May, 1995

SRC Research Report

134a

From Quadrangular Sets to the Budget Matroids

Lyle Ramshaw
and
James B. Saxe

digital

Systems Research Center
130 Lytton Avenue
Palo Alto, California 94301

Systems Research Center

The charter of SRC is to advance both the state of knowledge and the state of the art in computer systems. From our establishment in 1984, we have performed basic and applied research to support Digital's business objectives. Our current work includes exploring distributed personal computing on multiple platforms, networking, programming technology, system modelling and management techniques, and selected applications.

Our strategy is to test the technical and practical value of our ideas by building hardware and software prototypes and using them as daily tools. Interesting systems are too complex to be evaluated solely in the abstract; extended use allows us to investigate their properties in depth. This experience is useful in the short term in refining our designs, and invaluable in the long term in advancing our knowledge. Most of the major advances in information systems have come through this strategy, including personal computing, distributed systems, and the Internet.

We also perform complementary work of a more mathematical flavor. Some of it is in established fields of theoretical computer science, such as the analysis of algorithms, computational geometry, and logics of programming. Other work explores new ground motivated by problems that arise in our systems research.

We have a strong commitment to communicating our results; exposing and testing our ideas in the research and development communities leads to improved understanding. Our research report series supplements publication in professional journals and conferences. We seek users for our prototype systems among those with whom we have common interests, and we encourage collaboration with university researchers.

Robert W. Taylor, Director

From Quadrangular Sets to the Budget Matroids

Lyle Ramshaw and James B. Saxe

May, 1995

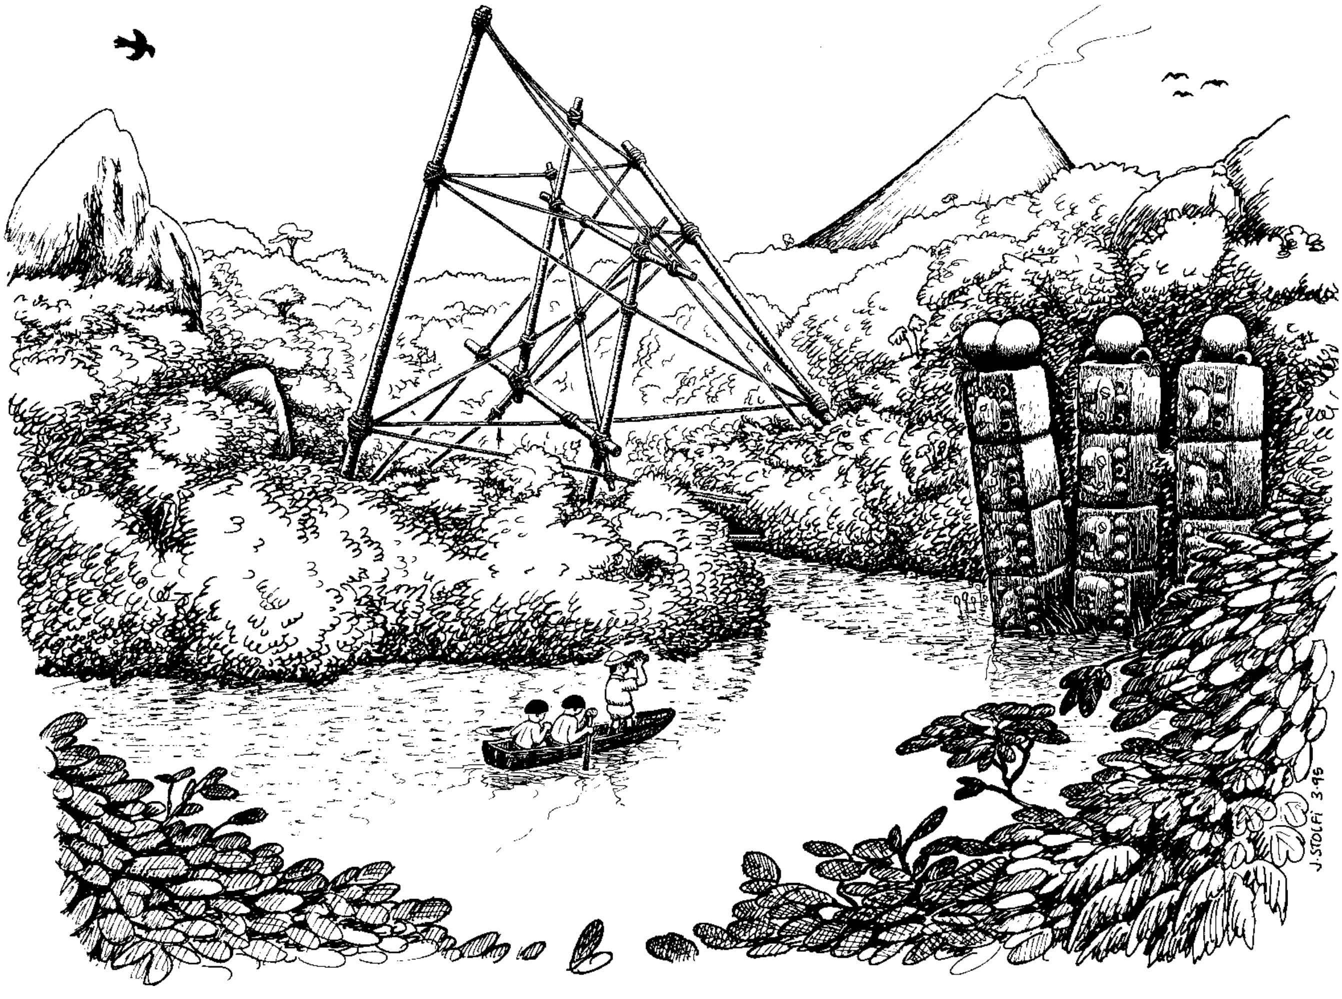

To Bob, who made it all possible,
to Jorge, who made it all better, and
to Kim, who made it all worthwhile.

Contents

List of Figures

Preface

0.1 The central question

For work in computer-aided geometric design, as for countless other purposes, we must understand the structure of the univariate polynomials. In particular, we must be able to test for linear dependence. If we know the $n + 1$ coefficients of each of $n + 1$ polynomials of degree n, we can test for dependence simply by computing an $(n + 1)$-by-$(n + 1)$ determinant. But there are times — such as when using polar forms [43] to study splines — when what we know about each polynomial is its n roots, rather than its $n + 1$ coefficients. At such times, we can compute the coefficients from the roots, up to an irrelevant scale factor, and then proceed as before. For example, in the quadratic case, the three nonzero polynomials

$$\begin{aligned}
f_1(X) &= r_1(X - a_1)(X - b_1) = r_1(X^2 - (a_1 + b_1)X + a_1b_1)\\
f_2(X) &= r_2(X - a_2)(X - b_2) = r_2(X^2 - (a_2 + b_2)X + a_2b_2)\\
f_3(X) &= r_3(X - a_3)(X - b_3) = r_3(X^2 - (a_3 + b_3)X + a_3b_3)
\end{aligned}$$

are linearly dependent just when the 3-by-3 determinant

$$\begin{vmatrix} 1 & a_1 + b_1 & a_1b_1 \\ 1 & a_2 + b_2 & a_2b_2 \\ 1 & a_3 + b_3 & a_3b_3 \end{vmatrix}$$

is zero, the numbers (r_i) being the irrelevant scale factors. In the cubic case, the four nonzero polynomials

$$\begin{aligned}
g_1(X) &= r_1(X - a_1)(X - b_1)(X - c_1)\\
g_2(X) &= r_2(X - a_2)(X - b_2)(X - c_2)\\
g_3(X) &= r_3(X - a_3)(X - b_3)(X - c_3)\\
g_4(X) &= r_4(X - a_4)(X - b_4)(X - c_4)
\end{aligned}$$

are linearly dependent just when

$$\begin{vmatrix} 1 & a_1 + b_1 + c_1 & a_1b_1 + a_1c_1 + b_1c_1 & a_1b_1c_1 \\ 1 & a_2 + b_2 + c_2 & a_2b_2 + a_2c_2 + b_2c_2 & a_2b_2c_2 \\ 1 & a_3 + b_3 + c_3 & a_3b_3 + a_3c_3 + b_3c_3 & a_3b_3c_3 \\ 1 & a_4 + b_4 + c_4 & a_4b_4 + a_4c_4 + b_4c_4 & a_4b_4c_4 \end{vmatrix} = 0.$$

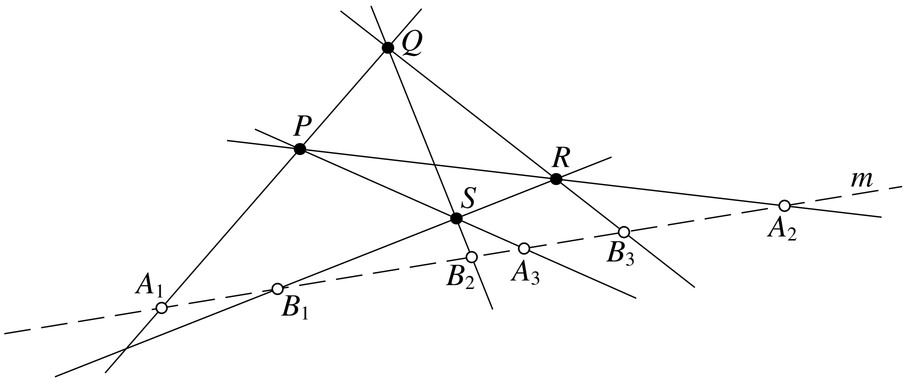

Figure 0.1: A complete quadrangle with vertices (P, Q, R, S) and the quadrangular set $\{\{A_i, B_i\}\}_{i\in[1..3]}$ formed by stabbing that quadrangle with the line m.

In the quadratic case, projective geometry provides an elegant alternative test. Given four points (P, Q, R, S) in the plane with no three collinear, as shown in Figure 0.1, the six lines that join those four points in pairs form an instance of the configuration called *the complete quadrangle*. The line PQ is called *opposite* to the line RS, and similarly for the pairs $\{PR, QS\}$ and $\{PS, QR\}$. Three pairs of points along a line m are said to form a *quadrangular set* when they are the intersections of m with the three pairs of opposite lines of some instance of the complete quadrangle. For example, in Figure 0.1, the three pairs $\{A_1, B_1\}$, $\{A_2, B_2\}$, and $\{A_3, B_3\}$ form a quadrangular set. Once we fix a coordinate system for the line m, any pair of points $\{A, B\}$ on m gives rise to a pair of coordinates $\{a, b\}$, and those coordinates, in turn, are the roots of a quadratic polynomial $f(X) = r(X - a)(X - b)$ that is uniquely determined, up to the irrelevant scale factor r. It is a classical result that, given three pairs of points on m, the three polynomials so determined are linearly dependent just when the three pairs of points form a quadrangular set.

Here is the central question that sparked this research: Is there a projective configuration that provides an analogous geometric test for the linear dependence of four cubic polynomials? Since that configuration would be a cubic analog of the complete quadrilateral, we shall call it *the complete cubangle*. (Don't worry: We use that horrid name only in this preface.) Does the complete cubangle exist?

Let's flesh out the analogy between the quadratic and cubic cases, so that we can see what properties the complete cubangle should have. An instance of the complete quadrangle consists of three pairs of lines in the plane, certain triples of which are constrained to be concurrent (that is, to pass through a common point). By analogy, an instance of the complete cubangle should consist of four triples of lines in the plane — or, more likely, four triples of planes in 3-space — constrained by certain required incidences. Let us say that four triples of points along a line m in 3-space form a *cubangular set* when they are the intersections of m with the four triples of planes of some instance of the complete cubangle. Once we fix a coordinate system for the line m, any triple of points $\{A, B, C\}$ on m determines a triple of coordinates $\{a, b, c\}$ and hence determines a cubic polynomial $g(X) = r(X - a)(X - b)(X - c)$ uniquely, up to the scale factor r. The analogy requires that, given four triples of points on m, the four cubic polynomials so determined are linearly dependent just when the four triples of points form a cubangular set.

Ta-da! The complete cubangle does exist — that is, there is a configuration that fulfills this analogy. An instance of the complete cubangle consists of twelve planes, two lines, and thirteen points in 3-space, constrained by various incidences.

And that's not all! The complete cubangle is just one member of a large family of configurations, some familiar and some novel, which we can define using the notion of a *budget partition*. The complete cubangle is associated with the budget partition $(2, 1, 1)$, while other budget partitions give rise to other configurations. Defining these configurations and studying their properties involves an intriguing combination of old-fashioned geometry and modern combinatorics. From geometry, we use projective frames, projective transformations, and null systems; from combinatorics, we use matroids, minors, and representations.

0.2 A duality warning

One beautiful aspect of projective geometry is the principle of duality, which lets us interchange the concepts 'point' and 'hyperplane', provided that we also interchange the concepts 'lies on' and 'passes through'. The complete cubangle consists of four triples of planes in 3-space, constrained by their incidences with two lines and thirteen points. The dual configuration consists of four triples of points in 3-space, constrained by their incidences with two lines and thirteen planes. As we discuss in Section 2.6, it turns out to be more convenient to focus on the dual configuration. That is why Chapter 1 begins by discussing, not the complete quadrangle, but its dual, which is called *the complete quadrilateral*.

0.3 The accompanying videotape

This monograph is being published as Research Report 134*a* from the Systems Research Center of the Digital Equipment Corporation. Report 134*b* is a 46-minute videotape, titled *Introducing the Budget Configurations*. The videotape animates the configurations in 3-space that are associated with the budget partitions $(2, 1, 1)$ and $(1, 1, 1, 1)$, the former being the dual of the complete cubangle. One goal of the videotape is to give non-mathematicians some idea of what this work is all about. The other goal is to convince you that you want to read this monograph. If you are already convinced, there is no need to watch the tape — though you might enjoy doing so.

0.4 Acknowledgments

First, I want to thank Jorge Stolfi for lots of help and lots of ideas. Jorge's searching questions sparked Sections 5.4, 10.3, and 10.4. He gets the bulk of the credit for Section 4.5 and every ounce of the credit for Section 11.4. On top of all that, he drew the cartoon.

I thank and praise my colleagues Allan Heydon and Greg Nelson for writing Juno-2 [18], the constraint-based drawing editor that I used to produce both the figures in this monograph and the 2D portions of the videotape. In a similar vein, I thank and praise Marc A. Najork and Marc H. Brown for writing the Anim3D library [35] and the Obliq-3D system [36], which I used for the 3D portions of the videotape.

I am grateful also to Jürgen Richter-Gebert for his insightful comments on a draft of this monograph.

Finally, my thanks go to Jim Saxe for contributing his formidable expertise and intuition to our joint investigation of this pretty mathematics. While we explored the mathematics jointly, the words and pictures in the monograph and videotape are mine, so the blunders that no doubt remain in them are all my fault.

Lyle Ramshaw
`ramshaw@pa.dec.com`
May 17, 1995

Chapter 1

Preview

The preface describes the central question out of which this monograph grew. This preview summarizes the results of the monograph, to give you some sense, in advance, of where we are heading. It's not going to be until Chapter 2 that we finally get around to introducing the material in depth.

1.1 Prerequisites

In order to understand this monograph, you have to know something about projective geometry and something about matroids.

Projective geometry is a standard topic. Good places to start are the books by Coxeter [8] and by Samuel [48]. A valuable supplement, particularly as regards computations using homogeneous coordinates and Plücker coordinates, is Stolfi's book [50] — although Stolfi discusses oriented projective geometry, which we don't need, because our matroids are not oriented [5].

As for matroids, the books by Oxley [37] and by Welsh [52] are fine sources. Because matroid theory isn't yet as standard a topic as projective geometry, we take the time, in Chapter 3, to review the definition of a matroid and some of the elementary properties of matroids.

1.2 Configurations and dependent polynomials

The *complete quadrilateral*, an instance of which is shown in Figure 2.2 on page 7, is a projective configuration with six points and four lines. It is a standard result that the complete quadrilateral characterizes the dependence of three quadratic polynomials. As for what it means to 'characterize dependence', two theorems are involved, which we shall call the Projection Theorem and the Witness Theorem. (Warning: Those names are not standard, as discussed in Section 2.5.)

Given any instance of the complete quadrilateral and given any point in the plane O, we can project the six vertices of the quadrilateral from O to get six lines through O. The *Projection Theorem* says that the three univariate, quadratic polynomials that have the slopes of those six lines as their roots are always linearly dependent.

Conversely, suppose that we take any six numbers that are the roots of three linearly dependent quadratics, we take any point O in the plane, and we draw the six lines through O with those slopes. The *Witness Theorem* says that there exists an instance of the complete quadrilateral whose six vertices lie on those six lines, and which hence *witnesses* to the linear dependence of those three quadratics. Furthermore, there is a natural sense in which the witnessing quadrilateral is unique: Any two witnessing quadrilaterals are related by a projective transformation of the plane that fixes the point O and fixes every line through O.

Here is the first half of the good news in a nutshell:

> There is a projective configuration $B_{2,1,1}$ in 3-space — consisting of twelve points, two lines, and thirteen planes — that characterizes the dependence of four cubic polynomials in the same way that the complete quadrilateral in the plane characterizes the dependence of three quadratic polynomials.

In particular, there are cubic analogs of the Projection Theorem and the Witness Theorem that hold for the configuration $B_{2,1,1}$. Note that we project an instance of the configuration $B_{2,1,1}$ from a line o in 3-space, rather than from a point O in the plane. The twelve vertices of the configuration determine twelve planes through the line o, the slopes of which are the roots of four cubic polynomials that are always linearly dependent.

1.3 The budget matroids

A *matroid* is an algebraic structure that describes the pattern of incidences — the collinearities, coplanarities, and the like — that are required of the points in a projective configuration (it being understood that every incidence that is not required is forbidden). So, roughly speaking, 'matroid' is the modern term for 'projective configuration'. If the matroid M describes the incidences of some configuration C, then a *representation* of M is the same thing as an instance of C; and a matroid is *representable* when it has representations.

The incidences of the complete quadrilateral are described by a matroid that we shall call $B_{2,1}$. The justification for that name comes from the other half of the good news:

> For each partition $b = b_1 + \cdots + b_k$ of an integer b into at least two positive parts, we can define an associated matroid, which we shall call the *budget matroid* $B_{b_1,\ldots,b_k}$. A surprising number of these budget matroids are representable. In particular, the complete quadrilateral

is described by the matroid $B_{2,1}$, while the configuration in 3-space proclaimed above is described by the matroid $B_{2,1,1}$.

In more detail, let $b = b_1 + \cdots + b_k$ be some partition of an integer b into at least two positive parts. The budget matroid $B_{b_1,\ldots,b_k}$ has rank b and has, as its ground set, a matrix

$$\begin{pmatrix} E_{11} & E_{12} & \ldots & E_{1k} \\ E_{21} & E_{22} & \ldots & E_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ E_{b1} & E_{b2} & \ldots & E_{bk} \end{pmatrix}$$

consisting of bk points. The rules for independence in the matroid $B_{b_1,\ldots,b_k}$ force the b points in the j^{th} column to lie in a common flat subspace of dimension b_j, for each j in $[1\,..\,k]$. Also, for each possible way of choosing b points so that precisely one is chosen from each row and so that, for each j, precisely b_j are chosen from the j^{th} column, the b points so chosen must lie in a common hyperplane. Here are geometric descriptions of the representations of the budget matroids of rank $b = 3$:

- $B_{2,1}$: the six vertices of a complete quadrilateral in the plane.
- $B_{1,1,1}$: the nine points of a Pappus configuration in the plane, as shown in Figure 4.1 on page 39.

Here are similar descriptions for rank $b = 4$:

- $B_{3,1}$: the four vertices of a tetrahedron in 3-space, together with the four points where a line cuts its four faces.
- $B_{2,2}$: the eight vertices of a Möbius pair of tetrahedra — that is, two tetrahedra with the vertices of each lying on the faces of the other.
- $B_{2,1,1}$: twelve points in 3-space, four on a plane, four each on each of two lines, and with 12 other coplanarities in a certain pattern. This is the configuration that characterizes the dependence of four cubic polynomials.
- $B_{1,1,1,1}$: sixteen points in 3-space, four each on each of four lines and with 24 coplanarities in a certain pattern.

We prove two general results about the representability of the budget matroids: Every budget matroid $B_{m,n}$ with two parts is representable over the rationals, as is every matroid $B_{m,1,1}$ with three parts, two of which are ones. The latter result can be interpreted as generalizing Pappus's Theorem from the plane to $(m+1)$-space. In the particular case $m = 2$, the latter result gives us one construction for representations of the budget matroid $B_{2,1,1}$.

To show that the matroid $B_{2,1,1}$ characterizes the dependence of cubic polynomials, we develop a different construction for its representations — one that makes the choices in a different order. This second construction exploits the properties of *null systems*, the little-known cousins of the well-known *polar systems* that are associated with quadric hypersurfaces. We use the same null-system machinery also to construct representations of $B_{1,1,1,1}$, the most complicated of the budget matroids of rank 4.

1.4 The bad news about the quartic case

It would be satisfying to find, for each n, a projective configuration in n-space that geometrically characterized the dependence of $n+1$ polynomials, each of degree n. Since the budget matroid $B_{2,1}$ in the plane — better known as the complete quadrilateral — characterizes the dependence of three quadratics and the matroid $B_{2,1,1}$ in 3-space characterizes the dependence of four cubics, it is natural to ask whether the matroid $B_{2,1,1,1}$ in 4-space perhaps characterizes the dependence of five quartics. Unfortunately, the answer seems to be no. The matroid $B_{2,1,1,1}$ is representable, but it has so few representations that they characterize a property of five quartics that is stronger than linear dependence.

Chapter 2

Introduction

2.1 The quadratic case

The quadratic polynomial $X^2 - (u + v)X + uv$ has the pair of scalars $\{u, v\}$ as its roots.[1] Let us say that three pairs of scalars $\big\{\{u_1, v_1\}, \{u_2, v_2\}, \{u_3, v_3\}\big\}$ form a 2-*dependent block* when the three quadratic polynomials with those pairs as their roots are linearly dependent, that is, when the determinant

$$\begin{vmatrix} 1 & u_1 + v_1 & u_1 v_1 \\ 1 & u_2 + v_2 & u_2 v_2 \\ 1 & u_3 + v_3 & u_3 v_3 \end{vmatrix}$$

is zero. Note that the order of the two scalars $\{u_i, v_i\}$ in a pair doesn't affect the 2-dependency of the block, nor does the order of the three pairs.

By the way, you can think of the scalars that we deal with, such as u_i and v_i, as either rational numbers, real numbers, or complex numbers, as suits your fancy. More formally, we fix some field of scalars in which to carry out our numeric computations and over which to build our projective spaces. Unless otherwise stated, we require only that the scalar field have characteristic zero — that property being all that we need for the bulk of our arguments. Of course, when we want to illustrate a geometric construction with a picture, it is simplest to assume that the field of scalars is the real numbers.

The algebraic condition of 2-dependence corresponds to a simple geometric condition involving a conic curve in the plane. Let O be a point in the plane, and let $\big\{\{a_1, b_1\}, \{a_2, b_2\}, \{a_3, b_3\}\big\}$ be three pairs of lines through O, as shown (over the real numbers) in Figure 2.1. Each of the lines has a scalar slope, and when those slopes form a 2-dependent block of scalars, we shall refer to the block of lines itself as 2-*dependent*. To test a block of lines for 2-dependence, let k be any

[1] The expression $\{u, v\}$ denotes a pair of roots, even when $u = v$; that is, the curly braces here indicate a *suite* [45] or *multiset* or *bag*, rather than a set. But please ignore all degenerate cases, as much as you can, until Section 2.4.

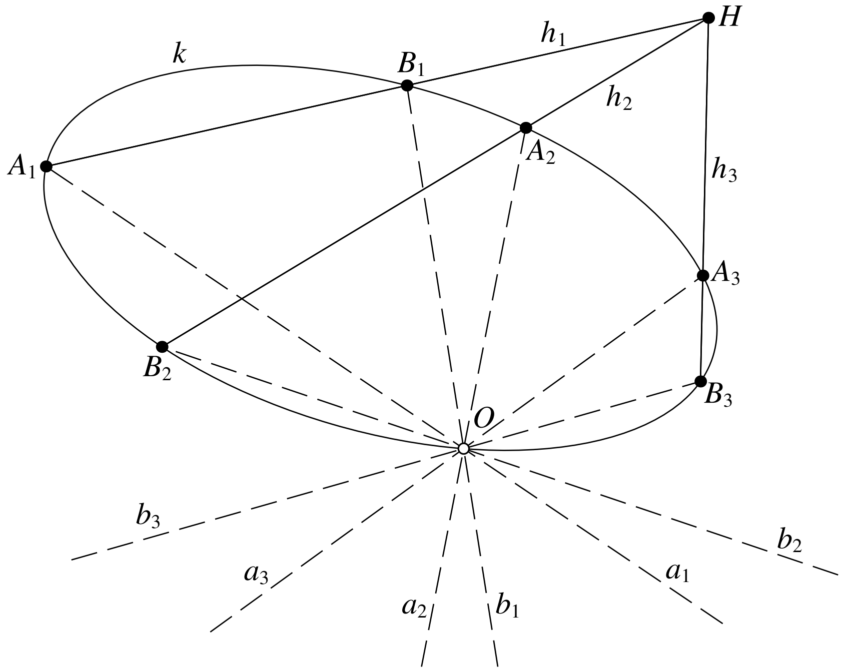

Figure 2.1: A 2-block of lines, all passing through a common point O, whose 2-dependence is demonstrated by the concurrence of three chords of a conic curve k that also passes through the point O.

conic passing through the point O. Any line through O intersects the conic k at O itself and at one other point. Let $h_i := A_i B_i$ be the chord of the conic k joining the residual intersections of the lines a_i and b_i, for i in $[1 \,..\, 3]$. It is straightforward to check (see Exercise 2.1-4) that the block of lines $\big\{\{a_1, b_1\}, \{a_2, b_2\}, \{a_3, b_3\}\big\}$ is 2-dependent just when the three chords h_1, h_2, and h_3 of the conic k are concurrent — that is, pass through a common point H.

It is a fascinating result of classical projective geometry that 2-dependence can be characterized in an even simpler way, at the cost of some loss of symmetry. The four auxiliary lines a, q, r, and s in Figure 2.2 intersect at the six points

$$\begin{aligned} A_1 &:= a \cap q \qquad & B_1 &:= r \cap s \\ A_2 &:= a \cap r \qquad & B_2 &:= q \cap s \\ A_3 &:= a \cap s \qquad & B_3 &:= q \cap r. \end{aligned}$$

These six points and four lines are called *a complete quadrilateral*. More precisely, they are an instance of the classical configuration called *the complete quadrilateral*; that is, the name of the pattern is used also for instances of that pattern. For i in $[1 \,..\, 3]$, the two points A_i and B_i are called *opposite*, since they lie on no common line. Because of the loss of symmetry, as we discuss shortly, we shall specify a complete quadrilateral by giving its six points $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ as

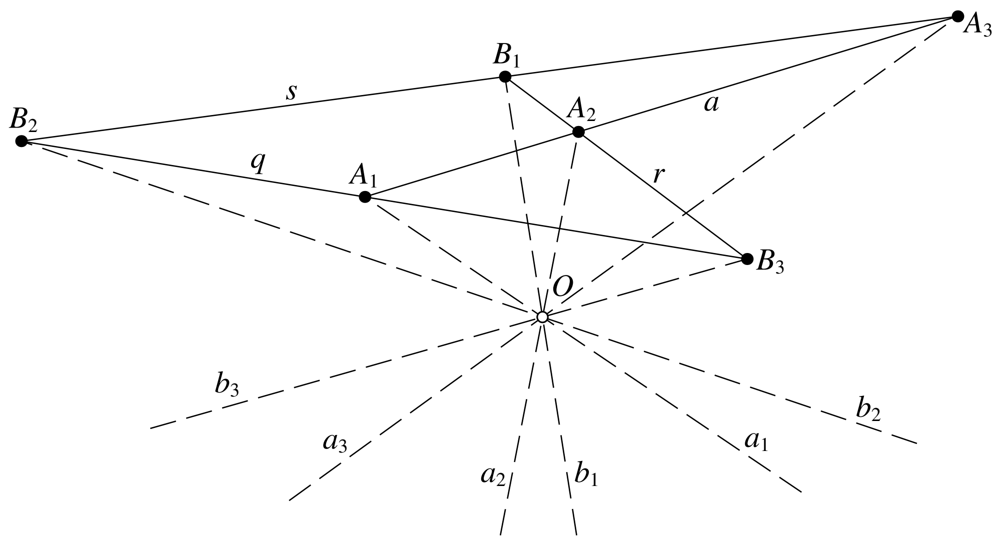

Figure 2.2: The same 2-block of lines as in the previous figure, but with their 2-dependence demonstrated, this time, by a witnessing complete quadrilateral.

an *ordered block*, that is, as an *unordered* triple of *ordered* pairs. So, for each pair of opposite vertices (A_i, B_i), we know which of them is the A-point and which is the B-point. The A-points of the three pairs are collinear, along the auxiliary line a; and the A-point of any pair is collinear with the B-points of the other two pairs, along some other of the auxiliary lines.

Given an ordered block of points that form a complete quadrilateral and given any point O in the plane, we can project the three pairs of points from the point O — which is called, in this context, the *center* of the projection — to get an ordered block of lines $\{(a_1, b_1), (a_2, b_2), (a_3, b_3)\}$ through O. It turns out that this block of lines is always 2-dependent. We call this result the *Projection Theorem*, and we discuss its proof in Section 2.5.

Conversely, given any ordered block of lines through a common point O that is 2-dependent, there exists a complete quadrilateral each of whose six vertices lies on the appropriate line and which hence *witnesses* to the 2-dependence. Of course, this witnessing quadrilateral is not unique. It is clear that any projective transformation of the plane that fixes the point O and fixes every line through O carries any witnessing quadrilateral to another quadrilateral that is also a witness. But the witnessing quadrilateral is unique up to such transformations; that is, any witnessing quadrilateral can be mapped to any other by a projective transformation of the plane that fixes both the point O and every line through O. We call this converse result the *Witness Theorem*, and we discuss its proof also in Section 2.5.

Figure 2.2 is simpler than Figure 2.1 in many respects; for instance, Figure 2.2 has one less point and has no curves. But it is important to realize that some of the inherent symmetries of the problem, which are preserved in Figure 2.1, are broken

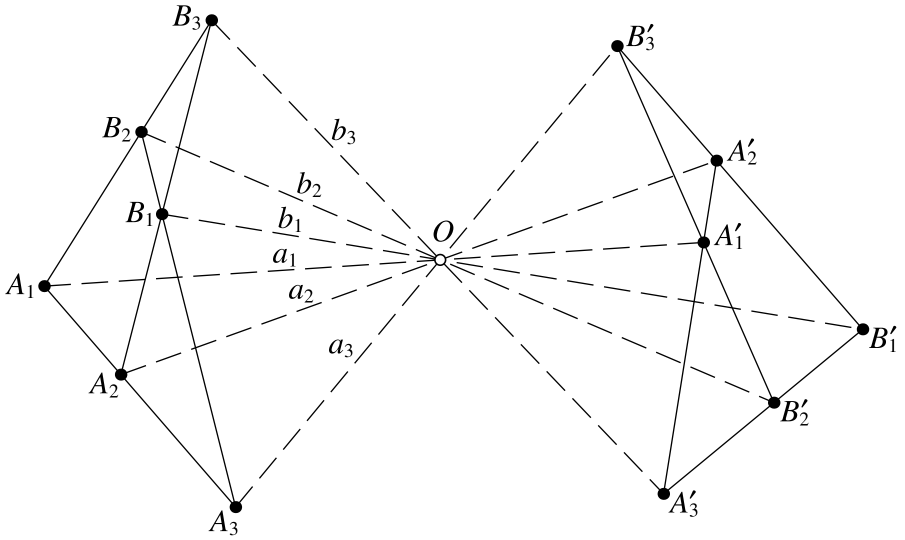

Figure 2.3: An unordered block of lines and two complete quadrilaterals that witness to the 2-dependence of that block, but under different orderings thereof.

in Figure 2.2. (Coxeter [9] relates this broken symmetry to the axioms for projective geometry.) Suppose that we arbitrarily select the line a_1 from the pair $\{a_1, b_1\}$ and the line a_2 from the pair $\{a_2, b_2\}$. Even after we have made those selections, the two lines a_3 and b_3 in Figure 2.1 play completely symmetric roles. But not so in Figure 2.2: The point A_3 on the line a_3 is collinear with A_1 and A_2, while the point B_3 is not. Thus, once we have ordered two of the three pairs in Figure 2.2, the third pair acquires an order as well. We can swap any even number of pairs without changing anything. If we swap an odd number of pairs, we still don't affect the 2-dependency of the block of slopes; but we do change the structure of the witnessing quadrilaterals.

Something must be done about this broken symmetry, lest the uniqueness claim in the Witness Theorem fail. Figure 2.3 shows a 2-dependent, unordered block of lines $\big\{\{a_1, b_1\}, \{a_2, b_2\}, \{a_3, b_3\}\big\}$ through a point O. It also shows two complete quadrilaterals, one on the left and one on the right. The one on the left witnesses to the 2-dependence of the obvious ordered block $\{(a_1, b_1), (a_2, b_2), (a_3, b_3)\}$, while the one on the right witnesses, instead, to the 2-dependence of the ordered block $\{(a_1, b_1), (a_2, b_2), (b_3, a_3)\}$, in which the single pair (a_3, b_3) has been swapped. Note that no projective transformation that fixes O and every line through O can possibly map the left quadrilateral to the right one. Any such transformation would have to map $A_1 \mapsto A'_1$, $A_2 \mapsto A'_2$, and $A_3 \mapsto B'_3$; but the points A_1, A_2, and A_3 are collinear, while the points A'_1, A'_2, and B'_3 are not.

Thus, when talking about witnessing quadrilaterals, we cannot allow the three pairs of a block to be swapped arbitrarily. We shall respond to this situation by the simple expedient of forbidding swapping altogether. That explains why we defined a complete quadrilateral to be an ordered block of points. When a complete quadrilateral witnesses to the 2-dependence of a block of lines or slopes, we treat that block also as ordered.

Exercise 2.1-1 Count the degrees of freedom in the Witness Theorem, so as to verify that its claims of existence and uniqueness are at least plausible.

[Answer: Fix a center point O in the plane, and consider 2-dependent blocks of lines through O and the complete quadrilaterals that witness to their 2-dependence. There are five degrees of freedom in choosing such a 2-dependent block: six scalar slopes, related by one equation. There are eight degrees of freedom in choosing a complete quadrilateral: two in each of its four lines. Thus, for each 2-dependent block of lines through O, there must be a 3-parameter family of quadrilaterals that witness to the 2-dependence of that block. The Witness Theorem says that any such witness can be mapped to any other by a projective transformation of the plane that fixes both the point O and every line through O. There are eight degrees of freedom in an arbitrary projective transformation of the plane — any four points, no three collinear, to any other four such points. Of those eight, it takes four to fix two chosen lines through O, which fixes the point O as well. It takes one more to fix all of the remaining lines through O. The three degrees of freedom that are left are just enough to map any witness to any other.]

Exercise 2.1-2 Given five lines a_1, a_2, b_1, b_2, and b_3 through a point O in the plane, use a complete quadrilateral to construct the unique sixth line a_3 through O that makes the ordered block $\{(a_i, b_i)\}_{i\in[1..3]}$ of lines 2-dependent.

[Answer: Choose B_i arbitrarily on b_i, for i in $[1\,..\,3]$, thus using up the three degrees of freedom that are involved in choosing a witness. Then construct $A_1 := a_1 \cap B_2B_3$, $A_2 := a_2 \cap B_1B_3$, $A_3 := A_1A_2 \cap B_1B_2$, and $a_3 := OA_3$.]

Exercise 2.1-3 In this exercise, the field of scalars must be ordered; for simplicity, let's use the real numbers. A pair $\{u_i, v_i\}$ of real numbers is said to *separate* another pair $\{u_j, v_j\}$ when the product $(u_j - u_i)(u_j - v_i)(v_j - u_i)(v_j - v_i)$ is negative. Let $\big\{\{u_1, v_1\}, \{u_2, v_2\}, \{u_3, v_3\}\big\}$ be six distinct, real numbers that form a 2-dependent block. Show that either each pair separates both of the other pairs or else no two pairs separate each other.

[Hint: In Figure 2.1, does the point H lie inside or outside of the conic k? Alternatively, in Figure 2.2, the auxiliary lines a, q, r, and s divide the real projective plane into seven regions: four triangles (two of which are finite) and three quadrilaterals (one of which is finite). Does the center point O lie in a triangle or in a quadrilateral?]

Exercise 2.1-4 Verify that the three chords h_1, h_2, and h_3 of the conic in Figure 2.1 are concurrent just when the six lines $\big\{\{a_1, b_1\}, \{a_2, b_2\}, \{a_3, b_3\}\big\}$ through O form a 2-dependent block.

[Hint: We can choose our coordinate system on the projective plane so that the point O is the origin and the conic k is the standard parabola $Y = X^2$. The line through O with slope t then cuts the conic k at the point $(X, Y) = (t, t^2)$. Letting the scalars u_i and v_i, for i in $[1 \,..\, 3]$, denote the slopes of the lines a_i and b_i, it follows that the chord h_i is the line with the equation $Y - (u_i + v_i)X + u_i v_i = 0$. Deduce from this that the three chords are concurrent just when the six slopes are 2-dependent.]

2.2 Generalizing to $n > 2$

To what extent does this theory generalize to degrees n greater than 2?

Let's define an *unordered n-block* — of scalars, points, lines, planes, or whatever — to be an unordered $(n+1)$-tuple of unordered n-tuples. An *ordered n-block*, on the other hand, is an unordered $(n + 1)$-tuple of ordered n-tuples.

It is easy to define n-dependence algebraically, for any n. We say that an unordered n-block of scalars $\big\{\{u_{ij}\}_{j\in[1..n]}\big\}_{i\in[0..n]}$ is *n-dependent* when the $n+1$ polynomials

$$\begin{array}{c}
(X - u_{01})(X - u_{02}) \cdots (X - u_{0n}) \\
(X - u_{11})(X - u_{12}) \cdots (X - u_{1n}) \\
\vdots \qquad \vdots \qquad \vdots \qquad \vdots \\
(X - u_{n1})(X - u_{n2}) \cdots (X - u_{nn})
\end{array}$$

of degree n that have those n-tuples as their roots are linearly dependent. We can test this by forming an $(n + 1)$-by-$(n + 1)$ matrix whose i^{th} row consists of the elementary symmetric polynomials in the entries u_{i1} through u_{in} of the i^{th} tuple and checking whether the determinant vanishes. For example, the unordered 3-block of scalars $\big\{\{u_i, v_i, w_i\}\big\}_{i\in[0..3]}$ is 3-dependent just when

$$\begin{vmatrix}
1 & u_0 + v_0 + w_0 & u_0v_0 + u_0w_0 + v_0w_0 & u_0v_0w_0 \\
1 & u_1 + v_1 + w_1 & u_1v_1 + u_1w_1 + v_1w_1 & u_1v_1w_1 \\
1 & u_2 + v_2 + w_2 & u_2v_2 + u_2w_2 + v_2w_2 & u_2v_2w_2 \\
1 & u_3 + v_3 + w_3 & u_3v_3 + u_3w_3 + v_3w_3 & u_3v_3w_3
\end{vmatrix} = 0.$$

Before we can test for n-dependence geometrically, we first have to decide how to encode scalars geometrically; we're going to use the slopes of the hyperplanes in a pencil of hyperplanes in n-space. Recall that a *pencil* consists of all of the hyperplanes that pass through a fixed flat subspace of dimension $n-2$, which is called the *center* of the pencil. So the hyperplanes in a pencil form a one-parameter family, with slope as the parameter. For example, when $n = 2$, we test the 2-dependence of a 2-block of lines in the plane, all passing through a common center point. When $n = 3$, we test the 3-dependence of a 3-block of planes in 3-space, all passing through a common center line. (In this case, where the center of the pencil is a line, that line is also called the *axis*.) Let's say that an n-block of hyperplanes, all in a common pencil, is *n-dependent* just when the slopes of those hyperplanes form an n-dependent block of scalars. As we discuss in Section 2.5, it doesn't matter what scale we use to measure slopes.

It is easy to test for n-dependence geometrically using a curve. The role that is played by a conic when $n = 2$ is played, in the general case, by a rational normal curve of degree n in projective n-space [15, 47]. For example, when $n = 3$, the relevant curve is a twisted cubic in 3-space — one example of which is the curve given parametrically by $t \mapsto (t, t^2, t^3)$. Let $\big\{\{\alpha_i, \beta_i, \gamma_i\}\big\}_{i \in [0..3]}$ be an unordered 3-block of planes, all through a common line o. Choose a twisted cubic k that intersects the center line o at two points — that is, so that o is a chord of k. Any plane through o intersects k at three points: the two points where o intersects k and one residual intersection. Given a triple of planes $\{\alpha, \beta, \gamma\}$ through o, let A, B, and C denote the residual intersections of those three planes with the twisted cubic k, and let η be the plane $\eta := ABC$. The block of planes $\big\{\{\alpha_i, \beta_i, \gamma_i\}\big\}_{i \in [0..3]}$ is 3-dependent just when the four planes $\eta_i := A_i B_i C_i$ determined by the residual intersections are concurrent.

But it isn't at all clear how to take the test for 2-dependence that is based on complete quadrilaterals and to generalize that into a test for n-dependence. Is there a projective configuration in n-space that characterizes n-dependence in the same way that the complete quadrilateral in the plane characterizes 2-dependence? For example, when $n = 3$, is there some pattern of collinearities and coplanarities that we can impose on twelve points in 3-space so that, for any twelve points with those incidences and for any additional line o, projecting the twelve points from o will result in twelve planes whose slopes are 3-dependent? Conversely, given any twelve planes whose slopes are 3-dependent, does there always exist an instance of this hypothetical configuration that witnesses to that 3-dependence? The main goal of this monograph is to study such a configuration — one for which cubic analogs of both the Projection and Witness Theorems hold.

Recall that using a complete quadrilateral to test for 2-dependence breaks some of the inherent symmetry, which is why we introduced ordered 2-blocks. Since the symmetry gets broken already when $n = 2$, it seems likely that it will be broken also when $n > 2$. Hence, just as we defined an instance of the complete quadrilateral to be an ordered 2-block of points in the plane, we expect an instance of its cubic analog to be an ordered 3-block of points in 3-space. And we expect the Witness Theorem in the cubic case to talk about such an instance as witnessing to the 3-dependence of an ordered 3-block of planes, all in a common pencil.

Exercise 2.2-1 When does the 1-block $\big\{\{u_0\}, \{u_1\}\big\}$ formed by two singletons of scalars have the property of 1-dependence?

[Answer: When $u_0 = u_1$.]

Exercise 2.2-2 Suppose that we fix all but one of the scalars in the $(n + 1)$-by-n

matrix

$$\begin{bmatrix} u_{01} & u_{02} & \dots & u_{0n} \\ u_{11} & u_{12} & \dots & u_{1n} \\ u_{21} & u_{22} & \dots & u_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ u_{n1} & u_{n2} & \dots & u_{nn} \end{bmatrix}.$$

Show that there is typically a unique value for the final scalar that makes the rows of the resulting matrix form an n-dependent block. What can happen in atypical cases?

[Hint: Consider the cases

$$\begin{bmatrix} -1 & -1 \\ 1 & 1 \\ 0 & u \end{bmatrix} \qquad \text{and} \qquad \begin{bmatrix} 0 & 0 \\ 0 & 1 \\ 0 & u \end{bmatrix}.$$

Keep in mind, in the former case, that not all lines have finite slope.]

Exercise 2.2-3 Justify the geometric test for 3-dependence based on a twisted cubic curve, hence generalizing Exercise 2.1-4 to the cubic case.

[Hint: We can choose our coordinate system on projective 3-space so that the center line o is the Z-axis and the twisted cubic k is given parametrically by the function $t \mapsto (t, t^2, t^3)$. Note that the line o intersects the cubic k at the origin, where $t = 0$, and also at the point at infinity on the Z-axis, where $t = \infty$. The equation of a plane through o has the form $Y = tX$, and we can treat the t in this equation as the slope of the plane. The plane $Y = tX$ with slope t then has the point (t, t^2, t^3) as its residual intersection with the cubic k. Letting u_i, v_i, and w_i, for i in $[0 \,..\, 3]$, denote the slopes of the planes α_i, β_i, and γ_i, show that the plane η_i determined by the three residual intersections has the equation

$$Z - (u_i + v_i + w_i)Y + (u_i v_i + u_i w_i + v_i w_i)X - u_i v_i w_i = 0.$$

Deduce from this that the four planes (η_i) are concurrent just when the four triples of slopes are 3-dependent.]

2.3 Configurations versus constructions

Our goal is to find a configuration in 3-space that characterizes 3-dependence in the same way that the complete quadrilateral in the plane characterizes 2-dependence. To clarify what that means, we should discuss several things that could be our goal, but aren't. For one thing, it is not our goal simply to devise a geometric construction that tests for n-dependence.

When $n = 2$, given three pairs of lines through a common point, we can test them for 2-dependence geometrically as follows: We use five of them to construct a complete quadrilateral, as in Exercise 2.1-2, and we then check to see if the sixth comes out properly. The resulting complete quadrilateral witnesses either to the 2-dependence of the six slopes or to their lack of 2-dependence. The case $n = 3$ is similar. Given four triples of planes through a common line, we can use any eleven of them to construct an instance of the configuration $B_{2,1,1}$ that witnesses either to the 3-dependence of the twelve slopes or to their lack of 3-dependence.

But there are lots of constructions that test for n-dependence; devising such a construction is too easy to be a worthwhile goal. When $n = 2$, the construction based on complete quadrilaterals is probably the simplest, and it is hence featured in textbooks.[2] But we make no claim that, when $n = 3$, the construction based on the configuration $B_{2,1,1}$ is the simplest or the best. It may be in the running for the shortest, but it involves lots of degenerate cases that would require special treatment, as discussed in Section 9.6. Our central interest is the configuration itself, not the associated construction.

Just how easy is it to test for n-dependence with a construction? The simplest constructions in the plane are those that can be carried out using only a straightedge. There is no standard name for the geometric tool in n-space that is analogous to a straightedge, but it is clear what it should do; we shall call it a *flat-side*. (To carry out a geometric construction in the plane, you look through your pile of scrap lumber for a strip with a straight edge; to carry out a construction in 3-space, you look instead for a board with a flat side.) In the plane, a straightedge lets us mark a line ℓ in such a way that the unique point common to ℓ and to a second, marked line is distinguished as doubly marked. In 3-space, a flat-side lets us mark a plane π in such a way that the line common to π and to a second marked plane is distinguished as doubly marked, while the point common to π and to two other marked planes is distinguished as triply marked. A flat-side in n-space lets us mark hyperplanes in an analogous way.

It is well known that addition and multiplication can be carried out geometrically, using a straightedge in the plane. Hence, we can test n-dependence geometrically for any n with just a straightedge, by mimicking the algebraic definition: Construct the determinant of the $(n + 1)$-by-$(n + 1)$ matrix — that is, construct a segment whose length represents that determinant — and test it for zero. We can shorten the construction quite a bit by moving from the plane to n-space. Instead of constructing the determinant itself, we set up a coordinate system in n-space, we construct the $n + 1$ points whose homogeneous coordinates are given by the rows of the matrix, and we then use a flat-side to test whether those $n + 1$ row-points lie on a common hyperplane.

We can devise another construction that tests for n-dependence by exploiting a rational normal curve of degree n. Of course, with a flat-side as our only tool, we can't draw the rational normal curve itself. But there is no need to do so. Using

[2] Actually, it is the dual construction, based on complete quadrangles, that is generally featured in textbooks, such as Coxeter [8]. See Section 2.6.

only a flat-side, we can choose such a curve and we can construct the residual intersections of the $n(n+1)$ given hyperplanes with that chosen curve, as discussed in the following exercises; and that is enough.

Exercise 2.3-1 Implement the test for 2-dependence based on a conic curve — the one shown in Figure 2.1 — using only a straightedge.

[Hint: Choose four points O', P_1, P_2, and P_3 so that no three of the five points (O, O', P_1, P_2, P_3) are collinear. There is a unique projective correspondence between the pencil of lines through O and the pencil through O' in which the line $p_i := OP_i$ corresponds to the line $p'_i := O'P_i$, for i in $[1\,..\,3]$. The intersections of corresponding lines in those two pencils trace out a conic k. Let a'_1 be the line in the O' pencil whose cross ratio with respect to (p'_1, p'_2, p'_3) is the same as the cross ratio of a_1 with respect to (p_1, p_2, p_3). The residual intersection A_1 of the line a_1 with the conic k is the point $A_1 := a_1 \cap a'_1$.]

Exercise 2.3-2 Using only a flat-side, implement the test for 3-dependence based on a twisted cubic curve.

[Hint: Given twelve planes through the line o, choose two lines o' and o'' that are skew to each other and to o, and choose three points P_1, P_2, and P_3 in general position. There are unique projective correspondences between the three pencils of planes with axes o, o', and o'' in which the planes $\pi_i := \mathrm{Span}(o, P_i)$, $\pi'_i := \mathrm{Span}(o', P_i)$, and $\pi''_i := \mathrm{Span}(o'', P_i)$ correspond, for i in $[1\,..\,3]$. The intersections $\pi \cap \pi' \cap \pi''$ of corresponding planes in the three pencils trace out a twisted cubic through P_1, P_2, and P_3 of which the lines o, o', and o'' are chords.]

2.4 Dealing with degeneracies

The claims made in the previous sections are true generically, but some of them may fail in degenerate cases, where a case is *degenerate* when some polynomial relation among the input parameters that usually doesn't hold happens to hold. For example, we treated slopes above as finite scalars, and they usually are; but there do exist lines with infinite slope. There are fundamentally two ways to deal with a degenerate case: Either we handle it or we outlaw it.

We *handle* a degenerate case by extending our definitions in some clever way, after which that case is no longer degenerate. For example, there is a standard way to handle the degenerate case of lines with infinite slope. We represent a slope, not as a single scalar u, but as the ratio $u^{\uparrow} : u^{\downarrow}$ of two scalars[3] $u^{\uparrow} := \Delta y$ and $u^{\downarrow} := \Delta x$, which are called *homogeneous coefficients* — 'homogeneous' because multiplying them both by the same nonzero scalar doesn't change their ratio. Either of the homogeneous coefficients can be zero, but not both, and the slope of a vertical line is represented by the ratio $u^{\uparrow} : u^{\downarrow} = 1 : 0$. In this way, we can extend the algebraic

[3] Which can be euphoniously read 'u high' and 'u low'.

definition of n-dependence to handle infinities — that is, we can extend it from the affine line to the projective line. In the quadratic case, the polynomial with the two ratios $u^{\uparrow} : u^{\downarrow}$ and $v^{\uparrow} : v^{\downarrow}$ as its two roots is $(u^{\downarrow}X - u^{\uparrow})(v^{\downarrow}X - v^{\uparrow})$, and the three polynomials

$$\begin{aligned}
(u_1^{\downarrow}X - u_1^{\uparrow})(v_1^{\downarrow}X - v_1^{\uparrow}) &= u_1^{\downarrow}v_1^{\downarrow}X^2 - (u_1^{\downarrow}v_1^{\uparrow} + u_1^{\uparrow}v_1^{\downarrow})X + u_1^{\uparrow}v_1^{\uparrow} \\
(u_2^{\downarrow}X - u_2^{\uparrow})(v_2^{\downarrow}X - v_2^{\uparrow}) &= u_2^{\downarrow}v_2^{\downarrow}X^2 - (u_2^{\downarrow}v_2^{\uparrow} + u_2^{\uparrow}v_2^{\downarrow})X + u_2^{\uparrow}v_2^{\uparrow} \\
(u_3^{\downarrow}X - u_3^{\uparrow})(v_3^{\downarrow}X - v_3^{\uparrow}) &= u_3^{\downarrow}v_3^{\downarrow}X^2 - (u_3^{\downarrow}v_3^{\uparrow} + u_3^{\uparrow}v_3^{\downarrow})X + u_3^{\uparrow}v_3^{\uparrow}
\end{aligned}$$

are linearly dependent just when

$$\begin{vmatrix}
u_1^{\uparrow}v_1^{\uparrow} & u_1^{\uparrow}v_1^{\downarrow} + u_1^{\downarrow}v_1^{\uparrow} & u_1^{\downarrow}v_1^{\downarrow} \\
u_2^{\uparrow}v_2^{\uparrow} & u_2^{\uparrow}v_2^{\downarrow} + u_2^{\downarrow}v_2^{\uparrow} & u_2^{\downarrow}v_2^{\downarrow} \\
u_3^{\uparrow}v_3^{\uparrow} & u_3^{\uparrow}v_3^{\downarrow} + u_3^{\downarrow}v_3^{\uparrow} & u_3^{\downarrow}v_3^{\downarrow}
\end{vmatrix} = 0. \tag{2.4-1}$$

In forming this determinant, we dropped the minus signs in the middle column, as we have been doing all along; we also swapped the first and third columns, to make the highs come before the lows in each row, as is the case in the ratio $u^{\uparrow} : u^{\downarrow}$.

