the plane $BCEH$, and vice versa. The all-black plane is mapped onto the plane $ACEG$, and vice versa. The all-light plane is mapped onto the plane $BDFH$, and vice versa.

Thus all planes are accounted for. We make two observations:

— *Every plane is mapped onto a plane*, and so if we relabel the cube as above, the Cubics won't notice the difference!

— There is a face-plane mapped onto an opposite-edge-plane, and there is an opposite-edge-plane mapped onto the all-black plane. This implies that the Cubics cannot see any difference between these three types of planes.

14.1.1 Fanoan philosophers have long been troubled by the difference between points and lines. There are many similarities (for example, there are 7 of each); why are they different? Represent each line by a new point; for each old point, take the 3 lines through it, and connect the 3 new points representing them by a (new) line. What structure do you get?

14.1.2 The Fanoans call a set of 3 points a *circle* if they are not on a line. For example, the 3 vertices in figure 14.1 form a circle. They call a line a *tangent* to

the circle, if it contains exactly one point of the circle. Show that at every point of a circle there is exactly one tangent.

14.1.3 The Fanoans call a set of 4 points a *hypercircle* if no 3 of them are a line. Prove that the 3 points not on a hypercircle form a line, and vice versa.

14.1.4 Representatives of the 7 points in the Fano plane often vote on different issues. In votes where everyone has to vote yes or no, they have a strange rule to count ballots, however: it is not the majority who wins, but "line wins": if all 3 points on a line want something, then this is so decided. Show that (a) it cannot happen that contradictory decisions are reached because the points on another line want the opposite, and (b) in every issue there is a line whose points want the same, and so decision is reached.

14.1.5 Prove that the Tictactoe plane, extended with elements at infinity, satisfies all properties (a)–(d).

14.1.6 In response to the Tictacs' explanation about how they could extend their world with infinite elements, the Fanoans decided to declare one of their lines the "line at infinity," and the points on this line "points at infinity." The remaining 4 points and 6 lines form a really tiny plane. Will property (g) of parallel lines be valid in this geometry?

14.1.7 We want to verify the claim of the Fanoans that all their points are alike, and rearrange the points of the Fano plane so that the middle point becomes (say) the top point, but lines remain lines. Describe a way to do this.

14.1.8 Every point of the Cube space is contained in 7 lines and 7 planes. Is this numerical similarity with the Fano plane a coincidence?

14.2 Finite Affine and Projective Planes

It is time to leave our excursion to imaginary worlds and introduce mathematical names for the structures we studied above. If we have a finite set V whose elements are called *points*, and some of its subsets are called *lines*, and (a), (b), and (c) above are satisfied, then we call it a *finite projective plane*. The Fano plane (named after the Italian mathematician Gino Fano) is one projective plane (we'll see that it has the least possible number, 7, of points), and another one was constructed by the Tictac theoretical physicists by adding to their world 4 points and a line at infinity.

We have heard the proof from Fanoan scientists that every line in a finite projective plane has the same number of points; for reasons that should become clear soon, this number is denoted by $n + 1$, where the positive integer n is called the *order* of the plane. So the Fano plane has order 2, and the extended Tictactoe plane has order 3. The Fanoans also know that

if a finite projective plane has order n, then $n + 1$ lines go through every point of it.

14.2.1 Prove that a finite projective plane of order n has $n^2 + n + 1$ points and $n^2 + n + 1$ lines.

We can also study structures consisting of points and lines, where (a) and the "Axiom of Parallels" (g) are assumed; to exclude trivial (ugly?) examples, we assume that each line has at least 2 points. Such a structure is called a *finite affine plane*.

The "Axiom of Parallels" implies that all lines parallel to a given line L are also parallel to each other (if two of them had a point p in common, then we would have two lines through p parallel to L). So all lines parallel to L form a "parallel class" of mutually parallel lines, which cover every point in the affine plane.

Affine versus projective planes. The construction used by the Tictacs to extend their plane can be used in general. To every parallel class of lines we append a new "point at infinity" and create a new "line at infinity" going through all points at infinity. Then (a) remains satisfied: Two "ordinary" points are still connected by a line (the same line as before), two "infinite" points are connected by a line (the "infinite" line), and an ordinary and an infinite point are connected by a line (the parallel class belonging to the infinite point contains a line through the given ordinary point). It is even easier to see that we do not have two lines through any pair of points.

Furthermore, (b) is satisfied: Two ordinary lines intersect each other unless they are parallel, in which case they share a point at infinity; an ordinary line intersects the infinite line at its point at infinity. We leave it to you to check (c), (d), and (e) (condition (f) does not hold for every finite projective plane; it is a special feature of the Fano plane and some other projective planes).

The construction in Exercise 14.1.6 is again quite general. We can take any finite projective plane, and call any of its lines along with the points on this line "infinite." The remaining points and lines form a finite affine plane. So in spite of the rivalry between the Fanoans and Tictacs, finite affine planes and projective planes are essentially the same structures.

To sum up, we have the following theorem.

Theorem 14.2.1 *Every finite affine plane can be extended to a finite projective plane by adding new points and a single new line. Conversely, from every projective plane we can construct an affine plane by deleting any line and its points.*

A projective plane of order n has $n + 1$ points on each line; the corresponding affine plane has n. We call this number the *order* of the affine plane. (So this turns out to be more natural for affine planes than for projective planes. We'll see soon why we chose the number of points on a line

of the affine plane, rather than on a line of the projective plane, to be called the order.)

We have seen (Exercise 14.2.1) that the projective plane has $n^2 + n + 1$ points. To get the affine plane, we delete the $n + 1$ points on a line, so the affine plane has n^2 points.

Coordinates. We have discussed two finite planes (affine or projective; we know it does not matter much): the Fano and the Tictactoe planes. Are there any others?

Coordinate geometry gives the solution: Just as we can describe the Euclidean plane using real coordinates, we can describe finite affine planes using the strange arithmetic of prime fields from Section 6.8. Let us fix a prime p, and consider the "numbers" (elements of the prime field) $\overline{0}, \overline{1}, \ldots, \overline{p-1}$.

