

Branch-Width and Well-Quasi-Ordering in Matroids and Graphs¹

James F. Geelen

*Department of Combinatorics and Optimization, University of Waterloo, Waterloo,
Ontario N2L 3G1, Canada*

A. M. H. Gerards

*CWI, Postbus 94079, 1090 GB Amsterdam, The Netherlands; and Department of Mathematics
and Computer Science, Eindhoven University of Technology, Postbus 513,
5600 MB Eindhoven, The Netherlands*

and

Geoff Whittle

*School of Mathematical and Computing Sciences, Victoria University,
P.O. Box 600, Wellington, New Zealand*

Received July 11, 2000

We prove that a class of matroids representable over a fixed finite field and with bounded branch-width is well-quasi-ordered under taking minors. With some extra work, the result implies Robertson and Seymour's result that graphs with bounded tree-width (or equivalently, bounded branch-width) are well-quasi-ordered under taking minors. We will not only derive their result from our result on matroids, but we will also use the main tools for a direct proof that graphs with bounded branch-width are well-quasi-ordered under taking minors. This proof also provides a model for the proof of the result on matroids, with all specific matroid technicalities stripped off. © 2002 Elsevier Science (USA)

Key Words: matroids; graphs; minors; finite fields; connectivity; submodularity; branch-width; tree-width; well-quasi-ordering.

1. INTRODUCTION

We prove the following result.

(1.1) THEOREM. *Let \mathbb{F} be a finite field and n be an integer. Then each*

¹This research was carried out in November 1999 when the authors visited the Fields Institute at the University of Toronto. It was partially supported by the Natural Sciences and Engineering Council of Canada and by a grant from the Marsden Fund of New Zealand.

infinite collection of \mathbb{F} -representable matroids with branch-width at most n has two members such that one is isomorphic to a minor of the other.

The finiteness of the field is crucial here. In Section 7 we shall prove:

(1.2) *Let \mathbb{F} be an infinite field. Then there exists an infinite collection of \mathbb{F} -representable matroids, all with branch width 3, none isomorphic to a minor of another.*

Broadly speaking, a graph or matroid has “small width” if it decomposes across a set of noncrossing separations into small parts. Two standard notions of “width” for graphs, introduced by Robertson and Seymour [5, 6], are tree-width and branch-width. Robertson and Seymour [6] show that these are equivalent, in that a family of graphs has bounded tree-width if and only if it has bounded branch-width. We work with branch-width (to be defined in Sections 2 and 4), which extends in a natural way to matroids (see Section 5).

Theorem 1.1 is an analogue of the following theorem of Robertson and Seymour [5].

(1.3) **THEOREM.** *Let n be an integer. Then each infinite set of graphs with branch-width at most n has two members such that one is isomorphic to a minor of the other.*

To introduce the ideas used in the proof of Theorem 1.1, we provide a direct proof of Theorem 1.3 in Section 4. Of course, as a 3-connected graph is uniquely determined by its cycle matroid, the restriction of Theorem 1.3 to 3-connected graphs immediately follows from Theorem 1.1. On the other hand, by itself Theorem 1.1 says nothing for an infinite set of trees. However, with some extra work, it is possible to derive Theorem 1.3 from Theorem 1.1. We give that derivation in Section 6, even though it is as long as the direct proof in Section 4.

Robertson and Seymour’s proof of Theorem 1.3 relies on a “linked tree-decomposition theorem” of Thomas [7]. One of the main contributions of this paper is a “linked branch-decomposition theorem,” Theorem 2.1. This is a general theorem about symmetric submodular functions and, as such, applies to both matroid and graph connectivity. Our proof of Theorem 2.1 is modeled on Thomas’ proof, but many of the technicalities in his proof are avoided by considering branch-width instead of tree-width. (Bellenbaum and Diestel (see Diestel [1]) derived a short proof of Thomas’ result on tree-width.)

A *quasi-order* is a pair (X, \preceq) , where X is a set and \preceq is a reflexive and transitive binary relation on X . For example, the relation “ A is isomorphic to a minor of B ” is a quasi-order on any set of graphs or

matroids. (The distinction between equality and isomorphism of graphs or matroids is irrelevant in this paper. Hence, we will often use “minor” when meaning “isomorphic to a minor.”) A *well-quasi-order* is a quasi-order (X, \preceq) with the property that for each infinite sequence (x_0, x_1, \dots) in X there exist integers i and j such that $i < j$ and $x_i \preceq x_j$. An *antichain* is a collection of pairwise incomparable elements. A sequence (x_0, x_1, \dots) is *strictly descending* if $x_{i+1} \preceq x_i$ and $x_i \not\preceq x_{i+1}$ for $i \geq 0$. Note that the “minor orderings” for graphs and matroids admit no infinite strictly descending sequences. For a quasi-order with no infinite strictly descending sequences, being well-quasi-ordered is equivalent to having no infinite antichain.

In their fundamental series of papers on graph minors, Robertson and Seymour prove the following remarkable result.

(1.4) GRAPH MINORS THEOREM. *Graphs are well-quasi-ordered under taking minors.*

An important corollary is:

(1.5) GENERAL KURATOWSKI THEOREM. *For each surface S there exist graphs G_1, \dots, G_n such that a graph can be embedded in S if and only if it has none of G_1, \dots, G_n as a minor.*

While the proof of the Graph Minors Theorem still remains mysterious to many, the proof of the General Kuratowski Theorem is now surprisingly accessible. Theorem 1.3 shows that a counterexample to Theorem 1.5 would contain graphs with arbitrarily high branch-width. Diestel, Yu, Gorbunov, Jensen, and Thomassen [2] have a straightforward proof that graphs with high branch-width contain large “grid” minors. Finally, Thomassen [8] has an easy proof that a minor-minimal graph that does not embed in S does not contain a large grid minor.

2. BRANCH DECOMPOSITIONS

A function λ defined on the collection of subsets of a finite *ground set* S is called *submodular* if $\lambda(A) + \lambda(B) \geq \lambda(A \cap B) + \lambda(A \cup B)$ for each $A, B \subseteq S$. We call λ *symmetric* if $\lambda(A) = \lambda(S \setminus A)$ for each $A \subseteq S$. The symmetric submodular functions considered in this paper are the connectivity functions of graphs and matroids (see Sections 4 and 5). For disjoint subsets A and B of S , we denote by $\lambda(A, B)$ the minimum of $\lambda(X)$ over all $X \subseteq S$ containing A and disjoint from B . Clearly, if λ is symmetric on S , then $\lambda(A, B) = \lambda(B, A)$ for each pair of disjoint subsets A and B of S .

A *branch-decomposition* of a symmetric submodular function λ on a finite set S is a cubic tree T (that is, all degrees are 1 or 3) such that S is

contained in the set of leaves of T . The *set displayed* by a given subtree of T is the set of elements of S in that subtree. A set of elements of S is *displayed* by an edge e of T if it is displayed by one of the two components of $T \setminus e$. The *width* $\lambda(e)$ of an edge e in T is the λ -value of either one of the two sets displayed by e . The *width* of a branch decomposition is the maximum of the widths of its edges and the *branch-width* of a symmetric submodular function is the minimum of the widths of all its branch decompositions. See Fig. 2 in Section 4 for a branch decomposition of the connectivity function of the triangular prism.

Note that we allow a branch decomposition to have some leaves that do not correspond to elements of the ground set of the symmetric submodular function; this is for technical reasons in some of the proofs below. We call such leaves *unlabeled*. Branch decompositions with unlabeled leaves are easily turned into branch decompositions with the same width but no unlabeled leaves: just delete the unlabeled leaves and replace pairs of series edges by single edges until the tree is cubic again.

