

The Complexity of Proving the Discrete Jordan Curve Theorem

PHUONG NGUYEN, McGill University
STEPHEN COOK, University of Toronto

The Jordan Curve Theorem (JCT) states that a simple closed curve divides the plane into exactly two connected regions. We formalize and prove the theorem in the context of grid graphs, under different input settings, in theories of bounded arithmetic that correspond to small complexity classes. The theory $V^0(2)$ (corresponding to $AC^0(2)$) proves that any set of edges that form disjoint cycles divides the grid into at least two regions. The theory V^0 (corresponding to AC^0) proves that any sequence of edges that form a simple closed curve divides the grid into exactly two regions. As a consequence, the Hex tautologies and the st-connectivity tautologies have polynomial size $AC^0(2)$ -Frege-proofs, which improves results of Buss which only apply to the stronger proof system TC^0 -Frege.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Proof Theory

General Terms: Theory

Additional Key Words and Phrases: Jordan Curve Theorem, bounded arithmetic, bounded reverse mathematics

ACM Reference Format:

Nguyen, P. and Cook, S. 2011. The Complexity of Proving the Discrete Jordan Curve Theorem ACM Trans. Comput. Logic V, N, Article A (January YYYY), 24 pages.
DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

1.1. Proof Complexity Background

This paper is a contribution to “Bounded Reverse Mathematics” [CN10; Ngu08], a theme whose goal is to formalize and prove discrete versions of mathematical theorems in weak theories of bounded arithmetic. (*Reverse Mathematics* is a program introduced by Friedman and Simpson (see [Sim99]) to classify mathematical theorems according to the strength of the axiomatic theories needed to prove them.) Razborov’s simplified proof (in the theory V_1^1) of Hastad’s Switching Lemma [Raz95] demonstrates the advantage of formalizing non-trivial arguments by reducing the complexity of the concepts needed in the proof, and can be regarded as an early important example of this theme. Here we are concerned with theories which capture reasoning in complexity classes in the low end of the hierarchy

$$AC^0 \subset AC^0(2) \subset TC^0 \subseteq NC^1 \subseteq L \subseteq P \quad (1)$$

The class AC^0 (problems expressible by polynomial size constant depth Boolean circuits with unbounded fanin AND gates and OR gates) can compute binary addition,

Author’s addresses: P. Nguyen, School of Computer Science, McGill University, H3A 2A7, Montréal, Québec, Canada; S. Cook, Department of Computer Science, University of Toronto, M5S 3G4, Toronto, Ontario, Canada.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 1529-3785/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

but not binary multiplication, and cannot determine the parity of the number of input 1's. The class $\text{AC}^0(2)$ strengthens AC^0 by allowing parity gates with unbounded fanin. TC^0 allows threshold gates and can compute binary multiplication, while NC^1 has the computing power of polynomial size Boolean formulas. L stands for deterministic log space, and P for polynomial time.

Our theories are “second order” (as are those in Simpson’s book [Sim99]), or more properly two-sorted first order. The first sort is the set \mathbb{N} of natural numbers, and the second sort is the set of finite subsets of \mathbb{N} . We think of a finite subset $X \subset \mathbb{N}$ as a finite bit string $X(0)X(1)\dots$, where $X(i)$ is 1 or 0 depending on whether $i \in X$. A function F on bit strings is *definable* in a two-sorted theory \mathcal{T} if its graph $Y = F(X)$ is expressible by a bounded existential formula $\varphi(X, Y)$ such that

$$\mathcal{T} \vdash \forall X \exists! Y \varphi(X, Y)$$

The complexity class associated with \mathcal{T} is given by the set of functions definable in \mathcal{T} . We have a theory for each of the complexity classes in (1), and these theories form the hierarchy [CN10]

$$\mathbf{V}^0 \subset \mathbf{V}^0(2) \subseteq \mathbf{VTC}^0 \subseteq \mathbf{VNC}^1 \subseteq \mathbf{VL} \subseteq \mathbf{TV}^0 \quad (2)$$