Projective geometry adroitly handles the many degenerate cases that are associated with parallelism in a similar way. It extends n-dimensional affine space into n-dimensional projective space by moving from n affine coordinates $(X_1, \ldots, X_n)$ to $n+1$ *homogeneous coordinates* $[x_0, x_1, \ldots, x_n]$, where $X_i = x_i : x_0$. Following Stolfi [50], we shall use square brackets to delimit the homogeneous coordinates of a point and angle brackets to delimit the *homogeneous coefficients* of a hyperplane. So the hyperplane $\langle c_0, c_1, \ldots, c_n\rangle$ passes through the point $[x_0, x_1, \ldots, x_n]$ just when $c_0x_0 + c_1x_1 + \cdots + c_nx_n = 0$. Dividing through by the weight coordinate x_0 of the point converts this equation into $c_0 + c_1X_1 + \cdots + c_nX_n = 0$, which is the inhomogeneous equation of that same hyperplane.

If we don't know how to handle a degenerate case, we can *outlaw* it. For example, consider defining 2-dependence using a conic curve, as in Figure 2.1. Some of the degenerate cases that arise are straightforward to handle:

- One of the lines through O, say a_1, might be tangent to the conic k — in which case we set $A_1 := O$.
- The two lines in a pair, say a_1 and b_1, might coincide — in which case the chord h_1 becomes a tangent.
- Two of the pairs of lines might coincide, say $\{a_1, b_1\} = \{a_2, b_2\}$ because $a_1 = b_2$ and $a_2 = b_1$, causing the two chords h_1 and h_2 to coincide — in which case the three chords are bound to concur.

But there is one degenerate case that must simply be outlawed: the case in which the conic k through the point O factors as the union of two lines.

Sometimes, it is simpler to outlaw a degenerate case, even though it could be handled. For example, we want the Projection Theorem to hold; so we must define the notion 'complete quadrilateral' in such a way that projecting the six vertices of any complete quadrilateral, as in Figure 2.2, gives six lines whose slopes are 2-dependent. We clearly can't allow all four of the auxiliary lines a, q, r, and s to coincide and the six points of intersection $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ to vary arbitrarily and independently along that common line. But lesser degeneracies wouldn't hurt: three of the lines being concurrent, two of them coinciding, even three of them coinciding. Despite this, we shall opt for simplicity — and follow the textbooks — by insisting, as part of our definition of a complete quadrilateral, that no three of its four lines are concurrent. In a similar way, when defining our cubic analog of the complete quadrilateral, we shall stipulate that every incidence that is not required is forbidden.

In the quadratic Projection and Witness Theorems, there are also issues of degeneracy about the location of the complete quadrilateral with respect to the center point O of the projection. The approach taken in the textbooks is to outlaw any 2-block of lines in which two lines in different pairs coincide. Exercise 2.4-3 below explains why that prohibition suffices to eliminate all of the troublesome degeneracies. Adopting that approach and fixing a center point O leads to the following tidy results:

Projection Theorem If none of the four lines of a complete quadrilateral passes through the center point O, then projecting the three pairs of vertices of that quadrilateral from O gives an ordered, 2-dependent block of lines in which no line in any pair coincides with any line in any other pair.

Witness Theorem Conversely, suppose that we are given any ordered, 2-dependent block of lines through O in which no line in any pair coincides with any line in any other pair. Then, there exists a complete quadrilateral whose six vertices lie on the appropriate lines and none of whose four lines passes through O. Furthermore, any two such witnessing quadrilaterals are related by a unique projective transformation of the plane that fixes the point O and every line through O.

Note that the two lines in a single pair are allowed to coincide. They do so precisely when the center point O lies on a diagonal of the complete quadrilateral, the *diagonals* of a complete quadrilateral being the three lines that join the opposite pairs of vertices. Indeed, the center point O may lie at the intersection of two of the three diagonals, as shown in Figure 2.4, and this is an important special case. In the pencil of lines through a point O, a line f is called the *harmonic conjugate* [10] of a line g with respect to two distinct lines a and b just when the block of lines $\big\{\{a, a\}, \{b, b\}, \{f, g\}\big\}$ is 2-dependent.

In the cubic Projection and Witness Theorems, degeneracy is a harder problem: There are more degenerate cases, they are harder to recognize, and they cause a

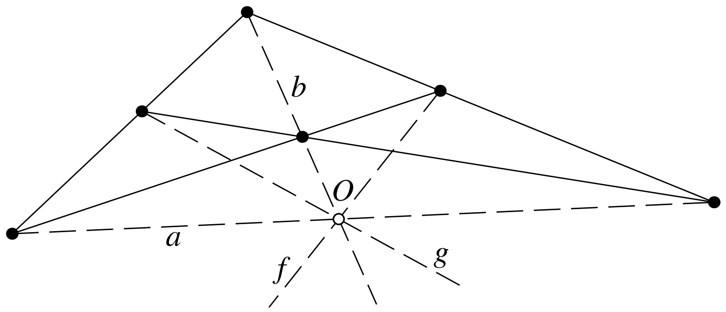

Figure 2.4: Two pairs of lines $\{a, b\}$ and $\{f, g\}$ through a common point O that form a *harmonic set*; that is, each line of each pair is the harmonic conjugate of its mate with respect to the two lines of the other pair.

wider variety of troubles. Note that it is easy to recognize the 2-dependent blocks that might cause trouble, because, as shown in Exercise 2.4-3, they all have at least two entries that are equal. In contrast, a 3-dependent block can cause trouble even when all twelve of its entries are distinct — see Exercise 2.4-4 for an example. We discuss some of the degenerate cases of the cubic Projection and Witness Theorems in Section 9.6, and we speculate on how they might be handled. But, in proving those theorems, we take the coward's way out by restricting ourselves to the *generic case* — those instances in which no troublesome degeneracies arise.

Exercise 2.4-2 Convince yourself that an ordered block $\{(A_i, B_i)\}_{i\in[1..3]}$ of points in some projective space forms a complete quadrilateral in the strict sense — that is, where we forbid every incidence that is not required — just when

- none of the $\binom{6}{2} = 15$ pairs of points coincides;
- every one of the $\binom{6}{4} = 15$ quadruples of points is coplanar; and
- among the $\binom{6}{3} = 20$ triples of points, just the following four are collinear: $\{A_1, A_2, A_3\}$, $\{A_1, B_2, B_3\}$, $\{B_1, A_2, B_3\}$, and $\{B_1, B_2, A_3\}$ — those four corresponding to the four auxiliary lines a, q, r, and s in Figure 2.2.

Exercise 2.4-3 Suppose that

$$\begin{bmatrix} u & v \\ w & x \\ y & z \end{bmatrix}$$

is a 2-dependent block of scalars; so the determinant

$$\Delta := \begin{vmatrix} 1 & u+v & uv \\ 1 & w+x & wx \\ 1 & y+z & yz \end{vmatrix}$$

is zero. If either entry in some pair equals either entry in some other pair, show that a second equality also holds, and show that any entry not involved in either equality can be varied freely, while holding all other entries fixed, without destroying the 2-dependence. Conversely, if there exists any entry that can be varied freely without destroying the 2-dependence, show that some entry in some pair equals some entry in some other pair.

[Hint: For the first part, we can assume without loss of generality that $u = w$, in which case the determinant Δ factors as $\Delta = (x-v)(y-u)(z-u)$. Since $\Delta = 0$, some other equality must hold, and any entry not involved in either equality can be varied freely.

For the second part, solving the equation $\Delta = 0$ for z yields

$$z = \frac{uvw + uvx - uvy + wxy - uwx - vwx}{xy + wy - wx + uv - vy - uy}.$$

Suppose that z can be varied freely, which means that both the numerator N and denominator D of that fraction are zero. It then follows that

$$(w + x - u - v)N - (wx - uv)D = (w - u)(w - v)(x - u)(x - v) = 0,$$

so some entry in the first pair equals some entry in the second pair.]

Exercise 2.4-4 Note that, in the 2-dependent block

$$\begin{bmatrix} 0 & 1 \\ 0 & 1 \\ 0 & 1 \end{bmatrix},$$

we can vary both of the elements in any single row freely and independently, without destroying the 2-dependence. Show that the analogous property holds of the following 3-dependent block, even though all of its entries are distinct:

$$\begin{bmatrix} 1 & 9 & -10 \\ -1 & -9 & 10 \\ 5 & 6 & -11 \\ -5 & -6 & 11 \end{bmatrix}$$

That is, show that the 4-by-4 matrix formed from the coefficients of the cubic polynomials with those triples of roots not only has determinant zero, but actually has rank only 2.

2.5 Proving the quadratic case

Now that we know precisely what the quadratic cases of the Projection and Witness Theorems say, we should discuss how to prove them.

The Projection and Witness Theorems do not appear, as such, in standard texts on projective geometry, for the excellent reason that 2-dependence is an algebraic notion, not a geometric one. But those books do provide pieces which we can assemble into a proof. Coxeter [12] shows that the six lines $\{\{a_i, b_i\}\}_{i\in[1..3]}$ through a point O are the projections of the six vertices of a complete quadrilateral just when there exists an *involution* — that is, a self-inverse projectivity — of the pencil of lines through O that swaps a_i with b_i, for each i. Maxwell [33] shows that such an involution exists precisely when the three chords of the conic in Figure 2.1, that is, the lines $h_i := A_i B_i$ for i in $[1\,..\,3]$, are concurrent. And we already showed, in Exercise 2.1-4, that the three chords (h_i) are concurrent precisely when the six slopes have the algebraic property of 2-dependence. That chain of reasoning suffices to prove the Projection Theorem and the existence half of the Witness Theorem.

That assembled proof has a couple of unfortunate aspects: It doesn't establish the uniqueness half of the Witness Theorem, and it relies on the concept of an involution — a concept that would be somewhat clumsy to generalize to the cubic case. Fortunately, it isn't hard to prove the Projection Theorem and both halves of the Witness Theorem directly. Let's do that, since proving the quadratic case directly will help to prepare us for proving the cubic case, in Chapter 9.

We begin with an overdue lemma about n-dependence. When we are testing an n-block of hyperplanes for n-dependence, it doesn't matter what scale we use to measure the slopes of the hyperplanes in that pencil; that is, we can take any three distinct hyperplanes in that pencil — whichever ones we like — and assign, to them, the slopes 0, 1, and ∞.

Lemma 2.5-1 *The relation of n-dependence is projectively invariant, meaning the following. Let $\{\{u_{ij}\}_{j\in[1..n]}\}_{i\in[0..n]}$ be any unordered n-block, and form a derived n-block $\{\{v_{ij}\}_{j\in[1..n]}\}_{i\in[0..n]}$ by setting*

$$v_{ij} := \frac{au_{ij} + b}{cu_{ij} + d},$$

where a, b, c, and d are fixed scalars with $ad - bc$ nonzero. Then, the derived block is n-dependent just when the original block is.

Proof If we rewrite each entry u_{ij} of the original block as a ratio $u^{\uparrow}_{ij} : u^{\downarrow}_{ij}$ of homogeneous coefficients, we can express the corresponding entry v_{ij} of the derived block as the ratio $v^{\uparrow}_{ij} : v^{\downarrow}_{ij}$, where $v^{\uparrow}_{ij} := au^{\uparrow}_{ij} + bu^{\downarrow}_{ij}$ and $v^{\downarrow}_{ij} := cu^{\uparrow}_{ij} + du^{\downarrow}_{ij}$.

The original block is n-dependent just when the polynomials

$$U_i(X) := (u^{\downarrow}_{i1}X - u^{\uparrow}_{i1}) \cdots (u^{\downarrow}_{in}X - u^{\uparrow}_{in}),$$

for i in $[0\,..\,n]$, are linearly dependent. In a similar way, the derived block is n-dependent just when the polynomials

$$V_i(Y) := (v^{\downarrow}_{i1}Y - v^{\uparrow}_{i1}) \cdots (v^{\downarrow}_{in}Y - v^{\uparrow}_{in})$$

are linearly dependent. But we have

$$V_i(Y) = (a - cY)^n\, U_i\left(\frac{dY - b}{a - cY}\right),$$

and neither substituting the expression $(dY - b)/(a - cY)$ for the variable X nor multiplying through by the fixed polynomial $(a - cY)^n$ has any effect on the linear dependence. □

Theorem 2.5-2 (Projection Theorem, quadratic case) *Let O be any point in the plane, and let the six points $\{(A_i, B_i)\}_{i\in[1..3]}$ be the vertices of any complete quadrilateral, none of whose four lines passes through O. The six lines $\{(a_i, b_i)\}_{i\in[1..3]}$ joining O to the vertices of the quadrilateral then form a block that is 2-dependent and in which no line in any pair coincides with any line in any other pair.*

Proof In order for a line in one pair to coincide with either of the two lines in either of the other two pairs, the center O of the projection would have to lie on one of the four lines of the quadrilateral, which is forbidden. So it suffices to show 2-dependence, which we shall do using analytic geometry, after choosing our coordinate system with some care.

The four points B_1, B_2, B_3, and O are four points in the plane, with no three collinear. So we can choose our coordinate system so that those four points have the homogeneous coordinates

$$\begin{aligned} B_1 &= [1, 0, 0] \\ B_2 &= [0, 1, 0] \\ B_3 &= [0, 0, 1] \\ O &= [1, 1, 1]. \end{aligned}$$

That makes three of the four lines of the quadrilateral have simple coefficients. Recall that, by our convention, the point with homogeneous coordinates $[w, x, y]$ lies on the line with homogeneous coefficients $\langle c, d, e\rangle$ just when $cw + dx + ey = 0$. So the line $q = A_1B_2B_3$ of the quadrilateral has the homogeneous coefficients $q = \langle 1, 0, 0\rangle$, while we similarly have $r = B_1A_2B_3 = \langle 0, 1, 0\rangle$ and $s = B_1B_2A_3 = \langle 0, 0, 1\rangle$. The fourth line $a = A_1A_2A_3$ of the quadrilateral, however, has homogeneous coefficients about which we know very little; let's say that $a = \langle f, g, h\rangle$.

It is easy to calculate the homogeneous coordinates of the points (A_i) in terms of the homogeneous coefficients f, g, and h of the line a; we have

$$\begin{aligned} A_1 &= [0, h, -g] \\ A_2 &= [-h, 0, f] \\ A_3 &= [g, -f, 0]. \end{aligned}$$

We can then calculate the homogeneous coefficients of the six lines that connect the center point O to the six vertices of the quadrilateral:

$$\begin{aligned} a_1 &= \langle -g-h, g, h\rangle \qquad & b_1 &= \langle 0, 1, -1\rangle \\ a_2 &= \langle f, -f-h, h\rangle \qquad & b_2 &= \langle -1, 0, 1\rangle \\ a_3 &= \langle f, g, -f-g\rangle \qquad & b_3 &= \langle 1, -1, 0\rangle. \end{aligned}$$

Our next task is to choose a definition of slope, in the pencil of lines through the center point O; by Lemma 2.5-1, it doesn't matter which definition we choose. Let's adopt, as our measure of slope, the first homogeneous coefficient divided by the last — so the three distinct lines b_1, b_2, and b_3 have the slopes 0, -1, and ∞. (The middle coefficient doesn't appear explicitly in this recipe, but it is not being ignored; we apply this recipe only to lines through the point O, whose three homogeneous coefficients sum to zero.)

We then test for 2-dependence by plugging those slopes into the 3-by-3 determinant in Equation 2.4-1, getting

$$\begin{vmatrix} 0 & g+h & -h \\ -f & f-h & h \\ f & -f-g & 0 \end{vmatrix},$$

and simple algebra shows that this determinant is identically zero. □

Theorem 2.5-3 (Witness Theorem, quadratic case) *Let O be any point in the plane and let $\{(a_i, b_i)\}_{i\in[1..3]}$ be any ordered, 2-dependent block of lines through O in which no line in any pair coincides with any line in any other pair. Then, there exist complete quadrilaterals $\{(A_i, B_i)\}_{i\in[1..3]}$, each of whose six vertices lies on the corresponding line and none of whose four lines passes through O. Furthermore, any two such quadrilaterals are related by a projective transformation of the plane that fixes O and every line through O.*

Proof Since no line in any pair coincides with any line in any other pair, we deduce from Exercise 2.4-3 that the condition of 2-dependence determines each line uniquely, given the other five lines. In particular, the line a_3 is so determined.

To construct a witnessing quadrilateral, we proceed as in Exercise 2.1-2. We choose a point B_i on the line b_i, for each i in $[1..3]$, being careful only that the three points B_1, B_2, and B_3 are not collinear and that none of them coincides with O. We then construct the points $A_1 := a_1 \cap B_2B_3$ and $A_2 := a_2 \cap B_1B_3$, neither of which can possibly coincide with O. Since the lines a_1 and a_2 are distinct, we deduce that the line $a := A_1A_2$ does not pass through O. We construct the point $A_3 := a \cap B_1B_2$ and the line $a'_3 := OA_3$. The six points $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ form a complete quadrilateral, none of whose four lines passes through O. Hence, from the Projection Theorem, it follows that the ordered block of lines $\{(a_1, b_1), (a_2, b_2), (a'_3, b_3)\}$ is 2-dependent. Since the relation of 2-dependence determines

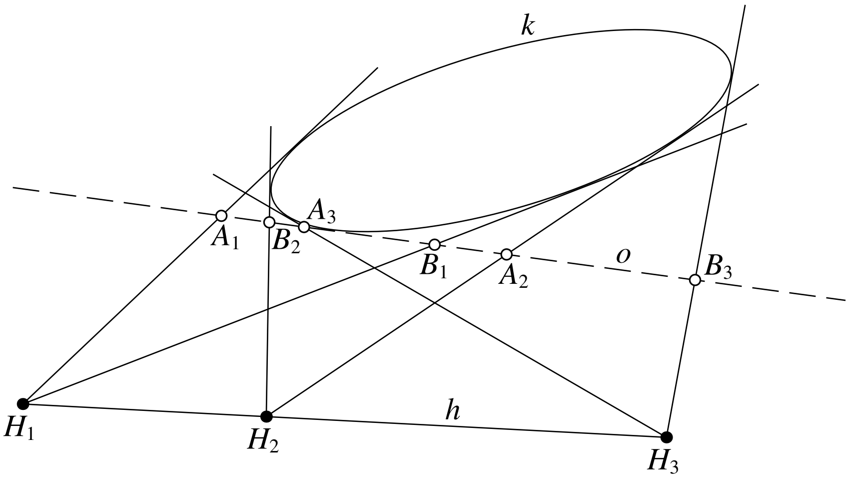

Figure 2.5: A 2-block of points along a line o whose 2-dependence is demonstrated by the collinearity of the three points (H_i), each of which is the intersection of a pair of tangents of a conic curve k, of which the line o is also a tangent.

the line a_3 uniquely, given the other five lines, the lines a_3 and a_3' must coincide, so the quadrilateral that we have constructed is a witness.

The only freedom of choice that we had, in carrying out this construction, was to select the point B_i on the line b_i, for i in $[1\,..\,3]$. It follows that every witnessing quadrilateral can be produced by that construction, provided that the points (B_i) are chosen correctly. If $\{(A_i', B_i')\}_{i\in[1..3]}$ and $\{(A_i'', B_i'')\}_{i\in[1..3]}$ are any two witnessing quadrilaterals, the unique projective transformation of the plane that maps the four points (O, B_1', B_2', B_3') to (O, B_1'', B_2'', B_3'') fixes O and every line through O and also carries the first witnessing quadrilateral to the second. □

2.6 The dual point of view

Figures 2.1 and 2.2 show two geometric tests of 2-dependence, one based on a conic curve and the other based on a complete quadrilateral. In both cases, the six scalars being tested are encoded as the slopes of six lines through a common point. Recall that duality in the projective plane interchanges the concepts of 'point' and 'line'. The duals of the situations in Figures 2.1 and 2.2 — shown in Figures 2.5 and 2.6 — are also geometric tests of 2-dependence. But the six scalars being tested in Figures 2.5 and 2.6 are the coordinates of six points along a common line.

In Figure 2.5, we have an unordered 2-block $\big\{\{A_1, B_1\}, \{A_2, B_2\}, \{A_3, B_3\}\big\}$ of points, all lying on a common line o. There is a conic k of which o is a tangent line and there are three collinear points H_1, H_2, and H_3 with the property that the tangents a_i and b_i from H_i to the conic k intersect o at A_i and B_i. In Fig-

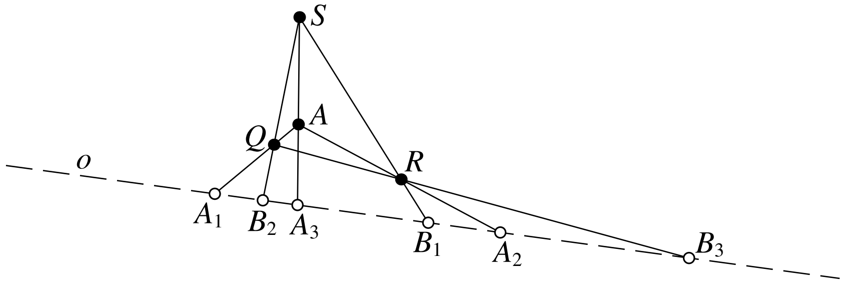

Figure 2.6: The same 2-block of points along the same line o as in the previous figure, but here demonstrated to be a quadrangular set by the existence of the complete quadrangle with vertices (A, Q, R, S).

ure 2.6, we interpret the same six points along the same line o as an ordered 2-block $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$. There are four auxiliary points A, Q, R, and S with the property that the points $\{(A_i, B_i)\}_{i\in[1..3]}$ are located where the line o stabs the lines

$$\begin{aligned} a_1 &:= AQ & b_1 &:= RS \\ a_2 &:= AR & b_2 &:= QS \\ a_3 &:= AS & b_3 &:= QR. \end{aligned}$$

The four auxiliary points and the three pairs of opposite lines joining them are called a *complete quadrangle*.[4] Three pairs of points on a common line have the conic-based property of Figure 2.5 if and only, as in Figure 2.6, they are a cross section of a complete quadrangle, and they are then called a *quadrangular set*.

Figures 2.5 and 2.6 geometrically characterize the same algebraic notion of 2-dependence that is characterized in Figures 2.1 and 2.2: If we choose any projective coordinate system on the line o, the six points $\big\{\{A_1, B_1\}, \{A_2, B_2\}, \{A_3, B_3\}\big\}$ form a quadrangular set just when their coordinates along o form a 2-dependent block of scalars. Thus, it makes no essential difference which of the dual geometric situations we discuss. The situation in Figures 2.5 and 2.6 is the one generally presented in textbooks. For our purposes, however, it is more convenient to talk about the situation in Figures 2.1 and 2.2.

Why is that? Our goal is to study projective configurations that characterize n-dependence, which means that we have to make a choice. As in Figure 2.2, we can test the n-dependence of the slopes of an n-block of hyperplanes by requiring that those hyperplanes result from projecting the points of some instance of a configuration C. Of course, the configuration C will involve various flat subspaces of various dimensions; but it must involve an n-block of points, and we

[4] Indeed, Figure 2.6 is quite similar to Figure 0.1, except that the point P has been relabeled A, the line m has been relabeled o, and the separation properties — as discussed in Exercise 2.1-3 — of the pairs $\{A_i, B_i\}$ are different. In Figure 2.6, each pair separates both of the other pairs, while, in Figure 0.1, no two pairs separate each other.

are likely to think of its other flats as built up by joining certain subsets of those points. Alternatively, as in Figure 2.6, we can test the n-dependence of the coordinates of an n-block of points by requiring that those points result from stabbing the hyperplanes of some instance of the dual configuration C^*. The configuration C^* also involves various flats of various dimensions; but it must involve an n-block of hyperplanes, and we are likely to think of its other flats as formed by intersecting certain subsets of those hyperplanes. The two situations are dual, and hence completely equivalent. But it is somewhat simpler to talk about building a configuration by working up, starting from an n-block of points, rather than by working down, starting from an n-block of hyperplanes. This is especially true when matroids are involved, since the tradition is to represent a matroid by mapping the elements of its ground set to points, rather than to hyperplanes.

Don't get confused between the duality of projective geometry and the duality of matroid theory. The duality of projective geometry interchanges points and hyperplanes, as we have been discussing. The duality of matroid theory associates, with each matroid of rank r on s points, a dual matroid of rank $s-r$ on those same s points. We are going to define a family of matroids, and each of those matroids does have a dual. But we shall have no further occasion to mention those dual matroids in this monograph. As that fact suggests, we aren't going to be appealing to any deep results of matroid theory; instead, we use matroids merely as a formal framework in which to talk about projective configurations.

Exercise 2.6-1 Given a 3-block $\big\{\{\alpha_i, \beta_i, \gamma_i\}\big\}_{i\in[1..4]}$ of planes in 3-space, all passing through a common line o, we discussed in Section 2.2 how to test the resulting block of slopes for 3-dependence geometrically, using a twisted cubic curve. Dually, given a 3-block $\big\{\{A_i, B_i, C_i\}\big\}_{i\in[1..4]}$ of points in 3-space, all lying on a common line o, describe how to test the resulting block of coordinates for 3-dependence geometrically, using a twisted cubic.

[Hint: Choose your twisted cubic so that o is the line where two osculating planes intersect.]

2.7 Auxiliary points in configurations

We can now clarify our goal a bit further. We want to find a configuration in 3-space that characterizes the notion of 3-dependence in the same way that the complete quadrilateral characterizes 2-dependence. It's going to turn out that we don't want our goal configuration to include any auxiliary points, and we don't want its twelve key points to be coplanar.

Figure 2.7 shows the ten points and seven lines of Figure 2.6, with the six points along the line o further projected from a new point O. Since the coordinates of the six points $\{(A_i, B_i)\}_{i\in[1..3]}$ along o form a 2-dependent block, the slopes of the six lines $\{(a_i, b_i)\}_{i\in[1..3]}$ through O also form a 2-dependent block.

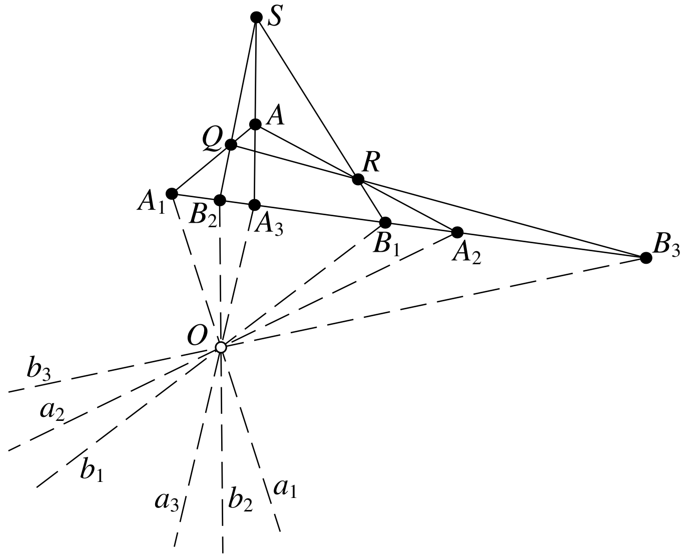

Figure 2.7: A configuration with six, collinear key points and four auxiliary points that tests 2-blocks of lines for 2-dependence.

Compare Figure 2.7 with Figure 2.2. In both cases, we have projected some six points of a configuration from a new point O, getting six lines through O whose slopes are 2-dependent. Thus, we can think of the ten solid points and the seven solid lines in Figure 2.7 as forming a configuration that, like the complete quadrilateral, characterizes 2-dependence.

Are the Projection and Witness Theorems true for the solid configuration in Figure 2.7? The Projection Theorem holds, as does the existence half of the Witness Theorem. But the uniqueness half of the Witness Theorem fails. It fails because Figure 2.7 involves the four auxiliary points (A, Q, R, S), in addition to the six key points whose projections have 2-dependent slopes. We take this as evidence that auxiliary points are bad, and we hereby restrict our search to configurations in 3-space that have no auxiliary points. That is, we demand that the twelve key points of our goal configuration be constrained only by collinearities and coplanarities among themselves, without any help from auxiliary points.

The configuration in Figure 2.7 cheats in a second way: Its six key points are collinear. This is cheating because it trivializes the projection operation. As we move the projection point O in Figure 2.7 around, the resulting block of slopes does vary. But it follows trivially, from the 2-dependence of the block of coordinates of the key points along their common line, that the resulting block of slopes is always 2-dependent. For the complete quadrilateral in Figure 2.2, where the key points are not collinear, moving the projection point O makes the block of slopes vary in a richer way. To avoid this second flavor of cheating, we restrict our search to configurations in 3-space whose twelve key points are not collinear or even coplanar, but instead span all of 3-space.

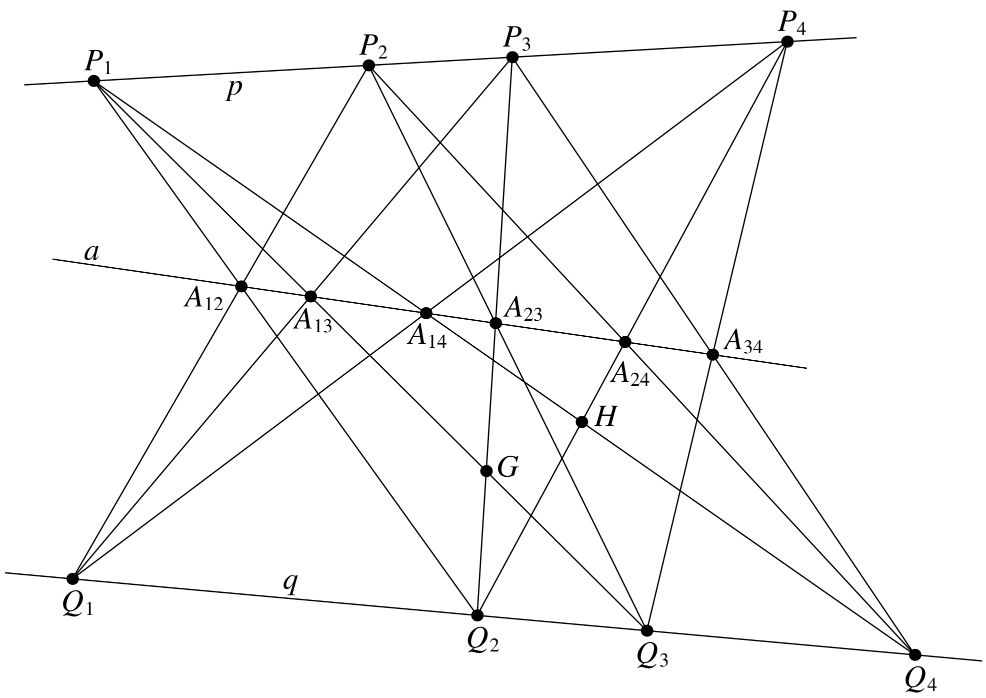

Figure 2.8: A configuration with six, collinear key points and eight auxiliary points that could test 2-blocks of lines for 2-dependence.

Exercise 2.7-1 There are lots of configurations that characterize 2-dependence, if we allow the six key points to be collinear and we allow auxiliary points to exist. Figure 2.8 shows an interesting example. The auxiliary points P_1 through P_4 are collinear, as are the auxiliary points Q_1 through Q_4. The six key points are the cross-joins $A_{ij} := P_i Q_j \cap P_j Q_i$, for $1 \le i < j \le 4$, which we also require to be collinear. (This requirement implies that the two quadruples (P_1, P_2, P_3, P_4) and (Q_1, Q_2, Q_3, Q_4) have the same cross ratio.) Prove that the 2-block $\big\{\{A_{12}, A_{34}\}, \{A_{13}, A_{24}\}, \{A_{14}, A_{23}\}\big\}$ formed by the six key points is 2-dependent.

By the way, the common line a on which all six of the key points lie is called the *axis* [11] of the projectivity from the line p to the line q that maps the four points (P_1, P_2, P_3, P_4) to the four points (Q_1, Q_2, Q_3, Q_4).

[Hint: One strategy applies Pappus's Theorem to the hexagon $P_1 Q_3 P_4 Q_2 P_3 Q_4$ to deduce that the three points $G := P_1 Q_3 \cap P_3 Q_2$, $H := P_1 Q_4 \cap P_4 Q_2$, and A_{34} are collinear. Then consider the complete quadrangle whose four vertices are P_1, Q_2, G, and H.]

2.8 Matroids

When is it that six points form an instance of the complete quadrilateral? Any four of the six points must be coplanar and certain triples of them must be collinear. In addition, because we have agreed to forbid every incidence that is not required, the remaining triples must not be collinear and no two of the six points may coincide. Thus, we can think of the complete quadrilateral as a set of abstract points and a rule that says, for each subset of those points, whether the k points in that subset must or must not lie in a common $(k-2)$-flat.[5] In this monograph, we shall refer to any such rule as a *projective configuration*. Warning: In classical projective geometry, such a rule has to have lots of other properties, including a high degree of numeric symmetry, before it can be called a configuration. We discuss that narrow sense of the word 'configuration' in Section 4.5.

In describing the incidences that are required by a configuration, there is no need to use separate terms, such as 'collinear' and 'coplanar', for sets of different sizes; the single term 'incident' will do for all. Call k points in some projective space *mutually incident* when they lie in a common $(k-2)$-flat. Algebraically, this means that the k vectors of homogeneous coordinates are linearly dependent. So four points are mutually incident when they are coplanar; three points, when they are collinear; two points, when they coincide; and one point, when it is *indeterminate* — that is, all of its homogeneous coordinates are equal to zero.[6] When k points are not mutually incident, we shall call them *mutually skew*.

Note that every superset of a mutually incident set is also mutually incident; for example, if the three points A, B, and C are collinear, then the four points A, B, C, and D must be coplanar. To make this property hold also when the smaller set is empty, we make the convention that the zero points in the empty set $\emptyset$ are mutually skew, rather than mutually incident; this agrees with the standard convention that the empty set of vectors is linearly independent.

An *instance* of a configuration in some projective space maps each abstract point of the configuration to a concrete point in that projective space so that precisely the required incidences hold. Note that, since we have defined the notion of a configuration quite broadly, it is easy to come up with configurations that have no instances. For example, one way to guarantee that no instances exist is to foolishly require that the set of points S be mutually incident while forbidding the set T from being mutually incident, for some $T \supset S$. Another way is to foolishly require that the empty set $\emptyset$ be mutually incident.

We are mostly interested in configurations that do have instances. Hence, when

[5]We use the term 'm-flat' to mean a flat subspace of dimension m; thus, a line is a 1-flat and a plane is a 2-flat. Be warned that many authors writing about matroids use 'm-flat' to mean a flat of rank m, which has dimension $m-1$.

[6]It is often desirable, as Stolfi [51] points out, to augment projective space with a unique null object of each dimension. The traditional names 'indeterminate point', 'indeterminate line', and so forth for these null objects have the advantage that they encode the dimension in a very natural way. Using the term 'indeterminate point' has the severe disadvantage of making it forever afterwards unclear, when we assume that P is a point in a projective space, whether we are allowing P to be indeterminate or not. In this monograph, fortunately, the indeterminate point arises only while we are reviewing matroids and their representations. It never arises thereafter because the particular matroids that we define have no single-element dependent sets — that is, no loops.

designing our configurations, it behooves us to avoid the two sorts of foolishness just mentioned. There is a third, less obvious sort of foolishness that we should also avoid, and any configuration that avoids all three is called a *matroid*. More formally, as we review in Chapter 3, a matroid is a configuration that satisfies a certain three axioms. Those axioms are not strong enough to guarantee that every matroid has instances; but they do eliminate three sorts of foolishness, at least, and they thereby ensure that the combinatorial structure of the configuration has some nice properties. What we are calling an *instance* of a configuration is called, in matroid theory, a *representation* of the matroid. Not every matroid is representable; but every configuration that has any instances is a representable matroid.

The cubic analog of the complete quadrilateral, which is the configuration for which we are searching, turns out to correspond to a particular representable matroid, which we shall call $B_{2,1,1}$. As that fancy name suggests, the matroid $B_{2,1,1}$ belongs to a family of matroids: the *budget matroids*, which we define in Chapter 4. There is a budget matroid $B_{b_1,\ldots,b_k}$ associated with each partition $b = b_1 + \cdots + b_k$ of an integer b into nonnegative parts, at least two of which are positive. The complete quadrilateral is the budget matroid $B_{2,1}$ — and hence that matroid characterizes 2-dependence. In an analogous way, the matroid $B_{2,1,1}$ characterizes 3-dependence. Unfortunately, this pattern doesn't continue: The matroid $B_{2,1,1,1}$ characterizes some property that is stronger than 4-dependence.

Each of the parts b_j in the partition $b = b_1 + \cdots + b_k$ that determines the budget matroid $B_{b_1,\ldots,b_k}$ actually plays a double role. From the point of view of matroid theory, there is no reason not to separate those two roles, assigning them to independent parameters b_j and d_j. Chapter 5 does this, thereby defining a larger family of matroids, which we call the *budgetary matroids*. From a geometric point of view, however, this generalization seems to be of little value. Indeed, many of the budgetary matroids turn out not to be representable at all.

The representability of the budget matroids makes for a happier story. If the budget matroid $B_{b_1,\ldots,b_k}$ is representable, it is fairly easy to see that the budget matroids that result from reducing or eliminating one or more parts in the partition $b = b_1 + \cdots + b_k$ are also representable. So, in proving representability, the easy cases are those with few parts and with small parts. In Chapter 6, we show that every budget matroid of the form $B_{m,n}$, with precisely two parts, is representable over the rational numbers and hence over every field of characteristic zero. In Chapter 7, we show the same result for the budget matroids of the form $B_{m,1,1}$. Neither proof is very difficult, but both are fairly long. The latter proof involves a theorem that can be interpreted as generalizing Pappus's Theorem to higher dimensions.

Since the budget matroid $B_{m,1,1}$ is representable for every positive m, the particular matroid $B_{2,1,1}$ is representable. Unfortunately, the scheme that we develop for constructing representations of the matroids $B_{m,1,1}$ in Chapter 7 doesn't seem to be helpful in studying the relationship between the matroid $B_{2,1,1}$ and the concept of 3-dependency. To get a handle on that relationship, we need a different construction, one that makes its choices in a different order. That second construction is based on null systems.

Exercise 2.8-1 Let P be a finite set of points in a projective space. Show that the points in P are mutually skew just when, for all ways of partitioning P into two subsets $P = Q \cup R$, the flats $\mathrm{Span}(Q)$ and $\mathrm{Span}(R)$ are skew in the standard sense of the word 'skew' — that is, they do not meet.

2.9 Null systems

The Principle of Duality in projective geometry tells us that the entire structure of projective n-space is preserved if we interchange, for each k, the notion of a k-flat with the notion of an $(n-k-1)$-flat, provided that we also interchange the notion 'lies in' with the notion 'passes through'. If S is a projective n-space, its *dual space*, written S^*, is the n-space whose points are the hyperplanes of S, and vice versa. The double dual S^{**} is thus S, once again.

A *polarity* is a self-inverse (or *involutive*) projective transformation from a projective space S to its dual S^*. For example, if S is a plane, a polarity maps each point P of S to a line P^* of S and maps each line ℓ to a point ℓ^* in such a way that:

- For all points P, we have $P^{**} = P$.
- For all lines ℓ, we have $\ell^{**} = \ell$.
- For all points P and lines ℓ, the point P lies on the line ℓ if and only if the line P^* passes through the point ℓ^*.

The line P^* is called the *polar line* of the point P in the chosen polarity, and the point ℓ^* is called the 'polar point' or *pole* of the line ℓ. In 3-space, a polarity maps a point P to its *polar plane* P^*, a line ℓ to its *polar line* ℓ^*, and a plane π to its 'polar point' or *pole* π^*.

Polarities come in two distinct types, called — somewhat confusingly — *polar systems* and *null systems* [23].

Polar systems are the ones most often discussed in textbooks [8, 48];. They are intimately related to quadric hypersurfaces, and we are going to review this relationship algebraically in Section 8.1. For now, let's recall the geometric intuition by considering a nondegenerate conic k in the real projective plane. Given any point P outside of the conic k, there are two tangents to k that pass through P, and there is a unique line P^* that joins the two points of tangency. Conversely, given any line ℓ that intersects the conic at two points, we can intersect the tangents at those two points to get back to the point ℓ^*. This mapping from P to P^* can be extended into a polarity on the entire plane, and such polarities are called *polar systems*. Note that, in the polar system associated with a conic, a point P lies on its own polar line P^* just when P lies on the conic.

In 3-space, it is nondegenerate quadric surfaces that have associated polar systems. The polar plane of a point P is the plane P^* containing all of the points of the quadric at which the tangent cone from P to the quadric touches the quadric — a point P on the outside of a non-ruled quadric, such as an ellipsoid, is the easiest case to visualize. Just as in the planar case, a point P lies on its polar plane P^* just when P lies on the quadric.

Null systems are a different type of polarity, perhaps even prettier than polar systems. But they don't get talked about very often, largely because they exist only in spaces of odd dimension — so there are no null systems in the plane. The distinguishing feature of a null system is that *every* point P lies on its polar hyperplane P^*, not just the points on some quadric hypersurface.

One geometric situation in 3-space in which null systems arise is the study of twisted cubic curves. Given a twisted cubic and given some point P, there are at most three points on the cubic at which the osculating plane to the cubic passes through P. When there are three such points, the plane P^* that they determine always passes through P; for further discussion of this fascinating fact and its relationship to the notion of harmonic conjugacy, see Exercise 9.5-2. This mapping from the point P to the plane P^* can be extended to a null system on the entire 3-space, and thus twisted cubics are one source of null systems.

In Chapter 8, we review polarities in general, using linear algebra, and we show that every polarity is either a polar system or a null system. We show how to define a null system in 3-space by using a *skew-Pappian hexagon*: a hexagon whose vertices lie, alternately, on two skew lines. In the special case in which those two lines are both coordinate axes, we give explicit formulas for computing in the associated null system.

Exercise 2.9-1 Suppose that the point P lies inside the nondegenerate conic k in the real projective plane. Give several recipes for finding the line P^* that is polar to P in the polar system associated with k.

[Answer: One recipe is to find the two conjugate, complex lines through P that are tangent to the complexification of k. The line P^* is the real line that joins the two conjugate, complex points of tangency. Alternatively, we can stick to the real numbers by drawing any two lines ℓ and m through P, each of which must intersect the conic k at two points. The polar P^* is the line joining ℓ^* to m^*.]

2.10 The cubic case and beyond

By exploiting null systems, we can construct representations of the budget matroid $B_{2,1,1}$ in a different way, a way that helps to clarify the relationship with the concept of 3-dependency. Chapter 9 uses this approach to prove the cubic analogs of the Projection Theorem and the Witness Theorem, thus verifying that the budget matroid $B_{2,1,1}$ does indeed characterize 3-dependence.

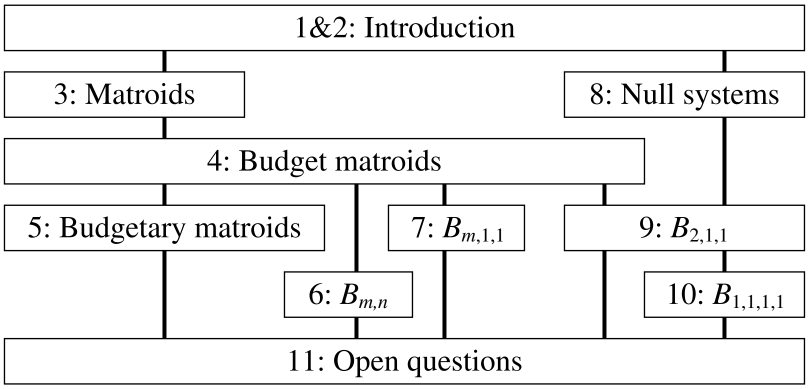

Figure 2.9: The dependencies between the chapters in this monograph.

As it happens, our machinery based on null systems also yields a way to construct representations of the budget matroid $B_{1,1,1,1}$, as we discuss in Chapter 10.

We close in Chapter 11 by discussing some open questions.

Figure 2.9 gives a rough indication of the dependencies between the chapters in this monograph. The most significant choice comes after Chapter 4, at which point a reader might reasonably jump to any one of following four chapters.

Chapter 3

A review of matroids

3.1 The axioms

We shall define matroids by talking about their independent sets — in particular, by giving the axioms that the family of independent sets must satisfy. When matroid theory is applied to the study of projective configurations, the word 'independent' means 'required to be mutually skew', while 'dependent' means 'required to be mutually incident'.

A *matroid* is a finite set G, called the *ground set*, together with a family $\mathcal{I}$ of subsets of G, called the family of *independent sets*, subject to the following three axioms:

Empty-Set Axiom The empty set is independent: $\emptyset \in \mathcal{I}$.

Subset Axiom Every subset of an independent set is independent: If $Y \in \mathcal{I}$ and $X \subseteq Y$, then $X \in \mathcal{I}$.

Augmentation Axiom Given two independent sets of different sizes, there is always an element of the larger, not already in the smaller, that can be added to the smaller without destroying its independence: If $X \in \mathcal{I}$ and $Y \in \mathcal{I}$ with $|Y| > |X|$, then there exists $e \in Y \setminus X$ such that $X \cup \{e\} \in \mathcal{I}$. (Note that we use a backslash to denote the operation 'set minus'.)

We claimed, in Section 2.8, that any configuration that has any instances is a matroid. In support of this claim, we pointed out that the Empty-Set Axiom and the Subset Axiom must hold of any such configuration. To finish the job, note that the Augmentation Axiom holds as well. For example, suppose that the sets $X = \{P, Q, R\}$ and $Y = \{A, B, C, D\}$ are both independent in some instance of a configuration. Then, the four points A, B, C, and D do not lie in any single plane, so at least one of them must lie outside of the plane PQR; that one can be added to the set X without destroying its independence.

3.2 Elementary notions

One of the pretty things about matroids is that there are simple systems of axioms that describe them from lots of different points of view. So far, we have been talking about a matroid in terms of its independent sets. Alternatively, we could talk about its bases, its circuits, its rank function, its flats, its spanning sets, or what have you. We shan't be needing those other axiom systems, but we shall be using some of the concepts on which they are based. For more about these topics, see Chapter 1 of Oxley [37].

Let S be any subset of the ground set of a matroid M — so S is not necessarily independent. An independent subset $T \subseteq S$ is called *maximal* when adding any element of $S \setminus T$ to T would destroy its independence. Given two independent subsets $X \subseteq S$ and $Y \subseteq S$ of different sizes, the Augmentation Axiom tells us that the smaller cannot be maximal; so all of the maximal, independent subsets of a set S have the same cardinality. That common cardinality is called the *rank* of the set S. For example, in the complete quadrilateral shown in Figure 2.2, both the independent set $\{A_1, A_2\}$ and the dependent set $\{A_1, A_2, A_3\}$ have rank 2.

In the special case where the set S is the entire ground set of the matroid, there is some special terminology. A maximal independent subset of the entire matroid is called a *base*, and the common cardinality of all of the bases is called the *rank* of the matroid. For example, the complete quadrilateral is a matroid of rank 3, with $\{B_1, B_2, B_3\}$ and $\{A_1, A_2, B_1\}$ as two of its bases.

Next, we turn from maximal independent sets to minimal dependent sets. A dependent set in a matroid is called a *circuit* when removing any element from it would destroy its dependence.[1] In a typical matroid, there are circuits of various cardinalities; for example, in the complete quadrilateral once again, the sets $\{A_1, A_2, A_3\}$ and $\{A_1, A_2, B_1, B_2\}$ are both circuits.

A subset of a matroid is called a *flat* when adding any element to it would increase its rank. The entire matroid is trivially a flat, whose rank equals the rank of the matroid. A *hyperplane* is a flat whose rank is one less than the rank of the matroid, while the words *line* and *plane* are used for flats of ranks 2 and 3. For example, in the complete quadrilateral, the set $\{A_1, A_2, A_3\}$ is both a line and a hyperplane.

It follows from the axioms that the intersection of any two flats is always a flat, so there is a unique smallest flat containing an arbitrary set S. That flat is called the *span* of S. For example, in the complete quadrilateral, the span of the set $\{A_1, A_2\}$ is the line $\{A_1, A_2, A_3\}$. Another way to define a base is as a minimal set whose span is the entire matroid. Another way to define a hyperplane is as a maximal set whose span is less than the entire matroid.

[1] While we are using matroids to encode the incidence patterns of points in projective configurations, they are used also to encode the cycle structure of graphs, and that latter context is where the term 'circuit' comes from.

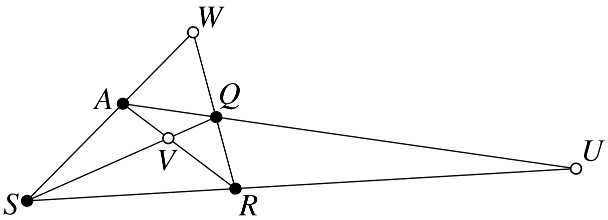

Figure 3.1: A complete quadrangle, together with its three diagonal points.

3.3 Representations

A *representation* of a matroid M is a function that maps the elements of the ground set of M to points in some projective space in such a way that a set of k elements is independent precisely when the k points that represent those elements are mutually skew — that is, when the k vectors of homogeneous coordinates of those points are linearly independent.

A representation of a matroid is allowed to map certain of its elements to the indeterminate point — the null object of dimension 0, all of whose homogeneous coordinates are zero. Indeed, if a singleton set $\{e\}$ is dependent in the matroid, the element e must be mapped to the indeterminate point by any representation. Such an element e is called a *loop*.[2] On the other hand, elements of a matroid that are not loops must be mapped to determinate points by any representation.

A representation of a matroid is also allowed to map distinct elements to the same point. Suppose that the two-element set $\{e, f\}$ in a matroid is dependent, but that neither e nor f is a loop — which means that the set $\{e, f\}$ is also a circuit. The element e must then be mapped to the same determinate point as f by any representation. In this situation, the two elements e and f are called *parallel*.[3]

Whether a matroid is representable or not can depend on the choice of the scalar field, the field over which our projective spaces are built. There are matroids that are representable over any field, matroids that are representable over some fields but not over others, and matroids that are not representable over any field. In this monograph, we are primarily interested in representability over the standard fields of characteristic zero: the rationals, the reals, and the complexes; but Exercises 3.3-1, 4.2-2, and 6.2-3 are exceptions to that rule.

Exercise 3.3-1 Figure 3.1 shows a complete quadrangle, lying in the real projective plane, along with its three *diagonal points* $U := AQ \cap RS$, $V := AR \cap QS$, and $W := AS \cap QR$, the points where the three pairs of opposite sides intersect. Note that the diagonal points are not collinear. Verify that they would be collinear just if the characteristic of the scalar field were 2. Using this, find two matroids of rank 3 on seven points, the first of which is representable just over those fields

[2]'Loop' is another term that is motivated by the application of matroids to graph theory.

[3]Another term motivated by graph theory.

of characteristic 2, while the second is representable just over those fields whose characteristic is not 2. These two matroids are known as *Fano* and *non-Fano*.

[Hint: We can choose a coordinate system for the plane in which the vertices of the quadrangle have the homogeneous coordinates $A = [1, 1, 1]$, $Q = [1, 0, 0]$, $R = [0, 1, 0]$, and $S = [0, 0, 1]$. The diagonal points then have the coordinates $U = [0, 1, 1]$, $V = [1, 0, 1]$, and $W = [1, 1, 0]$, so they are collinear just when $2 = 0$. The non-Fano matroid has, as its dependent sets, all sets of size at least four, along with the six triples of points that are collinear in Figure 3.1. In the Fano matroid, the triple $\{U, V, W\}$ is also dependent.]

