



City Research Online

City, University of London Institutional Repository

Citation: Daviaud, L., Johnson, M. and Kambites, M. (2018). Identities in upper triangular tropical matrix semigroups and the bicyclic monoid. *Journal of Algebra*, 501, pp. 503-525. doi: 10.1016/j.jalgebra.2017.12.032

This is the accepted version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/21291/>

Link to published version: <http://dx.doi.org/10.1016/j.jalgebra.2017.12.032>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

IDENTITIES IN UPPER TRIANGULAR TROPICAL MATRIX SEMIGROUPS AND THE BICYCLIC MONOID

LAURE DAVIAUD¹, MARIANNE JOHNSON² and MARK KAMBITES³

ABSTRACT. We establish necessary and sufficient conditions for a semigroup identity to hold in the monoid of $n \times n$ upper triangular tropical matrices, in terms of equivalence of certain tropical polynomials. This leads to an algorithm for checking whether such an identity holds, in time polynomial in the length of the identity and size of the alphabet. It also allows us to answer a question of Izhakian and Margolis, by showing that the identities which hold in the monoid of 2×2 upper triangular tropical matrices are exactly the same as those which hold in the bicyclic monoid. Our results extend to a broader class of “chain structured tropical matrix semigroups”; we exhibit a faithful representation of the free monogenic inverse semigroup within such a semigroup, which leads also to a representation by 3×3 upper triangular matrix semigroups, and a new proof of the fact that this semigroup satisfies the same identities as the bicyclic monoid.

1. INTRODUCTION

Over the last few years there has been considerable interest in the structure of the monoid $M_n(\mathbb{T})$ of tropical (max-plus) matrices under multiplication, with lines of investigation including characterisations of Green’s relations [10, 13, 21, 22], structural properties of subsemigroups and maximal subgroups [7, 9, 19], as well as connections between the algebraic properties of elements (or indeed subsemigroups) and their actions upon tropically convex sets [18, 24, 34].

A natural question, especially in view of d’Alessandro and Pasku’s proof that finitely generated subsemigroups of $M_n(\mathbb{T})$ have polynomial growth [7], is whether $M_n(\mathbb{T})$ satisfies a semigroup identity. Izhakian and Margolis [20] were the first to consider this question, in the 2×2 case; they showed that $M_2(\mathbb{T})$ does indeed satisfy an identity. A key step of their proof was establishing an identity for the upper triangular submonoid $UT_2(\mathbb{T})$, which is also of interest in its own right. Specifically, they showed that $UT_2(\mathbb{T})$ satisfies the celebrated identity $AB^2A^2BAB^2A = AB^2ABA^2B^2A$, which was shown by Adjan [1] to hold in the bicyclic monoid \mathcal{B} . Since \mathcal{B} embeds in $UT_2(\mathbb{T})$, every identity satisfied in the latter must also hold in the former. In view of this and their results, Izhakian and Margolis posed the natural question

¹ Faculty of Mathematics, Informatics and Mechanics, University of Warsaw, Poland. Email ldaviaud@mimuw.edu.pl. Work supported by the LIPA project, funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No 683080).

²School of Mathematics, University of Manchester, Manchester M13 9PL, UK. Email Marianne.Johnson@maths.manchester.ac.uk.

³School of Mathematics, University of Manchester, Manchester M13 9PL, UK. Email Mark.Kambites@manchester.ac.uk.

of whether the converse holds, that is, whether $UT_2(\mathbb{T})$ and \mathcal{B} satisfy exactly the same identities, or equivalently (by Birkhoff's HSP theorem [3]), whether they generate the same variety. Further evidence that this might be the case was recently provided by Chen, Hu, Luo and Sapir [6], who showed that the variety generated by $UT_2(\mathbb{T})$ shares with that generated by \mathcal{B} the property of not admitting a finite basis of identities.

In Section 3 we develop an exact characterisation of identities which hold in $UT_2(\mathbb{T})$, by associating k tropical polynomial equations in k variables to each identity on k letters. This gives a simple and algorithmically efficient (as we shall see in Section 8) method to check whether any given identity holds. In Section 4, by considering the embedding of \mathcal{B} into $UT_2(\mathbb{T})$, we use this to show that an identity which fails to hold in $UT_2(\mathbb{T})$ must also fail in \mathcal{B} , thus answering the above-mentioned question of Izhakian and Margolis [20] and establishing that the bicyclic monoid \mathcal{B} and the upper triangular tropical matrix monoid $UT_2(\mathbb{T})$ satisfy exactly the same semigroup identities (Theorem 4.1).

It follows by Birkhoff's HSP Theorem, of course, that \mathcal{B} and $UT_2(\mathbb{T})$ generate the same variety, and hence share all properties of monoids which are visible in the variety generated. For example, by a theorem of Shneerson [32], \mathcal{B} has infinite axiomatic rank, so $UT_2(\mathbb{T})$ must also have. (In particular, since infinite axiomatic rank implies the non-existence of a finite basis, we can deduce the results of [6] from our theorem plus [32].)

There is an interesting relationship between these results and work of Pastijn [26], which studies identities in the bicyclic monoid by reduction to checking linear inequalities or, equivalently, properties of certain polyhedral complexes. Once we have established that identities in \mathcal{B} are equivalent to identities in $UT_2(\mathbb{T})$, it becomes evident that our methods (for identities in $UT_2(\mathbb{T})$) are related to Pastijn's (for identities in \mathcal{B}). However, since our methods are used to establish the correlation in the first place, we cannot hope to deduce our results from Pastijn's work.

In higher dimensions, understanding identities seems to be hard. Izhakian and Margolis have conjectured that $M_n(\mathbb{T})$ satisfies an identity for every n ; evidence supporting this includes their own results on $M_2(\mathbb{T})$, and the proof by d'Alessandro and Pasku that finitely generated subsemigroups of $M_n(\mathbb{T})$ all have polynomial growth. More recently Shitov [30] has produced an identity for $M_3(\mathbb{T})$, over a 2-letter alphabet and with 1,795,308 letters on each side. Identities satisfied by semigroups of matrices satisfying non-singularity conditions were considered in [17, 30]. In [15] Izhakian claimed a family of identities for $UT_n(\mathbb{T})$ for each n ; in fact the proof contains a technical error and at least some of the claimed identities do not hold, but they can be corrected [16, 35] to give examples of identities which do hold in these semigroups. Further examples of identities for the upper triangular case have been constructed by Okninski [25] and by Cain *et al* [5]; the latter work provides an additional motivation for studying upper triangular tropical matrices by exhibiting an embedding of the *plactic monoid* of rank 3 into the direct product $UT_3(\mathbb{T}) \times UT_3(\mathbb{T})$.

In Section 5 we extend the ideas of Section 3 to characterise identities which hold in $UT_n(\mathbb{T})$, and indeed in a more general class of *chain structured*

tropical matrix semigroups (to be defined below), in terms of equivalence of certain tropical polynomial functions. Section 6 establishes some structural results about chain structured tropical matrix semigroups. Section 7 considers the free monogenic inverse semigroup \mathcal{I} ; we exhibit a natural embedding of \mathcal{I} into a chain structured tropical matrix semigroup which in turn embeds in $UT_3(\mathbb{T})$. By combining with the results of Section 5 we recover an alternative proof that \mathcal{I} satisfies the same identities as $UT_2(\mathbb{T})$ (which via Theorem 4.1, is equivalent to the known fact that it satisfies the same identities as the bicyclic monoid).

Finally, Section 8 applies the results of Sections 3 and 5 to the algorithmic problem of deciding whether a given identity holds in $UT_n(\mathbb{T})$. These results allow us to reduce the problem to (real) linear programming, and hence to show that, for a fixed n , it can be solved in time polynomial in the size of the identity and the alphabet over which it is defined. In particular, by the results of Section 4, this gives a polynomial-time algorithm to check whether a given identity holds in the bicyclic monoid. (The work of Pastijn [26] also yields an algorithm for this problem in the bicyclic monoid, which also amounts to a reduction to linear programming; although Pastijn does not analyse the complexity of his algorithm, we believe it is similar to ours.)

Semigroups of tropical matrices are closely related to *max-plus automata*, a computational model defined as a quantitative extension of finite automata. (*Weighted automata* more generally were introduced by Schützenberger in [29]; these machines perform an automatic computation on words over a finite alphabet, associating with each word a value or ‘weight’ in a given semiring.) Semigroups of tropical matrices and max-plus automata provide two perspectives on the same underlying mathematical structures; the questions raised about them by algebraists and computer scientists are often different, but the study of identities is natural from both points of view; from a computational perspective it addresses the question of which pairs of distinct inputs can be distinguished by a computational model. This common ground has fruitfully been explored in the case of finite (unweighted) automata, leading to the equational and topological development of profinite theory [27, 11]. In a companion article [8] the first two authors investigate the computational power of two-state max-plus automata, and use the ideas developed there to construct minimal length identities for $M_2(\mathbb{T})$. A full understanding of all identities in $M_n(\mathbb{T})$ for $n \geq 3$ seems still quite distant.