Let f and g be two edges in a branch decomposition T of λ , let F be the set displayed by the component of $T \setminus f$ not containing g , and let G be the set displayed by the component of $T \setminus g$ not containing f . Let P be the shortest path in T containing f and g . Each edge on P displays a subset of S that contains F and is disjoint from G . So the widths of the edges of P are upper bounds for $\lambda(F, G)$. We call f and g *linked* if $\lambda(F, G)$ is equal to the minimum width of an edge on P . A branch decomposition is *linked* if all its edge pairs are linked. The following result is an analogue of Thomas' result [7] on linked tree decompositions of a graph.

(2.1) THEOREM. *An integer-valued symmetric submodular function with branch-width n has a linked branch decomposition of width n .*

Proof. Let λ be an integer-valued symmetric submodular function with branch-width n . For each branch decomposition T of λ we define T_k to be the forest in T induced by the edges with width at least k . (Edge induced subgraphs have no isolated nodes.) For a graph H we denote by $e(H)$ the number of edges in H and by $c(H)$ the number of components of H . If T and R are two branch decompositions of λ we write $T < R$ if there exists a number k such that either $e(T_k) < e(R_k)$ or $e(T_k) = e(R_k)$ and $c(T_k) > c(R_k)$, and such that for each $k' > k$, $e(T_{k'}) = e(R_{k'})$ and $c(T_{k'}) = c(R_{k'})$. This defines a partial order on the branch decompositions of λ . Choose a minimal element T in this partial order. Note that T has width n . We claim that T is linked. Assume not. Choose an unlinked pair of edges f and g in T . Clearly, $f \neq g$. Let F be the set displayed by the component of $T \setminus f$ not containing g , and G the set displayed by the component of $T \setminus g$ not




FIG. 1. Proof of Theorem 2.1.

containing f . Let x be the end vertex of f and y be the end vertex of g such that the xy -path P in T does not contain f or g .

We say that a subset X of S *splits* a subset Y of S if $Y \cap X$ and $Y \setminus X$ are both nonempty. Note that X splitting Y does not imply Y splitting X . Choose a subset A of $S \setminus G$ containing F with $\lambda(A) = \lambda(F, G)$ such that A splits as few subsets of S displayed by edges in T as possible. Define a new tree \hat{T} as follows (see Fig. 1): take a copy T^+ of the component of $T \setminus g$ containing f and a copy T^- of the component of $T \setminus f$ containing g ; connect T^+ with T^- by a new edge a joining the copy of y in T^+ to the copy of x in T^- .

We turn \hat{T} into a branch decomposition of λ as follows: Each element s of S (which is a leaf of T) is identified with its copy in T^+ if $s \in A$ and with its copy in T^- otherwise.

(2.1.1) *Let e be an edge in T and \hat{e} be one of its copies in \hat{T} . Then $\lambda(\hat{e}) \leq \lambda(e)$, with equality only if e has at most one copy in $\hat{T}_{\lambda(A)+1}$.*

In order to prove this, by symmetry, we may assume that \hat{e} lies in T^+ . Let W be the set displayed by the component of $T \setminus e$ not containing y .

Then, $\lambda(e) = \lambda(W)$ and $\lambda(\hat{e}) = \lambda(W \cap A)$. Combined with submodularity this yields $\lambda(\hat{e}) + \lambda(W \cup A) \leq \lambda(e) + \lambda(A) = \lambda(e) + \lambda(F, G) \leq \lambda(e) + \lambda(W \cup A)$. Hence $\lambda(\hat{e}) \leq \lambda(e)$, with equality only if $\lambda(W \cup A) = \lambda(A)$.

Suppose from now on that $\lambda(\hat{e}) = \lambda(e)$. Then $\lambda(W \cup A) = \lambda(A) = \lambda(F, G)$, hence by the choice of A we know that $W \cup A$ splits at least as many sets displayed by edges in T as A does. As the sets displayed by edges in T are pairwise either disjoint or comparable by inclusion, it is straightforward to show that this means that A does not split W . So at least one of $W \cap A$ and $W \setminus A$ is empty. Note that by combining symmetry and submodularity, $2\lambda(B) = \lambda(B) + \lambda(S \setminus B) \geq \lambda(\emptyset) + \lambda(S) = 2\lambda(\emptyset)$ for each $B \subseteq S$. So either $\lambda(W \cap A) \leq \lambda(A)$ or $\lambda(W \setminus A) \leq \lambda(A)$. Recall that $\lambda(\hat{e}) = \lambda(W \cap A)$ and note that if e has a second copy e^* in \hat{T} , and so in T^- , then $\lambda(e^*) = \lambda(W \cup A) = \lambda(A)$ if $e \in P$ and $\lambda(e^*) = \lambda(W \setminus A)$ if $e \notin P$. Hence, at most one of \hat{e} and e^* lies in $\hat{T}_{\lambda(A)+1}$. Thus (2.1.1) follows.

Let p the smallest integer greater than $\lambda(A)$ such that $e(T_k) = e(\hat{T}_k)$ for $k > p$. For each $k \geq p$, it follows from (2.1.1) that each edge of T_k is copied at most once in \hat{T}_k . Moreover, $\lambda(a) = \lambda(A)$, hence $a \notin \hat{T}_k$ for $k > \lambda(A)$. So if $k \geq p$, then $e(T_k) \geq e(\hat{T}_k)$ and $c(T_k) \leq c(\hat{T}_k)$ whenever $e(T_k) = e(\hat{T}_k)$. However, $\hat{T} \not\prec T$, so in fact $e(T_k) = e(\hat{T}_k)$ and $c(T_k) = c(\hat{T}_k)$ for $k \geq p$. Thus also T_p and \hat{T}_p have the same number of edges, which by definition of p implies that $p = \lambda(A) + 1$. Moreover, as $c(\hat{T}_{\lambda(A)+1}) = c(T_{\lambda(A)+1})$, each component of $T_{\lambda(A)+1}$ is copied entirely and as one component in $\hat{T}_{\lambda(A)+1}$. In particular, this is the case for the component of $T_{\lambda(A)+1}$ containing $P \cup \{f, g\}$, which lies entirely in $T_{\lambda(A)+1}$. This is absurd: f has a copy only in T^+ , g has a copy only in T^- , and a is not in $T_{\lambda(A)+1}$. So T is linked, indeed. ■

3. ROBERTSON AND SEYMOUR'S "LEMMA ON TREES"

In order to prove their result on well-quasi-ordering of graphs with bounded tree-width, Robertson and Seymour [5] invoke a "lemma on trees," which they prove in the same paper. The proof of this lemma on trees extends a simple proof by Nash-Williams [4] of the result of Kruskal [3] that forests are well-quasi-ordered under taking minors or actually, more strongly, by "topological containment." We also use Robertson and Seymour's lemma on trees. To make this paper self-contained, we include a proof of that lemma. We need some definitions.

A *rooted tree* is a finite directed tree where all but one of the vertices have indegree 1. A *rooted forest* is a collection of countably many, vertex disjoint rooted trees. Its vertices with indegree 0 are called *roots* and those with outdegree 0 are called *leaves*. The edges leaving a root are *root edges*

and those entering a leaf are *leaf edges*. If S is a set of edges in rooted forest F , then $u_F(S)$ denotes the set of those edges in F whose tail is a head of an edge in S .

An n -edge labeling of a graph G is a map from the edges of G to the set $\{0, \dots, n\}$. Let λ be an n -edge labeling of a rooted forest F and e and f be edges in F . We say that e is λ -linked to f if F contains a directed path P starting with e and ending with f such that $\lambda(g) \geq \lambda(e) = \lambda(f)$ for each edge g on P .

(3.1) LEMMA ON TREES. *Let F be a forest with an n -edge labeling λ . Moreover, let \preceq be a quasi-order on the edges of F with no infinite strictly descending sequence and such that $e \preceq f$ whenever f is λ -linked to e . If the edges of F are not well-quasi-ordered by \preceq then there exists an infinite antichain A of edges of F such that $(u_F(A), \preceq)$ is a well-quasi-order.*

Proof. Assume that this is false and let F , n , and λ form a counterexample with n minimal. This means that any n -edge labeled forest with no label equal to 0 satisfies the lemma (just subtract 1 from all labels). Moreover, any n -edge labeled forest in which the edges labeled 0 are well-quasi-ordered satisfies the lemma, as otherwise deleting these edges would yield a forest contradicting the lemma in spite of the fact that none of its labels is 0.

