A STRONG GEOMETRIC HYPERBOLICITY PROPERTY FOR DIRECTED GRAPHS AND MONOIDS

ROBERT GRAY[1] and MARK KAMBITES[2]

ABSTRACT. We introduce and study a strong "thin triangle" condition for directed graphs, which generalises the usual notion of hyperbolicity for a metric space. We prove that finitely generated left cancellative monoids whose right Cayley graphs satisfy this condition must be finitely presented with polynomial Dehn functions, and hence word problems in $\mathcal{NP}$. Under the additional assumption of right cancellativity (or in some cases the weaker condition of bounded indegree), they also admit algorithms for more fundamentally semigroup-theoretic decision problems such as Green's relations $\mathcal{L}$, $\mathcal{R}$, $\mathcal{J}$, $\mathcal{D}$ and the corresponding pre-orders.

In contrast, we exhibit a right cancellative (but not left cancellative) finitely generated monoid (in fact, an infinite class of them) whose Cayley graph is a essentially a tree (hence hyperbolic in our sense and probably any reasonable sense), but which is not even recursively presentable. This seems to be strong evidence that no geometric notion of hyperbolicity will be strong enough to yield much information about finitely generated monoids in absolute generality.

1. INTRODUCTION

Over the past half century, combinatorial group theory has been increasing dominated by ideas from geometry. One of the most successful aspects is the theory of *word hyperbolic groups*, in which a simple, combinatorial notion of negative curvature for a group Cayley graph is used to give tight control on the geometric, combinatorial and even computational structure of the group [20]. The question naturally arises of whether these geometric techniques are particular to groups, or if they apply to a wider class of monoids. There are numerous equivalent characterisations of word hyperbolic groups, which lead to different (non-equivalent) ways in which one could define a word hyperbolic monoid. For example, a finitely presented group is word hyperbolic exactly if it has linear *Dehn function*, and this property can also be studied for monoids [10, 29]. A beautiful theorem of Gilman [16] characterises hyperbolic groups as those which admit *context-free multiplication*

tables: several authors have studied the class of monoids satisfying this and similar conditions [9, 11, 14, 22]. Another approach is to treat a monoid Cayley graph as an undirected graph, and require that it be a hyperbolic [9, 13, 15] metric space. In general, the directional information in a monoid or semigroup Cayley graph is of crucial importance in understanding the algebraic structure (and most especially the *ideal* structure) of the monoid, and it is unreasonable to expect much control from any condition on the Cayley graph which disregards the directed structure. For example, it is easily seen that any finitely generated semigroup with a zero element will have undirected Cayley graph of bounded diameter 2. However, this approach does seem to have some merit in classes of semigroups with a very restricted ideal structure, such as *completely simple semigroups* [15].

Here we propose a new way in which hyperbolicity conditions can be extended to cancellative monoids, in a way which is geometric but does not artificially impose a metric structure on the monoid. Specifically, we introduce (in Section 2 below) and study a strong "directed thin triangle" condition for *directed* graphs, which generalises the usual notion of hyperbolicity for a metric space. We prove (in Section 4) that a finitely generated left cancellative monoid whose Cayley graph satisfies this condition must be finitely presented with polynomial Dehn function, and hence solvable (indeed, non-deterministic polynomial time) word problem. Under additional right cancellativity assumptions, we also show (Section 6) that such a monoid admits algorithms for more fundamentally semigroup-theoretic decision problems, including Greens' equivalence relations $\mathcal{L}$, $\mathcal{R}$, $\mathcal{J}$, and $\mathcal{D}$ and corresponding pre-orders. These results suggest that, for at least the class of left cancellative monoids, our thin triangle condition carries with it some genuinely "hyperbolic" structure. Truly geometric methods for left cancellative monoids (see for example [19]) are of considerable interest, for example because the word problem for one-relation monoids (arguably the most significant open question in semigroup theory) is reducible to the left cancellative case [2]. The relationship between one-relator presentations and hyperbolicity for groups has been studied in for example [23].

We have noted that previous notions of hyperbolicity for monoids appear to be too weak, in the sense that they do not suffice to control the behaviour of the monoid in other ways. In contrast, our new definition yields a great deal of control on the monoid but may perhaps be too strong, in that there are examples of "well-behaved" cancellative monoids which one might intuitively expect to be hyperbolic but which do not satisfy this condition. It remains a very interesting and important open question whether there is an intermediate notion which encapsulates a wider class of cancellative monoids while still yielding information comparable with that obtained for groups.

All of the above discussion applies to cancellative (and in some respects left cancellative) monoids. In the more general case of monoids without a cancellativity condition, we suspect there is no meaningful definition of hyperbolicity. In Section 7, we give an elementary way to construct finitely generated monoids (right cancellative but not left cancellative) whose Cayley

graphs are essentially trees (and hence likely to be hyperbolic in any reasonable geometric sense), but which need not even be recursively presentable. This seems to be very strong evidence that no geometric hyperbolicity-type condition on a Cayley graph will will impose much control on a general monoid.

2. DIRECTED GRAPHS AND THIN TRIANGLES

The main objects of study in this paper are directed graphs, which we allow to have loops and multiple directed edges. We shall view directed graphs chiefly as sets of vertices with an asymmetric, partially defined "distance" function given by setting $d(u, v)$ to be the shortest length of a directed path from u to v if there is such a path, or ∞ otherwise. (In particular, when viewed in this way, it is unimportant whether the digraph has loops and/or multiple edges.) With this convention a directed graph is a *semimetric space* of the kind studied in [17, 18, 19].

In more detail, let $\mathbb{R}^\infty$ denote the set $\mathbb{R}^{\geq 0} \cup \{\infty\}$ of non-negative real numbers with ∞ adjoined. We equip it with the obvious order, addition and multiplication, leaving 0∞ undefined. Now let X be a set. A function $d : X \times X \to \mathbb{R}^\infty$ is called a *semimetric* on X if:

 (i) $d(x, y) = 0$ if and only if $x = y$; and
 (ii) $d(x, z) \leq d(x, y) + d(y, z)$

for all $x, y, z \in X$. A set equipped with a semimetric on it is a *semimetric space*. In particular any directed graph is a semimetric space, where the distance between two vertices defined to be the length of the shortest directed path between them, or ∞ if there is no such path.

If X is a directed graph, the *underlying undirected graph* of X is the graph with the same vertex set and an edge from x to y whenever there is an edge from x to y or an edge from y to x. In particular, if in the directed graph there are multiple directed edges connecting a given pair of vertices, then in the underlying undirected graph there will be a single undirected edge joining this pair of vertices.

Now let X be a directed graph (or semimetric space), $x_0 \in X$ be a point, and r be a positive real number or ∞. The *out-ball* of radius r based at x_0 is

$$\overrightarrow{\mathcal{B}}_r(x_0) = \{y \in X \mid d(x_0, y) \leq r\}.$$

Dually, the *in-ball* of radius r based at x_0 is defined by

$$\overleftarrow{\mathcal{B}}_r(x_0) = \{y \in X \mid d(y, x_0) \leq r\},$$

and the *strong ball* of radius r based at x_0 is

$$\mathcal{B}_r(x_0) = \overrightarrow{\mathcal{B}}_r(x_0) \cap \overleftarrow{\mathcal{B}}_r(x_0).$$

If S is a set of points then $\overrightarrow{\mathcal{B}}_r(S)$, $\overleftarrow{\mathcal{B}}_r(S)$ and $\mathcal{B}_r(x_0)$ are defined to be the unions of the appropriate out-balls, in-balls and strong balls respectively around points in S.

A directed graph (or semimetric space) is called δ-*bounded* if no finite distance in the space exceeds δ (although there may be points at distance ∞ in one or both directions). It is called *bounded* if it is δ-bounded for some finite δ.

FIGURE 1. A schematic illustration of a δ-thin directed geodesic triangle in a directed graph: each side of the triangle is contained in the union of the δ-outball around the side that meets its initial vertex, and the δ-inball around the side that meets its terminal vertex.

For $n \geq 0$, a *path* [of length n in a directed graph X] is a sequence $[x_0, \ldots, x_n]$ of vertices such that X has an edge from x_i to x_{i+1} for $0 \leq i < n$. The vertices x_0 and x_n are the *start* and *end* of the path, respectively, and are denoted ιp and τp. The path is called a *geodesic* if $n = d(x_0, x_n)$. A path is called *simple* if it does not contain a repeated vertex.

Now let $p = [x_0, \ldots, x_n]$ and $q = [y_0, \ldots, y_m]$ be paths in X. We say that p and q are *parallel* if $x_0 = y_0$ and $x_n = y_m$, and we write $p \parallel q$. The paths p and q are *composable* if $x_n = y_0$, in which case we define $p \circ q$ to be the path $[x_0, \ldots, x_{n-1}, y_0, \ldots, y_m]$. A *geodesic triangle* in X is an ordered triple (p, q, r) of geodesics such that p and q are composable, and $p \circ q$ is parallel to r.

Definition 2.1. *Let δ be a non-negative real number. A geodesic triangle in a directed graph is δ-**thin** if each side x is contained in the union of the out-balls of radius δ around points on the side of the triangle meeting the start of x, together with the in-balls of radius δ around points on the side of the triangle meeting the end of x.*

This definition is illustrated schematically in Figure 1. Note that although the three sides of a directed triangle play different roles, the definition treats them in a uniform way. Made explicit, the definition says that a directed geodesic triangle (p, q, r) is δ-thin if:

- every vertex in r is contained within the union of the out-balls of radius δ around vertices in p, and the in-balls of radius δ around vertices in q; and
- every vertex in p is contained within the union of the out-balls of radius δ around vertices in r, and the in-balls of radius δ around vertices in q; and

- every vertex in q is contained within the union of the out-balls of radius δ around vertices in p, and the in-balls of radius δ around vertices in r.

Definition 2.2. *A directed graph is called strongly δ-hyperbolic if all of its geodesic triangles are δ-thin.*

Given an undirected graph Y we say that this graph is strongly δ-hyperbolic if the digraph obtained by replacing each edge of Y by a pair of oppositely oriented directed edges is strongly δ-hyperbolic. With this definition an undirected graph is strongly δ-hyperbolic if and only if it is δ-hyperbolic in the classical sense (viewed as a metric space).

The following propositions give some basic examples of strongly hyperbolic directed graphs.