2. PRELIMINARIES

2.1. Tropical polynomials. The *tropical semiring* \mathbb{T} is the commutative idempotent semiring whose elements are drawn from the set $\mathbb{R} \cup \{-\infty\}$, and whose binary operations are defined by $a \oplus b = \max(a, b)$ and $a \otimes b = a + b$, where $-\infty$ should be thought of as the “zero” element of the semiring, satisfying $-\infty \oplus a = a$ and $-\infty \otimes a = -\infty$ for all $a \in \mathbb{T}$. Often it will be convenient to work without this zero element; in order to distinguish this algebraic structure from the field \mathbb{R} we shall write \mathbb{FT} to denote the “finitary tropical semiring”, namely the set \mathbb{R} , under the operations \oplus and \otimes . There is an obvious total order on \mathbb{T} in which $-\infty$ is the least element. We denote the supremum of a bounded-above set of elements $X \subset \mathbb{T}$ by $\bigoplus_{x \in X} x$.

By a *formal tropical polynomial* in variables from a (countable) set X we mean an element of the commutative polynomial semiring $\mathbb{T}[X]$, that is, a finite formal sum in which each term is a formal product of a coefficient from \mathbb{FT} and formal powers of finitely many of the variables of X , considered up to the commutative and distributive laws and the idempotent addition in \mathbb{T} . The order on \mathbb{T} induces a partial order on $\mathbb{T}[X]$, in which $g \preceq f$ if and only if each term of g appears as a term of f with coefficient greater than or equal to the corresponding coefficient in g . A set of polynomials is *bounded* if it is bounded above in this order; any bounded set of polynomials has a supremum which is a polynomial. We note that evaluation at any fixed values of the variables gives a supremum-preserving semiring morphism from $\mathbb{T}[X]$ to \mathbb{T} .

Each formal tropical polynomial naturally defines a function from \mathbb{T}^X to \mathbb{T} , by interpreting all formal products and formal sums tropically. Unlike with classical polynomials over a field, however, two distinct formal tropical polynomials may define the same function. For example, $x^{\otimes 2} \oplus x \oplus 1$ and $x^{\otimes 2} \oplus 1$ are distinct formal polynomials but define the same function because x can never exceed both $x^{\otimes 2}$ and 1. We say that two formal polynomials are *equivalent* if they represent the same function. We note that it is easy to see that two tropical polynomials $f, g : \mathbb{T}^X \rightarrow \mathbb{T}$ will be equivalent if and only if they agree on \mathbb{FT}^X .

We say that a tropical polynomial f is *0-flat* if it is a sum of products of variables (without coefficients). An elementary but important property of 0-flat polynomials is the following relationship with classical scaling:

Lemma 2.1. *If $f \in \mathbb{T}[x_1, \dots, x_k]$ is 0-flat and $\lambda \in \mathbb{R}$ is non-negative then $f(\lambda x_1, \dots, \lambda x_k) = \lambda f(x_1, \dots, x_k)$ for all evaluations of x_1, \dots, x_k in \mathbb{T} .*

Proof. This is immediate from the fact that scaling by a non-negative value distributes over both maximum and classical addition. \square

2.2. Tropical semigroups. It is easy to see that the set of all $n \times n$ matrices with entries in \mathbb{T} (respectively \mathbb{FT}) forms a monoid (respectively semigroup) under the matrix multiplication induced from the operations \oplus and \otimes . We denote these semigroups by $M_n(\mathbb{T})$ and $M_n(\mathbb{FT})$. We also write $UT_n(\mathbb{T})$ (respectively $UT_n(\mathbb{FT})$) to denote the subsemigroup of $M_n(\mathbb{T})$ consisting of those upper-triangular matrices in $M_n(\mathbb{T})$ whose entries on and above the main diagonal lie in \mathbb{T} (respectively \mathbb{FT}).

More generally, let (Γ, ρ) be a (not necessarily finite) set equipped with a reflexive, transitive binary relation. Consider the set $\mathbb{T}^{\Gamma \times \Gamma}$ of functions from $\Gamma \times \Gamma$ to \mathbb{T} ; we think of its elements as matrices with rows and columns indexed by Γ , and use the notation $M_{i,j}$ for $M(i, j)$. We define

$$\begin{aligned} \Gamma(\mathbb{T}) &= \{A \in \mathbb{T}^{\Gamma \times \Gamma} \mid A \text{ is bounded above and } A_{i,j} \neq -\infty \implies i\rho j\}, \\ \Gamma(\mathbb{FT}) &= \{A \in \mathbb{T}^{\Gamma \times \Gamma} \mid A \text{ is bounded above and } A_{i,j} \neq -\infty \iff i\rho j\}. \end{aligned}$$

We define a multiplication on $\Gamma(\mathbb{T})$ (and on $\Gamma(\mathbb{FT})$) as for matrices by:

$$(M \otimes N)_{i,j} = \bigoplus_{k \in \Gamma} M_{i,k} \otimes N_{k,j}.$$

The boundedness of M and N ensures the product is defined and bounded, while the transitivity of Γ ensures that appropriate entries of the product are $-\infty$. The operation is also easily seen to be associative, so it gives $\Gamma(\mathbb{T})$ and $\Gamma(\mathbb{FT})$ the structure of semigroups.

Notice that taking $\Gamma = \{1, \dots, n\}$ with the complete binary relation [respectively, the obvious partial order] we have $\Gamma(\mathbb{T}) = M_n(\mathbb{T})$ and $\Gamma(\mathbb{FT}) = M_n(\mathbb{FT})$ [respectively, $\Gamma(\mathbb{T}) = UT_n(\mathbb{T})$ and $\Gamma(\mathbb{FT}) = UT_n(\mathbb{FT})$]. More generally, if Γ is finite then fixing a bijection with a set $\{1, \dots, n\}$ yields an embedding of $\Gamma(\mathbb{T})$ into $M_n(\mathbb{T})$. If Γ is a finite partial order then an appropriate choice of bijection embeds $\Gamma(\mathbb{T})$ in some $UT_n(\mathbb{T})$.

2.3. Semigroup identities. We write \mathbb{N}_0 and \mathbb{N} respectively for the natural numbers with and without 0. If Σ is a finite alphabet, then Σ^+ will denote the free semigroup on Σ , that is, the set of finite, non-empty words over Σ under the operation of concatenation. For $w \in \Sigma^+$ and $s \in \Sigma$ we write $|w|$ for the length of w and $|w|_s$ for the number of occurrences of the letter s in w . For $1 \leq i \leq |w|$ we write w_i to denote the i th letter of w . The *content* of w is the map $\Sigma \rightarrow \mathbb{N}_0, s \mapsto |w|_s$.

Recall that a (*semigroup*) *identity* is a pair of words, usually written “ $u = v$ ”, in the free semigroup Σ^+ on an alphabet Σ . We say that the identity *holds* in a semigroup S (or that S *satisfies* the identity) if every morphism from Σ^+ to S maps u and v to the same element of S . If a morphism maps u and v to the same element we say that it *satisfies* the given identity in S ; otherwise it *falsifies* it. (*Monoid* identities are defined similarly using the free monoid Σ^* in place of Σ^+ ; we shall work primarily with semigroup identities, but it is easy to deduce corresponding results about monoid identities from these.)

For some purposes it is easier to work with matrices over \mathbb{T} , while for others it is easier to work with matrices over \mathbb{FT} . The following proposition ensures that we can choose to work with whichever is easier at each stage.

Proposition 2.2. *Let (Γ, ρ) be a set equipped with a reflexive, transitive binary relation. Then $\Gamma(\mathbb{T})$ and $\Gamma(\mathbb{FT})$ satisfy exactly the same semigroup identities.*

Proof. Since $\Gamma(\mathbb{FT})$ is a subsemigroup of $\Gamma(\mathbb{T})$, any identity satisfied in the latter is also satisfied in the former. To prove the converse we shall work with a semigroup $\Gamma(\mathbb{T}[x])$, which is constructed by exactly the same process as $\Gamma(\mathbb{T})$ but using the formal polynomial semiring $\mathbb{T}[x]$ in place of \mathbb{T} ; its elements are thus bounded-above functions $\Gamma \times \Gamma \rightarrow \mathbb{T}[x]$, with respect to the order \preceq defined on $\mathbb{T}[x]$. For any evaluation of x in \mathbb{T} , we obtain a supremum-preserving semiring morphism from $\mathbb{T}[x]$ to \mathbb{T} , which in turn it induces a semigroup morphism from $\Gamma(\mathbb{T}[x])$ to $\Gamma(\mathbb{T})$.

Returning to the proof of the converse, suppose an identity $u = v$ over alphabet Σ is *not* satisfied in $\Gamma(\mathbb{T})$. Let $\phi : \Sigma^+ \rightarrow \Gamma(\mathbb{T})$ be a morphism falsifying the identity, and choose i and j with $\phi(u)_{i,j} \neq \phi(v)_{i,j}$.