Our base theory is \mathbf{V}^0 , where the definable functions are those in AC^0 . Thus \mathbf{V}^0 can define $F_+(X, Y) = X + Y$ (binary addition) but \mathbf{V}^0 cannot define $\text{parity}(X)$ (the number of ones in $X \bmod 2$). The theory $\mathbf{V}^0(2)$ is associated with $\text{AC}^0(2)$, and can define $\text{parity}(X)$ but not $F_\times(X, Y) = X \cdot Y$ (binary multiplication). The theory \mathbf{VTC}^0 can define $F_\times(X, Y)$ but not any function that is not in TC^0 .

All these theories are finitely axiomatizable, and have the same finite vocabulary.

We are interested in finding the weakest theory that can prove a given universal combinatorial principle. The best known result here is due to Ajtai [Ajt88], which (stated in our terms) says that the Pigeonhole Principle $\text{PHP}(n, X)$ (asserting there is no one-one map from $\{0, 1, \dots, n\}$ to $\{0, 1, \dots, n-1\}$) is not a theorem of \mathbf{V}^0 . It is known that \mathbf{VTC}^0 proves $\text{PHP}(n, X)$, but it is an open question whether $\mathbf{V}^0(2)$ proves $\text{PHP}(n, X)$.

The study of the proof complexity of combinatorial principles is often formulated in terms of propositional proof systems, rather than theories such as (2). In fact there are propositional proof systems corresponding to each of the theories in (2), so that we have a three-way correspondence between complexity classes, theories, and proof systems as follows:

class	AC^0	$\text{AC}^0(2)$	TC^0	NC^1	P
theory	\mathbf{V}^0	$\mathbf{V}^0(2)$	\mathbf{VTC}^0	\mathbf{VNC}^1	\mathbf{TV}^0
system	$\text{AC}^0\text{-Frege}$	$\text{AC}^0(2)\text{-Frege}$	$\text{TC}^0\text{-Frege}$	Frege	eFrege

For example a Frege system is a standard Hilbert-style propositional proof system in which a formal proof is a sequence of propositional formulas which are either axioms or follow from earlier formulas from rules. In an $\text{AC}^0\text{-Frege}$ proof the formulas must have depth at most d , where d is a parameter. In an $\text{AC}^0(2)\text{-Frege}$ proof the formulas are allowed parity gates, and in a $\text{TC}^0\text{-Frege}$ proof the formulas are allowed threshold gates.

There is a simple correspondence between Σ_0^B formulas $\varphi(x, X)$ (i.e. bounded formulas in the language of the theories, with no second-order quantifiers) and a polynomial size family $\langle \varphi(x, X)[n], n \in \mathbb{N} \rangle$ of propositional formulas such that the propositional formulas are all valid iff $\forall x \forall X \varphi(x, X)$ holds in the standard model. Further, for each theory \mathcal{T} and associated proof system $S_{\mathcal{T}}$ there is a simple translation which takes

each Σ_0^B formula $\varphi(x, X)$ provable in \mathcal{T} into a polynomial size family of $S_{\mathcal{T}}$ -proofs of the tautologies $\varphi(x, X)[n]$.

For example, in the case of the Pigeonhole Principle, the Σ_0^B formula $PHP(x, X)$ translates into a family PHP_n^{n+1} of tautologies, in which the variables have the form p_{ij} , $0 \leq i \leq n, 0 \leq j < n$, and p_{ij} is intended to assert that i gets mapped to j . Ajtai [Ajt88] proved that the tautologies PHP_n^{n+1} do not have polynomial size AC^0 -Frege proofs. From this it follows that $PHP(x, X)$ is not provable in V^0 , as we mentioned earlier.