3.4 Minors

There are two ways to make a matroid smaller, one obvious and the other less so. Let M be a matroid, let H be a subset of its ground set G, and suppose that we want to build a matroid M' on the smaller ground set $G' := G \setminus H$. The obvious approach is to take, as the independent sets of M', precisely those subsets of G' that, viewed as subsets of G, are independent in M. This process is called *deleting* the elements of H, and the resulting matroid is written $M\backslash H$. In the less obvious approach, we choose some maximal independent subset I of H — it follows from the axioms that it doesn't matter which we choose — and we let a subset X' of G' be independent in M' precisely when the subset $X := X' \cup I$ of G is independent in M. This process is called *contracting*[4] the elements of H, and the resulting matroid is written M/H. A matroid M' is called a *minor* of a matroid M if we can get from M to M' by some sequence of deletions and contractions.

It is easy to see that deletions preserve representability. If H is any subset of the ground set of a matroid M, any representation of M can be converted into a representation of the deleted matroid $M\backslash H$ by simply forgetting about the points to which the elements of H were mapped.

It is only slightly trickier to show that contractions also preserve representability. Let M be a representable matroid, let H be a subset of its ground set G, and fix some representation φ of M, say in a projective space S of dimension n. The points $\varphi(H)$ to which the elements of H are mapped span a certain flat F in S, say of dimension k. Consider all flats of dimension $k + 1$ in S that contain all of the points in $\varphi(H)$ and hence include F. Those $(k + 1)$-flats themselves are the 'points' of a projective space S' of dimension $n - k - 1$. For example, the lines in 3-space through a fixed point form a projective plane, while the planes in 3-space through a fixed line form a projective line.[5] We can construct a representation φ' of the contracted matroid M/H in the space S' by mapping each element e of $G \setminus H$ to the $(k+1)$-flat $\varphi'(e) := \mathrm{Span}(\varphi(H \cup \{e\}))$. Typically, the span of the points in $\varphi(H \cup \{e\})$ is a $(k + 1)$-flat in S that includes F, so it is a point of S'. When e is

[4] Yet another term motivated by graph theory.

[5] In these examples, it doesn't matter whether the ambient 3-space is projective or affine.

a loop in the contracted matroid M/H, the point $\varphi(e)$ lies in the flat F, and we set $\varphi'(e)$ to be the indeterminate point of S'.

Exercise 3.4-1 Let M be a matroid of rank r. The element e of M does not belong to any base just when e is a loop. In this case, show that deleting e and contracting e produce the same matroid $M\backslash\{e\} = M/\{e\}$, whose rank is r. On the other hand, suppose that the element e of M belongs to every base; such an element e is called a *coloop*. In this case also, show that deleting e and contracting e produce the same matroid $M\backslash\{e\} = M/\{e\}$, but its rank is $r-1$. If the element e is neither a loop nor a coloop, show that deleting e gives a matroid $M\backslash\{e\}$ of rank r, while contracting e gives a matroid $M/\{e\}$ of rank $r - 1$.

Chapter 4

The budget matroids

4.1 Initial examples

Each budget matroid $B_{b_1,\ldots,b_k}$ is associated with a partition $b = b_1 + \cdots + b_k$ of a *total budget b* into k *column budgets* b_1 through b_k. To motivate the definition of the budget matroids, let's consider a couple of examples of partitions and their associated matroids.

We begin with the partition $3 = 1 + 1 + 1$. The budget matroid $B_{1,1,1}$ turns out to be the Pappus matroid — the matroid whose representations are instances of the configuration that arises in Pappus's Theorem.

As shown in Figure 4.1, let A_1, A_2, and A_3 be collinear points in the plane, let B_1, B_2, and B_3 also be collinear,[1] and, whenever $\{i, j, k\} = \{1, 2, 3\}$, let $C_i := A_j B_k \cap A_k B_j$. Pappus's Theorem tells us that the three intersection points C_1, C_2, and C_3 of the opposite sides of the hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ are also collinear. If

[1] I apologize for overusing the letter 'B': When a B has a single subscript, it denotes a point; when it has multiple subscripts, it denotes a budget matroid.

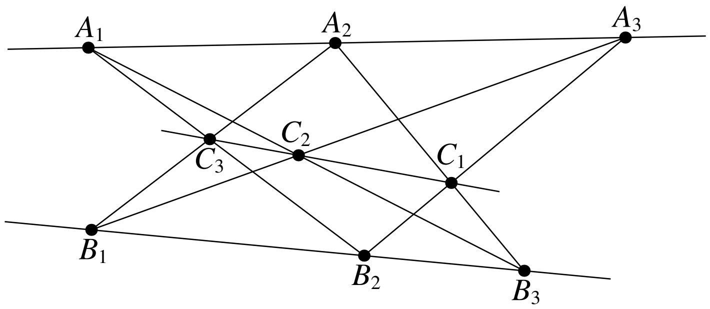

Figure 4.1: An instance of the Pappus configuration.

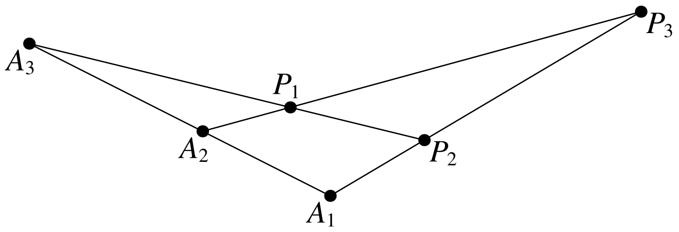

Figure 4.2: An instance of the complete quadrilateral — as in Figure 2.2, but with the vertices (B_i) relabeled (P_i).

we arrange the nine points of the Pappus configuration in a matrix, like this,

$$\begin{array}{c} \begin{array}{ccc} 1 & 1 & 1 \end{array} \\ \begin{pmatrix} A_1 & B_1 & C_1 \\ A_2 & B_2 & C_2 \\ A_3 & B_3 & C_3 \end{pmatrix} \end{array},$$

we can describe the nine lines of the configuration as follows:

- The three points in a column are collinear. (But the three points in a row are not so constrained.)

- For each of the six possible ways of choosing one point from each row so that one is taken, also, from each column, the three points so chosen are collinear.

By the way, we often write the j^{th} column budget b_j at the head of the j^{th} column, as a reminder. In this example, all three column budgets are 1.

As a second example, consider the partition $3 = 2 + 1$. A representation of the budget matroid $B_{2,1}$ is, it turns out, precisely a complete quadrilateral, as shown in Figure 4.2. If we arrange the six points of such a quadrilateral in the matrix

$$\begin{array}{c} \begin{array}{cc} 2 & 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix} \end{array},$$

its four lines can be described as follows:

- The three points in the A column are collinear, while the three points in the P column are constrained only to be coplanar.

- For each of the three possible ways of choosing one point from each row so that two are taken from the P column and one from the A column, the three points so chosen are collinear.

As for naming the points in a representation of a budget matroid, our standard convention will be to number the rows and to associate the letters A, B, C, and so forth with the columns, from left to right. But for matroids of the special form $B_{m,1,1,\ldots,1}$, where all of the column budgets except the first are ones, it is often convenient to use a special letter for the first column and to start with A for the second column. Here, for the matroid $B_{2,1}$, we have used the letter P for the first column, on the grounds that 'P' stands for '(co)planar'.

Note that we are failing to exploit all of the available symmetry in the two examples above. The four lines of a complete quadrilateral play entirely symmetric roles. But to describe that quadrilateral as a representation of the matroid $B_{2,1}$, we must choose one of its four lines to become the A column. In a similar way, to describe a Pappus configuration as a representation of the matroid $B_{1,1,1}$, we must choose one of its three families of three disjoint lines to become the three columns.

Failing to exploit a symmetry is often a mistake, but the unexploited symmetries in those two cases are sporadic accidents. Consider a larger budget matroid, say the matroid $B_{1,1,1,1}$ associated with the partition $4 = 1 + 1 + 1 + 1$:

$$\begin{array}{c} \begin{array}{cccc} 1 & 1 & 1 & 1 \end{array} \\ \begin{pmatrix} A_1 & B_1 & C_1 & D_1 \\ A_2 & B_2 & C_2 & D_2 \\ A_3 & B_3 & C_3 & D_3 \\ A_4 & B_4 & C_4 & D_4 \end{pmatrix} \end{array}.$$

A representation of $B_{1,1,1,1}$ consists of sixteen points in 3-space with the following two properties:

- The four points in each column are collinear.
- For each of the 24 possible ways of choosing one point from each row so that one is taken also from each column, the four points so chosen are coplanar.

Since lines and planes are flats of different dimensions, there is no chance that some additional symmetry of the configuration might intermix the four lines required by the first property with the 24 planes required by the second. (For more about when such intermixing is possible and when it isn't, see the analysis of the automorphisms of the budget matroids in Exercises 4.3-2 and 4.3-3.)

Exercise 4.1-1 (from Jorge Stolfi, and in preparation for Sections 10.3 and 10.4) If we write a point in the projective plane as a triple $[w, x, y]$ of homogeneous coordinates, verify that the nine points

	A	B	C
1	$[1, 0, 0]$	$[0, 1, 0]$	$[0, 0, 1]$
2	$[p, q, q]$	$[q, p, q]$	$[q, q, p]$
3	$[q, p, p]$	$[p, q, p]$	$[p, p, q]$

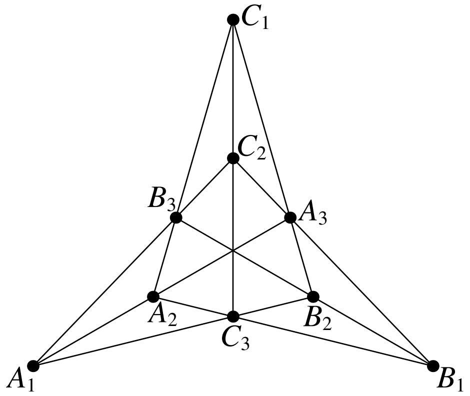

Figure 4.3: A representation of the Pappus matroid $B_{1,1,1}$ in which all six permutations of the A, B, and C columns can be achieved via Euclidean symmetries.

form a Pappus configuration — that is, they represent the budget matroid $B_{1,1,1}$ — whenever the ratio p/q does not lie in $\{-2, -1, -\frac{1}{2}, 0, 1, \infty\}$. What goes wrong for each of the six forbidden ratios?

Figure 4.3 shows an embedding in the Euclidean plane of the case $p/q = 3$. Note that the three lines $A_1A_2A_3$, $B_1B_2B_3$, and $C_1C_2C_3$ are concurrent. Indeed, those lines will concur for any ratio p/q, at the point $[1, 1, 1]$. That concurrence is not required in a representation of the matroid $B_{1,1,1}$ — indeed, it did not happen in Figure 4.1 — but it is allowed. It is allowed despite the fact that, for any subset S of the nine points of the matroid, whenever the points in S are not required to be mutually incident, they are required to be mutually skew.

These Pappus configurations are particularly symmetric: Any permutation of the three homogeneous coordinates w, x, and y gives us a projective transformation of the plane that is a symmetry of the configuration in which the rows are preserved and the columns are correspondingly permuted. Furthermore, when the configuration is embedded in the Euclidean plane as in Figure 4.3, those six projective transformations become Euclidean transformations — in fact, the symmetries of the equilateral triangle $\triangle A_i B_i C_i$.

[Hint: When $p/q = -2$, the points A_2, B_2, and C_2 all lie on the line with homogeneous coefficients $\langle 1, 1, 1\rangle$ — which, in Figure 4.3, is the line at infinity.]

4.2 The definition

With those three examples to guide us, we can now define the budget matroids. We're going to call them 'matroids' right from the start, even though we shan't get around to proving that the matroid axioms are satisfied until Section 5.3.

Let $b = b_1 + \cdots + b_k$ be some partition of a nonnegative integer b, called the *total budget*, into k nonnegative parts, called the *column budgets*. For reasons that we discuss in Exercise 4.4-3 and Section 5.1, we require that at least two of the column budgets be positive, which implies that $k \geq 2$. We are going to define a *budget matroid* associated with this partition, which we shall denote by $B_{b_1,\ldots,b_k}$. A representation of the matroid $B_{b_1,\ldots,b_k}$ involves bk points in projective $(b-1)$-space, which we think of as organized into a b-by-k matrix:

$$\begin{array}{c} \begin{array}{cccc} \scriptstyle b_1 & \scriptstyle b_2 & & \scriptstyle b_k \end{array} \\ \begin{pmatrix} E_{11} & E_{12} & \ldots & E_{1k} \\ E_{21} & E_{22} & \ldots & E_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ E_{b1} & E_{b2} & \ldots & E_{bk} \end{pmatrix} \end{array}$$

Note that we have written the j^{th} column budget at the head of the j^{th} column, as we did in the three examples above. The column budgets constrain the locations of the points in two ways:

- The b points in the j^{th} column are constrained to lie in a common flat of dimension b_j.
- For each of the $\binom{b}{b_1 \ldots b_k}$ possible ways of choosing one point from each row so that, for all j in $[1\,..\,k]$, precisely b_j points are taken from the j^{th} column, the b points so chosen are mutually incident — that is, rather than spanning the entire ambient $(b-1)$-space, they lie in a common hyperplane.

Let X be a subset of the matrix (E_{ij}) of points that contains precisely one point from each row and contains, for each j in $[1\,..\,k]$, precisely b_j points from the j^{th} column. We shall call such a set X *perfect*, on the grounds that it meets each column budget perfectly. Note that the multinomial coefficient $\binom{b}{b_1 \ldots b_k}$ counts the perfect sets. The second rule above requires that, in any representation of the budget matroid $B_{b_1,\ldots,b_k}$, every perfect set of points must be mutually incident. Hence, in the budget matroid $B_{b_1,\ldots,b_k}$ itself, every perfect sets of elements is going to be dependent.

To define the matroid $B_{b_1,\ldots,b_k}$, we need a set of elements and a rule for independence. The ground set of the matroid $B_{b_1,\ldots,b_k}$ is the set of entries of a b-by-k matrix

$$\mathbf{e} := \begin{array}{c} \begin{array}{cccc} \scriptstyle b_1 & \scriptstyle b_2 & & \scriptstyle b_k \end{array} \\ \begin{pmatrix} e_{11} & e_{12} & \ldots & e_{1k} \\ e_{21} & e_{22} & \ldots & e_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ e_{b1} & e_{b2} & \ldots & e_{bk} \end{pmatrix} \end{array},$$

which we shall refer to as the *ground matrix*. We define a subset X of the ground matrix $\mathbf{e}$ to be *perfect* when X contains precisely one element from each row and, for each j, contains precisely b_j elements from the j^{th} column. We define a subset X of $\mathbf{e}$ to be *independent* when it satisfies the following three rules:

Ambient Rule X contains at most b elements overall;

Column Rule X contains at most $b_j + 1$ of the b elements in the j^{th} column, for each j in $[1\,..\,k]$; and

Perfect Rule X has no perfect subsets.

What is the effect of these three rules?

When talking about the Pappus configuration $B_{1,1,1}$ and the complete quadrilateral $B_{2,1}$, we implicitly assumed that all of the points involved lay in a common plane. In the general case, we want all of the points in any representation of the budget matroid $B_{b_1,\ldots,b_k}$ to lie in an ambient projective space of dimension $b - 1$ and hence of rank b. The Ambient Rule guarantees this by making any set with more than b elements dependent.

Furthermore, we want the b points in the j^{th} column to lie in a common flat of dimension b_j and hence of rank $b_j + 1$. The Column Rule guarantees this in an analogous way.

As for the Perfect Rule, if a set X has a proper subset that is perfect, then X must contain more than b elements, so X also violates the Ambient Rule. Thus, in the presence of Ambient Rule, the effect of the Perfect Rule is just to make the perfect sets themselves dependent, as we intended to do.

The special case of a zero column budget deserves comment. For one thing, if $b_j = 0$, then the Column Rule forces all of the elements of the j^{th} column to lie in a common 0-flat — that is, to coincide, say at some point Z_j. But more importantly, none of the elements of the j^{th} column belong to any perfect set. Hence, the point Z_j does not participate in any incidences that involve any other points; instead, it lies in the ambient space in general position with respect to the other points of the representation. Exercise 4.2-2 shows that zero column budgets are not of much geometric interest.

Exercise 4.2-1 Construct a representation of the budget matroid $B_{1,1}$.

[Answer: Map the four elements of the ground matrix to the points

$$\begin{array}{c} \begin{array}{cc} 1 & 1 \end{array} \\ \begin{pmatrix} P & Q \\ Q & P \end{pmatrix} \end{array},$$

where P and Q are any two distinct points along a line.]

Exercise 4.2-2 Construct a representation of the budget matroid $B_{2,1,0,0,\ldots,0}$ in the projective plane, where there are, say, a thousand columns whose budgets are zero. Note that you wouldn't succeed if the projective plane were built over a small finite field.

[Answer: Map the six elements

$$\begin{pmatrix} e_{11} & e_{12} \\ e_{21} & e_{22} \\ e_{31} & e_{32} \end{pmatrix}$$

of the first two columns of the ground matrix to the six vertices

$$\begin{array}{c} \begin{matrix} \scriptstyle 2 & \scriptstyle 1 \end{matrix} \\ \begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix} \end{array}$$

of some complete quadrilateral. For $j > 2$, map all three elements e_{1j}, e_{2j}, and e_{3j} of the j^{th} column to a common point Z_j in the plane, where the points Z_3 through Z_{1002} are chosen so that

- none of them lies on the line determined by any two vertices of the quadrilateral (neither on a line of the quadrilateral itself nor on one of its diagonals),
- none of the lines determined by any two of them passes through any vertex of the quadrilateral, and
- no three of them are collinear.]

Exercise 4.2-3 Recall that the *rank* of a matroid is the common cardinality of all of its maximal independent sets. Show that the rank of the budget matroid $B_{b_1,\ldots,b_k}$ is the total budget $b := b_1 + \cdots + b_k$.

[Hint: The Ambient Rule guarantees that the rank is at most b. To show that it is at least b, it suffices to construct some independent set with b elements. One choice is to take the top b_j elements from the j^{th} column, for all j in $[1 \,..\, k]$. Note that the resulting set is not perfect, because we have required that at least two of the column budgets be positive.]

Exercise 4.2-4 Recall that a *circuit* in a matroid is a minimal dependent set. Show that the circuits of the budget matroid $B_{b_1,\ldots,b_k}$ are precisely

perfect circuits the perfect sets,

column circuits the subsets — if any — of the j^{th} column of size $b_j + 2$, and

ambient circuits those sets of size $b + 1$ that do not include any perfect circuit or any column circuit as a subset.

Exercise 4.2-5 A matroid is called *simple* when all of its dependent sets have size at least 3; that is, there are no circuits of size 1 or size 2 — no loops and no parallel pairs. Which budget matroids are simple?

[Answer: No budget matroid has any loops. A budget matroid has parallel pairs either when its total budget is 2 or when some column budget is 0.]

4.3 Further examples

Let's explore the world of budget matroids a bit, by taking various partitions $b = b_1 + \cdots + b_k$ and seeing what matroids $B_{b_1,\ldots,b_k}$ we get. Since we have agreed that zero column budgets are not of much geometric interest, let's consider partitions into at least two parts, all of which are positive.

For each partition $b = b_1 + \cdots + b_k$ that we consider, we want to get some geometric intuition for what a representation of the budget matroid $B_{b_1,\ldots,b_k}$ looks like. Among other things, we want to count the degrees of freedom that are involved. If M is any representable matroid, let us denote by $\#(M)$ the number of degrees of freedom that are involved in choosing a representation of the matroid M, sitting in a fixed projective space of the appropriate dimension — which is one less than the rank of M. We know from Exercise 4.2-3 that the rank of the budget matroid $B_{b_1,\ldots,b_k}$ is simply its total budget $b := b_1 + \cdots + b_k$, so the appropriate dimension for the ambient space when computing the freedom $\#(B_{b_1,\ldots,b_k})$ is $b - 1$. For example, a representation of the budget matroid $B_{2,1}$ is a complete quadrilateral, lying in some plane, so we have $\#(B_{2,1}) = 8$ — two degrees of freedom in each of the quadrilateral's four lines. If the matroid M is not representable, we shall write $\#(M) = \perp$.

Warning: This counting of degrees of freedom is one situation where the properties of the scalar field might make a difference. For example, in Chapter 10, we study a construction for a generic representation of the budget matroid $B_{1,1,1,1}$ that has twenty free, scalar parameters. But one step of that construction involves solving a quadratic equation; so the construction may fail, if we attempt to carry it out over a scalar field in which some scalars don't have square roots. Thus, we have $\#(B_{1,1,1,1}) = 20$ over the complex numbers, but it is not clear whether that same result holds over the rationals or over the reals. In those cases where it makes a difference, let's agree to count degrees of freedom over the complex numbers. Indeed, as we discuss in Section 11.1, what we really mean by the informal phrase 'degrees of freedom' is the dimension of some variety, and it is easiest to talk about varieties and their dimensions over a field of scalars, like the complex numbers, that is algebraically closed.

The only budget matroid of rank 2 whose column budgets are all positive is $B_{1,1}$. As we saw in Exercise 4.2-1, a representation of $B_{1,1}$ consists of two points on a line, so we have $\#(B_{1,1}) = 2$.

There are two budget matroids of rank 3 whose column budgets are all positive: the complete quadrilateral $B_{2,1}$ and the Pappus configuration $B_{1,1,1}$. We pointed out a moment ago that $\#(B_{2,1}) = 8$. It is just as easy to calculate that $\#(B_{1,1,1}) = 10$: five degrees of freedom in the three collinear A-points and five more in the B-points, after which the C-points are uniquely determined.

The Pappus configuration differs from the complete quadrilateral in an important respect: the Pappus constraints are redundant. It takes twelve degrees of freedom to choose six arbitrary points in the plane, and, if those points are to form a complete quadrilateral, four triples of them must be collinear. Each collinearity costs one degree of freedom, so — because those four constraints are nonredundant — we are left with $\#(B_{2,1}) = 12 - 4 = 8$ degrees of freedom. In the Pappus case, on the other hand, it takes eighteen degrees of freedom to choose nine arbitrary points in the plane, and the Pappus configuration has nine collinearities, each of which naively costs one degree of freedom. But we do not have $\#(B_{1,1,1}) = 18 - 9 = 9$. Instead, enforcing all nine collinearities together costs only eight degrees of freedom, since Pappus's Theorem tells us that any eight of the collinearities together imply the ninth.[2] Hence, we have $\#(B_{1,1,1}) = 18 - 8 = 10$, in agreement with our earlier count. Redundancies of this type become more and more prevalent, as the budgets increase.

4.3.1 The partition $4 = 3 + 1$

We continue our explorations by moving to rank 4. A representation of the budget matroid $B_{3,1}$ consists of eight points in 3-space:

$$\begin{array}{c} \begin{array}{cc} \scriptstyle 3 & \scriptstyle 1 \end{array} \\ \begin{pmatrix} T_1 & A_1 \\ T_2 & A_2 \\ T_3 & A_3 \\ T_4 & A_4 \end{pmatrix} \end{array}.$$

The four T-points span the entire 3-space, so they are the vertices of a tetrahedron, while the four A-points are collinear, along a line a. There are four perfect constraints, which require the point A_i to lie in the plane spanned by $\{T_j, T_k, T_l\}$, whenever $\{i, j, k, l\} = \{1, 2, 3, 4\}$. Thus, the A-points are located where the line a cuts the faces of the T tetrahedron. We have $\#(B_{3,1}) = 16$, since there are three degrees of freedom in each vertex T_i of the tetrahedron and an additional four in the choice of the line a.

Note that the complete quadrilateral of Figure 4.2, which represents the matroid $B_{2,1}$, can be described in a analogous way: The three P-points are the vertices of a triangle and the three A-points are located where some line a cuts the sides of that triangle. This pattern extends to the matroid $B_{m,1}$ for any m. The $m + 1$ points in the first column are the vertices of an m-simplex, and the points in the second column are located where some line cuts the $m + 1$ facets of that simplex.

[2] The ninth is implied by the other eight over any commutative field. One can also construct projective planes over division rings of scalars that are not commutative, and Pappus's Theorem does not hold in such planes. In this monograph, however, we always require that the field of scalars be commutative.

4.3.2 The partition $4 = 2 + 2$

A representation of the budget matroid $B_{2,2}$ consists of eight points in 3-space,

$$\begin{array}{cc} 2 & 2 \\ \end{array}$$
$$\begin{pmatrix} A_1 & B_1 \\ A_2 & B_2 \\ A_3 & B_3 \\ A_4 & B_4 \end{pmatrix},$$

with no three collinear and with precisely the following eight sets of four coplanar — the two columns and the six perfect sets:

$$\begin{array}{c} A_1, A_2, A_3, A_4 \\ B_1, B_2, B_3, B_4 \\[4pt] A_1, A_2, B_3, B_4 \\ A_1, B_2, A_3, B_4 \\ A_1, B_2, B_3, A_4 \\ B_1, A_2, A_3, B_4 \\ B_1, A_2, B_3, A_4 \\ B_1, B_2, A_3, A_4 \end{array}$$

A simple way to describe those eight sets is to note that each set has one point from each row of the matrix and an even number of points from each column. Looked at in this way, the eight sets fall into three classes, since the three possible column counts are $(4, 0)$, $(0, 4)$, and $(2, 2)$. But we can reduce the number of classes from three to two if we relabel the eight points with one pair swapped, as follows:

$$\begin{array}{ll} C_1 := A_1 & D_1 := B_1 \\ C_2 := A_2 & D_2 := B_2 \\ C_3 := A_3 & D_3 := B_3 \\ C_4 := B_4 & D_4 := A_4 \end{array}$$

In terms of the C and D labels, the coplanar sets are those that consist of one point from each row and an odd number from each column. So the possible column counts are $(3, 1)$ and $(1, 3)$ — that is, the eight sets have the form $\{C_i, C_j, C_k, D_l\}$ and $\{D_i, D_j, D_k, C_l\}$, for $\{i, j, k, l\} = \{1, 2, 3, 4\}$.

Thus, we have a C-tetrahedron and a D-tetrahedron with the property that, for each l in $[1 \,..\, 4]$, the vertex D_l lies on the face plane of the C-tetrahedron that is opposite the vertex C_l, and vice versa. Two tetrahedra positioned in this way are called a *Möbius pair*, that configuration being one of the ones studied in classical projective geometry [21]. The Möbius pair matroid $B_{2,2}$ is our third example, along with the complete quadrilateral $B_{2,1}$ and the Pappus configuration $B_{1,1,1}$, of a budget matroid whose representations are already famous on other grounds.

The eight coplanarities in a Möbius pair have a redundancy of much the same sort as the nine collinearities in a Pappus configuration: If any seven of the eight coplanarities hold, the eighth must hold as well. (We review the proof of this standard result in Section 6.1.) Because of this redundancy, it costs only seven degrees of freedom to make all eight coplanarities hold, so we have $\#(B_{2,2}) = 24-7 = 17$.

Exercise 4.3-1 Construct an explicit Möbius pair in Euclidean 3-space.

[One answer: Let the four A-points be the vertices of an equilateral triangle, together with its centroid, lying in some plane α. Note that, whenever $\{i, j, k, l\} = \{1, 2, 3, 4\}$, the lines $A_i A_j$ and $A_k A_l$ in α are orthogonal. Let the B-points form a congruent figure, located in a plane β parallel to α, but rotated by ninety degrees. The lines $A_i A_j$ and $B_k B_l$ will then be parallel, so the perfect set $\{A_i, A_j, B_k, B_l\}$ will be coplanar.]

Exercise 4.3-2 What automorphisms are there of the budget matroid $B_{b_1,\ldots,b_k}$? For simplicity, let's restrict ourselves to cases in which all of the column budgets are positive.

An automorphism must permute the bk elements of the ground matrix in such a way that the dependency of sets is preserved. We can always permute the rows of the ground matrix as wholes arbitrarily, so every budget matroid has $b!$ automorphisms. In addition, we can clearly interchange any two columns that have the same budgets. Call an automorphism *trivial* when it simply combines some permutation of the rows with some budget-preserving permutation of the columns. (If we were allowing zero column budgets, a trivial automorphism could also permute arbitrarily the elements of any column whose budget is zero.)

Show that each of the four budget matroids $B_{1,1}$, $B_{2,1}$, $B_{1,1,1}$, and $B_{2,2}$ has a nontrivial automorphism.

[Hint: For $B_{1,1}$, swap one diagonal; for $B_{2,1}$, swap two of the three rows; for $B_{1,1,1}$, cyclically rotate the i^{th} row i places; for $B_{2,2}$, swap two of the four rows.]

Exercise 4.3-3 Let $B_{b_1,\ldots,b_k}$ be a budget matroid in which all of the column budgets are positive and in which at most one of them exceeds $b-3$; in particular, let's say that $b_1 \geq \cdots \geq b_k \geq 1$ and that $b_2 \leq b-3$. Note that this condition rules out precisely the four budget partitions discussed in the previous exercise. Show that all of the automorphisms of the matroid $B_{b_1,\ldots,b_k}$ are trivial.

[Hint: Let's call a flat of the budget matroid $B_{b_1,\ldots,b_k}$ a *full flat* when it has rank less than b and cardinality precisely b. Each perfect set is a full flat of rank $b-1$ — that is, a full hyperplane. In addition, the b points in the j^{th} column form a full flat of rank $b_j + 1$. Those are the only full flats. The full flat formed by the first column may be a hyperplane; but for any j in $[2 \,..\, k]$, we have $b_j + 1 < b - 1$, so the full flat formed by the j^{th} column is not a hyperplane. Hence, the elements of the j^{th} column must be mapped, as a set, to the elements of the j'^{th} column, for some j' with $b_j = b_{j'}$. Since every column except the first is mapped, as a set, to some column with the same budget, that property must hold for the first column also, by the process of elimination.

Now, consider some element e_{ij}. In any column other than the j^{th}, the element of that column in the i^{th} row is uniquely distinguished by the property that there are no perfect sets — and hence no full hyperplanes — that contain both that element and e_{ij}. So the automorphism must transform the rows as wholes.]

4.3.3 The partition $4 = 2 + 1 + 1$

The last two budget matroids of rank 4 are the ones that we shall spend much of this monograph studying.

A representation of the matroid $B_{2,1,1}$ consists of twelve points in 3-space, for which we shall use the column letters P, A, and B:

$$\begin{array}{c} \begin{array}{ccc} 2 & 1 & 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 & B_1 \\ P_2 & A_2 & B_2 \\ P_3 & A_3 & B_3 \\ P_4 & A_4 & B_4 \end{pmatrix} \end{array}$$

The four P-points are coplanar; the four A-points are collinear, as are the four B-points; and the twelve perfect sets, of the form $\{P_i, P_j, A_k, B_l\}$ for $\{i, j, k, l\} = \{1, 2, 3, 4\}$, are coplanar.

There are 36 degrees of freedom involved in choosing twelve arbitrary points in 3-space. Each perfect coplanarity costs 1 degree of freedom. The coplanarity of the four P-points also costs 1, while the collinearities of the A-points and the B-points cost 4 each. Thus, if the perfect constraints and the column constraints were all nonredundant, we would have $\#(B_{2,1,1}) = 36 - 12 - 1 - 4 - 4 = 15$. In fact, we shall see later — by two different constructions — that $\#(B_{2,1,1}) = 19$. Thus, there are four dimensions' worth of redundancy among the constraints, up from the one dimension's worth in the Pappus and Möbius-pair cases.

4.3.4 The partition $4 = 1 + 1 + 1 + 1$

As we mentioned in Section 4.1, a representation of the budget matroid $B_{1,1,1,1}$ consists of sixteen points in 3-space:

$$\begin{array}{c} \begin{array}{cccc} 1 & 1 & 1 & 1 \end{array} \\ \begin{pmatrix} A_1 & B_1 & C_1 & D_1 \\ A_2 & B_2 & C_2 & D_2 \\ A_3 & B_3 & C_3 & D_3 \\ A_4 & B_4 & C_4 & D_4 \end{pmatrix}. \end{array}$$

The four points in each column are collinear, while the 24 perfect sets are coplanar.

Choosing sixteen arbitrary points in 3-space involves 48 degrees of freedom. Naively, we would expect the coplanarity of each perfect set to cost 1 and the collinearity of each column to cost 4, leaving only $48 - 24 - 16 = 8$ degrees of freedom. But, in fact, we shall find that $\#(B_{1,1,1,1}) = 20$. Thus, the redundancy among the constraints amounts to a full twelve degrees of freedom — the highest that we've seen so far. This high level of redundancy makes a representation of $B_{1,1,1,1}$ a tightly interwoven geometric structure.

4.4 Budget minors of budget matroids

There is a natural sense in which a big budget matroid has lots of little budget matroids sitting inside it — as minors. If one budget matroid is a minor of another, we'll call the former a *budget minor*. We'll state two simple results about budget minors in this section, though we shan't prove them until Section 5.6, where we can prove slightly more general versions with no more work.

Proposition 4.4-1 *Eliminating a zero part from the budget partition gives a budget minor. That is, if $M := B_{b_1,\ldots,b_k}$ is any budget matroid and we let M' denote the matroid $M' := B_{b_1,\ldots,b_k,0}$, then M is a minor of M'.*

Exercise 4.2-2 provides an example of this: The complete quadrilateral $B_{2,1}$ is a budget minor of the matroid $B_{2,1,0,\ldots,0}$. To get from the latter to the former, we simply delete all of the elements in the columns with zero budgets.

Proposition 4.4-2 *Reducing the parts of the budget partition also gives a budget minor, as long as at least two parts remain positive. That is, if $M := B_{b_1,\ldots,b_k}$ and $M' := B_{b'_1,\ldots,b'_k}$ are two budget matroids with $b_j \le b'_j$ for all j in $[1 \,..\, k]$, then M is a minor of M'.*

This is more interesting. For an example, let's check that the complete quadrilateral $B_{2,1}$ is a budget minor of the matroid $B_{3,1}$. Consider some representation

$$\begin{array}{cc} 3 & 1 \\ \end{array}$$
$$\begin{pmatrix} T_1 & A_1 \\ T_2 & A_2 \\ T_3 & A_3 \\ T_4 & A_4 \end{pmatrix}.$$

of $B_{3,1}$, sitting in 3-space. Suppose that we put our eye at some vertex T_i of the tetrahedron, and we look out at the six points that are not in the i^{th} row. The lines through our eye form the 'points' of a projective plane. The face of the tetrahedron opposite the vertex T_i projects down into a triangle in that plane, while the line a projects down into a line. Since the points $\{T_i, T_j, T_k, A_l\}$ are coplanar, for any

$\{i, j, k, l\} = \{1, 2, 3, 4\}$, the lines T_iT_j, T_iT_k, and T_iA_l are also coplanar — which means that, in the plane of lines through T_i, the projected images of the points T_j, T_k, and A_l are collinear. So, what we see is a complete quadrilateral.

Thinking about the matroid $B_{3,1}$ itself, rather than a representation of it, what we have done is to contract the i^{th} element of the T column, while deleting the other element of the i^{th} row. More generally, in the situation of Proposition 4.4-2, we contract some $b'_j - b_j$ elements of the j^{th} column, for each j, chosen so that no two contracted elements come from the same row, and we delete all of the other elements of any row that contains a contracted element.

By the way, when working with budget minors, it may seem unfortunate that we have required at least two of the column budgets of a budget matroid to be positive. For example, what do we see if we put our eye at the point A_4, in the representation of $B_{3,1}$ above? Answer: The points T_1, T_2, and T_3 appear to be collinear, while the points A_1, A_2, and A_3 appear to coincide, but don't appear to lie on the line through the three T's. This is a representation of a perfectly respectable matroid of rank 3, which it would be quite natural to denote $B_{3,0}$. Let's say that an *all-or-nothing matroid* is the matroid $B_{b_1,\ldots,b_k}$ that results from a budget partition in which precisely one part is positive, while all of the others are zero. We could have defined the class of budget matroids to include the all-or-nothing matroids. But doing so costs quite a bit in complexity, as the following exercise suggests — enough that the costs seem to outweigh the benefits.

Exercise 4.4-3 As samples of the complexities associated with the all-or-nothing matroids, show that:

- For $m \geq 1$, the matroid B_m has rank $m - 1$, rather than m.
- For $m \geq 2$, the matroid $B_{m,0}$ has rank m, all right, but it has $\#(B_{m,0}) = m^2-2$ degrees of freedom, as opposed to the $m^2 - 3$ predicted by the formula in Theorem 6.2-2, the $B_{m,n}$ Representation Theorem.

4.5 Projective configurations in the narrow sense

Lots of the budget matroids are representable, and their representations are intriguing geometric structures. Since the study of projective configurations was once quite a popular field, it is a bit surprising that the budget matroids weren't well studied long ago. One reason why the classical geometers may have overlooked the budget matroids is that they defined the term 'configuration' more narrowly than we. As we discuss in this section, they required a configuration to have a high level of symmetry, at least in a certain numeric sense. The budget matroids don't have that symmetry, and that may explain why they were overlooked.

Actually, whether a configuration does or does not have a certain symmetry can depend upon how we interpret that configuration. The budget matroids don't have the required numeric symmetry when they are interpreted in the obvious way. But Jorge Stolfi came up with a devious way to reinterpret a budget matroid, thereby converting it into a structure that does meet the classical definition of a projective configuration. We'll discuss Stolfi's trick in this section also. Keep in mind that the Stolfi Trick is mostly of historical and nomenclatural interest. Applying the Stolfi Trick to a budget matroid obfuscates its structure, and that generally isn't helpful. But it is worth noting that the classical demand for numeric symmetry did not slam the door completely on the budget matroids; there was still a tiny crack through which they could have slipped.

4.5.1 The required numeric symmetry

What is a 'configuration', in the narrow sense? Classically, an arrangement in the plane involving p points and q lines is called a *configuration* of type (p_r, q_s) if each of the p points lies on r of the lines and each of the q lines passes through s of the points [20]. Note that the four parameters p, q, r, and s must satisfy $pr = qs$, since both of those products count the total incidences between points and lines. For example, a complete quadrilateral has six points, each lying on two lines, and has four lines, each passing through three points; so it is a configuration of type $(6_2, 4_3)$ — and we have $6 \cdot 2 = 4 \cdot 3 = 12$.

In a space of any dimension, we can speak of a configuration of points and hyperplanes of type (p_r, q_s). Consider 3-space, for example. A Möbius pair of tetrahedra is a configuration of points and planes of type $(8_4, 8_4)$: eight points, each lying on four planes, and eight planes, each passing through four points. Classical geometry also talks about configurations of flats other than points and hyperplanes, such as configurations of points and lines in 3-space. But a configuration always involves flats of only two different dimensions, say small and large, with every small flat lying on the same number of large flats and every large flat passing through the same number of small flats. There is no such thing as a configuration of points, lines, and planes in 3-space, for example — at least, not in this classical sense of the word 'configuration'.

Unfortunately, the budget matroid $B_{b_1,\ldots,b_k}$ involves noteworthy flats of various dimensions. Each perfect set is mutually incident, its b points lying on a common hyperplane, of dimension $b-2$; there are $\binom{b}{b_1 \,\ldots\, b_k}$ such noteworthy flats. But the b points in the j^{th} column also lie on a common flat, which is also noteworthy; and its dimension is b_j, which is typically less than $b-2$. That suggests that representations of budget matroids typically are not configurations in the narrow sense.

4.5.2 The Stolfi Trick

But Jorge Stolfi realized that a little trickery can get around this problem. We can encode a column flat of dimension b_j by using some number n_j of hyperplanes. These n_j hyperplanes are in general position except that they all include the column flat, and there are enough of them so that the intersection of all of them is precisely the column flat.

What constraints are there on the numbers (n_j)? Each hyperplane reduces the dimension of their intersection by at most 1, so we must have $n_j \geq b - b_j - 1$, in order to reduce the dimension all the way from $b-1$ down to b_j. On the other hand, there is typically no upper limit on n_j. But there are upper limits in two special cases. If $b_j = b - 1$, so that the b points in the j^{th} column span the entire ambient $(b-1)$-space, we must have not only $n_j \geq 0$, but actually $n_j = 0$. Similarly, if $b_j = b - 2$, so that the points in the j^{th} column lie in a common hyperplane, we must have not only $n_j \geq 1$, but actually $n_j = 1$ — since we don't want two hyperplanes in our projective configuration to be required to coincide.

By the way, just as we don't want two hyperplanes to be required to coincide, we also don't want two points to be required to coincide; such a thing would surely have troubled the classical projective geometers. Hence, we shall restrict ourselves to budget partitions $(b_1, \ldots, b_k)$ in which all of column budgets are positive.

How close are we to having a configuration of points and hyperplanes? The hyperplane corresponding to each perfect set contains precisely b points. Each of the n_j hyperplanes that we are using to encode the j^{th} column flat also contains precisely b points: the b points in the j^{th} column. So we are in pretty good shape. All that remains is to arrange that every point lies on the same number of hyperplanes, and we can do that by choosing the counts (n_j) appropriately.

In fact, we can typically choose the counts (n_j) in an infinite variety of ways and still do that. Consider the budget matroid $B_{5,3,1}$, for example — let's optimistically assume that this matroid is representable, sitting in 8-space. Applying the Stolfi Trick, each point in the first column lies in the n_1 hyperplanes that encode the first column's 5-flat, as well as in $\binom{8}{4\,3\,1} = 280$ of the hyperplanes that correspond to perfect sets. Each point in the second column lies in n_2 column hyperplanes, plus $\binom{8}{5\,2\,1} = 168$ of the perfect hyperplanes. And each point in the third column lies in n_3 column hyperplanes, plus $\binom{8}{5\,3\,0} = 56$ perfect hyperplanes. To achieve the required numeric symmetry, it suffices to arrange that $n_1+280 = n_2+168 = n_3+56$, subject to the side conditions $n_1 \geq 3$, $n_2 \geq 5$, and $n_3 \geq 7$. The solution is to set $n_1 := m + 3$, $n_2 := m + 115$, and $n_3 = m + 227$, for any nonnegative integer m. Thus, we can view a representation of $B_{5,3,1}$ as a configuration of points and hyperplanes in 8-space of type $(27_{m+283}, (3m + 849)_9)$, for any nonnegative m.

For a budget matroid $B_{b_1,\ldots,b_k}$ in which no column budget b_j exceeds $b-3$, this process cannot fail, since the side conditions put only lower bounds on the counts (n_j). But if $b_j \geq b - 2$ for some j, we have to analyze more carefully, because some of the side conditions are then equalities. The cases that require extra care are all contained in the three families $B_{b-1,1}$, $B_{b-2,2}$, and $B_{b-2,1,1}$, and those families are dealt with in the first three exercises below. So we conclude the following.

Proposition 4.5-1 (Stolfi Trick) *Let $b = b_1 + \cdots + b_k$ be any partition of the integer b into at least two positive parts. There exists at least one sequence $(n_1, \ldots, n_k)$*

of nonnegative counts with the property that a representation of the budget matroid $B_{b_1,\ldots,b_k}$ in $(b-1)$-space can be interpreted as a projective configuration of points and hyperplanes in the narrow sense by the trick of encoding its j^{th} column flat as the intersection of n_j otherwise unconstrained hyperplanes.

Exercise 4.5-2 For any $b \geq 3$, show that there are unique column counts n_1 and n_2 that make the Stolfi Trick work for the budget matroid $B_{b-1,1}$. The case $b = 2$ works also, in a sense, but it is a delicate matter to discuss a configuration of points and hyperplanes in the line, where hyperplanes are the same things as points.

[Answer: The counts n_1 and n_2 must satisfy the relation $n_1 + b - 1 = n_2 + 1$ and the side conditions $n_1 = 0$ and $n_2 \geq b - 2$. In addition, when $b \leq 3$, the side condition $n_2 \geq b - 2$ tightens to the equality $n_2 = b - 2$. But n_2 works out to be $n_2 = b - 2$ in any case.]

Exercise 4.5-3 For any $b \geq 3$, show that there are unique column counts n_1 and n_2 that make the Stolfi Trick work for the budget matroid $B_{b-2,2}$.

[Answer: We must satisfy the relation $n_1 + (b-1)(b-2)/2 = n_2 + b - 1$ and the side conditions $n_1 = 1$ and $n_2 \geq b - 3$. In addition, when $b \leq 4$, the side condition $n_2 \geq b - 3$ tightens to the equality $n_2 = b - 3$. Solving, we find that $n_2 = (b-2)(b-3)/2$, which is at least $b - 3$ for all integers b and equals $b - 3$ for $b = 3$ or $b = 4$.]

Exercise 4.5-4 For any $b \geq 3$, show that there are unique column counts n_1, n_2, and n_3 that make the Stolfi Trick work for the budget matroid $B_{b-2,1,1}$.

[Answer: We must satisfy the relation $n_1 + (b-1)(b-2) = n_2 + b - 1 = n_3 + b - 1$ and the side conditions $n_1 = 1$, $n_2 \geq b-2$, and $n_3 \geq b-2$. In addition, when $b = 3$, the side conditions $n_2 \geq 1$ and $n_3 \geq 1$ tighten to the equalities $n_2 = n_3 = 1$. Solving, we find that $n_2 = n_3 = (b-2)^2$, which is at least $b-2$ for all integers b and equals $b - 2 = 1$ when $b = 3$.]

Exercise 4.5-5 For example, suppose that we use the Stolfi Trick to interpret a representation of the budget matroid $B_{2,1,1}$ as a configuration in the narrow sense. What do we get?

[Answer: a configuration of points and planes in 3-space of type $(12_7, 21_4)$. We have $n_1 = 1$ and $n_2 = n_3 = 4$, so each of the two column lines is encoded as the intersection of four planes.]

Chapter 5

The budgetary matroids

In the budget matroids, the j^{th} column budget b_j is used for two purposes: as the dimension of the common flat in which all of the points of the j^{th} column must lie and as the number of points from the j^{th} column that any perfect set must contain. Those two parameters of the j^{th} column don't have to be the same; but the cases in which they are the same seem to have the most intriguing geometric properties.

We shall reserve the name 'budget matroid' for those cases in which those two parameters of each column are the same. We shall use the clumsy term 'budgetary matroid' for the more general situation in which those two parameters of each column are fairly independent. The budget matroids are an interesting subclass of the budgetary matroids — interesting because so many of them are representable.

In this chapter, we generalize from the budget matroids to the budgetary matroids, but we then argue, by looking at a slew of examples, that the budgetary matroids are pretty boring. If you get really bored, remember that our study of the budget matroids in general picks up again in Chapter 6, while Chapter 8 is where we prepare for the study of the budget matroid $B_{2,1,1}$ in particular.

5.1 The parameters of a budgetary matroid

To define a budgetary matroid, we need both some partition $b = b_1 + \cdots + b_k$ of a total budget b into k column budgets $(b_1, \ldots, b_k)$ and also a separate sequence of integers $(d_1, \ldots, d_k)$. A representation of the budgetary matroid $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ is going to involve bk points:

$$\begin{array}{c}
\begin{array}{cccc} {\scriptstyle d_1} & {\scriptstyle d_2} & & {\scriptstyle d_k} \\ {\scriptstyle b_1} & {\scriptstyle b_2} & & {\scriptstyle b_k} \end{array} \\
\begin{pmatrix} E_{11} & E_{12} & \ldots & E_{1k} \\ E_{21} & E_{22} & \ldots & E_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ E_{b1} & E_{b2} & \ldots & E_{bk} \end{pmatrix}
\end{array}$$

The column budget b_j is the number of points from the j^{th} column that any perfect set must contain; so the number of rows in the matrix of points is still the total budget b and the number of perfect sets is still $\binom{b}{b_1 \ldots b_k}$. The new parameter d_j, which we shall call the j^{th} *column dimension*, is the dimension of the span of the b points in the j^{th} column. A budget matroid is a budgetary matroid in which the dimensions equal the budgets, so $d_j = b_j$ for all j.

What constraints should we place on the parameters (d_j) and (b_j), the dimensions and the budgets? Let's study that question starting with a blank slate. In particular, let's try to justify why, when defining the budget matroids, we required at least two of the column budgets to be positive.

First, if the total budget b were zero, then the empty set would be perfect and hence dependent, which would violate the Empty-Set Axiom. We want to end up with a matroid, so we require that b be positive, which implies that k must be positive as well.

For any j in $[1\,..\,k]$, the b points in the j^{th} column of the matrix can't possibly span a flat of dimension greater than $b - 1$. We want the parameter d_j to measure the dimension of that span; so we require that $d_j < b$, for all j.

In the opposite direction, how small a value of d_j is permissible? If $d_j < b_j - 1$ for any j, the whole notion of a perfect set becomes irrelevant. In order to be perfect, a set has to contain b_j elements from the j^{th} column. But any set that contains more than $d_j + 1$ elements from a flat of dimension d_j is already dependent, on those grounds alone — so, when $d_j + 1 < b_j$, the constraint on perfect sets never comes into play. For example, suppose that we start with the complete quadrilateral

$$\begin{array}{c} \begin{array}{cc} 2 & 1 \\ 2 & 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix} \end{array},$$

which is the budget matroid $B^{2,1}_{2,1}$. (Note that each column is now headed with both its dimension and its budget.) What would happen if we reduced the dimension d_1 for the first column from 2 to 0? The resulting structure, if it were legal, would be called $B^{0,1}_{2,1}$. In a representation

$$\begin{array}{c} \begin{array}{cc} 0 & 1 \\ 2 & 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix} \end{array}$$

of this structure, the three points in the first column would have to lie in a common 0-flat, that is, would have to coincide: $P_1 = P_2 = P_3$. Once that was true, a perfect set such as $\{P_1, P_2, A_3\}$ and a non-perfect set such as $\{P_1, P_2, A_1\}$ would both be dependent, and for the same reason, so the essential character of the budgetary matroids would be lost. We therefore require that $d_j \geq b_j - 1$, for all j.

We now pause to consider a special case. Suppose that all of the column budgets other than the j^{th} are zero, so $b_j = b$; we called this the *all-or-nothing case*, in Section 4.4. The constraint $d_j \geq b_j - 1$ then implies that $d_j \geq b - 1$. But, when $b_j = b$, the b elements in the j^{th} column form a perfect set — indeed, they form the unique perfect set. They are hence dependent, so the most that they can do is to span a flat of dimension $b - 2$. We want d_j to measure the dimension of that span, which contradicts the inequality $d_j \geq b - 1$. We respond to this contradiction by outlawing the all-or-nothing case; we require that $b_j < b$ for all j, which means that $k \geq 2$ and that there are at least two columns with positive budgets.

Returning to the main thread, we are already requiring that $d_j \geq b_j - 1$, for all j; but if $d_j = b_j - 1$ for any j, the resulting structure typically isn't a matroid. Since we want matroids, we shall actually require the stronger inequality $d_j \geq b_j$. For an example of what goes wrong if we don't, suppose that we tried to reduce the dimension d_1 in the complete quadrilateral from 2 to 1, getting the structure $B^{1,1}_{2,1}$:

$$\begin{array}{cc} 1 & 1 \\ 2 & 1 \end{array}$$
$$\begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix}$$

The only effect of reducing d_1 from 2 to 1 is to make the three P-points be collinear, that is, to make the set $\{P_1, P_2, P_3\}$ dependent. The sets $X := \{P_1, P_2\}$ and $Y := \{P_2, P_3, A_3\}$ are independent in the original matroid $B^{2,1}_{2,1}$, and they remain independent when d_1 is reduced, since neither X nor Y contains all three P-points. In order for the Augmentation Axiom to hold, there must exist some element of Y that we can add to X without destroying independence. But there is no such element: The element P_2 is in X already, the set $X \cup \{P_3\} = \{P_1, P_2, P_3\}$ is dependent because $d_1 = 1$, and the set $X \cup \{A_3\} = \{P_1, P_2, A_3\}$ is dependent because it is perfect. Analogous examples of sets X and Y can be constructed in almost all cases in which $d_j = b_j - 1$ for some j, as shown in Exercise 5.1-1.

In summary, we henceforth require the following: There must be at least two columns with positive budgets and, for all j, the column dimension d_j must lie in the left-closed, right-open interval $[b_j \,..\, b)$.

Exercise 5.1-1 Given column dimensions (d_m) and budgets (b_m) with $d_m \geq b_m - 1$ for all m but with $d_j = b_j - 1$ and $0 < b_j < b$ for some j, show that the resulting structure $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ is not a matroid. In particular, find sets X and Y that violate the Augmentation Axiom.

As for the side conditions $0 < b_j < b$, we are already requiring that $b_j < b$. The case $b_j = 0$ and $d_j = -1$ does not cause a violation of the matroid axioms, but we are going to outlaw that case anyway, just to simplify things. If we did make it legal to set d_j to -1 when b_j was 0, doing so would force all of the points in the j^{th} column to be indeterminate — that is, in the resulting matroid, all of those elements would be loops.