Define a morphism $\psi : \Sigma^+ \rightarrow \Gamma(\mathbb{T}[x])$ by for each $a \in \Sigma$ setting

$$\psi(a)_{k,l} = \begin{cases} \phi(a)_{k,l} & \text{if } \phi(a)_{k,l} \neq -\infty \\ x & \text{if } k\rho l \text{ but } \phi(a)_{k,l} = -\infty \\ -\infty & \text{otherwise.} \end{cases}$$

For any $x \in \mathbb{T}$, because evaluation at x induces a semigroup morphism from $\Gamma(\mathbb{T}[x])$ to $\Gamma(\mathbb{T})$, we may define a semigroup morphism $\psi_x : \Sigma^+ \rightarrow \Gamma(\mathbb{T})$ by $\psi_x(w)_{i,j} = [\psi(w)_{i,j}](x)$. If $x \in \mathbb{FT}$ then the image of ψ_x is contained in $\Gamma(\mathbb{FT})$, while if $x = -\infty$ we recover the original morphism ϕ .

Consider now the polynomials $\psi(u)_{i,j}$ and $\psi(v)_{i,j}$. Since at $x = -\infty$ they evaluate to $\phi(u)_{i,j}$ and $\phi(v)_{i,j}$ respectively, which are different, they must have different constant terms. It follows that for x small enough but finite, they evaluate to different values. Thus, for small enough x , ψ_x is a morphism to $\Gamma(\mathbb{FT})$ falsifying the identity $u = v$. \square

Proposition 2.2 has the following immediate consequence:

Corollary 2.3. *For all $n \in \mathbb{N}$,*

- (i) $M_n(\mathbb{T})$ and $M_n(\mathbb{FT})$ satisfy exactly the same semigroup identities;
and
- (ii) $UT_n(\mathbb{T})$ and $UT_n(\mathbb{FT})$ satisfy exactly the same semigroup identities.

The following two elementary results will allow us at various points to simplify the structure of and number of parameters in the matrices we must consider.

Lemma 2.4. *Let $\phi : \Sigma^+ \rightarrow M_n(\mathbb{T})$ be a morphism, $\mu_s \in \mathbb{FT}$ for each $s \in \Sigma$, and define a new morphism $\psi : \Sigma^+ \rightarrow M_n(\mathbb{T})$ by $\psi(s) = \mu_s \otimes \phi(s)$ for all $s \in \Sigma$. If u and v are words with the same content and $\phi(u) = \phi(v)$, then $\psi(u) = \psi(v)$.*

Proof. Since tropical scaling commutes with matrix multiplication it is easy to see that

$$\psi(u) = \left(\bigotimes_{s \in \Sigma} \mu_s^{|u|_s} \right) \otimes \phi(u) = \left(\bigotimes_{s \in \Sigma} \mu_s^{|v|_s} \right) \otimes \phi(v) = \psi(v),$$

where the powers are to be interpreted tropically, and the middle equality holds because $\phi(u) = \phi(v)$ and u and v have the same content. \square

Lemma 2.5. *Suppose an identity $u = v$ over alphabet Σ is satisfied by all morphisms $\phi : \Sigma^+ \rightarrow UT_n(\mathbb{FT})$ with the property that $\phi(s)_{n,n} = 0$ for all $s \in \Sigma$. If u and v have the same content then the identity $u = v$ holds in $UT_n(\mathbb{T})$.*

Proof. Every morphism to $UT_n(\mathbb{FT})$ can clearly be obtained from one of the given form by the construction in Lemma 2.4, which since u and v have the same content means it satisfies the identity. So $UT_n(\mathbb{FT})$ and hence also (by Corollary 2.3) $UT_n(\mathbb{T})$ satisfy the identity. \square

3. IDENTITIES IN $UT_2(\mathbb{T})$.

In this section we shall give an elementary but powerful characterisation of the identities which hold in $UT_2(\mathbb{T})$ (or equivalently, by Corollary 2.3, in $UT_2(\mathbb{FT})$).

Let w be a word over an alphabet Σ . For $s \in \Sigma$ and $0 \leq i \leq |w|$ we write $\lambda_s^w(i)$ for the number of occurrences of the letter s within the first i letters of the word w (so in particular $\lambda_s^w(0) = 0$ for all $s \in \Sigma$). For each $t \in \Sigma$ we define a 0-flat formal tropical polynomial having variables x_s for each $s \in \Sigma$ as follows:

$$f_t^w = \bigoplus_{w_i=t} \bigotimes_{s \in \Sigma} x_s^{\lambda_s^w(i-1)},$$

where of course the powers are to be interpreted tropically, and a maximum over the empty set is taken to be $-\infty$. By a slight abuse of notation, given $\underline{x} \in \mathbb{FT}^\Sigma$ we shall also write x_s to denote the image of s under \underline{x} , and $f_t^w(\underline{x})$ to denote the corresponding evaluation of f_t^w at the variables $x_s \in \mathbb{FT}$.

We shall see shortly (Theorem 3.3) that the values of the polynomials f_t^w exactly characterise the identities which hold in $UT_2(\mathbb{T})$. The following lemma explains how they arise from the multiplication in $UT_2(\mathbb{FT})$:

Lemma 3.1. *Suppose a morphism $\phi : \Sigma^+ \rightarrow UT_2(\mathbb{T})$ is such that $\phi(s)_{1,1} \in \mathbb{FT}$ and $\phi(s)_{2,2} = 0$ for all $s \in \Sigma$, say*

$$\phi(s) = \begin{pmatrix} x_s & x'_s \\ -\infty & 0 \end{pmatrix}, \text{ where } x_s, x'_s \in \mathbb{T} \text{ with } x_s \neq -\infty.$$

Then for any word $w \in \Sigma^+$,

$$\phi(w)_{1,2} = \bigoplus_{s \in \Sigma} x'_s \otimes f_s^w(\underline{x})$$

Proof. This is a straightforward calculation, using the definitions of matrix multiplication and of the polynomial functions f_s^w . \square

The next lemma says that the polynomial functions f_t^w are together sufficient to characterise the content of the word w .

Lemma 3.2. *If $f_t^w(\underline{x}) = f_t^v(\underline{x})$ for all $t \in \Sigma$ and $\underline{x} \in \mathbb{FT}^\Sigma$ then w and v have the same content.*

Proof. It is easy to see that f_t^w is the constant $-\infty$ function if and only if w does not contain any occurrences of the letter t . Thus, we may assume that exactly the same letters from Σ occur in w and v .

Let $z \in \Sigma$, and suppose that z occurs in both w and v . Define $\underline{x} \in \mathbb{FT}^\Sigma$ by $x_z = 1$ and $x_s = 0$ for $s \in \Sigma \setminus \{z\}$. Then

$$f_z^w(\underline{x}) = \bigoplus_{w_i=z} \bigotimes_{s \in \Sigma} x_s^{\lambda_s^w(i-1)} = \bigoplus_{w_i=z} \lambda_z^w(i-1)$$

where the last equality is because of the values of \underline{x} . Since for fixed z and w the function $\lambda_z^w(i)$ is non-decreasing with i , it is clear that the maximum on the right is attained when i is the position of the final z in w . With this value of i , we see that $f_z^w(\underline{x}) = \lambda_z^w(i-1) = |w|_z - 1$. By the same argument, $f_z^v(\underline{x}) = |v|_z - 1$. But by assumption $f_z^w(\underline{x}) = f_z^v(\underline{x})$, so we must have $|w|_z = |v|_z$. \square

We are now ready to prove the main theorem of this section.

Theorem 3.3. *The identity $w = v$ over alphabet Σ is satisfied in $UT_2(\mathbb{T})$ if and only if the tropical polynomials f_t^w and f_t^v are equivalent for all $t \in \Sigma$.*

Proof. Suppose first that we can choose $\underline{x} \in \mathbb{FT}^\Sigma$ and $t \in \Sigma$, such that $f_t^w(\underline{x}) \neq f_t^v(\underline{x})$. Define a morphism

$$\phi : \Sigma^+ \rightarrow UT_2(\mathbb{T}), \quad a \mapsto \begin{pmatrix} x_a & x'_a \\ -\infty & 0 \end{pmatrix}$$

where $x'_t = 0$ and $x'_s = -\infty$ for each $s \neq t$. Then by Lemma 3.1,

$$\phi(w)_{1,2} = \bigoplus_{s \in \Sigma} x'_s \otimes f_s^w(\underline{x}) = f_t^w(\underline{x})$$

and similarly

$$\phi(v)_{1,2} = \bigoplus_{s \in \Sigma} x'_s \otimes f_s^v(\underline{x}) = f_t^v(\underline{x}) \neq \phi(w)_{1,2}$$

so the morphism ϕ falsifies the identity in $UT_2(\mathbb{T})$.