In general, the propositional proof systems can be regarded as nonuniform versions of the corresponding theories (more precisely the $\forall\Sigma_0^B$ -consequences of the theories). Showing that a given Σ_0^B formula $\varphi(x, X)$ is provable in a theory \mathcal{T} establishes that the tautology family $\varphi(x, X)[n]$ has polynomial size $S_{\mathcal{T}}$ proofs. However the converse is false in general: The tautologies $\varphi(x, X)[n]$ might have polynomial size $S_{\mathcal{T}}$ proofs even though $\varphi(x, X)$ is not provable in \mathcal{T} .

In the present paper our main results are positive and uniform: we show various principles are provable in various theories, and polynomial size upper bounds on the proof size of the corresponding tautologies follow as corollaries.

1.2. Discrete Planar Curves

We are concerned with principles related to the Jordan Curve Theorem (JCT), which asserts that a simple closed curve divides the plane into exactly two connected components. The authors were inspired by a talk by Thomas Hales [Hal05; Hal07] explaining his computer-verified proof of the theorem (involving 44,000 proof steps), which in turn is based on Thomassen's five-page proof [Tho92]. The latter proof starts by proving $K_{3,3}$ is not planar, which in turn implies the JCT.

Hales first proves the JCT for grid graphs, and this is the setting for the present paper. A grid graph has its vertices among the planar grid points $\{(i, j) \mid 0 \leq i, j \leq n\}$ and its edges among the horizontal and vertical lines connecting adjacent grid points.

Buss [Bus06] has extensive results on the propositional proof complexity of grid graphs, and nicely summarizes what was known on the subject before the present paper. The st-connectivity principle states that it is not possible to have a red path of edges and a blue path of edges which connect diagonally opposite corners of the grid graph unless the paths intersect. In this paper we focus on the following two ways of expressing this principle as a family of tautologies: the harder tautologies $STCONN(n)$ [Bus06] express the red and blue edges as two sets, with the condition that every node except the corners has degree 0 or 2 (thus allowing disjoint cycles as well as paths). The easier tautologies $STSEQ(n)$ express the paths as sequences of edges.

In 1997 Cook and Rackoff [CR97] showed that the easier tautologies $STSEQ(n)$ expressing st-connectivity have polynomial size TC^0 -Frege-proofs. Their proof is based on winding numbers. Buss [Bus06] improved this by showing that the harder tautologies $STCONN(n)$ also have polynomial size TC^0 -Frege-proofs. Buss's proof shows how the red and blue edges in each column of the grid graph determine an element of a certain finitely-generated group. The first and last columns determine different elements, but assuming the red and blue paths do not cross, adjacent columns must determine the same element. This leads to a contradiction.

The Hex tautologies, proposed by Urquhart [Urq01], assert that every completed board in the game of Hex has a winner. [Bus06] shows that the Hex Tautologies can be reduced to the hard st-connectivity tautologies $STCONN(n)$, and hence also have polynomial size TC^0 -Frege-proofs.

1.3. New Results

We work in the uniform setting, formalizing proofs of principles in the theories V^0 and $V^0(2)$, which imply upper bounds on the propositional proof complexity of the principles. In Section 3 we show that $V^0(2)$ proves the part of the discrete JCT asserting a closed curve divides the plane into at least two connected components, for the (harder) case in which the curve and paths are given as sets of edges. The proof is inspired by Buss's TC^0 -Frege proof of $STCONN(n)$ and is based on the idea that a vertical line passing through a grid curve can detect which regions are inside and outside the curve by the parity of the number of horizontal edges it intersects. It follows that $V^0(2)$ proves the st-connectivity principle for edge sets.

As a corollary we conclude that the $STCONN(n)$ tautologies and the Hex tautologies have polynomial size $AC^0(2)$ -Frege proofs, thus strengthening Buss's [Bus06] result that is stated for the stronger TC^0 -Frege system. Our result is stronger in two senses: the proof system is weaker, and we show the existence of uniform proofs by showing the st-connectivity principle is provable in $V^0(2)$. In fact, showing provability in a theory such as $V^0(2)$ is often easier than directly showing its corollary that the corresponding tautologies have polynomial size proofs. This is because we can use the fact that the theory proves the induction scheme and the minimization scheme for formulas expressing concepts in the corresponding complexity class.