[Hint: Let P be any perfect set, let i and i' be row indices with the property that E_{ij} is in P but $E_{i'j}$ is not, and let j' be the unique column index with $E_{i'j'}$ in P. It suffices to take $X := P \setminus \{E_{i'j'}\}$ and $Y := P \cup \{E_{i'j}\} \setminus \{E_{ij}\}$.]

5.2 The definition

The budgetary matroids are just like the budget matroids, except that, in defining independence, the Column Rule talks about the column dimension d_j, rather than the column budget b_j. Let's spell everything out, though, to avoid any confusion.

Let $(b_1, \ldots, b_k)$ be any nonnegative integers of which at least two are positive, let b — the *total budget* — denote the sum $b := b_1 + \cdots + b_k$, and let d_j, for each j in $[1\,..\,k]$, be any integer in the half-open interval $[b_j\,..\,b)$. We hereby define the *budgetary matroid* $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$, which has, for each j, the *budget* b_j and the *dimension* d_j associated with its j^{th} column. The elements of the ground set are the entries of a b-by-k matrix

$$\mathbf{e} := \begin{array}{c} \begin{matrix} d_1 & d_2 & & d_k \\ b_1 & b_2 & & b_k \end{matrix} \\ \begin{pmatrix} e_{11} & e_{12} & \ldots & e_{1k} \\ e_{21} & e_{22} & \ldots & e_{2k} \\ \vdots & \vdots & \vdots & \vdots \\ e_{b1} & e_{b2} & \ldots & e_{bk} \end{pmatrix} \end{array},$$

which we shall refer to as the *ground matrix*. We define a subset X of the ground matrix $\mathbf{e}$ to be *perfect* when X contains precisely one element from each row and, for each j, contains precisely b_j elements from the j^{th} column. A subset X of $\mathbf{e}$ is *independent* when it has the following three properties:

Ambient Rule X contains at most b elements overall;

Column Rule X contains at most $d_j + 1$ of the b elements in the j^{th} column, for each j in $[1\,..\,k]$; and

Perfect Rule X has no perfect subsets.

In the special case in which the budgets and the dimensions are equal — that is, when $b_j = d_j$ for all j in $[1\,..\,k]$ — the budgetary matroid $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ is a budget matroid, and we shall omit the dimensions when we write it: $B_{b_1,\ldots,b_k} = B^{b_1,\ldots,b_k}_{b_1,\ldots,b_k}$.

Let's consider $B^{2,2}_{2,1}$ as a simple example of a budgetary matroid. It is just like the complete quadrilateral $B_{2,1}$, except that the dimension d_2 associated with the second column has been increased from 1 to 2. A representation of the matroid $B^{2,2}_{2,1}$ consists of six points

$$\begin{array}{c} \begin{matrix} 2 & 2 \\ 2 & 1 \end{matrix} \\ \begin{pmatrix} P_1 & A_1 \\ P_2 & A_2 \\ P_3 & A_3 \end{pmatrix} \end{array}$$

in the plane, where the three perfect sets are required to be collinear, but the three A-points are forbidden from being collinear. So we have a triangle $\triangle P_1P_2P_3$, together with three noncollinear A-points, one on each side of the P triangle. Note that relaxing the collinearity constraint on the A-points gives us back one degree of freedom: While $\#\big(B^{2,1}_{2,1}\big) = 8$, we have $\#\big(B^{2,2}_{2,1}\big) = 9$.

Exercise 5.2-1 Given some complete quadrilateral Q, whose six points represent the budget matroid $B_{2,1}$, describe what three points must be added to Q in order to represent each of the budgetary matroids $B^{2,1,1}_{2,1,0}$ and $B^{2,1,2}_{2,1,0}$.

[Answer: In the first case, the three new points must be collinear, but no two may coincide; in the second case, they must not be collinear. In both cases, the three new points must lie in the same plane as the quadrilateral Q, but none of them may lie on the line determined by any two points of Q and the line determined by any two of them must not pass through any point of Q.]

5.3 They really are matroids

It is finally time to verify that the budgetary matroids — and among them, the budget matroids — really are matroids.

Proposition 5.3-1 *Let $b = b_1 + \cdots + b_k$ be a partition of the total budget b into nonnegative, integral column budgets (b_j), at least two of which are positive. For each j in $[1\,..\,k]$, let the column dimension d_j lie in the half-open interval $[b_j\,..\,b)$. For each i in $[1\,..\,b]$, let R_i denote the i^{th} row of the b-by-k ground matrix $\mathbf{e} := (e_{ij})_{i\in[1..b],\, j\in[1..k]}$, and, for each j in $[1\,..\,k]$, let C_j denote its j^{th} column. We say that a subset X of the ground matrix $\mathbf{e}$ is* perfect *when*

- $|X \cap R_i| = 1$ *for all i in $[1\,..\,b]$ and*
- $\big|X \cap C_j\big| = b_j$ *for all j in $[1\,..\,k]$.*

Finally, we say that a subset X of $\mathbf{e}$ is independent *when:*

Ambient Rule $|X| \le b$;

Column Rule $\big|X \cap C_j\big| \le d_j + 1$, *for all j in $[1\,..\,k]$; and*

Perfect Rule *X has no perfect subsets.*

The ground set $\mathbf{e}$, when equipped with this family of independent subsets, forms a matroid, which we denote $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$.

Proof The Empty-Set Axiom is easy: The empty set has zero elements overall, zero elements in each column, and is not perfect, so it satisfies all three rules.

The Subset Axiom is also easy. Deleting an element from an independent set doesn't increase the number of elements overall, doesn't increase the number of elements in any column, and doesn't add any new subsets that might be perfect.

The hard part is verifying the Augmentation Axiom. It is helpful to introduce some nomenclature. Given some set $X \subseteq \mathbf{e}$, we call the i^{th} row of X *empty* when $X \cap R_i = \emptyset$, and we call the j^{th} column of X

- *tight* when $\left|X \cap C_j\right| = d_j + 1$;
- *snug* when $\left|X \cap C_j\right| \in [b_j \mathinner{..} d_j]$;
- *loose* when $\left|X \cap C_j\right| < b_j$; and
- *baggy* when $\left|X \cap C_j\right| < b_j - 1$.

With this nomenclature, the Column Rule says that, for every j, the j^{th} column of X is either tight, snug, or loose. Some of the loose columns may actually be baggy. Note that, if a set X has either some row that is empty or some column that is loose, then X can't have any perfect subsets, so X satisfies the Perfect Rule.

Let X and Y be independent sets with $|Y| > |X|$. We want to find an element e_{ij} in $Y \setminus X$ that can be added to X without destroying independence. We have $|X| < |Y| \leq b$, so adding any single element to X won't cause the augmented X to have too many elements overall, and hence to violate the Ambient Rule. But we must choose the new element with some care, lest the augmented X violate the Column Rule or the Perfect Rule.

We are going to consider five cases, in sequence. At the outset, note that there must be at least one column of X that is loose, since $|X| < b$.

Case 1: Suppose that, for some j, the j^{th} column of X is snug and the j^{th} column of Y contains some element that isn't in X; that is, suppose that $b_j \leq \left|X \cap C_j\right| \leq d_j$ and that $(Y \setminus X) \cap C_j \neq \emptyset$. In this case, we can choose any element in $(Y \setminus X) \cap C_j$ and add that element to X without destroying independence. The j^{th} column of X either remains snug or becomes tight, so the Column Rule still holds. And any column of X that was loose — and there was at least one such — remains loose in the augmented X, so the Perfect Rule also holds.

In what follows, we can assume that the first case does not pertain; that is, for every j for which the j^{th} column of X is snug, we have $(Y \setminus X) \cap C_j = \emptyset$, which implies that $\left|Y \cap C_j\right| \leq \left|X \cap C_j\right|$. So Y is either snug or loose in such columns. On the other hand, if the j^{th} column of X is tight, we have $\left|X \cap C_j\right| = d_j + 1$ and $\left|Y \cap C_j\right| \leq d_j + 1$, the latter because Y is independent. But Y has more elements than X overall. Therefore, for one or more columns j where X is loose, we have $(Y \setminus X) \cap C_j \neq \emptyset$.

Case 2: Suppose that X has two loose columns. We choose a column m of X that is loose and in which $(Y \setminus X) \cap C_m \neq \emptyset$, and we add any element of $(Y \setminus X) \cap C_m$ to X. Adding an element to a loose column certainly doesn't violate the Column

Rule. The Perfect Rule holds because, while we may be moving the m^{th} column from loose to snug, there is some other column of X that was loose and that remains loose.

Case 3: This case is very similar to Case 2. Suppose that X has only one loose column, the m^{th}, but that column is actually baggy. We must have $(Y \setminus X) \cap C_m \neq \emptyset$, so we add any element of $(Y \setminus X) \cap C_m$ to X. The m^{th} column remains loose even after the augmentation, so both the Column and Perfect Rules hold.

In what follows, we can assume that none of the first three cases pertains. Thus, there is precisely one column that is loose in X, say the m^{th}, but this column is not baggy. It follows from this and the constraint $|X| < b$ that, in fact, $|X| = b - 1$, $|Y| = b$, and, for all $j \neq m$, we have $\left|X \cap C_j\right| = b_j$ and $(Y \setminus X) \cap C_j = \emptyset$.

Case 4: Suppose that, in the m^{th} column, Y has at least two elements that X doesn't have; that is, $|(Y \setminus X) \cap C_m| \geq 2$. Since $|X| = b-1$, we know that X has at least one empty row; choose one of them. Because we have at least two elements of $Y \setminus X$ in the m^{th} column to pick from, we can find one to add to X that does not spoil our chosen empty row.

Case 5: In the remaining case, we can assume that Case 4 also fails to pertain, which means that $|(Y \setminus X) \cap C_m| = 1$. Since $\left|(Y \setminus X) \cap C_j\right| = 0$ for all j different from m, we deduce that $|Y \setminus X| = 1$. Since we also have $|Y| - |X| = 1$, it follows that $X \subset Y$. We can hence take the unique element of $Y \setminus X$ and add it to X without destroying independence. □

5.4 Weakening the rules for independence

Given column budgets (b_j) and dimensions (d_j), the budgetary matroid $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ arises out of the interplay among the three rules that define the notion of independence: the Ambient Rule, the Column Rule, and the Perfect Rule. In this section, we attempt to justify those three rules by considering what would happen if we omitted one of the three.

What would happen if we kept the Ambient and Perfect Rules, but dropped the Column Rule? Then, every set with less than b elements would be independent, every set with more than b elements would be dependent, and, among the sets with b elements, precisely the perfect sets would be dependent. Note that, if a set X has $b - 1$ elements, there is at most one perfect set that includes X. From this alone, it is straightforward to check that the resulting structure would be a matroid of a standard type, called a *paving matroid* [39].

If we kept the Ambient and Column Rules, but dropped the Perfect Rule, we would again get a matroid of a standard type: a *truncation* of a *direct sum* of *uniform matroids* [40].

So what about omitting the Ambient Rule? Since the Ambient Rule was arguably the least well motivated of the three rules to begin with, it is particularly interesting to study what happens when we omit that rule — or perhaps just weaken it somewhat. We'll call a weakened version of the Ambient Rule *successful* when the structure that results from the Column Rule, the Perfect Rule, and that weakened Ambient Rule is still a matroid.

If we weaken the Ambient Rule too far, we lose all hope of being successful. In particular, whenever the total budget b is at least 3, there are some sets of size $b+1$ that satisfy both the Column Rule and the Perfect Rule but that must be classified as dependent, in order for the resulting structure to be a matroid. To construct such a set, let P be any perfect set in the structure $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ with $b \geq 3$, and let f, g, and h be any three elements of P with f and g in different columns. Let f' be the element that is in the same row as f and the same column as g, and let g' be the fourth vertex of that rectangle. Consider applying the Augmentation Axiom to the sets $X := P \cup \{f'\} \setminus \{g\}$ and $Y := P \cup \{f', g'\} \setminus \{h\}$. Here are pictures of what the sets P, X, and Y might look like, with the budget partition $(b_1, b_2) = (2, 1)$:

$$\begin{array}{c} \begin{matrix} 2 & \;1 \end{matrix} \\ \begin{pmatrix} f & f' \\ g' & g \\ h & \text{-} \end{pmatrix} \end{array} : \qquad \begin{array}{c} \begin{matrix} 2 & \;1 \end{matrix} \\ \begin{pmatrix} P & \text{-} \\ \text{-} & P \\ P & \text{-} \end{pmatrix} \end{array}, \quad \begin{array}{c} \begin{matrix} 2 & \;1 \end{matrix} \\ \begin{pmatrix} X & X \\ \text{-} & \text{-} \\ X & \text{-} \end{pmatrix} \end{array}, \quad \begin{array}{c} \begin{matrix} 2 & \;1 \end{matrix} \\ \begin{pmatrix} Y & Y \\ Y & Y \\ \text{-} & \text{-} \end{pmatrix} \end{array}.$$

We have $|X| = b$ and $|Y| = b + 1$. Neither X nor Y has any perfect subsets, since each has an empty row. And both X and Y satisfy the Column Rule, since $d_j \geq b_j$ for all j. So, if we weaken the Ambient Rule far enough, the set Y will be independent, as well as X. But adding either of the elements of $Y \setminus X = \{g, g'\}$ to X results in a set that has a perfect subset, and hence violates the Perfect Rule. If we want to satisfy the Augmentation Axiom, we must preserve enough of the Ambient Rule to classify the set Y as dependent.

When the parameters (b_j) and (d_j) satisfy $b_j = d_j > 0$ for all j — that is, for those budgetary matroids that are budget matroids and all of whose column budgets are positive, which are the cases of primary interest — much more is true: There is no way to weaken the Ambient Rule at all and still be successful. To see this, recall from Exercise 4.2-4 that the circuits (the minimal dependent sets) of the budget matroid $B_{b_1,\ldots,b_k}$ are of three types:

perfect circuits the perfect sets,

column circuits the subsets of column j of size $b_j + 2$, and

ambient circuits those sets of size $b + 1$ that do not include, as a subset, any perfect circuit or column circuit.

Let's refer to the perfect and column circuits as the *key circuits*, since they will remain circuits, even if the Ambient Rule is somehow weakened. Suppose that M is some other matroid on the same ground set as the budget matroid $B_{b_1,\ldots,b_k}$, and suppose that all of the key circuits of $B_{b_1,\ldots,b_k}$ are circuits also in M. As the next exercise shows, it follows from those assumptions alone that the rank of M

is at most b. Hence, in these cases, weakening the Ambient Rule however slightly would mean that we would no longer have a matroid.

Exercise 5.4-1 Show that the Ambient Rule can't be successfully weakened when $b_j = d_j > 0$ for all j. In particular, let M be any matroid whose ground set is the ground matrix of the budget matroid $B_{b_1,\ldots,b_k}$. Show that, if every key circuit of the matroid $B_{b_1,\ldots,b_k}$ is a circuit also in M, then the rank of M is at most the total budget $b = b_1 + \cdots + b_k$.

[Hint: Let P be any perfect set, let j be any column index, let f be an element of the j^{th} column that belongs to P, let g be an element of the j^{th} column that does not belong to P, and let X denote the set $X := P \cup \{g\} \setminus \{f\}$. Because $|X| = b$, it suffices to show that all of the elements of the ground matrix $\mathbf{e}$ lie in the flat $\mathrm{Span}(X)$, where that span is computed in the matroid M. Since the perfect set P is a circuit in M, the element f depends on the elements in the set $P \setminus \{f\}$, which is a subset of X; hence, f lies in $\mathrm{Span}(X)$. The set $X \cup \{f\} = P \cup \{g\}$ exceeds the dimension for the j^{th} column by precisely 1; it follows that every element of the j^{th} column depends on some subset of $X \cup \{f\}$, and hence lies in $\mathrm{Span}(X)$. Show next, by considering new perfect sets, that every element in the same row as f belongs to $\mathrm{Span}(X)$. Conclude that, in fact, all of the elements of the ground matrix $\mathbf{e}$ must belong to $\mathrm{Span}(X)$.]

Exercise 5.4-2 Show that the Ambient Rule can be successfully weakened, even when $b_j = d_j$ for all j, if at least one of the column budgets is zero. In particular, after reviewing Exercise 4.2-2, consider those matroids M on the same ground set as the budget matroid $B_{2,1,0,0,\ldots,0}$ and with the property that every key circuit of $B_{2,1,0,0,\ldots,0}$ is a circuit also in M. Find such a matroid M whose rank is 1003.

[Hint: Put the points Z_3 through Z_{1002} in general position with respect to the plane of the complete quadrilateral.]

Exercise 5.4-3 Show that the Ambient Rule can sometimes be successfully weakened, even when b_j is positive for all j, if the dimensions (d_j) aren't equal to the budgets (b_j). In particular, consider the budgetary matroid $B^{2,2,2}_{1,1,1}$, whose key circuits are precisely the six perfect sets. Find a matroid M of rank 4 on the same 3-by-3 ground matrix

$$\begin{array}{ccc} 2 & 2 & 2 \\ 1 & 1 & 1 \end{array}$$
$$\begin{pmatrix} A_1 & B_1 & C_1 \\ A_2 & B_2 & C_2 \\ A_3 & B_3 & C_3 \end{pmatrix}$$

as $B^{2,2,2}_{1,1,1}$ in which all six of those perfect sets are circuits.

[Hint: Let the circuits of M be precisely the perfect sets, the 2-by-2 submatrices, and the unions of a row and a column. To check that M is a matroid, we can either verify the circuit-based axioms for a matroid or — probably easier —

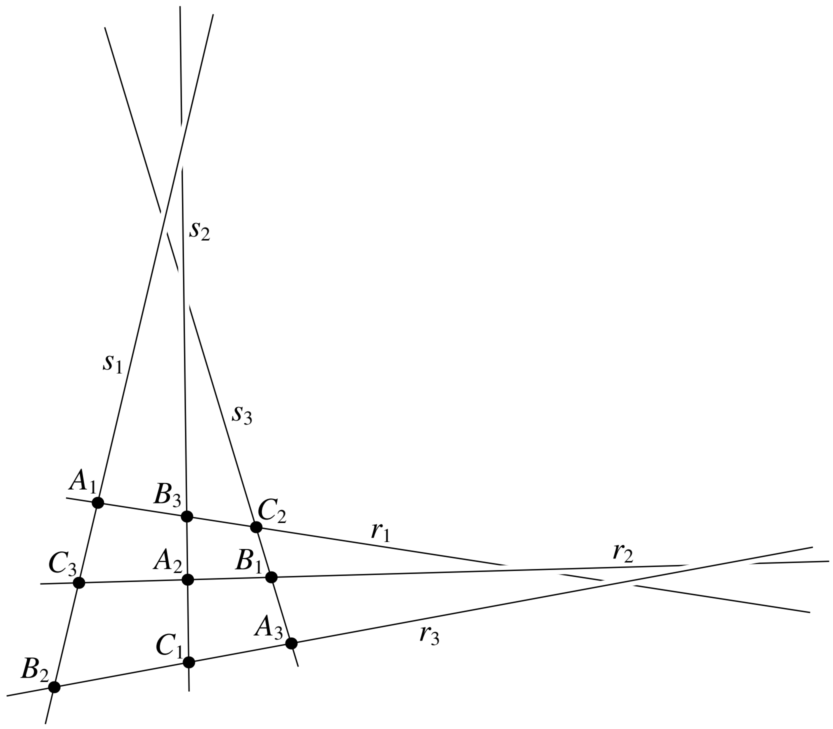

Figure 5.1: A representation in 3-space of a matroid of rank 4 that has, as circuits, all six of the key circuits of the budgetary matroid $B^{2,2,2}_{1,1,1}$. (Imagine the points A_1 and A_3 as being above the plane of the paper, while B_2 and C_2 are below it.)

construct a representation. To do the latter, let r_1, r_2, and r_3 be three skew lines in 3-space and let s_1, s_2, and s_3 be three of their common transversals, as shown in Figure 5.1. (We are going to study this configuration of nine points and six lines in Section 9.3, where we call it a *grid*.) We set $A_i := r_i \cap s_i$, $B_i := r_{i+1} \cap s_{i-1}$, and $C_i := r_{i-1} \cap s_{i+1}$ for i in $[1\,..\,3]$, where all subscripts are interpreted modulo 3.

By the way, a matroid theorist [42] would think of this matroid M as the dual of the cycle matroid of the complete bipartite graph $K_{3,3}$. The nine edges of the graph $K_{3,3}$ correspond to the nine points in Figure 5.1, while its six vertices correspond to the three lines (r_i) and the three lines (s_i).]

5.5 The representability of budgetary matroids

What makes the budget matroids interesting, as a subclass of the budgetary matroids, is that so many of them are representable. Turning that statement around, what makes the budgetary matroids boring — at least, to a geometer — is that so few of them are representable. Our goal in this section is to report on some numeric experiments that support that latter claim.

To introduce those numeric experiments, let's take each of the example budget matroids $B_{b_1,\ldots,b_k}$ that we considered in Chapter 4, and let's study the various budgetary matroids $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$ that result from relaxing the dimension constraints on the columns by various amounts. We shall refer to a matroid of the form $B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$, where $d_j > b_j$ for at least one j in $[1\,..\,k]$, as a *budgetary relaxation* of the budget matroid $B_{b_1,\ldots,b_k}$. Warning: The term 'relaxation' has a technical meaning in matroid theory that is similar in spirit to this, but is definitely not the same concept; see Exercise 5.5-1.

The matroid $B_{1,1} = B^{1,1}_{1,1}$ doesn't have any budgetary relaxations, so there is nothing to say.

As for the complete quadrilateral $B_{2,1} = B^{2,1}_{2,1}$, the only way to get a budgetary relaxation is to raise the dimension d_2 from 1 to 2. As we discussed in Section 5.2, a representation of the budgetary relaxation $B^{2,2}_{2,1}$ consists of a triangle together with one point on each of its sides, and we have $\#\big(B^{2,2}_{2,1}\big) = 9$.

The Pappus configuration $B_{1,1,1}$ is a more interesting case. Starting from $B^{1,1,1}_{1,1,1}$, we can raise either one, two, or all three of the column dimensions (d_j) from 1 to 2, and the results are as follows:

$$\#\big(B^{1,1,1}_{1,1,1}\big) = 10 \qquad \#\big(B^{1,2,2}_{1,1,1}\big) = 11$$
$$\#\big(B^{1,1,2}_{1,1,1}\big) = \bot \qquad \#\big(B^{2,2,2}_{1,1,1}\big) = 12.$$

Recall that the Pappus configuration $B^{1,1,1}_{1,1,1}$ itself has nine collinearities, but they are redundant: Requiring any eight of them forces the ninth to hold as well, by Pappus's Theorem, so we have $\#\big(B^{1,1,1}_{1,1,1}\big) = 18 - 8 = 10$. If we relax a single column dimension from 1 to 2, the resulting budgetary matroid $B^{1,1,2}_{1,1,1}$ requires eight of the Pappian collinearities, but forbids the ninth, and hence is not representable over any field. Indeed, in matroid theory, the relaxation $B^{1,1,2}_{1,1,1}$ is called *non-Pappus*, and it serves as a standard example of a matroid whose unrepresentability has nothing to do with the characteristic of the field of scalars. If we relax two or three columns, however, the matroids return to representability; there is no redundancy, so the degrees of freedom are easy to count.

5.5.1 The partition $4 = 3 + 1$

Moving on to rank 4, recall that a representation of the budget matroid $B_{3,1}$ consists of the four vertices of a tetrahedron in 3-space, together with the four points where a line cuts its four faces. Recall also that $\#(B_{3,1}) = 16$.

What happens if we raise the dimension d_2 of the budget matroid $B_{3,1} = B^{3,1}_{3,1}$, to get a budgetary relaxation? There are no redundancies among the constraints, so all possible relaxations are representable and counting the degrees of freedom is easy. If we relax the constraint on the four points of the second column from collinearity to coplanarity, the cost of that column constraint goes down from 4 to

1, so the overall freedom goes up to $\#\big(B^{3,2}_{3,1}\big) = 19$. If we let the four points in the second column points span the entire ambient 3-space, we have $\#\big(B^{3,3}_{3,1}\big) = 20$.

5.5.2 The partition $4 = 2 + 2$

Recall that a representation of the budget matroid $B_{2,2}$ consists of the eight vertices of a Möbius pair of tetrahedra. Recall also that the eight coplanarity constraints on those eight points are redundant: If any seven of the eight coplanarities hold, the eighth must hold as well. So we have $\#(B_{2,2}) = 24 - 7 = 17$.

Our choices, in making a budgetary relaxation, are to increase either one or both of the column dimensions from 2 to 3. If we increase just one of them, the resulting matroid $B^{2,3}_{2,2}$ is unrepresentable, since it requires seven of the Möbius-pair coplanarities, while forbidding the eighth. So we have $\#\big(B^{2,3}_{2,2}\big) = \perp$. If we increase both, the six coplanarity constraints remaining in the relaxation $B^{3,3}_{2,2}$ are nonredundant, and we have $\#\big(B^{3,3}_{2,2}\big) = 24 - 6 = 18$.

5.5.3 The partition $4 = 2 + 1 + 1$

Recall that a representation of the budget matroid $B_{2,1,1}$ consists of twelve points

$$\begin{array}{c} \begin{matrix} \scriptstyle 2 & \scriptstyle 1 & \scriptstyle 1 \end{matrix} \\ \begin{pmatrix} P_1 & A_1 & B_1 \\ P_2 & A_2 & B_2 \\ P_3 & A_3 & B_3 \\ P_4 & A_4 & B_4 \end{pmatrix} \end{array}$$

in 3-space. Recall also that a naive count of the degrees of freedom would lead one to guess that $\#(B_{2,1,1}) = 36 - 21 = 15$, while the truth is that $\#(B_{2,1,1}) = 19$.

This high level of redundancy among the constraints causes lots of the budgetary relaxations of the budget matroid $B^{2,1,1}_{2,1,1}$ to be unrepresentable. Here is what seems to be the situation:

$$\begin{array}{ll}
\#\big(B^{2,1,1}_{2,1,1}\big) = 19 & \#\big(B^{3,1,1}_{2,1,1}\big) = \perp \\
\#\big(B^{2,1,2}_{2,1,1}\big) = \perp & \#\big(B^{3,1,2}_{2,1,1}\big) = \perp \\
\#\big(B^{2,1,3}_{2,1,1}\big) = \perp & \#\big(B^{3,1,3}_{2,1,1}\big) \doteq 20 \\
\#\big(B^{2,2,2}_{2,1,1}\big) \doteq 22 & \#\big(B^{3,2,2}_{2,1,1}\big) \doteq 22 \\
\#\big(B^{2,2,3}_{2,1,1}\big) = \perp & \#\big(B^{3,2,3}_{2,1,1}\big) \doteq 23 \\
\#\big(B^{2,3,3}_{2,1,1}\big) \doteq 23 & \#\big(B^{3,3,3}_{2,1,1}\big) \doteq 24
\end{array}$$

Warning: The number 19 and the five unrepresentable cases are facts, proved later in this monograph (Theorem 7.3-1 and Exercises 7.2-2 and 7.2-4). The numbers in the other six cases are the authors' educated guesses, supported by numeric experiments — hence the use of the dotted equal sign '$\doteq$'.

Assuming that all of the numbers are correct, what do they mean? The final relation $\#(B^{3,3,3}_{2,1,1}) \doteq 24$ tells us that 24 degrees of freedom remain if we require only the 12 perfect constraints, so those 12 constraints are nonredundant. If we add to those 12 the constraint that any single column be coplanar, the freedom drops to 23, so the new constraint is also nonredundant: $\#(B^{2,3,3}_{2,1,1}) \doteq \#(B^{3,2,3}_{2,1,1}) \doteq 23$.

Once the four P-points are coplanar, however, a theorem of projective geometry comes into play: When the twelve perfect coplanarities hold and the P column is coplanar, Theorem 7.2-1 tells us that the A and B columns must have spans of the same dimension. It follows that the budgetary relaxations $B^{2,1,2}_{2,1,1}$, $B^{2,1,3}_{2,1,1}$, and $B^{2,2,3}_{2,1,1}$ are unrepresentable. Also, if we require the P and A columns to be coplanar, the B column is forced to be coplanar as well; thus, the freedom for the case $\#(B^{2,2,2}_{2,1,1})$ is 22, rather than 21. Considering this case naively, one might suspect that, given the perfect coplanarities, the three column coplanarities are mutually redundant, with any one following from the other two. But coplanarities are quartic equations on the homogeneous coordinates, not linear ones; so naive reasoning about redundancy can fail. It turns out that requiring both the A and the B columns to be coplanar does not force the P column to be coplanar: Instead of having $\#(B^{3,2,2}_{2,1,1}) = \perp$, we have $\#(B^{3,2,2}_{2,1,1}) \doteq 22$. That is, suppose that we require the perfect coplanarities and the coplanarities of the A and B columns. The resulting system of 14 quartic equations seems to define a reducible variety with at least two components of dimension 22. Throughout one of those components, the P column is also coplanar; but the P column is not coplanar at a generic point of another of those components.

We have three budgetary relaxations of $B_{2,1,1}$ left to consider, in all of which the P points are not coplanar. It costs four degrees of freedom to make the A-points collinear, and those four constraints are independent of the twelve perfect constraints; so, if that is all that we require, we have $\#(B^{3,1,3}_{2,1,1}) \doteq 20$. But Exercise 7.2-4 shows that, once the A-points are collinear, requiring the B-points to be coplanar — which would naively reduce the freedom to 19 — either forces the B-points to be collinear and the P-points to be coplanar or else forces some degeneracy. So the two remaining relaxations are unrepresentable: $\#(B^{3,1,1}_{2,1,1}) = \#(B^{3,1,2}_{2,1,1}) = \perp$.

Exercise 5.5-1 Let M be a matroid and let X be a subset of M that is both a circuit and a hyperplane. So X itself is dependent, every proper subset of X is independent, no element not in X lies in $\mathrm{Span}(X)$, and $|X| = \mathrm{rank}(M)$. It is well known [38] that we can build a matroid M' that differs from M only in that the set X is independent in M' — and hence X is a base of M'. The process of going from M to M' is called *relaxing the circuit-hyperplane* X. For example, the budgetary matroid $B^{3,1,1}_{2,1,1}$ is the result of relaxing the circuit-hyperplane $\{P_1, P_2, P_3, P_4\}$ in the budget matroid $B_{2,1,1}$.

Let the matroid $B^{d_1,\dots,d_k}_{b_1,\dots,b_k}$ be some budgetary relaxation of the budget matroid $B^{b_1,\dots,b_k}_{b_1,\dots,b_k}$. When is it the case that we can get from the latter to the former by relaxing some set of circuit-hyperplanes?

[Answer: Only occasionally, because the j^{th} column of the budget matroid $B_{b_1,\dots,b_k}$ is a circuit-hyperplane only when $b_j = b - 2$. That relationship holds only for the first columns of $B_{m,2}$ and $B_{m,1,1}$, for any positive m. (It holds also for the second columns of $B_{2,1}$ and $B_{2,2}$ and for all three columns of $B_{1,1,1}$, but those cases are symmetric.)]

5.5.4 The partition $4 = 1 + 1 + 1 + 1$

Recall that a representation of the budget matroid $B_{1,1,1,1}$ consists of sixteen points in 3-space:

$$\begin{array}{cccc} 1 & 1 & 1 & 1 \\ \end{array}$$
$$\begin{pmatrix} A_1 & B_1 & C_1 & D_1 \\ A_2 & B_2 & C_2 & D_2 \\ A_3 & B_3 & C_3 & D_3 \\ A_4 & B_4 & C_4 & D_4 \end{pmatrix},$$

The four points in each column are collinear, while the 24 perfect sets are coplanar. A naive count of the degrees of freedom would lead one to guess that $\#(B_{1,1,1,1}) = 48 - 40 = 8$, while the truth — at least, over the complex numbers — is that $\#(B_{1,1,1,1}) = 20$. So the redundancy accounts for 12 degrees of freedom.

The more redundancy there is among the constraints of a budget matroid, the more likely it is that a random budgetary relaxation of that matroid will be unrepresentable. Here is what seems to be the story for the matroid $B_{1,1,1,1}$:

$$\begin{array}{lll}
\#\big(B^{1,1,1,1}_{1,1,1,1}\big) = 20 & & \\
\#\big(B^{1,1,1,2}_{1,1,1,1}\big) = \bot & & \\
\#\big(B^{1,1,1,3}_{1,1,1,1}\big) = \bot & & \\
\#\big(B^{1,1,2,2}_{1,1,1,1}\big) = \bot & & \\
\#\big(B^{1,1,2,3}_{1,1,1,1}\big) = \bot & & \\
\#\big(B^{1,1,3,3}_{1,1,1,1}\big) = \bot & & \\
\#\big(B^{1,2,2,2}_{1,1,1,1}\big) \doteq \bot & \#\big(B^{2,2,2,2}_{1,1,1,1}\big) \doteq 21 & \\
\#\big(B^{1,2,2,3}_{1,1,1,1}\big) \doteq \bot & \#\big(B^{2,2,2,3}_{1,1,1,1}\big) \doteq \bot & \\
\#\big(B^{1,2,3,3}_{1,1,1,1}\big) \doteq \bot & \#\big(B^{2,2,3,3}_{1,1,1,1}\big) \doteq 22 & \\
\#\big(B^{1,3,3,3}_{1,1,1,1}\big) \doteq \bot & \#\big(B^{2,3,3,3}_{1,1,1,1}\big) \doteq 23 & \#\big(B^{3,3,3,3}_{1,1,1,1}\big) \doteq 24
\end{array}$$

Warning: Only the first six lines in this table are proven facts (Proposition 10.1-2 and Exercise 5.6-3); the rest merely report on the authors' numeric experiments — but let's view the entire table as correct, in the following paragraph.

The last line $\#\big(B^{3,3,3,3}_{1,1,1,1}\big) \doteq 24$ implies that the 24 perfect constraints themselves are nonredundant. The gap between this and the first line $\#\big(B^{1,1,1,1}_{1,1,1,1}\big) = 20$ is only 4. Indeed, if we require even one of the four columns to be collinear, that uses up all four of those degrees of freedom, apparently forcing all four columns to be collinear. The remaining subtle point is that requiring any three columns to be coplanar seems to force the last column to be coplanar as well; so we have $\#\big(B^{2,2,2,3}_{1,1,1,1}\big) \doteq \perp$ and $\#\big(B^{2,2,2,2}_{1,1,1,1}\big) \doteq 24 - 3 = 21$.

5.5.5 The partition $5 = 2 + 1 + 1 + 1$

We have now covered all of the budget matroids of rank at most 4. Note that, as the number k of parts in the partition $b = b_1 + \cdots + b_k$ increases and as the parts themselves increase, the percentage of the budgetary relaxations $B^{d_1,\dots,d_k}_{b_1,\dots,b_k}$ that are representable goes down. This trend continues aggressively: In larger examples, it is often the case that none of the budgetary relaxations are representable.

One example that will be important to us in Chapter 11 is the budget matroid $B_{2,1,1,1}$, a representation of which consists of 20 points in 4-space:

$$\begin{array}{c} \begin{array}{cccc} \scriptstyle 2 & \scriptstyle 1 & \scriptstyle 1 & \scriptstyle 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 & B_1 & C_1 \\ P_2 & A_2 & B_2 & C_2 \\ P_3 & A_3 & B_3 & C_3 \\ P_4 & A_4 & B_4 & C_4 \\ P_5 & A_5 & B_5 & C_5 \end{pmatrix} \end{array}.$$

This matroid is definitely representable (Section 11.5 gives a rational representation), and numeric experiments suggest that $\#(B_{2,1,1,1}) \doteq 30$; but it seems that all of the budgetary relaxations of $B_{2,1,1,1}$ — up through and including $B^{4,4,4,4}_{2,1,1,1}$ — are unrepresentable.

Here is some of the reasoning behind that claim; for more details, see Section 11.6. To determine that $\#(B_{2,1,1,1}) \doteq 30$, we find some rational representation R of the matroid $B_{2,1,1,1}$ and we compute the Jacobian matrix of the constraints at the point R. For the representation R that the authors studied, that Jacobian turns out to have rank 50. Since choosing 20 arbitrary points in 4-space involves 80 degrees of freedom, we conclude that the local dimension at R of the variety of representations is $80 - 50 = 30$, so 30 is a good guess for the freedom $\#(B_{2,1,1,1})$. Note that $\binom{5}{2\,1\,1\,1} = 60$ of the constraints are perfect constraints. Since the rank of the entire Jacobian is only 50, this means that the perfect constraints themselves must be redundant. To determine how redundant they are, by themselves, we compute the rank of just those sixty rows of the Jacobian matrix; and we again get 50. So, in the neighborhood of the representation R, some 50 of the perfect constraints imply the other ten perfect constraints and all of the column constraints. If the behavior at R is typical, it follows that all of the budgetary relaxations of $B_{2,1,1,1}$ are unrepresentable.

5.5.6 Some larger partitions

So it seems likely that none of the budgetary relaxations of the matroid $B_{2,1,1,1}$ are representable; but that matroid has four columns. It is worth noting that the same phenomenon arises also when there are only three columns, once the budgets for those three columns get high enough. On the other hand, that phenomenon probably doesn't arise when there are only two columns.

In the 3-column case, the authors found a rational representation R of the budget matroid $B_{3,3,1}$. The rank of the Jacobian of all of the constraints, at R, turned out to be 68. Since a representation consists of 21 points in 6-space, which means 126 degrees of freedom, we conclude that $\#(B_{3,3,1}) \doteq 58$ would be a good guess. Furthermore, the rank at R of just the $\binom{7}{3\ 3\ 1} = 140$ rows of the Jacobian that came from the perfect constraints was also 68. This suggests that none of the budgetary relaxations of $B_{3,3,1}$ — up through and including $B^{6,6,6}_{3,3,1}$ — is representable.

As for the case of only 2 columns, we get the most perfect constraints, for a fixed total budget, when that total budget is divided in half. The authors investigated the following cases:

$$\begin{array}{llll} \#\big(B^{2,2}_{2,2}\big) = 17 & \#\big(B^{3,3}_{3,3}\big) = 42 & \#\big(B^{4,4}_{4,4}\big) = 77 & \#\big(B^{5,5}_{5,5}\big) = 122 \\ \#\big(B^{3,3}_{2,2}\big) = 18 & \#\big(B^{5,5}_{3,3}\big) \doteq 44 & \#\big(B^{7,7}_{4,4}\big) \doteq 83 & \#\big(B^{9,9}_{5,5}\big) \doteq 134 \end{array}$$

The facts in the upper row are from Theorem 6.2-2; the guesses in the lower row result from calculating the ranks of appropriate Jacobians. Even though the number of perfect constraints is growing quite rapidly, as we move from left to right, the redundancy among those constraints grows rapidly enough that the column constraints continue to play an important role — they do not become consequences of the perfect constraints.

But whatever the fine details might turn out to be, the big picture seems clear. Starting with Chapter 6 of this monograph, we shall focus almost exclusively on the budget matroids, since they seem so much more likely to be representable than their budgetary relaxations.

5.6 Budgetary minors of budgetary matroids

It is often the case that one budgetary matroid is a minor of another, in which case we refer to the former as a *budgetary minor*. This relationship is worth noting because, when M is a minor of M' and M is not representable, we get an easy proof that M' is not representable either. In this section, we use appropriate deletions and contractions to prove two results about which budgetary matroids are minors of which others.

Our first proposition says that deleting a zero column budget from the partition of a budgetary matroid M' always leads to a budgetary minor M, regardless of the dimension of the deleted column. In this case, the ground matrix of the budgetary minor M has the same number of rows as M', but one less column.

Proposition 5.6-1 *Let M be a budgetary matroid $M := B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$, and let $b := b_1 + \cdots + b_k$ be its total budget. For any value of d_{k+1} in the range $[0 \mathinner{.\,.} b)$, the budgetary matroid $M' := B^{d_1,\ldots,d_k,d_{k+1}}_{b_1,\ldots,b_k,0}$ has M as a minor.*

Proof In the larger matroid M', let C_{k+1} denote the elements in the last column, and let X be any set of elements with $X \cap C_{k+1} = \emptyset$. It is easy to check that X is independent in M' if and only if that same set X is independent in the smaller matroid M. Thus, we have $M = M' \backslash C_{k+1}$. □

Our second proposition says that reducing a column budget by one in the partition of a budgetary matroid M' also leads to a budgetary minor M, as long as the dimension of the affected column is reduced by one also, the dimensions of the other columns are reduced by one, if necessary, to keep them less than the reduced total budget, and at least two column budgets remain positive after the reduction. In this case, the ground matrix of the budgetary minor M has the same number of columns as M', but one less row.

Proposition 5.6-2 *Let M be a budgetary matroid $M := B^{d_1,\ldots,d_k}_{b_1,\ldots,b_k}$, and let m in the interval $[1 \mathinner{.\,.} k]$ be the index of a column in M. Define the augmented column budgets (b'_j) by the rules*

$$\begin{aligned} b'_m &:= b_m + 1 \\ b'_j &:= b_j \quad \text{for } j \neq m; \end{aligned}$$

so the m^{th} column budget goes up by 1, while all others stay the same, and hence the total budget $b' = b + 1$ goes up by 1. Also, define the augmented column dimensions (d'_j) by the rules

$$\begin{aligned} d'_m &:= d_m + 1 \\ d'_j &:= d_j && \text{when } j \neq m \text{ and } d_j \leq b - 2 \\ d'_j &\in \{b-1, b\} && \text{when } j \neq m \text{ and } d_j = b - 1; \end{aligned}$$

so the m^{th} column dimension goes up by 1 as well, while all others stay the same, except that any column dimension that starts out at $b - 1$, which is the old upper bound, may either stay the same or increase by 1 to b, the new upper bound. The augmented budgetary matroid $M' := B^{d'_1,\ldots,d'_k}_{b'_1,\ldots,b'_k}$ then has M as a minor.

Proof We contract some element of the m^{th} column of the matroid M' — say the last, $e_{b'm}$ — and we delete all of the other elements in that same row. To show that this works, let X denote any subset of the smaller matroid M. We must show that X is independent in M just when the set $X' := X \cup \{e_{b'm}\}$ is independent in M'.

The set X' has precisely one more element overall than does X. But the total budget b' is also one larger than the total budget b; so the set X satisfies the Ambient Rule in M just when X' satisfies that same rule in M'.

The Column Rule in the m^{th} column is similar. Letting C_j denote the j^{th} column of M', we have $|X' \cap C_m| = |X \cap C_m| + 1$, because of the new element $e_{b'm}$ in X'. But the column dimension d'_m also exceeds the dimension d_m by 1.

As for the Column Rule in the j^{th} column for some $j \neq m$, the sets $X \cap C_j$ and $X' \cap C_j$ have the same elements. If $d_j = d'_j$, then the Column Rule holds for X in M just when it holds for X' in M'. If $d_j = b - 1$ and $d'_j = b$, then the Column Rule holds both for X in M and for X' in M'.

It remains to consider the Perfect Rule. If the set X has a perfect subset Y in the matroid M, then the set X' has the perfect subset $Y' := Y \cup \{e_{b'm}\}$ in M'. Conversely, if the set X' has any perfect subset Y', the element $e_{b'm}$ must belong to Y', since $e_{b'm}$ is the only element of X' in the last row; and the subset $Y := Y' \setminus \{e_{b'm}\}$ of X is then perfect in M. $\square$

Exercise 5.6-3 The budgetary relaxation $B^{1,1,2}_{1,1,1}$ of the Pappus matroid $B_{1,1,1}$ is not representable over any field. Using just that fact and the two results of this section, which budgetary relaxations of the matroids $B_{2,1,1}$ and $B_{1,1,1,1}$ can you conclude are not representable?

[Answer: In the former case, $B^{2,1,2}_{2,1,1}$, $B^{2,1,3}_{2,1,1}$, and $B^{3,1,1}_{2,1,1}$; in the latter case, $B^{1,1,1,2}_{1,1,1,1}$, $B^{1,1,1,3}_{1,1,1,1}$, $B^{1,1,2,2}_{1,1,1,1}$, $B^{1,1,2,3}_{1,1,1,1}$, and $B^{1,1,3,3}_{1,1,1,1}$. Comparing with the data in Sections 5.5.3 and 5.5.4, we see that, in both cases, about half of the budgetary relaxations that are unrepresentable can be shown to be so simply because they have the non-Pappus matroid $B^{1,1,2}_{1,1,1}$ as a budgetary minor.]

Chapter 6

Representing the matroid $B_{m,n}$

The budget matroids are an interesting subclass of the budgetary matroids because so many of the former are representable. Indeed, to the best of the authors' current knowledge, it might be the case that every budget matroid is representable over the rational numbers. We shan't come close to proving that wild conjecture. But we shall show that two infinite families of budget matroids are representable over the rationals. In this chapter, we consider the budget matroids $B_{m,n}$ that have just two columns.

In addition to demonstrating representability, we shall also count the degrees of freedom involved, showing that $\#(B_{m,n}) = m^2+3mn+n^2-3$. Recall that $\#(B_{m,n})$ denotes the number of degrees of freedom involved in choosing a representation of $B_{m,n}$, lying in some fixed projective space of the appropriate dimension, which is $m + n - 1$.

6.1 The case $m = n = 2$ of a Möbius pair

The case $n = 1$ (or, symmetrically, the case $m = 1$) is easy. A representation of the budget matroid $B_{m,1}$ consists of an m-simplex in m-space along with the $m+1$ points where a line cuts the facets of the simplex. There are m degrees of freedom in each of the $m + 1$ vertices of the simplex and an additional $2m - 2$ degrees of freedom in the choice of the line, for a total of $m^2 + 3m - 2$, which agrees with the formula above.

Thus, the first tricky case is the case $B_{2,2}$ of a Möbius pair of tetrahedra. Recall that a representation of the matroid $B_{2,2}$ consists of eight points

$$\begin{array}{cc} 2 & 2 \\ \end{array}$$
$$\begin{pmatrix} A_1 & B_1 \\ A_2 & B_2 \\ A_3 & B_3 \\ A_4 & B_4 \end{pmatrix}$$

in 3-space, with the A-points, the B-points, and the six perfect sets coplanar.

We can construct such a representation in a fixed 3-space as follows. Choose the plane β, in which the B-points will lie — that's three degrees of freedom — and choose the four points (B_i) in the plane β so that no three are collinear — that's another eight. Choose a line c in the plane β that intersects the six sides of the complete quadrangle with vertices (B_i) in six distinct points; that is, the line c must not pass through any of the four vertices of the quadrangle nor through any of the three diagonal points $B_iB_j \cap B_kB_l$, where $\{i, j\} \cup \{k, l\} = \{1, 2, 3, 4\}$. And choose a plane α through the line c, but distinct from β, so that $c = \alpha \cap \beta$. There are two degrees of freedom in the choice of the line c and one more in the choice of the plane α; so the current total freedom is fourteen.

It remains to choose the points (A_i) in the plane α. But we must be careful, when making these choices, to arrange that the six perfect sets end up being coplanar. Note that each perfect set has the form $\{A_i, A_j, B_k, B_l\}$, where $i < j$, $k < l$, and $\{i, j, k, l\} = \{1, 2, 3, 4\}$. Making those four points coplanar means arranging that the two points $A_iA_j \cap c$ and $B_kB_l \cap c$ on the line c coincide, as shown in Figure 6.1. Let $C_{ij;kl}$ denote the point $B_kB_l \cap c$. Choose the point A_1 arbitrarily in the plane α, but not on c; and choose the point A_2 arbitrarily on the line $A_1C_{12;34}$, but not coincident with either A_1 or $C_{12;34}$. Everything else is forced: We must choose A_3 to be the point where the lines $A_1C_{13;24}$ and $A_2C_{23;14}$ intersect, and we must choose A_4 similarly to be the point where $A_1C_{14;23}$ meets $A_2C_{24;13}$. This guarantees that all six perfect sets are coplanar except possibly for $\{A_3, A_4, B_1, B_2\}$. But the six points $\big\{\{C_{12;34}, C_{34;12}\}, \{C_{13;24}, C_{24;13}\}, \{C_{14;23}, C_{23;14}\}\big\}$ on the line c result from stabbing a complete quadrangle in the plane β; so they form a quadrangular set, and any one of them is determined by the other five. The four points (A_i) in the plane α determine another complete quadrangle, which the line c stabs at five of the same points; so the sixth point must also be shared, that is, the line A_3A_4 must meet c at the point $C_{34;12}$.

Since choosing the points A_1 and A_2 involves an additional three degrees of freedom, the grand total is seventeen, which agrees with the formula above.

Note that Figure 6.1 has a lot in common with Figure 2.3. Figure 2.3 shows a 2-block of lines, all through a common point, and shows two complete quadrilaterals that witness to the 2-dependence of the six slopes. Dually, Figure 6.1 shows a 2-block of points, all on a common line, and shows two complete quadrangles that witness to the 2-dependence of the six coordinates. In Figure 2.3, there is no projective transformation that takes the first complete quadrilateral to the second while fixing all lines through O, because one of the three pairs gets swapped. In Figure 6.1, there is no projective map from the plane α to the plane β that carries the first quadrangle to the second while fixing every point along the line c, because all three pairs get swapped. To see concretely that no such projective map exists, note that the lines in the plane α passing through the three points $C_{12;34}$, $C_{13;24}$, and $C_{14;23}$ concur at A_1, while the lines in the plane β passing through those three points are not concurrent.

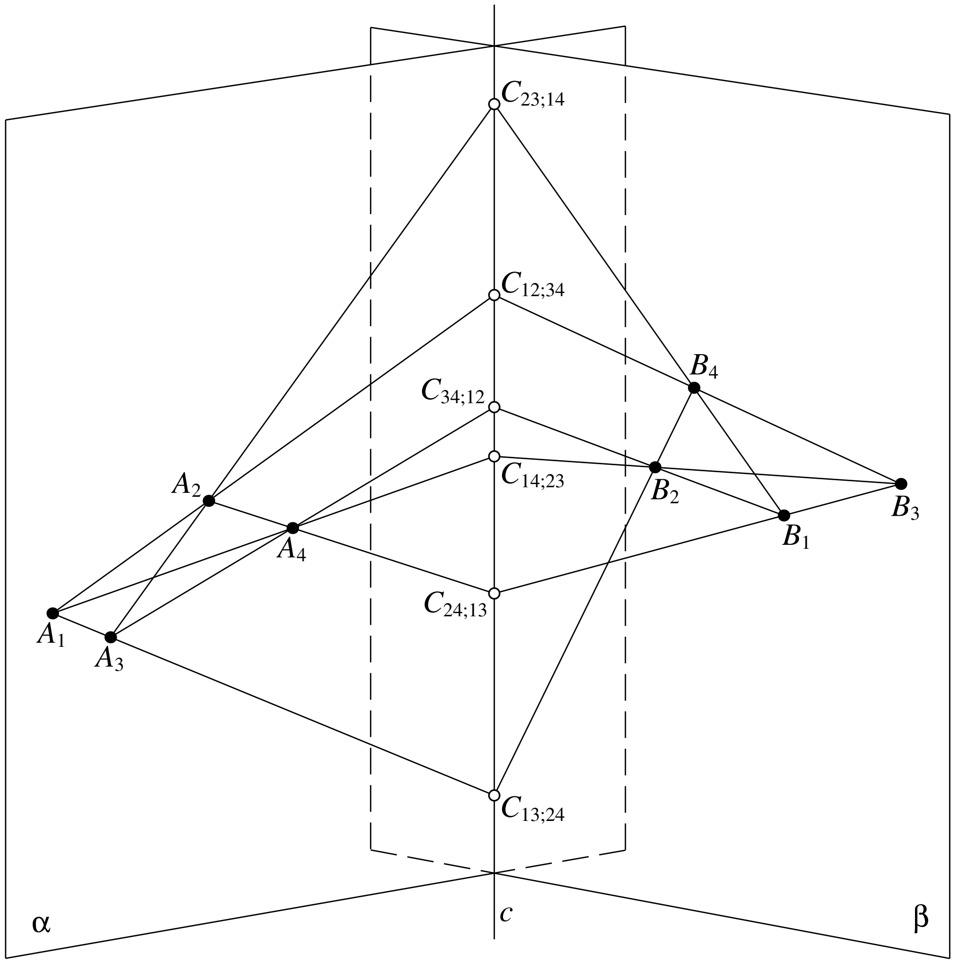

Figure 6.1: Constructing a representation of the budget matroid $B_{2,2}$ in 3-space.

We have shown, so far, that the eight points (A_i) and (B_i) constructed as above will have all of the incidences that are required for a representation of the budget matroid $B_{2,2}$. But will they be free of forbidden incidences? To check this, since subsets of a mutually skew set are automatically mutually skew, it suffices to check that all of the bases of the matroid $B_{2,2}$ are mapped, by our construction, to sets of four points that are not coplanar.