In the Euclidean plane, every point has two coordinates, so let's do the same here: Let the points of the plane be all pairs $(\overline{u}, \overline{v})$. This gives us p^2 points.

We have to define the lines. In the Euclidean plane, these are given by linear equations, so let's do the same here: For every equation

$$\overline{a}x + \overline{b}y = \overline{c},$$

we take the set of all pairs $(\overline{u}, \overline{v})$ for which $x = \overline{u}$, $y = \overline{v}$ satisfies the equation, and introduce a line containing all these points. To be precise, we have to assume that the above equation is proper, in the sense that at least one of $\overline{a}$ and $\overline{b}$ is different from 0.

Now we have to verify that (a) through any two points there is exactly one line, (b) for any line and any point outside it, there is exactly one line through the point that is parallel to the line, and (c) there are at least 2 points on each line. We'll not go through this proof, which is not difficult, but lengthy. It is more important to realize that *all this works because it works in the Euclidean plane, and we can do arithmetic in prime fields just as with real numbers.*

This construction provides an affine plane for every prime order (from this we can construct a projective plane for every prime order). Let's see what we get from the smallest prime field, the 2-element field. This will have $2^2 = 4$ points, given by the four pairs $(0,0), (0,1), (1,0), (1,1)$. The lines will be given by linear equations, of which there are six: $x = 0$, $x = 1$, $y = 0$, $y = 1$, $x + y = 0$, $x + y = 1$. Each of these lines goes through 2 points; for example, $y = 1$ goes through $(0,1)$ and $(1,1)$, and $x+y = 0$ goes through $(0,0)$ and $(1,1)$. So we get the very trivial affine plane (already familiar from Exercise 14.1.6) consisting of 4 points and 6 lines. If we extend this to a projective plane, we get the Fano plane.

Figure 14.8 shows the affine plane of order 5 obtained this way (we don't show all the lines; there are too many).

14.2.2 Show that the Cube space can be obtained by 3-dimensional coordinate geometry from the 2-element field.
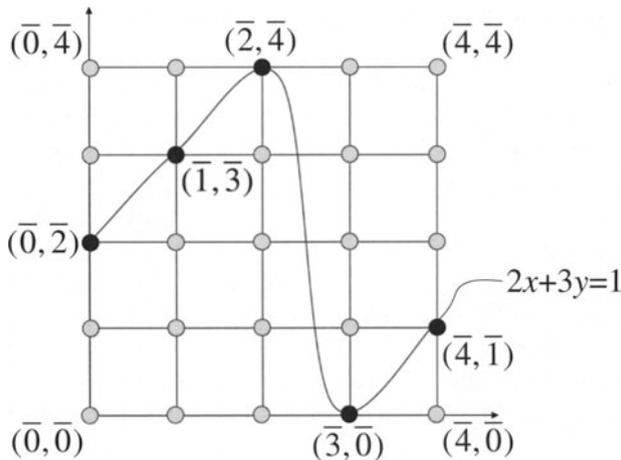
FIGURE 14.8. An affine plane of order 5. Only one line is shown besides the trivial "vertical" and "horizontal" lines.

What are the possible orders of planes? The construction with coordinates above shows that for every prime number there is a finite affine (or projective) plane of that order. Using similar but more involved algebra, one can construct projective planes for every order that is a higher power of a prime number (so for 4, 8, 9 etc.)

Theorem 14.2.2 *For every order that is a power of a prime (including the primes themselves) there is a finite affine (as well as a finite projective) plane of that order.*

   The smallest positive integer that is not a prime power is 6, and Gaston Tarry proved in 1901 that no finite plane of order 6 exists. The next one is 10; the nonexistence of a finite projective plane of order 10 was proved in 1988 by Lam, Thiel and Swiercz based on an extensive use of computers. Nobody has ever found a projective plane whose order is not a prime power, but the question whether such a plane exists is unsolved.

 14.2.3 Suppose that we want to verify the nonexistence of a finite projective plane of order 10 by computer, by simple "brute force": We check that no matter how we specify the appropriate number of subsets of points as lines, one of the conditions (a), (b), or (c) will not hold. How may possibilities do we have to try? About how long would this take?

14.3   Block Designs

The inhabitants of a town like to form clubs. They are socially very sensitive (almost as sensitive as the Fanoans), and don't tolerate any inequalities.

Therefore, they don't allow larger and smaller clubs (because they are afraid that larger clubs might suppress smaller ones). Furthermore, they don't allow some people to be members of more clubs than others, since those who are members of more clubs would have larger influence than the others. Finally, there is one further condition: Each citizen A must behave "equally" toward citizens B and C, A can not be in a tighter relationship with B than C. So A must meet B in the same number of clubs as he/she meets C.

We can formulate these strongly democratic conditions mathematically as follows. The town has v inhabitants; they organize b clubs; every club has the same number of members, say k; everybody belongs to exactly r clubs, and for any pair of citizens, there are exactly λ clubs where both of them are members.

The structure of clubs discussed in the previous paragraphs is called a *block design*. Such a structure consists of a set of v elements, together with a family of k-element subsets of this set (called *blocks*) in such a way that every element occurs in exactly r blocks, and every pair of elements occurs in λ blocks jointly. We denote the number of blocks by b. It is clear that block designs describe what we were discussing when talking about the clubs in the town. In the sequel, sometimes we use this everyday description and sometimes the block design formulation.

Let us see some concrete examples for block designs. One example is given by the Fano plane (Figure 14.1). The points represent the inhabitants of the town, and 3 people form a club if they are on a line.

Let us check that this configuration is indeed a block design: Every club consists of 3 elements (so $k = 3$). Every person belongs to exactly 3 clubs (which means that $r = 3$). Finally, any pair of people belongs to exactly one club by (a) (which means $\lambda = 1$). Thus our configuration is indeed a block design (the number of elements is $v = 7$, the number of blocks is $b = 7$).