In Section 4 we prove the surprising result that when the input curve and paths are presented as sequences of grid edges then even the very weak theory V^0 proves the Jordan Curve Theorem. This is the most technically interesting result in this paper. The key idea in the proof is to show (using only AC^0 -concepts) that in every column of the grid, the horizontal edges of the curve alternate between pointing right and pointing left. It follows that V^0 proves the st-connectivity principle for sequences of edges. As a corollary we conclude that the $STSEQ(n)$ tautologies have polynomial size AC^0 -Frege-proofs. This strengthens the early result [CR97] (based on winding numbers) that $STSEQ(n)$ have polynomial size TC^0 -Frege-proofs.

This is the full version of [NC07]. We have extended Section 5 substantially, and added Sections 6, 7.

2. PRELIMINARIES

The material in this section is from [Coo05; CN10; NC05].

2.1. Complexity Classes and Reductions

It will be convenient to define the relevant complexity classes AC^0 and $AC^0(2)$ in a form compatible with our theories, so we start by giving the syntax of the latter. We use a two-sorted language with variables x, y, z, \dots ranging over \mathbb{N} and variables X, Y, Z, \dots ranging over finite subsets of \mathbb{N} (interpreted as bit strings). Our basic two-sorted vocabulary \mathcal{L}_A^2 includes the usual symbols $0, 1, +, \cdot, =, \leq$ for arithmetic over \mathbb{N} , the length function $|X|$ on strings, the set membership relation \in , and string equality $=_2$ (where we usually drop mention of the subscript 2). The function $|X|$ denotes 1 plus the largest element in the set X , or 0 if X is empty (roughly the length of the corresponding string). We will use the notation $X(t)$ for $t \in X$, and we will think of $X(t)$ as the t -th bit in the string X .

Number terms are built from the constants $0, 1$, variables x, y, z, \dots , and length terms $|X|$ using $+$ and \cdot . The only *string terms* are string variables X, Y, Z, \dots . The atomic formulas are $t = u$, $X = Y$, $t \leq u$, $t \in X$ for any number terms t, u and string variables X, Y . Formulas are built from atomic formulas using \wedge, \vee, \neg and both number and string quantifiers $\exists x, \exists X, \forall x, \forall X$. Bounded number quantifiers are defined as usual,

and the bounded string quantifier $\exists X \leq t \varphi$ stands for $\exists X (|X| \leq t \wedge \varphi)$ and $\forall X \leq t \varphi$ stands for $\forall X (|X| \leq t \supset \varphi)$, where X does not occur in the term t .

Σ_0^B is the set of all \mathcal{L}_A^2 -formulas in which all number quantifiers are bounded and with no string quantifiers. Σ_1^B (corresponding to *strict* $\Sigma_1^{1,b}$ in [Kra95]) formulas begin with zero or more bounded existential string quantifiers, followed by a Σ_0^B formula. These classes are extended to Σ_i^B , $i \geq 2$, (and Π_i^B , $i \geq 0$) in the usual way.

We use the notation $\Sigma_0^B(\mathcal{L})$ to denote Σ_0^B formulas which may have symbols from the vocabulary \mathcal{L} in addition to the basic vocabulary \mathcal{L}_A^2 .

Two-sorted complexity classes contain relations $R(\vec{x}, \vec{X})$ (and possibly number-valued functions $f(\vec{x}, \vec{X})$ or string-valued functions $F(\vec{x}, \vec{X})$), where the arguments $\vec{x} = x_1, \dots, x_k$ range over \mathbb{N} , and $\vec{X} = X_1, \dots, X_\ell$ range over finite subsets of \mathbb{N} . In defining complexity classes using machines, the number arguments x_i are presented in unary notation (a string of x_i ones), and the arguments X_i are presented as bit strings. Thus the string arguments are the important inputs, and the number arguments are small auxiliary inputs useful for indexing the bits of strings.