Conversely, suppose that $f_t^w(\underline{x}) = f_t^v(\underline{x})$ for all $t \in \Sigma$ and all $\underline{x} \in \mathbb{FT}^\Sigma$. By Lemma 3.2, w and v have the same content. Hence, by Lemma 2.5 it suffices to show that the identity $w = v$ is satisfied by morphisms $\phi : \Sigma^+ \rightarrow UT_2(\mathbb{FT})$ such that $\phi(s)_{2,2} = 0$ for all $s \in S$. Let ϕ be such a morphism, and for each $s \in \Sigma$ write $x_s = \phi(s)_{1,1}$ and $x'_s = \phi(s)_{1,2}$.

It is immediate from the properties of ϕ that $\phi(w)_{2,2} = 0 = \phi(v)_{2,2}$. It is also easy to see that for any word u ,

$$\phi(u)_{1,1} = \bigotimes_{i=1}^{|u|} \phi(u_i)_{1,1} = \sum_{s \in \Sigma} |u|_s \phi(s)_{1,1}.$$

Since w and v have the same content, it follows that $\phi(w)_{1,1} = \phi(v)_{1,1}$. Finally, Lemma 3.1 gives

$$\phi(w)_{1,2} = \bigoplus_{t \in \Sigma} x'_t \otimes f_t^w(\underline{x}) = \bigoplus_{t \in \Sigma} x'_t \otimes f_t^v(\underline{x}) = \phi(v)_{1,2}$$

so that $\phi(w) = \phi(v)$ as required. \square

We conclude this section by noting that checking whether a pair of words on an alphabet containing just two letters defines an identity on $UT_2(\mathbb{T})$ amounts to a comparison of four pairs of tropical polynomials in just one variable.

Corollary 3.4. *Let Σ be a two-letter alphabet and $w, v \in \Sigma^+$. Then the identity $w = v$ holds in $UT_2(\mathbb{T})$ if and only if $f_z^w(x, 1) = f_z^v(x, 1)$ and $f_z^w(x, -1) = f_z^v(x, -1)$ for all $x \in \mathbb{FT}$, and all $z \in \Sigma$.*

Proof. The direct implication follows immediately from Theorem 3.3. Conversely, suppose that $f_z^w(x, 1) = f_z^v(x, 1)$ and $f_z^w(x, -1) = f_z^v(x, -1)$ for all $z \in \Sigma$ and all $x \in \mathbb{FT}$. Since the polynomials f_z^w and f_z^v are 0-flat, for any $c, d \in \mathbb{R}$ with $d \neq 0$ applying Lemma 2.1 we have

$$f_z^w(c, d) = |d| f_z^w\left(\frac{c}{|d|}, \frac{d}{|d|}\right) = |d| f_z^v\left(\frac{c}{|d|}, \frac{d}{|d|}\right) = f_z^v(c, d)$$

where $|d|$ denotes the absolute value of d so that $d/|d|$ is 1 or -1 . Since the functions of the form $f_z^w(c, -), f_z^v(c, -) : \mathbb{R} \rightarrow \mathbb{R}$ are continuous it follows also that $f_z^w(c, 0) = f_z^v(c, 0)$ for all $c \in \mathbb{R}$, so by Theorem 3.3 the identity $w = v$ holds in $UT_2(\mathbb{T})$. \square

4. THE BICYCLIC MONOID

The *bicyclic monoid* $\mathcal{B} = \langle p, q \mid pq = 1 \rangle$ is an inverse monoid which is ubiquitous in almost all areas of infinite semigroup theory, and indeed in many other areas of mathematics. Identities in the bicyclic monoid have been extensively studied; Adjan [1] established that the identity $AB^2A^2BAB^2A = AB^2ABA^2B^2A$ is a minimal length identity for this monoid. Shleifer [31] has shown (with computer assistance) that there is just one other identity of minimal length. Jones [23] has shown that a semigroup variety contains a simple semigroup that is not bisimple if and only if it contains the bicyclic semigroup. Motivated by this work, Pastijn [26] gave an elegant geometric characterisation of the identities which hold in the bicyclic monoid in terms of the properties of associated polyhedral complexes. Most recently Shneerson [33] has provided a family of interesting examples of semigroups of cubic growth which satisfy the same identities as the bicyclic monoid.

In this section we apply Theorem 3.3 to prove that the identities satisfied by $UT_2(\mathbb{T})$ are precisely those satisfied by the bicyclic monoid, thus establishing that $UT_2(\mathbb{T})$ and several other associated semigroups all generate the same variety as \mathcal{B} .

Theorem 4.1. *The bicyclic monoid \mathcal{B} and the upper triangular tropical matrix monoid $UT_2(\mathbb{T})$ satisfy exactly the same semigroup identities.*

Proof. It is well known and easy to see that every element of the bicyclic monoid can be written uniquely in the form $q^i p^j$ for some $i, j \in \mathbb{N}_0$. Izhakian and Margolis [20] noted that the map:

$$\rho : \mathcal{B} \rightarrow UT_2(\mathbb{T}), \quad q^i p^j \mapsto \begin{pmatrix} i - j & i + j \\ -\infty & j - i \end{pmatrix}$$

is a semigroup embedding of \mathcal{B} into $UT_2(\mathbb{T})$; of course, it also gives an embedding into $UT_2(\mathbb{FT})$.

It is immediate from the existence of this embedding that any identity which holds in $UT_2(\mathbb{T})$ must hold in \mathcal{B} . It remains to show that any identity which holds in \mathcal{B} must also hold in $UT_2(\mathbb{T})$. We shall show the contrapositive. To this end, suppose that a semigroup identity $w = v$ over finite alphabet Σ does not hold in $UT_2(\mathbb{T})$. Clearly we may assume that w and v have the same content; indeed if $z \in \Sigma$ is such that $|w|_z \neq |v|_z$ then the identity is falsified in \mathcal{B} by (for example) the morphism sending z to p and every other letter to the identity element.

By Theorem 3.3 there exists $\underline{x} \in \mathbb{FT}^\Sigma$ and $t \in \Sigma$ such that $f_t^w(\underline{x}) \neq f_t^v(\underline{x})$. Suppose without loss of generality that $f_t^w(\underline{x}) > f_t^v(\underline{x})$. Since both functions are piecewise linear in the components of \underline{x} , there is an open neighbourhood (in the usual Euclidean topology on \mathbb{FT}^Σ considered as a set of finite dimensional vectors over \mathbb{R}) of \underline{x} in which every vector satisfies the given inequality. This neighbourhood must contain a vector in \mathbb{Q}^Σ so by

replacing \underline{x} with such a vector, we may assume without loss of generality that the components of \underline{x} are all rational.

Let d denote the least common multiple of the denominators in the lowest forms of the entries x_s for $s \in \Sigma$. Since f_t^w and f_t^v are 0-flat, it follows from Lemma 2.1 that $f_t^w(2d\underline{x}) > f_t^v(2d\underline{x})$. Hence, by replacing \underline{x} with $2d\underline{x}$, we may further assume, again without loss of generality, that the components of \underline{x} are even integers.

Now for each $s \in \Sigma$, choose a non-negative even integer x'_s greater than x_s , in such a way that x'_t is very large relative to all other x'_s . Define a morphism $\psi : \Sigma^+ \rightarrow UT_2(\mathbb{F}\mathbb{T})$ by

$$\psi(s) = \begin{pmatrix} x_s & x'_s \\ -\infty & 0 \end{pmatrix} \text{ for all } s \in \Sigma.$$

By Lemma 3.1 we have

$$\psi(w)_{1,2} = \bigoplus_{s \in \Sigma} x'_s \otimes f_s^w(\underline{x}) = x'_t \otimes f_t^w(\underline{x}),$$

provided x'_t was chosen sufficiently large, and similarly

$$\psi(v)_{1,2} = \bigoplus_{s \in \Sigma} x'_s \otimes f_s^v(\underline{x}) = x'_t \otimes f_t^v(\underline{x}) \neq \psi(w)_{1,2},$$

where the inequality is because $f_t^w(\underline{x}) \neq f_t^v(\underline{x})$. So the morphism ψ falsifies the identity in $UT_2(\mathbb{F}\mathbb{T})$.

Next for each $s \in \Sigma$ let $i_s = \frac{1}{2}x'_s$, and $j_s = \frac{1}{2}(x'_s - x_s)$. Then i_s is non-negative (because x'_s is non-negative) and integer (because x'_s is even), while j_s is non-negative (because $x'_s > x_s$) and integer (because x'_s and x_s are even). Define a morphism $\phi : \Sigma^+ \rightarrow \text{im } \rho \subseteq UT_2(\mathbb{T})$ by

$$\phi(s) = \rho(q^{i_s} p^{j_s}) = \begin{pmatrix} i_s - j_s & i_s + j_s \\ -\infty & j_s - i_s \end{pmatrix} = (j_s - i_s) \otimes \begin{pmatrix} x_s & x'_s \\ -\infty & 0 \end{pmatrix}$$

where the last equality follows from the definitions of i_s and j_s . Because of the last equality, we see that ψ can be obtained from ϕ by the construction in Lemma 2.4, with $\mu_s = i_s - j_s$ for each s . Hence, by the contrapositive of Lemma 2.4, ϕ falsifies the identity.

