30. + (Basic knowledge of linear algebra required.) Let V be a finite dimensional vector space, and let us associate a positive cost $c(v)$ to each vector $v \in V$. Let us now select a basis for V in the greedy way: First, select the cheapest vector, then select the cheapest vector that is linearly independent of the first chosen vector, and so on. In step i, choose the cheapest vector that is not in the subspace spanned by the $i-1$ vectors that were previously chosen. Stop when $dim\ V$ vectors are selected.

    Will the obtained basis B indeed be a minimum-cost basis?

31. Let A be a 1000×1000 matrix whose entries are equal to 0 or 1 so that there are at least 20000 entries equal to 1. Prove that at least one of the following two statements holds.
    (a) The matrix A has a row that contains at least 100 entries equal to 1, or
    (b) the matrix A has at least 100 rows so that each of them contains at least ten entries equal to 1.

6.7 Solutions to exercises

1. We prove this by induction on n, the initial case of $n = 2$ being easy to verify. Take an edge of G with endpoints A and B. If the sum of the degrees of A and B is at least $2n+1$, then A and B have a common neighbor C, and we are done. Otherwise, remove A, B, and all the edges adjacent to them. The remaining graph has $2n-2$ vertices and at least $n^2+1-(2n-1) = (n-1)^2+1$ edges, so by induction, it contains a triangle.

2. (a) (Part (c) implies this part, but we also provide a different proof.) Take all bipartite subgraphs of G in all possible ways, and count their edges. In other words, count all pairs (e, B) where e is an edge of the bipartite subgraph B of G.

    Let n be the number of vertices of G. On one hand, each edge of G is part of 2^{n-2} bipartite subgraphs, so the number of these pairs is $e(G)2^{n-2}$. Indeed, the number of ways to split the set of $n-2$ vertices disjoint from a given edge into an ordered pair of two subsets is 2^{n-2}. On the other hand, the number of bipartite subgraphs B of G that have at least one edge is $2^{n-1}-1$. Therefore, by the pigeonhole principle, there must be at least one B that has at least $\frac{e(G)2^{n-2}}{2^{n-1}-1} > \frac{e(G)}{2}$ edges.

    (c) (This implies (b).) Partition the vertex set into two blocks, Q and R, and disregard the edges within each block. If the obtained bipartite graph B satisfies the criterion of part (c), then we are done. If not, then there is a vertex x that violates that criterion. That means that more than half of the edges adjacent to x are in the same block as x. Therefore, moving x into the *other block*, the number of edges of B increases. Keep doing this as long as there is a vertex violating the criterion of part (c). The procedure eventually has to stop, since the number of edges of B cannot grow for an infinitely long time, and it does grow in each step. When the procedure stops, there are no vertices x violating the criterion, which proves the statement.

3. It suffices to prove that, on any interval $[a,b]$ within the allowed limits, the function $g(x)=\binom{x}{r}+\binom{a+b-x}{r}$ has a unique minimum at $x=(a+b)/2$. Set $n=a+b$ to simplify the formulae. Then we need to prove that the unique root of $g'(x)$ is at $x=n/2$, and that at that root $g'(x)$ changes from being negative to being positive.

   Apply induction on r, the case of $r=1$ being true. Note that, for symmetry reasons, we can assume that $x\leq n/2$. Also note that

$$\begin{aligned} g'(x) &= \left(\binom{x}{r-1}\cdot\frac{x-r+1}{r}+\binom{n-x}{r-1}\cdot\frac{n-x-r+1}{r}\right)' \\ &= \frac{\binom{x}{r-1}-\binom{n-x}{r-1}}{r}+\binom{x}{r-1}'\cdot\frac{x-r+1}{r} \\ &+ \binom{n-x}{r-1}'\cdot\frac{n-x-r+1}{r}. \end{aligned}$$

   It is now routine to verify that $x=n/2$ is indeed a root of $g'(x)$. Note that $\binom{x}{r}$ is increasing if $x>r-1$. This, with the induction hypothesis and the above inequality, imply that $g'(x)>0$ if $x>n/2$, and $g'(x)<0$ if $x<n/2$. This completes the proof.