Proposition 2.3. *Any δ-bounded directed graph is strongly δ-hyperbolic.*

Proof. Since δ is an upper bound on the finite distances in the graph, it is an upper bound on the length of geodesics in the graph. It follows that every vertex within a directed geodesic triangle is contained with an out-ball of radius δ around its start point, which suffices to show that every directed geodesic triangle is δ-thin. $\square$

Proposition 2.4. *Let X be a directed graph whose underlying undirected graph is a tree. Then X is strongly 0-hyperbolic.*

Proof. Let Y be the undirected graph underlying X, and let (p, q, r) be a directed geodesic triangle in X. Then in particular p, q and r represent simple paths in X, and hence also in Y. Since Y is a tree and $p \circ q$ is parallel to r, it follows that there is a prefix p' of p and a suffix q' of q such that $r = p'q'$. But then every vertex of r lies on either p or q. Similar arguments show that every vertex of p also lies on q or r, and every vertex of q also lies on p or r. This suffices to show that X is strongly 0-hyperbolic. $\square$

Note however that it is *not* the case in general that if the underlying undirected graph of a digraph X is hyperbolic then X itself must be strongly hyperbolic. A counterexample may easily be constructed by taking a digraph X that is not strongly hyperbolic, and considering the digraph X^0 obtained from X by adding an extra vertex z and an edge from every vertex to z. Then by Proposition 2.6 below the digraph X^0 will not be strongly hyperbolic, while the underlying undirected graph of X^0 will be strongly hyperbolic, since it has bounded diameter.

Associated with any semimetric space X is a natural preorder $\lesssim$ relation given by $x \lesssim y$ if and only if $d(y, x) < \infty$. Let $\sim$ denote the equivalence relation given by $x \sim y$ if and only if $x \lesssim y$ and $y \lesssim x$. We call the $\sim$-classes the *strongly connected components* of X.

Proposition 2.5. *If X is a strongly δ-hyperbolic directed graph, then every strongly connected component of X is strongly δ-hyperbolic.*

Proof. Let Y be a strongly connected component of a strongly δ-hyperbolic graph X. First note that since no directed paths between vertices of Y pass outside Y, distances between vertices in Y are the same in Y as in X. It

FIGURE 2. An infinite directed grid, giving an example of a
directed graph that is strongly δ-hyperbolic (since there are
no directed geodesic triangles) but such that the underlying
undirected graph is not δ-hyperbolic.

follows that every directed geodesic triangle in Y is also a directed geodesic
triangle in X, and hence is δ-thin in X. Since distances are the same in X
as in Y, it follows that every directed geodesic triangle in Y is δ-thin in Y.
Thus, Y is strongly δ-hyperbolic. $\square$

Proposition 2.6. *Let X be a directed graph, and let X^0 be the directed
graph obtained from X by adding an extra vertex z and an edge from every
vertex of X^0 to the vertex z. If X is strongly δ-hyperbolic then X^0 is strongly
$\max(1, \delta)$-hyperbolic. Conversely, if X^0 is strongly δ-hyperbolic then X is
strongly δ-hyperbolic.*

Proof. Suppose X is strongly δ-hyperbolic, and let (p, q, r) be a directed
geodesic triangle in X^0. If (p, q, r) lies entirely in X then it is δ-thin in X,
and hence in X^0. Otherwise, it contains the vertex z, and since there are
no non-loop edges out of z, it follows that the common end-point of q and r
is z. But since q and r are geodesics and there are edges from every vertex
of X to the vertex z, it follows that q and r are sides of length 1, and the
only vertices on them are the vertices of the triangle. Thus, every vertex
on q is either on p or r, and every point on r is either on p or q. Finally,
every point on p is contained in an in-ball of radius 1 about z, which lies on
q. Thus, the triangle is 1-thin. Hence, X^0 is strongly ϕ-hyperbolic where
$\phi = \max(1, \delta)$.

The converse follows from Proposition 2.5, since X is a strongly connected
component in X^0. $\square$

3. TRIANGLE AND POLYGON INEQUALITIES

One of the difficulties of working with semimetric spaces is the limited
nature of the triangle inequality. If (p, q, r) is a geodesic triangle then cer-
tainly $|r| \le |p| + |q|$, but we do not automatically have an upper bound on
$|p|$ in terms of $|q|$ and $|r|$, or on $|q|$ in terms of $|p|$ and $|r|$. In many cases

of interest, however, hyperbolicity allows us to acquire such bounds, albeit rather weaker than the conventional triangle inequality.

Recall that a semimetric space is called *quasi-metric* if there is a constant λ such that $d(x, y) \leq \lambda d(y, x) + \lambda$ for all points x and y. Equivalently, a space is quasi-metric if it is quasi-isometric to a metric space [18].

Lemma 3.1. *Let X be a locally finite, strongly δ-hyperbolic directed graph with indegree and outdegree bounded by α. Then there is a constant λ, depending only on (and polynomial-time computable from) δ and α, such that whenever p, q, and r are the sides of a directed geodesic triangle in X (in no particular order), we have*

$$|p| \leq \lambda(|q| + |r|).$$

Proof. If p is the hypotenuse the result follows from the standard triangle inequality for semimetric spaces. Now suppose that r is the hypotenuse and that $\tau p = \iota q$ (the third remaining case being dual to this one). If $|r| = |q| = 0$ then, since p is geodesic, it follows that $|p| = 0$ and the result holds. So we may now suppose that either $|r| \geq 1$ or $|q| \geq 1$. Since p is geodesic it has no repeated vertices. Also since X is strongly δ-hyperbolic it follows that the vertices of p are all contained in the set $\overrightarrow{\mathcal{B}}_\delta(r) \cup \overleftarrow{\mathcal{B}}_\delta(q)$. We may assume that the graph X has at least one directed edge, since otherwise the result is trivially true, and so in particular $\alpha \geq 1$. As the outdegree is bounded by α it follows that the outdegree is also bounded by $\alpha + 1$ and so

$$|\overrightarrow{\mathcal{B}}_\delta(r)| \leq (|r| + 1) \sum_{i=0}^{\delta} ((\alpha + 1)^i) = (|r| + 1) \left(\frac{(\alpha + 1)^{\delta+1} - 1}{(\alpha + 1) - 1} \right).$$

(We take $\alpha + 1$ rather than α here just to avoid having to deal with the case $\alpha = 1$ seperately.) Also, the indegree is bounded by α, so one obtains a similar bound for $|\overleftarrow{\mathcal{B}}_\delta(q)|$, which combined with the above inequality gives

$$|p| \leq (|r| + 1 + |q| + 1) \left(\frac{(\alpha + 1)^{\delta+1} - 1}{(\alpha + 1) - 1} \right) \leq \left(\frac{3((\alpha + 1)^{\delta+1} - 1)}{(\alpha + 1) - 1} \right) (|q| + |r|),$$

completing the proof. $\square$

Corollary 3.2. *Let X be a locally finite strongly δ-hyperbolic directed graph with indegree and outdegree bounded by α. Then there is a constant λ, depending only on (and polynomial-time computable from) δ and α, such that every strongly connected component of X is $(\lambda, 0)$-quasi-metric.*

Proof. Let λ be the constant given by Lemma 3.1. Now suppose x and y belong to the same strongly connected component of X. Let p be a geodesic from x to y, q a geodesic from y to x, and e the (geodesic) empty path at x. Then the triple (x, y, e) is a geodesic triangle, and so Lemma 3.1 yields

$$d(x, y) = |p| \leq \lambda(|q| + |e|) = \lambda(d(y, x) + 0) = \lambda d(y, x).$$

$\square$

Definition 3.3 (Directed geodesic n-gon)**.** *A directed geodesic n-gon in a directed graph Γ is an n-tuple $(p_1, \ldots, p_{n-1}, q)$ of geodesic paths such that $p = p_1 \circ p_2 \circ \cdots \circ p_{n-1}$ is defined, and $p \parallel q$.*

Theorem 3.4 (Polygon quasi-inequality). *Let X be a locally finite strongly δ-hyperbolic directed graph with indegree and outdegree bounded by α. Then there is a constant K, depending only on (and polynomial-time computable from) α and δ, such that every side length of a directed geodesic n-gon is bounded above by K times the sum of the other side lengths.*

Proof. Let λ be the constant given by Lemma 3.1 and let $K = \max(\lambda^2, 1)$.

The case $n = 3$ is immediate from Lemma 3.1. Let $(p_1, \ldots, p_{n-1}, q)$ be a directed geodesic n-gon. Since q is geodesic it is immediate from the triangle inequality that $|q| \leq |p_1| + \cdots + |p_n| \leq K(|p_1| + \cdots + |p_n|)$. It remains only to prove that the sides p_i satisfy the claimed bound when $n \geq 4$.

We treat first the case where $n = 4$ and $i = 2$. (This case is special because our general strategy will involve choosing two composable sides of the polygon excluding p_i, and this is the only case where this approach is impossible.) In this case, let r be a geodesic from ιp_1 to τp_2. Then (r, p_3, q) and (p_1, p_2, r) are geodesic triangles, so applying Lemma 3.1 twice we have $|r| \leq \lambda(|p_3| + |q|)$ and

$$|p_2| \leq \lambda(|p_1| + |r|) \leq \lambda(|p_1| + \lambda(|p_3| + |q|)) \leq K(|p_1| + |p_3| + |q|)$$

as required.

We establish the result in the remaining cases by induction on n. Consider, then, a geodesic n-gon $(p_1, \ldots, p_{n-1}, q)$ where either $n > 4$, or $n = 4$ but $i \neq 2$, and suppose the claim holds for all geodesic $(n-1)$-gons.