We required the six points $(C_{ij;kl})$ to be distinct. This implies that neither A_3 nor A_4 ends up along the line c, so any base consisting of one A-point and three B-points is non-coplanar. It also implies that no three of the four points A_i are collinear in the plane α, so any base consisting of three A-points and one B-point is also non-coplanar. The remaining bases have the form $\{A_i, A_j, B_{k'}, B_{l'}\}$ where $i < j$ and $k' < l'$, but the sets of row indices $\{i, j\}$ and $\{k', l'\}$ are not disjoint. The line A_iA_j meets the line c at the point $C_{ij;kl}$, while the line $B_{k'}B_{l'}$ meets the line c at the distinct point $C_{i'j';k'l'}$. It follows that the base $\{A_i, A_j, B_{k'}, B_{l'}\}$ is also non-coplanar.

6.2 The general case

Extending these arguments from the case $B_{2,2}$ to the case $B_{m,n}$ for any positive m and n involves a fair amount of fussy detail, but doesn't require much in the way of new ideas. The only interactions between the A-points, lying in their m-flat, and the B-points, lying in their n-flat, occur along the line where those two flats meet. And what we have, lying along that line, are lots of overlapping quadrangular sets, each of which can be analyzed separately.

We begin with an easy lemma that we shall use also in Chapter 7.

Lemma 6.2-1 *Let P_1 through P_n be n mutually skew points in a projective n-space S, and let $\pi = \operatorname{Span}(P_1, \ldots, P_n)$ be the hyperplane of S that they span. Let A_1 through A_n be any n points of S, none lying in the hyperplane π. There exists a unique point B that lies in all n of the hyperplanes*

$$H_i := \operatorname{Span}(P_1, \ldots, \widehat{P_i}, \ldots, P_n, A_i),$$

for i in $[1\,..\,n]$ — that is, the intersection $\bigcap_{i\in[1..n]} H_i = \{B\}$ is a flat of dimension 0 — and the intersection point B does not lie in π.

Proof Any n hyperplanes intersect in at least a point, so the challenge is to show that the intersection $\beta := \bigcap_{i\in[1..n]} H_i$ contains just one point. Note that the points P_1 through P_n are the vertices of an $(n-1)$-simplex in the hyperplane π, whose n facets are the $(n-2)$-flats

$$\pi \cap H_i = \operatorname{Span}(P_1, \ldots, \widehat{P_i}, \ldots, P_n).$$

Since no point is common to all of the facets of a simplex, it follows that the intersection $\pi \cap \beta$ is empty. Since the flat β is skew to the hyperplane π, it must consist of a single point B that does not lie in π. $\square$

Theorem 6.2-2 ($B_{m,n}$ Representation) *For any positive m and n, the budget matroid $B_{m,n}$ is representable over the rationals, and the number of degrees of freedom involved in choosing a representation, lying in a fixed projective space of dimension $m+n-1$, is $\#(B_{m,n}) = m^2 + 3mn + n^2 - 3$.*

Proving this theorem will occupy us for the rest of Chapter 6. If you get bogged down in the proof, feel free to skip on to Chapter 7.

6.2.1 The initial choices

Fix a projective space S of dimension $m+n-1$. To construct a representation of $B_{m,n}$ in the space S, we first pick an n-flat β, lying in S; this involves $(n+1)(m-1)$ degrees of freedom. Next, we choose the points (B_i), for i in $[1 \mathinner{..} m+n]$, lying in the flat β, which involves $(m+n)n$ degrees of freedom. Of course, every set of at least $n+2$ of the B-points must be mutually incident, since they all lie in a common n-flat. But we constrain our choices of the (B_i) so that those are the only incidences; that is, every set of at most $n+1$ of the B-points is mutually skew.

Given a subset $J \subseteq [1 \mathinner{..} m+n]$, let us denote by β_J the subspace $\beta_J :=$ Span$(\{B_j \mid j \in J\})$ of β. And let us use the term *t-set* to describe any subset $J \subseteq [1 \mathinner{..} m+n]$ that has t elements. Because we have chosen the B-points to be in general position within the n-flat β, we have $\dim(\beta_J) = |J| - 1$ whenever $|J| \le n+1$. In particular, for any $(n+1)$-set J, we have $\beta_J = \beta$. For any n-set J, we have $\dim(\beta_J) = n-1$, so the flat β_J is a hyperplane in β. If J and J' are two distinct n-sets, the two hyperplanes β_J and $\beta_{J'}$ must be distinct, since $\beta_{J \cup J'} = \beta$. It follows that their intersection $\beta_J \cap \beta_{J'}$ has dimension $n-2$.

Choose a line c in the flat β that is skew to the $(n-2)$-flat $\beta_J \cap \beta_{J'}$, for all pairs of distinct n-sets (J, J'). And choose any m-flat α that intersects the n-flat β precisely in the line c. The choice of c involves $2(n-1)$ degrees of freedom and the choice of α involves $(m-1)(n-1)$, so our running total is now $3mn + n^2 - 2$.

For every n-set J, the hyperplane β_J in β cannot include the entire line c, since $\beta_J \cap \beta_{J'}$ is skew to c for any n-set J' distinct from J. So the hyperplane β_J intersects c at a unique point. By analogy with the case $m = n = 2$ above, we should denote that point of intersection by $C_{I;J}$, where I denotes the m-set $I := [1 \mathinner{..} m+n] \setminus J$ that is complementary to the n-set J. But, to save writing, we shall instead denote it simply by C_I. If $(I; J)$ and $(I'; J')$ are any two distinct partitions of the set $[1 \mathinner{..} m+n]$ into an m-set and an n-set, the fact that $\beta_J \cap \beta_{J'}$ is skew to c implies that the points $C_I := \beta_J \cap c$ and $C_{I'} := \beta_{J'} \cap c$ are distinct. So there are $\binom{m+n}{m}$ distinct points (C_I) along the line c, one for each partition $(I; J)$ of $[1 \mathinner{..} m+n]$ into an m-set I and an n-set J.

6.2.2 Goals for choosing the A-points

Our next task is to choose the points (A_i), lying in the m-flat α. We want the incidence structure of the A-points in α to be completely analogous to the structure of the B-points in β. Our first goal in this regard is that we want every set of at most $m+1$ of the A-points to be mutually skew. That is, letting α_I denote the subspace $\alpha_I := \mathrm{Span}(\{A_i \mid i \in I\})$ of α, we want $\dim(\alpha_I) = |I| - 1$ whenever $|I| \leq m+1$. Since subsets of a mutually skew set are mutually skew, it suffices to make sure, for every $(m+1)$-set I, that $\alpha_I = \alpha$. Our second goal is that, for any m-set I, the flat α_I — which we have already said should be a hyperplane — should intersect the line c at the single point C_I.

If those two goals are met, then the incidence structure of the A-points in α will be completely analogous to that of the B-points in β. For example, given any two distinct m-sets I and I', the $(m-2)$-flat $\alpha_I \cap \alpha_{I'}$ will be skew to c, since the hyperplanes α_I and $\alpha_{I'}$ in α intersect c precisely in the distinct points C_I and $C_{I'}$.

Our plan for achieving these goals involves an induction on the weight of the subset I, where we define the *weight* $w(I)$ of a set $I \subseteq [1 \,..\, m+n]$ to be the number of its elements that exceed m; so we have $w(I) := \big|I \cap [m+1 \,..\, m+n]\big|$. We therefore formalize our two goals as follows:

Goal1 The property $G_1(t)$ holds when, for all $(m+1)$-sets I with weight $w(I) = t$, we have $\alpha_I = \alpha$.

Goal2 The property $G_2(t)$ holds when, for all m-sets I with weight $w(I) = t$, we have $\dim(\alpha_I) = m-1$ and $\alpha_I \cap c = \{C_I\}$.

Note that we have included, as part of Goal2, the assertion that the flat α_I is a hyperplane. If we were able to prove $G_1(t)$ for all t before worrying about Goal2, then this aspect of Goal2 would follow automatically. But we are actually going to establish the two goals using a simultaneous induction. The base cases are $G_1(1)$, $G_2(0)$, and $G_2(1)$. The inductive steps show that

- For all $t \geq 2$, the property $G_2(t-1)$ implies $G_1(t)$.
- For all $t \geq 2$, the properties $G_1(t-1)$, $G_1(t)$, $G_2(t-2)$, and $G_2(t-1)$ together imply $G_2(t)$.

6.2.3 Choosing the A-points

Choose the points A_1 through A_{m-1} to be mutually skew in the flat α and so that the $(m-2)$-flat that they span is skew to the line c. That involves $(m-1)m$ degrees of freedom. The m points $\{A_1, \ldots, A_{m-1}, C_{[1..m]}\}$ are then mutually skew, so their span

$$\pi := \mathrm{Span}(A_1, \ldots, A_{m-1}, C_{[1..m]})$$

is a hyperplane in α. We choose the point A_m to lie in the hyperplane π, subject to some nondegeneracy conditions that we discuss in a second. The choice of A_m involves a final $m-1$ degrees of freedom, for a grand total of $m^2+3mn+n^2-3$, as claimed.

Choosing the point A_m in the hyperplane π forces the set $\{A_1,\ldots,A_m,C_{[1..m]}\}$ to be mutually incident, but we constrain the choice of A_m so that no proper subset of that set is mutually incident. Among other things, this implies that the m points $\{A_1,\ldots,A_m\}$ span the entire hyperplane π, so we have $\pi = \mathrm{Span}(A_1,\ldots,A_m)$. Since we chose the $(m-2)$-flat $\mathrm{Span}(A_1,\ldots,A_{m-1})$ to be skew to the line c, the hyperplane π can't intersect c in more than a single point; so we have $\pi\cap c = \{C_{[1..m]}\}$.

We have finished making choices, but we haven't finished constructing all of the A-points. Let k be an index in the interval $[m+1\mathinner{..}m+n]$. Our choice of the point A_k is constrained by our need to arrange things so that, for each i in $[1\mathinner{..}m]$, the m points $\{A_1,\ldots,\widehat{A_i},\ldots,A_m,A_k\}$ will span a hyperplane in α whose intersection with the line c consists of the single point $C_{[1..m]\cup\{k\}\setminus\{i\}}$. Turning this around, we must choose the point A_k to lie in the hyperplane

$$H_{i,k} := \mathrm{Span}(A_1,\ldots,\widehat{A_i},\ldots,A_m,C_{[1..m]\cup\{k\}\setminus\{i\}}).$$

Note that the flat $H_{i,k}$ is a hyperplane, and not something smaller, because the point $C_{[1..m]\cup\{k\}\setminus\{i\}}$ does not lie in $\pi = \mathrm{Span}(A_1,\ldots,A_m)$. We are now in perfect shape to appeal to Lemma 6.2-1, from which we conclude that the hyperplanes $H_{i,k}$, for i in $[1\mathinner{..}m]$, intersect in a unique point. We call that point A_k. Lemma 6.2-1 also tells us that the point A_k does not lie in π.

6.2.4 The base cases hold

As a result of choosing the A-points in this way, we claim that the three base cases $G_1(1)$, $G_2(0)$, and $G_2(1)$ of the induction hold.

To see that $G_1(1)$ holds, note that any $(m+1)$-set I of weight 1 has the form $I = [1\mathinner{..}m]\cup\{k\}$, for some k in $[m+1\mathinner{..}n]$. We have $\alpha_I = \mathrm{Span}(A_1,\ldots,A_m,A_k) = \mathrm{Span}(\pi\cup\{A_k\}) = \alpha$, since Lemma 6.2-1 tells us that A_k does not lie in π.

As for $G_2(0)$, the unique m-set of weight 0 is the set $[1\mathinner{..}m]$, and we have $\alpha_{[1..m]} = \mathrm{Span}(A_1,\ldots,A_m) = \pi$. We have already observed that π is a hyperplane in α and that $\pi\cap c = \{C_{[1..m]}\}$.

We claim that $G_2(1)$ also holds. An m-set I of weight 1 has the form $I := [1\mathinner{..}m]\cup\{k\}\setminus\{i\}$ for some i in $[1\mathinner{..}m]$ and some k in $[m+1\mathinner{..}m+n]$. Since the point A_k does not lie in the hyperplane π, the flat $\mathrm{Span}(A_1,\ldots,A_m,A_k)$ is all of α; so the flat $\mathrm{Span}(A_1,\ldots,\widehat{A_i},\ldots,A_m,A_k) = \alpha_I$ is a hyperplane. Our choice of A_k guaranteed that this hyperplane contains the point C_I. It remains to show that the hyperplane α_I does not contain the entire line c, which we shall do by showing that the $(m-2)$-flat $\mathrm{Span}(A_1,\ldots,\widehat{A_i},\ldots,A_m)$ is skew to c. The case $i=m$ is easy, since, in choosing the points A_1 through A_{m-1}, we constrained the $(m-2)$-flat $\mathrm{Span}(A_1, \ldots, A_{m-1})$ to be skew to c. So suppose that $i < m$. We constrained the choice of the point A_m to guarantee that the $(m-2)$-flat $\mathrm{Span}(A_1, \ldots, \widehat{A_i}, \ldots, A_m)$ does not contain the point $C_{[1..m]}$. But that flat cannot contain any other point on the line c either, because $C_{[1..m]}$ is the only point on the line c that lies in the entire hyperplane $\pi = \mathrm{Span}(A_1, \ldots, A_m)$. Thus, the flat $\mathrm{Span}(A_1, \ldots, \widehat{A_i}, \ldots, A_m)$ is skew to c, for any i in $[1 .. m]$, which completes the proof of $G_2(1)$.

6.2.5 The easy inductive step

We claim that $G_2(t-1)$ implies $G_1(t)$, for any $t \geq 2$.

Let I be any $(m+1)$-set of weight t; we must show that $\alpha_I = \alpha$, that is, that the set $\{A_i\}_{i \in I}$ is mutually skew. Let k and l be two indices in the set $I \cap [m+1 .. m+n]$. Note that this set has cardinality at least 2, since $w(I) = t \geq 2$. The sets $I \setminus \{k\}$ and $I \setminus \{l\}$ are both m-sets of weight $t-1$. From the assumption $G_2(t-1)$, we conclude that $\dim(\alpha_{I \setminus \{k\}}) = m-1$ and that the hyperplane $\alpha_{I \setminus \{k\}}$ intersects c precisely in the point $C_{I \setminus \{k\}}$; and similarly with k replaced with l. Since the points $C_{I \setminus \{k\}}$ and $C_{I \setminus \{l\}}$ are distinct, it follows that the flat α_I is all of α.

6.2.6 The hard inductive step

We claim also — and this is the heart of the whole proof, and the place where quadrangular sets get involved — that $G_1(t-1)$, $G_1(t)$, $G_2(t-2)$, and $G_2(t-1)$ together imply $G_2(t)$, for any $t \geq 2$.

Let I be any m-set of weight t; we must show that $\dim(\alpha_I) = m-1$ and that $\alpha_I \cap c = \{C_I\}$. Let k and l be any two indices in the set $I \cap [m+1 .. m+n]$, which must exist because $w(I) = t \geq 2$. Let i and j be any two indices in the set $[1 .. m] \setminus I$, which must exist for the same reason. Let P denote the set $P := I \setminus \{k, l\}$, and let Q denote the set $Q := [1 .. m+n] \setminus (I \cup \{i, j\})$. So we have $|P| = m-2$, $|Q| = n-2$, and the set $[1 .. m+n]$ is partitioned into the three parts $[1 .. m+n] = P \cup Q \cup \{i, j, k, l\}$.

Consider the four index sets $P \cup \{i, j, k\}$, $P \cup \{i, j, l\}$, $P \cup \{i, k, l\}$, and $P \cup \{j, k, l\}$. Each of these is an $(m+1)$-set; the first two have weight $t-1$, while the latter two have weight t. From the assumptions $G_1(t-1)$ and $G_1(t)$, we conclude that the flat $\alpha_{P \cup \{i,j,k\}}$ coincides with α, and similarly for the other three index sets. In addition to salting those facts away for the future, it follows from the relation $\alpha_{P \cup \{i,k,l\}} = \alpha$ that the flat $\alpha_I = \alpha_{P \cup \{k,l\}}$ is a hyperplane — that is, $\dim(\alpha_I) = m-1$ — which is one of the claims that we are trying to show.

The set of all $(m-2)$-flats in α that include the fixed $(m-3)$-flat α_P forms a projective plane U. Projecting any point R in $\alpha \setminus \alpha_P$ from the flat α_P produces a unique 'point' $\mathrm{Span}(\alpha_P \cup \{R\})$ in this plane U, and we shall denote that 'point' as $\tilde{R} := \mathrm{Span}(\alpha_P \cup \{R\})$. In particular, we are interested in the four 'points' $\tilde{A}_i$, $\tilde{A}_j$, $\tilde{A}_k$, and $\tilde{A}_l$ in the plane U. The fact that the flat $\alpha_{P \cup \{i,j,k\}}$ coincides with α tells us that the three 'points' $\tilde{A}_i$, $\tilde{A}_j$, and $\tilde{A}_k$ in the plane U are not collinear, and similarly for the other three triples. Thus, the four 'points' $\tilde{A}_i$, $\tilde{A}_j$, $\tilde{A}_k$, and $\tilde{A}_l$ determine a nondegenerate complete quadrangle in the plane U; call that quadrangle D.

Next, we consider the five index sets $P \cup \{i, j\}$, $P \cup \{i, k\}$, $P \cup \{i, l\}$, $P \cup \{j, k\}$, and $P \cup \{j, l\}$, omitting the sixth choice $I = P \cup \{k, l\}$. All six are m-sets. Of the five that we are considering, the first has weight $t-2$, while the others have weight $t-1$. From the assumption $G_2(t-2)$, we conclude that $\dim(\alpha_{P\cup\{i,j\}}) = m-1$ (as we already knew) and that the hyperplane $\alpha_{P\cup\{i,j\}}$ in α cuts the line c at the unique point $C_{P\cup\{i,j\}}$. From the assumption $G_2(t-1)$, we make similar conclusions about the four pairs $\{i, k\}$, $\{i, l\}$, $\{j, k\}$, and $\{j, l\}$. The fact that the hyperplanes $\alpha_{P\cup\{i,k\}}$ and $\alpha_{P\cup\{j,k\}}$ cut the line c at distinct points implies that the $(m-2)$-flat $\alpha_{P\cup\{k\}}$ is skew to c. It follows that the hyperplane $\alpha_I = \alpha_{P\cup\{k,l\}}$ intersects c at a unique point, say X. We need to show that $X = C_I$.

If we project the hyperplane $\alpha_{P\cup\{i,j\}}$ down from the flat α_P into the projective plane U, we deduce that the 'line' $\tilde{A}_i\tilde{A}_j$ in U that joins the 'point' $\tilde{A}_i$ to the 'point' $\tilde{A}_j$ cuts the 'line' $\tilde{c} := \{\tilde{R} \mid R \in c\}$ at the unique 'point' $\tilde{C}_{P\cup\{i,j\}}$. Similar results hold with the pair $\{i, j\}$ replaced by any of the other four pairs. Since the 'line' $\tilde{c}$ cuts five of the six sides of the complete quadrangle D at five distinct 'points', we deduce that $\tilde{c}$ does not pass through any vertex of the quadrangle D; so the quadrangle D cuts out, along $\tilde{c}$, three pairs of 'points' that form a quadrangular set: $\big\{\{\tilde{C}_{P\cup\{i,j\}}, \tilde{X}\}, \{\tilde{C}_{P\cup\{i,k\}}, \tilde{C}_{P\cup\{j,l\}}\}, \{\tilde{C}_{P\cup\{i,l\}}, \tilde{C}_{P\cup\{j,k\}}\}\big\}$.

But there is also a complete quadrangle E sitting in the n-flat β, and we can reason about all six of its sides, since there is no induction going on over there. More precisely, the set of all $(n-2)$-flats in β that include the fixed $(n-3)$-flat β_Q forms a projective plane V. Each point R in $\beta \setminus \beta_Q$ determines a unique 'point' $\bar{R} := \operatorname{Span}(\beta_Q \cup \{R\})$ in the plane V, and the four points $\bar{B}_i$, $\bar{B}_j$, $\bar{B}_k$, and $\bar{B}_l$ are the vertices of a complete quadrangle E in V. The side $\bar{B}_i\bar{B}_j$ of the quadrangle E cuts the 'line' $\bar{c} := \{\bar{R} \mid R \in c\}$ at the 'point' $\bar{C}_{P\cup\{k,l\}}$, and similarly for the other five sides. (Don't be confused by our abbreviated notation; remember that, if we weren't lazy, we would be denoting the point $C_{P\cup\{k,l\}}$ as $C_{P\cup\{k,l\};Q\cup\{i,j\}}$.) It follows that the six points $\big\{\{\bar{C}_{P\cup\{i,j\}}, \bar{C}_{P\cup\{k,l\}}\}, \{\bar{C}_{P\cup\{i,k\}}, \bar{C}_{P\cup\{j,l\}}\}, \{\bar{C}_{P\cup\{i,l\}}, \bar{C}_{P\cup\{j,k\}}\}\big\}$ form a quadrangular set along the line $\bar{c}$. Since five points of a quadrangular set uniquely determine the sixth, we deduce that $X = C_{P\cup\{k,l\}}$, which completes the induction.

6.2.7 Checking the incidences

It remains to show that the A-points and B-points, chosen as above, do represent the budget matroid $B_{m,n}$.

Did we achieve every required incidence? The A-points lie in the m-flat α and the B-points lie in the n-flat β. Also, for every partition $(I; J)$ of $[1 \,..\, m+n]$ into an m-set I and an n-set J, the perfect set of points $\{A_i \mid i \in I\} \cup \{B_j \mid j \in J\}$ is mutually incident, because the flats α_I and β_J both meet the line c at the point C_I.

To show that we have avoided every forbidden incidence, it suffices to check that we have mapped every base of the matroid $B_{m,n}$ to a mutually skew set of $m+n$ points — a set that spans all of S. There are three types of bases. A base of the first type has the form $\{A_i \mid i \in I\} \cup \{B_j \mid j \in J\}$ where I is any $(m-1)$-set and J is any $(n+1)$-set. We have $\beta_J = \beta$, while α_I is some $(m-2)$-flat in α. Note that adding one new element to the set I to produce an m-set $I' \supset I$ gives us a hyperplane $\alpha_{I'}$ in α that cuts the line c at the point $C_{I'}$. Since the point $C_{I'}$ varies, depending upon which m-set $I' \supset I$ we choose, the $(m-2)$-flat α_I must be skew to c. It follows that $\mathrm{Span}(\alpha_I \cup \beta_J) = S$. The second type of base is symmetric, with $m+1$ points in the A column and $n-1$ in the B column. A base of the third and final type has the form $\{A_i \mid i \in I\} \cup \{B_j \mid j \in J'\}$ where I is an m-set and J' is an n-set, but the partitions $(I; J)$ and $(I'; J')$ of $[1\,..\,m+n]$ are distinct. The span of such a base includes the flats α_I and $\beta_{J'}$, so it contains the two distinct points C_I and $C_{I'}$ along c, so it includes the entire line c, so it includes both α and β, so it coincides with S. And that, at long last, finishes the proof of the $B_{m,n}$ Representation Theorem.

Exercise 6.2-3 (For matroid mavens) We have finally finished proving that the budget matroid $B_{m,n}$ is representable over the rationals. If it were also binary (that is, representable over the Galois field of order 2), it would follow from standard results that it was regular (that is, representable over all fields). And indeed, the budget matroids $B_{1,1}$ and $B_{2,1}$ are both binary and regular. But show that those two are the only budget matroids $B_{b_1,\ldots,b_k}$ that are binary.

[Hint: Recall from Exercise 4.2-4 what the circuits of a budget matroid are. It is a standard result that a matroid is binary if and only if the symmetric difference $C \triangle D$ of any two distinct circuits C and D always includes a circuit [41]. Let X be the subset of $B_{b_1,\ldots,b_k}$ of size b that contains, for each j, precisely the top b_j elements in the j^{th} column. Show that one can typically find two distinct elements f and g so that $C := X \cup \{f\}$ and $D := X \cup \{g\}$ are circuits (of the type called ambient circuits in Exercise 4.2-4), and note that the symmetric difference $C \triangle D = \{f, g\}$ is too small to include any circuit. Why are the cases $B_{1,1}$ and $B_{2,1}$ atypical?]

Chapter 7

Representing the matroid $B_{m,1,1}$

As the number of parts in the partition $b = b_1 + \cdots + b_k$ increases and as the parts themselves increase, it rapidly becomes quite unclear whether or not the budget matroid $B_{b_1,\ldots,b_k}$ is representable. But for the simplest of the budget matroids with three columns, those of the form $B_{m,1,1}$, we can still get a general representability result over the rational numbers. Already in this case, however, avoiding forbidden incidences is much harder than it was in the case of the budget matroids $B_{m,n}$, with only two columns.

7.1 The case $m = 2$ of the matroid $B_{2,1,1}$

Let's begin by considering the matroid $B_{2,1,1}$ — the one that we are eventually going to show characterizes the dependence of four cubic polynomials. Recall that a representation of $B_{2,1,1}$ consists of twelve points

$$\begin{array}{ccc} 2 & 1 & 1 \\ \end{array}$$
$$\begin{pmatrix} P_0 & A_0 & B_0 \\ P_1 & A_1 & B_1 \\ P_2 & A_2 & B_2 \\ P_3 & A_3 & B_3 \end{pmatrix}$$

in 3-space. The four A-points lie on a common line a, the four B-points lie on a common line b, and the four P-points lie on a common plane π. Furthermore, each of the twelve perfect sets $\{P_i, P_j, A_k, B_l\}$ for $\{i, j, k, l\} = \{0, 1, 2, 3\}$ must be coplanar. Note that, in this chapter, it is convenient to index the four rows starting at 0, rather than at 1.

Suppose that we choose any plane π in 3-space and four points (P_i) in the plane π with no three collinear. We choose a line a whose intersection with π — call it $A_\pi := a \cap \pi$ — is not on any line joining two of the P-points. And we choose four distinct points (A_i) on the line a, none coinciding with A_π. Whenever $\{i, j, k, l\} = \{0, 1, 2, 3\}$, the point B_l is restricted to lie in the three planes $P_i P_j A_k$, $P_i P_k A_j$, and

$P_k P_j A_i$. By Lemma 6.2-1, those three planes intersect at a unique point B_l, which does not lie in the plane π. But are the four points (B_l) so determined collinear? It turns out that they are; in fact, much more is true.

Note that the A-points did not have to be collinear, in order to enable the construction of the B-points. Given any sequence of four A-points, none lying in π, we get a sequence of four B-points, none lying in π; and the relationship between the sequences (A_i) and (B_i) is symmetric. It turns out that the spans of the two sequences (A_i) and (B_i) always have the same dimension. Furthermore, this result generalizes from 3-space to projective space of any dimension.

7.2 The n-Space Pappus Theorem

Theorem 7.2-1 (n-Space Pappus) *Let S be a projective space of dimension n, and let π be a hyperplane in S. Let $(P_0, \ldots, P_n)$ be $n+1$ points lying in π, but with no n of them mutually incident; so the points (P_i) form a projective frame for the hyperplane π. We say that two sequences $(A_0, \ldots, A_n)$ and $(B_0, \ldots, B_n)$ of points in $S \setminus \pi$ are* compatible *just when, for all $n(n+1)$ pairs (i, j) with i and j in $[0\,..\,n]$ and $i \neq j$, the $n+1$ points*

$$\{P_0, \ldots, \widehat{P_i}, \ldots, \widehat{P_j}, \ldots, P_n, A_i, B_j\}$$

lie in a common hyperplane. There is a unique sequence (B_i) of points in $S \setminus \pi$ that is compatible with any given sequence (A_i) of such points, and we have the equality $\dim(\operatorname{Span}(A_i)) = \dim(\operatorname{Span}(B_i))$.

Note that, when $n = 2$, this theorem is essentially Pappus's Theorem in the plane: Given three distinct points (P_0, P_1, P_2) along a line, the collinearity of the three points (A_0, A_1, A_2) implies the collinearity of the three compatible points (B_0, B_1, B_2), where compatibility means that the triple of points $\{P_i, A_j, B_k\}$ is collinear whenever $\{i, j, k\} = \{0, 1, 2\}$. Thus, this theorem constitutes one possible generalization of Pappus's Theorem from the plane to n-dimensional space.

Proof The existence of a unique sequence (B_i) compatible with a given sequence (A_i) is simple to see. For any j in $[0.\,.n]$, the n points $\{P_0, \ldots, \widehat{P_j}, \ldots, P_n\}$ are mutually skew by assumption, spanning the hyperplane π. Since none of the points (A_i) lies in π, we conclude from Lemma 6.2-1 that the point B_j is uniquely determined by the n constraints that it lie in each of the hyperplanes

$$\operatorname{Span}(P_0, \ldots, \widehat{P_i}, \ldots, \widehat{P_j}, \ldots, P_n, A_i),$$

for i in $[0\,..\,n] \setminus \{j\}$, and the point B_j so determined does not lie in π.

The interesting result is that $\dim(\operatorname{Span}(A_i)) = \dim(\operatorname{Span}(B_i))$. We shall use analytic geometry to show this, choosing our coordinate system carefully. Since the points (P_i) form a projective frame for the hyperplane π, we can choose our coordinate system for points $[x_0, \ldots, x_n]$ in the n-space S so that the hyperplane π has the equation $x_0 = 0$ and so that the homogeneous coordinates of the points (P_i) are given by the rows of the matrix

$$\mathbf{p} := \begin{array}{c} \\ P_0 \\ P_1 \\ P_2 \\ P_3 \\ \vdots \\ P_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & n \end{array} \\ \begin{pmatrix} 0 & 1 & 1 & 1 & \ldots & 1 \\ 0 & 1 & 0 & 0 & \ldots & 0 \\ 0 & 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \ldots & 1 \end{pmatrix} \end{array}.$$

Note that we have chosen P_0 to be the unit point of the projective frame, hence singling it out from the other P-points. Treating the index 0 specially, in this way, builds an asymmetry into our coordinate system, and that seems unavoidable.

Since none of the A-points are allowed to lie in the hyperplane π, we can arrange that their coordinates are given by the rows of the matrix

$$\mathbf{a} := \begin{array}{c} \\ A_0 \\ A_1 \\ A_2 \\ A_3 \\ \vdots \\ A_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & n \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & \ldots & 0 \\ 1 & c_{11} & c_{12} & c_{13} & \ldots & c_{1n} \\ 1 & c_{21} & c_{22} & c_{23} & \ldots & c_{2n} \\ 1 & c_{31} & c_{32} & c_{33} & \ldots & c_{3n} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & c_{n1} & c_{n2} & c_{n3} & \ldots & c_{nn} \end{pmatrix} \end{array}.$$

Giving the point A_0 the simple coordinates $[1, 0, 0, \ldots, 0]$ builds a second asymmetry into our chosen coordinate system, an asymmetry between the A-points and the upcoming B-points. This second asymmetry could be eliminated, but only by splitting the difference, which would introduce annoying factors of $\frac{1}{2}$.

Claim: The coordinates of the compatible points (B_j) are then given by the rows of the matrix $\mathbf{b}$:

$$\mathbf{b} := \begin{array}{c} \\ B_0 \\ B_1 \\ B_2 \\ B_3 \\ \vdots \\ B_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & n \end{array} \\ \begin{pmatrix} 1 & c_{11} & c_{22} & c_{33} & \ldots & c_{nn} \\ 1 & 0 & c_{22}-c_{21} & c_{33}-c_{31} & \ldots & c_{nn}-c_{n1} \\ 1 & c_{11}-c_{12} & 0 & c_{33}-c_{32} & \ldots & c_{nn}-c_{n2} \\ 1 & c_{11}-c_{13} & c_{22}-c_{23} & 0 & \ldots & c_{nn}-c_{n3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & c_{11}-c_{1n} & c_{22}-c_{2n} & c_{33}-c_{3n} & \ldots & 0 \end{pmatrix} \end{array}$$

To verify this claim, let (i, j) be any pair with i and j in $[0 \,..\, n]$ and $i \neq j$; we must show that the points

$$\{P_0, \ldots, \widehat{P_i}, \ldots, \widehat{P_j}, \ldots, P_n, A_i, B_j\}$$

lie in a common hyperplane — that is, that the matrix formed by replacing the i^{th} and j^{th} rows of $\mathbf{p}$ with A_i and B_j has determinant zero. If neither i nor j is 0, this boils down to checking that the matrix

$$\begin{array}{cc} & \begin{array}{ccc} 0 & i & j \end{array} \\ \begin{array}{c} P_0 \\ A_i \\ B_j \end{array} & \begin{pmatrix} 0 & 1 & 1 \\ 1 & c_{ii} & c_{ij} \\ 1 & c_{ii} - c_{ij} & 0 \end{pmatrix} \end{array}$$

has zero determinant, which it does; the common hyperplane in this case has the homogeneous coefficients

$$\langle \overset{0}{c_{ij} - c_{ii}}, 0, \ldots, 0, \overset{i}{1}, 0, \ldots, 0, \overset{j}{-1}, 0, \ldots, 0 \rangle.$$

The special cases $i = 0$ and $j = 0$ are even easier:

$$\begin{array}{cc} & \begin{array}{cc} 0 & j \end{array} \\ \begin{array}{c} A_0 \\ B_j \end{array} & \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \end{array} \quad \text{and} \quad \begin{array}{cc} & \begin{array}{cc} 0 & i \end{array} \\ \begin{array}{c} A_i \\ B_0 \end{array} & \begin{pmatrix} 1 & c_{ii} \\ 1 & c_{ii} \end{pmatrix} \end{array}.$$

It remains to show that $\operatorname{rank}(\mathbf{a}) = \operatorname{rank}(\mathbf{b})$, which we can do by converting one matrix into the transpose of the other using elementary row and column operations. Start with the matrix $\mathbf{b}$. Subtracting the top row from each row below it gives us

$$\begin{pmatrix} 1 & c_{11} & c_{22} & c_{33} & \cdots & c_{nn} \\ 0 & -c_{11} & -c_{21} & -c_{31} & \cdots & -c_{n1} \\ 0 & -c_{12} & -c_{22} & -c_{32} & \cdots & -c_{n2} \\ 0 & -c_{13} & -c_{23} & -c_{33} & \cdots & -c_{n3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & -c_{1n} & -c_{2n} & -c_{3n} & \cdots & -c_{nn} \end{pmatrix}.$$

Adding appropriate multiples of the leftmost column to each column to its right then gives us the matrix

$$\begin{pmatrix} 1 & -1 & -1 & -1 & \cdots & -1 \\ 0 & -c_{11} & -c_{21} & -c_{31} & \cdots & -c_{n1} \\ 0 & -c_{12} & -c_{22} & -c_{32} & \cdots & -c_{n2} \\ 0 & -c_{13} & -c_{23} & -c_{33} & \cdots & -c_{n3} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & -c_{1n} & -c_{2n} & -c_{3n} & \cdots & -c_{nn} \end{pmatrix},$$

which differs from the transpose of $\mathbf{a}$ only by some unimportant signs. $\square$

Exercise 7.2-2 Show that the three budgetary relaxations $B^{2,1,2}_{2,1,1}$, $B^{2,1,3}_{2,1,1}$, and $B^{2,2,3}_{2,1,1}$ of the budget matroid $B_{2,1,1}$ are not representable over any field. Recall that the first two of these were shown to be unrepresentable also in Exercise 5.6-3.

[Answer: Letting $n = 3$ in the n-Space Pappus Theorem tells us at once that the column dimensions of the A and B columns must be the same.]

Exercise 7.2-3 What about those budgetary relaxations $B^{3,?,?}_{2,1,1}$ of the budget matroid $B_{2,1,1}$ in which the P column dimension is 3 — that is, the P-points are not coplanar? Note that the twelve perfect coplanarities still hold, and they still suffice to determine the B-points from the P-points and the A-points. Find the homogeneous coordinates of the B-points, assuming that the P-points are given by the rows of the matrix

$$\mathbf{p} := \begin{array}{c} P_1 \\ P_2 \\ P_3 \\ P_4 \end{array} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

while the A-points are given by

$$\mathbf{a} := \begin{array}{c} A_1 \\ A_2 \\ A_3 \\ A_4 \end{array} \begin{pmatrix} a_{11} & a_{12} & a_{13} & a_{14} \\ a_{21} & a_{22} & a_{23} & a_{24} \\ a_{31} & a_{32} & a_{33} & a_{34} \\ a_{41} & a_{42} & a_{43} & a_{44} \end{pmatrix}.$$

[Answer: The B-points are given by the rows of the matrix

$$\mathbf{b} := \begin{array}{c} B_1 \\ B_2 \\ B_3 \\ B_4 \end{array} \begin{pmatrix} a_{21}a_{31}a_{41} & a_{22}a_{31}a_{41} & a_{21}a_{33}a_{41} & a_{21}a_{31}a_{44} \\ a_{11}a_{32}a_{42} & a_{12}a_{32}a_{42} & a_{12}a_{33}a_{42} & a_{12}a_{32}a_{44} \\ a_{11}a_{23}a_{43} & a_{13}a_{22}a_{43} & a_{13}a_{23}a_{43} & a_{13}a_{23}a_{44} \\ a_{11}a_{24}a_{34} & a_{14}a_{22}a_{34} & a_{14}a_{24}a_{33} & a_{14}a_{24}a_{34} \end{pmatrix}. \Bigg]$$

Exercise 7.2-4 Using the preceding exercise, show that the budgetary matroids $B^{3,1,1}_{2,1,1}$ and $B^{3,1,2}_{2,1,1}$ are not representable over any field. Note that the first of these was shown to be unrepresentable also in Exercise 5.6-3.

For context, recall that $\#\big(B^{3,1,3}_{2,1,1}\big) = 20$: three degrees of freedom in each of the four P-points, three each in A_1 and A_2, and one each in A_3 and A_4, after which the B-points are determined.

[Hint: Replace the points A_3 and A_4, which are generic points in 3-space in the previous exercise, with generic points on the line A_1A_2 — say, $A_3 = u_3A_1 + v_3A_2$ and $A_4 = u_4A_1 + v_4A_2$ — and verify that $\det(\mathbf{b})$ then factors as

$$\det(\mathbf{b}) = u_3u_4v_3v_4(u_3v_4 - u_4v_3) \prod_{1\le i<j\le 4} (a_{1i}a_{2j} - a_{1j}a_{2i}).$$

Thus, the B-points can be coplanar only when there is some degeneracy: either two of the A-points coincide or the line containing the four A-points intersects one of the six edges of the P-point tetrahedron.]

7.3 The general case

Theorem 7.3-1 ($B_{m,1,1}$ Representation) *For any positive m, the budget matroid $B_{m,1,1}$ is representable over the rational numbers, and the number of degrees of freedom involved in choosing a representation, lying in a fixed projective space S of dimension $m+1$, is $\#(B_{m,1,1}) = m^2 + 6m + 3$.*

The proof of this theorem, like the proof of the $B_{m,n}$ Representation Theorem, is long. The good news is that you aren't as likely to get bogged down this time. The bad news is that we are going to use the concepts introduced in this proof, in Chapter 9, to analyze what can go wrong when constructing a representation of the matroid $B_{2,1,1}$ that witnesses to the 3-dependence of four cubic polynomials. So skipping this proof will limit the depth of your understanding of Chapter 9.

7.3.1 The plan of attack

We use n as an abbreviation for $m+1$. We index the three columns with the letters P, A, and B and the $m+2$ rows with the integers $[0\,..\,m+1] = [0\,..\,n]$.

The basic plan of attack is obvious. We choose any hyperplane π in S; that involves $m+1$ degrees of freedom. We choose the points (P_i) for i in $[0\,..\,m+1]$ to form a projective frame for π; that involves $m(m+2)$ degrees of freedom. We choose some line a, not lying in the hyperplane π — another $2m$ degrees of freedom. Finally, we choose $m+2$ points (A_i) on a, none lying in π, which involves a final $m+2$ degrees of freedom, for a total of m^2+6m+3. Invoking the n-Space Pappus Theorem with $n := m+1$, we let (B_i) be the unique sequence of points that is compatible with the sequence (A_i), and we observe that the points (B_i) will be collinear. The resulting P-points, A-points, and B-points have all of the incidences required for a representation of $B_{m,1,1}$. The hard part is to show that they don't have any forbidden incidences. Indeed, they will have forbidden incidences in special cases; but they won't in the generic case.

7.3.2 Choosing the coordinate system

As in the proof of the n-Space Pappus Theorem, we can choose our coordinate system on the space S so that the homogeneous coordinates of the P-points are given by the rows of the matrix

$$\mathbf{p} := \begin{array}{c} \\ P_0 \\ P_1 \\ P_2 \\ P_3 \\ \vdots \\ P_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \dots & n \end{array} \\ \begin{pmatrix} 0 & 1 & 1 & 1 & \dots & 1 \\ 0 & 1 & 0 & 0 & \dots & 0 \\ 0 & 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \dots & 1 \end{pmatrix} \end{array},$$

while the point A_0 has the coordinates $A_0 = [1, 0, 0, \ldots, 0]$. When proving the n-Space Pappus Theorem, the point A_i had the coordinates $A_i = [1, c_{i1}, \ldots, c_{in}]$, for i in $[1 \mathrel{..} n]$. But we here want to constrain the points (A_i) to lie on a common line a. One convenient way to do that is as follows.

Let $A_\pi := a \cap \pi$ be the point where the line a cuts the hyperplane π, and let $[0, v_1, \ldots, v_n] = A_\pi$ be some set of homogeneous coordinates for the point A_π. Now, the point A_1 lies on the line $a = A_0 A_\pi$. Hence, for any finite, nonzero scalar u_1 that we choose, there is a unique projective coordinate system for the line a in which the three distinct points A_0, A_π, and A_1 are assigned the coordinates 0, ∞, and u_1. In that coordinate system for the line a, the points A_2 through A_n have finite, nonzero coordinates, which we shall denote u_2 through u_n.

There is a unique way to extend our chosen coordinate system for the hyperplane π to a coordinate system for the entire space S so that, in addition to the point A_0 getting the coordinates $[1, 0, 0, \ldots, 0]$, the point A_1 gets the coordinates $[1, u_1 v_1, \ldots, u_1 v_n]$. In that unique coordinate system for S, the point A_i, for i in $[2 \mathrel{..} n]$, must have the coordinates $[1, u_i v_1, \ldots, u_i v_n]$, where the scalars (u_i) are as defined above. So the homogeneous coordinates of the A-points are given by the rows of the matrix

$$\mathbf{a} := \begin{array}{c} \\ A_0 \\ A_1 \\ A_2 \\ A_3 \\ \vdots \\ A_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & n \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & \ldots & 0 \\ 1 & u_1 v_1 & u_1 v_2 & u_1 v_3 & \ldots & u_1 v_n \\ 1 & u_2 v_1 & u_2 v_2 & u_2 v_3 & \ldots & u_2 v_n \\ 1 & u_3 v_1 & u_3 v_2 & u_3 v_3 & \ldots & u_3 v_n \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & u_n v_1 & u_n v_2 & u_n v_3 & \ldots & u_n v_n \end{pmatrix} \end{array}.$$

Note that, for i in $[1 \mathrel{..} n]$, we have the inhomogeneous equation $A_i = A_0 + u_i A_\pi$.

Given that those are the A-points, the same argument as in the proof of the n-Space Pappus Theorem shows that the coordinates of the B-points are given by the rows of the matrix

$$\mathbf{b} := \begin{array}{c} \\ B_0 \\ B_1 \\ B_2 \\ B_3 \\ \vdots \\ B_n \end{array} \begin{array}{c} \begin{array}{cccccc} 0 & 1 & 2 & 3 & \ldots & n \end{array} \\ \begin{pmatrix} 1 & u_1 v_1 & u_2 v_2 & u_3 v_3 & \ldots & u_n v_n \\ 1 & 0 & u_2(v_2 - v_1) & u_3(v_3 - v_1) & \ldots & u_n(v_n - v_1) \\ 1 & u_1(v_1 - v_2) & 0 & u_3(v_3 - v_2) & \ldots & u_n(v_n - v_2) \\ 1 & u_1(v_1 - v_3) & u_2(v_2 - v_3) & 0 & \ldots & u_n(v_n - v_3) \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & u_1(v_1 - v_n) & u_2(v_2 - v_n) & u_3(v_3 - v_n) & \ldots & 0 \end{pmatrix} \end{array}.$$

Note that these points (B_i) are indeed collinear. Letting B_π denote the point $B_\pi := [0, u_1, \ldots, u_n]$ in the hyperplane π, the point B_i, for i in $[1 \mathrel{..} n]$, lies on the line $b := B_0 B_\pi$, at a position determined by the parameter v_i. In fact, we have the inhomogeneous equation $B_i = B_0 - v_i B_\pi$.

Just as in the proof of the n-Space Pappus Theorem, our chosen coordinate system treats the A-points and the B-points asymmetrically. But that asymmetry manifests itself only in the coordinates that are assigned to points outside of the hyperplane π. In particular, the asymmetry does not affect the parameters (u_i) and (v_i). We have $A_\pi = a \cap \pi = [0, v_1, \ldots, v_m]$ and, symmetrically, we have $B_\pi = b \cap \pi = [0, u_1, \ldots, u_m]$. Alternatively, we can define the scalars $(u_1, \ldots, u_m)$ as the coordinates of the points $(A_1, \ldots, A_m)$ in some projective coordinate system for the line a that assigns the coordinates 0 and ∞ to A_0 and A_π. Symmetrically, the scalars $(v_1, \ldots, v_m)$ are the coordinates of the points $(B_1, \ldots, B_m)$ in some coordinate system for the line b that assigns the coordinates 0 and ∞ to B_0 and B_π.

Each of the parameter vectors $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$ is homogeneous, and thus represents $n - 1 = m$ degrees of freedom. If we think of the parameters (u_i) as the homogeneous coordinates of the point B_π, they are obviously homogeneous. Of course, we introduced the scalar u_i, instead, as the coordinate of the point A_i in some coordinate system for the line a that assigns coordinates 0 and ∞ to A_0 and A_π. But the parameters (u_i) are also homogeneous when thought of in that way: The ratio $u_i : u_j$ is determined by the cross ratio (A_i, A_j, A_0, A_π), but we can rescale all n parameters $(u_1, \ldots, u_n)$ without changing anything.

Note that the two parameter vectors $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$ determine the entire configuration, up to a projective transformation of the ambient space S. The m degrees of freedom in each of the homogeneous parameter vectors combine with the $(m+2)^2 - 1 = m^2 + 4m + 3$ degrees of freedom in a projective transformation of the $(m + 1)$-space S to account for the grand total of $m^2 + 6m + 3$ degrees of freedom.

7.3.3 The residual matrix of a base

We claim that, for generic choices of the two parameter vectors $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$, the resulting configuration will be free of forbidden incidences and will hence represent the budget matroid $B_{m,1,1}$. It suffices to consider each base of $B_{m,1,1}$ and to rule out pairs of parameter vectors that make the $m+2$ points in that base mutually incident. Our hope is that each base will rule out only those pairs of parameter vectors that satisfy some algebraic relation. The bad thing — which we must show does not happen — would be for one of the bases to rule out all pairs of parameter vectors.

We can classify the bases of the matroid $B_{m,1,1}$ by the number of points that they have in each column. Since no independent set is allowed to have more than $b_j + 1$ points in the j^{th} column, no base can contain more than two A-points or more than two B-points. Thus, for m large, the bulk of the points in any base are P-points. We test the points of a base for incidence by forming an $(m+2)$-by-$(m+2)$ matrix whose rows are their coordinates and testing the determinant for zero. Since the bulk of the points in any base are P-points, the bulk of the rows of this matrix will have the form $[0, \ldots, 0, 1, 0, \ldots, 0]$, with only one nonzero entry. When computing the determinant, we can delete each such row and delete each column in which any such row has its single 1. What remains — let's call it the *residual matrix* — is of bounded size.

7.3.4 The determinants of the residual matrices

To cut down on the number of cases in what follows, let's introduce $u_0 := 0$ and $v_0 := 0$ as alternative names for zero. We can then rewrite the matrices $\mathbf{a}$ and $\mathbf{b}$ so that the top row looks just like all of the rest:

$$\mathbf{a} = \begin{pmatrix} 1 & u_0v_1 & u_0v_2 & u_0v_3 & \ldots & u_0v_n \\ 1 & u_1v_1 & u_1v_2 & u_1v_3 & \ldots & u_1v_n \\ 1 & u_2v_1 & u_2v_2 & u_2v_3 & \ldots & u_2v_n \\ 1 & u_3v_1 & u_3v_2 & u_3v_3 & \ldots & u_3v_n \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & u_nv_1 & u_nv_2 & u_nv_3 & \ldots & u_nv_n \end{pmatrix}$$

and

$$\mathbf{b} = \begin{pmatrix} 1 & u_1(v_1 - v_0) & u_2(v_2 - v_0) & u_3(v_3 - v_0) & \ldots & u_n(v_n - v_0) \\ 1 & u_1(v_1 - v_1) & u_2(v_2 - v_1) & u_3(v_3 - v_1) & \ldots & u_n(v_n - v_1) \\ 1 & u_1(v_1 - v_2) & u_2(v_2 - v_2) & u_3(v_3 - v_2) & \ldots & u_n(v_n - v_2) \\ 1 & u_1(v_1 - v_3) & u_2(v_2 - v_3) & u_3(v_3 - v_3) & \ldots & u_n(v_n - v_3) \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & u_1(v_1 - v_n) & u_2(v_2 - v_n) & u_3(v_3 - v_n) & \ldots & u_n(v_n - v_n) \end{pmatrix}.$$

Using the convention that $u_0 = v_0 = 0$, we can capture the determinants of the residual matrices associated with all possible bases in Table 7.1. Each of the eight rows in this table describes a possible vector of column populations for a base. For example, in the third row, any m of the P-points and any two of the A-points together form a base. The fourth row is special, in that, when choosing m of the P-points, one A-point, and one B-point, we must avoid choosing precisely one point with each of the possible indices $[0\,..\,m+1]$ — that would give us a perfect set, which is not a base. We specify which A-points and which B-points are chosen by listing their indices, in the obvious way. But, for P-points, we use the process of elimination, listing instead the indices of the points that are *not* chosen. Claim: In each row of Table 7.1, the formula in the rightmost column gives the determinant of the $(m+2)$-by-$(m+2)$ matrix formed by assembling the coordinate vectors of the $m+2$ chosen points.