4. We claim that $\chi(G')=\chi(G)$ if G has at least one edge. (Otherwise, $\chi(G)=1$ and $\chi(G')=2$.) Indeed, if f is a proper coloring of G, then let $f(x)+1$ be the color of x', where addition is modulo $\chi(G)$. That is, if $f(x)=\chi(G)$, then $f(x')=1$, and leave the color of x unchanged. This gives a proper coloring of G' with $\chi(G)$ colors.

5. (a) Assume without loss of generality that $k=\chi(G)\leq\chi(H)$. Let f be any k-coloring of G. Let $f(g)$ be the color of all vertices (g,h). This results in a proper k-coloring of $G\times H$ since two vertices of this graph are adjacent only if their G-coordinates are adjacent in G.

   (b) Yes. The induced subgraph of $G\times G$ whose vertices are of the form (g,g) is isomorphic to G, so we do need $\chi(G)$ colors to color that subgraph.

6. Color the vertices with colors from [4] the *greedy way* as follows: Go in nonincreasing order of vertex degrees, and for each vertex, use the

smallest color that keeps the coloring proper. If $d_1 \geq d_2 \geq \cdots \geq d_{10}$ are the degrees, then for vertex i the number of forbidden colors in the above coloring is $\min(d_i, i-1)$. Indeed, vertex i has d_i neighbors, of which at most $i-1$ are colored before i. In the graph at hand, $\max_{i=1}^{10} \min(d_i, i-1) = 3$, and the maximum is taken at the fourth vertex. So there are never more than three forbidden colors, proving that $\chi(G) \leq 4$.

7. In order to see that $\chi(G) = 4$ is possible, take three vertex-disjoint edges and a triangle. This disconnected graph has chromatic number three because of the triangle it contains. Now connect all the nine vertices of this graph to a tenth vertex, creating a graph G with the prescribed degree sequence and chromatic number four.

   To see that $\chi(G) = 3$ is possible, take three vertex-disjoint copies of $K_{2,1}$ and connect all nine of their vertices to a new tenth vertex.

   Finally, if $\chi(G) = 2$ were possible, then there would be a bipartite graph with this degree sequence. That is impossible, since the vertex with degree nine would have to be connected to all other vertices.

8. The graph shown in Figure 6.15 shows that $\binom{n-1}{2} + 1$ edges are possible. On the other hand, more edges are not possible. Indeed, assume G has at least $\binom{n-1}{2} + 2$ edges and a cut-point x. Then $G - x$ has $n-1$ vertices, at least $\binom{n-2}{2} + 1$ edges, and is not connected. This is impossible, as Supplementary Exercise 5 asks you to prove.

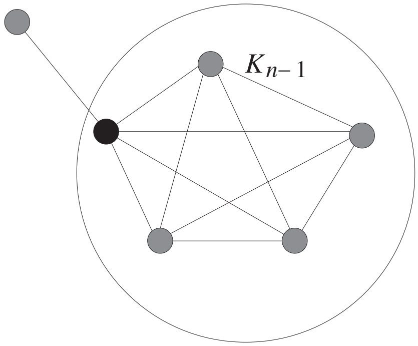

Figure 6.15
A connected, but not 2-connected, graph on n vertices and $\binom{n-1}{2}+1$ edges.

9. (a) Let us use colors $1, 2, \cdots, k+1$, and let us use them in a greedy way as follows: Start coloring the vertices of G one by one, and in each step use the smallest color that can be used without violating existing constraints, that is, the smallest color that does not occur among the colors of the neighbors of our current

vertex. As the maximum degree of G is D, at no step will there be more than D colors forbidden, so there will always be at least one color available for each vertex.

(b) The complete graphs K_n and the odd cycles C_{2m+1} have this property.