Suppose first that $i > 2$. Let r be a geodesic from ιp_1 to τp_2. Then by the triangle inequality $|r| \leq |p_1| + |p_2|$. Now $(r, p_2, \ldots, p_{n-1}, q)$ is a geodesic $(n-1)$-gon, so by the inductive hypothesis we have

$$
\begin{aligned}
|p_i| &\leq K\left(|r| + |p_3| + \cdots + |p_{i-1}| + |p_{i+1}| + \cdots + |p_{n-1}| + |q|\right) \\
&\leq K\left((|p_1| + |p_2|) + |p_3| + \cdots + |p_{i-1}| + |p_{i+1}| + \cdots + |p_{n-1}| + |q|\right).
\end{aligned}
$$

Finally suppose $i \leq 2$. By assumption, either $n > 4$, or else $n = 4$ but $i \neq 2$. It follows that we have $i \neq p_{n-1}$ and $i \neq p_{n-2}$. Now we can apply the same argument as in the previous case, but taking r this time to be a geodesic from ιp_{n-1} to τp_n.   □

We note that Lemma 3.1 and Theorem 3.4 can both fail if the hypothesis of local finiteness is dropped. Indeed, let X be the directed graph with vertex set $\mathbb{Z}$ and an edge from i to $i + 1$ for each $i \in \mathbb{Z}$, and let $Y = X^0$ be the graph obtained from X by the construction in Proposition 2.6. Then Y is strongly 1-hyperbolic by Propositions 2.4 and 2.6, but contains directed geodesic triangles with two sides of length 1 and the third arbitrarily long.

4. Monoids and Cayley Graphs

Let M be a monoid generated by a finite set S. Then M naturally has the structure of a directed graph (with vertex set M, and an edge from m to n for each $s \in S$ such that $ms = n$), and hence also of a semimetric space. This directed graph is called the *right Cayley graph* of M.

Definition 4.1. *A monoid M generated by a finite subset S is called strongly δ-hyperbolic (with respect to the generating set S) if its right Cayley graph with respect to S is strongly δ-hyperbolic. A monoid M is called strongly*
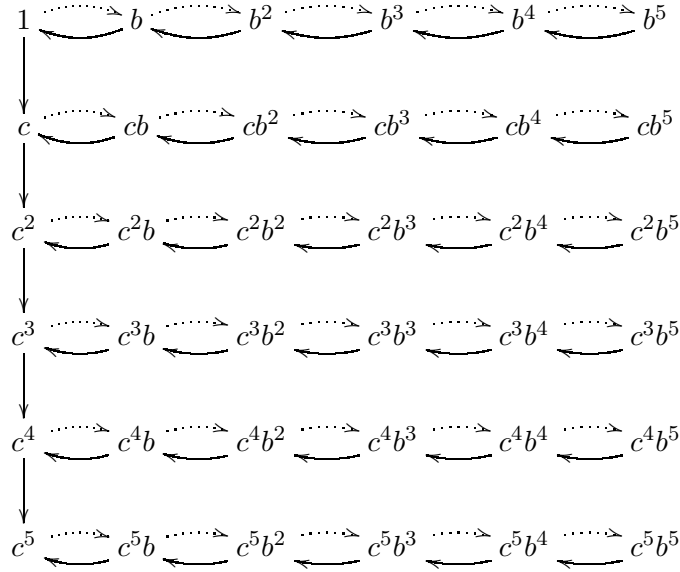
FIGURE 3. A partial view of the right Cayley graph of the bicyclic monoid $\langle b, c \mid bc = 1 \rangle$ where $\rightarrow$ corresponds to multiplication by c and $\dashrightarrow$ corresponds to multiplication by b.

hyperbolic if it is strongly δ-hyperbolic for some δ with respect to some finite generating set.

We do not presently know whether a strongly hyperbolic monoid is necessarily strongly hyperbolic with respect to every choice of finite generating set. This question is deserving of further study.

Our initial results about strongly δ-hyperbolic spaces immediately give a number of examples of monoids which are strongly hyperbolic.

Example 4.2. *Every finite monoid is strongly δ-hyperbolic with respect to every generating set, where δ is the length of the longest geodesic representative for an element of M. Indeed, it is easy to see that the Cayley graph is δ-bounded, so this follows from Proposition 2.3.*

Example 4.3. *Free monoids of finite rank are strongly 0-hyperbolic with respect to their free generating sets. Indeed, the underlying undirected graph of the Cayley graph is a tree, so this follows from Proposition 2.4.*

Example 4.4. *The bicyclic monoid $\mathcal{B} = \langle p, q \mid pq = 1 \rangle$ is strongly 0-hyperbolic with respect to the standard generating set $\{p, q\}$. Again, the underlying undirected graph of the Cayley graph (see Figure 3) is a tree, so this follows from Proposition 2.4.*

Example 4.5. *Polycyclic monoids of finite rank are strongly 1-hyperbolic with respect to their standard generating sets. Recall that for $n \geq 2$ the polycyclic monoid $\mathcal{P}_n$ of rank n is given by the presentation*

$$\langle p_1, \ldots, p_n, q_1, \ldots, q_n, z \mid p_i q_i = 1, p_i q_j = z = p_i z = q_i z = z p_i = z q_i \text{ for all } i \neq j \rangle.$$

The generator z represents a zero element. Let Y be the graph obtained from the Cayley graph of $\mathcal{P}_n$ (with respect to the generating set from the presentation above) by removing the vertex z and all edges incident with it. Then it is straightforward to verify that the underlying undirected graph is a tree, and hence is strongly 0-hyperbolic by Proposition 2.4. Moreover, the Cayley graph of $\mathcal{P}_n$ can be recovered from Y by the construction in Proposition 2.6, so by that proposition $\mathcal{P}_n$ itself is strongly 1-hyperbolic.

Example 4.6. *Word hyperbolic groups are strongly hyperbolic in our sense. Indeed, suppose G is word hyperbolic and choose a finite generating set S for G which is closed under the taking of inverses. Then S is also a monoid generating set for G, and the distance function on the monoid Cayley graph is the same as that on the group Cayley graph. The claim now follows from the usual "thin triangle" property of hyperbolic metric spaces.*

We can also expand our class of examples by showing closure under some elementary semigroup-theoretic constructions.

Proposition 4.7. *Let M be a monoid, and let M^0 be the monoid obtained from M by adjoining a new element 0 which acts as a zero element. Then M is strongly hyperbolic if and only if M^0 is strongly hyperbolic.*

Proof. Suppose M is strongly hyperbolic with respect to a generating set S. Then M^0 is generated by $S \cup \{0\}$. The Cayley graph of M^0 with respect to this generating set is clearly obtained from that of M by the construction in Proposition 2.6, and so by the proposition M^0 is strongly hyperbolic.

Conversely, suppose M^0 is strongly hyperbolic with respect to a generating set T. Since $M^0 \setminus \{0\}$ is a subsemigroup of M^0, we must have $0 \in T$. Now $T \setminus \{0\}$ is a generating set for M, and just as above the Cayley graph for M^0 with respect to T is again obtained from that for M with respect to $T \setminus \{0\}$ by the construction in Proposition 2.6. Thus, by the proposition, M is strongly hyperbolic $\qquad\square$

Proposition 4.8. *Let M be a finitely generated monoid and let I be an ideal of M. If M is strongly δ-hyperbolic then the Rees quotient M/I is strongly hyperbolic.*

Proof. Suppose that M is strongly hyperbolic with respect to a finite generating set S, and let 0 be the 0 element in the Rees quotient M/I. Then M/I is generated by the set $A = (S \cap (M \setminus I)) \cup \{0\}$. Let X be the Cayley graph of M/I with respect to A, and Y the Cayley graph of M with respect to S. Let $x_0 \in X$ be the vertex of X corresponding to the zero element 0 of M/I. Notice that for any two non-0 elements of M/I (that is, elements of $M \setminus I$), we have $d_X(a, b) = d_Y(a, b)$.

We claim that X is strongly $\max(1, \delta)$-hyperbolic. Let (p, q, r) be a directed geodesic triangle in X. If (p, q, r) does not involve the vertex 0 then the distances between vertices visited are the same in X as in Y; but Y is strongly δ-hyperbolic, so (p, q, r) is δ-thin in Y and hence in X. Otherwise, (p, q, r) contains the vertex 0 and, arguing as in the proof of Proposition 2.6, since in X there are no edges out of 0, it follows that the common endpoint of q and r is 0. But every point of (p, q, r) has a directed geodesic of length at most 1 to 0. It follows that (p, q, r) is 1-thin, completing the proof. $\quad\square$

Corollary 4.9. *Let M be a strongly hyperbolic monoid, and S a submonoid which is the complement of an ideal. Then S is strongly hyperbolic.*

Proof. Let $I = M \setminus S$. Then the Rees quotient M/I is isomorphic to S^0, so the result follows from Propostions 4.7 and 4.8. □

Recall that an element x of a monoid is called a *unit* if there is an element y such that $xy = yx = 1$; the set of all units forms a (maximal) subgroup of M.

Corollary 4.10. *The group of units of a cancellative strongly hyperbolic monoid is a hyperbolic group.*

Proof. It is well-known and easy to prove that the complement of the group of units in a cancellative monoid forms an ideal, so this follows from Corollary 4.9. □

Of course, the converse to the latter corollary does not hold in general: for example, the free commutative monoid of rank two is cancellative with trivial (hence hyperbolic) group of units, but is itself not a strongly hyperbolic monoid.

5. Directed 2-Complexes, Presentations and Dehn Functions

It is well known that, even if a monoid is given by a finite presentation, the word problem for the monoid may be undecidable. Markov [26] and Post [30] proved independently that the word problem for finitely presented monoids is undecidable in general; this result was extended by Turing [33] to cancellative semigroups, and then by Novikov and Boone to groups (see [24] for references). For classes of monoids that do have decidable word problem it is natural to consider the complexity of the word problem. For example, monoids which admit presentations by finite complete rewriting systems have solvable word problem, but there is no bound on the complexity of the word problem for such monoids; see [3]. On the other hand, automatic monoids have word problem that is solvable in quadratic time [12, Corollary 3.7]. As mentioned in the introduction, a finitely presented group is word hyperbolic exactly if it has linear Dehn function; recent results of Cain [8] show that word-hyperbolic semigroups (in the sense of Duncan and Gilman [14]) have word problem solvable in polynomial time. A remarkable result of Birget [4] characterises finitely generated semigroups with word problem in $\mathcal{NP}$ as exactly those embeddable in finitely presented semigroups with polynomial Dehn function. (An analogous statement for groups was proved later [7, 32].)

Our main aim in this section is to prove that finitely generated, left cancellative, strongly δ-hyperbolic monoids are finitely presented with polynomial Dehn functions, and therefore admit non-deterministic polynomial-time word problem solutions. Our proof is most easily and intuitively expressed in the language of *direct 2-complexes* [19, 21], so we begin by briefly recalling some definitions and results concerning these.