The Tictactoe plane gives another block design. Here we have nine points, so $v = 9$; the blocks are the lines, having 3 points on each (so $k = 3$), and there are 12 of them (so $b = 12$). Each point is contained in 4 lines ($r = 4$), and through each pair of points there is a unique line (so $\lambda = 1$).

A block design in which $\lambda \neq 1$ can be obtained from the Cube space if we take the planes as blocks. Clearly, we have $v = 8$ elements; each block contains $k = 4$ elements, the number of blocks is $b = 14$ and every block is contained in $r = 7$ planes. The crucial property is that every pair of points is contained in exactly 3 planes, so we have $\lambda = 3$.

There are some uninteresting, trivial block designs. For $k = 2$, there is only one block design on a given number v of elements: It consists of all $\binom{v}{2}$ pairs. The same construction gives a block design for every $k \geq 2$: We can take all k-subsets as blocks to get a block design with $b = \binom{v}{k}$, $r = \binom{v-1}{k-1}$, and $\lambda = \binom{v-2}{k-2}$. The most boring block design consists of a single block

($k = v$, $b = r = \lambda = 1$). This is so uninteresting that we exclude it from further consideration and don't call it a block design.

Parameters of block designs. Are there any relations among the numbers b, v, r, k, λ? One equation can be derived from the following. Suppose that every club gives a membership card to every one of its members; how many membership cards do they need? There are b clubs and every club has k members, so altogether there are bk membership cards. On the other hand, the town has v inhabitants and everybody has r membership cards, so counting this way we get vr membership cards. In counting the number of membership cards two ways we have to get the same number, so we get

$$bk = vr. \tag{14.1}$$

Let us find another relation. Imagine that the clubs want to strengthen the friendship among their members, so they require that every member should have a dinner one-on-one with each of his/her fellow club members in the clubhouse. How many dinners will a citizen C eat? We can count this in the following two ways:

First reasoning: There are $v - 1$ other citizens in the town, and each of them is in λ clubs jointly with C, so with every one of the other $v - 1$ citizens, C will have to eat λ dinners in the different clubhouses. This means altogether $\lambda(v - 1)$ dinners.

Second reasoning: C is a member in r clubs. Every club has $k - 1$ further members, so in every club C is a member of, C has to eat $k - 1$ dinners. Altogether this means $r(k - 1)$ dinners.

The result of the two counts must be equal:

$$\lambda(v - 1) = r(k - 1). \tag{14.2}$$

14.3.1 If a town has 924 clubs, each club has 21 members, and any 2 persons belong to exactly 2 clubs jointly, then how many inhabitants does the town have? How many clubs does each person belong to? (Don't be surprised: This is a very small town, and everybody belongs to many clubs!)

14.3.2 Show that the assumption that every person is in the same number of clubs is superfluous: It follows from the other assumptions we made about the clubs.

It follows that among the numbers b, v, r, k, λ we can specify at most three freely, and the other two are already determined by the relations (14.1) and (14.2). In fact, we cannot arbitrarily specify even three. Is it possible, for instance, that in a town of 500 people all the clubs have 11 members, and everybody belongs to 7 clubs? The answer is no; please don't read on, but try to prove it yourself (it is not too hard).

Here is our proof: If this were possible, then for the number of clubs b we would get from (14.1) that

$$b = \frac{v \cdot r}{k} = \frac{500 \cdot 7}{11}.$$

But this is not an integer, so these numbers cannot occur.

OK, there is a rather trivial answer to this problem: We must specify three of the given numbers so that in computing the other two using (14.1) and (14.2) we get integer values. But this is not the whole story. There is an important inequality that holds true in every block design, called Fisher's Inequality:

$$b \geq v. \tag{14.3}$$

The proof of this inequality uses mathematical tools that go beyond the scope of this book.

Unfortunately, one can find five numbers b, v, r, k, λ satisfying conditions (14.1), (14.2), and (14.3) for which no block design with these parameters exists. But we are running out of simple, easily checkable necessary conditions. For instance, there is no block design with parameters $b = v = 43$, $k = r = 7$, $\lambda = 1$ (since this block design would be a finite projective plane of order 6, which we know does not exist). These numbers satisfy (14.1), (14.2), and (14.3), and there is no simple way to rule out this block design (just a tedious study of many cases).

14.3.3 (a) Find an example of specific values for v, r, and k where computing b from (14.1) gives an integer value, but (14.2) leads to a contradiction. (b) Find an example of 5 integers b, v, k, r, λ ($b, k, v, r \geq 2$, $\lambda \geq 1$) that satisfy both (14.1) and (14.2), but $b < v$.

14.3.4 For every $v > 1$, construct a block design with $b = v$.

Club badges. In our town, every club has a badge. The town organizes a parade in which everybody participates and everybody is required to wear the badge of one of the clubs where he/she is a member. Can the badges be chosen so that no two persons wear the same badge?

We need, of course, that there are enough different badges, at least as many as citizens. That is, $b \geq v$. This is indeed guaranteed by Fisher's Inequality (14.3). But is this enough? We have to make sure that every citizen is wearing a badge of one of his or her clubs; not just different badges.

The question has some resemblance to the Marriage Theorem (Theorem 10.3.1) described in Chapter 10. To make use of this resemblance, we assign a bipartite graph to our block design (Figure 14.9). The lower set of points represents the people; the upper set of points represents the clubs. We connect point a to point X if citizen a is a member in club X (in Figure
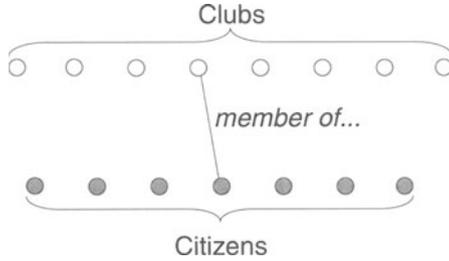
Clubs

member of...

Citizens

FIGURE 14.9. Representing club membership by a graph

14.9 we have drawn only one edge of the graph). We know the following properties of our graph: from every point below, exactly r edges go up, and from every upper point exactly k edges go down. Below we have v points, above we have b points. If we choose n points from the lower set (obviously, $n \leq v$), then we know that from these n points nr edges leave. Let us denote by m the number of other endpoints of these nr edges.