But $\text{im } \phi \subseteq \text{im } \rho$ which is isomorphic to \mathcal{B} , so the identity cannot hold in \mathcal{B} . \square

Corollary 4.2. *The subsemigroups of $UT_2(\mathbb{T})$ obtained by restricting the on-and-above diagonal entries to be respectively in \mathbb{Q} , \mathbb{Z} , \mathbb{N}_0 , $\mathbb{Q} \cup \{-\infty\}$, $\mathbb{Z} \cup \{-\infty\}$ and $\mathbb{N}_0 \cup \{-\infty\}$ all satisfy exactly the same identities as \mathcal{B} and $UT_2(\mathbb{T})$.*

Proof. Write $UT_2(X)$ for the given semigroups, where $X = \mathbb{Q}$, \mathbb{Z} , \mathbb{N}_0 , $\mathbb{Q} \cup \{-\infty\}$, $\mathbb{Z} \cup \{-\infty\}$ or $\mathbb{N}_0 \cup \{-\infty\}$. The semigroups $UT_2(\mathbb{Q})$, $UT_2(\mathbb{Z})$, $UT_2(\mathbb{Q} \cup \{-\infty\})$ and $UT_2(\mathbb{Z} \cup \{-\infty\})$ are all contained in $UT_2(\mathbb{T})$ and contain the image of the embedding ρ , which is isomorphic to \mathcal{B} .

The semigroups $UT_2(\mathbb{N}_0)$ and $UT_2(\mathbb{N}_0 \cup \{-\infty\})$ are again contained in $UT_2(\mathbb{T})$, and so satisfy all of the given identities. Conversely, notice that every matrix in $UT_2(\mathbb{Z})$ is a tropical scaling of a matrix in $UT_2(\mathbb{N}_0)$. By the previous paragraph, any identity which does not hold in $UT_2(\mathbb{T})$ is falsified by some morphism to $UT_2(\mathbb{Z})$, and now the contrapositive of Lemma 2.4

yields a morphism falsifying the identity with image contained in $UT_2(\mathbb{N}_0)$, and hence also in $UT_2(\mathbb{N}_0 \cup \{-\infty\})$. \square

The bicyclic monoid is well known (see for example [14, Section 1.6]) to be isomorphic to $\mathbb{N}_0 \times \mathbb{N}_0$ under the multiplication

$$(a, b)(c, d) = (a - b + \max(b, c), d - c + \max(b, c)).$$

Several natural extensions of this construction have been considered (see for example [12, 36, 35]). We write $\mathcal{B}_{\mathbb{Z}} := \mathbb{Z} \times \mathbb{Z}$, $\mathcal{B}_{\mathbb{Q}} := \mathbb{Q} \times \mathbb{Q}$ and $\mathcal{B}_{\mathbb{R}} = \mathbb{R} \times \mathbb{R}$ to denote the semigroups with completely analogous multiplication, noting the obvious embeddings $\mathcal{B} \subseteq \mathcal{B}_{\mathbb{Z}} \subseteq \mathcal{B}_{\mathbb{Q}} \subseteq \mathcal{B}_{\mathbb{R}}$.

Corollary 4.3. *The semigroups $\mathcal{B}_{\mathbb{Z}}$, $\mathcal{B}_{\mathbb{Q}}$ and $\mathcal{B}_{\mathbb{R}}$ satisfy exactly the same semigroup identities as \mathcal{B} and $UT_2(\mathbb{T})$.*

Proof. It is straightforward to check that the map

$$\phi : \mathcal{B}_{\mathbb{R}} \rightarrow UT_2(\mathbb{T}), (a, b) \mapsto \begin{pmatrix} a - b & a + b \\ -\infty & b - a \end{pmatrix}$$

is an embedding. Thus we have embeddings

$$\mathcal{B} \hookrightarrow \mathcal{B}_{\mathbb{Z}} \hookrightarrow \mathcal{B}_{\mathbb{Q}} \hookrightarrow \mathcal{B}_{\mathbb{R}} \hookrightarrow UT_2(\mathbb{T}).$$

Since \mathcal{B} and $UT_2(\mathbb{T})$ satisfy the same identities, the result follows. \square

5. $UT_n(\mathbb{T})$ AND CHAIN-STRUCTURED TROPICAL MATRIX SEMIGROUPS

In this section we extend the results of Section 3 to a more general class of semigroups $\Gamma(\mathbb{T})$ and $\Gamma(\mathbb{FT})$ (including $UT_n(\mathbb{T})$ and $UT_n(\mathbb{FT})$ for $n > 2$). We say that $\Gamma(\mathbb{T})$ and $\Gamma(\mathbb{FT})$ are *chain-structured tropical matrix semigroups* if Γ is a partial order with an upper bound on the length of chains. For the rest of this section, we assume that Γ has these properties, and let n be the length of the longest chain in Γ . (Thus there exists a sequence $v_1 \leq v_2 \leq \dots \leq v_k$ of k distinct elements of Γ if and only if $k \leq n$.)

By a *k-vertex walk* (or *walk of vertex length k*) in Γ we mean a k -tuple (v_1, \dots, v_k) such that $v_1 \leq v_2 \leq \dots \leq v_k$. A *k-vertex path* (or *path of vertex length k*) is a k -vertex walk in which consecutive vertices (and hence all vertices) are distinct.

Let w be a word over the alphabet Σ . For $0 \leq p < q \leq |w| + 1$ and $s \in \Sigma$ we define

$$\beta_s^w(p, q) = |\{i \in \mathbb{N} \mid p < i < q, w_i = s\}|$$

to be the number of occurrences of s lying strictly *between* w_p and w_q . For each $u \in \Sigma^*$ with $|u| \leq n - 1$ and each $(|u| + 1)$ -vertex path $\rho = (\rho_0, \rho_1, \dots, \rho_{|u|})$ in Γ , we define a (0-flat) formal tropical polynomial having variables $x(s, i)$ for each letter $s \in \Sigma$ and vertex $i \in \Gamma$ as follows:

$$f_{u, \rho}^w = \bigoplus_{s \in \Sigma} \bigotimes_{k=0}^{|u|} x(s, \rho_k)^{\beta_s^w(\alpha_k, \alpha_{k+1})},$$

where the sum ranges over all $0 = \alpha_0 < \alpha_1 < \dots < \alpha_{|u|} < \alpha_{|u|+1} = |w| + 1$ such that $w_{\alpha_k} = u_k$ for $k = 1, \dots, |u|$. Here the powers are to be interpreted tropically, and a maximum taken over the empty set is taken to be $-\infty$. Thus it is easy to see that $f_{u, \rho}^w \neq -\infty$ if and only if u is a scattered subword

of w of length equal to the path ρ . Note that taking u to be the empty word forces $\rho = (\rho_0)$ for some $\rho_0 \in \Gamma$ and hence $f_{u,\rho}^w = \bigotimes_{s \in \Sigma} x(s, \rho_0)^{|w|_s}$, which is completely determined by the content of w .

In general, it is clear that the number of choices for the α_i 's is bounded above by $|w|^{|u|}$, which in turn is bounded above by $|w|^{n-1}$. Thus, fixing Γ (and hence n), it is easy to verify that the number of terms in the formal tropical polynomial $f_{u,\rho}^w$ is polynomial in $|w|$; this will be important in Section 8.

Lemma 5.1. *Let $\phi : \Sigma^+ \rightarrow \Gamma(\mathbb{T})$ be a morphism, and define $\underline{x} \in \mathbb{T}^{\Sigma \times \Gamma}$ by*

$$\underline{x}(s, i) = \phi(s)_{i,i}.$$

Then for any word $w \in \Sigma^+$ and vertices $i, j \in \Gamma$ we have

$$\phi(w)_{i,j} = \bigoplus_{\substack{u \in \Sigma^*, \\ |u| \leq n-1}} \bigoplus_{\rho \in \Gamma_{i,j}^{|u|}} \left(\bigotimes_{k=1}^{|u|} \phi(u_k)_{\rho_{k-1}, \rho_k} \right) \otimes f_{u,\rho}^w(\underline{x}), \quad (1)$$

where $\Gamma_{i,j}^{|u|}$ denotes the set of all $(|u| + 1)$ -vertex paths from i to j in Γ .