In the uniform setting, the complexity class AC^0 has several equivalent characterizations [Imm99], including LTH (the log time hierarchy on alternating Turing machines) and FO (describable by a first-order formula using $<$ and *Bit* predicates). This motivates the following definition for the two-sorted setting (see the number/string input conventions above).

Definition 2.1. A relation $R(\vec{x}, \vec{X})$ is in AC^0 iff some alternating Turing machine accepts R in time $O(\log n)$ with a constant number of alternations.

The following result [Imm99; CN10] nicely connects AC^0 and our two-sorted \mathcal{L}_A^2 -formulas.

THEOREM 2.2 (Σ_0^B REPRESENTATION THEOREM). *A relation $R(\vec{x}, \vec{X})$ is in AC^0 iff it is represented by some Σ_0^B formula $\varphi(\vec{x}, \vec{X})$.*

In general, if \mathbf{C} is a class of relations (such as AC^0) then we want to associate a class \mathbf{FC} of functions with \mathbf{C} . Here \mathbf{FC} will contain string-valued functions $F(\vec{x}, \vec{X})$ and number-valued functions $f(\vec{x}, \vec{X})$. We require that these functions be p -bounded; i.e. for each F and f there is a polynomial $g(n)$ such that $|F(\vec{x}, \vec{X})| \leq g(\max(\vec{x}, |\vec{X}|))$ and $f(\vec{x}, \vec{X}) \leq g(\max(\vec{x}, |\vec{X}|))$.

We define the *bit graph* $B_F(i, \vec{x}, \vec{X})$ by

$$B_F(i, \vec{x}, \vec{X}) \leftrightarrow F(\vec{x}, \vec{X})(i)$$

Definition 2.3. If \mathbf{C} is a two-sorted complexity class of relations, then the corresponding functions class \mathbf{FC} consists of all p -bounded number functions whose graphs are in \mathbf{C} , together with all p -bounded string functions whose bit graphs are in \mathbf{C} .

For example, binary addition $F_+(X, Y) = X + Y$ is in FAC^0 , but binary multiplication $F_\times(X, Y) = X \cdot Y$ is not.

Definition 2.4. A string function is Σ_0^B -definable from a collection \mathcal{L} of two-sorted functions and relations if it is p -bounded and its bit graph is represented by a $\Sigma_0^B(\mathcal{L})$ formula. Similarly, a number function is Σ_0^B -definable from \mathcal{L} if it is p -bounded and its graph is represented by a $\Sigma_0^B(\mathcal{L})$ formula.

It is not hard to see that FAC^0 is closed under Σ_0^B -definability, meaning that if the bit graph of F is represented by a $\Sigma_0^B(\text{FAC}^0)$ formula, then F is already in FAC^0 .

B1. $x + 1 \neq 0$	B7. $(x \leq y \wedge y \leq x) \supset x = y$
B2. $x + 1 = y + 1 \supset x = y$	B8. $x \leq x + y$
B3. $x + 0 = x$	B9. $0 \leq x$
B4. $x + (y + 1) = (x + y) + 1$	B10. $x \leq y \vee y \leq x$
B5. $x \cdot 0 = 0$	B11. $x \leq y \leftrightarrow x < y + 1$
B6. $x \cdot (y + 1) = (x \cdot y) + x$	B12. $x \neq 0 \supset \exists y \leq x (y + 1 = x)$
L1. $X(y) \supset y < X $	L2. $y + 1 = X \supset X(y)$
SE. $[X = Y \wedge \forall i < X (X(i) \leftrightarrow Y(i))] \supset X = Y$	

Fig. 1. 2-BASIC

In order to define complexity classes such as $\text{AC}^0(2)$ and TC^0 we need to iterate Σ_0^B -definability to obtain the notion of AC^0 reduction.