We claim that $n \leq m$. Since every upper point has degree k, altogether this makes mk edges. The nr edges mentioned above are among these mk edges, and hence

$$nr \leq mk. \tag{14.4}$$

On the other hand, $bk = vr$ by (14.1). By $b \geq v$ we get $k \leq r$, and so

$$mk \leq mr. \tag{14.5}$$

From (14.4) and (14.5) we get

$$nr \leq mr,$$

and therefore $n \leq m$, as claimed.

So any n lower points are connected to at least n upper points. We invoke the Marriage Theorem, more precisely, its version stated in Exercise 10.3.2: We get that there exists a matching of the lower set into the upper set, i.e., there exist v independent edges that connect every lower point to a different point above. This matching tells every citizen which badge to wear.

14.4   Steiner Systems

We have seen that block designs with $k \leq 2$ are trivial; we take a closer look at the next case, $k = 3$. We also take the smallest possible value for λ, namely $\lambda = 1$. Block designs with $k = 3$ and $\lambda = 1$ are called *Steiner systems* (named after Jakob Steiner, a Swiss mathematician in the nineteenth century). The Fano plane and the Tictactoe plane are Steiner systems, but the Cube space is not.

How many inhabitants must a town have to allow a system of clubs that is a Steiner system? In other words, what conditions do we get for v if our block design is a Steiner system? We use equations (14.1) and (14.2), substituting the values $k = 3$ and $\lambda = 1$. We get

$$3b = vr \qquad \text{and} \qquad 2r = v - 1,$$

and hence

$$r = \frac{v - 1}{2} \tag{14.6}$$

and

$$b = \frac{v(v - 1)}{6}. \tag{14.7}$$

The numbers r and v must be integers, which imposes some conditions on v. Since the denominator in (14.7) is 6 and that in (14.6) is a divisor of 6, the condition imposed concerns the remainder of v upon division by 6. From (14.6) it follows that v must be an odd number, so if we divide it by 6, the remainders can be 1, 3, or 5. This means that v can be written in the forms $6j + 1$, $6j + 3$, or $6j + 5$, where j is an integer. Furthermore, v can not be of the form $6j + 5$, because then by (14.7) we get

$$b = \frac{(6j + 5)(6j + 4)}{6} = 6j^2 + 9j + 3 + \frac{1}{3},$$

which is not an integer.

So v must be of the form $6j + 1$ or $6j + 3$. Taking into consideration that we must have $v > k = 3$, we see that one can have a Steiner system only in towns where the number of inhabitants is $v = 7, 9, 13, 15, 19, 21, \dots$ etc.

For these numbers one can construct Steiner systems indeed. In the case $v = 7$ we already have seen the Fano plane, and for $v = 9$, the Tictactoe plane. The general construction is quite involved, and we don't describe it here.

14.4.1 Show that for for $v = 7$, the Fano plane is the only Steiner system. (Of course, 7 citizens can set up their clubs in many different ways, by "switching identities." We can think of 7 chairs, with triples forming the clubs specified. The citizens can choose chairs in many different ways.)

14.4.2 Does the Fisher inequality give any further condition on the number of elements in a Steiner system?

Representing the clubs. Imagine that in a town of v people, where the clubs form a Steiner system, people become unhappy about the membership fees, and they create a committee whose task is to protest these high fees. The committee needs at least one member from every club. How many members does this protest committee need to have?

This problem sounds similar to the badge problem discussed at the end of section 14.3, but there are two differences: First, in the badge problem every citizen (his or her interest) was "represented" by one of the clubs the citizen belonged to, while here the clubs will be represented by one of their members; and second (and more significantly), one and the same person can represent several clubs.

Consider citizen Andrew, who does not belong to the committee. Andrew is a member of r clubs, and since the clubs form a Steiner system, we have

$$r = \frac{v-1}{2}$$

(see equation (14.6)).

Every club in which Andrew is a member has two other members, and since Andrew is a member jointly in exactly one club with every other citizen, these $(v-1)/2$ clubs have no members in common other than Andrew. The protest committee must have a member from each of them, and since this member is not Andrew (since he is not a member of the committee), the committee has at least $\frac{v-1}{2}$ members. This means that one needs quite a large committee—almost half of the citizens!

And even this is only a lower bound! Can it be realized, or is there some other argument that gives perhaps an even larger lower bound on the size of the committee? In the case of the Fano plane, our lower bound is $(7-1)/2 = 3$, and indeed, the three points of any line represent every line (any two lines intersect). In the case of the Tictactoe plane, our lower bound is $(9-1)/2 = 4$, but there is no obvious choice of a committee of 4 representing every line; in fact, there is no such choice at all!

This can be seen by a tedious case distinction, but let us offer a nice argument that takes care of many other cases as well. We claim the following surprising fact:

Theorem 14.4.1 *If there exists a committee of size $\frac{v-1}{2}$ representing every club, then this committee itself is also a Steiner system.*

To be more precise, the elements of this new Steiner system will be the v_1 people in the committee, and its blocks will be those clubs for which all three members belong to the committee (such clubs will be called *privileged*).

Proof. To prove that this is indeed a Steiner system, first we show that every two committee members are contained in a privileged club. Suppose not, say Bob and Carl are two committee members who are not jointly contained in a privileged club. This means that the third member of the club to which both Bob and Carl belong (in the original Steiner system) is not on the committee. We may assume that this third member is Andrew. But then in the argument above we see that each club containing Andrew has at least one representative in the committee, and one club has two (Bob and

Carl). This implies that the committee has at least $(v-3)/2+2 = (v+1)/2$ members, a contradiction.