Let N be the set of edges in F with label 0. Note that an edge $e \in N$ is λ -linked to an edge $f \in N$ if and only if there exists a directed path in F starting with e and ending with f . A sequence (a_1, a_2, \dots) is called *independent* if $a_i \not\preceq a_j$ whenever $i < j$. Choose an infinite independent sequence (a_1, a_2, \dots) in N with the following property:

(3.1.1) *For each k , if a_k is λ -linked to some $e \in N \setminus \{a_k\}$, the sequence (a_1, \dots, a_{k-1}, e) cannot be extended to an infinite independent sequence in N .*

It is straightforward to prove that such a sequence does exist. Moreover, no two elements of (a_0, a_1, \dots) are λ -linked. As F is a counterexample and $\{a_1, a_2, \dots\}$ contains an infinite antichain, the set $u_F(a_1, a_2, \dots)$ is not well-quasi-ordered. The maximal subforest, R , of F with all root edges in $u_F(a_1, a_2, \dots)$ is a counterexample as well, as each set of edges A in R satisfies $u_R(A) = u_F(A)$. So, R contains an infinite independent sequence (b_1, b_2, \dots) of edges labeled 0. By construction of R , for each integer j there exists a unique integer $s(j)$ such that $a_{s(j)}$ is λ -linked to b_j . Choose ℓ with $s(\ell)$ minimal. Then for each $j \geq \ell$ and each $i < s(\ell)$ ($\leq s(j)$) we have that $b_j \preceq a_{s(j)}$ and $a_i \not\preceq a_{s(j)}$, so $a_i \not\preceq b_j$. Hence the infinite sequence $(a_1, \dots, a_{s(\ell)-1}, b_\ell, b_{\ell+1}, \dots)$ is independent as well. This contradicts (3.1.1), so the lemma follows. ■

Now we extract from the Lemma on Trees exactly what we need. A *binary forest* is a rooted orientation of a cubic forest with a distinction

between left and right outgoing edges. More precisely, we call a triple (F, ℓ, r) a *binary forest* if F is a rooted forest where the roots have out-degree 1 and ℓ and r are functions defined on the nonleaf edges of F , such that the head of each nonleaf edge e of F has exactly two outgoing edges, namely $\ell(e)$ and $r(e)$.

(3.2) **LEMMA ON CUBIC TREES.** *Let (F, ℓ, r) be an infinite binary forest with an n -edge labeling λ . Moreover, let \preceq be a quasi-order on the edges of F with no infinite strictly descending sequences, such that $e \preceq f$ whenever f is λ -linked to e . If the leaf edges of F are well-quasi-ordered by \preceq but the root edges of F are not, then F contains an infinite sequence (e_0, e_1, \dots) of nonleaf edges such that:*

- (i) $\{e_0, e_1, \dots\}$ is an antichain with respect to \preceq ;
- (ii) $\ell(e_0) \preceq \dots \preceq \ell(e_{i-1}) \preceq \ell(e_i) \preceq \dots$;
- (iii) $r(e_0) \preceq \dots \preceq r(e_{i-1}) \preceq r(e_i) \preceq \dots$.

Proof. Applying the Lemma on Trees (Lemma 3.1) to the rooted forest F , λ , and \preceq , we see that there exists an antichain A of edges such that $(u_F(A), \preceq)$ is a well-quasi-order. As the leaf edges of F are well-quasi-ordered, we may assume that A contains no leaf edge. It is straightforward to deduce now that A contains an infinite sequence (e_0, e_1, \dots) as claimed. ■

4. GRAPHS WITH BOUNDED BRANCH-WIDTH

Let $G = (V, E)$ be a graph. For $A \subseteq E$, we denote by $\Gamma_G(A)$ the set of vertices that are incident with an edge in A and also with an edge in $E \setminus A$. The *connectivity function* γ_G of G is defined by $\gamma_G(A) := |\Gamma_G(A)|$ for $A \subseteq E$. The *branch-decompositions* and *branch-width* of G are the branch-decompositions and the branch-width of the connectivity function of G , which is symmetric and submodular. See Fig. 2 for a branch-decomposition of the connectivity function of the triangular prism.

A *rooted graph* is a pair (G, X) where X is a subset of the vertex set of graph G . (Robertson and Seymour [5] use rooted graphs as well, but in their case X is an ordered set; here it is not.) The rooted graph obtained from (G, X) by *deleting* an edge uv is $(G \setminus uv, X)$. The rooted graph obtained by *contracting* an edge uv into a vertex w is $(G/uv, X')$, where $X' := X$ if u and v are not in X and $X' := (X \setminus \{u, v\}) \cup \{w\}$ otherwise. A *minor* of (G, X) is any rooted graph obtained by a sequence of deletions and contractions, and possibly by deleting isolated vertices that are not in X . The “minor ordering” on rooted graphs is clearly a quasi-order. (As for




FIG. 2. The triangular prism and a branch decomposition with width 4 of its connectivity function. Interchanging h and i in the branch decomposition yields a decomposition of width 3, the branch width of the triangular prism.

graphs and matroids we will also for rooted graphs often use “minor” when we mean “isomorphic to a minor.”)

Let (G, X) and (H, Y) be two rooted graphs with $|X| = |Y|$. Two graphs that are both obtained from (G, X) and (H, Y) by identifying the vertices in X one-to-one with the vertices in Y may be nonisomorphic (depending on which vertices are identified). However, up to isomorphism, there are only finitely many graphs—at most $|X|$ -factorial—that can be obtained by such identification. This is the crux of the proof of Theorem 1.3.

In proving Theorem 1.3, we will use branch decompositions that are linked. The benefit lies in the following variant of Menger’s theorem.

(4.1) *Let $E_1 \subseteq E_2$ be subsets of the edge set E of a graph G . For $i = 1, 2$, let G_i be the subgraph of G induced by E_i . If $\gamma_G(E_1) = \gamma_G(E_1, E \setminus E_2) = \gamma_G(E_2)$, then $(G_1, \Gamma_G(E_1))$ is a minor of $(G_2, \Gamma_G(E_2))$.*

Proof. By Menger’s theorem, the graph induced by $E_2 \setminus E_1$ contains a collection of $\gamma_G(E_1, E \setminus E_2)$ vertex disjoint paths from $\Gamma_G(E_1)$ to $\Gamma_G(E_2)$. Contracting these paths in $(G_2, \Gamma_G(E_2))$ and deleting all remaining edges in $E_2 \setminus E_1$ yields $(G_1, \Gamma_G(E_1))$. ■

Proof of Theorem 1.3. Let \mathcal{G} denote the set of graphs with branch-width at most n and assume it is not well-quasi-ordered by minor containment. For each $G \in \mathcal{G}$, let T_G be a linked branch decomposition of G with width at most n . We clearly may choose T_G such that at least one leaf corresponds to no element in G (otherwise, subdivide an edge of the tree and add a pendant edge to make it cubic again). Fix an unlabeled leaf r and orient T_G such that it becomes a rooted cubic tree with r as root. For an edge e of T_G , let E^e be the set of edges of G displayed by the component

of $T_G \setminus e$ not containing the root of T_G . Moreover, we define G^e to be the subgraph of G induced by E^e , the set $X^e := \Gamma_G(E^e)$, and $\lambda(e) := \gamma_G(E^e)$.

Let (F, ℓ, r) be the rooted binary forest composed of the rooted cubic trees T_G ($G \in \mathcal{G}$). We define a quasi-order \preceq on the edges of F as follows: If e, f are edges of F and the rooted graph (G^e, X^e) is isomorphic to a minor of the rooted graph (G^f, X^f) , then $e \preceq f$.