Let $P(\Gamma)$ denote the set of all directed paths in a directed graph Γ, including empty paths at each vertex. A *directed 2-complex* is a directed graph Γ equipped with a set F (called the *set of 2-cells*), and three maps

$\lceil \cdot \rceil : F \to P(\Gamma)$, $\lfloor \cdot \rfloor : F \to P(\Gamma)$, and $^{-1} : F \to F$ called *top, bottom,* and *inverse* such that

- for every $f \in F$, the paths $\lceil f \rceil$ and $\lfloor f \rfloor$ are parallel;
- $^{-1}$ is an involution without fixed points, and $\lceil f^{-1} \rceil = \lfloor f \rfloor$, $\lfloor f^{-1} \rfloor = \lceil f \rceil$ for every $f \in F$.

If K is a directed 2-complex, then the directed paths on K are called 1-*paths*. For every 2-cell $f \in F$, the vertices $\iota(\lceil f \rceil) = \iota(\lfloor f \rfloor)$ and $\tau(\lceil f \rceil) = \tau(\lfloor f \rfloor)$ are denoted $\iota(f)$ and $\tau(f)$, respectively.

An *atomic* 2-*path* is a triple (p, f, q), where p, q are 1-paths in K, and $f \in F$ such that $\tau(p) = \iota(f)$, $\tau(f) = \iota(q)$. If δ is an atomic 2-path then we use $\lceil \delta \rceil$ to denote $p \lceil f \rceil q$ and $\lfloor \delta \rfloor$ is denoted by $p \lfloor f \rfloor q$, these are the top and bottom 1-paths of the atomic 2-path. A *non-trivial* 2-path δ in K is then a sequence of atomic paths $\delta_1, \ldots, \delta_n$, where $\lfloor \delta_i \rfloor = \lceil \delta_{i+1} \rceil$ for every $1 \leq i < n$, and the length of this 2-path is n. The top and bottom 1-paths of δ, denoted $\lceil \delta \rceil$ and $\lfloor \delta \rfloor$ are then defined as $\lceil \delta_1 \rceil$ and $\lfloor \delta_n \rfloor$, respectively.

We use $\delta \circ \delta'$ to denote the composition of two 2-paths. We say that 1-paths p, q in K are homotopic if there exists a 2-path δ such that $\lceil \delta \rceil = p$ and $\lfloor \delta \rfloor = q$. Recall that a pair of paths $p, q \in P(\Gamma)$ are said to be parallel, written $p \parallel q$, if $\iota p = \iota q$ and $\tau p = \tau q$. We say that a directed 2-complex K is *directed simply connected* if for every pair of parallel paths $p \parallel q$, p and q are homotopic in K.

Let K be a directed 2-complex with underlying directed graph Γ and set of 2-cells F, and let $T = (p, q, r)$ be a directed triangle in Γ. Now let K' be the 2-complex obtained from K by adjoining one new element f to F satisfying $\lceil f \rceil = p \circ q$ and $\lfloor f \rfloor = r$. We call K' the directed 2-complex obtained from K by adjoining a 2-cell for the triangle T.

Definition 5.1 (Tessellation)**.** *Given a pair of parallel paths p and q in a directed graph, we say that a set $T_1, T_2, \ldots, T_r$ of directed geodesic triangles tessellates p and q if in the 2-complex K obtained by adjoining 2-cells for each T_i we have $p \sim_K q$. We say that a set of geodesic triangles tessellates a directed n-gon $(p_1, \ldots, p_{n-1}, r)$ if it tessellates the paths $p_1 \circ \cdots \circ p_{n-1}$ and r.*

Definition 5.2. *Given a directed triangle $T = (p, q, r)$ in a directed graph, we define the size $\Sigma(T)$ of T to be $|p| + |q|$.*

Our strategy for establishing our main result is to show that in a strongly δ-hyperbolic directed graph we can tessellate the "gap" between two parallel paths with (polynomially many, as a function of the path lengths) geodesic triangles of bounded size. We begin by showing that every pair of parallel paths can be tessellated by geodesic triangles (of not necessarily bounded size).

Lemma 5.3. *Let Γ be a directed graph. Every pair of parallel paths $p \parallel q$ in Γ can be tessellated by $|p| + |q| + 1$ directed geodesic triangles of size at most $2(|p| + |q|)$.*

Proof. We prove the result by induction on $|p| + |q|$. The base case is trivial since any path of length 0 or 1 is automatically geodesic. For the induction step, if p and q are both geodesic then (p, q) itself naturally may be viewed

as a single directed geodesic triangle of the required size and we are done. Now suppose that p, say, is not geodesic. Decompose $p = p_\iota \circ p_e \circ p_\tau$ where $p_\iota \circ p_e$ is the shortest non-geodesic subpath of p and p_e is a directed edge. Let p' be a geodesic path from ιp to τp_e. Then $|p' \circ p_\tau| < |p|$ and by induction the pair $(p' \circ p_\tau, q)$ may be tessellated by $|p' \circ p_\tau| + |q| + 1$ directed geodesic triangles of size at most $2(|p' \circ p_\tau| + |q|) < 2(|p| + |q|)$. Taken together with the directed geodesic triangle (p_i, p_e, p') which also has size less than $2(|p| + |q|)$ we conclude that (p, q) may be tessellated by

$$(|p' \circ p_\tau| + 1) + |q| + 1 \leq |p| + |q| + 1$$

directed geodesic triangles of size at most $2(|p| + |q|)$. $\qquad \square$

Lemma 5.4. *Let Γ be a strongly δ-hyperbolic directed graph. Then every directed geodesic triangle T can be tessellated by five directed geodesic triangles (some of which may be trivial triangles with a single vertex) of size no more than $\frac{3}{4}\Sigma(T) + 2\delta + 1$.*

Proof. For clarity in this proof, we will use the convention that XY denotes a geodesic path from a vertex X to a vertex Y, while XYZ denotes the directed geodesic triangle (XY, YZ, XZ). (Of course geodesics are not unique; we will be careful to make clear where the choice is important.)

Let $T = PQR = (PQ, QR, PR)$ be a directed geodesic triangle. Let $|PQ| = k$, $|QR| = l$ so that $\Sigma(T) = k + l$. Suppose $k \geq l$ (the case $l \geq k$ being dual); it follows in particular that $l + \frac{1}{2}k \leq \frac{3}{4}(l + k)$. Let M be the vertex on the geodesic PQ satisfying $d(P, M) = \lfloor \frac{k}{2} \rfloor$ and $d(M, Q) = \lceil \frac{k}{2} \rceil$. Since Γ is strongly δ-hyperbolic there are now two cases to consider.

Case (a): $\overrightarrow{\mathcal{B}}_\delta(M)$ **intersects** QR. Let O be a point in $\overrightarrow{\mathcal{B}}_\delta(M)$ which lies on QR. Consider geodesics MO, PO, PM, MQ, QO and OR, chosen so that $QO \circ OR = QR$ and $PM \circ MQ = PQ$. Consider also the three geodesic triangles: $T_1 = MQO$, $T_2 = PMO$ and $T_3 = POR$ (see Figure 4). The sizes of these triangles are bounded as follows:

$$\begin{aligned}
\Sigma(T_1) &\leq \left\lceil \frac{k}{2} \right\rceil + l \leq \frac{3}{4}(k + l) + 1 = \frac{3}{4}\Sigma(T) + 1, \\
\Sigma(T_2) &\leq \left\lfloor \frac{k}{2} \right\rfloor + \delta \leq \frac{1}{2}\Sigma(T) + \delta, \\
\Sigma(T_3) &\leq \left\lfloor \frac{k}{2} \right\rfloor + \delta + l \leq \frac{3}{4}\Sigma(T) + \delta.
\end{aligned}$$

So in this case, our triangle is tesselated by the three triangles T_1, T_2 and T_3, which satisfy the required size bound.

Case (b): $\overleftarrow{\mathcal{B}}_\delta(M)$ **intersects** PR. Let O be a point in $\overleftarrow{\mathcal{B}}_\delta(M)$ which lies on PR. Consider geodesics OM, OQ, PM, MQ, PO and OR such that $PM \circ MQ = PQ$ and $PO \circ OR = PR$. Consider also the three geodesic triangles $T_1 = POM$, $T_2 = OMQ$ and $T_3 = OQR$ (see Figure 5). The

FIGURE 4. Proof of Lemma 5.4, Case (a).

triangles, T_2 and T_3 have size bounds:

$$\Sigma(T_2) \leq \delta + \left\lceil \frac{k}{2} \right\rceil \leq \frac{1}{2}\Sigma(T) + \delta + 1,$$

$$\Sigma(T_3) \leq \left\lceil \frac{k}{2} \right\rceil + \delta + l \leq \frac{3}{4}(k+l) + \delta + 1 = \frac{3}{4}\Sigma(T) + \delta + 1,$$

It is not immediately clear how to get a such a size bound on T_1, so we further subdivide it in the middle of the edge PO, at the point U. So set $x = |PO|$ and let U be the vertex on the geodesic PO satisfying $d(P,U) = \left\lfloor \frac{x}{2} \right\rfloor$ and $d(U,O) = \left\lceil \frac{x}{2} \right\rceil$. Now there are two subcases to consider.

Case (b)(i): $\overrightarrow{\mathcal{B}}_\delta(U)$ intersects OM. Let S be a point in $\overrightarrow{\mathcal{B}}_\delta(U)$ which lies on OM. Choose geodesics PU, UO, UM, US, SM and OS so that $PU \circ UO = PO$ and $OS \circ SM = OM$. Now the directed geodesic triangle POM is tessellated by the three triangles $Y_1 = PUM$, $Y_2 = USM$ and $Y_3 = UOS$ (see Figure 6) and the sizes of these three triangles are bounded as follows:

$$\Sigma(Y_1) \leq \left\lfloor \frac{x}{2} \right\rfloor + 2\delta \leq \left\lfloor \frac{k+l}{2} \right\rfloor + 2\delta \leq \frac{1}{2}\Sigma(T) + 2\delta,$$

$$\Sigma(Y_2) \leq 2\delta,$$

$$\Sigma(Y_3) \leq \left\lceil \frac{x}{2} \right\rceil + \delta \leq \left\lceil \frac{k+l}{2} \right\rceil + \delta \leq \frac{1}{2}\Sigma(T) + \delta + 1.$$

Case (b)(ii): $\overleftarrow{\mathcal{B}}_\delta(U)$ intersects PM. In this case choose S in $\overleftarrow{\mathcal{B}}_\delta(U)$ which lies on PM. Choose geodesics PU, UO, SU, SO, PS and SM so that $PU \circ UO = PO$ and $PS \circ SM = PM$. Now the directed geodesic

FIGURE 5. Proof of Lemma 5.4, Case (b).