So any two committee members belong to a privileged club. Since no two citizens belong to more than one club in common, no two committee members belong to more than one privileged club in common. So every pair of committee members belongs to exactly one privileged club in common. This proves that the committee is indeed a Steiner system.   □

Now, if the 9-element Steiner system could be represented by 4 elements, then we would get a Steiner system on 4 elements, which we already know cannot exist! We get similarly that for Steiner systems on $13, 21, 25, 33, \ldots$ points, more than half of the citizens are needed to represent every club.

14.4.3 Suppose that a Steiner system on v elements contains a subset S of $(v - 1)/2$ elements such that those triples of the original system that belong totally to S form a Steiner system. Prove that in this case S forms a representative committee (so every triple of the original Steiner system contains an element of S).

Gender balance. The inhabitants of our town want to set up their clubs so that in addition to forming a Steiner system, they should be "gender-balanced." Ideally, they would like to have the same number of males and females in each club. But realizing that this cannot happen (3 being an odd number), they would settle for less: They require that every club must contain at least one male and at least one female.

In mathematical terms, we have a Steiner system, and we want to color the elements with 2 colors (red and blue, corresponding to "female" and "male") in such a way that no block (club) gets only one color. Let us call such a coloring a *good* 2-coloring of the Steiner system.

Is this possible? Let's start with the first nontrivial special case, the case $v = 7$. We have seen (Exercise 14.4.1) that the only Steiner system in this case is the Fano plane. After a little experimentation we can convince ourselves that there is no way to 2-color this system. In fact, we have stated this already in exercise 14.1.2: If a good 2-coloring were possible, then in any case where the males vote one way and the females the other way, the "line rule" would not provide a clear-cut decision.

But it is not only the Fano plane for which no good 2-coloring is possible:

Theorem 14.4.2 *No Steiner system has a good 2-coloring.*

Proof.   After our preparations, this is not hard to prove. Suppose that we have found a good 2-coloring (thus every triple contains differently colored elements). Then the set of red points and the set of blue points each represent all clubs, so (by our discussion above) both sets must contain at least $\frac{v-1}{2}$ elements. Altogether this gives $v - 1$ points, so there is only one further point, which is blue or red; say it is red. Then the $(v - 1)/2$ blue

points form a representative committee, and as we have seen, it is itself a Steiner system. But then any club that is a block in this smaller Steiner system contains only blue points, contradicting the fact that we supposed we had a good coloring.   □

14.4.4 Show that if we allow 3 colors, then both the Fano plane and the Tictactoe plane can be colored so that every block gets at least two colors (but not necessarily all three).

The Schoolgirls' Walking Schedule. A teacher has a group consisting of 9 schoolgirls. She takes them for a walk every day; they walk in three lines, with three girls in each line. The teacher wants to arrange the walks so that after several days, every girl should have walked with every other girl in one line exactly once.

14.4.5 How many days do they need for that?

If you've solved the previous exercise, then you know how many days they need, but is it possible to arrange the walk as the teacher wishes? Trying to make a plan from scratch is not easy.

An observation that helps is the following. Call a triple of girls a *block* if they walk in one line at any time. This way we get a Steiner system. We already know a Steiner system on 9 elements, the Tictactoe plane.

Are we done? No, because the problem asks for more than simply a Steiner system: We have to specify which blocks (triples) form lines on each day. The order of the days is clearly irrelevant, so what we need is a splitting of the 12 triples into 4 classes such that each class consists of three disjoint triples (thus giving a walk plan for a day). If we take a careful look at the Tictactoe plane, then we notice that this is exactly how it is constructed: a set of 3 parallel lines gives a walk plan for one day.

Thomas Kirkman, an English amateur mathematician, asked this question about 15 girls (then the girls need 7 days to complete a walking plan). The question for 15 girls remained unsolved for several years, but finally mathematicians found a solution. Obviously, once you have the right plan, it is easy to check the correctness of it, but there are many possible plans to try.

To find a perfect walking plan for the general case when $v = 6j + 3$, instead of 9 or 15, turned out to be a much harder question. It was solved only more than 100 years later, in 1969, by the Indian-American mathematician Ray-Chaudhuri. One should notice that from this result it follows that there exists a Steiner system for every $v = 6j + 3$; even to prove this simpler fact is quite hard.

There is a related question: Suppose that the teacher wants a plan in which every three girls walk together in a line exactly once. It is not hard to see that such a plan would last for $\binom{15}{3}/5 = 91$ days. The triples that

walk together form a block design again, but now this is what we called above "trivial" (all triples out of 15 points). So this problem appears easier than the Kirkman Schoolgirl Problem, but its solution (in general, with v schoolgirls walking in lines of k) took even more time: it was solved in 1974 by the Hungarian mathematician Zsolt Baranyai.

14.5   Latin Squares

Look at the little 4×4 tables below. Each of them has the property that on every field we have one of the numbers 0, 1, 2, and 3, in such a way that no number occurs in any row or in any column, more than once. A table with this property is called a (4×4) *Latin square*.

$$
\begin{array}{|c|c|c|c|}
\hline
0 & 1 & 2 & 3 \\ \hline
1 & 2 & 3 & 0 \\ \hline
2 & 3 & 0 & 1 \\ \hline
3 & 0 & 1 & 2 \\ \hline
\end{array}
\qquad
\begin{array}{|c|c|c|c|}
\hline
0 & 2 & 1 & 3 \\ \hline
2 & 1 & 3 & 0 \\ \hline
1 & 3 & 0 & 2 \\ \hline
3 & 0 & 2 & 1 \\ \hline
\end{array}
\qquad (14.8)
$$

$$
\begin{array}{|c|c|c|c|}
\hline
0 & 1 & 2 & 3 \\ \hline
1 & 0 & 3 & 2 \\ \hline
2 & 3 & 0 & 1 \\ \hline
3 & 2 & 1 & 0 \\ \hline
\end{array}
\qquad
\begin{array}{|c|c|c|c|}
\hline
0 & 1 & 2 & 3 \\ \hline
3 & 2 & 1 & 0 \\ \hline
1 & 0 & 3 & 2 \\ \hline
2 & 3 & 0 & 1 \\ \hline
\end{array}
\qquad (14.9)
$$

It is easy to construct many Latin squares with any number of rows and columns (see Exercise 14.5.2). Once we have a Latin square, it is easy to make many more from it. We can reorder the rows, reorder the columns, or permute the numbers $0, 1, \ldots$ occurring in them. For example, if we replace 1 by 2 and 2 by 1 in the first Latin square in (14.8), we get the second Latin square.

