Graphical representations of clutters[*]

M. H. Dinitz, J. M. Gold, T. C. Sharkey and L. Traldi
Department of Mathematics, Lafayette College
Easton, Pennsylvania 18042

Abstract

We discuss the use of K-terminal networks to represent arbitrary clutters. A given clutter has many different representations, and there does not seem to be any set of simple transformations that can be used to transform one representation of a clutter into any other. We observe that for $t \geq 2$ the class of clutters that can be represented using no more than t terminals is closed under minors, and has infinitely many forbidden minors.

1. Introduction

A *clutter* on a finite set S is a family of subsets of S, none of which contains any other. A graph naturally gives rise to many clutters, including the families of minimal edge- or vertex-cuts, edge- or vertex-sets of simple circuits, edge-sets of spanning trees, edge- or vertex-sets of simple paths between two given vertices, and so on.

A less familiar construction associates a clutter $C(G, K)$ to a K-*terminal network* (G, K) consisting of a graph G and a subset $K \subseteq V(G)$ of *terminals*: $C(G, K)$ is the clutter on $S = V(G) \setminus K$ which contains every minimal subset $M \subseteq S$ such that the full subgraph of G induced by $M \cup K$ is connected. The elements of $C(G, K)$ are *minpaths* of (G, K), and (G, K) is a *graphical representation* of $C(G, K)$. A K-terminal network may be thought of as a model of a real-world structure, perhaps a computer or telephone network; the terminals represent users of the network and the non-terminal vertices represent elements of the network which may or may not operate. Note that this interpretation does not explicitly allow for the failure of network elements represented by the edges of G; the possibility of such failures may be incorporated by inserting vertices of degree 2 in edges whose failure is possible. We refer the interested reader to [1] for a general discussion of K-terminal networks and network reliability.

Our interest in the construction of $C(G, K)$ does not lie in the modeling of real-world networks, however. Rather, we are interested in the following result of [12, 13].

Theorem 1. *Every clutter is $C(G, K)$ for some K-terminal network (G, K).*

Theorem 1 is not difficult to prove. The trivial clutters $\{\emptyset\}$ and $\emptyset$ are represented by edgeless 1-terminal and 2-terminal networks, respectively. A graphical representation of a nontrivial clutter C may be constructed using the *dual* or *blocker* C^*, defined by: $\emptyset^* = \{\emptyset\}$, $\{\emptyset\}^* = \emptyset$, and if $\emptyset \neq C \neq \{\emptyset\}$ then the elements of C^* are minimal among sets which intersect all the elements of C. A clutter C with $|C^*| = 1$ and $\emptyset \notin C^*$ is simply a collection of singletons; it is represented by a K-terminal network with $|K| = 2$ which has a non-terminal vertex for each element of an element of C, in which every non-terminal vertex is adjacent to both terminals and the two terminal vertices are not adjacent to each other. If C is a clutter on S and $|C^*| \geq 2$ then there is a K-terminal network with $C = C(G, K)$ which has $K = C^*$ and $V(G) \setminus K = S$; all of the non-terminal vertices of G are adjacent to each other, none of the terminal vertices are adjacent to each other, and each terminal vertex $B \in C^*$ is adjacent to each non-terminal vertex $v \in B$. That $C = C(G, K)$ follows from the fact that $C^{**} = C$ [3].

We call the graphical representations mentioned in the preceding paragraph *standard*. There is generally a great variety of other graphical representations of a given clutter. In Section 2 of the paper we discuss several simple ways to transform a graphical representation of a clutter into another representation of the same clutter. These transformations may be applied to an arbitrary graphical representation (G, K) of a clutter C to obtain a representation (G', K') with $|K'| \leq |K|$ which is *reduced* in the sense that there are no loops, every non-isolated non-terminal vertex appears in some minpath, no two terminals are adjacent or have the same neighbor-set, and every terminal's neighbor-set is an element of C^*.

In some contexts it happens that related combinatorial structures can be changed into each other incrementally, using simple transformations; consider the basis exchange property of matroids [5], or the Reidemeister moves in knot theory [6]. Examples indicate that there is no set of "simple" transformations which can be used to obtain all reduced graphical representations of a clutter from a given reduced representation, at least if "simple" is interpreted in a reasonable way. For instance in Theorem 3.9 we show that for $n \geq 4$ the clutter $U_{n-1,n} = \{(n-1)$-element subsets of an n-element set$\}$ has reduced representations (G, K) with $|K| = n$ and $|K| = \binom{n}{2}$, but

no reduced representation with $n < |K| < \binom{n}{2}$; consequently any transformation used to obtain one of these representations from another must involve the simultaneous introduction or removal of $\binom{n}{2} - n$ terminals.

The *terminal number* of a K-terminal network is $|K|$ and the *minimum terminal number* of a clutter, $\mathrm{term}(C)$, is the minimum of the terminal numbers of graphical representations of C; the standard representations show that in general $\mathrm{term}(C) \leq \max\{2, |C^*|\}$. We think of $\mathrm{term}(C)$ as a measure of the complexity of C.

If C is a clutter on S and S_1, S_2 are disjoint subsets of S then $(C/S_1) \setminus S_2 = (C \setminus S_2)/S_1$ is the clutter on $S \setminus (S_1 \cup S_2)$ consisting of the minimal subsets $N \subseteq S \setminus (S_1 \cup S_2)$ with the property that $N \cup S_1$ contains an element of C. This clutter is the *minor* of C obtained by *contracting* S_1 and *deleting* S_2. It is common to simplify notation when contracting or deleting single elements: C/x for $C/\{x\}$ and $C \setminus x$ for $C \setminus \{x\}$. Two simple properties of the minor operations are order-independence (i.e., $(((C/S_1) \setminus S_2)/S_3) \setminus S_4 = (C/(S_1 \cup S_3)) \setminus (S_2 \cup S_4)$) and duality (i.e., $((C/S_1) \setminus S_2)^* = (C^*/S_2) \setminus S_1$).