Let's verify the determinant in the seventh row, the case $(m-1, 1, 2)$, as an example; the other cases are similar. We form a base by taking all of the P-points

#P's {left}	#A's {taken}	#B's {taken}	determinant of residual matrix
$m+1$ $\{e\}$	1 $\{i\}$	0 $\{\}$	± 1
$m+1$ $\{e\}$	0 $\{\}$	1 $\{k\}$	± 1
m $\{e, f\}$	2 $\{i, j\}$	0 $\{\}$	$\pm(v_f - v_e)(u_j - u_i)$
m $\{e, f\}$	1 $\{i\}$	1 $\{k\}$	$\pm \begin{vmatrix} 1 & u_e & v_f \\ 1 & u_f & v_e \\ 1 & u_i & v_k \end{vmatrix}$
m $\{e, f\}$	0 $\{\}$	2 $\{k, l\}$	$\pm(u_f - u_e)(v_l - v_k)$
$m-1$ $\{e, f, g\}$	2 $\{i, j\}$	1 $\{k\}$	$\pm(u_j - u_i) \begin{vmatrix} 1 & v_e & u_e(v_e - v_k) \\ 1 & v_f & u_f(v_f - v_k) \\ 1 & v_g & u_g(v_g - v_k) \end{vmatrix}$
$m-1$ $\{e, f, g\}$	1 $\{i\}$	2 $\{k, l\}$	$\pm(v_k - v_l) \begin{vmatrix} 1 & u_e & v_e(u_e - u_i) \\ 1 & u_f & v_f(u_f - u_i) \\ 1 & u_g & v_g(u_g - u_i) \end{vmatrix}$
$m-2$ $\{e, f, g, h\}$	2 $\{i, j\}$	2 $\{k, l\}$	$\pm(u_j - u_i)(v_l - v_k) \begin{vmatrix} 1 & u_e & v_e & u_e v_e \\ 1 & u_f & v_f & u_f v_f \\ 1 & u_g & v_g & u_g v_g \\ 1 & u_h & v_h & u_h v_h \end{vmatrix}$

Table 7.1: Bases in a potential representation of the budget matroid $B_{m,1,1}$

except for P_e, P_f, and P_g and combining them with A_i and with B_k and B_l. We can assume that the indices in each group are sorted, with $e < f < g$ and $k < l$. Because of our convention that $u_0 = v_0 = 0$, we don't have to single out the subcases $i = 0$ and $k = 0$ as special. But the subcase $e = 0$ might still be special, so far as we know. When $e > 0$, the residual matrix is 4-by-4:

$$\begin{array}{c} \\ P_0 \\ A_i \\ B_k \\ B_l \end{array}
\begin{array}{c}
\begin{array}{cccc} 0 & e & f & g \end{array} \\
\begin{pmatrix}
0 & 1 & 1 & 1 \\
1 & u_i v_e & u_i v_f & u_i v_g \\
1 & u_e(v_e - v_k) & u_f(v_f - v_k) & u_g(v_g - v_k) \\
1 & u_e(v_e - v_l) & u_f(v_f - v_l) & u_g(v_g - v_l)
\end{pmatrix}
\end{array} \tag{7.3-2}$$

It is clear that the determinant of this matrix is divisible by the difference $v_l - v_k$, and it is straightforward to check that the quotient is the 3-by-3 determinant

$$\begin{vmatrix}
1 & u_e & v_e(u_e - u_i) \\
1 & u_f & v_f(u_f - u_i) \\
1 & u_g & v_g(u_g - u_i)
\end{vmatrix}$$

given in Table 7.1. In the subcase $e = 0$, the residual matrix is only 3-by-3:

$$\begin{array}{c} \\ A_i \\ B_k \\ B_l \end{array}
\begin{array}{c}
\begin{array}{ccc} e=0 & f & g \end{array} \\
\begin{pmatrix}
1 & u_i v_f & u_i v_g \\
1 & u_f(v_f - v_k) & u_g(v_g - v_k) \\
1 & u_f(v_f - v_l) & u_g(v_g - v_l)
\end{pmatrix}
\end{array} \tag{7.3-3}$$

But now the convention $u_0 = v_0 = 0$ comes to our rescue once again. Note that, if we set $e = 0$ in Matrix 7.3-2, the e^{th} column becomes $[1, 0, 0, 0]$, which means that we can delete the e^{th} column and the 0^{th} row. What is left is precisely Matrix 7.3-3. Therefore, the formula given for the case $(m - 1, 1, 2)$ in Table 7.1 works also in the special subcase $e = 0$. By similar arguments, all of the formulas in Table 7.1 work also when $e = 0$.

7.3.5 The six primitive degeneracies

It remains to verify, using the formulas in Table 7.1, that for no base of the matroid $B_{m,1,1}$ does the determinant of the residual matrix, viewed as a polynomial in the variables $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$, reduce to zero. Since we have to study each polynomial a little, to verify that it isn't zero, we might as well calculate what its irreducible factors are. Each irreducible factor of such a polynomial encodes a *primitive degeneracy* that we must avoid, by being careful when we choose the parameter vectors $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$. In fact, we might as well build a catalog of the various types of primitive degeneracies. It turns out that there are six types: six bad ways in which the positions of the A-points and the B-points, along the lines a and b, can be related to each other.

We ruled out the first type of primitive degeneracy a while back, before we thought about building a catalog. We shall say that a degeneracy of *Type 1a* arises when the point A_i lies in the hyperplane π, for some i in $[0\,..\,n]$, that is, when the points A_i and A_π coincide. Similarly, if $B_i = B_\pi$, then we have a degeneracy of *Type 1b*. Note that, as A_i approaches A_π, the parameter u_i approaches ∞. Thus, we have ruled out all degeneracies of Type 1 as soon as we assume that the parameters (u_i) and (v_i) are all finite.

A degeneracy of *Type 2a* arises when $A_i = A_j$ for some $i \neq j$ in $[0\,..\,n]$, and of *Type 2b* when $B_i = B_j$. Recall that $A_i = A_j$ just when $u_i = u_j$, so degeneracies of Type 2a correspond to linear factors of the form $u_j - u_i$, while those of Type 2b correspond to $v_j - v_i$. Such linear factors arise frequently in Table 7.1.

We now tackle Table 7.1, case by case. Four of the eight are easy. In the cases $(m+1, 1, 0)$ and $(m+1, 0, 1)$, the determinant of the residual matrix is ± 1, so there are no factors to consider. In the cases $(m, 2, 0)$ and $(m, 0, 2)$, there are factors, but they are all of the form $u_j - u_i$ or $v_j - v_i$, and hence reflect Type-2 degeneracies.

The case $(m, 1, 1)$ is the trickiest to analyze. Table 7.1 gives the formula

$$\begin{vmatrix} 1 & u_e & v_f \\ 1 & u_f & v_e \\ 1 & u_i & v_k \end{vmatrix} = u_f v_k - u_i v_e + u_i v_f - u_f v_f + u_e v_e - u_e v_k.$$

Note that, if $\{i, k\} = \{e, f\}$, this determinant does vanish. But that is all right, because the corresponding set of points is perfect, and hence is not a base. What other subcases are there? If $i = e$, two of the six terms in the determinant cancel, and the four that remain factor as the product $(u_f - u_e)(v_k - v_f)$. Since $i = e$ but $\{i, k\} \neq \{e, f\}$, we must have $k \neq f$; so both of these factors reflect Type-2 degeneracies. A similar analysis handles the subcases $i = f$, $k = e$, and $k = f$; so let us assume that $\{i, k\} \cap \{e, f\} = \emptyset$. The only remaining issue is whether i and k coincide or not. In either subcase, however, there is no cancellation: the six-term determinant is nonzero and irreducible, encoding a primitive degeneracy that we can and must avoid. We shall refer to these two types of degeneracies as *Type 3* and *Type 4*, according as $i = k$ or $i \neq k$.

Let's consider the case $(m - 1, 2, 1)$ next. Table 7.1 gives the determinant

$$\begin{vmatrix} 1 & v_e & u_e(v_e - v_k) \\ 1 & v_f & u_f(v_f - v_k) \\ 1 & v_g & u_g(v_g - v_k) \end{vmatrix},$$

multiplied by the factor of a Type-2 degeneracy, which we ignore. The issue here is whether k belongs to the set $\{e, f, g\}$. If $k = e$, for example, four of the twelve terms of the determinant cancel, and the eight that remain factor as the product of three Type-2 factors: $(u_f - u_g)(v_f - v_e)(v_g - v_e)$. A similar analysis handles the subcases $k = f$ and $k = g$. When $k \notin \{e, f, g\}$, there is no cancellation and the twelve-term determinant is irreducible, so it encodes another type of primitive degeneracy that we must avoid — which we shall call *Type 5a*.

The case $(m-1, 1, 2)$ is symmetric, including a twelve-term degeneracy of a type just like Type 5a, but with A and B interchanged. We shall call it *Type 5b*.

Finally, we come to the case $(m-2, 2, 2)$, for which Table 7.1 gives, ignoring Type-2 factors, the determinant

$$\begin{vmatrix} 1 & u_e & v_e & u_e v_e \\ 1 & u_f & v_f & u_f v_f \\ 1 & u_g & v_g & u_g v_g \\ 1 & u_h & v_h & u_h v_h \end{vmatrix}.$$

The four indices e, f, g, and h must be distinct, so no splitting into subcases is needed. The determinant has 24 terms and is irreducible; we shall refer to the degeneracies that result as *Type 6*.

Since none of the residual matrices has zero as its determinant, it follows that generic choices of the parameter vectors $[u_1, \ldots, u_n]$ and $[v_1, \ldots, v_n]$ will avoid all degeneracies, which completes the proof of $B_{m,1,1}$ Representation Theorem.

7.4 The primitive degeneracies geometrically

In the process of proving the $B_{m,1,1}$ Representation Theorem, we built a catalog of the six types of primitive degeneracies: the six things to watch out for, when deciding where to put the A-points and the B-points along their lines. Each primitive degeneracy can be defined either algebraically, in terms of the parameters (u_i) and (v_i), or geometrically, in terms of the A-points and B-points. Our goal in this section is to find the geometric characterizations for Types 3 through 6.

Given some representation of the budget matroid $B_{m,1,1}$, recall that one way to define the associated homogeneous parameters $[u_1, \ldots, u_{m+1}]$ is as the coordinates of the points $(A_1, \ldots, A_{m+1})$ in some projective coordinate system for the line a that gives the points A_0 and $A_\pi := a \cap \pi$ the coordinates 0 and ∞. Similarly, the parameters $[v_1 \ldots, v_{m+1}]$ are the coordinates of the points $(B_1, \ldots, B_{m+1})$ along the line b. But let's focus, not on the coordinate systems, but on the sequences of points $(A_0, A_1, \ldots, A_{m+1}; A_\pi)$ and $(B_0, B_1, \ldots, B_{m+1}; B_\pi)$ themselves. Note that each sequence consists of $m+3$ distinct, collinear points. If we fix any three of the points in the A sequence, the cross ratios of the other m points with respect to the fixed three constitute m degrees of freedom — the same m degrees of freedom that we have been dealing with algebraically by using the homogeneous coordinates $[u_1, \ldots, u_{m+1}]$ and $[v_1, \ldots, v_{m+1}]$.

Each type of primitive degeneracy can be described either algebraically or geometrically. For Type 1a, we have $u_i = \infty$ just when $A_i = A_\pi$. For Type 2a, we have $u_i = u_j$ just when $A_i = A_j$. In a similar way, the algebraic definitions of the degeneracies of Types 3 through 6 in Section 7.3.5 correspond to projectively invariant geometric properties of the A-points and B-points — in fact, to the condition that some four of the $m+3$ points $(A_0, \ldots, A_{m+1}; A_\pi)$ have the same cross ratio as some four of the points $(B_0, \ldots, B_{m+1}; B_\pi)$. To make it more convenient to express such a cross ratio, if i_1 through i_4 are any four distinct indices in the set $[0 \,..\, m+1] \cup \{\pi\}$, let us denote by $A_{(i_1,i_2,i_3,i_4)}$ the cross ratio of the four points $(A_{i_1}, A_{i_2}, A_{i_3}, A_{i_4})$ along the line a, and similarly for $B_{(i_1,i_2,i_3,i_4)}$.

Proposition 7.4-1 *For some positive m and for i in $[0 \,..\, m+1]$, suppose that the points (P_i), (A_i), and (B_i), all lying in some projective $(m+1)$-space S, come at least this close to representing the budget matroid $B_{m,1,1}$:*

- *The P-points form a projective frame for a hyperplane π in S.*
- *The A-points lie along a line a that intersects the hyperplane π in the single point $A_\pi := a \cap \pi$, and they are distinct from each other and from A_π.*
- *Similarly, the B-points lie along a line b that intersects π in the single point $B_\pi := b \cap \pi$, and they are distinct from each other and from B_π.*
- *The $(m+1)(m+2)$ perfect sets are mutually incident.*

The only degeneracies that might still arise are of the following types:

Type	*Cross ratios*	*Forbidden incidence*
3	$A_{(i,j,k,\pi)} = B_{(i,j,\pi,k)}$	$\{A_k, B_k\} \cup \mathcal{P} \setminus \{P_i, P_j\}$
4	$A_{(i,j,k,\pi)} = B_{(i,j,\pi,l)}$	$\{A_k, B_l\} \cup \mathcal{P} \setminus \{P_i, P_j\}$
5a	$A_{(i,j,k,\pi)} = B_{(i,j,k,l)}$	$\{A_*, A_*, B_l\} \cup \mathcal{P} \setminus \{P_i, P_j, P_k\}$
5b	$A_{(i,j,k,l)} = B_{(i,j,k,\pi)}$	$\{A_l, B_*, B_*\} \cup \mathcal{P} \setminus \{P_i, P_j, P_k\}$
6	$A_{(i,j,k,l)} = B_{(i,j,k,l)}$	$\{A_*, A_*, B_*, B_*\} \cup \mathcal{P} \setminus \{P_i, P_j, P_k, P_l\}$

The symbol $\mathcal{P} := \{P_0, \ldots, P_{m+1}\}$ denotes the set of all P-points. The symbol A_* means any A-point, and similarly for B_*. The indices i, j, k, and l in $[0 \,..\, m+1]$ must be distinct in Types 4 through 6, while i, j, and k must be distinct in Type 3.

Proof We pick up where the proof of the $B_{m,1,1}$ Representation Theorem left off. The points $(A_0, A_1, \ldots, A_{m+1}; A_\pi)$ have coordinates $(u_0 = 0, u_1, \ldots, u_{m+1}; \infty)$ in a valid, projective coordinate system for the line a. In a similar way, the points $(B_0, B_1, \ldots, B_{m+1}; B_\pi)$ have coordinates $(v_0 = 0, v_1, \ldots, v_{m+1}; \infty)$ along b. It is well known that, if the four points X_1 through X_4 have the coordinates x_1 through x_4 along one line while Y_1 through Y_4 have the coordinates y_1 through y_4 along some other line, the cross ratios $X_{(1,2,3,4)}$ and $Y_{(1,2,3,4)}$ are equal just when the determinant

$$\begin{vmatrix} 1 & x_1 & y_1 & x_1y_1 \\ 1 & x_2 & y_2 & x_2y_2 \\ 1 & x_3 & y_3 & x_3y_3 \\ 1 & x_4 & y_4 & x_4y_4 \end{vmatrix}$$

vanishes. Armed with this, it is straightforward to verify that the algebraic conditions that defined the degeneracies of Types 3 through 6 in Section 7.3.5 correspond to the geometric conditions listed above. □

By the way, in proving the $B_{m,1,1}$ Representation Theorem, we actually proved somewhat more than we claimed. We are now in a position to state that stronger result geometrically.

Proposition 7.4-2 *For some positive m, let $(S_0, \ldots, S_{m+1}; S_\pi)$ be a sequence of $m+3$ distinct points on a line s, let $(T_0, \ldots, T_{m+1}; T_\pi)$ be a similar sequence on a line t, and suppose that those two sequences are free of all degeneracies of Types 3 through 6. That is, for no three distinct indices i, j, and k in $[0\,..\,m+1]$ do we have $S_{(i,j,k,\pi)} = T_{(i,j,\pi,k)}$, and for no four distinct indices i, j, k, and l in $[0\,..\,m+1]$ do we have either $S_{(i,j,k,\pi)} = T_{(i,j,\pi,l)}$ or $S_{(i,j,k,\pi)} = T_{(i,j,k,l)}$ or $S_{(i,j,k,l)} = T_{(i,j,k,\pi)}$ or $S_{(i,j,k,l)} = T_{(i,j,k,l)}$. Then, there exist representations of the budget matroid $B_{m,1,1}$ in which the points $(A_0, \ldots, A_{m+1}; A_\pi)$ and $(B_0, \ldots, B_{m+1}; B_\pi)$ are projective images of these S-points and T-points. Furthermore, any two such representations differ from each other only by a projective transformation of the entire ambient $(m+1)$-space.*

Proof Choose a projective coordinate system for the line s that gives the points S_0 and S_π the coordinates 0 and ∞, and let $[u_1, \ldots, u_{m+1}]$ be the coordinates given to the points $(S_1, \ldots, S_{m+1})$. Determine the parameter vector $[v_1, \ldots, v_{m+1}]$ in a similar way from the T-points. In proving the $B_{m,1,1}$ Representation Theorem, we saw that, as long as all primitive degeneracies are avoided, the parameter vectors $[u_1, \ldots, u_{m+1}]$ and $[v_1, \ldots, v_{m+1}]$ determine a representation of $B_{m,1,1}$ that is unique up to a projective transformation of the ambient space. □

We shall exploit the $m = 2$ case of this proposition, in Chapter 9, in order to help us analyze the degeneracies that can arise in the cubic case of the Witness Theorem.

Exercise 7.4-3 Recall that the budget matroid $B_{1,1,1}$ is the Pappus matroid. Letting $m := 1$ in Proposition 7.4-1, which primitive degeneracies should we watch out for in constructing a Pappus configuration?

[Answer: Only Types 1 through 3 can occur, because Types 4 through 6 require four distinct row indices. Letting i, j, and k be the three row indices in some order, a Type-3 degeneracy occurs when the equality of cross ratios $A_{(i,j,k,\pi)} = B_{(i,j,\pi,k)}$ causes the three points P_k, A_k, and B_k in the k^{th} row to be collinear.]

Exercise 7.4-4 Primitive degeneracies of Types 3 through 6 can occur when constructing representations of the budget matroid $B_{2,1,1}$. Describe the forbidden incidence that arises in each type of degeneracy as concretely as possible. In particular, name the P-points that *are* involved, rather than those that are *not* involved.

[Answer: Let i, j, k, and l be the four row indices, in some order.

Type 3 If $A_{(i,j,k,\pi)} = B_{(i,j,\pi,k)}$, then the line $P_k P_l$ meets the line $A_k B_k$.

Type 4 If $A_{(i,j,k,\pi)} = B_{(i,j,\pi,l)}$, then the line $P_k P_l$ meets the line $A_k B_l$.

Type 5a If $A_{(i,j,k,\pi)} = B_{(i,j,k,l)}$, then the line $P_l B_l$ meets the line a.

Type 5b If $A_{(i,j,k,l)} = B_{(i,j,k,\pi)}$, then the line $P_l A_l$ meets the line b.

Type 6 If $A_{(i,j,k,l)} = B_{(i,j,k,l)}$, then the line a meets the line b.]

Exercise 7.4-5 Given a configuration $\{(P_i, A_i, B_i)\}_{i\in[0..3]}$ of points in 3-space that fails to represent the budget matroid $B_{2,1,1}$ precisely because of a Type 6 degeneracy — that is, the lines a and b fail to be skew — what geometric property do the six points $\{P_0, P_1, P_2, P_3, A_\pi, B_\pi\}$ in the plane π have?

[Answer: The equality $A_{(0,1,2,3)} = B_{(0,1,2,3)}$ means that the four homogeneous coordinates of the points $A_\pi = [0, v_1, v_2, v_3]$ and $B_\pi = [0, u_1, u_2, u_3]$ have the same cross ratio. The locus of points $X = [0, x_1, x_2, x_3]$ whose four coordinates have some fixed cross ratio is a conic in the plane π that passes through all four of the frame points P_0, P_1, P_2, and P_3. So the six listed points lie on a common conic.]

Chapter 8

Null systems

We have proved some general results about the representability of budget matroids, but our original goal was to find a configuration that characterizes the dependence of four cubic polynomials. It is going to turn out that the matroid $B_{2,1,1}$ provides that configuration. Unfortunately, we aren't yet equipped to prove that.

We shouldn't lose heart, since the results in Chapter 7 do give some positive indications. If we let $m = 2$ in the $B_{m,1,1}$ Representation Theorem, we find that the matroid $B_{2,1,1}$ is representable over the rational numbers, with $\#(B_{2,1,1}) = 19$. Representability is certainly a good sign, and nineteen degrees of freedom is also just the number that we would have been hoping for, had we paused to give the matter any thought.

Why is that? (The following accounting is the cubic analog of Exercise 2.1-1.) Let M be some matroid on 12 elements for which the cubic analog of the Witness Theorem holds. Fix some center line o in 3-space, and consider representations of M that witness to the dependence of 3-blocks of planes through the line o. There are eleven degrees of freedom in choosing such a 3-dependent block of planes: twelve slopes, but subject to one constraint. There are eight degrees of freedom in a projective transformation of 3-space that fixes the center line o and fixes every plane through o: fifteen degrees of freedom in an arbitrary projective transformation, minus six to fix two planes through o (which suffices to fix o itself), minus one more to fix all of the other planes through o. In order for the cubic Witness Theorem to hold for the matroid M, we must have $\#(M) = 11 + 8 = 19$.

The bad news is that the recipe for constructing representations of $B_{2,1,1}$ given in Chapter 7 isn't very helpful in exploring the relationship with 3-dependency. What we did in Chapter 7 was to choose the P-points and the A-points, after which we could easily construct the B-points. We shall refer to that construction as the *P-first construction*. But suppose that we are given a 3-dependent block of planes and that we would like to construct a representation of $B_{2,1,1}$ that witnesses to that 3-dependence. It isn't at all obvious how to constrain the choices of the P-points and A-points so that the B-points end up on the proper planes.

What we do instead, in Chapter 9, is to develop an alternative construction for representations of $B_{2,1,1}$, which we call the *P-last construction*. We choose the four A-points on a line a; the four B-points on a line b, skew to a; and the plane π, on which we want the four P points to lie. From that information, using the theory of null systems, we then construct the P-points, lying in the plane π. In this chapter, we prepare for that by reviewing the pretty theory of null systems.

8.1 Polarities in general

The *dual* of a projective n-space S is the projective n-space S^* whose k-flats are the $(n-k-1)$-flats of S. The *principle of duality* states that the dual space S^* is also a projective n-space. A *polarity* is a self-inverse (or *involutive*) isomorphism between a space S and its dual S^*, that is, a one-to-one, projective correspondence between the points of S and the points of S^*, which are the hyperplanes of S. Polarities come in two types, called *polar systems* and *null systems* [23], as we mentioned in Section 2.9. We now consider the linear algebra of this situation.

We shall represent a point in projective n-space as a column vector $\mathbf{c}$ of homogeneous coordinates, of height $n+1$; and we shall represent a hyperplane in n-space as a row vector $\mathbf{r}$ of homogeneous coefficients, of width $n+1$. The point $\mathbf{c}$ lies on the hyperplane $\mathbf{r}$ just when the dot product $\mathbf{rc}$ is zero.

Let M be an invertible $(n+1)$-by-$(n+1)$ matrix. Multiplication by M gives us a map from projective n-space to itself, taking the point $\mathbf{c}$ to the point $M\mathbf{c}$. Of course, because the coordinates are homogeneous, a nonzero scalar multiple uM of the matrix M gives us the same map. Such a map is called a *regular projective transformation*. Note that the hyperplane $\mathbf{r}$ is carried, by the projective transformation M, to the hyperplane $\mathbf{r}M^{-1}$, since we have $\mathbf{rc} = 0$ just when $(\mathbf{r}M^{-1})(M\mathbf{c}) = 0$.

We can also use invertible $(n+1)$-by-$(n+1)$ matrices to describe maps from projective n-space to its dual: The point $\mathbf{c}$ goes to the hyperplane $\mathbf{c}^{\mathbf{t}}M$, where the superscript '$\mathbf{t}$' means transpose. Note that the hyperplane $\mathbf{r}$ is carried, by this map, to the point $M^{-1}\mathbf{r}^{\mathbf{t}}$, since we have $\mathbf{rc} = 0$ just when $(\mathbf{c}^{\mathbf{t}}M)(M^{-1}\mathbf{r}^{\mathbf{t}}) = \mathbf{c}^{\mathbf{t}}\mathbf{r}^{\mathbf{t}} = (\mathbf{rc})^{\mathbf{t}} = 0$. Maps of this type are particularly simple — and are called *polarities* or *involutive correlations* — when they are equal to their own inverses; that is, mapping a point $\mathbf{c}$ to the hyperplane $\mathbf{c}^{\mathbf{t}}M$ and then mapping that hyperplane to the point $M^{-1}(\mathbf{c}^{\mathbf{t}}M)^{\mathbf{t}} = (M^{-1}M^{\mathbf{t}})\mathbf{c}$ gets us back to the same point $\mathbf{c}$ that we started with. For this to happen for all points $\mathbf{c}$, we must have $M^{-1}M^{\mathbf{t}} = vI$ for some nonzero scalar v, which means that $M^{\mathbf{t}} = vM$. Transposing both sides of this equation, we see that $M = vM^{\mathbf{t}}$, so $M = v^2M$. Since M is invertible and hence far from identically zero, we have $v^2 = 1$, so $v = \pm 1$.

Polarities with $v = 1$ have symmetric matrices $M = M^{\mathbf{t}}$ and are called *polar systems*; those with $v = -1$ have skew-symmetric matrices $M = -M^{\mathbf{t}}$ and are called *null systems*. While there are many parallels between the two theories, there are also important differences.

One difference arises when we consider whether or not a point lies on its own polar hyperplane. Fix a certain polarity, and let M be one of its associated matrices — so M is determined up to a scalar multiple. The polar hyperplane of the point $\mathbf{c}$, in the chosen polarity, has the homogeneous coefficients $\mathbf{c}^{\mathbf{t}}M$. So the point $\mathbf{c}$ lies on its own polar just when $\mathbf{c}^{\mathbf{t}}M\mathbf{c} = 0$. If the polarity is a polar system, then the matrix M is symmetric and this equation determines a certain hypersurface in the ambient projective space, the hypersurfaces that can be determined in this way being called *quadrics*. If the polarity is a null system, on the other hand, then the matrix M is skew-symmetric. For any point $\mathbf{c}$, it follows that the scalar $\mathbf{c}^{\mathbf{t}}M\mathbf{c}$ is equal to its own negative,

$$\mathbf{c}^{\mathbf{t}}M\mathbf{c} = (\mathbf{c}^{\mathbf{t}}M\mathbf{c})^{\mathbf{t}} = \mathbf{c}^{\mathbf{t}}M^{\mathbf{t}}\mathbf{c} = \mathbf{c}^{\mathbf{t}}(-M)\mathbf{c} = -(\mathbf{c}^{\mathbf{t}}M\mathbf{c}),$$

and hence must be zero.[1] Thus, in a null system, every point lies on its polar hyperplane.

Another important difference is that, while polar systems exist in spaces of any dimension, null systems exist only in spaces of odd dimension. If we take the determinant of both sides of the equation $M^{\mathbf{t}} = -M$, we find that $\det(M^{\mathbf{t}}) = \det(M) = (-1)^{n+1}\det(M)$. Since M has nonzero determinant, it follows that $(-1)^{n+1} = 1$, which means that n must be odd. In particular, there are no null systems in the plane, so most people who study projective geometry never have the pleasure of meeting a null system.

A third major difference involves the concept of equivalence, where two polarities on the same space are called *equivalent* when they differ only by a projective transformation. All null systems on a given space are equivalent, and this is true regardless of the properties of the field of scalars. If every scalar has a square root, then all polar systems are equivalent also; but the fewer scalars that have square roots, the more different equivalence classes of polar systems there are. For example, in 3-space over the real numbers — because negative scalars don't have square roots — there are three equivalence classes of polar systems and hence three different types of quadric surfaces [49]:

empty quadrics imaginary ellipsoids;

non-ruled quadrics ellipsoids, hyperboloids of two sheets, and elliptic paraboloids; and

ruled quadrics hyperboloids of one sheet and hyperbolic paraboloids.

To verify these claims about equivalence, note that two polarities are equivalent just when their matrices M and M' satisfy $M' = uP^{\mathbf{t}}MP$, for some nonzero scalar

[1] Over a field of characteristic 2, the equation $u = -u$ is an identity — it does not imply that $u = 0$. The standard cure for this, as Artin explains [4], is to define a matrix M to be *skew-symmetric* (or *alternating*) only when both $M = -M^{\mathbf{t}}$ and all of the diagonal elements of M are zero. The latter condition follows from the former when the characteristic is not 2. This whole issue doesn't affect us, however, since we are assuming that our scalar field is of characteristic zero.

u and invertible matrix P, the matrix P giving the projective transformation that takes one polarity into the other. Over any field of scalars, it is a standard result of matrix algebra [4, 26] that, for any invertible, skew-symmetric matrix M, there exists an invertible matrix P with

$$P^{\mathsf{t}}MP = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix},$$

where the two instances of I here are identity matrices of the same size. It follows from this that all null systems on a given space are equivalent.

In contrast, the extent to which an invertible, symmetric matrix M can be put into some standard form depends upon the existence of square roots in the field of scalars. Over the complex numbers, we can always find an invertible matrix P with $P^{\mathsf{t}}MP$ equal to the identity matrix. Over the real numbers, we can get the matrix $P^{\mathsf{t}}MP$ to be diagonal with each diagonal entry equal to ± 1, but we can't change the number of diagonal entries of each sign; that result is called *Sylvester's Law of Inertia* [3, 27]. We can, however, let $u = -1$ in the formula $M' = uP^{\mathsf{t}}MP$, thereby reversing all of the signs at once. It follows that, for a 4-by-4, invertible, symmetric matrix M over the real numbers, there are three standard forms, corresponding to the three types of quadric surfaces listed above:

empty quadrics all four diagonal elements of the same sign;

non-ruled quadrics three diagonal elements of one sign and the remaining one of the other sign; and

ruled quadrics two diagonal elements of each sign.

Over the rational numbers, things are more complicated still; Exercise 8.1-2 shows that, in any space of odd dimension, there are infinitely many different equivalence classes of polar systems. Fortunately, we are interested in null systems, so we don't need to explore the complexities of polar systems any further.

Exercise 8.1-1 Show that, on the line, not only are any two null systems equivalent, they are actually equal. In this unique null system, what is the polar P^* of a point P? Give the matrix of this null system.

[Answer: The polar P^* is the same point as P, but viewed as a hyperplane; there is no choice about this, since the point P and its polar hyperplane P^* must be mutually incident — which, in this case, means coincident. The matrix is $\left(\begin{smallmatrix} 0 & w \\ -w & 0 \end{smallmatrix}\right)$ for some nonzero w. Note that the point on the line with homogeneous coordinates $\left[\begin{smallmatrix} u \\ v \end{smallmatrix}\right]$, when viewed as a hyperplane, has homogeneous coefficients $\langle v, -u\rangle$.]

Exercise 8.1-2 For any odd n, construct an infinite family of polar systems on n-space over the rational numbers, no two of which are equivalent.

[Hint: When n is odd, it follows from the equation $M' = uP^{\mathsf{t}}MP$ that the ratio $\det(M')/\det(M)$ is a perfect square [24].]

8.2 Null systems in 3-space

The matrix of a null system in projective 3-space is 4-by-4, invertible, and skew-symmetric. Note that there are six degrees of freedom in such a matrix, one for each entry above the diagonal. Since an overall scalar multiple doesn't matter, we deduce that there are five degrees of freedom in the choice of a null system in 3-space. In this section, we study the geometry of such a null system. One good source for this material is Maxwell [28].

Recall from the previous section that all null systems in 3-space are equivalent, and thus have the same geometry, up to our choice of coordinate system. The five degrees of freedom determine only how the null system sits in 3-space.

There isn't much to say about points or planes: The polar plane of any point P is a plane P^* containing P; and the pole of any plane π is a point π^* lying on π. The relationship between a line ℓ and its polar line ℓ^* is more interesting.

Proposition 8.2-1 *Given any null system in 3-space, the polar ℓ^* of a line ℓ is either skew to ℓ or coincides with ℓ.*

In the former case, we shall call the pair of lines $\{\ell, \ell^*\}$ a *skew-polar pair*. In the latter case, where $\ell = \ell^*$, the line ℓ is called *self-polar*.

Proof For any line ℓ, whenever the point P lies on ℓ, the plane P^* must pass through the line ℓ^*. Thus, as P slides along ℓ, the polar plane P^* rotates around ℓ^*, and does so in such a way that P^* always passes through P. In typical cases, the two lines ℓ and ℓ^* are skew, and hence form a skew-polar pair $\{\ell, \ell^*\}$.

Could it happen that a line ℓ and its polar ℓ^* intersected without coinciding? No. As the point P slid along ℓ in such a case, the polar plane P^* would have to remain fixed at the unique plane spanned by ℓ and ℓ^*. But distinct points must have distinct polars.

We conclude that, whenever ℓ and ℓ^* intersect, they must actually coincide. In this case, as P slides along ℓ, the plane P^* rotates around that same line $\ell = \ell^*$. Thus, for any self-polar line ℓ, the null system provides a projective correspondence between the range of points along ℓ and the pencil of planes through ℓ. □

By the way, the geometry of a polar system in 3-space is quite different from that of a null system [30]. In a polar system, just because a line ℓ intersects its polar ℓ^* does not mean that the two must coincide. Indeed, a line intersects its polar just when it is tangent to the quadric surface associated with the polar system, while it coincides with its polar only when it lies entirely inside — that is, is a generating line of — that quadric. But back to null systems.

Proposition 8.2-2 *Given any null system in 3-space, let ℓ be any line and let P be any point on ℓ. The line ℓ is self-polar just when ℓ lies in the plane P^*.*

Proof The line ℓ^* always lies in the plane P^*. If the line ℓ is self-polar, then $\ell = \ell^*$, so the line ℓ lies in P^* also. Conversely, suppose that ℓ does lie in P^*. Since both ℓ and ℓ^* lie in P^*, they are not skew, so we know from Proposition 8.2-1 that they coincide — that is, ℓ is self-polar. □

If we fix a point P in space, we deduce that the lines through P that are self-polar are precisely those lying in the plane P^*. Thus, while there are 4 degrees of freedom in the choice of an arbitrary line in 3-space, there are 3 degrees of freedom in the choice of a self-polar line. In fact, the set of self-polar lines forms a *general linear complex*,[2] and knowing this linear complex is equivalent to knowing the null system. Indeed, Maxwell talks more about linear complexes [31] than he does about null systems as such [32].

As a corollary of Proposition 8.2-2, we can show that there are no triangles with all three edges self-polar.

Proposition 8.2-3 *If the points A, B, and C are distinct and all three of the lines AB, AC, and BC are self-polar with respect to some fixed null system, then A, B and C are collinear.*

Proof If A, B, and C were not collinear, the polar of each vertex of the triangle ABC would have to be the plane ABC, since the two triangle edges that meet at that vertex are self-polar. But distinct points must have distinct polars. □

Our next result gives us an easy way to prove that a line is self-polar.

Proposition 8.2-4 *Given any null system in 3-space, if the lines $\{\ell, \ell^*\}$ are a skew-polar pair, then any line m that meets both ℓ and ℓ^* is self-polar.*

Proof Let m be a line that meets both ℓ and ℓ^*, say at the points A and B respectively. Since m passes through A, the line m^* must lie in the plane A^*, which we know is the unique plane through ℓ^* that passes through A. Similarly, since m passes through B, the line m^* must lie in the plane B^*, which is the unique plane through ℓ that passes through B. It follows that m^* must coincide with the intersection $A^* \cap B^*$, which is precisely the line $m = AB$. Thus, m is self-polar. □

So suppose that the two opposite edges $\ell = AC$ and $\ell^* = BD$ of some tetrahedron $ABCD$ form a skew-polar pair. It follows from Proposition 8.2-4 that each of the other four edges AB, BC, CD, and DA of the tetrahedron is self-polar. We close this section with a converse to that result.

Proposition 8.2-5 *Given any null system in 3-space, let $ABCD$ be a skew quadrilateral all four of whose edges are self-polar lines. Then, the two diagonal lines $\{AC, BD\}$ form a skew-polar pair.*

[2] The word 'general' here means 'nonspecial', a *special linear complex* being the 3-parameter family of lines in 3-space that meet a fixed line.

Proof Since the vertex A lies on the two lines AB and AD, its polar plane A^* must pass through both $(AB)^* = AB$ and $(AD)^* = AD$; so $A^* = ABD$. In a similar way, the polar of any vertex of the quadrilateral is the plane spanned by the two adjacent edges. The polar of the diagonal AC is then the intersection $A^* \cap C^* = ABD \cap CBD$, which is precisely the opposite diagonal BD. □

Exercise 8.2-6 Given a null system — or a polar system, for that matter — in 3-space, show that two lines ℓ and m meet but do not coincide just when their polars ℓ^* and m^* meet but do not coincide.

[Hint: This is almost trivial. Two lines in 3-space meet just when they pass through a common point and just when they lie in a common plane.]

8.3 Skew-polar hexagons

We pointed out in Section 2.9 that one way to specify a null system in 3-space is to give a twisted cubic curve. Another very pretty way to specify a null system, as discussed in Exercise 8.3-2 below, is to give a Möbius pair of tetrahedra. For our purposes, though, it is most convenient to specify a null system by giving a certain type of skew hexagon.

Fix some null system in 3-space. Knowing any skew-polar pair of lines $\{\ell, \ell^*\}$ already goes a long way towards specifying that null system. Indeed, it specifies four out of the five degrees of freedom. It tells us, right away, the polars of all of the points lying on either ℓ or ℓ^*. Furthermore, given a point P on neither ℓ nor ℓ^*, there is a unique line m through P that meets both ℓ and ℓ^*; it is called the *common transversal*. Since the common transversal m touches both ℓ and ℓ^*, it must be self-polar, by Proposition 8.2-4. Thus, the polar plane P^* must pass through the line m. The one degree of freedom that remains rotates the plane P^* about the line m. We could tie down that one degree of freedom by specifying, for some such point P, its polar plane P^*. But we'll do something more symmetric.

Let us say that a hexagon $A_1B_2A_3B_1A_2B_3$ in 3-space is *skew-polar* with respect to a certain null system when its six vertices are distinct and when each side forms, with the opposite side, a skew-polar pair of lines — that is, each of the three pairs of lines $\{A_1B_2, A_2B_1\}$, $\{A_1B_3, A_3B_1\}$, and $\{A_2B_3, A_3B_2\}$ is skew-polar. Those requirements may not seem very strict, but it turns out that skew-polar hexagons are quite constrained.

First, every line joining two nonadjacent vertices of a skew-polar hexagon is self-polar. For example, because the opposite edges A_1B_2 and A_2B_1 form a skew-polar pair, it follows from Proposition 8.2-4 that the other four edges A_1A_2, B_1B_2, A_1B_1 and A_2B_2 of the tetrahedron $A_1A_2B_1B_2$ are self-polar. Repeating this argument twice more finishes the job.

In particular, the lines A_1A_2, A_1A_3, and A_2A_3 are all self-polar. We conclude from Proposition 8.2-3 that the points A_1, A_2, and A_3 lie on a common, self-polar line, which we shall call a. By a similar argument, the points B_1, B_2, and B_3 lie on a self-polar line b. Since the opposite sides A_1B_2 and A_2B_1 of the hexagon are skew, it follows that the lines a and b are also skew.

We have now uncovered all of the constraints on a skew-polar hexagon, as we can show by working backwards to construct one. Let a and b be any two skew lines, each of which is self-polar. For each point A on a, the polar plane A^* passes through a, and hence meets b in a unique point B. The line AB passes through A and lies in the plane A^*, so it is self-polar. It follows that the correspondence taking A to B is symmetric; that is, the unique point where the polar plane B^* meets the line a is A once again. Let A_1, A_2, and A_3 be any three distinct points on a, and let B_1, B_2, and B_3 be the corresponding points on b. We claim that the hexagon $A_1B_2A_3B_1A_2B_3$ is skew-polar. To see this, note that all four edges of the skew quadrilateral $A_1B_1B_2A_2$ are self-polar. It follows from Proposition 8.2-5 that the diagonals A_1B_2 and A_2B_1, which are a pair of opposite edges of the hexagon, form a skew-polar pair.

For future reference, note that there are nine degrees of freedom in the choice of a hexagon that is skew-polar for a given null system: three in the self-polar line a, three more in b, and a final three in the points A_1, A_2, and A_3 along a.

Exercise 8.3-1 Given a null system in 3-space, consider a cyclic list $\ell_1\ell_2^*\ell_3\ell_1^*\ell_2\ell_3^*$ of six lines in which opposite pairs of lines are skew-polar pairs and adjacent pairs are not skew. Show, from these assumptions alone, that the six lines are the edges of a skew-polar hexagon.

[Hints: Prove first that no adjacent pair of lines can coincide — for example, it cannot be the case that $\ell_1 = \ell_2^*$. We can then make the definitions

$$\begin{aligned} A_1 &:= \ell_1 \cap \ell_2^* \qquad & B_1 &:= \ell_1^* \cap \ell_2 \\ A_2 &:= \ell_2 \cap \ell_3^* \qquad & B_2 &:= \ell_2^* \cap \ell_3 \\ A_3 &:= \ell_3 \cap \ell_1^* \qquad & B_3 &:= \ell_3^* \cap \ell_1 \end{aligned}$$

and switch to thinking about the hexagon in terms of its vertices $A_1B_2A_3B_1A_2B_3$, rather than its edges. Conclude by proving that the six vertices are all distinct. Both proofs use Proposition 8.2-4.]

Exercise 8.3-2 Fix any null system in 3-space and fix any tetrahedron S. Show that the polar planes of the vertices of S are the faces of a tetrahedron T that forms, with S, a Möbius pair. Conversely, given any Möbius pair of tetrahedra in 3-space, show that there is a unique null system in which the faces of each tetrahedron are the polars of the vertices of the other tetrahedron. It follows from this that, if we fix any tetrahedron S, null systems in 3-space are in one-to-one correspondence with tetrahedra that form a Möbius pair with S. (Note that there are five degrees of freedom in either case.)

[Hint: The vertices of S are not coplanar, so the faces of T won't be concurrent. If A, B, and C are three of the vertices of S, then the three corresponding faces A^*,

B^*, and C^* of T will intersect at the vertex $A^* \cap B^* \cap C^* = (ABC)^*$ of T. Thus, the faces of S are the polars of the vertices of T, just as the faces of T are the polars of the vertices of S, which means that S and T form a Möbius pair. For the second half of the problem, see Griffiths and Harris [13].]

8.4 Skew-Pappian hexagons

If a hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ is skew-polar for any null system, its six vertices must be distinct and alternate vertices must lie on two skew lines. We shall call any hexagon in 3-space with those two properties *skew-Pappian*. In other words, a skew-Pappian hexagon is an ordered 2-block $\{(A_1, B_1), (A_2, B_2), (A_3, B_3)\}$ of points in 3-space, where A_1, A_2, and A_3 are distinct points along a line a and B_1, B_2, and B_3 are distinct points along a line b that is skew to a. Our goal in this section is to show that, given any skew-Pappian hexagon, there is a unique null system for which it is skew-polar.

Note that the degrees of freedom work out correctly, at least. There are fourteen degrees of freedom in a skew-Pappian hexagon: four in each of the two lines a and b and one in each of the six vertices, along its proper line. On the other hand, it takes five degrees of freedom to choose a null system in 3-space and it takes nine more to choose a hexagon that is skew-polar for that null system, so we again get a total of fourteen.

Lemma 8.4-1 *Given two skew-Pappian hexagons, there exists a projective transformation of 3-space that maps the first to the second.*

More precisely, let $A_1 B_2 A_3 B_1 A_2 B_3$ be a skew-Pappian hexagon and let π be a plane that passes through the line $A_1 B_1$, but does not pass through either the line $a := A_1 A_2 A_3$ or the line $b := B_1 B_2 B_3$. Let $A'_1 B'_2 A'_3 B'_1 A'_2 B'_3$ and π' be a second skew-Pappian hexagon and plane, with the analogous properties. Then, there exists a unique projective transformation of 3-space that both maps the first hexagon to the second and maps the plane π to the plane π'.

Proof Let C be the point where the line $A_3 B_3$ cuts the plane π, and define C' in an analogous way. No four of the five points (A_1, A_2, B_1, B_2, C) are coplanar, and the same holds for the primed versions $(A'_1, A'_2, B'_1, B'_2, C')$; so there is a unique projective transformation of 3-space that takes the former five points to the latter five. That transformation clearly takes the plane $\pi = A_1 B_1 C$ to the plane $\pi' = A'_1 B'_1 C'$. Furthermore, under that transformation, the common transversal from C to the lines a and b must map to the common transversal from C' to the lines a' and b'; so the point A_3 must map to A'_3 and B_3 must map to B'_3. □

Proposition 8.4-2 *Given any skew-Pappian hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ in 3-space, there is a unique null system $N = N((A_1, B_1), (A_2, B_2), (A_3, B_3))$ for which the hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ is skew-polar.*

Proof To show that some such null system exists, let N' be any fixed null system in that 3-space and let $A_1' B_2' A_3' B_1' A_2' B_3'$ be some hexagon that is skew-polar for the null system N'. We saw in Section 8.3 that the hexagon $A_1' B_2' A_3' B_1' A_2' B_3'$ must be skew-Pappian. Lemma 8.4-1 then tells us that there exists a projective transformation of 3-space that removes the primes, mapping the hexagon $A_1' B_2' A_3' B_1' A_2' B_3'$ to the given hexagon $A_1 B_2 A_3 B_1 A_2 B_3$. If we apply that projective transformation to the null system N', the result is a null system N for which the given hexagon is skew-polar.

It remains to show that the null system N is uniquely determined by the requirement that the hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ be skew-polar for N. To see why, let P be a generic point in 3-space. The two opposite sides $A_1 B_2$ and $A_2 B_1$ of the hexagon form a skew-polar pair, so, by Proposition 8.2-4, the common transversal c_3 from P to those two lines is self-polar. In a similar way, the common transversals c_2 and c_1 from P to the skew-polar pairs $(A_1 B_3, A_3 B_1)$ and $(A_2 B_3, A_3 B_2)$ are also self-polar. Thus, given only the hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ and the generic point P, we can construct three lines through P that lie in the plane P^*. It follows that the given hexagon uniquely determines the null system N. □

Of course, it had better be the case that those three concurrent lines c_1, c_2, and c_3 are coplanar — and indeed, the argument above proves that they are. But we can also show their coplanarity more directly, using Pappus's Theorem. Suppose that we put our eye at the generic point P and we look out at the skew-Pappian hexagon $A_1 B_2 A_3 B_1 A_2 B_3$. In the projective plane of lines through our eye, the hexagon that we see looks *Pappian*, that is, of the type that arises in Pappus's Theorem (as shown in Figure 4.1). The collinearity of the apparent intersections of the three pairs of opposite sides, when viewed from P, exactly corresponds to the fact that their three common transversals c_1, c_2, and c_3 are coplanar.

Exercise 8.4-3 Given a skew-Pappian hexagon $A_1 B_2 A_3 B_1 A_2 B_3$ in 3-space and a generic plane π, construct the pole π^* of the plane π in the associated null system $N((A_1, B_1), (A_2, B_2), (A_3, B_3))$.

[Answer: Let c_3 be the unique line in the plane π that is a common transversal of the lines $A_1 B_2$ and $A_2 B_1$ — that is, the line joining the point $A_1 B_2 \cap \pi$ to the point $A_2 B_1 \cap \pi$. Determine lines c_2 and c_1 in a similar way. The coplanar lines c_1, c_2, and c_3 are concurrent by the dual of Pappus's Theorem, and the pole π^* is the point of concurrence.]

8.5 The homogeneous coordinates of a pole

Given the homogeneous coordinates of the vertices of a skew-Pappian hexagon, it will be helpful to be able to compute the associated null system explicitly — that is, to go back and forth between the homogeneous coordinates of a pole point and the homogeneous coefficients of its polar plane. Finding formulas that worked in full generality would be a bit clumsy, because of the skew-Pappian condition: The vertices of a skew-Pappian hexagon must alternate between two skew lines, and that puts constraints on the coordinates of the vertices. For our purposes, though, it will suffice to deal with a very special case, one in which the skew-Pappian collinearities are guaranteed by zero coordinates.

Lemma 8.5-1 *The pole of the plane $\pi = \langle 1, 0, 0, -1\rangle$ in the null system determined by the skew-Pappian hexagon $A_1B_2A_3B_1A_2B_3$ with vertices*

$$\begin{aligned} A_1 &:= [1, a_1, 0, 0] & B_1 &:= [0, 0, b_1, 1] \\ A_2 &:= [1, a_2, 0, 0] & B_2 &:= [0, 0, b_2, 1] \\ A_3 &:= [1, a_3, 0, 0] & B_3 &:= [0, 0, b_3, 1] \end{aligned}$$

is the point π^* with homogeneous coordinates

$$\pi^* = \left[\begin{vmatrix} 1 & a_1 & b_1 \\ 1 & a_2 & b_2 \\ 1 & a_3 & b_3 \end{vmatrix}, \begin{vmatrix} 1 & a_1 & a_1b_1 \\ 1 & a_2 & a_2b_2 \\ 1 & a_3 & a_3b_3 \end{vmatrix}, \begin{vmatrix} 1 & a_1b_1 & b_1 \\ 1 & a_2b_2 & b_2 \\ 1 & a_3b_3 & b_3 \end{vmatrix}, \begin{vmatrix} 1 & a_1 & b_1 \\ 1 & a_2 & b_2 \\ 1 & a_3 & b_3 \end{vmatrix} \right].$$

Proof We verify that the point π^* results from the construction in Exercise 8.4-3. As a check of plausibility, note that the first and last homogeneous coordinates of the point π^* are the same, so the point π^* does lie on the plane π.

The points on the line A_1B_2 have the form $[u, ua_1, vb_2, v]$ for some ratio $u : v$. Thus, the line A_1B_2 intersects the plane π at the point $[1, a_1, b_2, 1]$. Similarly, the line A_2B_1 meets the plane π at the point $[1, a_2, b_1, 1]$. The construction of Exercise 8.4-3 tells us that the point π^* lies on the line in the plane π that joins those two intersections. This happens precisely if the matrix

$$\begin{pmatrix} 1 & a_1 & b_2 & 1 \\ 1 & a_2 & b_1 & 1 \\ \begin{vmatrix} 1 & a_1 & b_1 \\ 1 & a_2 & b_2 \\ 1 & a_3 & b_3 \end{vmatrix} & \begin{vmatrix} 1 & a_1 & a_1b_1 \\ 1 & a_2 & a_2b_2 \\ 1 & a_3 & a_3b_3 \end{vmatrix} & \begin{vmatrix} 1 & a_1b_1 & b_1 \\ 1 & a_2b_2 & b_2 \\ 1 & a_3b_3 & b_3 \end{vmatrix} & \begin{vmatrix} 1 & a_1 & b_1 \\ 1 & a_2 & b_2 \\ 1 & a_3 & b_3 \end{vmatrix} \end{pmatrix}$$

has rank only 2, which is indeed the case: The third row is equal to $(a_3-a_1)(b_3-b_2)$ times the second row minus $(a_3 - a_2)(b_3 - b_1)$ times the first row. $\square$

Chapter 9

On $B_{2,1,1}$ and 3-dependency

9.1 Constructing the P-points last

Using the theory of null systems, we can construct representations of the matroid $B_{2,1,1}$ in a new order: Choose the A-points, the B-points, and the plane π; then construct the P-points. We shall call this the *P-last construction*, to contrast it with the *P-first construction* used in proving the $B_{m,1,1}$ Representation Theorem. Since the four rows play completely symmetric roles in the P-last construction — there is no need to treat one row specially, as there was in the P-first case — we return to indexing the rows starting with 1:

$$\begin{array}{c} \begin{array}{ccc} 2 & 1 & 1 \end{array} \\ \begin{pmatrix} P_1 & A_1 & B_1 \\ P_2 & A_2 & B_2 \\ P_3 & A_3 & B_3 \\ P_4 & A_4 & B_4 \end{pmatrix} \end{array}$$

Proposition 9.1-1 (The P-Last Construction) *In any representation of the budget matroid $B_{2,1,1}$, the point P_1 is the pole of the plane $\pi := \mathrm{Span}(P_1, P_2, P_3, P_4)$ in the unique null system $N_1 := N((A_2, B_2), (A_3, B_3), (A_4, B_4))$ for which the skew-Pappian hexagon $A_2B_3A_4B_2A_3B_4$ is skew-polar, and analogous results hold for the points P_2, P_3, and P_4. We shall denote these results by writing $P_i = \pi^{*i}$, where the superscript '$*i$' means 'polar in the null system N_i'.*

Conversely, suppose that we choose any four collinear A-points in 3-space, any four collinear B-points, and any plane π. For generic choices of those nineteen parameters, the four A-points will be distinct, the four B-points will be distinct, and the lines a and b on which they lie will be skew. For i in $[1 \,..\, 4]$, we can then define the null system N_i, and we can construct the point $P_i := \pi^{*i}$. For generic choices of the nineteen parameters, once again, the resulting twelve points $\{(P_i, A_i, B_i)\}_{i\in[1..4]}$ will represent the matroid $B_{2,1,1}$.

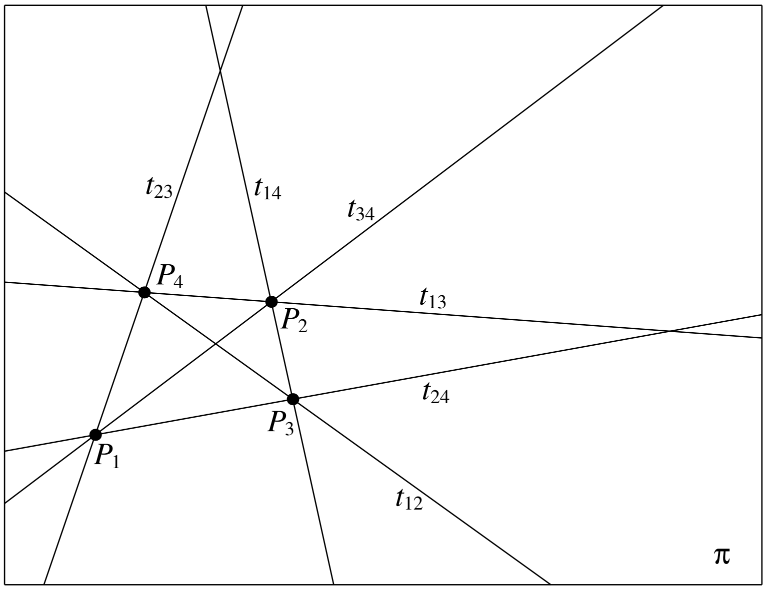

Figure 9.1: The complete quadrangle that arises in the plane π when the P-last construction is used to build a representation of the budget matroid $B_{2,1,1}$.