14.5.1 How many 4×4 Latin squares are there? What is the answer if we don't consider two Latin squares different if one of them can be obtained from the other by permuting rows, columns, and numbers?

14.5.2 Construct an $n \times n$ Latin square for every $n > 1$.

Let us have a closer look at the Latin squares in (14.9). If we place these two squares on top of each other, in every field we get an ordered pair of numbers: The first element of the pair comes from the appropriate field of the first square, and the second element from the appropriate field of the

second:

0,0	1,1	2,2	3,3
1,3	0,2	3,1	2,0
2,1	3,0	0,3	1,2
3,2	2,3	1,0	0,1

(14.10)

Do you notice something about this composite square? Every field contains a different pair of numbers! From this it follows that each of the possible $4^2 = 16$ pairs occurs exactly once (Pigeonhole Principle). If two Latin squares have this property, we call them *orthogonal*. One may check the orthogonality of two Latin squares in the following way: We take all the fields in the first Latin square that contain 0, and we check the same fields on the second square, to see whether they contain different integers. We do the same with 1, 2, etc. If the squares pass all these checks, then the first square is orthogonal to the second one, and vice versa.

14.5.3 Find two orthogonal 3×3 Latin squares.

Magic squares. If we have two orthogonal Latin squares, we may very easily construct from them a *magic square*. (In a magic square the sums of the numbers in every row and every column are equal.) Consider the pairs in the fields in (14.10). Replace each pair (a, b) by $\overline{ab} = 4a + b$ (in other words, consider $\overline{ab}$ as a two-digit number in base 4). Writing our numbers in decimal notation, we get the magic square shown in (14.11).

0	5	10	15
7	2	13	8
9	12	3	6
14	11	4	1

(14.11)

(This is a magic square indeed: Every row and column sum is 30.) From any two orthogonal Latin squares we can get a magic square using the same method. In every row (and also in every column) the numbers 0, 1, 2, 3 occur exactly once in the first position and exactly once in the second position, so in every row (and column) the sum of the elements is exactly

$$(0 + 1 + 2 + 3) \cdot 4 + (0 + 1 + 2 + 3) = 30,$$

as required in a magic square.

14.5.4 In our magic square we have the numbers 0 through 15, instead of 1 through 16. Try to make a magic square from (14.10) formed by the numbers 1 through 16.

14.5.5 The magic square constructed from our two orthogonal Latin squares is not "perfect", because in a perfect magic square the sums on the diagonals are

the same as the row and column sums. From which orthogonal Latin squares can we make perfect magic squares?

Is there a 4×4 Latin square that is orthogonal to both of our Latin squares making up (14.10)? The answer is yes; try to construct it yourself before looking at (14.12). It is interesting to notice that these three Latin squares consist of the same rows, but in different orders.

$$
\begin{array}{|c|c|c|c|}
\hline
0 & 1 & 2 & 3 \\
\hline
2 & 3 & 0 & 1 \\
\hline
1 & 0 & 3 & 2 \\
\hline
3 & 2 & 1 & 0 \\
\hline
\end{array}
\qquad (14.12)
$$

14.5.6 Prove that there does not exist a fourth 4×4 Latin square orthogonal to all three Latin squares in (14.9) and (14.12).

14.5.7 Consider the Latin square (14.13). It is almost the same as the previous one in (14.12); but (prove!) there does not exist any Latin square orthogonal to it. So Latin squares that look similar can be very different!

$$
\begin{array}{|c|c|c|c|}
\hline
0 & 1 & 2 & 3 \\
\hline
1 & 3 & 0 & 2 \\
\hline
2 & 0 & 3 & 1 \\
\hline
3 & 2 & 1 & 0 \\
\hline
\end{array}
\qquad (14.13)
$$

Latin squares and finite planes. There is a very close connection between Latin squares and finite affine planes. Consider an affine plane of order n; pick any class of parallel lines, and call them "vertical"; pick another class and call them "horizontal". Enumerate the vertical lines arbitrarily, and also the horizontal lines arbitrarily. Thus we can think of the points of the plane as entries of an $n \times n$ table in which every row as well as every column is a line (this is the way we presented the Tictactoe plane at the beginning of this Chapter).

Now consider an arbitrary third parallel class of lines, and again, label the lines arbitrarily $0, 1, \ldots, n-1$. Each entry of the table (point in the plane) belongs to exactly one line of this third parallel class, and we can write the label of this line in the field. So all the 0 entries will form a line of the plane, all the 1 entries a different, but parallel, line, etc.

Since any two nonparallel lines have exactly one point in common, the line of 0's will meet every row (and similarly every column) exactly once, and the same holds for the lines of 1's, 2's, etc. This implies that the table we constructed is a Latin square.

This is not too exiting so far, since Latin squares are easy to construct. But if we take a fourth parallel class, and construct a Latin square from it,

then *these two Latin squares will be orthogonal!* (This is just a translation of the fact that every line from the third parallel class intersects every line from the fourth exactly once.) The affine plane has $n + 1$ parallel classes; two of these were used to set up the table, but the remaining $n - 1$ provide $n - 1$ mutually orthogonal Latin squares.

From the Tictactoe plane we get two orthogonal 3×3 Latin squares this way (not surprisingly, they are just the ones found directly in exercise 14.5.3). From the affine plane of order 5 constructed earlier, we get 4 mutually orthogonal Latin squares, as shown below.