As noted in [13], the minor operations are compatible with graphical representations. If $C = C(G, K)$ then a (G', K) with $(C/S_1) \setminus S_2 = C(G', K)$ may be obtained by removing each non-terminal vertex $v \in S_2$ and all edges incident on v, and replacing each non-terminal vertex $v \in S_1$ with edges connecting all the pairs of neighbors of v. To motivate these representations of deletion and contraction, recall that we may think of (G, K) as a network whose function is to provide communication among the elements of K, and whose non-terminal vertices are vulnerable to failure. In $(C/S_1) \setminus S_2$ each non-terminal vertex $v \in S_2$ has failed, and each non-terminal vertex $v \in S_1$ has become invulnerable to failure and hence is logically equivalent to a clique of its neighbors. (G, K) and (G', K) have the same terminals, so we conclude that $\mathrm{term}(C) \geq \mathrm{term}(C')$ for any minor C' of C. It follows that for each fixed $t \geq 2$ the class of clutters satisfying $\mathrm{term}(C) \leq t$ is closed under minors, and hence is determined by a family of forbidden minors.

Theorem 2. *For every $t \geq 2$ the family of clutters satisfying* $\mathrm{term}(C) \leq t$ *has infinitely many forbidden minors.*

We have not completely determined any of these families of forbidden minors, but in Section 3 we present forbidden minors for various minimum terminal numbers; in particular we show that if $t \geq 2$ then the forbidden minors for $\mathrm{term}(C) \leq t$ include all the degenerate projective planes J_s with $s \geq t$.

In Section 5 we briefly discuss the properties of 2-terminal clutters.

2. Clutter-preserving transformations

It is not unusual for nonisomorphic K-terminal networks to represent the same clutter. For instance, each of the following transformations of a K-terminal network (G, K) does not affect $C(G, K)$. We denote by $N(v)$ the set of vertices adjacent to v, excluding v itself.

1. A loop at any vertex may be adjoined or deleted.

2. If v and w are non-terminal vertices with a common terminal neighbor, an edge between v and w may be adjoined or deleted.

3. If v is a non-terminal vertex which does not appear in any minpath of (G, K) then any edge incident on v may be removed; conversely an edge incident on such a v may be adjoined, so long as the edge does not create a minpath involving v.

4. If (G, K) has two terminal vertices with precisely the same neighbors then one of these terminals may be removed; conversely a new terminal may be introduced with the same neighbors as an existing terminal.

5. If two terminal vertices are adjacent, the edge connecting them may be contracted; that is, the two terminals may be combined into a single terminal adjacent to all the neighbors of the original terminals. Conversely, a terminal vertex τ may be replaced by two adjacent terminals τ_1 and τ_2 such that $((N(\tau_1) \setminus \{\tau_2\})) \cup (N(\tau_2) \setminus \{\tau_1\})) = N(\tau)$.

6. If τ_1 and τ_2 are terminal vertices such that neither is adjacent to any terminal and $N(\tau_1) \subseteq N(\tau_2)$, edges connecting all pairs of neighbors of τ_2 may be adjoined and then τ_2 may be removed.

7. If τ is a terminal cutpoint and $G \setminus \{\tau\}$ has terminals in separate components, edges connecting all pairs of neighbors of τ may be adjoined and τ may then be removed.

8. Given a terminal τ and a $B \in C(G, K)^*$ such that $B \subseteq N(\tau)$, edges connecting all pairs of neighbors of τ may be adjoined and then τ may be replaced with a terminal τ' such that $B = N(\tau')$.

We prove that transformations 6 and 8 do not affect $C(G, K)$, and leave the other proofs to the reader. Observe first that adjoining edges connecting all pairs of neighbors of τ_2 does not affect $C(G, K)$ because it is a combination of instances of transformation 2; we presume that all these edges are present in (G, K). Let (G', K') be obtained from (G, K)

4

by removing τ_2. Suppose $M \in C(G, K)$; then the full subgraph H of G induced by $M \cup K$ is connected. The full subgraph H' of G' induced by $M \cup K'$ is also connected, because a path in H which does not end at τ_2 has a corresponding path in H', in which any appearance of τ_2 has been replaced either by an appearance of τ_1 or by an edge connecting two neighbors of τ_2. If $M' \in C(G', K')$ then the full subgraph H' of G' induced by $M' \cup K'$ is connected; the full subgraph H of G induced by $M' \cup K$ is also connected because every neighbor of τ_1 is adjacent to τ_2. This verifies that transformation 6 does not affect $C(G, K)$. In the situation of transformation 8, adjoining τ' to (G, K) does not affect $C(G, K)$, because every $M \in C(G, K)$ intersects B and hence contains a non-terminal vertex adjacent to τ'. Transformation 6 may then be applied to adjoin edges connecting all pairs of neighbors of τ and remove τ.

Transformations 1–8 may be applied to a graphical representation (G, K) of a clutter C to obtain a representation (G', K') with $|K'| \le |K|$ which is *reduced* in the sense mentioned in the introduction: there are no loops, every non-isolated non-terminal vertex appears in some minpath, no two terminals are adjacent or have the same neighbor-set, and every terminal's neighbor-set is an element of C^*. We call two reduced representations *equivalent* if the same elements of C^* appear as neighbor-sets of terminals in both.

Transformations 1–8 cannot generally be used to obtain an inequivalent reduced representation from the standard one. A ninth transformation may sometimes be used to obtain inequivalent reduced representations.

9. Suppose $B \in C(G, K)^*$ and every two elements of B are either adjacent or connected by a path whose internal vertices are all terminals. Then a terminal τ with $N(\tau) = B$ may be adjoined.

Transformations 1–9 are certainly not "complete" in the sense of being adequate to generate all the graphical representations of a given clutter from any one. For example, we prove in Theorem 3.9 below that the uniform clutters $U_{n-1,n}$ have many different reduced representations; transformations 1–9 cannot be used to change any one of these reduced representations into an inequivalent one.