Proof. Let i and j be vertices. Using the definition of the functions $f_{u,\rho}^w$, the value given to \underline{x} and the distributivity of multiplication \otimes over addition \bigoplus , the right-hand-side of (1) is equal to

$$\bigoplus_{\substack{u \in \Sigma^*, \\ |u| \leq n-1}} \bigoplus_{\rho \in \Gamma_{i,j}^{|u|}} \left(\bigotimes_{k=1}^{|u|} \phi(u_k)_{\rho_{k-1}, \rho_k} \right) \otimes \left(\bigotimes_{s \in \Sigma} \bigotimes_{k=0}^{|u|} (\phi(s)_{\rho_k, \rho_k})^{\beta_s^w(\alpha_k, \alpha_{k+1})} \right)$$

where the unlabelled sum ranges over all $0 = \alpha_0 < \alpha_1 < \dots < \alpha_{|u|} < \alpha_{|u|+1} = |w| + 1$ such that $w_{\alpha_k} = u_k$ for $k = 1, \dots, |u|$. Notice that we are summing over all possible words u of length less than n , and then over all scattered subwords of w equal to u . Thus, we are simply summing over all scattered subwords of w of length less than n , so the above is equal to:

$$\bigoplus_{\rho \in \Gamma_{i,j}^l} \bigoplus_{\substack{l \\ k=1}} \left(\bigotimes_{k=1}^l \phi(w_{\alpha_k})_{\rho_{k-1}, \rho_k} \right) \otimes \left(\bigotimes_{s \in \Sigma} \bigotimes_{k=0}^l (\phi(s)_{\rho_k, \rho_k})^{\beta_s^w(\alpha_k, \alpha_{k+1})} \right)$$

where the unlabelled sum is over all $0 = \alpha_0 < \alpha_1 < \dots < \alpha_{l+1} = |w| + 1$ for some $l \leq n - 1$ and we set $u = w_{\alpha_0} \dots w_{\alpha_l}$.

Now to each term in the above sum, defined by a choice of α_i 's and a $\rho \in \Gamma_{i,j}^l$, we can associate a $(|w| + 1)$ -vertex walk $(\sigma_0 = i, \dots, \sigma_{|w|} = j)$ in Γ whose underlying path is ρ and which transitions to vertex ρ_k after α_k steps. Clearly every $(|w| + 1)$ -vertex walk from i to j arises exactly once in this way, and so effectively we are summing over all such walks. In each term, the content of the left-hand parentheses gives a factor $\phi(w_q)_{\sigma_{q-1}, \sigma_q}$ when $q = \alpha_k$ for some k , while from the definition of the functions β_s^w , the content of the right-hand parentheses gives a factor $\phi(w_q)_{\sigma_{q-1}, \sigma_q}$ for each q not of this form. Thus, the above is simply equal to:

$$\bigoplus_{q=1}^{|w|} \phi(w_q)_{\sigma_{q-1}, \sigma_q}$$

where the supremum is taken over all $(|w|+1)$ -vertex walks $(i = \sigma_0, \sigma_1, \dots, \sigma_{|w|} = j)$ in Γ . But by the definition of multiplication in $\Gamma(\mathbb{T})$, this is easily seen to be equal to $(\phi(w_1) \otimes \dots \otimes \phi(w_{|w|}))_{i,j} = \phi(w)_{i,j}$. \square

We are now ready to prove the main theorem of this section.

Theorem 5.2. *Let Γ be a partial order with maximum chain length n . Then the identity $w = v$ over alphabet Σ is satisfied in $\Gamma(\mathbb{T})$ if and only if for every $u \in \Sigma^*$ with $|u| \leq n-1$ and every path ρ of length $|u|$ in Γ the tropical polynomials $f_{u,\rho}^w$ and $f_{u,\rho}^v$ are equivalent.*

Proof. Suppose first that $f_{u,\rho}^w(\underline{x}) \neq f_{u,\rho}^v(\underline{x})$ for some choice of $u \in \Sigma^+$, ρ (from i to j) and $\underline{x} \in \mathbb{F}\mathbb{T}^{\Sigma \times \Gamma}$. Define a morphism $\phi : \Sigma^+ \rightarrow \Gamma(\mathbb{T})$ by

$$\phi(s)_{p,p} = \underline{x}(s,p) \in \mathbb{F}\mathbb{T}, \text{ for all } p \in \Gamma \text{ and } s \in \Sigma; \text{ and}$$

$$\phi(s)_{p,q} = \begin{cases} 0 & \text{if } s = u_i, p = \rho_{i-1}, q = \rho_i \\ -\infty & \text{otherwise.} \end{cases}$$

Then by Lemma 5.1,

$$\phi(v)_{i,j} = f_{u,\rho}^v(\underline{x}) \neq f_{u,\rho}^w(\underline{x}) = \phi(w)_{i,j},$$

and so the morphism ϕ falsifies the identity in $\Gamma(\mathbb{T})$.

Conversely, suppose that $f_{u,\rho}^w$ and $f_{u,\rho}^v$ are equivalent for all u, ρ with $u \in \Sigma^*$, and ρ a path of length $|u|$ through Γ . By Proposition 2.2 it suffices to show that the identity $w = v$ is satisfied by every morphism $\phi : \Sigma^+ \rightarrow \Gamma(\mathbb{F}\mathbb{T})$, so let ϕ be such a morphism. Define $\underline{x} \in \mathbb{F}\mathbb{T}^{\Sigma \times \Gamma}$ by $\underline{x}(s,i) = \phi(s)_{i,i}$.

Since ϕ is a morphism to $\Gamma(\mathbb{F}\mathbb{T})$, we know that $\phi(w)_{i,j} = -\infty = \phi(v)_{i,j}$ whenever $i \not\leq j$.

On the other hand, if $i \leq j$ then Lemma 5.1 gives

$$\phi(w)_{i,j} = \bigoplus_{\substack{u \in \Sigma^* \\ |u| \leq n-1}} \bigoplus_{\rho \in \Gamma_{i,j}^{|u|}} \left(\bigotimes_{k=1}^{|u|} \phi(u_k)_{\rho_{k-1}, \rho_k} \right) \otimes f_{u,\rho}^w(\underline{x}) = \phi(v)_{i,j}.$$

\square

Theorem 5.3. *Let Γ be a partial order with maximum chain length n . Then $\Gamma(\mathbb{T})$ satisfies exactly the same semigroup identities as $UT_n(\mathbb{T})$.*

Proof. Let $\rho = (\rho_1, \dots, \rho_n)$ be a maximal length path in Γ , and let $\tau_i = (\rho_1, \dots, \rho_i)$ for $i = 1, \dots, n$. We note that if ρ' and ρ'' are paths of the *same length*, then $f_{u,\rho''}^w$ can be obtained from $f_{u,\rho'}^w$ by the change of variables $x(s, \rho'_i) \mapsto x(s, \rho''_i)$. Thus $f_{u,\rho'}^w = f_{u,\rho''}^w$ if and only if $f_{u,\rho'}^w = f_{u,\rho''}^v$. It follows from this observation together with Theorem 5.2 that $w = v$ in $\Gamma(\mathbb{T})$ if and only if $f_{u,\tau_i}^w = f_{u,\tau_i}^v$ for all words $u \in \Sigma^*$ of length i for all $i = 0, \dots, n-1$. In particular, we note that $w = v$ in $\Gamma(\mathbb{T})$ if and only if $w = v$ in $UT_n(\mathbb{T})$. \square

We remark that Theorem 5.3 can also be deduced as a corollary of Theorem 6.1 below (the proof of which is independent) together with Birkhoff's HSP Theorem [3].

6. CHAIN-STRUCTURED SEMIGROUPS AS DIVISORS OF $UT_n(\mathbb{T})$

As in the previous section, let Γ be a partial order with finite maximum chain length n . Theorem 5.3 says that the chain-structured tropical matrix semigroup $\Gamma(\mathbb{T})$ satisfies exactly the same identities as the upper triangular tropical matrix semigroup $UT_n(\mathbb{T})$. By Birkhoff's HSP Theorem [3], this means that $\Gamma(\mathbb{T})$ must be realisable as a homomorphic image of a subsemigroup of a direct power of $UT_n(\mathbb{T})$. In this section we show how to construct an explicit realisation of $\Gamma(\mathbb{T})$ in this way.

For each $k = 1, \dots, n$ and each k -vertex path ρ in Γ , let Δ_ρ be an isomorphic copy of the semigroup $UT_k(\mathbb{T})$ but with rows and columns indexed by the vertices occurring in ρ , and the order of vertices in ρ used to determine the upper triangular structure in the obvious way (so that for any $M \in \Delta_\rho$ we have $M_{i,j} = -\infty$ if j comes strictly before i in ρ).

Let Δ be the direct product of the semigroups Δ_ρ as ρ varies over all directed paths in Γ . Given $D \in \Delta$ we write D_ρ for projection of D onto the ρ -coordinate (so $D_\rho \in \Delta_\rho$). For $i, j \in \rho$ we write $D_{\rho,i,j}$ for the (i, j) entry of D_ρ .

Notice that since each Δ_ρ is isomorphic to $UT_k(\mathbb{T})$ for some $k \leq n$, Δ embeds naturally in a direct power of $UT_n(\mathbb{T})$.

We define a function (which we do not claim to be morphism)

$$\psi : \Gamma(\mathbb{T}) \rightarrow \Delta, \quad \psi(M)_{\rho,i,j} = M_{ij},$$

which is not surjective; indeed its image need not even be a subsemigroup of Δ . Let Ψ be the subsemigroup of Δ generated by the image of ψ and let

$$\varphi : \Psi \rightarrow \Gamma(\mathbb{T}), \quad \varphi(D)_{i,j} = \sup\{D_{\rho,i,j} : i, j \in \rho\}.$$

It is immediate from the definitions that ψ is an injective function, and that $\varphi(\psi(M)) = M$ for all $M \in \Gamma(\mathbb{T})$. Moreover, the boundedness conditions on $\Gamma(\mathbb{T})$ ensure that φ is well-defined on Ψ . We show that φ is a surjective semigroup morphism.