10. (a) We claim that, if G is not complete and is not an odd cycle, then $\chi(G) \leq k$.

In order to prove this, note that, if G has a vertex v of degree less than $k = D(G)$, then we are done by induction on the number of vertices. Indeed, remove v to get $G - v$. Then, by the induction hypothesis, $G - v$ has a proper coloring with at most k colors. As v has fewer than k neighbors in G, there is a color that can be used for v. This will provide a proper coloring of G with at most k colors.

It remains to prove the statement for the case when all vertices of G have degree k, in other words, when G is *k-regular*. We can assume that $k \geq 3$, otherwise either the claim is trivially true, or $G = K_2$ or $G = C_{2m+1}$.

In this case, take two vertices b and c in G whose distance is two. Two such vertices must exist, otherwise G would be complete, or not k-regular. Let a be a common neighbor of b and c. Since G is 3-connected, $G - b - c$ is connected. Since $G - b - c$ is connected, it has a subgraph T that is a tree (a *spanning tree* of $G - b - c$, in the language described in the Notes section of Chapter 5). Direct the edges of this tree toward a. Now color the vertices of G by colors from $[k]$ as follows: First, color b, then c, and then color the remaining vertices in any order so that, if the distance of x to a is more than the distance from y to a (along the directed edges of T), then x gets colored before y. So a gets colored last. See Figure 6.16 for an example.

In each step, use the smallest possible color that keeps the coloring proper. Then b and c both get colored 1. The other vertices different from a have at most $k-1$ colored neighbors when they themselves get colored. Finally, a has k colored neighbors when it gets colored, but two of them, b and c, have the same color. Therefore, all vertices can be colored using only colors from $[k]$, as claimed.

(b) Now assume that G is not 3-connected. If G is not 2-connected, then it has a vertex w so that $G - w$ is not connected. Let A be a connected component of $G - w$, and let $B = G - w - A$. Then by induction on the number of vertices, we can assume that both $A \cup w$ and $B \cup w$ have proper k-colorings. Permuting the colors so that w gets the same colors in both graphs yields a proper coloring of G.

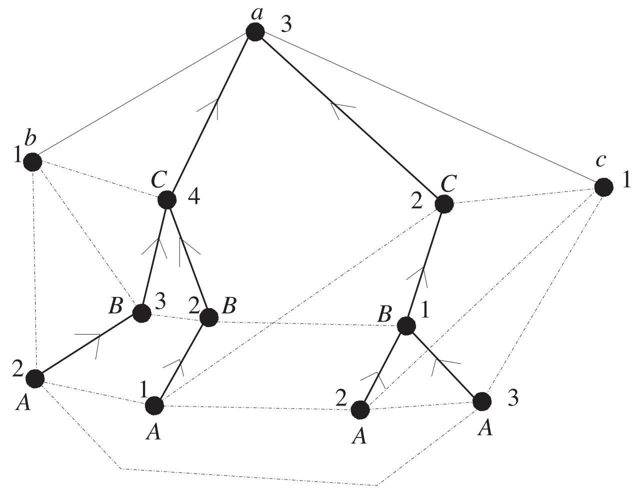

Figure 6.16
In this 4-regular graph G, the edges of T are the thick lines. After coloring b and c, we color the vertices in class A, then class B, then class C, then vertex a. Within each class, any order is allowed, but in this example, we proceeded from left to right.

What is left to consider is the case when G is 2-connected but not 3-connected. That means that G has two vertices, x and y, so that $G - x - y$ is not connected. In other words, there exist two induced subgraphs G_1 and G_2 of G so that $G_1 \cup G_2 = G$ and $G_1 \cap G_2 = \{x, y\}$. Note that both x and y must have at least one neighbor in each G_i, otherwise G would not be 2-connected.