Proof Given any representation of $B_{2,1,1}$, consider what happens in the plane π. The four P-points lie in π, and no three of them are collinear. Thus, they determine a complete quadrangle, as shown in Figure 9.1. For $\{i, j, k, l\} = \{1, 2, 3, 4\}$, note that the quadruples $\{P_i, P_j, A_k, B_l\}$ and $\{P_i, P_j, A_l, B_k\}$ are both coplanar. Thus, the line $P_i P_j$ is the unique common transversal in the plane π of the skew lines $A_k B_l$ and $A_l B_k$ — call this common transversal t_{kl}. In particular, the point P_i lies on all three of the common transversals t_{jk}, t_{jl}, and t_{kl}. We argued in Exercise 8.4-3 that those three transversals are indeed concurrent and that the point $P_i = \pi^{*i}$ where they concur is the pole of the plane π in the null system N_i determined by the skew-Pappian hexagon $A_j B_k A_l B_j A_k B_l$.

Conversely, given the A-points, the B-points, and the plane π, suppose that we construct $P_i := \pi^{*i}$ for i in $[1 \,..\, 4]$; will the result be a valid representation?

It is easy to see that every set that should be mutually incident will be. The four P-points are coplanar because the pole of the plane π in any null system always lies in π. The four points in the perfect set $\{P_i, P_j, A_k, B_l\}$ are coplanar because P_i is placed at the point of concurrence of the three common transversals t_{jk}, t_{jl}, and t_{kl}, while P_j is placed at the point of concurrence of t_{ik}, t_{il}, and t_{kl} — so both P_i and P_j lie on the line t_{kl}, which meets the line $A_k B_l$.

To show that no forbidden incidences arise in generic cases, it suffices to check that none arise in some particular case, and any representation of the matroid $B_{2,1,1}$ provides such a case. Readers who made it through Chapter 7 already know that the matroid $B_{2,1,1}$ is representable; the rest can verify for themselves that the twelve points given in Exercise 9.1-2 constitute a representation. □

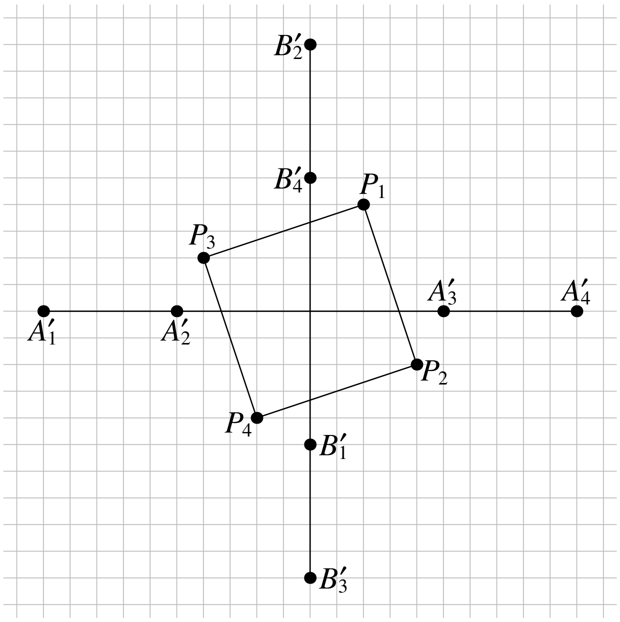

Figure 9.2: The projections onto the XY-plane of the twelve vertices of a particularly symmetric representation of the budget matroid $B_{2,1,1}$.

Exercise 9.1-2 Verify that the following is a representation of the budget matroid $B_{2,1,1}$, sitting in Euclidean (X, Y, Z)-space:

	P	A	B
1	$(2, 4, 0)$	$(-10, 0, 1)$	$(0, -5, -1)$
2	$(4, -2, 0)$	$(-5, 0, 1)$	$(0, 10, -1)$
3	$(-4, 2, 0)$	$(5, 0, 1)$	$(0, -10, -1)$
4	$(-2, -4, 0)$	$(10, 0, 1)$	$(0, 5, -1)$

The plane π is the plane $Z = 0$, which we take to be horizontal; the line a is parallel to the X axis, but one unit above it; the line b is parallel to the Y axis, but one unit below it. Figure 9.2 shows the (X, Y)-plane π, together with the points A'_i and B'_i, the vertical projections of A_i and B_i onto the plane π. Note that this representation is carried to itself by the following symmetry of order 4: Map each point (X, Y, Z) to the point $(-Y, X, -Z)$, permute the row indices through the 4-cycle $(1, 3, 4, 2)$, and swap the A and B columns.

[Hint: For any i and j, the line $A_i B_j$ meets the plane π at the midpoint of the segment $A'_i B'_j$. To show that the perfect sets are coplanar, check that, whenever i, j, and k are distinct, the point P_k lies on the line joining the midpoint of $A'_i B'_j$ to the midpoint of $A'_j B'_i$. To show that there are no forbidden incidences, it suffices to verify the absence of the primitive degeneracies listed in Exercise 7.4-4.]

Exercise 9.1-3 Figure 9.3 is a close-up of the rope-and-pole structure found by the

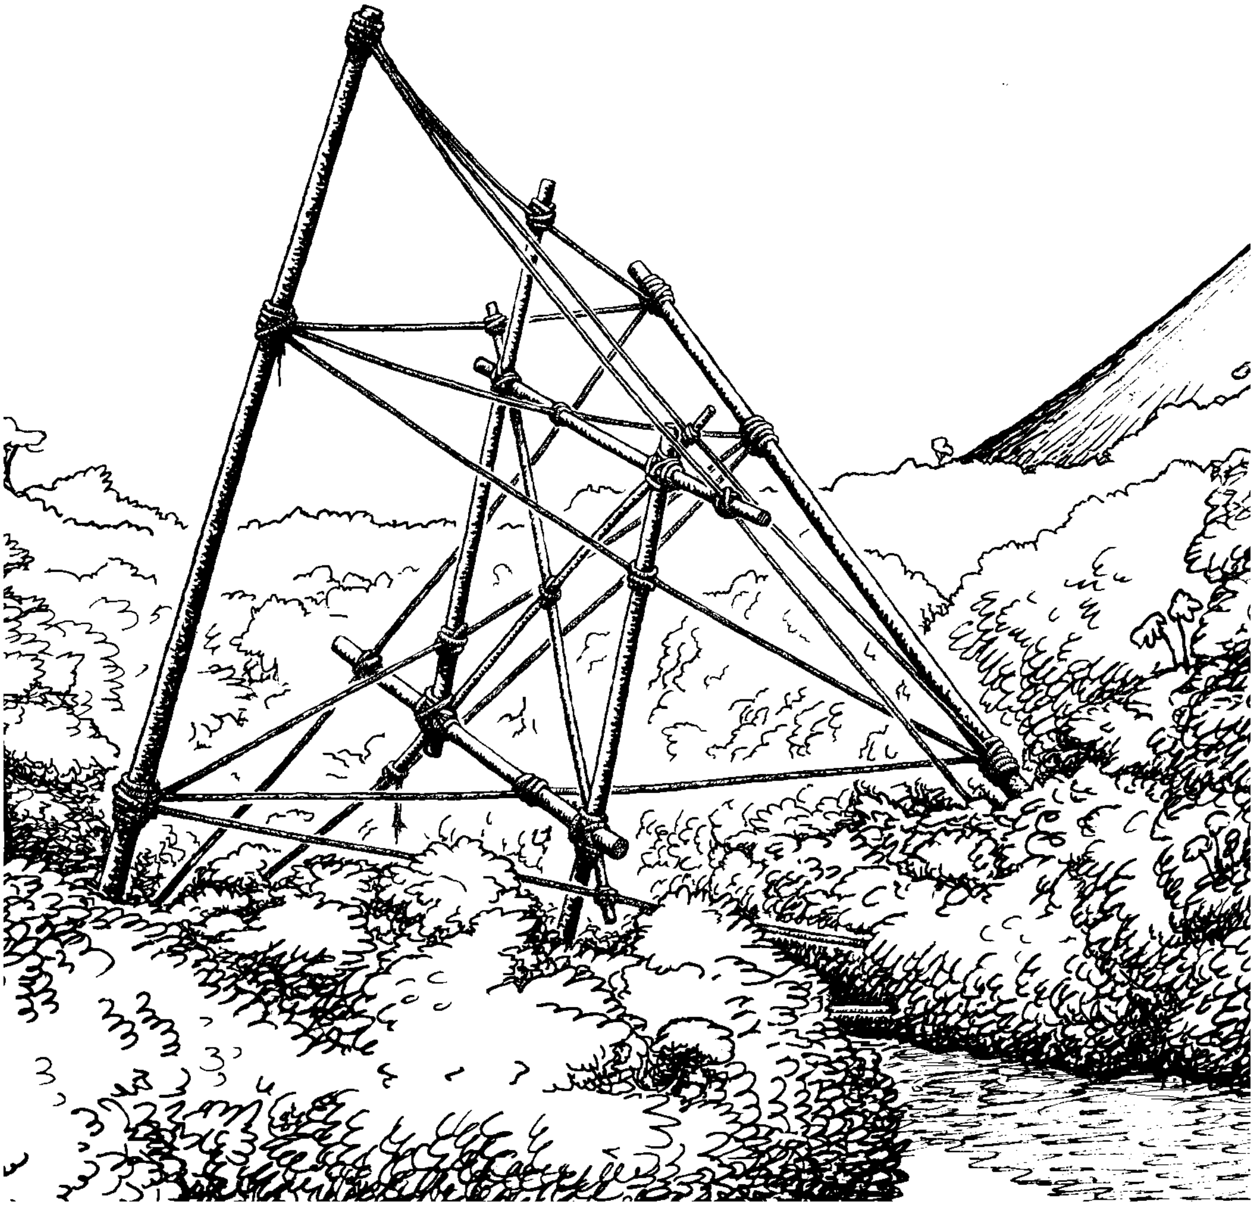

Figure 9.3: A representation of the budget matroid $B_{2,1,1}$ in the Brazilian jungle.

explorer in Jorge Stolfi's cartoon on the title page. Label the twelve largest knots to demonstrate the correspondence between this structure and the representation of $B_{2,1,1}$ in the preceding exercise. (Two of the twelve are hidden by jungle foliage.) Convince yourself that each of the twelve ropes demonstrates one of the twelve perfect coplanarities.

[Hint: Label the four knots of the central square P_1, P_2, P_3, and P_4, going from top to bottom down the page. The four knots on the left-hand pole can then be labeled in numeric order from A_1 on the bottom to A_4 on top, while the four knots on the right-hand pole get the labels B_2, B_4, B_1, and B_3, from top to bottom. The coplanarity of the perfect set $\{P_1, P_2, A_3, B_4\}$ is demonstrated by the rope from A_3 to B_4, which is tied, at its midpoint, to the beam that joins P_1 to P_2.]

9.2 Degenerate cases in the P-last construction

This section is just for those readers who made it to the end of Chapter 7 — in particular, to Proposition 7.4-2. Using that result, we can analyze exactly what to watch out for, in carrying out the P-last construction.

Suppose that we start the P-last construction with certain points (A_i) and (B_i), lying on the lines a and b, and with a certain plane π. Clearly, we must require that the four A-points be distinct and that none of them lies in the plane π. Letting $A_\pi := a \cap \pi$ denote the point where the line a cuts the plane π, we conclude that the five points $(A_1, A_2, A_3, A_4; A_\pi)$ must be distinct. In a similar way, the five points $(B_1, B_2, B_3, B_4; B_\pi)$ must also be distinct. Proposition 7.4-2 (or, more concretely, Exercise 7.4-4) then tells us precisely what else to watch out for: the primitive degeneracies of Types 3 through 6.

Each primitive degeneracy involves some cross ratio of four of the A-points happening to coincide with some cross ratio of four B-points. We shan't bother to review Types 3 through 5 here, except to note that the cross ratios that coincide in those cases involve either A_π or B_π or both. So those cases can be avoided simply by choosing the plane π carefully. A Type-6 degeneracy happens, on the other hand, when the cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ are equal — so the plane π is not involved. That makes a Type-6 degeneracy more troublesome, since perturbing the plane π is not enough to remove it.

If we avoid all primitive degeneracies of Types 3 through 6 in choosing our A-points and B-points (including A_π and B_π), we know from Proposition 7.4-2 that there do exist representations of $B_{2,1,1}$ whose A-points and B-points are projective images of ours. Fix any such representation. To avoid name conflicts, let's write its points as (A'_i) and (B'_i) and its lines as a' and b'. By Lemma 8.4-1, there is a unique projective transformation of 3-space that takes the skew-Pappian hexagon $A'_\pi B'_1 A'_2 B'_\pi A'_1 B'_2$ and the plane π' to the hexagon $A_\pi B_1 A_2 B_\pi A_1 B_2$ and the plane π. Since cross ratios along the lines a and b are preserved, that transformation must also take A'_i to A_i and B'_i to B_i for i in $\{3, 4\}$. We conclude that there is a representation of $B_{2,1,1}$ that has its A-points, its B-points, and its plane π just where we put ours, and the P-last construction will produce that representation.

On a related topic, it is interesting to contrast the behaviors of the P-first and P-last constructions when given input parameters with a single, primitive degeneracy of one of the six types. Of course, if the input parameters are degenerate, then the output configuration has forbidden incidences. But do the outputs of the P-first and P-last constructions have the same forbidden incidences? It depends on the type of degeneracy.

The easy cases are Types 3, 4, and 5. Those primitive degeneracies don't cause any problem for either construction, and the two constructions, when given the same degenerate input, produce the same degenerate output.

Type 2 is almost as simple. Having, say, $A_1 = A_2$ is no problem for the P-first construction. It is a problem for the P-last construction, since, for i in $\{3, 4\}$, the hexagon $A_1 B_2 A_i B_1 A_2 B_i$ is not skew-Pappian, and hence the null systems N_3 and N_4 are not well defined. But if we allow the P-last construction to determine its output by taking the limit as A_1 approaches A_2, it produces the same degenerate output that the P-first construction produces.

Type 1 is a more interesting case. Having, say, $A_1 = A_\pi$ is no problem for the P-last construction. But the way in which its output degenerates, in that case, is to have the four points P_2, P_3, P_4, and A_1 become collinear, in the plane π. That could never happen in the P-first construction, since the four P-points are fixed, at the outset, to be a projective frame for π. As A_1 approaches A_π in the P-first construction, what happens, instead, is that the three points B_2, B_3, and B_4 all converge to the point P_1.

Type 6 is also an interesting case. In the P-first construction, the bad thing that happens when $A_{(1,2,3,4)} = B_{(1,2,3,4)}$ is that the lines a and b intersect. That can't happen in the P-last construction, since the lines a and b are fixed, at the outset, to be skew. What happens instead, in the P-last construction, is that the projective correspondence from the line a to the line b that takes A_i to B_i for i in $[1\,..\,3]$ also takes A_4 to B_4. Therefore, the four null systems N_1 through N_4 coincide, and the four P-points defined by $P_i := \pi^{*i}$ coincide at the pole of the plane π in that single null system.

9.3 From frames to grids

To prepare for the Projection and Witness Theorems, we need to prove some easy facts about projective transformations of 3-space. One convenient way to do that is to introduce a configuration with nine points and six lines that we shall call a *grid*. Grids can play a role analogous to the role played by projective frames. Recall that a *projective frame* for 3-space consists of five points, with no four coplanar. Any projective frame can be mapped to any other by a unique projective transformation, and we are going to show that the same holds for grids.

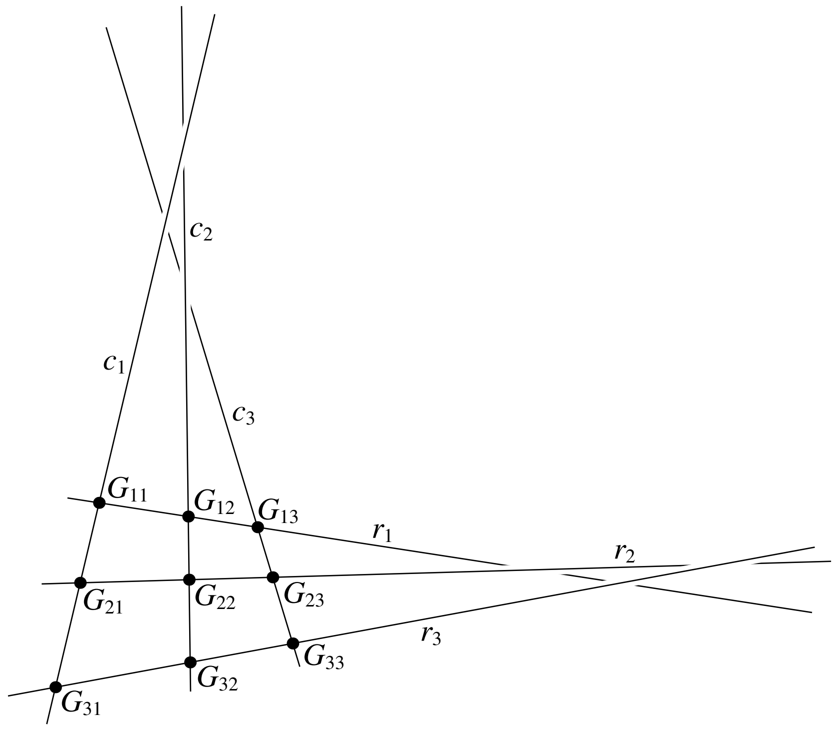

Figure 9.4: An example of a grid in 3-space: three skew lines and three of their common transversals. (Note that this figure is a relabeling of Figure 5.1.)

A *grid* consists of three skew lines in 3-space, along with three of their common transversals, as shown in Figure 9.4. Note that the three common transversals will also be skew. A grid has nine pairs of lines that intersect, each such pair determining a point. Focusing on those points, we can think of a grid, alternatively, as a 3-by-3 matrix $\mathbf{g} = (G_{ij})_{i,j\in[1..3]}$ of points in 3-space for which

- no two of the nine points coincide;
- three of the nine points are collinear just when they form a row or a column; and
- four or five of the nine points are coplanar, as we shall show shortly, just when they can be covered by the union of a row and a column.

Proposition 9.3-1 *Given any grid* $\mathbf{g} = (G_{ij})_{i,j\in[1..3]}$ *in* 3*-space, the five points* $(G_{11}, G_{12}, G_{21}, G_{22}, G_{33})$ *form a projective frame. Conversely, given any projective frame* $(G_{11}, G_{12}, G_{21}, G_{22}, G_{33})$ *for* 3*-space, there is a unique quadruple of points* $(G_{13}, G_{23}, G_{31}, G_{32})$ *that completes that frame into a grid.*

Proof Given any grid, any two distinct row indices i and j, and any two distinct column indices k and l, the quadrilateral $G_{ik}G_{jk}G_{jl}G_{il}$ is not planar, since the i^{th}

and j^{th} row lines are skew, as are the k^{th} and l^{th} column lines. It follows immediately that the four points G_{11}, G_{12}, G_{21}, and G_{22} are not coplanar. The four points G_{11}, G_{12}, G_{21}, and G_{33} are not coplanar either, because, if they were, then the four points G_{11}, G_{13}, G_{31}, and G_{33} would be coplanar as well. The remaining three cases are similar, so the five points $(G_{11}, G_{12}, G_{21}, G_{22}, G_{33})$ do form a projective frame.

Conversely, suppose that $(G_{11}, G_{12}, G_{21}, G_{22}, G_{33})$ is any projective frame for 3-space. Construct some grid $\mathbf{h} = (H_{ij})$, somewhere in that 3-space. We have just shown that the five points $(H_{11}, H_{12}, H_{21}, H_{22}, H_{33})$ form a projective frame, so there is a unique projective transformation that carries that H frame to the given, G frame. The image of the grid $\mathbf{h}$ under that projective transformation is also a grid; so there is at least one way to complete the G frame into a grid $\mathbf{g}$.

It remains to show that this completion $\mathbf{g}$ is unique. The first two row lines $r_1 := G_{11}G_{12}$ and $r_2 := G_{21}G_{22}$ are clearly uniquely determined, as are the first two column lines $c_1 := G_{11}G_{21}$ and $c_2 := G_{12}G_{22}$. As for the third row line r_3, it must be the unique common transversal from G_{33} to the two skew lines c_1 and c_2; and the third column line c_3 is determined in a similar way. □

Proposition 9.3-2 *Given any two grids, there is a unique projective transformation of* 3*-space that carries the first to the second.*

Proof Use Proposition 9.3-1 to convert the grids to frames. □

The reason that we are interested in grids is that they allow us to deduce the existence of some projective transformations that we are going to need later.

Corollary 9.3-3 *Let r_1, r_2, and r_3 be three skew lines and let π be a plane that does not contain any common transversal of those three lines. Let r'_1, r'_2, r'_3, and π' be another three lines and plane with the same properties. There exists a unique projective transformation of* 3*-space that takes each line r_i to the line r'_i and takes the plane π to the plane π'.*

Proof It cannot be the case that one of the lines, say r_1, lies in the plane π. If this did happen, the plane π would have to intersect both r_2 and r_3 only in points, and the line joining those two points of intersection would be a common transversal of the three lines (r_i), lying in π.

It follows that, for i in $[1 \,..\, 3]$, the line r_i cuts the plane π at a unique point, which we shall call G_{ii}. Let c_i be the common transversal from G_{ii} to the two skew lines r_j and r_k, where $\{i, j, k\} = \{1, 2, 3\}$. The three lines (r_i) and their three transversals (c_i) form a grid (G_{ij}).

Forming a grid (G'_{ij}) in an analogous way, there is a unique projective transformation of 3-space that maps the first grid to the second. That transformation must take each line r_i to the line r'_i and the plane $\pi = G_{11}G_{22}G_{33}$ to the plane $\pi' = G'_{11}G'_{22}G'_{33}$. □

Corollary 9.3-4 *If the triple of lines (o, a, b) is skew and the triple (o, a', b') is also skew, there exists a unique projective transformation of 3-space that fixes the line o and every plane through o, while taking the lines a and b to a' and b'.*

Proof Choose three distinct planes through o, say π_1, π_2, and π_3. In each plane π_i, there is a unique common transversal of the three lines o, a, and b. The three lines (o, a, b) together with those three common transversals form a grid.

Repeat this construction, starting with the three lines (o, a', b') and using the same three planes π_1, π_2, and π_3 through o. The result is a second grid.

The unique projective transformation that maps the first grid to the second fixes all three of the planes (π_i), so it fixes all planes through o. □

Exercise 9.3-5 Recall from Section 4.5 what it means to be a 'configuration' in the narrow sense of that term. A grid is obviously a configuration of points and lines in 3-space of type $(9_2, 6_3)$. Show that a grid can also be viewed as a configuration of points and planes of type $(9_5, 9_5)$.

[Answer: Each intersecting pair of lines spans a plane that passes through five points, and each point lies on five such planes.]

9.4 The Projection Theorem

Once we have formulas, in terms of some parameters, for all twelve of the points in a generic representation of $B_{2,1,1}$, the cubic case of the Projection Theorem reduces to algebra: showing that the determinant of a certain matrix is zero. Indeed, we could probably have proved the Projection Theorem for $B_{2,1,1}$ based on the P-first construction and the formulas in terms of the parameters $[u_1, u_2, u_3]$ and $[v_1, v_2, v_3]$ in the proof of the $B_{m,1,1}$ Representation Theorem. But we can get more symmetric formulas from the P-last construction, in terms of some new parameters $(a_1, a_2, a_3, a_4; b_1, b_2, b_3, b_4)$; so we shall do that instead.

Theorem 9.4-1 (Projection Theorem, cubic case) *Given any representation*

$$\begin{array}{ccc} 2 & 1 & 1 \\ \end{array}$$
$$\begin{pmatrix} P_1 & A_1 & B_1 \\ P_2 & A_2 & B_2 \\ P_3 & A_3 & B_3 \\ P_4 & A_4 & B_4 \end{pmatrix}$$

of the matroid $B_{2,1,1}$ in some 3-space and given a generic line o in that 3-space, the twelve planes joining o to those four triples of points form a 3-dependent block.

Proof We restrict the line o to be generic only so that we don't have to consider such degenerate cases as $o = a$, where projecting any one of the four A-points from the line o yields the indeterminate plane.

We proceed by analytic geometry, choosing our coordinate system on points $[w, x, y, z]$ with some care. To make Lemma 8.5-1 applicable, we want the line a to be the line $y = z = 0$, the line b to be the line $w = x = 0$, and the plane π to be the plane $w = z$, which has homogeneous coefficients $\langle 1, 0, 0, -1\rangle$. We also want the center line o of the projection to have simple coordinates — in fact, to be the line $y - w = z - x = 0$, whose points have the form $[u, v, u, v]$. To verify that we have enough freedom to achieve all of those goals at once, we apply Corollary 9.3-3.

Since the twelve given points represent the matroid $B_{2,1,1}$, the lines a and b are skew, and neither of them lies in the plane π. Hence, there is a unique common transversal of the lines a and b that lies in the plane π. Since the line o is generic, we can assume that o is skew to that common transversal. It follows that no common transversal of a, b, and o lies in the plane π.

Those same geometric properties hold of the lines and plane with the coordinates that we want a, b, o, and π to have. The three lines $y = z = 0$, $w = x = 0$, and $y-w = z-x = 0$ are indeed skew, since combining the equations for any two of them yields only the indeterminate solution $w = x = y = z = 0$. The plane $w = z$ does not contain any common transversal of those three lines, since the three points $[0, 1, 0, 0]$, $[0, 0, 1, 0]$, and $[1, 1, 1, 1]$ at which the three lines $y = z = 0$, $w = x = 0$, and $y - w = z - x = 0$ cut the plane $w = z$ are not collinear.

It follows from Corollary 9.3-3 that there is a unique coordinate system for 3-space in which the three lines (a, b, o) and the plane π have the coordinates that we want them to have. Where are the twelve points of the given representation in that coordinate system? The point $A_\pi := a \cap \pi$ has the homogeneous coordinates $[0, 1, 0, 0]$. The point A_i, for i in $[1\,..\,4]$, lies on the line a but is distinct from A_π; so we must have $A_i = [1, a_1, 0, 0]$, for some finite scalar a_i. Similarly, we have $B_\pi = [0, 0, 1, 0]$ and $B_i = [0, 0, b_i, 1]$, for some finite scalar b_i. We are going to use the scalars $(a_1, a_2, a_3, a_4; b_1, b_2, b_3, b_4)$ as our parameters. Note that these parameters involve eight degrees of freedom, while the homogeneous parameter vectors $[u_1, u_2, u_3]$ and $[v_1, v_2, v_3]$ that we used in analyzing the P-first construction involve only four. We need an additional four degrees of freedom in this analysis because we have constrained our choice of coordinate system to put the center line o of the projection in a fixed, simple place.

Lemma 8.5-1 now tells us the coordinates of the P-points. When $\{i, j, k, l\} = \{1, 2, 3, 4\}$, we have

$$P_i = \left[\begin{vmatrix} 1 & a_j & b_j \\ 1 & a_k & b_k \\ 1 & a_l & b_l \end{vmatrix}, \begin{vmatrix} 1 & a_j & a_j b_j \\ 1 & a_k & a_k b_k \\ 1 & a_l & a_l b_l \end{vmatrix}, \begin{vmatrix} 1 & a_j b_j & b_j \\ 1 & a_k b_k & b_k \\ 1 & a_l b_l & b_l \end{vmatrix}, \begin{vmatrix} 1 & a_j & b_j \\ 1 & a_k & b_k \\ 1 & a_l & b_l \end{vmatrix} \right].$$

Let's abbreviate those four coordinates as $P_i = [w_i, x_i, y_i, z_i]$ — so $w_i = z_i$. Note that swapping any two of the three indices j, k, and l would negate all four of the coordinates; but that doesn't matter, because they are homogeneous.

Given some point $[w, x, y, z]$, what is the slope of the plane joining that point to the center line o? It is easy to check that the homogeneous coefficients of the joining plane are $\langle z-x, w-y, x-z, y-w\rangle$: The point $[w, x, y, z]$ lies on that plane because $(z-x)w+(w-y)x+(x-z)y+(y-w)z = 0$, while every point $[u, v, u, v]$ on the line o lies on that plane because $(z-x)u+(w-y)v+(x-z)u+(y-w)v = 0$. As for what slope to assign to that joining plane, we showed in Lemma 2.5-1 that the notion of n-dependence is projectively invariant. Hence, we can take any three distinct planes we like, in the pencil of planes through o, and assign to those planes the slopes 0, 1, and ∞. To be definite, let's measure the slope of the plane $\langle z-x, w-y, x-z, y-w\rangle$ by using the ratio $z-x : y-w$ of the first homogeneous coefficient to the last.

Measuring slopes in this way, the plane joining the line o to the point $A_i = [1, a_i, 0, 0]$ has slope $-a_i : -1$, which is the same as $a_i : 1$. The plane joining o to $B_i = [0, 0, b_i, 1]$ has slope $1 : b_i$. And the plane joining o to $P_i = [w_i, x_i, y_i, z_i]$ has slope $z_i - x_i : y_i - w_i$. To build the i^{th} row of the 4-by-4 matrix that tests 3-dependence, we take those three ratios and combine them as in the elementary symmetric polynomials — that is, by the process that takes the three simple ratios $t^\uparrow : t^\downarrow$, $u^\uparrow : u^\downarrow$, and $v^\uparrow : v^\downarrow$ into the compound ratio

$$t^\uparrow u^\uparrow v^\uparrow \ : \ t^\uparrow u^\uparrow v^\downarrow + t^\uparrow u^\downarrow v^\uparrow + t^\downarrow u^\uparrow v^\uparrow \ : \ t^\uparrow u^\downarrow v^\downarrow + t^\downarrow u^\uparrow v^\downarrow + t^\downarrow u^\downarrow v^\uparrow \ : \ t^\downarrow u^\downarrow v^\downarrow.$$

Doing so for each of the four rows in turn, we conclude that the twelve slopes are 3-dependent just if this determinant is zero:

$$\begin{vmatrix} a_1(z_1-x_1) & (1+a_1b_1)(z_1-x_1)+a_1(y_1-w_1) & b_1(z_1-x_1)+(1+a_1b_1)(y_1-w_1) & b_1(y_1-w_1) \\ a_2(z_2-x_2) & (1+a_2b_2)(z_2-x_2)+a_2(y_2-w_2) & b_2(z_2-x_2)+(1+a_2b_2)(y_2-w_2) & b_2(y_2-w_2) \\ a_3(z_3-x_3) & (1+a_3b_3)(z_3-x_3)+a_3(y_3-w_3) & b_3(z_3-x_3)+(1+a_3b_3)(y_3-w_3) & b_3(y_3-w_3) \\ a_4(z_4-x_4) & (1+a_4b_4)(z_4-x_4)+a_4(y_4-w_4) & b_4(z_4-x_4)+(1+a_4b_4)(y_4-w_4) & b_4(y_4-w_4) \end{vmatrix}$$

That determinant is zero for the simple reason that the first and third rows of the matrix have the same sum as the second and fourth. To see this, start by considering the term a_1z_1 in the upper-left entry. The determinant

$$\begin{vmatrix} a_1 & 1 & a_1 & b_1 \\ a_2 & 1 & a_2 & b_2 \\ a_3 & 1 & a_3 & b_3 \\ a_4 & 1 & a_4 & b_4 \end{vmatrix}$$

is zero, since its first and third columns are equal; if we expand by cofactors down the first column, we find that $a_1z_1 - a_2z_2 + a_3z_3 - a_4z_4 = 0$. A similar argument starting with the determinant

$$\begin{vmatrix} a_1 & 1 & a_1 & a_1b_1 \\ a_2 & 1 & a_2 & a_2b_2 \\ a_3 & 1 & a_3 & a_3b_3 \\ a_4 & 1 & a_4 & a_4b_4 \end{vmatrix} = 0$$

shows that $a_1x_1 - a_2x_2 + a_3x_3 - a_4x_4 = 0$ as well; so the alternating sums of each of the two terms in the first column is zero. In the second column, of the six terms in each entry, similar arguments handle four, leaving the two terms $a_ib_iz_i + a_iy_i$. The alternating sums of those two terms are not separately zero, but they cancel each other out. Using cofactors down the first column again, the alternating sum $a_1b_1z_1 - a_2b_2z_2 + a_3b_3z_3 - a_4b_4z_4$ is the determinant

$$\begin{vmatrix} a_1b_1 & 1 & a_1 & b_1 \\ a_2b_2 & 1 & a_2 & b_2 \\ a_3b_3 & 1 & a_3 & b_3 \\ a_4b_4 & 1 & a_4 & b_4 \end{vmatrix},$$

while the sum $a_1y_1 - a_2y_2 + a_3y_3 - a_4y_4$ is the determinant

$$\begin{vmatrix} a_1 & 1 & a_1b_1 & b_1 \\ a_2 & 1 & a_2b_2 & b_2 \\ a_3 & 1 & a_3b_3 & b_3 \\ a_4 & 1 & a_4b_4 & b_4 \end{vmatrix}.$$

Those two determinants are negatives of each other because we can convert one matrix into the other by swapping the first and third columns. □

Exercise 9.4-2 In the P-last construction, suppose that the cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ are distinct, which means that the determinant of the matrix

$$\mathbf{m} := \begin{pmatrix} 1 & a_1 & b_1 & a_1b_1 \\ 1 & a_2 & b_2 & a_2b_2 \\ 1 & a_3 & b_3 & a_3b_3 \\ 1 & a_4 & b_4 & a_4b_4 \end{pmatrix}$$

is nonzero, and suppose that the plane π is generic. In analyzing the P-last construction in Section 9.2, we showed that the four P-points given by $P_i := \pi^{*i}$ will then form a projective frame for the plane π. Verify this directly and algebraically, using the formulas from Lemma 8.5-1 for the homogeneous coordinates $[w_i, x_i, y_i, z_i]$ of the point P_i. In particular, whenever $\{i, j, k, l\} = \{1, 2, 3, 4\}$, show that P_i, P_j, and P_k are not collinear because

$$\begin{vmatrix} x_i & y_i & z_i \\ x_j & y_j & z_j \\ x_k & y_k & z_k \end{vmatrix} = \pm \det(\mathbf{m})^2.$$

[Hint: Given any n-by-n matrix $\mathbf{s}$, let's borrow a term from the theory of determinants [34] and refer to the matrix of cofactors of $\mathbf{s}$ as the *adjugate*[1] of $\mathbf{s}$, and let's

[1] Artin [2] uses the name 'adjoint' for the transpose of the adjugate, but that conflicts with the Hermitian adjoint, which is the transpose of the complex conjugate.

write it $(\operatorname{adg}\mathbf{s})$. Cramer's Rule [2] tells us that $\mathbf{s}(\operatorname{adg}\mathbf{s})^{\mathbf{t}} = (\operatorname{adg}\mathbf{s})^{\mathbf{t}}\mathbf{s} = (\det\mathbf{s})I$, from which it is easy to calculate that $(\operatorname{adg}(\operatorname{adg}\mathbf{s})) = (\det\mathbf{s})^{n-2}\mathbf{s}$. The coordinates (x_i), (y_i), and (z_i) appear, up to sign, as entries in the adjugate matrix $(\operatorname{adg}\mathbf{m})$, which means that the determinant that we want to evaluate is, up to sign, an entry in the matrix $(\operatorname{adg}(\operatorname{adg}\mathbf{m}))$.]

9.5 The Witness Theorem

The place where the P-last construction really shines is in proving the cubic case of the Witness Theorem.

Theorem 9.5-1 (Witness Theorem, cubic case) *Let $\{(\rho_i, \alpha_i, \beta_i)\}_{i\in[1..4]}$ be some 3-dependent, ordered block of planes in 3-space, all passing through a common line o.*[2] *In generic cases, there exists a representation $\{(P_i, A_i, B_i)\}_{i\in[1..4]}$ of the budget matroid $B_{2,1,1}$ each of whose twelve vertices lies on the corresponding plane and none of whose vertices lies on the line o itself — and which hence witnesses to the 3-dependence of the planes. Furthermore, this representation is unique in the sense that any witnessing representation can be carried to any other by a projective transformation of 3-space that fixes the line o and fixes every plane through o.*

Proof To construct a witnessing representation, choose the lines a and b to be any two lines that are skew to each other and skew to o. Let A_i denote the point $A_i := \alpha_i \cap a$ and B_i denote $B_i := \beta_i \cap b$, for i in $[1\,..\,4]$. In generic cases, the points A_i and B_i will be distinct and the two cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ will also be distinct. The P-last construction tells us how to finish up building a representation of $B_{2,1,1}$: We choose some plane π and, whenever $\{i, j, k, l\} = \{1, 2, 3, 4\}$, we define the point $P_i := \pi^{*i}$ to be the pole of the plane π in the null system $N_i := N((A_j, B_j), (A_k, B_k), (A_l, B_l))$ determined by the skew-Pappian hexagon $A_j B_k A_l B_j A_k B_l$.

How should we choose the plane π? Note that the point $P_i = \pi^{*i}$ will lie on the plane ρ_i just when the plane π passes through the point $R_i := \rho_i^{*i}$. So we want π to pass through all four of the points (R_i). The assumed 3-dependence of the block of planes had better imply that the four points (R_i) are coplanar.

To see that it does, choose any three of them — say R_1, R_2, and R_3 — and choose π to be the plane $\pi := R_1 R_2 R_3$ that those three points determine. The resulting configuration of twelve points $\{(P_i, A_i, B_i)\}$ will, in generic cases, be a representation of $B_{2,1,1}$, and it will have all of its vertices on the appropriate input planes, except that P_4 might not lie on ρ_4. Let ρ_4' be the unique plane through o and P_4. It follows from the Projection Theorem that the slopes of the twelve planes $\{(\rho_1, \alpha_1, \beta_1), (\rho_2, \alpha_2, \beta_2), (\rho_3, \alpha_3, \beta_3), (\rho_4', \alpha_4, \beta_4)\}$ will be 3-dependent. But, in

[2] We refer to the first plane in the i^{th} triple as ρ_i, rather than as π_i, in order to avoid confusion with the plane $\pi := \operatorname{Span}(P_1, P_2, P_3, P_4)$ below.

generic cases once again, there is only one slope that the twelfth plane can have, given the other eleven slopes, that makes the twelve slopes 3-dependent. So we must have $\rho_4 = \rho'_4$, and we are done with the construction.

As for uniqueness, the only free choices that we made were the choices of the skew lines a and b. By Corollary 9.3-4, there is a unique projective transformation of 3-space that fixes the line o, fixes every plane through o, and takes the two lines a and b, skew to each other and to o, to any other two such lines a' and b'. Thus, any witnessing representation differs from any other only by a projective transformation. □

From the proof of the Witness Theorem, we can extract a flat-side geometric construction that solves the following problem: Given eleven planes through a line o — say α_i and β_i for i in $[1 \,..\, 4]$ and ρ_i for i in $[1 \,..\, 3]$ — construct the unique twelfth plane ρ_4 through o that makes the block $\{(\rho_i, \alpha_i, \beta_i)\}_{i\in[1..4]}$ 3-dependent. We choose two lines a and b that are skew to each other and skew to o, and we let A_i be the point where the line a cuts the plane α_i, and similarly for B_i. For i in $[1 \,..\, 3]$, we construct the point $R_i := \rho_i^{*i}$, the pole of the plane ρ_i in the null system determined by the skew-Pappian hexagon $A_j B_k A_l B_j A_k B_l$. Finally, letting π denote the plane $\pi := R_1 R_2 R_3$, we construct the point $P_4 := \pi^{*4}$. The plane ρ_4 that achieves 3-dependence is the unique plane through o that passes also through P_4. Note that, if the plane ρ_4 is our only goal, there is no need to construct the vertices P_1, P_2, and P_3 of the witnessing $B_{2,1,1}$ representation.

Exercise 9.5-2 (Cubic harmonic conjugacy) Here is an ordered block of scalars that is 3-dependent:

$$\begin{bmatrix} -1 & -1 & -1 \\ 4 & 4 & 4 \\ 5 & 5 & 5 \\ -5 & -2 & 1 \end{bmatrix}$$

Choose a line o in 3-space, find the twelve planes through o with those slopes, and carry out the construction for a witnessing representation of $B_{2,1,1}$. Is what results a valid representation, or are these slopes a degenerate case? What is special about the resulting configuration and its relationship with the center line o?

[Answer: The result is a valid representation of $B_{2,1,1}$, but it has the special property that the first three row planes $P_1A_1B_1$, $P_2A_2B_2$, and $P_3A_3B_3$ intersect, not in a point, as is typically the case, but in an entire line. In addition, the center line o of the projection coincides with that line of intersection.

Commentary: This situation is the cubic analog of harmonic conjugacy. Recall that two scalars u and v are *quadratic harmonic conjugates* of two distinct scalars p and q just when the block

$$\begin{bmatrix} p & p \\ q & q \\ u & v \end{bmatrix}$$

is 2-dependent. In a similar way, we say that three scalars u, v, and w are *cubic harmonic conjugates* of three distinct scalars p, q, and r just when the block

$$\begin{bmatrix} p & p & p \\ q & q & q \\ r & r & r \\ u & v & w \end{bmatrix}$$

is 3-dependent. In the example above, the scalars $(-5, -2, 1)$ are cubic harmonic conjugates of the distinct scalars $(-1, 4, 5)$.

The new wrinkle that arises in the cubic case is that only those representations of $B_{2,1,1}$ that have a special property can serve as witnesses to the harmonic conjugacy of two triples. In the quadratic case, given any representation of $B_{2,1}$ — that is, given any complete quadrilateral — there are three places where we can put the center point O so as to end up, as in Figure 2.4, with two pairs of lines through O that are harmonic conjugates. In the cubic case, on the other hand, in order to end up with two triples of planes that are harmonic conjugates, we must start with a representation of $B_{2,1,1}$ in which three of the four row planes lie in a common pencil; and then we must choose the center line o of the projection to coincide with the axis of that pencil.

For any degree n, harmonic conjugacy gives a binary relation on unordered n-tuples of distinct scalars, and that relation is always symmetric. For example, in the cubic case, suppose that both of the triples $\{p, q, r\}$ and $\{u, v, w\}$ consist of distinct scalars. Then, when either of the blocks

$$\begin{bmatrix} p & p & p \\ q & q & q \\ r & r & r \\ u & v & w \end{bmatrix} \qquad \text{or} \qquad \begin{bmatrix} u & u & u \\ v & v & v \\ w & w & w \\ p & q & r \end{bmatrix}$$

is 3-dependent, they both are.

But the relation of n-ic harmonic conjugacy is reflexive only when the degree n is odd. For example, the 1-block $\begin{bmatrix} p \\ p \end{bmatrix}$ is always 1-dependent and the 3-block

$$\begin{bmatrix} p & p & p \\ q & q & q \\ r & r & r \\ p & q & r \end{bmatrix}$$

is always 3-dependent. Indeed, if we interpret the 3-dependence of that latter block geometrically, we can demonstrate a property of twisted cubic curves that we mentioned back in Section 2.9. The cubic polynomial $(X - p)(X - q)(X - r)$ is the intersection of the osculating planes to the twisted cubic curve $F_t := (X - t)^3$ at the three points F_p, F_q, and F_r. We mentioned in Section 2.9 that the intersection point $(X-p)(X-q)(X-r)$ always lies in the plane $F_p F_q F_r$ determined by the three points of osculation.

When the degree n is even, on the other hand, the relation of n-ic harmonic conjugacy is not reflexive. For example, the 2-block

$$\begin{bmatrix} p & p \\ q & q \\ p & q \end{bmatrix}$$

is 2-dependent only when p and q are equal. We can interpret that fact geometrically also. The quadratic polynomial $(X-p)(X-q)$ is the intersection of the tangent lines to the conic curve $G_t := (X-t)^2$ at the points G_p and G_q. But, for p and q distinct, that intersection point never lies on the chord $G_p G_q$.]

9.6 Degenerate cases in projection and witnessing

One unpleasant aspect of our cubic Projection and Witness Theorems is that we did not deal with the degenerate cases. It is an obvious open question to do better — to either outlaw or handle each possible degenerate case. What challenges would we face, were we to tackle this question, and what tools could we use to try to meet those challenges?

Let's fix a center line o, for the remainder of this section. We would have three primary tools.

First, we could outlaw some of the 3-dependent ordered blocks of planes

$$\begin{bmatrix} \rho_1 & \alpha_1 & \beta_1 \\ \rho_2 & \alpha_2 & \beta_2 \\ \rho_3 & \alpha_3 & \beta_3 \\ \rho_4 & \alpha_4 & \beta_4 \end{bmatrix}$$

through the center line o, declaring that those blocks were too degenerate to be worthy of consideration. For example, we would almost surely want to outlaw those blocks in which all twelve planes coincided, since that would force all twelve points of any witnessing configuration to lie in a common plane. Let's call the blocks that we don't choose to outlaw the *legal blocks*.

Second, we could liberalize somewhat our notion of the relevant configurations of twelve points. We probably want to demand, of our twelve points, all of the incidences that are demanded for a representation of the budget matroid $B_{2,1,1}$. But we might want to allow some incidences that the matroid $B_{2,1,1}$ forbids. Let's refer to a configuration of twelve points that satisfies the resulting weaker rules as a 3-*tester*.

Third, we must set up some rules about the relationship between a 3-tester and the center line o from which we intend to project its twelve points. Is any point of the 3-tester allowed to lie on the line o? If so, what constraints are placed on the slope of the indeterminate plane that results? Is the line that joins two points of the 3-tester allowed to meet the line o? And so forth. If a 3-tester is located, with respect to the line o, so as to satisfy our requirements, we'll call it *well-located*.

By exploiting those three tools, we would try to arrange that:

- Projecting, from the line o, the twelve vertices of any 3-tester that is well-located with respect to o results in a 3-block of planes through o that is both 3-dependent and legal.

- Conversely, given any 3-block of planes through the line o that is both legal and 3-dependent, there exists a 3-tester that is well-located with respect to o and whose vertices lie on the appropriate planes, and which hence witnesses to the 3-dependence. Furthermore, this witnessing 3-tester is unique up to a projective transformation that fixes the line o and all planes through o.

Here are four of the challenges that we would face in doing so.

Challenge 1: There is a 10-dimensional family of 3-dependent blocks of planes through the line o in which the cross ratios $\alpha_{(1,2,3,4)}$ and $\beta_{(1,2,3,4)}$ of the planes in the A and B columns are equal. These blocks cause trouble in the Witness Theorem, since there are no representations of the matroid $B_{2,1,1}$ in which the cross ratios of the A-points and B-points are equal. One way to meet this challenge would be to declare such 3-blocks of planes illegal. But the authors suspect that it would be better to allow 3-testers in which the lines a and b intersect — that is, 3-testers with Type-6 degeneracies. It is not difficult to check that the Projection Theorem continues to hold for such 3-testers.

Ignoring Challenge 1 for a moment, let's fix the planes (α_i) and (β_i) in such a way that the cross ratios $\alpha_{(1,2,3,4)}$ and $\beta_{(1,2,3,4)}$ are distinct. There is a 3-parameter family of columns $(\rho_1, \rho_2, \rho_3, \rho_4)$ that make the block of planes 3-dependent. As the column $(\rho_1, \rho_2, \rho_3, \rho_4)$ varies over this family, the plane $\pi := R_1 R_2 R_3 R_4$ that appears in the P-last construction varies over all possible planes. Thus, lots of degenerate cases arise. Some are pretty mild, such as degeneracies of Types 1, 3, 4, or 5. But others are worse.

Challenge 2: Suppose that we choose the planes (α_i) and (β_i) arbitrarily and that we choose an additional plane μ, also in the pencil through o. We then choose each plane ρ_i so that the cross ratios $(\rho_i, \alpha_j, \alpha_k, \alpha_l)$ and $(\mu, \beta_j, \beta_k, \beta_l)$ are equal, where $\{i, j, k, l\} = \{1, 2, 3, 4\}$. The resulting block is always 3-dependent — so such blocks form a 9-parameter family. But, if we try to carry out the P-last construction to get a witnessing configuration, the four points (R_i) will lie, no three collinear, on a plane π that passes through the line a, and the four points (P_i) will all lie along a.

Challenge 2 looks harder to meet than Challenge 1. If we liberalize our notion of a 3-tester enough to allow such configurations, we run into trouble in the Projection Theorem: Since the P-points and the A-points all lie along one line, the perfect coplanarities hold trivially, so each of the twelve points can slide arbitrarily along its line. It may be possible to deal with this problem by making the plane π a separate component of a 3-tester and by requiring, in addition to the dependencies of the matroid $B_{2,1,1}$, that the P-points satisfy the relations $P_i = \pi^{*i}$, for the specified plane π. But it is far from clear what we should do if the lines a and b also happen to intersect, as in Challenge 1 above, which means that the null systems N_i are not well-defined.

Challenge 3: There is a 9-dimensional family of 3-dependent blocks in which the plane ρ_4 is not uniquely determined by the other eleven planes. What typically happens, if we apply the P-last construction to such a case, is that the point P_4 ends up lying on the center line o. That isn't too bad. The geometry is trying, as best it can, to tell us that we can take ρ_4 to be any plane in the pencil through o without destroying 3-dependence. Note that the two degrees of freedom that are lost by requiring this special property of the block of planes, taking us from 11 down to 9, match the two degrees of freedom that are lost by requiring that the point P_4 lie on the center line o.

Challenge 4: There is an 8-dimensional family of 3-dependent blocks in which all three of the planes in any single row can be varied arbitrarily without destroying 3-dependence; we studied one such block in Exercise 2.4-4. What typically happens, if we apply the P-last construction to such a block, is that the four points (R_i) end up collinear along a line r, leaving the plane π with an extra degree of freedom. Choosing a generic plane π through the line r results in a representation of $B_{2,1,1}$ that does witness to the 3-dependence of the twelve planes; but different choices of π give different representations. Thus, the uniqueness claim in the Witness Theorem fails; rather than too few witnesses, we have too many.

But enough, already, about degenerate cases. Let's return to happier topics.

Exercise 9.6-1 In the situation of Challenge 4, it turns out that the line r, on which all four of the points (R_i) given by $R_i := \rho_i^{*i}$ lie, is a common transversal of the four lines (o^{*i}). This means that, as we vary the plane π in the pencil with axis r, there are four special cases: the planes given by $\pi_i := \text{Span}(r, o^{*i})$, for i in $[1 \mathinner{..} 4]$. What is special about the witnessing representation of $B_{2,1,1}$ that results from choosing π to be π_i?

[Answer: When the plane π passes through the line o^{*i}, the point $P_i = \pi^{*i}$ lies on the line o. Thus, the resulting representation witnesses — in the sense of Challenge 3 — to the 3-dependence, not only of the given block, but also of any block obtained from it by varying the single plane ρ_i.]

Chapter 10

The budget matroid $B_{1,1,1,1}$

The final budget matroid of rank 4 is $B_{1,1,1,1}$. Recall that a representation of $B_{1,1,1,1}$ consists of sixteen points

$$\begin{array}{cccc} 1 & 1 & 1 & 1 \\ \end{array}$$
$$\begin{pmatrix} A_1 & B_1 & C_1 & D_1 \\ A_2 & B_2 & C_2 & D_2 \\ A_3 & B_3 & C_3 & D_3 \\ A_4 & B_4 & C_4 & D_4 \end{pmatrix}$$

in 3-space, where the four points in each column are collinear, lying along the four skew lines a, b, c, and d, and where the four points in any of the $4! = 24$ perfect sets are coplanar. The resulting structure has lots of combinatorial symmetries; note that the symmetric group on four letters acts independently on the rows and on the columns.

As it happens, the null-system machinery that we have developed to study representation of the matroid $B_{2,1,1}$ can also be used to construct representations of $B_{1,1,1,1}$. That result seems worth exploring here, even though the matroid $B_{1,1,1,1}$ isn't relevant to the concept of n-dependence for any n.

Be warned that one step in our construction for a representation of $B_{1,1,1,1}$ is to solve a quadratic equation. That makes the success of the construction contingent upon the existence of square roots in the field of scalars.