We next check that these objects satisfy all conditions in Lemma 3.2. It follows from (4.1) that $e \preceq f$ whenever f is λ -linked to e . Clearly the quasi-order \preceq has no infinite strictly descending sequences. The leaf edges of F are well-quasi-ordered by \preceq , as each of them corresponds to a rooted graph with at most one edge. The root edges are not well-quasi-ordered by \preceq as the associated rooted graphs are (G, \emptyset) with $G \in \mathcal{G}$. So indeed, (F, ℓ, r) , λ , and \preceq satisfy all the conditions of Lemma 3.2.

Consequently, there exists an infinite sequence (e_0, e_1, \dots) of nonleaf edges of F satisfying (i), (ii), and (iii) of Lemma 3.2. Each $X^{\ell(e_i)}$ and each $X^{r(e_i)}$ has at most n elements. So, by taking an infinite subsequence of (e_0, e_1, \dots) , we may assume that the sets $X^{\ell(e_i)}$ all have the same cardinality and also the sets $X^{r(e_i)}$ all have the same cardinality.

By (ii) of (3.2), for each $i = 1, 2, \dots$ we can label each vertex in $X^{\ell(e_i)}$ by a different *left label* from $\{1, \dots, n\}$ such that for each $i < j$, $G^{\ell(e_i)}$ can be obtained as a minor of $G^{\ell(e_j)}$ in such a way that a vertex in $X^{\ell(e_i)}$ goes to the vertex in $X^{\ell(e_j)}$ with the same left label. By (iii) and (iv) of (3.2), we can assign in a similar way *right labels* from $\{1, \dots, n\}$ to the vertices in $X^{r(e_1)}, X^{r(e_2)}, \dots$. Note that vertices in $X^{\ell(e_j)} \cap X^{r(e_j)}$ obtain both a right and a left label. As the left and right labels all come from the same finite set $\{1, \dots, n\}$ there has to exist an index i and an index $j > i$ such that the following two properties hold:

(4.2) *The set of left/right label pairs that are assigned to vertices in $X^{\ell(e_i)} \cap X^{r(e_i)}$ is the same as the set of these pairs assigned to vertices in $X^{\ell(e_j)} \cap X^{r(e_j)}$.*

(4.3) *The set of left (right) labels assigned to X^{e_i} equals the set of left (right) labels assigned to X^{e_j} .*

For each nonleaf edge e of F , G^e can be seen as obtained from $G^{\ell(e)}$ and $G^{r(e)}$ by identifying the vertices in $X^{\ell(e)} \cap X^{r(e)}$. Hence, by the definition of the left and right labels, (4.2) implies that G^{e_i} can be obtained as a minor of G^{e_j} such that each vertex $X^{\ell(e_j)} \cap X^{r(e_j)}$ goes to a vertex in $X^{\ell(e_i)} \cap X^{r(e_i)}$ with the same left and/or right label. Combining this with (4.3), we see that (G^{e_i}, X^{e_i}) is a minor of (G^{e_j}, X^{e_j}) . In other words, $e_i \preceq e_j$. As this contradicts (i) of Lemma 3.2, Theorem 1.3 follows. ■

5. MATROIDS, CONNECTIVITY AND BRANCH-WIDTH

In this section we define branch-width for matroids and prove our main result, Theorem 1.1. The proof goes along the same lines as the proof of Theorem 1.3.

Matroid Connectivity—Tutte's Linking Theorem

If M is a matroid on ground set S and with rank function r_M , then the *connectivity* $\lambda_M(A)$ of a subset A of S is defined by $\lambda_M(A) := r_M(A) + r_M(S \setminus A) - r_M(S) + 1$. The *branch decompositions* and *branch-width* of M are the branch decompositions and the branch-width of the connectivity function λ_M , which is symmetric and submodular.

Note that when A is a set of edges of a graph $G = (V, E)$, and $M(G)$ is the cycle matroid of G , then $\lambda_{M(G)}(A)$ is equal to $\gamma_G(A) + c(E) - c(A) - c(E \setminus A) + 1$, where $c(B)$ denotes the number of components of the graph induced by the edges in subset B of E . (So, when E , A , and $E \setminus A$ induce connected graphs, $\lambda_{M(G)}(A) = \gamma_G(A)$.) In spite of this slight difference between the two connectivity functions, they are similar enough to consider matroid connectivity as an extension of graph connectivity.

One of the main ingredients in proving Theorem 1.3 is (4.1), a variant of Menger's theorem. In proving Theorem 1.1, its role will be taken over by the following, not so well-known, result of Tutte [9]. It generalizes Menger's theorem to matroids.

(5.1) TUTTE'S LINKING THEOREM. *Let X and Y be disjoint subsets of a matroid M . Then $\lambda_M(X, Y) \geq n$ if and only if there exists a minor M' of M with ground set $X \cup Y$ such that $\lambda_{M'}(X) \geq n$.*

Because it links connectivity with the existence of minors, this result plays a central role in our proof of Theorem 1.1. In order to keep the paper self-contained, we include its proof. It uses two easy results, which we derive first. The following very useful inequality relates the connectivities in a matroid with the connectivities in its minors:

(5.2) *If x is an element of M and A and B are subsets of $S \setminus \{x\}$, then*

$$\lambda_{M \setminus x}(A) + \lambda_{M/x}(B) \geq \lambda_M(A \cap B) + \lambda_M(A \cup B \cup \{x\}) - 1.$$

Proof. This inequality is an immediate consequence of the definition of the connectivity function and the submodularity of the rank function,

$$\begin{aligned}
 & \lambda_{M \setminus x}(A) + \lambda_{M/x}(B) \\
 &= r_{M \setminus x}(A) + r_{M \setminus x}(S \setminus (A \cup \{x\})) - r_{M \setminus x}(S \setminus \{x\}) + 1 \\
 & \quad + r_{M/x}(B) + r_{M/x}(S \setminus (B \cup \{x\})) - r_{M/x}(S \setminus \{x\}) + 1 \\
 &= r_M(A) + r_M(S \setminus (A \cup \{x\})) - r_M(S \setminus \{x\}) \\
 & \quad + r_M(B \cup \{x\}) + r_M(S \setminus B) - r_M(S) - r_M(\{x\}) + 2 \\
 &\geq r_M(A) + r_M(B \cup \{x\}) + r_M(S \setminus (A \cup \{x\})) + r_M(S \setminus B) - 2r_M(S) + 1 \\
 &\geq r_M(A \cap B) + r_M(A \cup B \cup \{x\}) + r_M(S \setminus (A \cup B \cup \{x\})) \\
 & \quad + r_M(S \setminus (A \cap B)) - 2r_M(S) + 1 \\
 &= \lambda_M(A \cap B) + \lambda_M(A \cup B \cup \{x\}) - 1. \quad \blacksquare
 \end{aligned}$$

The following fact characterizes when the connectivity of a set is preserved in a minor.

(5.3) *Let $X, C,$ and D be disjoint subsets of the ground set S of a matroid M . Then $\lambda_{M \setminus D/C}(X) \leq \lambda_M(X)$, with equality if and only if $r_M(X \cup C) = r_M(X) + r_M(C)$ and $r_M(S \setminus X) + r_M(S \setminus D) = r_M(S) + r_M(S \setminus (X \cup D))$.*

Proof. Recall that $r_{M \setminus D/C}(A) = r_M(A \cup C) - r_M(C)$ for each $A \subseteq S \setminus (C \cup D)$. Hence

$$\begin{aligned}
 & \lambda_M(X) - \lambda_{M \setminus D/C}(X) \\
 &= r_M(X) + r_M(S \setminus X) - r_M(S) + 1 \\
 & \quad - [r_{M \setminus D/C}(X) + r_{M \setminus D/C}(S \setminus (X \cup D \cup C)) - r_{M \setminus D/C}(S \setminus (D \cup C)) + 1] \\
 &= r_M(X) + r_M(S \setminus X) - r_M(S) \\
 & \quad - r_M(X \cup C) + r_M(C) - r_M(S \setminus (X \cup D)) + r_M(C) + r_M(S \setminus D) - r_M(C) \\
 &= [r_M(X) + r_M(C) - r_M(X \cup C)] \\
 & \quad + [r_M(S \setminus X) + r_M(S \setminus D) - r_M(S \setminus (X \cup D)) - r_M(S)].
 \end{aligned}$$

As each of the two square-bracketed forms is nonnegative, (5.3) follows. \blacksquare

We now prove Tutte's Linking Theorem (Theorem 5.1).