0	1	2	3	4
1	2	3	4	0
2	3	4	0	1
3	4	0	1	2
4	0	1	2	3

0	1	2	3	4
2	3	4	0	1
4	0	1	2	3
1	2	3	4	0
3	4	0	1	2

0	1	2	3	4
3	4	0	1	2
1	2	3	4	0
4	0	1	2	3
0	1	2	3	4

0	1	2	3	4
4	0	1	2	3
3	4	0	1	2
2	3	4	0	1
1	2	3	4	0

This nice connection between Latin squares and affine planes works both ways: If we have $n - 1$ mutually orthogonal Latin squares, we can use them to construct an affine plane in a straightforward way. The points of the plane are the fields in an $n \times n$ table. The lines are the rows and columns, and for every number in $\{0, 1, \ldots, n - 1\}$ and every Latin square, we form a line from those fields that contain this number.

Recall that we have constructed finite planes only of prime order, even though we remarked that they exist for all prime power orders. We can now settle at least the first of the missing orders: Just use this reverse construction to get an affine plane of order 4 from our three mutually orthogonal 4×4 Latin squares in (14.9) and (14.12).

14.6   Codes

We are ready to talk about some *real* applications of the ideas discussed in this Chapter. Suppose that we want to send a message through a noisy channel. The message is (as usual) a long string of bits (0's and 1's), and "noisy" means that some of these bits may be corrupted (changed from 0 to 1 or vice versa). The channel itself could be radio transmission, telephone, internet, or just your compact disc player (in which case the "noise" may be a piece of dirt or a scratch on the disc).

How can we cope with these errors and recover the original message? Of course, a lot depends on the circumstances. Can we ask for a few bits to be resent, if we notice that there is an error? In internet protocols we can; in transmissions from a Mars probe or in listening to music on compact discs, we can't. So in some cases it is enough if we can *detect* whether there is an error in the message we receive, while in other cases we have to be able to *correct* the error just from the received message itself.

The simplest solution is to send the message twice, and check whether any of the bits arrives differently in the two messages (we can repeat each bit immediately, or repeat the whole message; it does not make any difference at this point). This is called a *repetition code*. Certainly, if a bit does not match, then we know something is wrong, but of course we don't know whether the first or second copy of the bit was wrong. So we *detect* an error, but cannot *correct* it. (Of course, it may happen that *both* copies of the bit are corrupted; we have to make the assumption that the channel is not too noisy, so that the probability that this happens is small. We'll come back to this issue.) An easy way to strengthen this is to send the message three times. Then we can also correct the error (for every bit, take as correct the version that arrives at least twice), or in the very noisy situation, at least detect it even if 2 (but not 3) copies of a bit are corrupted.

There is another simple way of detecting errors: a parity check. This is the simple trick of appending a bit to each string of a given length (say, after 7 bits) so that the number of 1's in the extended message is even (so we append a 0 if the number of 1's is already even, and append a 1 if it is odd). The recipient can look at the received block of 8 bits (a byte), and check whether it has an even number of 1's. If so, we consider it OK; otherwise, we know that it contains an error. (Again, in a very noisy channel errors may remain undetected: If two bits of the 8 are changed, then the parity check does not reveal it.)

Here is an example of how a string (namely, 10110010000111) is encoded in these two ways:

    11001111000011000000000111111   (repetition code)
    1011001000001111   (parity check)

These solutions are not cheap; their main cost is that the messages become longer. In the repetition code, the increase is 100%; in the parity check, it is about 14%. If the errors are so rare that we can safely assume that only one bit in (say) every 127 is corrupted, then it suffices to append a parity check bit after every 127 bits, at a cost of less than 1%. (Note that the repetition code can be thought of as inserting a parity check bit after every single bit!)

Is this the best way? To answer this question, we have to make an assumption about the noisiness of the channel. So we assume that we are sending a message of length k, and that there are no more than e errors (corrupted bits). We cannot use all strings of length k to send messages

(since then any error would result in another possible message, and the error could not be detected). The set of strings that we use is called a *code*. So a code is a set of 0-1 strings of length k. For $k = 8$ (one byte), the repetition code consists of the following 16 strings:

00000000, 00000011, 00001100, 00001111, 00110000, 00110011,

00111100, 00111111, 11000000, 11000011, 11001100, 11001111,

11110000, 11110011, 11111100, 11111111;

the parity check code consists of all strings of length 8 in which the number of 1's is even (there are $2^7 = 128$ of these, and so we don't list them all here).

We have seen that the parity check code is 1-*error-detecting*, and so is the repetition code. What is the strongest code on 8 bits (detecting the largest number of errors)? The answer is easy: the code consisting of the two codewords

00000000,   11111111

is 7-error-detecting: All 8 bits must be corrupted before we can be fooled. But this code comes with a very high price tag: What it means is that we resend every bit 8 times.

We can construct a more interesting code from the Cube space. This has 8 points, corresponding to the 8 bits. Let us fix an ordering of the points, say $ABCDEFGH$ in Figure 14.7. Every plane P in the geometry will provide a codeword: We send a 1 if the corresponding point lies in the plane P. For example, the bottom-face plane gives the codeword 11110000; the black plane gives the codeword 10100101. We also add the words 00000000 and 11111111, to get a total of 16 codewords.

How good is this code? How many bits must be corrupted before one codeword is changed into another? Assume that the two codewords come from two planes P and Q, which by property (B) of the Cube space are either parallel or intersect each other in a line (i.e., in two points).

First suppose that these planes are parallel. For example, if they are the "black" and "light" planes, then the two codewords they provide are

10100101

01011010.

We wrote the codewords above each other, so that it should be easier to make the following observation: The two planes have no point in common, which (according to the way we constructed the code) means that in no position can both of them have a "1". Since each of them has four 1's, it follows that in no position can both of them have a "0" either. So all 8 bits must be changed before one of them becomes the other.

Second, suppose that the two planes intersect in two points. For example, the "black" plane and the "bottom" plane give the codewords

$$10100101,$$
$$11110000.$$

The two codewords will have two common 1's, and (since each has four 1's) two common 0's. So 4 bits must be changed before one of them becomes the other.

The two further codewords that we added as a kind of an afterthought, all-0 and all-1, are easy to check: We must change 4 bits in them to get a codeword coming from a plane, and 8 bits in them to get one from the other.