FIGURE 6. Proof of Lemma 5.4, Case (b)(i).

triangle POM is tessellated by the three triangles $Y_1 = PSU$, $Y_2 = SUO$ and $Y_3 = SOM$ (see Figure 7) and the sizes of these three triangles are

FIGURE 7. Proof of Lemma 5.4, Case (b)(ii).

bounded as follows:

$$\Sigma(Y_1) \leq \left\lfloor \frac{k}{2} \right\rfloor + \delta \leq \frac{1}{2}\Sigma(T) + \delta,$$

$$\Sigma(Y_2) \leq \delta + \left\lceil \frac{x}{2} \right\rceil \leq \delta + \left\lceil \frac{k+l}{2} \right\rceil \leq \frac{1}{2}\Sigma(T) + \delta + 1,$$

$$\Sigma(Y_3) \leq \left\lceil \frac{x}{2} \right\rceil + 2\delta \leq \left\lceil \frac{k+l}{2} \right\rceil + 2\delta \leq \frac{1}{2}\Sigma(T) + 2\delta + 1.$$

Thus, in both case (b)(i) and case (b)(ii), our triangle is tessellated by the triangles T_2, T_3, Y_1, Y_2 and Y_3 which satisfy the required size bound. $\qquad\square$

Theorem 5.5. *Let Γ be a strongly δ-hyperbolic directed graph and $C > 8\delta + 4$ a constant. Then every directed geodesic triangle T in Γ can be tessellated by*

$$5 \left(\frac{\Sigma(T)}{C - 8\delta - 4} \right)^{\log_{\frac{4}{3}} 5}$$

or fewer geodesic triangles of size C or less.

Proof. Let T be a directed geodesic triangle in Γ. Then by Lemma 5.4, T is tessellated by $T_1, \ldots, T_5$ where for all i

$$\Sigma(T_i) \leq \frac{3}{4}\Sigma(T) + (2\delta + 1).$$

We iterate this procedure, at each stage subdividing every triangle from the previous stage into five triangles in this way. Define a sequence t_0 of natural numbers by $t_0 = \Sigma(T)$ and

$$t_{i+1} = \frac{3}{4}t_i + (2\delta + 1).$$

for $i \leq n$. A simple induction argument applying Lemma 5.4 shows that t_i is an upper bound on the size of the triangles obtained in the ith iteration.

Now for each $k \in \mathbb{N}$ we have

$$t_k = \left(\frac{3}{4}\right)^k t_0 + \left(\sum_{i=0}^{k-1} \left(\frac{3}{4}\right)^i\right)(2\delta + 1) \leq \left(\frac{3}{4}\right)^k t_0 + 4(2\delta + 1).$$

Let $D = C - 8\delta - 4$ and let N be the integer part of $\log_{4/3} \frac{t_0}{D} + 1$. Then rearranging we have

$$N \geq \log_{\frac{4}{3}} \frac{t_0}{D}, \text{ so } \left(\frac{4}{3}\right)^N \geq \frac{t_0}{D}, \text{ so } \left(\frac{3}{4}\right)^N t_0 \leq D = C - 8\delta - 4$$

$$\text{and hence } t_N \leq \left(\frac{3}{4}\right)^N t_0 + 8\delta + 4 \leq C.$$

Thus, after N iterations, we have tessellated T with triangles of size C or less. Moreover, since at each stage we subdivide each triangle into at most five triangles, the number of triangles in this tessellation is bounded above by 5^N, where

$$5^N \leq 5^{\left(\log_{\frac{4}{3}} \frac{t_0}{D}\right)+1} = 5 \times 5^{\log_{\frac{4}{3}} \frac{t_0}{D}} = 5 \times \left(\frac{t_0}{D}\right)^{\log_{\frac{4}{3}} 5}$$

$$\leq 5 \left(\frac{\Sigma(T)}{C - 8\delta - 4}\right)^{\log_{\frac{4}{3}} 5}$$

as required. $\square$

Our main aim with Theorem 5.5 was to give a reasonably concise argument for the existence of a polynomial bound on the number of triangles of fixed size required to tessellate a geodesic triangle, rather than to optimise the degree of the polynomial. The figure of $\log_{4/3} 5$ (which is approximately 5.6) can probably be lowered significantly at the expense of lengthening the proof, either by analysing more precisely the properties of the subdivision given by Lemma 5.4, or by considering alternative subdivisions.

Given two functions $f, g : \mathbb{N} \to \mathbb{N}$ we write $f \prec g$ if there exists a constant a such that $f(j) \leq ag(aj) + aj$ for all j. The functions f and g are said to be of the same *type*, written $f \sim g$, if $f \prec g$ and $g \prec f$.

Now fix a monoid presentation $\langle A \mid R \rangle$. If u and v are equivalent words then the *area* $A(u, v)$ is the smallest number of applications of relations from R necessary to transform u into v. The *Dehn function* of a presentation $\langle A \mid R \rangle$ is the function $\delta : \mathbb{N} \to \mathbb{N}$ given by

$$\delta(n) = \max\{A(u, v) \mid u, v \in A^*, u \equiv_R v, |u| + |v| \leq n\}.$$

The Dehn function is a measure of the complexity of transformations between equivalent words. The Dehn function depends on the presentation, but if δ and γ are Dehn functions of different finite presentations for the same monoid then $\delta \sim \gamma$ (see [25, 31]).

A corollary of the above result is the following.

Theorem 5.6. *Let M be a strongly δ-hyperbolic left cancellative monoid. Then M is finitely presented with Dehn function bounded above by a polynomial of degree $\log_{\frac{4}{3}} 5 + 1$.*

Proof. Let Γ be the Cayley graph of M, with respect to a generating set which makes M strongly δ-hyperbolic. Choose an integer $C > 8\delta + 4$.

Suppose (p, q) is a pair of parallel paths in Γ. Then by Lemma 5.3, (p, q) can be tessellated by at most $|p| + |q| + 1$ geodesic triangles of size at most $2(|p| + |q|)$. By Theorem 5.5, each of these may be tessellated by at most

$$5 \left(\frac{\Sigma(T)}{C - 8\delta - 4} \right)^{\log_{\frac{4}{3}} 5}$$

geodesic triangles of size at most C. Thus, (p, q) can be tessellated by at

$$5(|p| + |q| + 1) \left(\frac{2(|p| + |q|)}{C - 8\delta - 4} \right)^{\log_{\frac{4}{3}} 5}$$

directed geodesic triangles of size at most C.

Each such triangle will correspond to a face in $K_C(\Gamma)$, so this shows that $K_C(\Gamma)$ is simply connected with Dehn function bounded above by a polynomial of degree $\log_{\frac{4}{3}} 5 + 1$. It follows by the results of [19] that M is finitely presented with Dehn function bounded above by a polynomial of this degree. $\square$

Theorem 5.7. *Let M be a finitely generated, left cancellative, strongly δ-hyperbolic monoid. Then the word problem for M lies in $\mathcal{NP}$.*

Proof. By Theorem 5.6 M is finitely presented with polynomial Dehn function. Let $\langle A \mid R \rangle$ be a finite presentation, and $p : \mathbb{N} \to \mathbb{N}$ a polynomial upper bound on the corresponding Dehn function. Now given words $u, v \in A^*$, one may check non-deterministically if $u = v$ in M by guessing a sequence of relation applications of length $p(|u| + |v|)$ which can be applied to u, and seeing if the result of applying them is v. $\square$

6. Deciding Green's Relations

The statements of our main results for monoids so far have been direct analogues of known results in the group case, although the proofs have been rather more involved. If geometric techniques are to have more than a very limited application in semigroup theory, it is important that they give insight into aspects of the structure theory of semigroups which do not arise in groups such as, for example, the ideal structure of a semigroup. Recall that *Green's relations* are a collection of equivalence relations and pre-orders (reflexive, transitive binary relations) defined on any monoid (or semigroup) which encapsulate the structure of its principle left, right and two-sided ideals and maximal subgroups. They are a key tool in modern semigroup theory, playing a pivotal role in almost every area of the subject.

If S is a monoid then we define pre-orders $\leq_{\mathcal{R}}$, $\leq_{\mathcal{L}}$, $\leq_{\mathcal{J}}$ and equivalence relations $\mathcal{R}$, $\mathcal{L}$, $\mathcal{J}$, $\mathcal{H}$ and $\mathcal{D}$ by

- $a \leq_{\mathcal{R}} b$ if and only if $aS \subseteq bS$;
- $a \leq_{\mathcal{L}} b$ if and only if $aS \subseteq bS$;
- $a \leq_{\mathcal{J}} b$ if and only if $SaS \subseteq SbS$;
- $a\mathcal{R}b$ if and only if $aS = bS$ (that is, if $a \leq_{\mathcal{R}} b$ and $b \leq_{\mathcal{R}} a$);
- $a\mathcal{L}b$ if and only if $Sa = Sb$ (that is, if $a \leq_{\mathcal{L}} b$ and $b \leq_{\mathcal{L}} a$);
- $a\mathcal{J}b$ if and only if $SaS = SbS$ (that is, if $a \leq_{\mathcal{J}} b$ and $b \leq_{\mathcal{J}} a$);

- $a\mathcal{H}b$ if and only if $a\mathcal{R}b$ and $a\mathcal{L}b$; and
- $a\mathcal{D}b$ if and only if there exists $c \in S$ with $a\mathcal{L}c$ and $c\mathcal{R}b$.

A monoid is called $\mathscr{J}$-trivial if the $\mathcal{J}$ relation (and hence also the $\mathcal{D}$, $\mathcal{L}$, $\mathcal{R}$ and $\mathcal{H}$ relations) are the identity relation.