Definition 2.5. We say that a string function F (resp. a number function f) is AC^0 -reducible to \mathcal{L} if there is a sequence of string functions F_1, \dots, F_n ($n \geq 0$) such that

$$F_i \text{ is } \Sigma_0^B\text{-definable from } \mathcal{L} \cup \{F_1, \dots, F_{i-1}\}, \text{ for } i = 1, \dots, n; \quad (3)$$

and F (resp. f) is Σ_0^B -definable from $\mathcal{L} \cup \{F_1, \dots, F_n\}$. A relation R is AC^0 -reducible to \mathcal{L} if there is a sequence F_1, \dots, F_n as above, and R is represented by a $\Sigma_0^B(\mathcal{L} \cup \{F_1, \dots, F_n\})$ formula.

We define the number function $\text{numones}(x, X)$ to be the number of elements of X which are less than x . We define mod_2 by

$$\text{mod}_2(x, X) = \text{numones}(x, X) \bmod 2$$

Definition 2.6. $\text{AC}^0(2)$ (resp. $\text{FAC}^0(2)$) is the class of relations (resp. functions) AC^0 -reducible to mod_2 .

We note that the classes TC^0 and FTC^0 can be defined as in the above definition from the function numones , although we will not need these classes here.

2.2. The Theories

Our base theory \mathbf{V}^0 [Coo05; CN10], called Σ_0^p -comp in [Zam96] and $I\Sigma_0^{1,b}$ (without #) in [Kra95] is associated with the complexity class AC^0 . The theory \mathbf{V}^0 uses the two-sorted vocabulary \mathcal{L}_A^2 described in Section 2.1, and is axiomatized by the set 2-BASIC given in Figure 1, together with the Σ_0^B -Comprehension scheme

$$\exists X \leq y \forall z < y (X(z) \leftrightarrow \varphi(z)),$$

where $\varphi(z)$ is any Σ_0^B formula not containing X (but may contain other free variables).

It is not hard to show that \mathbf{V}^0 proves the Σ_0^B -IND scheme

$$[\varphi(0) \wedge \forall x, \varphi(x) \supset \varphi(x + 1)] \supset \forall z \varphi(z), \quad (4)$$

where $\varphi(x)$ is any Σ_0^B -formula.

It follows from a Buss-style witnessing theorem that the Σ_1^B -definable function in \mathbf{V}^0 are precisely the functions in FAC^0 . Thus binary addition $F_+(X, Y)$ is Σ_1^B -definable in \mathbf{V}^0 but binary multiplication $F_\times(X, Y) = X \cdot Y$ is not. Simple properties of definable functions can usually be proved in \mathbf{V}^0 , including commutativity and associativity of binary addition.

The pigeonhole principle PHP_n^{n+1} can be formulated in \mathbf{V}^0 by a Σ_0^B formula $\text{PHP}(n, X)$, where $X(\langle i, j \rangle)$ asserts that pigeon i gets mapped to hole j . However it follows from Ajtai's Theorem [Ajt88] that \mathbf{V}^0 does not prove $\text{PHP}(n, X)$, nor does \mathbf{V}^0

prove the (weaker) surjective pigeonhole principle, in which it is assumed that every hole gets at least one pigeon.

It is sometimes convenient to work in the theory \bar{V}^0 , which is a universal conservative extension of V^0 with vocabulary $\mathcal{L}_{\text{FAC}^0}$ containing symbols (and defining axioms) for all FAC^0 -functions. The theory \bar{V}^0 proves the induction scheme (4), where now φ is any $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0})$ -formula.

The theory $V^0(2)$ [NC05; CN10] has the same vocabulary \mathcal{L}_A^2 as V^0 , and extends V^0 by adding the single axiom

$$\exists Y \delta_{\text{MOD}_2}(x, X, Y)$$

where

$$\delta_{\text{MOD}_2}(x, X, Y) \equiv \neg Y(0) \wedge \forall z < x, Y(z+1) \leftrightarrow (Y(z) \oplus X(z))$$

(here \oplus is exclusive or).