Consider the graphs $G_1 \cup (x, y)$ and $G_2 \cup (x, y)$. (We add the edge (x, y) to G_i if it is not there yet.) It follows from the last sentence of the previous paragraph that the degree of each vertex in these graphs is at most k; therefore, these graphs both have proper k-colorings by induction. In each of these two colorings, x and y have different colors since they are adjacent. Therefore, permuting the colors so that the colors of x agree in both colorings and the colors of y agree in both colorings, we get a proper 2-coloring of G.

To summarize, we proved that, if G is not complete and not an odd cycle, then $\chi(G) = k$, where k is the maximum degree in G. This is a famous theorem of Brooks [21].

11. Let A_i be the set of schedules in which skill i is only available in one shift. This yields

$$|A_i| = 2 \cdot 2^{n-10} = 2^{n-9},$$

where n is the number of workers at the company. Indeed, there are two ways to schedule the 10 workers having skill i in a *bad way* (schedule them all in the first shift or all in the second shift), and then there are 2^{n-10} ways to schedule the remaining $n-10$ workers. Therefore,

$$| \cup_i A_i| \leq \sum_i |A_i| = 500|A_i| = 500 \cdot 2^{n-9} < 2^n,$$

since $2^9 = 512$. So the number of schedules in which at least one skill is not represented in both shifts is less than the number 2^n of all schedules. Therefore, there has to be a good schedule.

12. Taking the hypergraph $\mathcal{F}^c$ consisting of the complements of the edges in $\mathcal{F}$, we see that $\mathcal{F}^c$ has to consist of edges no two of which are disjoint. Therefore, Theorems 6.29 and 6.30 imply $|\mathcal{F}^c| = |\mathcal{F}| \leq 2^{n-1}$.

13. If $n = 2m + 1$, let $\mathcal{F}$ consist of all edges of $[n]$ having at least $m+1$ vertices. Any two of these edges are too large to be disjoint. If $n = 2m$, let $\mathcal{F}$ consist of all edges of $[n]$ having at least $m+1$ vertices and all subsets of $[n]$ having m vertices and containing the vertex n.

    Note that, in both cases, we pick exactly one vertex from each pair created in the proof of Theorem 6.29.

14. No, it cannot. Let us assume that it can, and let X be a minimum-sized edge in $\mathcal{F}$. Let us assume without loss of generality that $X = [k]$. Then, for each $i \in [k]$, there has to be an edge Y_i in $\mathcal{F}$ so that Y_i does not contain i (otherwise the intersection of all edges in $\mathcal{F}$ contains i). However, this implies that $X \cap Y_1 \cap Y_2 \cdots Y_k = \emptyset$, contradicting our hypothesis.

    Note that this result has a geometric counterpart, called *Helly's theorem.* That theorem says that, if there are $n \geq k+1$ convex sets given in $k+1$-dimensional Euclidean space, and they have the property that every $k+1$ of them intersect, then all of them intersect.

15. Color the integer n red if it is divisible by exactly one of 7 and 17, and color it blue otherwise. We show that this coloring satisfies the requirements.

    Let us assume that our coloring contains a monochromatic arithmetic progression A of length 18. Assume first that our arithmetic progression A has a difference d that is divisible by 17. Then, for A to be monochromatic, either all or no terms of A would have to be divisible by 7. The first is impossible, as that would require d to be at least $7 \cdot 17 = 119$, forcing $a_{18} = 2142$. For the same reason, d

cannot be divisible by 7, which means that exactly one of the first seven terms of A will be divisible by 7. This rules out the possibility that A has no terms divisible by 7, so d must not be divisible by 17.

Therefore, exactly one of the first 17 terms of A will be divisible by 17. If d is divisible by 7, then we are done, since either all or no terms of A are divisible by 7, so there will always be two terms that differ modulo 17, but not modulo 7. If d is not divisible by 7, then there are terms in our sequence not divisible by either 7 or 17, and there are some divisible by only one of them, contradicting our assumption that A is monochromatic.

16. Assume without loss of generality that vertex A has at least three red edges adjacent to it. Let the other endpoints of these edges be B, C, and D. If there is a red edge between any two of these three vertices, then the endpoints of that edge and A form a red triangle. Otherwise BCD is a blue triangle.