Theorem 6.1. *With notation as above, φ is a surjective semigroup morphism from Ψ (which is a subsemigroup of Δ and hence embeds naturally in a direct power of $UT_n(\mathbb{T})$) onto $\Gamma(\mathbb{T})$.*

Proof. Since every element of Ψ has the form $\psi(M_1) \cdots \psi(M_k)$ for some $M_1, \dots, M_k \in \Gamma(\mathbb{T})$ and we know that $\varphi(\psi(M)) = M$ for all $M \in \Gamma(\mathbb{T})$, it suffices to take $M_1, \dots, M_k \in \Gamma(\mathbb{T})$ and show that

$$\varphi(\psi(M_1) \cdots \psi(M_k)) = M_1 \otimes \cdots \otimes M_k.$$

Consider the (i, j) entry of the left-hand side. By definition, this is the supremum over all paths ρ containing i and j of the (i, j) entry of $(\psi(M_1) \cdots \psi(M_k))_\rho \in \Delta_\rho$. Since Δ is a direct product, we have

$$(\psi(M_1) \cdots \psi(M_k))_\rho = (\psi(M_1)_\rho) \otimes \cdots \otimes (\psi(M_k)_\rho)$$

in Δ_ρ . By definition of matrix multiplication, the (i, j) entry of this is the supremum over all $(k+1)$ -vertex walks $(i = w_0, w_1, \dots, w_{k-1}, w_k = j)$ within the vertex set of ρ of

$$(M_1)_{w_0 w_1} \otimes \cdots \otimes (M_k)_{w_{k-1} w_k}.$$

But since every $(k + 1)$ -vertex walk is contained in the vertex set of *some* path ρ , this means that the (i, j) entry of the left-hand-side is the supremum over all $(k + 1)$ -vertex walks $(i = w_0, w_1, \dots, w_{k-1}, w_k = j)$ in Γ of

$$(M_1)_{w_0 w_1} \otimes \cdots \otimes (M_k)_{w_{k-1} w_k}.$$

But by the definition of matrix multiplication in $\Gamma(\mathbb{T})$, this is exactly the (i, j) entry of the right-hand side. \square

As mentioned above, the proof of Theorem 6.1 is entirely independent of Theorem 5.3, and in fact the latter can (as an alternative proof strategy) be deduced from the former.

7. THE FREE MONOGENIC INVERSE MONOID

In this section we consider tropical representations of the free one-generated object in the category of inverse monoids, which we denote \mathcal{I} . This monoid admits several representations, but most conveniently for our purposes, it is isomorphic (see for example [14, Chapter 5, Exercise 42]) to the set of triples of integers:

$$\{(i, j, k) \in \mathbb{Z}^3 : i, j \geq 0, -j \leq k \leq i\}$$

with multiplication given by:

$$(i, j, k)(i', j', k') = (\max(i, i' + k), \max(j, j' - k), k + k').$$

Consider now the set $\Gamma = \{1, 2, 3\}$ equipped with the partial order \preceq in which $1, 2 \preceq 3$ but 2 and 3 are incomparable. It is easy to see, from the above representation, that there is an embedding of semigroups

$$\mathcal{I} \hookrightarrow \Gamma(\mathbb{FT}) \subseteq UT_3(\mathbb{T}), \quad (i, j, k) \mapsto \begin{pmatrix} k & -\infty & i \\ -\infty & -k & j \\ -\infty & -\infty & 0 \end{pmatrix}.$$

Since Γ has maximum chain length 2, we can deduce by Theorem 5.3 that \mathcal{I} satisfies all the identities satisfied by $UT_2(\mathbb{T})$, and hence by Theorem 4.1, all identities satisfied in \mathcal{B} . (In fact this was known as a consequence of work of Scheiblich [28] which showed that \mathcal{I} can be embedded in a direct product of copies of \mathcal{B} .) The converse also holds: since \mathcal{B} is itself a monogenic inverse monoid, it is a homomorphic image of \mathcal{I} , and hence satisfies all identities satisfied in the latter.

The above representation of \mathcal{I} in $\Gamma(\mathbb{FT})$ is of course also a representation in $UT_3(\mathbb{T})$. In view of the preceding remarks it is natural to ask if one can go one step better and find a faithful representation of \mathcal{I} in $UT_2(\mathbb{FT})$. It transpires that this cannot be done.

Proposition 7.1. *The free monogenic inverse monoid \mathcal{I} embeds in $UT_n(\mathbb{T})$ if and only if $n \geq 3$.*

Proof. We have seen that \mathcal{I} embeds into $UT_3(\mathbb{T})$ which in turn embeds in $UT_n(\mathbb{T})$ for all $n \geq 3$. To see that \mathcal{I} does not embed into $UT_2(\mathbb{T})$ we consider the structure formed by the idempotents of each semigroup under the natural partial order given by $e \leq f$ if and only if $e = ef = fe$.

It follows from the results of [21] that every idempotent in $UT_2(\mathbb{T})$ takes one of the following four forms:

$$Z = \begin{pmatrix} -\infty & -\infty \\ -\infty & -\infty \end{pmatrix}, G_x = \begin{pmatrix} 0 & x \\ -\infty & 0 \end{pmatrix},$$

$$E_x = \begin{pmatrix} 0 & x \\ -\infty & -\infty \end{pmatrix}, F_x = \begin{pmatrix} -\infty & x \\ -\infty & 0 \end{pmatrix},$$

where $x \in \mathbb{T}$. It is straightforward to verify that:

- (i) $Z \leq X \leq G_{-\infty}$ for all idempotent elements X ;
- (ii) $G_x \leq G_y$ if and only if $y \leq x$; and
- (iii) idempotents of the form E_x and F_x are incomparable with all idempotents except Z and $G_{-\infty}$;

In particular it is not possible within $UT_2(\mathbb{T})$ to find two incomparable idempotents and two other idempotents which lie above both of them.

On the other hand, it is readily verified that the idempotents of \mathcal{I} are exactly the elements of form $(i, j, 0)$, and that $(i, j, 0) \leq (i', j', 0)$ if and only if $i \geq i'$ and $j \geq j'$. Thus, for example, $(2, 3, 0)$ and $(3, 2, 0)$ are incomparable idempotents which both lie below the idempotents $(2, 2, 0)$ and $(1, 1, 0)$. \square

8. COMPLEXITY OF CHECKING IDENTITIES

In this section we show how our results can be used to derive efficient algorithms for checking whether a given identity holds in $UT_n(\mathbb{T})$, and hence (by Theorem 5.3) in chain-structured matrix semigroups in general. To be precise, the results above (Theorem 3.3, Corollary 3.4 and Theorem 5.2) allow us to reduce the problem to checking whether formal tropical polynomials define the same function, and in this section we show how the latter problem can be reduced to (real) linear programming. For a fixed n the resulting algorithms run in time polynomial in the size of the alphabet and the length of the identity, although the degree of the polynomial rises (and hence the complexity increases exponentially) as n grows.

One special case is of particular importance: our Main Theorem implies that the algorithm for $UT_2(\mathbb{T})$ can also be used to check (in polynomial time) whether a given identity holds in the bicyclic monoid (and hence also the free monogenic inverse monoid). Another algorithm for this problem was given by Pastijn [26]; in fact although this algorithm arose from a different approach to an ostensibly different problem, it is very closely related to our own. No complexity analysis is given in [26], Pastijn's algorithm also essentially reduces to linear programming, and hence should also be implementable in polynomial time.

We begin by establishing some basic properties of many-variable tropical polynomials. The ideas concerned are well-known to max-plus algebraists in the one-variable case (see [4, Section 5.1]) and are increasingly studied by algebraic geometers in the many-variable case. However, we have not been able to locate the precise statements we require in the literature, so we present (without any particular claim of originality) a concise self-contained exposition.