3. Some forbidden minors

Any clutter C can be represented by a K-terminal network (G, K) with $\max\{2, |C^*|\}$ terminals, and many clutters can be represented by K-terminal networks with fewer than $\max\{2, |C^*|\}$ terminals. A clutter that actually

requires $\max\{2, |C^*|\}$ terminals seems especially interesting because it has no graphical representation which is essentially simpler than the standard one.

Some preliminary definitions and results will be convenient in our discussion of examples.

Definition 3.1. *Two vertices of a K-terminal network have* perfect communication *if they are equal, adjacent, or connected by a path whose internal vertices are all terminals.*

The terminology reflects the idea that the only elements of a K-terminal network which are vulnerable to failure are the non-terminal vertices; vertices which communicate perfectly may be connected with a path which is invulnerable or "perfect."

Lemma 3.2. *If two vertices have perfect communication then they appear in the same connected component of $G \setminus B$ for every $B \in C(G, K)^*$ which does not contain either of them.*

Proof: The assertion follows immediately from the fact that $B \in C(G, K)^*$ is a vertex cut consisting of non-terminal vertices. ∎

Corollary 3.3. *If all elements of $V(G) \setminus K$ have perfect communication, then $|K| \geq |C(G, K)^*|$.*

Proof: Let $B \in C^*$. All of the non-terminals in $G \setminus B$ are in one component of $G \setminus B$, so there is a component Γ_B of $G \setminus B$ with only terminal vertices. If $B' \neq B \in C(G, K)^*$ then Γ_B cannot intersect $\Gamma_{B'}$, for if $\tau \in \Gamma_B \cap \Gamma_{B'}$ then $\Gamma_B = \{\text{terminals } \sigma \text{ for which there is a } \sigma\tau \text{ path which contains only terminals}\} = \Gamma_{B'}$ and hence $B = \{\text{non-terminal neighbors of elements of } \Gamma_B\} = B'$. ∎

For $s > 1$, the degenerate projective plane J_s is the clutter $\{\{x_1, y\}, \{x_2, y\}, ..., \{x_s, y\}, \{x_1, x_2, ..., x_s\}\}$; note that $J_s^* = J_s$. These clutters are well-known as forbidden minors for binary clutters [7], for matroid ports [8], and for the width-length property [4]. They play a similar role here: J_s is a forbidden minor for $\mathrm{term}(C) \leq t$ whenever $2 \leq t \leq s$.

Theorem 3.4. *For every $s \geq 2$, $\mathrm{term}(J_s) = |J_s^*| = s + 1$. Every proper minor of J_s has a 2-terminal representation.*

Proof: Let (G, K) represent J_s. Consider $B = \{x_1, ..., x_s\} \in J_s^*$. Only one non-terminal vertex, y, remains in $G \setminus B$. But $B \in J_s^*$, so $G \setminus B$ has at least two components; at least one component must consist entirely of terminals. This component provides perfect communication among the elements of B. Now consider $\{x_i, y\}$, where $1 \leq i \leq s$. The non-terminals not in $\{x_i, y\}$ have perfect communication, so Lemma 3.2 implies that they all appear in one component of $G \setminus \{x_i, y\}$. There must be another component whose vertices are all terminals, and this component provides perfect communication between x_i and y. Corollary 3.3 implies that $|K| \geq |J_s^*|$.

$J_s/y = \{\{x_1\}, ..., \{x_s\}\}$ may be graphically represented with two non-adjacent terminals, each adjacent to every x_i. $J_s/x_j = \{\{y\}, \{x_i : 1 \leq i \leq s, i \neq j\}\}$ may be represented by a 2-terminal network, with y adjacent to both terminals and all the x_i other than x_j appearing on a path connecting the two terminals. $J_s \setminus y = \{\{x_1, ..., x_s\}\}$ may be represented by a 2-terminal network with all the x_i as vertices of a path connecting the two terminals. $J_s \setminus x_j = \{\{y, x_i\} : 1 \leq i \leq s, i \neq j\}$ may be represented by a 2-terminal network in which y is the only neighbor of one terminal, and all the x_i other than x_j are adjacent to y and the second terminal. ■

We denote by $U_{a,n}$ the *uniform* clutter consisting of the a-element subsets of an n-element set; note that $U_{a,n}^* = U_{n-a+1,n}$. $U_{0,n} = \{\emptyset\}$ may be represented with just one terminal. $U_{1,n}$ is represented by a K-terminal network (G, K) with two nonadjacent terminals, both of which are adjacent to all n non-terminal vertices. $U_{n,n}$ may be represented with two terminals, connected by a path of length n.

A lemma will be useful in analyzing the $U_{a,n}$ with $1 < a < n - 1$.

Lemma 3.5. *Suppose* $1 < a < n - 1$, *and let* (G, K) *be a K-terminal network representing* $U_{a,n}$. *If there is an* $X \subset S = V(G) \setminus K$ *such that* $|X| \geq a - 1$ *and every element of* X *has perfect communication with every other, then* $|K| \geq |U_{a,n}^*|$.

Proof: Suppose $|X| > a - 1$ and let $Y = S \setminus X$; then $|Y| < n - a + 1$. Consider any $D \subset S$ with $|D| = n - a + 1$ and $Y \subset D$; note that $D \in U_{a,n}^*$. Every non-terminal of $G \setminus D$ is an element of X and so has perfect communication with all the others; hence every non-terminal is in the same component of $G \setminus D$. Another component of $G \setminus D$ must consist entirely of terminals; it provides perfect communication among the elements of D. Iterating over all possible choices of D provides perfect communication within Y and between any element of Y and any element of X. Thus every non-terminal has perfect communication with every other, so Corollary 3.3 applies.