Green's relations $\mathcal{R}$ and $\mathcal{L}$. We shall now see how the triangle quasi-inequality for locally finite strongly δ-hyperbolic directed graphs, and more generally the polygon quasi-inequality for strongly δ-hyperbolic locally finite directed graphs, can be usefully applied to prove decidability results for Green's relations in strongly δ-hyperbolic monoids.

The questions of decidability, and the complexity of deciding, Green's relations have been considered for semigroups defined by finite complete rewriting systems [27], automatic monoids [28], word-hyperbolic semigroups [8], and for the Thompson-Higman monoids [5, 6]. In particular, it was shown in [27] that there exists monoids that are presented by finite, length-reducing, and confluent string-rewriting systems (and therefore in particular have solvable word problem) but have Green's relations $\mathcal{R}$ and $\mathcal{L}$ that are undecidable. Also, in [28] examples are given of finitely generated monoids M with word problem solvable in quadratic time but such that $\mathcal{R}$ (respectively $\mathcal{L}$) is undecidable. For strongly δ-hyperbolic monoids, the quasi-triangle inequality prevents this from happening.

Recall that a monoid M generated by a finite subset A has *(right) indegree bounded by a natural number α* if one cannot choose generator $a \in A$ and $\alpha + 1$ distinct elements $x_0, \ldots, x_\alpha \in M$ such that $x_0 a = x_1 a = \cdots = x_\alpha a$. It is easily seen that the property of having bounded indegree is independent of the choice of finite generating set, although the actual bound may vary. Having bounded indegree is also equivalent to saying that the right Cayley graph of M (with respect to any or every choice of finite generating set) has bounded valency. Notice that a right cancellative finitely generated monoid always has bounded indegree, and indeed bounded indegree is often viewed as a weak right cancellativity condition.

Theorem 6.1. *Let M be a finitely generated strongly δ-hyperbolic monoid with bounded indegree. Then the problems of deciding the $\mathcal{R}$-order $\leq_\mathcal{R}$ and $\mathcal{L}$-order $\leq_\mathcal{L}$ are reducible in non-deterministic linear time to the word problem for M.*

Proof. Let A be a finite generating set with respect to which M is strongly δ-hyperbolic. Let $w, u \in A^*$, and suppose $\alpha \in A^*$ is of minimal length such that $w\alpha = u$ in S. Let w' and u' be a geodesic words representing the same elements as w and u respectively. Then (w', α, u') labels a geodesic triangle in the Cayley graph of M with respect to A, so by Theorem 3.4, there is constant K, depending only on δ, the maximum indegree of Γ and $|A|$, such that

$$|\alpha| \leq K(|w'| + |u'|) \leq K(|w| + |u|).$$

Thus, given $w, u \in A^*$, to test non-deterministically if $u \leq_\mathcal{R} w$, it suffices to guess a word $\alpha \in A^*$ of length at most $K(|w| + |u|)$, and then test if $w\alpha = u$.

The proof for $\leq_\mathcal{L}$ is entirely similar. $\qquad\square$

Note that in Theorem 6.1 we do not require the monoid to be left cancellative, although we still have a weak *right* cancellativity assumption in the form of the bounded indegree hypothesis. Combining Theorem 6.1 with Theorem 5.6 we obtain the following.

Theorem 6.2. *Let M be a finitely generated left cancellative monoid of bounded indegree which is strongly hyperbolic. Then the $\mathcal{L}$-order and $\mathcal{R}$-order for M are both in $\mathcal{NP}$.*

Neither the left cancellativity nor the strong hyperbolicity condition in Theorem 6.2 can be dropped. In Section 7 below we shall see examples of finitely generated, strongly 0-hyperbolic which have unsolvable word problems and all of Green's equivalence relations trivial, and hence also unsolvable. Also, as mentioned above, it is well known that there exist finitely presented groups with unsolvable word problem. Let G be such a group given by a finite monoid presentation $\langle A|R \rangle$, and define $M = \langle A, h|R \rangle$ where h is a symbol not in A. Then M is a two-sided cancellative monoid (since it is the monoid free product of G and the free monoid of rank one, both of which are cancellative). Moreover, for all words $u, w \in A^*$ we have

$$hu\mathcal{L}hw \Leftrightarrow w = u \text{ in } G, \text{ which is undecidable,}$$

$$uh\mathcal{R}wh \Leftrightarrow w = u \text{ in } G, \text{ which is undecidable,}$$

and

$$huh\mathcal{J}hwh \Leftrightarrow huh\mathcal{D}hwh \Leftrightarrow w = u \text{ in } G, \text{ which is undecidable.}$$

Thus, none of the relations $\mathcal{R}$, $\mathcal{L}$, $\mathcal{J}$ or $\mathcal{D}$ is decidable in M.

Green's relations $\mathcal{J}$ and $\mathcal{D}$. Next we look at the relations $\mathcal{J}$ and $\mathcal{D}$. The following technical lemma will be used to study Green's $\mathcal{J}$-relation. Intuitively speaking, it says that, in a geodesic quadrangle (p, q, r, s), if the side r is sufficiently long then there will be a short path from p to r. This will be using for carving up geodesic quadrangles into smaller geodesic quadrangles.

Lemma 6.3. *Let Γ be a strongly δ-hyperbolic, locally finite directed graph with indegree and outdegree bounded by α, and let (p, q, r, s) be a geodesic quadrangle. Then there are polynomial-time computable constants $C_{\delta,|s|}$, depending on δ and $|s|$, and $D_{\alpha,\delta,|q|,|s|}$, depending on α, δ, $|q|$ and $|s|$, such that if $|r| > D_{\alpha,\delta,|q|,|s|}$ then there is a geodesic path t in Γ satisfying $\iota t \in p$, $\tau t \in r$, $d(\tau t, \tau s) = C_{\delta,|s|}$, and $|t| \leq 2\delta$.*

Moreover, if the graph Γ is fixed then $C_{\delta,|s|}$ may be chosen to increase monotonically with $|s|$, and the $D_{\alpha,\delta,|q|,|s|}$ may be chosen to increase monotonically with $|q|$ and $|s|$, and to be bounded above by a linear function of $|q| + |s|$.

Proof. Let $K_{\alpha,\delta}$ be the constant, given by Theorem 3.4, depending on α and δ, and having the property that for any geodesic quadrangle in Γ each side has length bounded by $K_{\alpha,\delta}$ times the sum of the length of the other three sides. Define

$$C_{\delta,|s|} = \delta + |s| + 1, \tag{1}$$

and

$$D_{\alpha,\delta,|q|,|s|} = K_{\alpha,\delta}(K_{\alpha,\delta}(|s| + \delta + C_{\delta,|s|}) + \delta + |q| + |s|). \tag{2}$$

Since $K_{\alpha,\delta}$ is polynomial-time computable, it is clear that these values are also computable in polynomial time. Moreover, if the graph (and hence α and δ) remain fixed then clearly the values increase monotonically with the remaining variables, and $D_{\alpha,\delta,|q|,|s|}$ can be bounded above by a linear function of $|q| + |s|$.

Let $x \in r$ with $d(x, \tau s) = C_{\delta,|s|}$. Let u be a geodesic path from ιs to τq, and consider the directed geodesic triangle (u, r, s). Since Γ is strongly δ-hyperbolic either $\overrightarrow{\mathcal{B}}_\delta(x) \cap s \neq \varnothing$ or else $\overleftarrow{\mathcal{B}}_\delta(x) \cap u \neq \varnothing$. However, the first of these possibilities cannot arise since

$$d(x, \tau s) = C_{\delta,|s|} > \delta + |s|,$$

and thus we conclude that $\overleftarrow{\mathcal{B}}_\delta(x) \cap u \neq \varnothing$, so we may choose $y \in u$ with $d(y, x) \leq \delta$. Now consider the directed geodesic triangle (p, q, u) and the point $y \in u$. Since Γ is strongly δ-hyperbolic, either (i) $\overrightarrow{\mathcal{B}}_\delta(y) \cap q \neq \varnothing$ or (ii) $\overleftarrow{\mathcal{B}}_\delta(y) \cap p \neq \varnothing$.

Suppose, seeking a contradiction, that (i) holds. Then $d(y, \tau q) \leq \delta + |q|$. Now consider a directed geodesic quadrangle formed by geodesics from ιs to y, y to x, x to τs and ιs to τs. By Theorem 3.4 we have

$$d(\iota s, y) \leq K_{\alpha,\delta}(|s| + \delta + C_{\delta,|s|}),$$

and therefore

$$|u| = d(\iota s, \tau q) = d(\iota s, y) + d(y, \tau q) \leq K_{\alpha,\delta}(|s| + \delta + C_{\delta,|s|}) + \delta + |q|.$$

Then, again applying the triangle quasi-inequality from Theorem 3.4, we know that

$$|r| \leq K_{\alpha,\delta}(|u| + |s|) \leq K_{\alpha,\delta}(K_{\alpha,\delta}(|s| + \delta + C_{\delta,|s|}) + \delta + |q| + |s|),$$

which is a contradiction, since r was assumed to satisfy

$$r > D_{\alpha,\delta,|q|,|s|}.$$

We deduce that (ii) $\overleftarrow{\mathcal{B}}_\delta(y) \cap p \neq \varnothing$, say $z \in \overleftarrow{\mathcal{B}}_\delta(y) \cap p$. Now the lemma follows simply by setting t to be a geodesic path from z to x. $\qquad \square$

The main application of the above lemma is to the proof of the following one, which will be key to establishing the main results concerning $\mathcal{D}$ and $\mathcal{J}$.

Lemma 6.4. *Let M be a left cancellative monoid which is strongly δ-hyperbolic with respect to a finite generating set A, and has indegree with respect to A bounded by α.*

Then for every pair $u, v \in A^*$ of geodesic words, there is a constant $F_{|A|,\alpha,\delta,|u|,|v|}$ (depending on, and polynomial-time computable from, $|A|$, α, δ, $|u|$, $|v|$) such that if $u \mathcal{J} v$ in M then there exist words $a, b \in A^*$ such that $aub = v$ and $|a|, |b| \leq F_{|A|,\alpha,\delta,|u|,|v|}$.

If, in addition, M is cancellative and $puq = v$ in M for some units p and q, then a and b may be chosen to represent units.

Moveover, for a fixed monoid M, $F_{|A|,\alpha,|u|,|v|}$ is monotonically increasing as a function of $|u|$ and $|v|$, and can be bounded above by a linear function of $|u| + |v|$.