Note that $\delta_{\text{MOD}_2}(x, X, Y)$ defines Y as a kind of parity vector for the first x bits of X , in the sense that if $\delta_{\text{MOD}_2}(x, X, Y)$ and $z \leq x$ then

$$Y(z) \leftrightarrow \text{mod}_2(z, X) = 1$$

The Σ_1^B -definable functions of $V^0(2)$ are precisely those in $\text{FAC}^0(2)$ (see Definition 2.6).

As in the case of V^0 , it is sometimes convenient to work in the theory $\bar{V}^0(2)$, which is a universal conservative extension of $V^0(2)$ with vocabulary $\mathcal{L}_{\text{FAC}^0(2)}$ containing symbols (and defining axioms) for all $\text{FAC}^0(2)$ -functions. The theory $\bar{V}^0(2)$ proves the induction scheme (4), where now φ is any $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(2)})$ -formula.

3. INPUT AS A SET OF EDGES

We start by defining the notions of (grid) points and edges, and certain sets of edges which include closed curves, or connect grid points. All of these notions are definable by Σ_0^B -formulas, and their basic properties can be proved in V^0 .

We assume a parameter n which bounds the x and y coordinates of points on the curve in question. Thus a *grid point* (or simply a *point*) p is a pair (x, y) where $0 \leq x, y \leq n$. We use a standard pairing function $\langle x, y \rangle$ to represent a point (x, y) , where

$$\langle x, y \rangle = (x + y)(x + y + 1) + 2y$$

The x and y coordinates of a point p are denoted by $x(p)$ and $y(p)$ respectively. Thus if $p = \langle i, j \rangle$ then $x(p) = i$ and $y(p) = j$. An (undirected) *edge* is a pair (p_1, p_2) (represented by $\langle p_1, p_2 \rangle$) of adjacent points; i.e. either $|x(p_2) - x(p_1)| = 1$ and $y(p_2) = y(p_1)$, or $x(p_2) = x(p_1)$ and $|y(p_2) - y(p_1)| = 1$. For a horizontal edge e , we also write $y(e)$ for the (common) y -coordinate of its endpoints.

Let E be a set of edges (represented by a set of numbers representing those edges). The *E-degree* of a point p is the number of edges in E that are incident to p .

Definition 3.1. A *curve* is a nonempty set E of edges such that the E -degree of every grid point is either 0 or 2. A set E of edges is said to *connect* two points p_1 and p_2 if the E -degrees of p_1 and p_2 are both 1 and the E -degrees of all other grid points are either 0 or 2. Two sets E_1 and E_2 of edges are said to *intersect* if there is a grid point whose E_i -degree is ≥ 1 for $i = 1, 2$.

Note that a curve in the above sense is actually a collection of one or more disjoint closed curves. Also if E connects p_1 and p_2 then E consists of a path connecting p_1 and p_2 together with zero or more disjoint closed curves. So a curve can divide the grid into more than two connected components. To say that there are at least two

connected components we will show that there are two points off the curve that cannot be connected without crossing the curve. An easy way of picking such a pair of points is to use the following definition. We will consider points which are “close” to the curve. It suffices to consider the case in which one point is above and one point is below an edge in E . (Note that the case in which one point is to the left and one point is to the right of E can be reduced to this case by rotating the $(n + 1) \times (n + 1)$ array of all grid points by 90 degrees.)

Definition 3.2. Two points p_1, p_2 are said to be *on different sides of E* if

$$\begin{aligned} &x(p_1) = x(p_2) \wedge |y(p_1) - y(p_2)| = 2 \\ &\text{the } E\text{-degree of } p_i = 0 \text{ for } i = 1, 2 \\ &\text{the } E\text{-degree of } p = 2 \end{aligned}$$

where p is the point with $x(p) = x(p_1)$ and $y(p) = \frac{1}{2}(y(p_1) + y(p_2))$. (See Figure 2.)