Suppose $|X| = a - 1$, and again let $Y = S \setminus X$; then $|Y| = n - a + 1$ and hence $Y \in U^*_{a,n}$. All of the non-terminals in $G \setminus Y$ are elements of X and therefore in the same component of $G \setminus Y$; any other component of $G \setminus Y$ must consist entirely of terminals, and such a component provides perfect communication among the elements of Y. If $|Y| > |X|$ the conclusion of the lemma is obtained by applying the argument of the preceding paragraph to Y in place of X.

Suppose $|Y| \leq |X|$, and choose a fixed subset $A \subset X$ with $|A| = n - a$; then $A \cup \{y\} \in U^*_{a,n}$ for every $y \in Y$. We claim that there is at least one $y_A \in Y$ such that the elements of $A \cup \{y_A\}$ have perfect communication. The elements of X communicate perfectly, so Lemma 3.2 tells us that for any $y \in Y$ there is one component of $G \setminus (A \cup \{y\})$ which contains all the elements of $X \setminus A$; choose $y_A \in Y$ so that $G \setminus (A \cup \{y_A\})$ has a component H which contains no vertices from X, and contains as few vertices from Y as possible given that it contains none from X. Suppose $y \in V(H) \cap Y$. Every component of $G \setminus (A \cup \{y_A\})$ contains a vertex adjacent to y_A, because $G \setminus A$ is connected; consequently all the components of $G \setminus (A \cup \{y_A\})$ other than H will be contained in the one component of $G \setminus (A \cup \{y\})$ which contains y_A. The non-terminal vertices which appear in the remaining component(s) of $G \setminus (A \cup \{y\})$ include only those from $V(H) \setminus \{y\}$, contradicting the minimality of H. It follows that $V(H) \cap Y = \emptyset$, and hence that H contains no non-terminal vertices, so H provides perfect communication among the elements of $A \cup \{y_A\}$.

Suppose $x \in X$ and consider $B = (Y \cup \{x\}) \setminus \{y_A\} \in U^*_{a,n}$. The component of $G \setminus B$ which contains y_A also contains $A \setminus \{x\}$, because the elements of A communicate perfectly with y_A; observe that $a < n - 1$ implies that $|A| = n - a > 1$, guaranteeing that $A \setminus \{x\} \neq \emptyset$. This component of $G \setminus B$ also contains the other elements of X, because they communicate perfectly with the elements of $A \setminus \{x\}$. Any other component of $G \setminus B$ contains only terminal vertices, and hence provides perfect communication among the elements of B. This shows that every $x \in X$ communicates perfectly with every element of Y other than y_A. The elements of Y communicate perfectly with each other, so there is perfect communication between any two elements of $X \cup (Y \setminus \{y_A\})$. This set is strictly larger than X, so the argument of the first paragraph may be applied to it. ∎

Theorem 3.6. *If* $1 < a < n - 1$ *then* $\text{term}(U_{a,n}) = |U^*_{a,n}| = \binom{n}{n-a+1}$.

Proof: Consider any terminal vertex τ in a graphical representation (G, K) of $U_{a,n}$. Let T be the set of terminals which are connected to τ by paths whose internal vertices are all terminals. Then $X = N(T)$ consists entirely

of non-terminals, and T provides perfect communication among the elements of X. Observe that T is separated from the rest of K in $G \setminus X$, so X contains an element of $U_{a,n}^*$; hence $|X| \geq n - a + 1$.

If $|V(G) \setminus (K \cup X)| = n - |X| \leq n - a + 1$ then $|X| \geq a - 1$ and Lemma 3.5 tells us that $|K| \geq |U_{a,n}^*|$.

Otherwise $|V(G) \setminus (K \cup X)| = n - |X| > n - a + 1$. Suppose $Y \subset V(G) \setminus (K \cup X)$ and $|Y| = n - a + 1$. The elements of X communicate perfectly through T, so they appear in one component of $G \setminus Y$; any other component of $G \setminus Y$ contains no elements of X. Choose such a Y so that a component H of $G \setminus Y$ which does not contain any member of X contains the smallest possible number of non-terminal vertices. Suppose $v \in V(H) \setminus K$ and $y \in Y$, and let $V = \{v\} \cup (Y \setminus \{y\})$. $G \setminus (Y \setminus \{y\})$ is connected, so every component of $G \setminus Y$ contains a vertex adjacent to y; hence the component of $G \setminus V$ which contains y contains every component of $G \setminus Y$ other than H. The vertex-set of any other component of $G \setminus V$ is contained in $V(H) \setminus \{v\}$, contradicting the minimality of H; hence there is no $v \in V(H) \setminus K$. Thus $V(H) \subseteq K$, and H provides perfect communication among the elements of Y.

Recall that $X \cup Y$ is a proper subset of S, because $n - |X| > n - a + 1$. Let $Z = S \setminus (X \cup Y)$, and suppose $\emptyset \neq A \subset X$ and $\emptyset \neq B \subset Y$ have $|A| + |B| = n - a$. If $z \in Z$ then $A \cup B \cup \{z\} \in U_{a,n}^*$; Lemma 3.2 tells us that all the elements of Y appear in a single component of $G \setminus (A \cup B \cup \{z\})$. Choose $z \in Z$ so that a component H of $G \setminus (A \cup B \cup \{z\})$ which doesn't contain any element of Y has the smallest possible number of non-terminals. Suppose $v \in V(H) \cap Z$. $G \setminus (A \cup B)$ is connected, so every component of $G \setminus (A \cup B \cup \{z\})$ contains a vertex adjacent to z; consequently the component of $G \setminus (A \cup B \cup \{v\})$ containing z contains all the components of $G \setminus (A \cup B \cup \{z\})$ other than H. The vertex-set of any other component of $G \setminus (A \cup B \cup \{v\})$ is contained in $V(H) \setminus \{v\}$, violating the minimality of H. We conclude that $V(H)$ contains no element of Z. $V(H)$ also contains no element of Y, so all the non-terminal vertices of H are elements of X.

If $V(H) \cap X = \emptyset$ then H contains no non-terminals, so it provides perfect communication among the elements of $A \cup B \cup \{z\}$.