Proof. Let $a, b \in A^*$ be such that (i) $aub = v$, (ii) if M is cancellative and it is possible to choose them so, a and b represent units, (iii) a is a geodesic word and (iv) $|b|$ is minimal subject to the preceding three conditions (and in particular is a geodesic word).

Since a, u, b and v are all geodesic words, $aub = v$ in M and M is left cancellative, they label the sides of a geodesic quadrangle in the right Cayley graph of M with respect to A; for brevity we identify the words a, u, b and v with the geodesic paths they label in this quadrangle.

The intuition of the proof is that we work our way from τv to τu marking points at regular intervals on the geodesic labelled by b. Each of these points will be the terminal vertex of a path labelled by a word t_k from a point of the geodesic labelled by a to that labelled by b, and each of the words t_k will have length bounded by 2δ. If b were excessively long then two of the words, t_i and t_j say, must coincide. Then using left cancellativity we can perform a cut and paste operation gluing t_i along t_j and in the process reduce the length of the word b. See Figure 8 for an illustration of the argument.

In more detail, without loss of generality we may suppose that δ is an integer. Let $W_{A,2\delta}$ be the number of words over A of length less than or equal to 2δ, so $W_{A,2\delta}$ is equal to $2\delta + 1$ if $|A| = 1$, and is equal to $(|A|^{2\delta+1} - 1)/(|A| - 1)$ otherwise. First we define

$$E_{|A|,\alpha,\delta,|u|,|v|} = D_{\alpha,\delta,|u|,|v|} + D_{\alpha,\delta,|u|,2\delta} + C_{\delta,|v|} + W_{A,2\delta}C_{\delta,2\delta}.$$

Notice that if the monoid and generating set are fixed (so that $|A|$, α and δ are constant) then $W_{A,2\delta}$ and $C_{\delta,2\delta}$ are constant, $D_{\alpha,\delta,|u|,|v|} + D_{\alpha,\delta,|u|,2\delta}$ is bounded above by a linear function in $|u| + |v|$ by Lemma 6.3 and since δ is constant, and $C_{\delta,|v|}$ is bounded above by a linear function in $|u| + |v|$, again by Lemma 6.3 and since δ is constant. Thus, $E_{|A|,\alpha,\delta,|u|,|v|}$ can be bounded above by a linear function of $|u| + |v|$.

We claim that

$$|b| \leq E_{|A|,\alpha,\delta,|u|,|v|}.$$

Indeed, suppose false for a contradiction. Decompose $b = b_1 \circ c_1$ where $|c_1| = C_{\delta,|v|}$. Then by considering the geodesic quadrangle (a, u, b, v) and applying Lemma 6.3, there is a geodesic path labelled by a word t_1 from a point of a to τb_1 with $|t_1| \leq 2\delta$. Let a_1 be the suffix of a leading from the start point of this path. Now starting from $i = 1$, we repeatedly take the geodesic quadrangle corresponding to (a_i, u, b_i, t_i), and write $b_i = b_{i+1} \circ c_{i+1}$ where $|c_{i+1}| = C_{\delta,|t_i|}$. So long as $|b_i| > D_{\alpha,\delta,|u|,|t_i|}$ we may use Lemma 6.3 to find a geodesic path, labelled by a word t_{i+1} of length at most 2δ, from a point on a_i to τb_i. Let a_{i+1} be the suffix of a_i leading from the start point of this path, and then repeat for the next value of i.

Notice that we can continue this process for at least $N = W_{A,2\delta} + 1$ steps, since

$$|b| > E_{|A|,\alpha,\delta,|u|,|v|} \geq D_{\alpha,\delta,|u|,|v|},$$

which is needed for the first step,

$$|b| > E_{|A|,\alpha,\delta,|u|,|v|} \geq C_{\delta,|v|} + D_{\alpha,\delta,|u|,2\delta} \geq C_{\delta,|v|} + D_{\alpha,\delta,|u|,|t_1|},$$

FIGURE 8. Proof of Lemma 6.4.

which allows the second step, and

$$|b| > E_{|A|,\alpha,\delta,|u|,|v|} \geq C_{\delta,|v|} + W_{A,2\delta} C_{\delta,2\delta} + D_{\alpha,\delta,|u|,2\delta}$$

$$\geq C_{\delta,|v|} + \sum_{i=1}^{k} C_{\delta,|t_i|} + D_{\alpha,\delta,|u|,|t_{k+1}|},$$

for all $k = 1, 2, \ldots, W_{A,2\delta}$, by the monotonically increasing nature of $C_{\delta,|s|}$ and $D_{\alpha,\delta,|u|,|v|}$, which allows us to carry out the the third up to the $(W_{A,2\delta} + 1)$th step.

The result, after N steps, is depicted in Figure 8. At this point we have N words t_1, t_2, $\ldots$, t_N all of length less than 2δ. By the choice of $W_{A,2\delta}$ and the pigeon-hole principle, it follows that there exists an $i < j$ such that $t_i = t_j$.

This gives rise to a decomposition $a = a'a''a'''$, where a' labels the path from ιa to ιt_i, a'' the path from ιt_i to ιt_j, and a''' the path from ιt_j to τa. Similarly, we write $b = b'b''b'''$, where b' labels the path from ιb to τt_j, b'' the path from τt_j to τt_i and b''' the path from τt_i to τb. Note that since $i \neq j$, $|b''| > 0$. Now we have $a'a''a'''ub' = a'a''t_j$ in the monoid, and hence by left cancellativity, $a'''ub' = t_j$. But now

$$(a'a''')u(b'b''') = a'(a'''ub')b''' = a't_jb''' = a't_ib''' = v.$$

But b'' is non-empty, so $|b'b'''| < |b|$. Moreover, if M is cancellative and b represents a unit then b, and hence $b'b'''$ will be entirely composed of generators representing units, so $b'b'''$ will also be a unit. This contradicts the minimality in the choice of b, and establishes the claim that $|b| \leq E_{|A|,\alpha,\delta,|u|,|v|}$.

Now, if we let $\beta = \max(\alpha, |A|)$ then by Theorem 3.4 (the polygon inequality) we have

$$|a| \leq K_{\beta,\delta}(|u| + |v| + |b|).$$

so it suffices to set

$$F_{|A|,\alpha,\delta,|u|,|v|} = (K_{\beta,\delta} + 1)(|u| + |v| + E_{|A|,\alpha,\delta,|u|,|v|}).$$

The fact that $F_{|A|,\alpha,\delta,|u|,|v|}$ is polynomial-time computable, monotonically increasing and bounded above by a linear function follow from the corresponding properties for $E_{|A|,\alpha,\delta,|u|,|v|}$. □

For our main result concerning the $\mathcal{D}$ relation, we shall need the following elementary fact about cancellative monoids, a proof of which we include for completeness.

Lemma 6.5. *Let M be a cancellative monoid and let $a, b \in M$. Then $a\mathcal{D}b$ if and only if there are units $q, r \in U(M)$ satisfying $qar = b$.*

Proof. If $a\mathcal{D}b$ then by definition there exists $c \in M$ such that $a\mathcal{L}c$ and $c\mathcal{R}b$. The former means there are $p, q \in M$ with $pc = a$ and $qa = c$, so $qpc = c = 1.c$ and $pqa = a = 1.a$, which by cancellativity implies $pq = qp = 1$, so q is a unit. A dual argument gives $b = cr$ for some unit r, and now $b = cr = qar$ as required. The converse is immediate. $\square$

We are now ready to prove our main theorem about $\mathcal{D}$ and $\mathcal{J}$.

Theorem 6.6. *Let M be a finitely generated, left cancellative monoid of bounded indegree which is strongly hyperbolic. Then the $\mathcal{J}$-order is in $\mathcal{NP}$.*
  If, moreover, M is cancellative, then the $\mathcal{D}$-relation is in $\mathcal{NP}$.

Proof. Let A be a finite generating set with respect to which M is strongly δ-hyperbolic, and α a corresponding bound on the indegree. Suppose $u, v \in A^*$ are such that $v \leq_{\mathcal{J}} u$ in M. Let u' and v' be geodesic words representing the same elements as u and v respectively. Then applying Lemma 6.4 we have $au'b = v'$ in M for some elements $a, b \in A^*$ with

$$|a|, |b| \;\leq\; F_{|A|, \alpha, \delta, |u'|, |v'|} \;\leq\; F_{|A|, \alpha, \delta, |u|, |v|},$$

where the last inequality follows from the monotonicity of $F_{|A|, \alpha, \delta, |u|, |v|}$ as a function of $|u| + |v|$. Since $u = u'$ and $v = v'$ in M, it follows also that $aub = v$ in M.

Thus, given words u and v, to check non-deterministically if $v \leq_{\mathcal{J}} u$ it suffices compute the constant $F = F_{|A|, \alpha, \delta, |u|, |v|}$, guess words a and b of length no more than F, and test if $aub = v$ in M.

By Lemma 6.4, the computation of F can be performed in polynomial time, and F itself is bounded by a polynomial function of $|u| + |v|$. Thus, the words a and b to be guessed have polynomial length. Finally, the test of whether $aub = v$ in M can be performed in non-deterministic polynomial time by Theorem 5.7. Thus, the whole procedure is possible in (non-deterministic) polynomial time, so the $\mathcal{J}$-order is in $\mathcal{NP}$.

For the case of $\mathcal{D}$ in a cancellative monoid, let B be the set of generators in A which represent units in M. It is easily seen that B^* is exactly the set of words in A^* representing units in M. Thus, by Lemmas 6.4 and 6.5, two words u and v represent $\mathcal{D}$-related elements if and only if there are words $a, b \in B^*$ with $|a|, |b| \leq F$ and $aub = v$ in M. So to check if $u\mathcal{D}v$ in M, we can use exactly the same procedure as above but considering only those a and b in B^*. $\square$

We note that the second part of Theorem 6.6 is not a trivial consequence of the first, since there are finitely presented cancellative monoids for which the relations $\mathcal{D}$ and $\mathcal{J}$ do not coincide. For instance consider the monoid M defined by the presentation

$$\langle\, x, y, a, b \mid axb = y, \; ayb = x \,\rangle.$$

It follows from results of Adjan [1], since the relations have neither left cycles or right cycles, that M is a cancellative monoid (in fact, it is group embeddable). It is a straightforward exercise to verify that in this monoid the elements represented by x and y are $\mathcal{J}$-related, but they are not $\mathcal{D}$-related (since the group of units is trivial, and the $\mathcal{R}$- and $\mathcal{L}$-relations are trivial).