17. We prove this by induction on $k + l$. If $k + l = 2$, that is, if $k = l = 1$, then $R(k, l) = 2$, and the statement is true. Assume that the statement is true for $k+l-1$, that is, that $R(k, l-1)$ and $R(k-1, l)$ exist, and prove the statement for $k + l$.

    It suffices to prove that there is *one* integer with the required property. Indeed, in a nonempty set of positive integers, there is always a smallest one.

    We claim that $n = R(k, l-1) + R(k-1, l) - 1$ is such an integer. Indeed, let v be any vertex of K_n, and color each edge of K_n red or blue. Then v has either at least $R(k, l-1)$ blue edges or at least $R(k-1, l)$ red edges adjacent to it. In the first case, let the endpoints of the blue edges adjacent to v form the complete graph X. Then X has $R(k, l-1)$ vertices, so it either contains a red K_k, and we are done, or it contains a blue K_{k-1}, and then we are done again, adding v to this copy of K_{k-1}. The second case can be handled in an analogous way.

18. Let us assume without loss of generality that vertex A has at least six red edges adjacent to it. Look at the other endpoints of these vertices. If there is any red edge between them, we are done. If not, they form a copy of K_6 in which each edge is blue or green. Now apply the result of the previous exercise.

19. If N is prefix-free, then there can be no moment of time during the decoding of a message when we may stop and get a codeword A, or may continue and get a codeword B, since that would mean that A is a prefix of B. Similarly, if N is not prefix-free, and F is a prefix of G, and we try to decode a message starting with the digits of F,

then we do not immediately know whether we should stop and get F, or continue and get G. Therefore, N is not instantaneous.

20. The idea of pattern avoidance for matrices, and the earliest results, come from [35].

    (a) By the pigeonhole principle, if there are $n + 1$ entries equal to 1, then two of them must be in the same line. Therefore, $f(n, B) = n$.

    (b) We prove a more general statement. Define $f(n, m, B)$ to be the maximum number of 1s that an $n \times m$ matrix with 0 and 1 entries can contain if it avoids B. We claim that, if B is the 2×2 identity matrix, then $f(n, m, B) = n + m - 1$. This many entries equal to 1 are possible, as can be seen, for instance, by filling up the first row and column with 1s. To see that there cannot be more 1s, use induction on $n + m$. For $n + m = 2$, the statement is true. Let us assume that the statement is true for $n + m - 1$, and let A be an $n \times m$ matrix avoiding B. Assume A has $n + m$ entries equal to 1. If A has a row or column that contains at most one 1, then we can omit that line, and be done by the induction hypothesis. Otherwise, each row and each column contains at least two 1s. Since the total number of 1s is $n+m$, this means that each row and each column contains *exactly* two 1s. That is, the number of all 1s is $2n = n+m = 2m$, forcing $n = m$.

    Then A is the sum of two distinct permutation matrices. (The reader is invited to prove this directly, possibly using bipartite graphs. A more general statement is proved in Lemma 10.6.) These matrices cannot both correspond to the decreasing permutation, so one must correspond to a permutation that contains a non-inversion. That matrix will then contain B.

21. (a) Proof by induction on n. As A_n avoids B, the only way A_{n+1} could contain B would be if there was a copy of B in A_{n+1} that starts in the upper left quadrant and ends in another quadrant. It is easy to see that this is impossible.

    (b) Let $g(n)$ be the number of 1s in A_n. Then $g(0) = 1$, and $g(n + 1) = 2g(n)+n$ for $n \geq 0$. Solving this recurrence, we get $2g(n) = 2^n + n2^{n-2}$. Since A_n is of size $2^n \times 2^n$, setting $m = 2^n$, we see that $f(m, B) = \Omega(m \log m)$.

22. This result, and the results of the next two exercises, were proved in a paper by Adam Marcus and Gábor Tardos [56], and were used to prove a 25-year-old conjecture.