Proof of (5.1). By Theorem 5.3, M' cannot exist if $\lambda_M(X, Y) < n$. We prove the converse statement by induction on $|S \setminus (X \cup Y)|$. If $S = X \cup Y$

the statement is trivial, so suppose that is not the case; let $x \in S \setminus (X \cup Y)$. If the minor M' as claimed in Theorem 5.1 does not exist then, by induction, there exist sets A and B in $S \setminus (Y \cup \{x\})$ both containing X such that $\lambda_{M \setminus x}(A) \leq n-1$ and $\lambda_{M/x}(B) \leq n-1$. Hence, by (5.2), either $\lambda(A \cap B)$ or $\lambda(A \cup B \cup \{x\})$ is at most $n-1$. In other words, $\lambda_M(X, Y) < n$. So Theorem 5.1 follows. ■

Represented Matroids—(Rooted) Configurations and Minors

Throughout this section \mathbb{F} is a fixed field. Typically representations of matroids over a field are described as matrices over the field, where each column corresponds to a matroid element. Here it is more convenient to represent matroids as “configurations.” A *configuration* is a finite set of labeled points in some linear space over \mathbb{F} . Like columns in a matrix, points in a configuration may coincide, but labels do not. So the labels just serve to distinguish between points whose locations coincide and make it possible to consider configurations as sets and not as multisets. Two configurations are *isomorphic* if there is a bijection between the labels that preserves the points.

We denote the linear span of a configuration A by $\langle A \rangle$ (considered as a space of unlabeled points). A *rooted configuration* is a pair (A, V) where A is a configuration and V is a subspace of $\langle A \rangle$. We will glue rooted configurations together by identifying parts of these subspaces, just as we glued rooted graphs together by identifying the specified subsets of their vertices.

A configuration A is a *minor* of a configuration A' if there exists a linear transformation \mathcal{L} from $\langle A' \rangle$ to $\langle A \rangle$ such that $\langle A \rangle = \mathcal{L}(\langle A' \rangle)$, $\ker(\mathcal{L})$ is the linear span of some subset of A' , and $A \subseteq \mathcal{L}(A')$. (Here $\mathcal{L}(A')$ is the configuration consisting of the labeled points $\mathcal{L}(a)$ ($a \in A'$) where each $\mathcal{L}(a)$ has the same label as a .) If \mathcal{L} satisfies all this, we write $A \stackrel{\mathcal{L}}{\leftarrow} A'$. If moreover $\mathcal{L}(V') = V$ for linear spaces V in $\langle A \rangle$ and V' in $\langle A' \rangle$, we say that (A, V) is a *minor* of (A', V') and we write $(A, V) \stackrel{\mathcal{L}}{\leftarrow} (A', V')$ or just $(A, V) \leftarrow (A', V')$.

As for matroids and (rooted) graphs, also for (rooted) configurations we often just write “minor” when we mean “isomorphic to a minor.” As of now we will refrain from mentioning labels explicitly; it would only clutter the exposition.

The matroid $M(A)$ *represented* by the configuration A is the matroid with ground set A in which independence is just linear independence over the field \mathbb{F} . Different configurations may represent the same matroid; for instance, multiplying a single vector by a nonzero member of \mathbb{F} changes the configuration, not the matroid. The following is obvious:

(5.4) If $A \xrightarrow{\mathcal{L}} A'$, then $M(A)$ is obtained from $M(A')$ by contracting a subset X of $\ker(\mathcal{L}) \cap A'$ that spans $\ker(\mathcal{L})$, adding back a loop for each member of X , and finally taking the restriction to those elements of A' mapped by \mathcal{L} to A . Conversely, for each minor M of $M(A')$ there exists a linear transformation \mathcal{L} and a configuration A such that M is equal to $M(A)$ and $A \xrightarrow{\mathcal{L}} A'$.

The following says that “ \leftarrow ” is a quasi-ordering of configurations.

(5.5) If $A \xrightarrow{\mathcal{L}} A'$ and $A' \xrightarrow{\mathcal{L}'} A''$, then $A \xrightarrow{\mathcal{L}\mathcal{L}'} A''$.

Proof. As, clearly, $\mathcal{L}\mathcal{L}'(\langle A'' \rangle) = \langle A \rangle$ and $A \subseteq \mathcal{L}\mathcal{L}'(A'')$, we only need to prove that $\ker(\mathcal{L}\mathcal{L}')$ is the linear span of some subcollection of A'' . For that, let $A'_0 \subseteq A'$, respectively $A''_0 \subseteq A''$, be sets spanning $\ker(\mathcal{L})$, respectively $\ker(\mathcal{L}')$, and choose $A''_1 \subseteq A''$ such that $\mathcal{L}'(A''_1) = A'_0$. By standard linear algebra, $\mathcal{L}'(\ker(\mathcal{L}\mathcal{L}')) = \ker(\mathcal{L}) \cap \mathcal{L}'(\langle A'' \rangle) = \langle A'_0 \rangle = \mathcal{L}'(\langle A''_1 \rangle)$. Hence, $\ker(\mathcal{L}\mathcal{L}') \subseteq \langle A''_1 \rangle + \ker(\mathcal{L}') = \langle A''_1 \rangle + \langle A''_0 \rangle = \langle A''_1 \cup A''_0 \rangle$. As $A''_1 \cup A''_0$ is clearly contained in $\ker(\mathcal{L}\mathcal{L}')$, we get that the latter is indeed the span of a subcollection of A'' , so (5.5) follows. ■

For a configuration A and $X \subseteq A$, we define $\Gamma_A(X) := \langle X \rangle \cap \langle A \setminus X \rangle$. Note that the dimension of this “separating” space is $\lambda_{M(A)}(X) - 1$. The following result is the translation of Tutte’s linking theorem (Theorem 5.1) to rooted configurations.

(5.6) If A_1 and A_2 are disjoint subcollections of A with $\lambda_{M(A)}(A_1) = \lambda_{M(A)}(A_1, A_2) = \lambda_{M(A)}(A_2)$, then $(A_1, \Gamma_A(A_1)) \leftarrow (A \setminus A_2, \Gamma_A(A \setminus A_2))$.

Proof. It follows from (5.1) that there exist disjoint collections C and D partitioning $A \setminus (A_1 \cup A_2)$ such that $\lambda_{M(A) \setminus D/C}(A_1) = \lambda_{M(A)}(A_1, A_2) = \lambda_{M(A)}(A_1)$, where we may assume D to be coindependent; that is, $\langle A \setminus D \rangle = \langle A \rangle$.

By Theorem 5.3, $r_M(A \setminus A_1) + r_M(A \setminus D) = r_M(A) + r_M(A \setminus (A_1 \cup D))$. Hence, as D is coindependent, $r_M(A \setminus A_1) = r_M(A \setminus (A_1 \cup D)) = r_M(A_2 \cup C)$, or equivalently, $\langle A \setminus A_1 \rangle = \langle A_2 \cup C \rangle$. Moreover, by symmetry, $\langle A \setminus A_2 \rangle = \langle A_1 \cup C \rangle$.