10.1 Generic representations

Lemma 10.1-1 *In any representation of the budget matroid $B_{1,1,1,1}$, the point D_4 lies on the line c^{*4}, the polar of the column line $c = C_1C_2C_3C_4$ in the null system $N_4 := N((A_1, B_1), (A_2, B_2), (A_3, B_3))$. Analogously, for any i in $[1\,..\,4]$, the point D_i lies on the line c^{*i} and the point C_i lies on the line d^{*i}.*

Proof Whenever $\{i, j, k\} = \{1, 2, 3\}$, the set $\{A_i, B_j, C_k, D_4\}$ is perfect and hence coplanar, as is the set $\{A_j, B_i, C_k, D_4\}$. So the line C_kD_4 meets both of the lines

$A_i B_j$ and $A_j B_i$. Since those two lines form a skew-polar pair in the null system N_4, we conclude from Proposition 8.2-4 that the line $C_k D_4$ is self-polar in N_4, from which it follows that the point C_k lies in the plane D_4^{*4}. This holds for any k in $[1 \,..\, 3]$, so the entire line c must lie in the plane D_4^{*4}. Hence, the point D_4 lies on the line c^{*4}. □

Lemma 10.1-1 states only one-sixth of the full truth, because we could use any pair of columns in building the null systems — there is nothing special about columns A and B. But it simplifies the notation a bit to assume that columns A and B are always the two that are used to build the null systems (N_i), and that restricted result is enough for our needs.

Proposition 10.1-2 *In generic cases and when the requisite square root exists in the field of scalars, the following process produces a representation in* 3*-space of the budget matroid* $B_{1,1,1,1}$*: Choose three skew lines* a, b, *and* c, *choose four distinct points* $(A_i)_{i\in[1..4]}$ *on the line* a, *and choose four distinct points* $(B_i)_{i\in[1..4]}$ *on* b. *For* $\{i, j, k, l\} = \{1, 2, 3, 4\}$, *let* c^{*i} *be the line that is polar to the line* c *in the null system* $N_i := N((A_j, B_j), (A_k, B_k), (A_l, B_l))$. *In generic cases, the four lines* (c^{*i}) *will be skew and, when the requisite square root exists, they will have two distinct common transversals; let* d *be one of them. For* i *in* $[1 \,..\, 4]$, *let* C_i *be the point where* c *meets* d^{*i} *and let* D_i *be the point where* d *meets* c^{*i}.

Furthermore, there exist nondegenerate, rational choices for the input parameters a, b, c, (A_i), *and* (B_i) *for which the requisite square root is also rational. So the matroid* $B_{1,1,1,1}$ *is representable over the rational numbers, and hence over any field of characteristic zero.*

Proof We first verify that, in generic cases, the four lines (c^{*i}) are skew. It suffices to check that they are skew in some particular case — for example, writing points in the form $[w, x, y, z]$, in the case in which

$$\begin{aligned} A_1 &= [1, 8, 0, 0] & B_1 &= [0, 0, -5, 1] \\ A_2 &= [1, -2, 0, 0] & B_2 &= [0, 0, 0, 1] \\ A_3 &= [1, 7, 0, 0] & B_3 &= [0, 0, 2, 1] \\ A_4 &= [1, 4, 0, 0] & B_4 &= [0, 0, -1, 1] \end{aligned}$$

and in which c is the line $y - w = z - x = 0$. The polar lines (c^{*i}) then turn out to be:

$$\begin{aligned} &c^{*1}: 64w - 22x - 108y = 31y + 22z - 108w = 0 \\ &c^{*2}: 58w - 9x - y = 7y + 9z - w = 0 \\ &c^{*3}: 2w - x - 4y = 3y + z - 4w = 0 \\ &c^{*4}: 24w - 2x - 28y = 41y + 2z - 28w = 0. \end{aligned}$$

And those four lines[1] are indeed skew.

Finding a common transversal of four skew lines involves solving a quadratic equation. If every scalar has a square root, then we can solve any quadratic, so any four skew lines have two common transversals. (In degenerate cases, the quadratic equation might be a perfect square, in which case the two transversals will coincide. What distinguishes these cases geometrically is that each of the four skew lines is tangent to the unique ruled quadric that contains the other three. In cases that are even more degenerate — in particular, when all four skew lines lie in a single ruled quadric — there is a one-parameter family of common transversals — to wit, the other family of generating lines of that same quadric. What happens to the quadratic equation, in such cases, is that all three of its coefficients are zero.) But if there exist scalars that don't have square roots, then four skew lines may not have any common transversals.

Fortunately, there are rational input parameters for which the requisite square root is rational. For example, the four polar lines (c^{*i}) above have the two rational lines $68w - 14x + 4y = 27y + 14z - 66w = 0$ and $68w - 14x - 66y = 27y + 14z + 4w = 0$ as their common transversals.

Next, we suppose that the input parameters have been chosen so that the four lines (c^{*i}) are skew and have two common transversals. By Lemma 10.1-1, we must choose the fourth column line d to be one of those two common transversals, in order to have any hope of arriving at a representation of $B_{1,1,1,1}$. Furthermore, we must place the point D_i, for i in $[1 \,..\, 4]$, where the line d meets the line c^{*i}. Since the line d is a common transversal of the four lines (c^{*i}), it follows from Exercise 8.2-6 that the line c will be a common transversal of the four lines (d^{*i}). So we can place the point C_i, symmetrically, where the line c meets the line d^{*i}. We claim that, in the resulting configuration, all sets of points that should be mutually incident are so. The column collinearities are obvious, so we consider a perfect set, such as $\{A_1, B_2, C_3, D_4\}$; are those four points coplanar?

We have placed the point C_3 where the line c meets the line d^{*3}. It follows that the plane C_3^{*3} passes through both of the lines c^{*3} and d. In particular, the plane C_3^{*3} passes through the point D_4; so the line C_3D_4 is self-polar in the null system N_3. From a similar argument focused on the point D_4, we conclude that the line C_3D_4 is also self-polar in the null system N_4. Thus, the line C_3D_4 lies in both of the planes C_3^{*3} and C_3^{*4}.

Could it be that the two planes C_3^{*3} and C_3^{*4} coincide? No. The line c^{*3} lies in

[1]People doing serious geometry in 3-space usually choose to represent lines using *Plücker line coordinates* [28, 50]. We represent a line here as the intersection of two planes only to avoid setting up the Plücker machinery. Note that, in the four lines (c^{*i}) and in their two common transversals below, the coefficient of x in the first equation is the negative of the coefficient of z in the second equation; those two coefficients come from the same Plücker coordinate. In the four lines (c^{*i}), the coefficient of y in the first equation is equal to the coefficient of w in the second; but those equalities result from our special choices of the lines a and b — that equality doesn't hold for most lines. Indeed, the two common transversals differ in that those two coefficients are interchanged.

the plane C_3^{*3} and the line c^{*4} lies in the plane C_3^{*4}, and we are assuming that the four lines (c^{*i}) are skew. Therefore, the two planes C_3^{*3} and C_3^{*4} are distinct, and the line C_3D_4 is their unique line of intersection.

Meanwhile, the lines A_1B_2 and A_2B_1 form a skew-polar pair in the null system N_3, since they are an opposite pair of edges of the skew-Pappian hexagon $A_1B_2A_4B_1A_2B_4$ that defines N_3. In a similar way, the same two lines form a skew-polar pair in the null system N_4. It follows from Proposition 8.2-4 that the common transversal from C_3 to the skew lines A_1B_2 and A_2B_1 is self-polar in both N_3 and N_4. Thus, that common transversal must lie in both of the planes C_3^{*3} and C_3^{*4}, so it must coincide with C_3D_4, their line of intersection. We conclude that the line C_3D_4 is a common transversal of the lines A_1B_2 and A_2B_1, so both of the perfect sets $\{A_1, B_2, C_3, D_4\}$ and $\{A_2, B_1, C_3, D_4\}$ are coplanar. The rest of the perfect sets are coplanar by similar arguments.

It remains to show that, in generic cases, the resulting configuration is free of forbidden incidences. It suffices to verify that there are no forbidden incidences in some particular case. Continuing the example above, suppose that we choose, as our line d, the first of the two common transversals, the line $68w - 14x + 4y = 27y + 14z - 66w = 0$. We then end up with the following sixteen points:

	A	B	C	D
1	$[1, 8, 0, 0]$	$[0, 0, -5, 1]$	$[1, 3, 1, 3]$	$[16, 76, -6, 87]$
2	$[1, -2, 0, 0]$	$[0, 0, 0, 1]$	$[13, 59, 13, 59]$	$[1, 6, 4, -3]$
3	$[1, 7, 0, 0]$	$[0, 0, 2, 1]$	$[1, 2, 1, 2]$	$[3, 14, -2, 18]$
4	$[1, 4, 0, 0]$	$[0, 0, -1, 1]$	$[1, 13, 1, 13]$	$[4, 20, 2, 15]$

There are $\binom{16}{4} = 1820$ sets of 4 of these points. Of those sets, 49 are coplanar because they contain at least 3 points from the first column, another 49 similarly from each of the other three columns, 24 are coplanar because they are perfect, and the remaining 1600 — the bases of the matroid $B_{1,1,1,1}$ — are all non-coplanar. □

Corollary 10.1-3 *Over a field, such as the complex numbers, in which all scalars have square roots, choosing a representation of the budget matroid $B_{1,1,1,1}$, sitting in a fixed 3-space, involves $\#(B_{1,1,1,1}) = 20$ degrees of freedom.*

Exercise 10.1-4 In a representation of the budget matroid $B_{1,1,1,1}$, show that the two lines A_1B_2 and A_2B_1 are the two common transversals of the four skew lines a, b, C_3D_4, and C_4D_3.

10.2 A projective compass

Our next goal is to convert the process that Proposition 10.1-2 uses to produce representations of the matroid $B_{1,1,1,1}$ into a geometric construction. In order to do that, we need some geometric tool more powerful than a flat-side; essentially, we need a tool that solves quadratic equations — at least, solves those quadratics that do have solutions. In Euclidean geometry over the real numbers, the standard tool for this job is the compass; so, it seems reasonable to call the tool that we need here a *projective compass*. There are lots of alternatives for precisely what a projective compass might do, but the different models all have equivalent power:

Model 1 Given three distinct points A, B, and C along a line ℓ and three more distinct points A', B', and C' along that same line ℓ, the projective compass constructs for us the two fixed points of the unique projectivity from ℓ to ℓ that maps (A, B, C) to (A', B', C').

Model 2 Given two pairs of points (A, B) and (C, D) along a line ℓ, the projective compass constructs for us the two fixed points of the unique involution of ℓ that swaps A with B and swaps C with D.

Model 3 Given five points and one line in the plane, the projective compass constructs for us the two intersections of the line with the unique conic through the five points.

Model 4 Given four skew lines in 3-space, the projective compass constructs for us their two common transversals.

Of course, if we are working over a field in which some scalars don't have square roots, then the projective compass may fail on some problem instances; that is true for any of the four models of a projective compass.

Exercise 10.2-1 Show that the four models of a projective compass have equivalent power.

[Hint: Implementing Model 2 with Model 1 is trivial.

To implement Model 3 with Model 2, suppose that we are given five points and a line ℓ. Construct some point P on the conic through the five given points and construct the tangent line m to the conic at P, both of which we can do with just a straightedge. Also with a straightedge, we can construct the second tangent m' to the conic that passes through the point $\ell \cap m$ and the point P' where m' touches the conic. The mapping that takes P to P' is an involution on the conic, and the fixed points of that involution give us the points where the line ℓ cuts the conic.

To implement Model 1 with Model 3, note first that we can think of a projectivity as acting either on the points in a range or on the lines in a pencil. Using the second form, suppose that we are given two triples (a, b, c) and (a', b', c') of lines through a point P and that we want to construct the two fixed lines of the projectivity of the P pencil that takes (a, b, c) to (a', b', c'). Let Q be some point and let ℓ be some line. Given any line e through P, let $e^\star$ be the line that joins Q to $e \cap \ell$. There is a unique projective correspondence from the P pencil to the Q pencil that takes (a, b, c) to $(a'^\star, b'^\star, c'^\star)$, and the intersections of corresponding pairs of lines

$(e, e'^{\star})$ trace out a conic. The two points where the line ℓ cuts that conic give us the two fixed lines of the original projectivity.

To implement Model 4 with Model 1, let a, b, c, and d be the four skew lines. The one-parameter family of all common transversals of a, c, and d determines one projective correspondence between c and d, while the similar family of all common transversals of b, c, and d determines another. If we map from c to d via one of those correspondences and then back from d to c via the other, we get a projectivity of the line c whose two fixed points give us the two common transversals.

Finally, to implement Model 1 with Model 4, let $P \mapsto P'$ be the given projectivity of the line ℓ, and choose two lines m and n that are skew to each other and to ℓ. For each point P on ℓ, let $P^{\star}$ denote the point on m that lies in the plane $\mathrm{Span}(P, n)$. As the point P varies along ℓ, the lines of the form $P'P^{\star}$ sweep out a ruled quadric. The lines ℓ and m are generators of that same quadric, from the other family. Let k be a third generator from that other family, skew to ℓ and m. The common transversals of k, ℓ, m, and n give us the two fixed points of the original projectivity.]

Which model of the projective compass is most convenient depends upon the application. When constructing a representation of $B_{1,1,1,1}$ in Proposition 10.1-2, we needed to find a common transversal of four skew lines. But there is an important special case of the same construction where what we need to do, instead, is to find the fixed points of a projectivity.

Suppose that the A-points and B-points in Proposition 10.1-2 are chosen so that the cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ are equal. This equality does not hold in general, but it might be desirable to have it hold in certain cases — for example, when trying to construct a representation of $B_{1,1,1,1}$ with lots of geometric symmetries. When those two cross ratios are equal, the four null systems (N_i) coincide, so the four lines (c^{*i}) all coincide with the polar c^* of the line c in that single null system. It follows from Lemma 10.1-1 that we must take d to be the line c^*. But where on the lines c and $d = c^*$ should we put the C-points and the D-points?

One way to figure out where is as follows: Choose any point along the line c, and tentatively call it C_1. Using one of the perfect coplanarities $\{A_3, B_4, C_1, D_2\}$ or $\{A_4, B_3, C_1, D_2\}$, figure out where D_2 should go along the line d, given that C_1 is as we have guessed. In a similar way, compute the points C_3, D_4, and C_1 again, each from the previous one. If the final point C_1 coincides with our initial guess for C_1, we were very lucky, and we have found a representation of $B_{1,1,1,1}$. Typically, the initial and final values of C_1 will not coincide. But the mapping that takes our initial guess for C_1 to the resulting final value is some projectivity of the line c. We can characterize that projectivity uniquely by carrying out this four-step tracing process three times, starting with three distinct, initial guesses for C_1. We then employ a Model-1 projective compass to construct the two fixed points of that projectivity — those two points being the two places where C_1 can go in a valid representation.

10.3 Representations with Euclidean symmetries

If all of the column budgets of a budget matroid are equal, the columns play symmetric roles. It is tempting to try to find representations of such matroids in which as many as possible of the combinatorial symmetries that permute the columns are modeled by geometric symmetries — that is, by projective transformations of the ambient space that map the representation to itself.

Exercise 4.1-1 considered this issue for the Pappus matroid $B_{1,1,1}$, in which the three columns play symmetric roles. For the representations of $B_{1,1,1}$ discussed in that exercise, all six permutations of the three columns can be achieved by projective transformations — even by Euclidean transformations — of the ambient plane. Furthermore, all six of those transformations leave the rows fixed. In this section and the next, we see how close we can come to that same behavior for the matroid $B_{1,1,1,1}$ in 3-space.

As a warm-up for the projective case, let's first consider representing the matroid $B_{1,1,1,1}$ in Euclidean 3-space so that as many as possible of the 24 column permutations are achieved by Euclidean transformations. We shan't do very well with Euclidean transformations — in particular, we shall handle only those 6 of the 24 permutations that leave one of the four columns fixed. But at least the resulting representations are easy to visualize, as demonstrated in the accompanying videotape [44].

The basic idea is as follows: Given a hyperboloid of one sheet with circular cross sections, we choose three of the four column lines of the representation — say b, c, and d — to be generating lines of that hyperboloid, so arranged that they are cyclically permuted by a 120° rotation around the axis of the hyperboloid — which we choose to be the remaining column line, a. One convenient way to set this up in coordinates is to choose, as the hyperboloid, the surface $XY+XZ+YZ+1=0$ in Euclidean (X, Y, Z)-space. That surface is rotationally symmetric about the main diagonal a and it includes, as one generating line, the line b parameterized by $t \mapsto (t, -1, 1)$. Cyclically permuting the three coordinates X, Y, and Z corresponds to rotating this hyperboloid by 120° about its axis, which also cyclically permutes the generating line $b\colon t \mapsto (t, -1, 1)$ and its two cyclic shifts $c\colon t \mapsto (1, t, -1)$ and $d\colon t \mapsto (-1, 1, t)$.

Given any point (X, Y, Z), let us refer to the sum of its coordinates $X+Y+Z$ as its *height*. Note that the planes of constant height are orthogonal to the main diagonal a, and that the hyperboloid $XY+XZ+YZ+1=0$ is symmetric under reflection through the plane of height 0. We envision a representation of $B_{1,1,1,1}$ of the form:

	A	B	C	D
1	(q_1, q_1, q_1)	$(p_1, -1, 1)$	$(1, p_1, -1)$	$(-1, 1, p_1)$
2	(q_2, q_2, q_2)	$(p_2, -1, 1)$	$(1, p_2, -1)$	$(-1, 1, p_2)$
3	(q_3, q_3, q_3)	$(p_3, -1, 1)$	$(1, p_3, -1)$	$(-1, 1, p_3)$
4	(q_4, q_4, q_4)	$(p_4, -1, 1)$	$(1, p_4, -1)$	$(-1, 1, p_4)$

The parameter p_i, for i in $[1\,..\,4]$, is the common height of the points B_i, C_i, and D_i, while the point A_i lies at height $3q_i$ along the main diagonal.

For arbitrary values of the parameters (p_i), there are unique values for the parameters (q_i) that make the 24 perfect coplanarities of the matroid $B_{1,1,1,1}$ hold. To determine the requisite value of q_i, we intersect the plane $B_j C_k D_l$ with the main diagonal a, where $\{i, j, k, l\} = \{1, 2, 3, 4\}$.

Exercise 10.3-1 Verify that the plane $B_j C_k D_l$ intersects the main diagonal at the point (q_i, q_i, q_i), where

$$q_i := \frac{(p_j + p_k + p_l) + p_j p_k p_l}{9 + (p_j p_k + p_j p_l + p_k p_l)}.$$

Because of the three-fold rotational symmetry, the formula for q_i clearly had to be symmetric under cyclic shifts of the indices (j, k, l). It turns out to be symmetric under arbitrary permutations of (j, k, l). Thus, if we set $A_i := (q_i, q_i, q_i)$ for each i, where q_i is as given in Exercise 10.3-1, the resulting configuration has all of the incidences that are required for a representation of the matroid $B_{1,1,1,1}$. It may also have forbidden incidences, of course, depending upon our choices for the parameters (p_i).

We can double the number of Euclidean symmetries, from three to six, by stipulating that $p_1 + p_4$ and $p_2 + p_3$ are both zero, from which it follows that $q_1 + q_4$ and $q_2 + q_3$ will also be zero. Under this stipulation, a 180° rotation about the line $X = Y + Z = 0$ — which maps the point (X, Y, Z) to the point $(-X, -Z, -Y)$ — then maps A_i to A_{5-i}, B_i to B_{5-i}, C_i to D_{5-i}, and D_i to C_{5-i}, for all i in $[1\,..4]$. That is, the A and B columns are held fixed while the C and D columns are swapped, but at the cost of uniformly swapping row 1 with row 4 and row 2 with row 3.

There is no way that a Euclidean symmetry of the configuration could avoid fixing the axis a of the hyperboloid; so we aren't going to be able to achieve more than 6 of the 24 column permutations with Euclidean symmetries. Hence, we might as well set the two remaining free parameters p_1 and p_2 to achieve some lesser goal, such as beauty. For example, we can set $p_1 := 3p_2$, so that the B-points, C-points, and D-points are equally spaced, along their lines. We can then choose the remaining parameter p_2 so that the A-points also end up equally spaced, along the main diagonal.

Exercise 10.3-2 Assuming that we set $p_1 := 3p_2$, $p_3 := -p_2$, and $p_4 := -3p_2$, find the positive, real values of p_2 that result in the four A-points being equally spaced, in some order, along the main diagonal a.

[Answer: The values $p_2 = (\sqrt{17} \pm 2\sqrt{2})/3$ make the four A-points equally spaced in the order (A_1, A_3, A_2, A_4). But the resulting configurations are not representations of $B_{1,1,1,1}$, because they have twelve forbidden coplanarities, of which $\{A_1, B_1, C_4, D_1\}$ is an example. Luckily, the values $p_2 = (\sqrt{53} \pm 2\sqrt{2})/3\sqrt{5}$ also make the A-points equally spaced, this time in the order (A_3, A_1, A_4, A_2). The configurations that result from these two values have no forbidden incidences, and hence do represent $B_{1,1,1,1}$. The representation shown in the videotape [44] has $p_2 = (\sqrt{53} - 2\sqrt{2})/3\sqrt{5} = 0.6636^+$.]

This is a convenient time to interject a comment about cross ratios. Recall that there is no representation of the budget matroid $B_{2,1,1}$ in which the two cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ along the two column lines are equal. Given that, one might wonder whether there is any constraint on the four cross ratios $A_{(1,2,3,4)}$, $B_{(1,2,3,4)}$, $C_{(1,2,3,4)}$, and $D_{(1,2,3,4)}$ in a representation of the matroid $B_{1,1,1,1}$. The answer seems to be no. By varying p_2 continuously in the construction of the previous exercise, we can keep the last three of those four fixed while varying the first continuously; so there can't be any algebraic constraint relating the four cross ratios. In that example, of course, the last three are not only fixed, but also equal. Generically, however, we know that no two of the four cross ratios are equal, since that is what happened in the example in Proposition 10.1-2. Indeed, in that example, the six possible cross ratios that we can get by taking the four A-points in some order together with the six similar cross ratios from the other three lines constitute 24 distinct scalars. Thus, generically, even if we allow ourselves to reorder the points, no cross ratio along any column line coincides with any cross ratio along any other column line. On the other hand, it can also happen that all four cross ratios are equal, with no reordering allowed, as we see in the next section.

Exercise 10.3-3 Are there representations of the budget matroid $B_{1,1,1,1}$ in which all sixteen points lie on a common quadric? In such a representation, the four column lines must be generating lines of that quadric, all drawn from the same family.

[Answer: Yes — at least, the authors' Newton-Raphson search converged on a convincing approximation to such a representation.]

Exercise 10.3-4 Continuing the theme of the preceding exercise, suppose that the sixteen points $\{(A_i, B_i, C_i, D_i)\}_{i\in[1..4]}$ form a representation of the matroid $B_{1,1,1,1}$ in which all sixteen points lie on a common quadric. Show that the four cross ratios $A_{(1,2,3,4)}$, $B_{(1,2,3,4)}$, $C_{(1,2,3,4)}$, and $D_{(1,2,3,4)}$ along the four column lines must be distinct, and hence this representation is not very symmetric.

[Hint: Given three skew lines, all of their common transversals form a one-parameter family that is called — at least, was once called [29] — a *regulus*. The lines of a regulus sweep out a ruled quadric surface, forming one of that surface's two families of generating lines. The other family of generating lines of the same ruled quadric is called the *complementary regulus*.

Getting back to the exercise, suppose, by way of contradiction, that the cross ratios $A_{(1,2,3,4)}$ and $B_{(1,2,3,4)}$ are equal. The four lines $\{A_i B_i\}_{i\in[1..4]}$ then belong to a common regulus whose complement — call it R — contains both a and b. The lines c and d are constrained in two ways: They are polars of each other in the single null system $N = N((A_i, B_i), (A_j, B_j), (A_k, B_k))$ that results from any subset

$\{i, j, k\} \subset \{1, 2, 3, 4\}$, and they also belong to some regulus R' that contains both a and b. Show that those two constraints imply that the lines (c, d) are harmonic conjugates of the lines (a, b) in the single regulus $R' = R$. But, if $R' = R$, the lines $(A_i B_i)$ meet the lines c and d, which cannot happen because, for example, the set $\{A_1, B_1, C_1, C_2\}$ is forbidden to be coplanar.]

10.4 Representations with projective symmetries

We now turn from Euclidean 3-space to projective 3-space, which gives us a much richer class of transformations to exploit. We withdraw our decision to set $p_1 := 3p_2$ — that decision was motivated only by a quest for beauty, after all — but we keep everything else the same as before, including the constraints $p_1 + p_4 = 0$ and $p_2 + p_3 = 0$. The key idea in the projective case is to restrict our choices of the parameters (p_i) so that the identity

$$p_1p_2p_3p_4 + (p_1p_2 + p_1p_3 + p_1p_4 + p_2p_3 + p_2p_4 + p_3p_4) + 9 = 0$$

holds. Note that this identity is symmetric in the four variables (p_i). If this identity holds, it follows from Exercise 10.3-1 that we have $q_i = -1/p_i$ for all i. Thus, converting from Euclidean coordinates (X, Y, Z) to homogeneous coordinates $[w, x, y, z]$, we can write our configuration in the form:

	A	B	C	D
1	$[-p_1, 1, 1, 1]$	$[1, p_1, -1, 1]$	$[1, 1, p_1, -1]$	$[1, -1, 1, p_1]$
2	$[-p_2, 1, 1, 1]$	$[1, p_2, -1, 1]$	$[1, 1, p_2, -1]$	$[1, -1, 1, p_2]$
3	$[-p_3, 1, 1, 1]$	$[1, p_3, -1, 1]$	$[1, 1, p_3, -1]$	$[1, -1, 1, p_3]$
4	$[-p_4, 1, 1, 1]$	$[1, p_4, -1, 1]$	$[1, 1, p_4, -1]$	$[1, -1, 1, p_4]$

Given a configuration of this form, what column permutations can we achieve via projective symmetries? The six column permutations that fix the A column can all be achieved by the same Euclidean transformations as before. In addition, the projective transformation $[w, x, y, z] \mapsto [x, -w, -z, y]$ swaps the point A_i with B_i and C_i with D_i, for all i in $[1\,..\,4]$. Thus, we can achieve all 24 column permutations via projective symmetries. The symmetries that achieve the even column permutations leave the rows fixed; those that achieve the odd permutations swap row 1 with row 4 and swap row 2 with row 3.

It remains to choose the parameters (p_i), being careful to avoid any forbidden incidences. So far, we have three constraints on those four parameters. Perhaps the prettiest way to tie down the final degree of freedom is to require that $p_1 = 1/p_2$, so that the set $\{p_1, p_2, p_3, p_4\}$ is closed under both negation and inversion. A little algebra reveals that we then have

$$
\begin{aligned}
p_1 &= +\sqrt{3} + \sqrt{2} \\
p_2 &= +\sqrt{3} - \sqrt{2} \\
p_3 &= -\sqrt{3} + \sqrt{2} \\
p_4 &= -\sqrt{3} - \sqrt{2};
\end{aligned}
$$

and some further algebra verifies that these choices do not lead to any forbidden incidences. The resulting representation of $B_{1,1,1,1}$ seems like a good candidate to win the projective beauty contest.

Chapter 11

Open questions

We wrap up by discussing some open questions about the budget matroids and their relationship to n-dependency.

11.1 Representability in general

Since what makes the budget matroids interesting is that so many of them are representable, the most tempting challenge about the budget matroids is to determine precisely which of them are representable.

Challenge 11.1-1 Determine which of the budget matroids are representable over the complex numbers. For those that are representable, count the degrees of freedom involved in choosing a representation — that is, compute the dimension of the set of representations. Also, if possible, show how to construct a generic representation using a flat-side as the only geometric tool — that is, find a rational parameterization of the set of representations.

One could ask the same questions for the budgetary matroids as well. But the evidence in Chapter 5 suggests that most of them are unrepresentable, so the answers are less likely to be interesting.

Of course, the set of representations of a budget matroid has to be more than just a set, in order for its dimension to be well defined. The books by Harris [14], by Reid [46], and by Cox, Little, and O'Shea[7] are fine places to learn algebraic geometry at the level that you need to follow the next few paragraphs. It is simplest to assume, at this point, that our field of scalars is algebraically closed.

Let $B_{b_1,\ldots,b_k}$ be a budget matroid and let S denote projective space of dimension $b-1$, where $b := b_1 + \cdots + b_k$ is the total budget. Each incidence that the matroid $B_{b_1,\ldots,b_k}$ requires is a polynomial constraint on the homogeneous coordinates of our bk points in S. Those constraints determine a certain variety W in the Cartesian product S^{bk} of bk copies of S. (If we like, we can use the Segre embedding [16] to view the Cartesian product S^{bk} as a subvariety of projective space of dimension

$b^{bk}-1$, thus realizing W also as a subvariety of that single, high-dimensional projective space — so W is a projective variety.) Each point in the variety W gives a way of mapping the elements of the matroid $B_{b_1,\dots,b_k}$ to points in S so that all dependent sets of elements map to sets of points that are mutually incident — but various independent sets of elements may also map to sets of points that are mutually incident. Such a map is called a *weak representation*.

The variety W of weak representations is typically reducible, but we can decompose it into its irreducible components. Each of the irreducible components V of W is either good or bad:

good The bulk of the points in V give true representations; only those in a proper subvariety of V have forbidden incidences.

bad All of the points in V give weak representations that share some forbidden incidence.

We define the freedom $\#(B_{b_1,\dots,b_k})$ to be the maximum of the dimensions of the good components — unless all of the components are bad, in which case we set $\#(B_{b_1,\dots,b_k}) := \perp$.

Note that there typically exist bad components whose dimension is larger than that of any good component. For example, consider the matroid $B_{1,1,1,1}$. On the good side, we have $\#(B_{1,1,1,1}) = 20$. On the other hand, one way to build a weak representation is to let the three column lines a, b, and c coincide, which is enough to ensure that all of the perfect coplanarities hold. There are four degrees of freedom in the choice of the line $a = b = c$, four more in the skew line d, and one in each of the sixteen points along its column line, for a total of 24. Thus, speaking loosely, there are many more weak representations than there are true representations — many more degenerate cases than nondegenerate ones.

Exercise 11.1-2 A simpler flavor of weak representation of the budget matroid $B_{1,1,1,1}$ has all sixteen of its points lying in a common plane. How many degrees of freedom are involved in choosing one of those?

[Answer: 27, assuming that we still take the ambient projective space to be a 3-space, as opposed to a plane.]

Let $B_{b_1,\dots,b_k}$ be a budget matroid that is representable, let W be its variety of weak representations, and consider those irreducible components V of W that are good. It seems reasonable to hope, even if several good components V exist, that only one of them will have the maximal dimension $\#(B_{b_1,\dots,b_k})$. Whenever that is true, we have a particular, irreducible variety V that we can think of, loosely, as being 'the variety of representations' of the matroid. We can then go on to ask whether that variety V is either *rational*, meaning that V is birationally isomorphic to projective space of dimension $\#(B_{b_1,\dots,b_k})$, or — failing that — perhaps *unirational*, meaning that there exists a dominant rational map from projective space of dimension $\#(B_{b_1,\ldots,b_k})$ to V. Roughly speaking, either of those two conditions means that we can express the coordinates of the points in a generic representation as rational functions of $\#(B_{b_1,\ldots,b_k})$ free parameters. For the variety V to be rational, the correspondence between vectors of parameter values and representations must be generically one-to-one, whereas even a many-to-one correspondence is enough to qualify V as unirational [17].

In the $B_{m,n}$ and $B_{m,1,1}$ Representation Theorems, we constructed generic representations of the budget matroids $B_{m,n}$ and $B_{m,1,1}$ using a flat-side as our only geometric tool. In such a construction, the points of the representation are chosen in some order. When each point is chosen, the required incidences of the matroid force that point to lie in a certain k-flat, which can be constructed from the previously chosen points using only a flat-side. (If the dimension k of that flat is zero, there isn't actually any choice about where the new point will go.) When all points have been chosen in their appropriate flats, all of the required incidences of the matroid are guaranteed to hold and, in generic cases, no forbidden incidences hold. A construction of this type gives a birational isomorphism between projective space of the appropriate dimension and the variety of representations of the matroid. That is, when such a flat-side construction exists, precisely one irreducible component of the variety of weak representations is good and that unique good component is a rational variety.

It would be delightful if all of the budget matroids were representable and if a generic representation of any budget matroid could be constructed using only a flat-side. But that is a lot to hope for.

11.2 The budget matroids with few columns

What are the smallest budget matroids whose representability is still open? There are several answers to that question, depending upon what ordering we impose on the budget matroids. Two reasonable choices are to order them by the number of columns or by the number of rows.

The $B_{m,n}$ Representation Theorem resolves the representability issues for all budget matroids with two columns: They are all representable, flat-side constructions always exist, and the freedom involved is $\#(B_{m,n}) = m^2+3mn+n^2-3$. The $B_{m,1,1}$ Representation Theorem gives similar good news about those budget matroids with three columns, two of whose column budgets are ones: Representable, flat-side constructions exist, and the freedom is $\#(B_{m,1,1}) = m^2 + 6m + 3$. The authors suspect that the latter result can probably be extended to the case where only one of the column budgets is fixed at 1.

Conjecture 11.2-1 *For every positive m and n, the budget matroid $B_{m,n,1}$ is representable over the rationals, a generic representation can be constructed using only a flat-side, and the number of degrees of freedom involved in choosing a representation is $\#(B_{m,n,1}) = (m + n + 2)^2 - 6$.*

The authors have checked this conjecture numerically for various cases, up through $B_{3,3,1}$. As for that particular case, we said in Section 5.5.6 that $\#(B_{3,3,1}) \doteq 58$, which is consistent with the conjecture.

One might be tempted to broaden Conjecture 11.2-1 to say that $\#(B_{m,n,p}) = (m + n + p + 1)^2 - 6$ for any three positive column budgets m, n, and p. But that broadened conjecture doesn't square with the authors' numeric experiments, which suggest that $\#(B_{2,2,2}) \doteq 42$, rather than 43.

Exercise 11.2-2 Calculate the freedom $\#(B_{m,n,0})$ and compare your result with that predicted by the broadened form of Conjecture 11.2-1. When does the broadened conjecture get the right answer?

[Answer: The $B_{m,n}$ Representation Theorem tells us that $\#(B_{m,n}) = m^2 + 3mn + n^2 - 3$; so we have $\#(B_{m,n,0}) = m^2 + 3mn + n^2 + m + n - 4$. The broadened conjecture predicts that $\#(B_{m,n,0}) = (m+n+1)^2 - 6$, which is correct only when $(m - 1)(n - 1) = 0$ — that is, only in those cases that are covered by the original, unbroadened form of Conjecture 11.2-1. Of course, zero column budgets might have to be treated as a special case.]

For budget matroids with four columns, very little is known. Indeed, that little can be summed up in the two equations $\#(B_{1,1,1,1}) = 20$ and $\#(B_{2,1,1,1}) \doteq 30$. And nothing at all is known about the representability of particular budget matroids with more than four columns.

11.3 The budget matroids of low rank

Instead of ordering the budget matroids by the number of columns, we can order them by the number of rows — that is, by rank.

The budget matroids of rank at most 4 are all representable over the complex numbers, and we know the freedom involved in each case:

$$\begin{array}{cc}
\#(B_{1,1}) = 2 & \#(B_{3,1}) = 16 \\
 & \#(B_{2,2}) = 17 \\
\#(B_{2,1}) = 8 & \#(B_{2,1,1}) = 19 \\
\#(B_{1,1,1}) = 10 & \#(B_{1,1,1,1}) = 20.
\end{array}$$

Furthermore, there is a flat-side construction for a generic representation in every case except possibly for $B_{1,1,1,1}$.

Challenge 11.3-1 Proposition 10.1-2 uses a flat-side and a projective compass to construct a generic representation of the budget matroid $B_{1,1,1,1}$. By making the choices in some different order, is it possible to achieve the same result with only a flat-side?

For the budget matroids of rank 5, we have the following freedoms:

$$\begin{aligned}
\#(B_{4,1}) &= 26\\
\#(B_{3,2}) &= 28\\
\#(B_{3,1,1}) &= 30\\
\#(B_{2,2,1}) &\doteq 30\\
\#(B_{2,1,1,1}) &\doteq 30\\
\#(B_{1,1,1,1,1}) &= ?.
\end{aligned}$$

The first three are covered by the $B_{m,n}$ and $B_{m,1,1}$ Representation Theorems, so those three cases also have flat-side constructions. The authors have found rational representations of $B_{2,2,1}$ and of $B_{2,1,1,1}$. Given such a representation, it is straightforward to use linear algebra to calculate the local dimension of the space of representations at that point, which is a good guess for the true freedom involved. But the authors have been unable to find a rational — or even an approximate, complex — representation of the matroid $B_{1,1,1,1,1}$.

Challenge 11.3-2 Is the budget matroid $B_{1,1,1,1,1}$ representable over the complex numbers? Over the rationals?

There is a subtlety about the way that a budget matroid relates to its minors that should be discussed here, since it can impact the search for a representation of $B_{1,1,1,1,1}$. In particular, a minor of a generic representation is not necessarily a generic representation of the minor.

One way to search for a representation of some matroid, such as $B_{1,1,1,1,1}$, is to use a numeric method on the system of nonlinear equations that encodes the required incidences. If we are lucky enough to find an approximate solution of the incidence equations, we can check that solution to determine whether any forbidden incidences also hold. The authors performed various experiments of this type, using a Newton-Raphson iteration as their numeric method.

The bane of this search strategy is degeneracy. Unless we do something clever to constrain the numeric search, it is overwhelmingly likely that whatever solution we converge to will have some forbidden incidences. Recall that there are typically many more weak representations than there are true representations.

How can we constrain the numeric search so as to improve the odds that it will converge to a solution that is free of forbidden incidences? Recall that a budget matroid of rank r has lots of budget matroids of rank $r-1$ as minors. One way to constrain the search is to fix the representation of one of those minors to be some particular, true representation of the smaller budget matroid. In some of the authors' experiments, this strategy worked splendidly.

But now we come to the subtle point: Just because any representation of the larger matroid includes, within it, a representation of the minor, it does not follow that any representation of the minor can be extended into a representation of the larger matroid. Instead, it may be the case that the bulk of the representations of the minor cannot be extended — the ones that can be extended being only those that satisfy certain algebraic constraints. In particular, the authors' search for a representation of $B_{1,1,1,1,1}$ may have been torpedoed at the start, because the particular representation of the minor $B_{1,1,1,1}$ that they chose to try to extend was not, in fact, extensible.

For concreteness, here is a blatant case of this phenomenon. We know from the $B_{m,n}$ Representation Theorem that the matroid $B_{9,9}$ is representable and that $\#(B_{9,9}) = 402$. Assuming that Conjecture 11.2-1 is true, the matroid $B_{9,9,1}$ is also representable, and $\#(B_{9,9,1}) = 394$. Given any representation of $B_{9,9,1}$, sitting in 18-space, if we place our eye at one of the points in the third column and look out at the points that are neither in our row nor in our column, we see a representation of $B_{9,9}$, sitting in the 17-space of lines through our eye. Since $394 < 402$, the representation of $B_{9,9}$ that we see must have some special properties. But what those special properties are is not clear.

Challenge 11.3-3 Which representations of the budget matroid $B_{1,1,1,1}$ are most likely to be extensible into representations of $B_{1,1,1,1,1}$?

11.4 Pushing points together

Before we turn to consider the relationship between the budget matroids and the notion of n-dependence, Jorge Stolfi contributes one more open problem about the budget matroids themselves.

The budget matroids whose column budgets are all ones are the most delicate to construct representations of, but they are also the most symmetric. In many cases, if we take two points in a representation of such a matroid and we push them together until they coincide, the resulting configuration consists of a representation of a simpler budget matroid of the same rank, along with some irrelevant trash. For example, consider the Pappus matroid $B_{1,1,1}$. If we push A_2 and A_3 together in the Pappus configuration of Figure 4.1, the point C_1 comes in to join them and we are left with a complete quadrilateral — that is, a representation of $B_{2,1}$ — formed by the six points

$$\begin{array}{cc} 2 & 1 \\ \multicolumn{2}{c}{\begin{pmatrix} A_1 & B_1 \\ C_2 & B_2 \\ C_3 & B_3 \end{pmatrix}} \end{array},$$

along with the trash point $A_2 = A_3 = C_1$.

Challenge 11.4-1 Is it always the case that, when two points in a representation of a budget matroid are pushed together in this way, what remains is a representation of another budget matroid of the same rank, but with fewer columns? Explore the extent to which the representations of different budget matroids of the same rank are degenerate cases of one another.

Someone tackling this challenge would probably benefit from studying the results of Bokowski and Sturmfels [6].

11.5 The matroids $B_{2,1,\ldots,1}$

We were led to study the budget matroids because the matroid $B_{2,1,1}$ characterizes the dependence of four cubic polynomials, just as $B_{2,1}$ characterizes the dependence of three quadratic polynomials. Sad to say, the obvious pattern does not continue: The matroid $B_{2,1,1,1}$ cannot characterize the dependence of five quartic polynomials in an analogous way, because the degrees of freedom don't work out properly.

Let's fix a plane σ in 4-space and then count. There are nineteen degrees of freedom in a 4-dependent block of 3-flats through σ — twenty scalar slopes, subject to one constraint. There are fifteen degrees of freedom in a projective transformation of 4-space that fixes the plane σ and fixes also every 3-flat through σ. To see why, note that there are twenty-four degrees of freedom in an arbitrary projective transformation of 4-space; but it costs four to fix one 3-flat through σ, four more to fix a second, and one more to fix all of the rest. Therefore, if a quartic analog of the Witness Theorem were to hold, the configuration involved must have $19+15=34$ degrees of freedom. But the authors' experiments, as reported above, indicate that there are only 30 degrees of freedom in a representation of $B_{2,1,1,1}$ — four less than the 34 required.

While the obvious quartic generalization of the Witness Theorem thus cannot hold, the quartic generalization of the Projection Theorem does seem to hold. For each of the rational representations of the matroid $B_{2,1,1,1}$ that the authors have investigated, projecting the twenty points of that representation from any plane σ always results in a block of twenty 3-flats whose slopes are 4-dependent.

By the way, here is an example rational representation of the matroid $B_{2,1,1,1}$, in case some readers would like to do their own experiments:

	P	A	B	C
1	$[12,16,76,-6,87]$	$[1,1,8,0,0]$	$[-7,0,0,-5,1]$	$[71,1,3,1,3]$
2	$[11,1,6,4,-3]$	$[-9,1,-2,0,0]$	$[-2,0,0,0,1]$	$[623,13,59,13,59]$
3	$[125,93,434,-62,558]$	$[0,1,7,0,0]$	$[0,0,0,2,1]$	$[86,1,2,1,2]$
4	$[97,4,20,2,15]$	$[-3,1,4,0,0]$	$[-3,0,0,-1,1]$	$[-79,1,13,1,13]$
5	$[1,0,0,0,0]$	$[-93,127,796,0,0]$	$[31,0,0,5,-13]$	$[991,71,483,71,483]$

Note that, if you put your eye at the point P_5 and look out at the sixteen points in the first four rows, what you see is the rational representation of $B_{1,1,1,1}$ that we used in proving Proposition 10.1-2.

Challenge 11.5-1 Prove the quartic case of the Projection Theorem. That is, given any representation of the budget matroid $B_{2,1,1,1}$ in projective 4-space and given a generic plane σ, prove that the slopes of the twenty 3-flats that join σ to the twenty points of the representation form a block of scalars that is 4-dependent.

Let's temporarily call an ordered 4-block of 3-flats through a plane σ in 4-space *nice* when a representation of $B_{2,1,1,1}$ does exist, each of whose twenty points lies on the appropriate 3-flat. Assuming that the quartic case of the Projection Theorem does hold, any nice block must be 4-dependent — the representation of $B_{2,1,1,1}$ serving as a witness to the 4-dependence of the block. But a nice block must have additional properties as well. Since any projective transformation of 4-space that fixes σ and fixes every 3-flat through σ clearly takes a witnessing representation to another such, each nice block has 15 dimensions' worth of witnesses. So there can be only $\#(B_{2,1,1,1}) - 15 \doteq 15$ degrees of freedom in a nice block. In addition to 4-dependence, there must be four more dimensions' worth of algebraic constraints on the twenty 3-flats in a nice 4-block.

Challenge 11.5-2 Determine all of the algebraic constraints that relate the slopes of the twenty 3-flats in a nice 4-block.

Challenge 11.5-3 Extend this theory to $n = 5$ and beyond, thereby revealing the connection between representations of the budget matroid $B_{2,1,\ldots,1}$ — where there are n columns, all but the first of whose budgets are ones — and ordered n-blocks of slopes that are 'nice'.

11.6 Characterizing 4-dependence

But wait a minute. Perhaps we are giving up too soon on our original problem. While the budget matroid $B_{2,1,1,1}$ characterizes some property — the property of being 'nice' — that is stronger than 4-dependence, perhaps there is some other configuration that characterizes 4-dependence itself. Since representations of $B_{2,1,1,1}$ do seem to yield 4-dependent blocks when projected, a natural way to try to design such a new configuration is to relax the constraints of the matroid $B_{2,1,1,1}$. Perhaps, if we relax constraints correctly, we can find a new matroid M, with $\#(M) = 34$, for which the Projection Theorem still holds.

One way to relax our constraints on the configuration would be to replace the budget matroid $B_{2,1,1,1}$ with one of its budgetary relaxations. Unfortunately, as we mentioned in Section 5.5.5, it seems that all of those budgetary relaxations are unrepresentable. All of them still insist on the sixty perfect incidences; and the rank of the Jacobian matrix of those sixty constraints, when evaluated at a rational representation of $B_{2,1,1,1}$, turns out to be 50. Thus, of the eighty degrees of freedom that are present in twenty arbitrary points in 4-space, fifty are taken away by the sixty perfect constraints themselves, leaving only thirty. But those same fifty are taken away by the combination of the sixty perfect constraints and all of the column constraints of the budget matroid $B_{2,1,1,1}$. Thus, it seems that the perfect constraints — in the presence of appropriate nondegeneracy conditions — imply the column constraints. It is this which makes all of the budgetary relaxations unrepresentable.

We might consider going further and relaxing away some of the perfect constraints also — perhaps fourteen of them: ten to remove the redundancy and four more to increase the freedom of the constrained configurations from 30 to 34. But there is good evidence of a different sort that no such approach could succeed.

Whatever we do, we want the Projection Theorem to continue to hold for our relaxed configurations. Given five quadruples of points in 4-space whose coordinates are independent variables, suppose that we constrain those twenty points by requiring that, when projected from each of m different random planes, the block of twenty slopes that results is 4-dependent. That gives us m nonlinear constraints on the twenty varying points — that is, m constraints on 80 variables. (In the authors' experiment, m was 70; but any m comfortably larger than 50 would do just as well, as we shall see shortly.)

If the twenty varying points happen to be located at the twenty vertices of a rational representation of $B_{2,1,1,1}$, then all m constraints are satisfied, because we are currently assuming the quartic case of the Projection Theorem. So we then compute the rank of the Jacobian matrix of the m constraints, at the point corresponding to some rational representation of $B_{2,1,1,1}$. We know that there must be at least 30 degrees of freedom in the variety that those m constraints define, since $\#(B_{2,1,1,1}) \doteq 30$. And lo, the rank of the Jacobian turns out to be 50; that is, there are precisely $80 - 50 = 30$ degrees of freedom, and no more.

This experiment gives good evidence that there is no way to get a configuration that characterizes precisely the notion of 4-dependence simply by relaxing the constraints of the matroid $B_{2,1,1,1}$. It doesn't matter which constraints we decide to relax. As soon as we relax enough constraints to increase the freedom in the constrained configurations, the Projection Theorem starts to fail.

It is still possible that there is some completely different configuration — still with no auxiliary points and whose twenty key points span all of 4-space — that does characterize precisely the notion of 4-dependence. The instances of that configuration would have to form a 34-dimensional variety, presumably completely different from the 30-dimensional variety of representations of $B_{2,1,1,1}$. One way to search for such a thing would be to take the system of m nonlinear equations above and to look for some solution of it where the rank of the Jacobian falls, not just from m down to 50, but all the way down to 46. Unfortunately, it isn't clear how to look for such a solution. And even if we found one, any theory that we then managed to build probably wouldn't have much to do with the budget matroids — so it is perhaps best if we leave that investigation for another time.

Bibliography

[1] Michael Artin. *Algebra*. Prentice Hall, 1991.

[2] Pages 29–30.

[3] Page 245.

[4] Pages 260–261.

[5] Anders Björner, Michel Las Vergnas, Bernd Sturmfels, Neil White, and Günter M. Ziegler. *Oriented Matroids*. Cambridge University Press, 1993.

[6] Jürgen Bokowski and Bernd Sturmfels. *Computational Synthetic Geometry*. Volume 1355 of Lecture Notes in Mathematics, Springer-Verlag, 1989.

[7] David Cox, John Little, and Donal O'Shea. *Ideals, Varieties, and Algorithms: An Introduction to Computational Algebraic Geometry and Commutative Algebra*. Springer-Verlag, 1992.

[8] H. S. M. Coxeter. *Projective Geometry*, second edition. Republished by Springer-Verlag, 1987.

[9] Pages 21 and 42.

[10] Page 28.

[11] Page 37.

[12] Page 46.

[13] Phillip Giffiths and Joseph Harris. *Principles of Algebraic Geometry*, pages 760–761. John Wiley & Sons, 1978.

[14] Joe Harris. *Algebraic Geometry: A First Course*. Springer-Verlag, 1992.

[15] Pages 9–14.

[16] Pages 25–31.

[17] Pages 78 and 87.

[18] Allan Heydon and Greg Nelson. The Juno-2 constraint-based drawing editor. Research report 131a from the Systems Research Center of the Digital Equipment Corp., 1994.

[19] D. Hilbert and S. Cohn-Vossen. *Geometry and the Imagination*. Second edition, translated from the German by P. Nemenyi, Chelsea, 1952.

[20] Pages 94–98, 133–134.

[21] Page 134.

[22] Kiyosi Itô, editor. *Encyclopedic Dictionary of Mathematics*, second edition. MIT Press, 1987.

[23] Article 343D, Projective Transformations in Projective Geometry.

[24] Article 348A, General Remarks about Quadratic Forms.

[25] Serge Lang. *Algebra*, third edition. Addison-Wesley, 1993.

[26] Pages 586–587.

[27] Page 577.

[28] E. A. Maxwell. *General Homogeneous Coordinates in Space of Three Dimensions*. Cambridge University Press, 1959.

[29] Pages 54–56 and 72–74.

[30] Pages 57–58.

[31] Pages 69–71 and 91–92.

[32] Pages 154–156.

[33] E. A. Maxwell. *The Methods of Plane Projective Geometry Based on the Use of General Homogeneous Coordinates*, pages 119–120. Cambridge University Press, 1963.

[34] Thomas Muir. *A Treatise on the Theory of Determinants*, revised and enlarged by William H. Metzler, page 165. Dover edition, 1960.

[35] Marc A. Najork and Marc H. Brown. A library for visualizing combinatorial structures. Research report 128a from the Systems Research Center of the Digital Equipment Corp., 1994.

[36] Marc A. Najork. Obliq-3D tutorial and reference manual. Research report 129 from the Systems Research Center of the Digital Equipment Corp., 1994.

[37] James G. Oxley. *Matroid Theory*. Oxford University Press (1992).

[38] Page 43.

[39] Page 76.

[40] Pages 267, 130, and 19, respectively.

[41] Page 304.

[42] Page 514.

[43] Lyle Ramshaw. Blossoms are polar forms. *CAGD* **6**(4): 323–358, 1989.

[44] Lyle Ramshaw. Introducing the budget configurations; a 46-minute videotape. Research report 134b from the Systems Research Center of the Digital Equipment Corp., 1995.

[45] Lyle Ramshaw. Suitening our nomenclature. *SIGACT News* **20**(1): 60–61, 1989.

[46] Miles Reid. *Undergraduate Algebraic Geometry*. Cambridge University Press, 1988.

[47] Pages 85–86.

[48] Pierre Samuel. *Projective Geometry*. Translated from the French by Silvio Levy, Springer-Verlag, 1988.

[49] Page 97.

[50] Jorge Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computations*. Academic Press, 1991.

[51] Pages 44–45.

[52] D. J. A. Welsh. *Matroid Theory*. Academic Press, 1976.