    We claim that the answer is $f(\frac{n}{k^2}, B)$. Call a block of A a *zero* block if all its entries are zeros. Form the matrix A' by replacing all zero blocks of A with a 0 and all nonzero blocks of A with a 1, to get the matrix A'. If A' contains B, then so does A, and the result follows. Note that we do need the fact that B is a permutation matrix, and as such, has only one 1 in each row and column.

23. Let us assume the contrary. There are $\binom{k^2}{k}$ possible k-tuples of rows in which a block can have nonzero elements, so, by the pigeonhole principle, this would mean that there is a k-tuple of rows that contains nonzero elements in k different blocks. However, that implies that A has a $k \times k$ submatrix M *of nonzero blocks*. Therefore, A contains *all* $k \times k$ matrices, including B. Indeed, M essentially simulates a $k \times k$ matrix in which each entry is 1, and the latter certainly contains all $k \times k$ matrices.

    Similarly, in any column of blocks, there can be at most $(k-1)k^2$ wide blocks.

24. First, if n is not divisible by k^2, then we just fill the last rows and columns of A with 1s so that the submatrix A_s of A that has not been filled by 1s has size $k^2 \cdot \lfloor n/k^2 \rfloor \times k^2 \cdot \lfloor n/k^2 \rfloor$. This increases the number of 1s by less than $2k^2 n = O(n)$. We can therefore assume from now on that n is divisible by k^2, that is, $A_s = A$. Indeed, the above argument shows that, if the statement were false, the excess of 1s would have to come from A_s.

    The previous two exercises show that, if A avoids B, and B is a permutation matrix of size $k \times k$, then there can be at most $f(n/k^2)$ nonzero blocks, and among those nonzero blocks, at most $(k-1)\binom{k^2}{k}$ in each row can be tall. Similarly, in each column of blocks, at most $(k-1)\binom{k^2}{k}$ blocks can be wide. If a block is not tall or wide, then the number of 1s in that block is at most $(k-1)^2$, since the block can contain 1s in at most $k-1$ rows and at most $k-1$ columns.

    We can now summarize the number of 1s in our matrix A.

    - There are at most $2\frac{n}{k^2}(k-1)\binom{k^2}{k}$ blocks that are tall or wide, and each of them contains at most k^4 entries equal to 1, for a total of fewer than $2nk^3\binom{k^2}{k}$ entries equal to 1, and
    - there are at most $f(n/k^2, B)$ nonzero blocks that are neither tall nor wide, and each of them contains at most $(k-1)^2$ entries 1, for a total of at most $(k-1)^2 f(n/k^2, B)$ entries 1, and
    - the rest must be zero blocks with no 1s.

    This leads to the recursive formula

    $$f(n, B) \leq 2nk^3\binom{k^2}{k} + (k-1)^2 f(n/k^2, B),$$

    from which it is routine to prove our statement by induction.

25. This result was first proved by Martin Klazar [49], and then his proof was somewhat simplified by Vincent Vatter and Doron Zeilberger (not formally published). This is the proof we present.

    Let q be a pattern of length k, and let p be an n-permutation that avoids q. Then the permutation matrix of p avoids the permutation matrix B_q of q. Now let $M_n(q)$ be the number of $n \times n$ matrices with 0 and 1 entries that avoid B_q. Then $S_n(q) \leq M_n(q)$. We will now find an upper bound for $M_n(q)$.