On the other hand, suppose $V(H) \cap X \neq \emptyset$; Lemma 3.2 tells us that $X \setminus A \subseteq V(H)$. $G \setminus (A \cup B)$ is connected and H is a component of $G \setminus (A \cup B \cup \{z\})$, so H must have a vertex u adjacent to z. If u may be chosen in X then we may also choose $x \neq u \in X \setminus A$, because $|X - A| \geq n - a + 1 - (n - a - 1) = 2$. If u cannot be chosen in X then u is a terminal. H is connected, so there is a path in H from u to an element

of X; by shortening the path if necessary we may assume that its internal vertices are all terminals. $|X - A| \geq 2$, so we may choose an $x \in X \setminus A$ which is not the element of X which appears on this path. Either way, u is a vertex of H which is adjacent to z, $x \in V(H) \cap X$, and u has perfect communication with an element of $V(H) \cap X$ in $H \setminus x$. $G \setminus (A \cup B)$ is connected, so every component of $G \setminus (A \cup B \cup \{z\})$ contains a vertex adjacent to z; consequently all the components of $G \setminus (A \cup B \cup \{z\})$ other than H are contained in the component of $G \setminus (A \cup B \cup \{x\})$ which contains z. $V(H) \setminus \{x\}$ contains u which is adjacent to z, and also contains an element of X which communicates perfectly with u in $H \setminus x$; because there is perfect communication among the elements of X, Lemma 3.2 tells us that every element of $(V(H) \cap X) \setminus \{x\}$ is contained in the component of $G \setminus (A \cup B \cup \{x\})$ which contains z. Any other component of $G \setminus (A \cup B \cup \{x\})$ contains only vertices from $V(H) \setminus (V(H) \cap X)$. The non-terminals appearing in H are all elements of X, so any other component of $G \setminus (A \cup B \cup \{x\})$ contains only terminal vertices and provides perfect communication among the elements of $A \cup B \cup \{x\}$.

We conclude that whether $V(H) \cap X = \emptyset$ or $V(H) \cap X \neq \emptyset$ there is perfect communication between any element of A and any element of B. We can repeat this for all nonempty $A \subset X$ and $B \subset Y$ such that $|A| + |B| = n - a$, and conclude that there is perfect communication between any element of X and any element of Y. There is perfect communication among the elements of X and also among the elements of Y, so there is perfect communication among the elements of $X \cup Y$.

Denote $X \cup Y$ by X_1. We may apply the argument above, starting with the second paragraph of the proof, to X_1 in place of X. We conclude that either $|V(G) \setminus (K \cup X_1)| = n - |X_1| \leq n - a + 1$ (in which case $|X_1| \geq a - 1$ and Lemma 3.5 tells us that $|K| \geq |U_{a,n}^*|$) or $|V(G) \setminus (K \cup X_1)| = n - |X_1| > n - a + 1$ (in which case the argument produces a $Y_1 \in U_{a,n}^*$ which is contained in $V(G) \setminus (K \cup X_1)$ and has the property that there is perfect communication among the elements of $X_2 = X_1 \cup Y_1$). Repeating as many times as necessary, we conclude that $|K| \geq |U_{a,n}^*|$.

The standard representation shows that $\mathrm{term}(U_{a,n}) \leq |U_{a,n}^*|$. $\blacksquare$

Corollary 3.7. *If $n \geq 6$, $2 < a < n - 2$ and $\max\{\binom{n-1}{n-a+1}, \binom{n-1}{n-a}\} < t < \binom{n}{n-a+1}$ then $U_{a,n}$ is a forbidden minor for $\mathrm{term}(C) \leq t$.*

Proof: Theorem 3.6 states that $\mathrm{term}(U_{a,n}) = |U_{a,n}^*| = \binom{n}{n-a+1}$. A deletion $U_{a,n} \setminus x$ consists of the a-element subsets of S which do not contain x, and hence is isomorphic to $U_{a,n-1}$. It follows that $\mathrm{term}(U_{a,n} \setminus x) =$

10

$\left|U^*_{a,n-1}\right| = \binom{n-1}{n-a}$. On the other hand, a contraction $U_{a,n}/x$ consists of the $(a-1)$-element subsets of $S \setminus \{x\}$, and hence is isomorphic to $U_{a-1,n-1}$. It follows that $\mathrm{term}(U_{a,n}/x) = \left|U^*_{a-1,n-1}\right| = \binom{n-1}{n-a+1}$. $\blacksquare$

Corollary 3.8. *If $n \geq 3$ then $U_{2,n}$ is a forbidden minor for $\mathrm{term}(C) \leq n-1$.*

Proof: If $n \geq 4$ then Theorem 3.6 states that $\mathrm{term}(U_{2,n}) = \left|U^*_{2,n}\right| = \binom{n}{n-1} = n$. It is a simple matter to determine directly that $\mathrm{term}(U_{2,2}) = 2$ and $\mathrm{term}(U_{2,3}) = 3$.

Observe that if $n \geq 3$ then a deletion $U_{2,n} \setminus x$ consists of the 2-element subsets of S which do not contain x, and hence is isomorphic to $U_{2,n-1}$. It follows that $\mathrm{term}(U_{2,n} \setminus x) = n-1$. On the other hand, a contraction $U_{a,n}/x$ consists of the 1-element subsets of $S \setminus \{x\}$, and hence is 2-terminal. $\blacksquare$

It is not hard to verify that all reduced graphical representations of $U_{1,2}$ and $U_{2,3}$ are standard. This property is not shared by the larger clutters $U_{n-1,n}$.

Theorem 3.9. *If $n \geq 4$ then $U_{n-1,n}$ has precisely one type of nonstandard reduced graphical representation: a cycle in which n terminals and n non-terminals appear alternately, which may have an edge between the two neighbors of any terminal. Consequently, $\mathrm{term}(U_{n-1,n}) = n$.*

Proof: Suppose (G, K) is a reduced graphical representation of $U_{n-1,n}$. $U^*_{n-1,n} = U_{2,n}$, so every terminal vertex in (G, K) is of degree 2.