Again by Theorem 5.3, $\langle A_1 \rangle \cap \langle C \rangle = \{0\}$. So there exists a linear transformation \mathcal{L} defined on $\langle A \rangle$ such that $\ker(\mathcal{L}) = \langle C \rangle$ and $\mathcal{L}(x) = x$ for each $x \in A_1$. It follows from $\mathcal{L}(\langle A \setminus A_2 \rangle) = \mathcal{L}(\langle A_1 \cup C \rangle) = \langle A_1 \rangle$, $A_1 = \mathcal{L}(A_1) \subseteq \mathcal{L}(A \setminus A_2)$, $\ker(\mathcal{L}) = \langle C \rangle$, and $C \subseteq A \setminus A_2$, that $A_1 \xrightarrow{\mathcal{L}} A \setminus A_2$.

So it remains to prove that $\mathcal{L}(\Gamma_A(A \setminus A_2)) = \Gamma_A(A_1)$, in other words, that $\mathcal{L}(\langle A \setminus A_2 \rangle \cap \langle A_2 \rangle) = \langle A_1 \rangle \cap \langle A \setminus A_1 \rangle$. In order to do so recall that each triple X, Y, Z of subspaces with $Z \subseteq X$ satisfies $(X \cap Y) + Z = X \cap (Y + Z)$. Hence, $(\langle A \setminus A_2 \rangle \cap \langle A_2 \rangle) + \langle C \rangle = (\langle A_1 \cup C \rangle \cap \langle A_2 \rangle) + \langle C \rangle = \langle A_1 \cup C \rangle \cap (\langle A_2 \rangle + \langle C \rangle) = \langle A_1 \cup C \rangle \cap \langle A_2 \cup C \rangle$. However, by the

symmetry between A_1 and A_2 , this means that $(\langle A \setminus A_2 \rangle \cap \langle A_2 \rangle) + \langle C \rangle = (\langle A \setminus A_1 \rangle \cap \langle A_1 \rangle) + \langle C \rangle$. Hence, as $\langle C \rangle = \ker(\mathcal{L})$ and as \mathcal{L} is the identity on A_1 , it follows that $\mathcal{L}(\langle A \setminus A_2 \rangle \cap \langle A_2 \rangle) = \mathcal{L}(\langle A_1 \rangle \cap \langle A \setminus A_1 \rangle) = \langle A_1 \rangle \cap \langle A \setminus A_1 \rangle$. ■

For later reference we rewrite (5.6) by replacing A_2 with $A \setminus A_2$. Note the resemblance with (4.1).

(5.7) *If $A_1 \subseteq A_2 \subseteq A$ with $\lambda_{M(A)}(A_1) = \lambda_{M(A)}(A_1, A \setminus A_2) = \lambda_{M(A)}(A_2)$, then $(A_1, \Gamma_A(A_1)) \leftarrow (A_2, \Gamma_A(A_2))$.*

We actually prove a variant of Theorem 1.1 for configurations. In proving Theorem 1.1 we may restrict ourselves to loopless matroids. So by (5.4), Theorem 1.1 is an immediate consequence of Theorem 5.8 below.

(5.8) THEOREM. *Let \mathbb{F} be a fixed finite field and $n \in \mathbb{N}$. Then each infinite set of configurations over \mathbb{F} with branch-width at most n has two members such that one is a minor of the other.*

Proof. Let \mathcal{A} be the collection of configurations over \mathbb{F} with branch-width at most n and assume that it is not well-quasi-ordered by minor-containment. As in the proof of Theorem 1.3, we will set up an appropriate framework that enables us to apply Lemma 3.2.

For each $A \in \mathcal{A}$, let T_A be a linked branch decomposition of A with width at most n . We clearly may choose T_A such that at least one leaf corresponds to no element in A . Fix an unlabeled leaf r and orient T_A such that it becomes a rooted cubic tree with r as root. For an edge e of T_A , let A^e be the set of elements of A displayed by the component of $T_A \setminus e$ not containing the root of T_A . Moreover, we define the subspace $X^e := \Gamma_A(A^e)$ and $\lambda(e) := \lambda_{M(A)}(A^e) = \dim(X^e) + 1$. We call (A^e, X^e) the *rooted configuration associated with e* .

Let (F, ℓ, r) be the rooted binary forest composed of the rooted cubic trees T_A ($A \in \mathcal{A}$). If e is a nonleaf edge of F , then

$$(5.8.1) \quad A^e = A^{\ell(e)} \cup A^{r(e)}, \quad X^e \subseteq X^{\ell(e)} + X^{r(e)}, \quad \text{and} \\ X^{\ell(e)} \cap X^{r(e)} = \langle A^{\ell(e)} \rangle \cap \langle A^{r(e)} \rangle.$$

Indeed, the first statement is obvious and the second one follows by $X^e = \langle A^e \rangle \cap \langle A \setminus A^e \rangle = \langle A^{\ell(e)} \cup A^{r(e)} \rangle \cap \langle (A \setminus A^{\ell(e)}) \cap (A \setminus A^{r(e)}) \rangle \subseteq (\langle A^{\ell(e)} \rangle + \langle A^{r(e)} \rangle) \cap \langle A \setminus A^{\ell(e)} \rangle \cap \langle A \setminus A^{r(e)} \rangle \subseteq (\langle A^{\ell(e)} \rangle \cap \langle A \setminus A^{\ell(e)} \rangle) + (\langle A^{r(e)} \rangle \cap \langle A \setminus A^{r(e)} \rangle) = X^{\ell(e)} + X^{r(e)}$. Finally, $\langle A^{\ell(e)} \rangle \cap \langle A^{r(e)} \rangle \subseteq \langle A^{\ell(e)} \rangle \cap \langle A \setminus A^{\ell(e)} \rangle = X^{\ell(e)} \subseteq \langle A^{\ell(e)} \rangle$ and, by symmetry, $\langle A^{\ell(e)} \rangle \cap \langle A^{r(e)} \rangle \subseteq X^{r(e)} \subseteq \langle A^{r(e)} \rangle$. So indeed, $X^{\ell(e)} \cap X^{r(e)} = \langle A^{\ell(e)} \rangle \cap \langle A^{r(e)} \rangle$, and (5.8.1) follows. ■

Finally, define the quasi-order \preceq on the edges of F as $e \preceq f$ if $(A^e, X^e) \leftarrow (A^f, X^f)$. We have constructed a binary forest (F, ℓ, r) with an n -edge labeling λ and a quasi-ordering \preceq on its edges.

We next check that these objects satisfy all conditions in Lemma 3.2. It follows from (5.7) that $e \preceq f$ whenever f is λ -linked to e . Clearly the quasi-order \preceq has no infinite strictly descending sequences. The leaf edges of F are well-quasi-ordered by \preceq , as each of them corresponds to a rooted configuration with at most one element. The root edges are not well-quasi-ordered as the associated rooted configurations are $(A, \{0\})$ with $A \in \mathcal{A}$. So indeed, (F, ℓ, r) , λ , and \preceq satisfy all conditions of Lemma 3.2.

Consequently, an infinite sequence (e_0, e_1, \dots) as claimed in Lemma 3.2 exists. To simplify notation, let (A_i^ℓ, X_i^ℓ) , (A_i^r, X_i^r) , and (A_i, X_i) be the rooted configurations associated with, respectively, $\ell(e_i)$, $r(e_i)$, and e_i .

For each $i = 0, 1, \dots$, the subspace $X_i^\ell + X_i^r$ has dimension at most $2(n-2)$. By replacing (e_0, e_1, \dots) with an infinite subsequence, we may assume that all subspaces $X_i^\ell + X_i^r$ have the same dimension. Hence, by isomorphically changing each $\langle A_i \rangle$ (and with it its subsets $A_i^\ell, X_i^\ell, A_i^r$, etc.) we may assume that in fact all $X_i^\ell + X_i^r$ are equal to one and the same linear space. As that latter space is a finite set containing each X_i^ℓ, X_i^r , and X_i , the triple (X_i^ℓ, X_i^r, X_i) can take only finitely many values. So some value, (X^ℓ, X^r, X) say, is repeated infinitely often. In other words, by replacing (e_0, e_1, \dots) with an infinite subsequence, we may assume that $X_i^\ell = X^\ell, X_i^r = X^r$, and $X_i = X$ for every i .