What is important from these is that *if we change up to 3 bits in any codeword, we get a string that is not a codeword.* In other words, this code is 3-error-detecting.
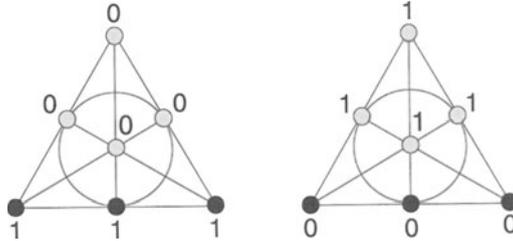

FIGURE 14.10. Two codewords from one line.

The Fano plane provides another interesting code. Again, let each point correspond to a position in the codewords (so the codewords will consist of 7 bits). Each line will provide two codewords, one in which we put 1's for the points on the line and 0's for the points outside, and one in which it is the other way around. Again, we add the all-0 and all-1 strings to get 16 codewords.

Instead of ordering the bits of the codeword, we can think of them as writing 0 or 1 next to each point of the Fano plane. Figure 14.10 illustrates the two codewords associated with a line.

Since we have 16 codewords again, but use only 7 bits, we expect less from these codes than from the codes coming from the Cube space. Indeed, these Fano codes can no longer detect 3 errors. If we start with the codeword defined by a line L (1 on the line and 0 elsewhere), and change the three 1's to 0's, then we get the all-0 string. But it is not only these two special codewords that cause the problem: Again, if we start with the same codeword, and flip the 3 bits on any other line K (from 1 to 0 at the intersection point of K and L, from 0 to 1 at the other two points of K), then we get a codeword coming from a third line (Figure 14.11).
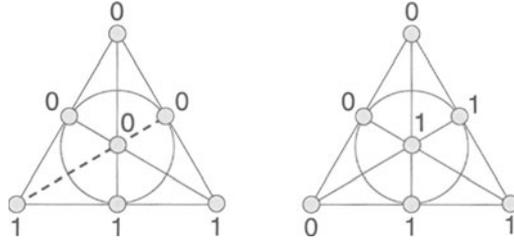
FIGURE 14.11. Three errors are too much for the Fano code: flipping the bits along the dotted line produces another valid codeword.

Going through an argument as above, we can see that

> *the Fano code is 2-error-detecting: We cannot get a valid codeword from another valid codeword by flipping 1 or 2 bits.*

This implies that

the Fano code is 1-error-correcting.

What this means is that not only can we detect if there is an erroneous bit, but we can correct it. Suppose that we receive a string that comes from a valid codeword a by changing a single bit. Could this string come from another valid codeword b? The answer is no: Otherwise, the codewords a and b would differ in only two places, which is not possible.

The codes we derived from the Fano plane and the Cube space are special cases of a larger family of codes called *Reed–Müller codes*. These are very important in practice. For example, the NASA Mariner probes used them to send back images from the Mars. Just as the Cube Code was based on 2-dimensional subspaces of a 3-dimensional space, the code used by the Mariner probes were based on 3-dimensional subspaces of a 5-dimensional space. They worked with blocks of size 32 (instead of 8 as we did), and could correct up to 7 errors in each block. The price was, of course, pretty stiff: There were only 64 codewords used, so to safely transmit 6 bits, one had to package them in 32 bits. But of course, the channel (the space between Earth and Mars) was very noisy!

Error-correcting codes are used all around us. Your CD player uses a more sophisticated error-correcting code (called the Reed–Solomon code) to produce a perfect sound even if the disk is scratched or dusty. Your Internet connection and digital phone use such codes to correct for noise on the line.

14.6.1 Prove that a code is d-error-correcting if and only if it is $(2d)$-error-detecting.

14.6.2 Show that every string of length 7 is either a codeword of the Fano code or arises from a unique codeword by flipping one bit (a code with this property is called *perfect 1-error-correcting*).

Review Exercises

14.6.3 Verify that the Tictactoe plane is the same as the affine plane over the 3-element field.

14.6.4 A lab has 7 employees. Everybody works 3 nights a week: Alice on Monday, Tuesday, and Thursday; Bob on Tuesday, Wednesday, and Friday; etc. Show that any two employees meet exactly one night a week, and for any two nights there is an employee who is working on both. What is the connection with the Fano plane?

14.6.5 The game SMALLSET (which is a simplified version of the commercial card game SET) is played with a deck of 27 cards. Each card has 1, 2, or 3 identical shapes; each shape can be a circle, triangle, or square, and it can be red, blue, or green. There is exactly one card with 2 green triangles, exactly one with 3 red circles, etc. A SET is a triple of cards such that the number of shapes on them is either all the same or all different; the shapes are either all the same or all different; and their colors are all the same or all different. The game consists of putting down 9 cards, face up, and recognizing and removing SETs as quickly as you can; if no SETs are left, 3 new cards are turned up. If no SETs are left and all the remaining cards are turned up, the game is over.

  (a) What is the number of SETs?

  (b) Show that for any two cards there is exactly one third card that forms with them a SET.

  (c) What is the connection between this game and the affine space over the 3-element field?

  (d) Prove that at the end of the game, either no cards or at least 6 cards remain.

14.6.6 How many points do the two smallest projective planes have?

14.6.7 Consider the prime field with 13 elements. For every two numbers x and y in the field, consider the triple $\{x + y, 2x + y, 3x + y\}$ of elements of the field. Show that these triples form a block design, and compute its parameters.

14.6.8 Determine whether there exists a block design with the following parameters:

  (a) $v = 15$, $k = 4$, $\lambda = 1$;

  (b) $v = 8$, $k = 4$, $\lambda = 3$;

  (c) $v = 16$, $k = 6$, $\lambda = 1$.

14.6.9 Prove that the Tictactoe plane is the only Steiner system with $v = 9$.

14.6.10 Consider the addition table of the "Days of the Week" number system in Section 6.8. Show that this table is a Latin square. Can you generalize this observation?

14.6.11 Describe the code you get from the projective plane over the 3-element field, analogously to the Fano code. How much error correction/detection does it provide?