Suppose P is a simple path $v_1, \tau_1, v_2, \tau_2, ..., \tau_{c-1}, v_c$ in G which involves non-terminals v_i and terminals τ_i. Note that $c \leq n$, for there are only n non-terminal vertices in G. We claim that there is a terminal adjacent to one of the ends of the path.

Suppose $c < n$, and let $Z = S \setminus \{v_1, ..., v_c\}$ contain all the non-terminals not in P. Every 2-element subset of S is a vertex cut of G, because $U^*_{n-1,n} = U_{2,n}$. If $z \in Z$ then a single component of $G \setminus \{v_1, z\}$ contains the path $\tau_1, v_2, \tau_2, ..., \tau_{c-1}, v_c$. Choose $z \in Z$ so that a component H of $G \setminus \{v_1, z\}$ which does not intersect P contains the smallest possible number of non-terminals. Suppose $u \in V(H) \setminus K$; then $u \in Z$. Consider the mincut $\{v_1, u\}$. Every component of $G \setminus \{v_1, z\}$ contains a vertex adjacent to z, because $G \setminus \{v_1\}$ is connected; hence all of the components of $G \setminus \{v_1, z\}$ other than H are contained in the component of $G \setminus \{v_1, u\}$ which contains z. Any other component of $G \setminus \{v_1, u\}$ contains only vertices from $V(H) \setminus \{u\}$,

11

contradicting the minimality of H. Therefore H contains no non-terminals. H is connected and (G, K) is reduced, so H is just a single terminal adjacent only to v_1 and z.

If $c = n$ then all the non-terminals appear in P. Consider the mincut $\{v_1, v_c\}$. One component of $G \setminus \{v_1, v_c\}$ contains the path $\tau_1, v_2, \tau_2, ..., \tau_{c-1}$, which includes all the non-terminal vertices of $G \setminus \{v_1, v_c\}$. Any other component of $G \setminus \{v_1, v_c\}$ cannot contain any non-terminals; (G, K) is reduced, so such a component is simply a single terminal adjacent to v_1 and v_c.

This completes the proof of the claim. Observe that we have actually proven a more detailed assertion: if $c < n$ there is a terminal adjacent to v_1 and some $z \in S \setminus \{v_1, ..., v_c\}$, and if $c = n$ there is a terminal adjacent to v_1 and v_c. If τ_1 is any terminal then τ_1 is of degree 2, and hence there is a path v_1, τ_1, v_2 in G; applying the assertion repeatedly, we extend this path until we obtain a cycle $v_1, \tau_1, v_2, \tau_2, ..., \tau_{c-1}, v_c, \tau_c, v_1$ with $c = n$.

To complete the proof we show that if (G, K) has any terminal which does not appear in this cycle Γ, or any edge which does not appear in Γ and does not connect two non-terminal vertices adjacent to a given terminal, then (G, K) is equivalent to the standard representation of $U_{n-1,n}$.

Suppose first that (G, K) has a terminal τ which does not appear in Γ; say $N(\tau) = \{v_1, v_i\}$ with $2 < i < n$. If $1 < a < i$ and $i < b \leq n$ then $\{v_a, v_b\} \in U^*_{n-1,n}$. The component of $G \setminus \{v_a, v_b\}$ which contains τ also contains paths within Γ which connect all the non-terminals other than v_a and v_b to either v_1 or v_i. Another component of $G \setminus \{v_a, v_b\}$ cannot contain any non-terminal vertex and hence must be simply a single terminal with neighbor-set $\{v_a, v_b\}$. Applying the same argument to the various values of a and b instead of 1 and i, we conclude that every pair of non-terminals $\{v_p, v_q\}$ is the neighbor-set of a terminal, i.e., (G, K) is equivalent to the standard representation.

Suppose now that every terminal of (G, K) appears in Γ. Suppose further that there is an edge in G between two non-terminal vertices v_i and v_j which are not adjacent to the same terminal. Then $\{v_{i-1}, v_{i+1}\}$ is not a cut, contradicting the fact that $\{v_{i-1}, v_{i+1}\} \in U^*_{n-1,n}$. ∎

Corollary 3.10. *If $n \geq 3$ then $U_{n-1,n}$ is a forbidden minor for* term$(C) \leq n - 1$.

Proof: Contracting one element from $U_{n-1,n}$ yields $U_{n-2,n-1}$, whose minimum terminal number is $n-1$. Deleting one element from $U_{n-1,n}$ yields $U_{n-1,n-1}$, whose minimum terminal number is 2. ∎

Corollary 3.11. *If $n \geq 5$ and $\binom{n-1}{3} < t < \binom{n}{3}$ then $U_{n-2,n}$ is a forbidden minor for* term$(C) \leq t$.

Proof: Theorem 3.6 implies term$(U_{n-2,n}) = \binom{n}{3}$. Contracting one element from $U_{n-2,n}$ yields $U_{n-3,n-1}$, and deleting one element yields $U_{n-2,n-1}$. The corollary follows, because term$(U_{n-3,n-1}) = \binom{n-1}{3} \geq$ term$(U_{n-2,n-1}) = n-1$. ∎

If $1 \leq k \leq n$ then the *circulant clutter* C_n^k contains all the sets of k consecutive elements of $\mathbb{Z}_n$. We use *consecutive* in the natural sense in $\mathbb{Z}_n$; for instance, $C_5^3 = \{\{0,1,2\}, \{1,2,3\}, \{2,3,4\}, \{3,4,0\}, \{4,0,1\}\}$. The particular circulant clutters C_n^2 with n odd are forbidden minors for the width-length inequality [4], and they turn out to be forbidden minors here as well.