By (i) and (ii) in Lemma 3.2, there exist maps \mathcal{L}_i and \mathcal{R}_i , such that for $i \geq 1$,

$$(5.8.2) \quad (A_{i-1}^\ell, X^\ell) \xleftarrow{\mathcal{L}_i} (A_i^\ell, X^\ell) \quad \text{and} \quad (A_{i-1}^r, X^r) \xleftarrow{\mathcal{R}_i} (A_i^r, X^r),$$

and by (iii) in Lemma 3.2 for each $i < j$,

$$(5.8.3) \quad (A_i, X) \leftarrow (A_j, X).$$

Consider, for each $i = 1, 2, \dots$, the restriction π_i of the product $\mathcal{L}_1 \cdots \mathcal{L}_i$ to X^ℓ and the restriction ψ_i of $\mathcal{R}_1 \cdots \mathcal{R}_i$ to X^r . The maps π_i and ψ_i are automorphisms of X^ℓ and X^r , respectively. As X^ℓ and X^r are finite sets, there exists an i and a $j > i$ such that the pair (π_i, ψ_i) is equal to the pair (π_j, ψ_j) . The restriction $\pi_i^{-1}\pi_j$ of $\mathcal{L} := \mathcal{L}_{i+1} \cdots \mathcal{L}_j$ to X^ℓ and the restriction $\psi_i^{-1}\psi_j$ of $\mathcal{R} := \mathcal{R}_{i+1} \cdots \mathcal{R}_j$ to X^r are identity maps. Clearly, $(A_i^\ell, X^\ell) \xleftarrow{\mathcal{L}} (A_j^\ell, X^\ell)$ and $(A_i^r, X^r) \xleftarrow{\mathcal{R}} (A_j^r, X^r)$.

The linear transformations \mathcal{L} and \mathcal{R} coincide on $X^\ell \cap X^r$, which is, by (5.8.1), the intersection $\langle A_j^\ell \rangle \cap \langle A_j^r \rangle$ of their domains. So \mathcal{L} and \mathcal{R} have a

uniquely defined common extension to a linear transformation, \mathcal{T} say, on the sum $\langle A_j \rangle = \langle A_j^\ell \rangle + \langle A_j^r \rangle$ of their domains.

$$(5.8.4) \quad A_i \stackrel{\mathcal{T}}{\leftarrow} A_j.$$

Obviously, $\mathcal{T}(\langle A_j \rangle) = \langle A_i \rangle$ and $\mathcal{T}(A_j) \subseteq A_i$. To prove that $\ker(\mathcal{T})$ is the span of a subcollection of A_j , it suffices to prove that $\ker(\mathcal{T})$ is equal to its subspace $\ker(\mathcal{L}) + \ker(\mathcal{R})$, as this space is the span of a subcollection of $A_j^\ell \cup A_j^r = A_j$. For this, let $x \in \ker(\mathcal{T})$. Then $x = x_\ell + x_r$ for some $x_\ell \in \langle A_j^\ell \rangle$ and $x_r \in \langle A_j^r \rangle$. As $\mathcal{L}(x_\ell) + \mathcal{R}(x_r) = \mathcal{T}(x) = 0$, the vector $\mathcal{L}(x_\ell) = \mathcal{R}(-x_r)$ lies in $\langle A_i^\ell \rangle \cap \langle A_i^r \rangle = X^\ell \cap X^r$. As \mathcal{L} is an identity map on X^ℓ , we have that $\mathcal{L}^2(x_\ell) = \mathcal{L}(x_\ell)$, so $x_\ell - \mathcal{L}(x_\ell) \in \ker(\mathcal{L})$. By symmetry, also $x_r - \mathcal{R}(x_r) \in \ker(\mathcal{R})$. Hence $x = (x_\ell - \mathcal{L}(x_\ell)) + (x_r - \mathcal{R}(x_r)) \in \ker(\mathcal{L}) + \ker(\mathcal{R})$. This completes the proof of (5.8.4).

Hence, as \mathcal{T} is the identity map on $X^\ell + X^r$, we obtain $(A_i, X) \stackrel{\mathcal{T}}{\leftarrow} (A_j, X)$, contradicting (5.8.3). So Theorem 5.8 follows. ■

6. GRAPHS REVISITED

It is tempting to consider Theorem 1.1 as an extension of Theorem 1.3. However, so far in this paper this has not been justified and, in fact, as mentioned in the Introduction it is not that obvious, either. However, below we shall see that with some extra work it is possible to derive Theorem 1.3 from Theorem 1.1. Unfortunately, this derivation is as long as the direct proof of Theorem 1.3 presented in Section 4.

An edge is called an *apex* if both of its end vertices are adjacent to all other vertices.

(6.1) LEMMA. *Let G be a graph with an apex edge g and H be a simple minor of G with an apex edge h . Then G contains a minor H' isomorphic to H such that g is an apex of H' .*

Proof. As g is an apex in G , it is not that hard to see that each simple minor of G/g is isomorphic to a minor of G/f , where f is any edge adjacent to g . Moreover, contraction in G of any edge other than g keeps the property that g is an apex. So by induction on the size of G , we may assume that no contraction in G preserves the property that H is isomorphic to a minor. Hence, by isomorphism, we may assume that H is a subgraph of G and has the same set of vertices as G . This means that h is an apex in G as well. As switching the names of apices h and g is an automorphism of G , Lemma 6.1 follows. ■

The following facts are easy to prove: adding two new vertices to a graph increases the branch-width by at most 2; the branch-width of a cycle

matroid of a graph is at most the branch-width of the graph; and graphs with at least one apex edge are, up to isomorphism, uniquely determined by their cycle matroids. Combining these facts with Lemma 6.1 and Theorem 1.1, we obtain the following special case of Theorem 1.3.

(6.2) *Any class of simple graphs with bounded branch-width is well-quasi-ordered under taking minors.*

If G is a graph, then \tilde{G} denotes the graph obtained from G by subdividing each edge (and loop) once. Two edges in \tilde{G} are *mates* if they come from subdividing the same edge of G .

(6.3) LEMMA. *If \tilde{H} is a minor of \tilde{G} and no component of H is a path, then H is isomorphic to a minor of G .*

Proof. Let G and H form a counterexample where G has a minimum number of edges. This means that if e and e' are mates in \tilde{G} , then \tilde{H} is not a minor of $\tilde{G}/e, e'$. Hence, in going from \tilde{G} to \tilde{H} we never contract an entire series class. This means that \tilde{G} contains a subdivision H' of \tilde{H} as a subgraph. This containment is proper, because otherwise G and H would be topologically the same, and thus, as \tilde{H} is a minor of \tilde{G} , H would be a minor of G . Let e be an edge of \tilde{G} not in H' .

Let e' be the mate of e and let v be the vertex of \tilde{G} incident to both e and e' . As G is a minimal counterexample, \tilde{H} is not a minor of $\tilde{G} \setminus e, e'$, $v = \tilde{G} \setminus e/e'$; hence, as \tilde{H} has no isolated vertices, \tilde{H} is also not a minor of $\tilde{G} \setminus e, e'$. Therefore e' is an edge of H' . Let P' be the shortest path in H' from v to a vertex w with degree at least 3 in H' . (P' exists as no component of \tilde{H} is a path.) As v is a vertex of \tilde{G} that subdivides one of the edges in G and w is not, the path P' has an odd number of edges. Let P be the part of \tilde{H} such that P' is the subdivision of P that appears in obtaining H' as a subdivision of \tilde{H} . The end vertices of P have degree different from 2 in \tilde{H} , hence P has an even number of edges. So P' is longer than P and thus H'/e' is also a subdivision of \tilde{H} , this in spite of the aforementioned fact that \tilde{H} is not a minor of $\tilde{G} \setminus e/e'$. ■

Subdividing all edges and loops in a graph yields a simple graph and does not increase the branch-width. Hence (6.2) and Lemma 6.3 imply the following result:

(6.4) *Any class of graphs with no component equal to a path and with bounded branch-width is well-quasi-ordered under taking minors.*

As each graph is the vertex disjoint union of a simple graph and a graph in which no component is a path, (6.2) and (6.4) imply Theorem 1.3.