    By the result of the previous exercise, we know that there exists a constant b so that $f(n, B_q) \leq bn$. Let A be an $n \times n$ matrix that avoids B_q, and cut A up into 2×2 blocks. (There will be smaller blocks at the end if n is odd.) If a block contains no 1s, then replace it with a 0; if it contains at least one 1, replace it with a 1. This results in a matrix of size $\lceil n/2 \rceil \times \lceil n/2 \rceil$ that is also B_q-avoiding and therefore has at most $f(\lceil n/2 \rceil, B_q) \leq b\lceil n/2 \rceil$ entries equal to 1. Furthermore, each of these ones comes from a nonzero 2×2 block of A. There are $2^4 - 1 = 15$ possibilities for each such block. Therefore, the number $M_n(q)$ of all possible matrices A we could have started with satisfies

    $$M_n(q) \leq 15^{b\lceil n/2 \rceil} \cdot M_{b\lceil n/2 \rceil}(q).$$

    Repeat this argument until the right-hand side becomes $M_1(q) = 2$. Then we get $M_n(q) \leq 15^{2bn}$, and therefore $S_n(q) \leq 15^{2bn}$.

26. Count the pairs (p, q'), where p is an n-permutation containing q, and q' is a copy of q in p. There are $\binom{n}{k}$ choices for the position of q', then there are $\binom{n}{k}$ choices for the entries in q', and finally there are $(n-k)!$ choices for the rest of p. So the number of such pairs is $(n-k)! \cdot \binom{n}{k}^2$. The pigeonhole principle shows that at least $1/n!$ of them, that is, at least

    $$\frac{(n-k)! \cdot \left(\frac{n!}{k!(n-k)!}\right)^2}{n!} = \frac{\binom{n}{k}}{k!}$$

    pairs, must belong to one permutation. Furthermore, there exist q-avoiding permutations, and no (p, q) pair belongs to them. Therefore, there exists a p with strictly *more* than $\frac{\binom{n}{k}}{k!}$ occurrences of q.

27. (a) Represent the jobs and the candidates by vertices, and connect candidates to jobs for which they are qualified by an edge. The result is a bipartite graph. If aX is an edge of this graph G, then write the number t_{aX} on this edge, where t_{aX} is the lowest salary for which candidate a will accept job X.

    Recall that *perfect matchings* of a graph were defined in Supplementary Exercise 34 of Chapter 5. The task of the manager is then to find the perfect matching of G in which the sum of the numbers written on the edge (the *cost* of the matching) is minimal.

(b) No, this greedy strategy will not always produce the best results. The graph shown in Figure 6.17 is a counterexample. The greedy algorithm provides a matching costing 13, while a matching costing 12 exists.

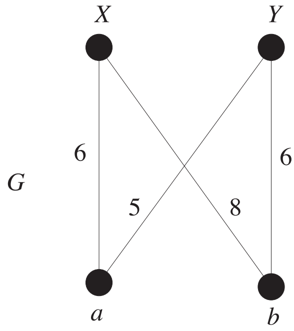

Figure 6.17
The greedy algorithm will not find the minimum-cost perfect matching of G.

28. Both A and B are forests. Because A has more edges, it has fewer components. So A must have an edge between vertices that are in different components of B, and that edge can be added to B without creating a cycle there.

29. Yes, in this case the greedy algorithm will work. Look at all pairs of towns in the county, take a complete graph whose vertices are the towns, and write the cost of construction of each road to the edge corresponding to that road. Using the terminology of the Notes section of Chapter 5, we have to find a *minimum-cost spanning tree* in this graph, that is, a tree on n vertices so that the sum of the numbers written on the edges of the tree is minimal.

    Let G_j be the graph the greedy algorithm builds in j steps. The definitions then imply that G_j is a forest and G_{n-1} is a tree. We claim that no other spanning tree of K_n will have a smaller cost than G. Indeed, assume that $c(T) < c(G)$, where $c(T)$ denotes the cost of T. Let $g_1, g_2, \cdots, g_{n-1}$ and $t_1, t_2, \cdots, t_{n-1}$ be the edges of G and T in nondecreasing order of their costs. As $c(T) < c(G)$, there has to be a smallest index i so that $c(g_i) > c(t_i)$.

    This is a contradiction, however. Indeed, this would imply that $c(g_i) > c(t_j)$ for all $j \leq i$. The previous exercise shows that at least one of the edges t_j, with $j \leq i$, can be added to G_{i-1} without creating a cycle. Therefore, g_i cannot be the edge that the greedy