Theorem 3.12. *If $n > 1$ is odd then the circulant clutter C_n^2 has no non-standard reduced representation.*

Proof: Any $B \in (C_n^2)^*$ must contain at least one of every two consecutive elements of $\mathbb{Z}_n$; otherwise it would miss an element of C_n^2. It follows that every $B \in (C_n^2)^*$ must contain at least one pair of consecutive elements of $\mathbb{Z}_n$, because n is odd. In addition, a $B \in (C_n^2)^*$ cannot contain three consecutive elements of $\mathbb{Z}_n$, because if $x_1, x_2, x_3 \in B$ are consecutive then $B \setminus x_2$ also intersects all the elements of C_n^2 and thus B is non-minimal. These two conditions — that B must contain one of every two consecutive elements of $\mathbb{Z}_n$ and cannot contain any three consecutive elements — completely characterize the elements of $(C_n^2)^*$.

Let (G, K) be a reduced graphical representation of C_n^2. Certainly (G, K) has at least one terminal τ_0; suppose $N(\tau_0) = B_0 \in (C_n^2)^*$. Let $B_1 = \{x + 1 : x \in B_0\}$, where we interpret $+$ in $\mathbb{Z}_n$. Clearly $B_1 \in (C_n^2)^*$, for it inherits the two characterizing properties from B_0. Every $x \notin B_1$ is an element of B_0, for if $x \notin B_1$ then $x + 1 \in B_1$ and hence $x \in B_0$. Therefore the component of $G \setminus B_1$ which contains τ_0 also contains all the non-terminal vertices of $G \setminus B_1$. Any other component of $G \setminus B_1$ must have only terminal vertices; (G, K) is reduced, so such a component must consist of a single, isolated terminal τ_1. Then $N(\tau_1) \subseteq B_1$ and hence $N(\tau_1) = B_1$; (G, K) is reduced, so there is a unique such τ_1.

We define B_2, B_3, ..., B_{n-1} in the corresponding fashion. From identical arguments, applied first to B_2 and then in succession to the others, we conclude that for each i there is a terminal τ_i with $N(\tau_i) = B_i$. It is not necessarily the case that these n blocker elements are distinct — for instance if $n = 9$ then $B_0 = \{0, 1, 3, 4, 6, 7\} = B_3$ — but this is irrelevant to our argument.

We now show that these terminals $\tau_0, ..., \tau_{n-1}$ provide perfect communication among the non-terminal vertices of G; suppose x is a non-terminal vertex. Some consecutive pair $\{a - 1, a\}$ is contained in B_0; then $\{x - 1, x\} = \{a - 1 + x - a, a + x - a\} \subseteq B_{x-a}$. Clearly x has perfect communication, through τ_{x-a}, with all other elements of B_{x-a}. Every $y \notin B_{x-a}$ is an element of B_{x-a+1}, for if $y \notin B_{x-a}$ then $y + 1 \notin B_{x-a+1}$ and B_{x-a+1} must contain at least one of y and $y + 1$. Since $x - 1 \in B_{x-a}$, $x \in B_{x-a+1}$, so x has perfect communication, through τ_{x-a+1}, with all other elements of B_{x-a+1}. It follows that x has perfect communication with every non-terminal vertex.

The theorem now follows from Corollary 3.3. ∎

The reader can easily verify that in contrast, if $n > 1$ is even then $\mathrm{term}(C_n^2) = 2$.

Corollary 3.13. *If $n > 1$ is odd and $2 \leq t < \left|(C_n^2)^*\right|$ then the circulant clutter C_n^2 is a forbidden minor for $\mathrm{term}(C) \leq t$.*

Proof: We leave it to the reader to verify that contracting or deleting a single element from C_n^2 results in a 2-terminal clutter. ∎

The circulant clutters C_n^2 have the property that $\left|(C_n^2)^*\right| = \left|(C_{n-2}^2)^*\right| + \left|(C_{n-3}^2)^*\right|$ for $n > 4$. It follows that $\left|(C_n^2)^*\right|$ is monotonically increasing for $n > 4$, and hence for every $t \geq 2$ the circulant clutters C_n^2 with n odd and sufficiently large are all forbidden minors for $\mathrm{term}(C) \leq t$. The only other family of clutters we know to have this property is the family of degenerate projective planes.
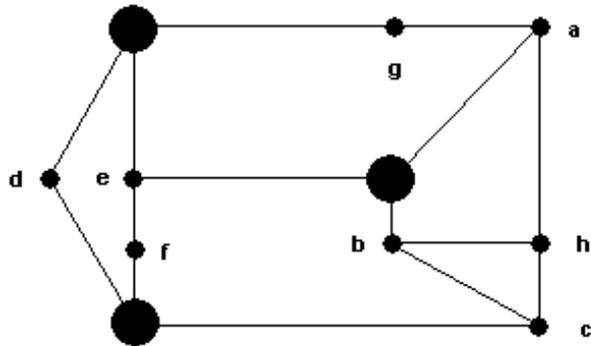
4. Some other examples

Different clutters have very different assortments of reduced graphical representations; we leave it to the interested reader to analyze the following examples.
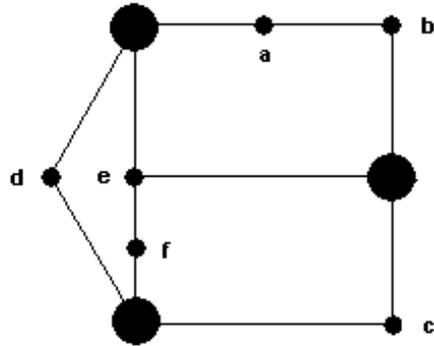
a clutter with six inequivalent reduced representations
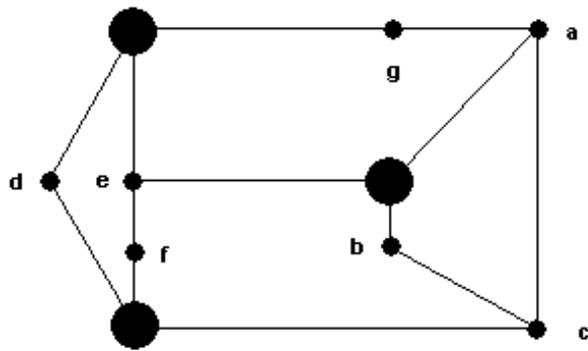
a clutter with three inequivalent reduced representations

a clutter with two inequivalent reduced representations

15

5. 2-terminal clutters

In this section we briefly summarize the special properties of clutters which can be represented using two terminals. Some of these properties have been studied extensively in the literature; see [2] for a survey.