7. Monoids which are Not Left Cancellative

The definition of strong δ-hyperbolicity makes sense for arbitrary directed graphs, and hence for arbitrary monoids, but most of our results so far have required a left cancellativity assumption on the monoid. One might reasonably ask what can be deduced about a more general finitely generated monoid, given only the information that it is strongly δ-hyperbolic. In this section we present a class of finitely generated monoids which are right cancellative but not left cancellative and which, although strongly 0-hyperbolic, need not even be recursively presentable. Indeed, the Cayley graph of each monoid is in most respects very like a tree, being a rooted directed tree with the addition of some duplicate edges, and so is likely to satisfy any reasonable geometric definition of hyperbolicity. We believe this provides very strong evidence that geometric hyperbolicity conditions cannot provide the same kind of information about general monoids that they do about groups, and perhaps about other restricted classes of monoids.

Given $w = a_1 a_2 \cdots a_r$ and $k \in \{1, \ldots, r\}$ write $w[k] = a_1 \cdots a_k$. For each subset I of $\mathbb{N}$ define

$$M_I = \langle\, a, b, c, d \mid ab^i c = ab^i d \ (i \in I) \,\rangle.$$

Lemma 7.1. *Let $w, u \in \{a, b, c, d\}^*$. If $w = u$ in M_I then $|w| = |u|$ and for all $k \in \{1, \ldots, |w|\}$ we have $w[k] = u[k]$ in M_I.*

Proof. This can be shown by a straightforward induction on the number of applications of relations required to transform w into u. It suffices just to consider the case where w and u are separated just by the application of a single relation, so $w = \alpha ab^i c \beta$, $u = \alpha ab^i d \beta$. But in this case the result clearly holds. $\square$

Corollary 7.2. *M_I is right cancellative and $\mathcal{J}$-trivial.*

Proof. Suppose $x, y, z \in M_I$ are such that $xz = yz$. Choose words u, v, w respectively to represent them, so that $uw = vw$ in M_I. Then by Lemma 7.1, $|uw| = |vw|$, so $|u| = |v|$. Now taking $k = |u| = |v|$ and using Lemma 7.1 again we have that $u = v$ in M_I, so $x = y$.

If $u \mathcal{J} v$ in M_I then we have $puq = v$ and $rvs = u$ in M_I for some words p, q, r and s. But now by Lemma 7.1 again, $|u| = |rvs| = |rpuqs|$, which means p, q, r and s are all the empty word and $u = v$ in M_I. $\square$

Proposition 7.3. *For every subset I of $\mathbb{N}$, M_I is a finitely generated strongly 0-hyperbolic monoid.*

Proof. Let Γ be the right Cayley graph of M_I with respect to $A = \{a, b, c, d\}$, let (p, q, r) be a geodesic directed triangle in Γ, and let s be a geodesic path in Γ from 1_M to ιp. Suppose that w_s, w_p, w_q and w_r are the words

labelling the paths s, p, q and r, respectively. By Lemma 7.1 we have $w_s w_p w_q[k] = w_s w_r[k]$ for all $k \in \{1, \ldots, |w_s w_r|\}$. From this it follows that the set of vertices visited by the path $p \circ q$ is equal to the set of vertices visited by the path r. It is then immediate that the geodesic triangle (p, q, r) is strongly 0-hyperbolic. □

Corollary 7.4. *There exists a finitely generated, right cancellative, $\mathcal{J}$-trivial strongly 0-hyperbolic monoid M which is not recursively presentable (and hence has word problem and all of Green's relations undecidable).*

Proof. Let I be a subset of the natural numbers which is not recursively enumerable, and set $M = M_I$. Then M is finitely generated by definition, right cancellative and $\mathcal{J}$-trivial by Corollary 7.2 and strongly 0-hyperbolic by Proposition 7.3. Suppose for a contradiction that M were recursively presentable. Then by enumerating a presentation and its consequences, we could enumerate all relations which hold in M. In particular we could enumerate those relations of the form $ab^i c = ab^i d$ which hold in M. But it is easily seen that such a relation holds if and only if $i \in I$, so this would allow us to enumerate I, contradicting the assumption that I is not recursively enumerable.

Since M is not recursively presentable, it does not have solvable word problem. And since M is $\mathcal{J}$-trivial, a decision process for any of Green's equivalence or pre-order relations would permit the solution of the word problem. □

References

[1] S. I. Adyan. On the embeddability of semigroups in groups. *Soviet Math. Dokl.*, 1:819–821, 1960.

[2] S. I. Adyan and G. U. Oganesyan. On the word and divisibility problems for semigroups with one relation. *Mat. Zametki*, 41:412–421, 458, 1987.

[3] G. Bauer and F. Otto. Finite complete rewriting systems and the complexity of the word problem. *Acta Inform.*, 21:521–540, 1984.

[4] J. C. Birget. Time-complexity of the word problem for semigroups and the Higman embedding theorem. *Internat. J. Algebra Comput.*, 8:235–294, 1998.

[5] J. C. Birget. The $\mathscr{R}$- and $\mathscr{L}$-orders of the Thompson-Higman monoid $M_{k,1}$ and their complexity. *Internat. J. Algebra Comput.*, 20:489–524, 2010.

[6] J. C. Birget. The Thompson-Higman monoids $M_{k,i}$: the $\mathscr{J}$-order, the $\mathscr{D}$-relation, and their complexity. *Internat. J. Algebra Comput.*, 21:1–34, 2011.

[7] J. C. Birget, A. Yu. Ol'shanskii, E. Rips and M. V. Sapir. Isoperimetric functions of groups and computational complexity of the word problem. *Ann. of Math. (2)*, 156:467–518, 2002.

[8] A. J. Cain. Decision problems for word-hyperbolic semigroups. arXiv: 1303.1763.

[9] A. J. Cain. Hyperbolicity of monoids presented by confluent monadic rewriting systems. *Beiträge zur Algebra und Geometrie*, 54:593–608, 2013.

[10] A. J. Cain and V. Maltcev. Finitely presented monoids with linear Dehn function need not have regular cross-sections. *Semigroup Forum*, 88:300–315, 2014.

[11] A. J. Cain and V. Maltcev. Context-free rewriting systems and word-hyperbolic structures with uniqueness. *Internat. J. Algebra. Comput.*, 22, 2012.

[12] C. M. Campbell, E. F. Robertson, N. Ruškuc, and R. M. Thomas. Automatic semigroups. *Theoret. Comput. Sci.*, 250(1–2):365–391, 2001.

[13] J. Cassaigne and P. V. Silva. Infinite words and confluent rewriting systems. *Internat. J. Algebra Comput.*, 19:443–490, 2009.

[14] A. Duncan and R. H. Gilman. Word hyperbolic semigroups. *Math. Proc. Cambridge Philos. Soc.*, 136(3):513–524, 2004.

[15] J. B. Fountain and M. Kambites. Hyperbolic groups and completely simple semi-groups. In *Semigroups and Languages*, pages 106–132. World Sci. Publishing, River Edge, NJ, 2004.

[16] R. H. Gilman. On the definition of word hyperbolic groups. *Math. Z.*, 242(3):529–541, 2002.

[17] R. Gray and M. Kambites. A Švarc-Milnor lemma for monoids acting by isometric embeddings. *Internat. J. Algebra Comput.*, 21:1135–1147, 2011.

[18] R. Gray and M. Kambites. Groups acting on semimetric spaces and quasi-isometries of monoids. *Trans. Amer. Math. Soc.*, 365:555–578, 2013.

[19] R. Gray and M. Kambites. Quasi-isometry and finite presentations for left cancellative monoids. *Internat. J. Algebra Comput.*, 23:1099–1114, 2013.

[20] M. Gromov. Hyperbolic groups. In *Essays in Group Theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[21] V. S. Guba and M. V. Sapir. Diagram groups and directed 2-complexes: homotopy and homology. *J. Pure Appl. Algebra*, 205(1):1–47, 2006.

[22] M. Hoffmann and R. M. Thomas. Notions of hyperbolicity in monoids. *Theoret. Comput. Sci.*, 411:799–811, 2010.

[23] S. V. Ivanov and P. E. Schupp. On the hyperbolicity of small cancellation groups and one-relator groups. *Trans. Amer. Math. Soc.*, 350:1851–1895, 1998.

[24] R. C. Lyndon and P. E. Schupp. *Combinatorial group theory*. Classics in Mathematics. Springer-Verlag, Berlin, 2001. Reprint of the 1977 edition.

[25] K. Madlener and F. Otto. Pseudo-natural algorithms for the word problem for finitely presented monoids and groups. *J. Symbolic Comput.*, 1:383–418, 1985.

[26] A. Markov. On the impossibility of certain algorithms in the theory of associative systems. *C. R. (Doklady) Acad. Sci. URSS (N.S.)*, 55:583–586, 1947.

[27] F. Otto. Some undecidability results for nonmonadic Church-Rosser Thue systems. *Theoret. Comput. Sci.*, 33:261–278, 1984.

[28] F. Otto. Some decidability and undecidability results on Green's relations for automatic monoids. *Semigroup Forum*, 75:521–536, 2007.

[29] F. Otto, A. Sattler-Klein, and K. Madlener. Automatic monoids versus monoids with finite convergent presentations. In *Rewriting techniques and applications (Tsukuba, 1998)*, volume 1379 of *Lecture Notes in Comput. Sci.*, pages 32–46. Springer, 1998.

[30] E. L. Post. Recursive unsolvability of a poblem of Thue. *J. Symbolic Logic*, 12:1–11, 1947.

[31] S. J. Pride. Geometric methods in combinatorial semigroup theory. *Semigroups, Formal Languages and Groups, Kluwer Academic Publishers, Dordrecht*, pages 215–232, 1995.

[32] M. V. Sapir, J. C. Birget and E. Rips. Isoperimetric and isodiametric functions of groups. *Ann. of Math. (2)*, 156:345–466, 2002.

[33] A. M. Turing. The word problem in semi-groups with cancellation. *Ann. of Math. (2)*, 52:491–505, 1950.