A term in a formal tropical polynomial is called *essential* if there is some value of the variables for which it is the only term to attain the maximum,

and *inessential* otherwise. For example, x is an inessential term in $x^{\otimes 2} \oplus x \oplus 1$. We shall need the following observation, which follows from [20, Proposition 1.5 and Lemma 1.6]:

Lemma 8.1. *Every formal tropical polynomial is equivalent to a unique essential polynomial.*

Lemma 8.2. *There is an algorithm which, given a formal tropical polynomial in k variables with m terms with one distinguished term, decides whether the distinguished term is essential in time polynomial in k and m .*

Proof. It suffices to check if there are real values of the variables which make the distinguished term simultaneously exceed each of the other $m - 1$ terms. Since each term of a tropical polynomial is a classical linear function, this is just a (classical) linear programming problem of checking the solvability of $m - 1$ linear inequalities in k real variables, which is solvable in polynomial time (see for example [2]). \square

Theorem 8.3. *Let $n \in \mathbb{N}$ be a fixed positive integer. Then there is an algorithm which, given an identity $v = w$ over a finite alphabet Σ , decides in time polynomial in $|v| + |w|$ and $|\Sigma|$ whether the identity holds in $UT_n(\mathbb{T})$.*

Proof. Suppose we are given an identity $v = w$ where v and w are words over a k -letter alphabet Σ and $|v| + |w| = m$. By Theorem 5.2 it suffices to compute the formal polynomials $f_{u,\rho}^v$ and $f_{u,\rho}^w$ for $u \in \Sigma^*$ with $|u| < n$ and ρ a $(|u|+1)$ -vertex path in the partially ordered set $\{1, \dots, n\}$, and check whether each $f_{u,\rho}^v$ is equivalent to the corresponding $f_{u,\rho}^w$. The number of pairs of polynomials to consider is bounded above by the number of ways to choose u (which since n is fixed is a polynomial function of $k = |\Sigma|$) times the number of paths of length n or less in $\{1, \dots, n\}$ (which since n is fixed is a constant). Moreover, it is easy to see from the definition that each tropical polynomial f_s^w can be computed in time polynomial in $|\Sigma|$ and $|w|$, and so in particular has polynomially many terms.

Now by Lemma 8.2 we may check in polynomial time which terms in our functions are essential; by discarding the inessential terms we obtain, again in polynomial time, essential formal polynomials representing the same functions. By Lemma 8.1 it suffices to check if these are the same as formal polynomials, which again can clearly be done in polynomial time. \square

We emphasise that the algorithm runs in polynomial time only for a fixed semigroup $UT_n(\mathbb{T})$; allowing n to grow results in exponential growth in the number of the pairs of formal polynomials which must be compared, so time complexity as a function of n is exponential.

In the case $n = 2$ we may slightly reduce the number of polynomials to be considered by using Theorem 3.3 in place of Theorem 5.2. In the special case that $n = 2$ and $|\Sigma| = 2$ (hence for two-letter identities over the bicyclic monoid) Corollary 3.4 allows us to reduce the problem to checking equivalence of four pairs of *single-variable* polynomial equations, each of size linear in $|v| + |w|$. Equivalence of single-variable tropical polynomials can be checked even more efficiently: in linear time, presuming a random access model of computation with unit time arithmetic and comparison of numbers [4, Algorithm 5.1.11]). Hence, 2-letter identities in $UT_2(\mathbb{T})$ (and hence by

Theorem 4.1, in \mathcal{B}) can be checked in linear time given in such a model of computation.

REFERENCES

- [1] S. I. Adjan, Defining relations and algorithmic problems for groups and semigroups. *Proceedings of the Steklov Institute of Mathematics*, No. 85 (1966). (Translated from the Russian by M. Greendlinger, American Mathematical Society, Providence, R.I. 1966.)
- [2] B. Aspvall, R. E. Stone, Khachiyan's Linear Programming Algorithm. *Journal of Algorithms* **1** (1980), 1–13.
- [3] G. Birkhoff, On the structure of abstract algebras. *Mathematical Proceedings of the Cambridge Philosophical Society* **31** (1935), 433–454.
- [4] P. Butkovič, Max-linear systems: theory and algorithms. Springer Monographs in Mathematics, Springer, London, 2010.
- [5] A. Cain, G. Klein, L. Kubat, A. Malheiro, J. Okniński, A note on identities in plactic monoids and monoids of upper-triangular tropical matrices. Preprint 2017, arXiv:1705.04596v1.
- [6] Y. Chen, X. Hu, Y. Luo, O. Sapir, The finite basis problem for the monoid of two-by-two upper triangular tropical matrices. *Bulletin of the Australian Mathematical Society* **94** (2016), no. 1, 54–64.
- [7] F. d'Alessandro, E. Pasku, A combinatorial property for semigroups of matrices. *Semigroup Forum* **67** (2003), no. 1, 22–30.
- [8] L. Daviaud, M. Johnson, Minimal identities in the semigroup of 2×2 tropical matrices. Proceedings of the 42nd International Symposium on Mathematical Foundations of Computer Science (MFSC 2017), *Leibniz International Proceedings in Informatics*, to appear.
- [9] S. Gaubert, On the Burnside problem for semigroups of matrices in the $(\max, +)$ algebra. *Semigroup Forum* **52** (1996), no. 3, 271–292.
- [10] S. Gaubert, Two lectures on max-plus algebra. Proceedings of the 26th Spring School of Theoretical Computer Science *Algres Max-Plus et applications en informatique et automatique* (1998).
- [11] M. Gehrke, S. Grigorieff and J.-É. Pin, Duality and equational theory of regular languages. Proceedings of the 35th International Colloquium on Automata, Languages and Programming (ICALP 2008, Part II), in *Lecture Notes in Computer Science* vol. 5126, L. A. et al. (ed.), Springer, Berlin, 2008, pp. 246–257.
- [12] V. Gould, Axiomatisability of free, projective and flat S-acts. Proceedings of Semigroups, acts and categories with applications to graphs (Tartu), in *Mathematics Studies of the Estonian Mathematical Society* **3** (2008), 41–56.
- [13] C. Hollings, M. Kambites, Tropical matrix duality and Green's \mathcal{D} relation. *Journal of the London Mathematical Society* **86** (2012), no. 2, 520–538.
- [14] J. M. Howie, Fundamentals of Semigroup Theory. London Mathematical Society Monographs (New Series) 12, Oxford University Press, New York, 1995.
- [15] Z. Izhakian, Semigroup identities in the monoid of triangular tropical matrices. *Semigroup Forum* **88** (2014), no. 1, 145–161.
- [16] Z. Izhakian, Erratum to: Semigroup identities in the monoid of triangular tropical matrices. *Semigroup Forum* **92** (2016), no. 3, p733.
- [17] Z. Izhakian, Semigroup identities of tropical matrix semigroups of maximal rank. *Semigroup Forum* **92** (2016), no. 3, 712–732.
- [18] Z. Izhakian, M. Johnson, M. Kambites, Pure dimension and projectivity of tropical polytopes. *Advances in Mathematics* **303** (2016), 1236–1263.
- [19] Z. Izhakian, M. Johnson, M. Kambites, Tropical matrix groups. *Semigroup Forum*, to appear.
- [20] Z. Izhakian, S. Margolis, Semigroup identities in the monoid of two-by-two tropical matrices. *Semigroup Forum* **80** (2010), no. 2, 191–218.
- [21] M. Johnson, M. Kambites, Multiplicative structure of 2×2 tropical matrices. *Linear Algebra and its Applications*, **435** (2011), 1612–1625.

- [22] M. Johnson, M. Kambites, Green's \mathcal{J} -order and the rank of tropical matrices. *Journal of Pure and Applied Algebra*, **217** (2013), 280–292.
- [23] P. R. Jones, Analogues of the bicyclic semigroup in simple semigroups without idempotents. *Proceedings of the Royal Society of Edinburgh Section A* **106** (1987), no. 1-2, 11–24.
- [24] G. Merlet, Semigroup of matrices acting on the max-plus projective space. *Linear Algebra and its Applications* **432** (2010), no. 8, 1923–1935.
- [25] J. Okniński, Identities of the semigroup of upper triangular tropical matrices. *Communications in Algebra* **43** (2015), no. 10, 4422–4426.
- [26] F. Pastijn, Polyhedral convex cones and the equational theory of the bicyclic semigroup. *Journal of the Australian Mathematical Society* **81** (2006), no. 1, 63–96.
- [27] J. Reiterman, The Birkhoff theorem for finite algebras. *Algebra Universalis* **14**, (1982), 1–10.
- [28] H. E. Scheiblich, A characterization of a free elementary inverse semigroup. *Semigroup Forum* **2** (1971), no. 1, 76–79.
- [29] M-P. Schützenberger, On the definition of a family of automata. *Information and Control* **4** (1961), 245–270.
- [30] Y. Shitov, A semigroup identity for 3×3 matrices. *Ars Mathematica Contemporanea*, **14** (2018), 15–23.
- [31] F. G. Shleifer, Looking for identities on a bicyclic semigroup with computer assistance. *Semigroup Forum* **41** (1990), no. 2, 173–179.
- [32] L. M. Shneerson, On the axiomatic rank of varieties generated by a semigroup or monoid with one defining relation. *Semigroup Forum* **39** (1989), no. 1, 17–38.
- [33] L. M. Shneerson, On growth, identities and free subsemigroups for inverse semigroups of deficiency one. *International Journal of Algebra and Computation* **25** (2015), no. 1-2, 233–258.
- [34] G. B. Shpiz, G. L. Litvinov, S. N. Sergeev, On common eigenvectors for semigroups of matrices in tropical and traditional linear algebra. *Linear Algebra and its Applications* **439** (2013), no. 6, 1651–1656.
- [35] M. Taylor, *On upper triangular tropical matrix semigroups, tropical matrix identities and \mathbb{T} -modules*. PhD Thesis, University of Manchester, 2017.
- [36] R. J. Warne, I -bisimple semigroups. *Transactions of the American Mathematical Society* **130** (1968) 367–386.