Menger's Theorem tells us that if a clutter has a 2-terminal representation then the maximum number of pairwise disjoint elements of the clutter is the minimum cardinality of an element of its blocker. That is, in the terminology of [10] a 2-terminal clutter *packs*. A more complicated-seeming property is the width-length inequality. According to [4, 11] this property defines a class of clutters which is closed under minors, and none of whose forbidden minors packs. Consequently any minor-closed class of clutters whose elements all pack also has the property that its elements all satisfy the width-length inequality; the 2-terminal clutters constitute such a class.

The clutter Q_6 of edge-sets of triangles in K_4 is $\{\{a, b, d\}, \{b, c, e\}, \{a, e, f\}, \{c, d, f\}\}$; it does not pack. Consider the clutter Q_6^+ obtained from Q_6 by adjoining a common element x to every minpath. Given any graphical representation of Q_6, we obtain a representation of Q_6^+ by replacing some terminal τ by a non-terminal vertex x with $N(x) = N(\tau)$ and a terminal τ' with $N(\tau') = \{x\}$; this shows that $\text{term}(Q_6) \geq \text{term}(Q_6^+)$. The reverse inequality follows from the fact that $Q_6 = Q_6^+/x$, so $\text{term}(Q_6^+) = \text{term}(Q_6) > 2$. This example shows that not all clutters which pack and satisfy the width-length inequality have 2-terminal representations, for Q_6^+ has both properties [2].

By the way, all the minors of Q_6 and its dual are 2-terminal. It turns out that Q_6 is a forbidden minor for $\text{term}(C) \leq t$, $2 \leq t \leq 6$, and Q_6^* is a forbidden minor for $\text{term}(C) \leq t$, $2 \leq t \leq 3$.

While studying this material we often jokingly quoted the motto "Everything is a forbidden minor," because so many of the clutters which have appeared in the literature turn out to be forbidden minors for $\text{term}(C) \leq t$ for some t. One exception is the clutter of lines in the Fano plane, $F_7 = \{\{a, b, d\}, \{b, c, e\}, \{c, d, f\}, \{d, e, g\}, \{e, f, a\}, \{f, g, b\}, \{g, a, c\}\}$. The reader might like to verify that Q_6 is the minor of F_7 obtained by deleting a single element and that $\text{term}(Q_6) = 7 = \text{term}(F_7)$.

If (G, K) is a 2-terminal network with $K = \{s, t\}$ then the clutter of st-paths in (G, K) is the clutter on $E(G)$ whose elements are the edge-sets of simple paths connecting s to t in G. As remarked in [10], these clutters are all 2-terminal.
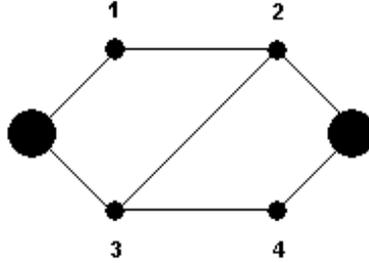
16

Figure 5.1: the clutter P_4

Proposition 5.1. *Clutters of st-paths constitute a proper subclass of the class of 2-terminal clutters.*

Proof: Let (G, K) be a 2-terminal network and C the associated clutter of *st*-paths, defined on $S = E(G)$. To represent C with a 2-terminal network we insert into each edge e of G a degree-2 vertex named e, and replace each non-terminal vertex v of G with the edges of a clique consisting of the newly introduced vertices e such that v is incident on the edge e in G.

To verify that not all 2-terminal clutters are clutters of *st*-paths, consider the 2-terminal network given in Figure 5.1. It represents the clutter $P_4 = \{\{1, 2\}, \{2, 3\}, \{3, 4\}\}$, which is not a clutter of *st*-paths [9]. ■

Acknowledgment

References

[1] C. J. Colbourn, *The Combinatorics of Network Reliability*, Oxford Univ. Press, Oxford, 1987.

[2] G. Cornuéjols, *Combinatorial Optimization: Packing and Covering*, SIAM, Philadelphia, PA, 2001.

[3] J. Edmonds and D. R. Fulkerson, *Bottleneck extrema*, J. Combinatorial Theory **8** (1970), 299-306.

[4] A. Lehman, *On the width-length inequality*, Math. Programming **17** (1979), 403-417.

[5] J. G. Oxley, *Matroid Theory*, Oxford Univ. Press, Oxford, 1992.

[6] K. Reidemeister *Knotentheorie*, Springer, Berlin, 1932.

[7] P. D. Seymour, *The forbidden minors of binary clutters*, J. London Math. Soc. (2) **12** (1976), 356-360.

[8] P. D. Seymour, *A forbidden minor characterization of matroid ports*, Quart. J. Math. Oxford (2) **27** (1976), 407-413.

[9] P. D. Seymour, *A note on the production of matroid minors*, J. Combinatorial Theory (B) **22** (1977), 289-295.

[10] P. D. Seymour, *The matroids with the max-flow min-cut property*, J. Combinatorial Theory (B) **23** (1977), 189-222.

[11] P. D. Seymour, *On Lehman's width-length characterization*, Polyhedral Combinatorics, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, vol. 1, Amer. Math. Soc., Providence, RI, 1990, pp. 107-117.

[12] L. Traldi, *Two universal reliability problems*, Congr. Numer. **132** (1998), 199-204.

[13] L. Traldi, *Visualizing clutters*, Congr. Numer. **154** (2002), 13-20.