7. SPIKES AND INFINITE ANTICHAINS OF MATROIDS

In this section we prove (1.2), which says that for each infinite field there exist infinite antichains of matroids of bounded branch-width that are representable over that field.

An n -spike, or just *spike*, is a matroid whose ground set can be partitioned into n two-element sets (the *legs* of the spike) such that each two different legs form a circuit as well as a cocircuit. To avoid pathological cases we only consider spikes with $n \geq 5$. The following facts are straightforward consequences of elementary matroid axioms: an n -spike has rank and corank equal to n ; each circuit of an n -spike is either the union of two legs or has n or $n+1$ elements; each n -element circuit of an n -spike is a circuit-hyperplane and shares exactly one element with each leg.

As each rank- n matroid is determined by its circuits of size at most n , a spike is determined by its legs and its circuit hyperplanes. Conversely, a collection \mathcal{P} of n disjoint 2-element sets and an n -uniform hypergraph \mathcal{C} comprise the legs and circuit hyperplanes of a spike if and only if each member of \mathcal{C} shares exactly one element with each member of \mathcal{P} and at most $n-2$ elements with each other member of \mathcal{C} . Among other things, this shows that a spike is isomorphic to its dual; an isomorphism is the map that swaps the elements within each leg.

We now briefly discuss the relevant properties of spikes: their connectivity, their spike-minors, and their linear representations.

(7.1) LEMMA. *Spikes have branch-width 3.*

Proof. To see this observe that if $0 < k < n$ and A is a union of k legs of an n -spike M , then $\lambda_M(A) = (k+1) + (n-k+1) - n + 1 = 3$. So, if T is any branch decomposition of M such that for each leg $\{a, b\}$ the two leaves a and b in T have a common neighbour in T , then T has width 3. (Note that the width of an edge adjacent to a leaf of any branch decomposition of any matroid is always at most 2.) ■

The collection of spikes contains infinite antichains. In order to see this we first consider how spikes can turn up as proper minors of spikes.

(7.2) LEMMA. *Let N and M be spikes such that N is a proper minor of M . Then N is a minor of $M/x \setminus y$ for some leg $\{x, y\}$ of M .*

Proof. By duality we may assume that, for some leg $\{a_1, a_2\}$ of M , spike N is a minor of M/a_1 but not of $M/a_1 \setminus a_2$. Let $\{b_1, b_2\}$ and $\{c_1, c_2\}$ be two other legs of M . Then $\{b_1, b_2, a_2\}$ and $\{c_1, c_2, a_2\}$ are triangles of

M/a_1 . As spikes have no triangles and N is not a minor of $M/a_1 \setminus a_2$, we may assume, by symmetry, that N is a minor of $M/a_1 \setminus b_1, c_1$. In the latter matroid, the elements b_2 and c_2 are in series. So, again by symmetry, we may assume that N is a minor $M/a_1 \setminus b_1, c_1/b_2$, and hence of $M/b_2 \setminus b_1$. ■

By Lemma 7.2 it is quite easy to construct antichains of spikes. For each $n \geq 5$, let S_n be the n -spike with exactly two circuit hyperplanes, one the complement in the ground set of S_n of the other.

(7.3) $\{S_n \mid n \geq 5\}$ is an antichain under taking minors.

Proof. Let $m > n \geq 5$. Assume that S_n is a minor of S_m . Hence, by Lemma 7.2, $S_n = S_m \setminus \{a_1, \dots, a_k\} / \{b_1, \dots, b_k\}$ for some collection of legs $\{a_1, b_1\}, \dots, \{a_k, b_k\}$ of S_m . As $\{a_1, \dots, a_k\}$ intersects at least one of the two circuit hyperplanes of S_m , this means that S_n would have at most one circuit hyperplane, a contradiction. ■

Representations of spikes can be easily described.

(7.4) Let $n \geq 5$ and M be an \mathbb{F} -representable n -spike with legs $\{a_1, b_1\}, \dots, \{a_n, b_n\}$ such that $\{a_1, \dots, a_n\}$ is independent. Then each representation of M over \mathbb{F} is equivalent, under row operations and column scaling, to the columns in a matrix $[I, J + D]$ where I is the $n \times n$ identity matrix, J is the $n \times n$ matrix with all entries equal to 1, and D is a diagonal matrix with diagonal entries $1/\alpha_1, \dots, 1/\alpha_n$ (with $\alpha_1, \dots, \alpha_n$ nonzero, of course). Moreover, for each $X \subseteq \{1, \dots, n\}$, the set $\{a_i \mid i \notin X\} \cup \{b_i \mid i \in X\}$ is a circuit, and so a circuit hyperplane, if and only if $\sum_{i \in X} \alpha_i = -1$.

We skip the proof; it is straightforward.

Now consider a spike S_n as defined above. Let $\{a_1, b_1\}, \dots, \{a_n, b_n\}$ be its legs and $\{a_1, \dots, a_{n-1}, b_n\}$ and $\{b_1, \dots, b_{n-1}, a_n\}$ be its two circuit hyperplanes. By (7.4), S_n is representable over a field if that field contains a subset $\{\alpha_1, \dots, \alpha_{n-2}\}$ such that no nonempty subset of $\{1, \alpha_1, \dots, \alpha_{n-2}\}$ adds up to zero (take in (7.4) these α_i 's together with $\alpha_{n-1} := -1 - (\alpha_1 + \dots + \alpha_{n-2})$ and $\alpha_n := -1$). It is obvious that for each infinite field \mathbb{F} such a set $\{\alpha_1, \dots, \alpha_{n-2}\}$ exists. Hence $\{S_n \mid n \geq 5\}$ is an infinite antichain of matroids that are representable over each infinite field and have branch-width 3. So (1.2) follows.

ACKNOWLEDGMENT

We thank Neil Robertson and Paul Seymour for valuable discussions that provided the impetus for this research.

REFERENCES

1. R. Diestel, Two short proofs concerning tree-decompositions, preprint, 2001.
2. R. Diestel, K. Yu. Gorbunov, T. R. Jensen, and C. Thomassen, Highly connected sets and the excluded grid theorem, *J. Combin. Theory Ser. B* **75** (1999), 61–73.
3. J. Kruskal, Well-quasi-ordering, the tree theorem, and Vázsonyi's conjecture, *Trans. Amer. Math. Soc.* **95** (1960), 210–225.
4. C. St. J. A. Nash-Williams, On well-quasi-ordering finite trees, *Proc. Cambridge Philos. Soc.* **59** (1963), 833–835.
5. N. Robertson and P. D. Seymour, Graph minors. IV. Tree-width and well-quasi-ordering, *J. Combin. Theory Ser. B* **48** (1990), 227–254.
6. N. Robertson and P. D. Seymour, Graph minors. X. Obstructions to tree-decomposition, *J. Combin. Theory Ser. B* **52** (1991), 153–190.
7. R. Thomas, A Menger-like property of tree-width: The finite case, *J. Combin. Theory* **48** (1990), 67–76.
8. C. Thomassen, A simpler proof of the excluded minor theorem for higher surfaces, *J. Combin. Theory Ser. B* **70** (1997), 306–311.
9. W. T. Tutte, Menger's theorem for matroids, *J. Res. Natl. Bur. Standards, B. Math. Math. Phys.* **69** (1965), 49–53.