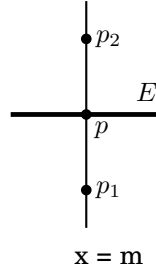


Fig. 2. p_1, p_2 are on different sides of E .

Now we show that any set of edges that forms at least one simple curve must divide the plane into at least two connected components. This is formalized in the following theorem.

THEOREM 3.3 (MAIN THEOREM FOR $V^0(2)$). *The theory $V^0(2)$ proves the following: Suppose that B is a set of edges forming a curve, p_1 and p_2 are two points on different sides of B , and that R is a set of edges that connects p_1 and p_2 . Then B and R intersect.*

3.1. The Proof of the Main Theorem for $V^0(2)$

In the following discussion we also refer to the edges in B as “blue” edges, and the edges in R as “red” edges.

We argue in $V^0(2)$, and prove the theorem by contradiction. Suppose to the contrary that B and R satisfy the hypotheses of the theorem, but do not intersect.

Notation A horizontal edge is said to be *on column k* (for $k \leq n - 1$) if its endpoints have x -coordinates k and $k + 1$.

Let $m = x(p_1) = x(p_2)$. W.l.o.g., assume that $2 \leq m \leq n - 2$. Also, we may assume that the red path comes to both p_1 and p_2 from the left, i.e., the two red edges that are incident to p_1 and p_2 are both horizontal and on column $m - 1$ (see Figure 3). (Note that if the red path does not come to both points from the left, we could fix this by effectively doubling the density of the points by doubling n to $2n$, replacing each edge in B or R by a double edge, and then extending each end of the new path by three (small) edges forming a “C” shape to end at points a distance 1 from the blue curve, approaching from the left.)

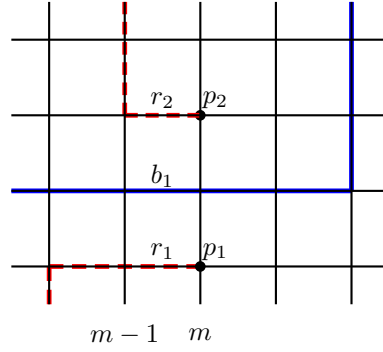


Fig. 3. The red (dashed) path must cross the blue (undashed) curve.

We say that edge e_1 *lies below* edge e_2 if e_1 and e_2 are horizontal and in the same column and $y(e_1) < y(e_2)$. For each horizontal red edge r we consider the parity of the number of horizontal blue edges b that lie below r . The following notion is definable in $V^0(2)$.

Notation An edge r is said to be an *odd edge* if it is red and horizontal and

$$\text{parity}(\{b : b \text{ is a horizontal blue edge that lies below } r\}) = 1$$

For example, it is easy to show in $V^0(2)$ that exactly one of r_1, r_2 in Figure 3 is an odd edge.

For each $k \leq n-1$, define using $\Sigma_0^B(\text{parity})$ -COMP the set

$$X_k = \{r : r \text{ is an odd edge in column } k\}$$

The Main Theorem for $V^0(2)$ follows from the lemma below as follows. We may as well assume that there are no edges in either B or R in columns 0 and $n-1$, so $\text{parity}(X_0) = \text{parity}(X_{n-1}) = 0$. On the other hand, it follows by $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0(2)})$ -IND using **b)** that $\text{parity}(X_0) = \text{parity}(X_{m-1})$ and $\text{parity}(X_m) = \text{parity}(X_{n-1})$, which contradicts **a)**.

LEMMA 3.4. *It is provable in $V^0(2)$ that*

- a)** $\text{parity}(X_{m-1}) = 1 - \text{parity}(X_m)$.
- b)** For $0 \leq k \leq n-2$, $k \neq m$, $\text{parity}(X_k) = \text{parity}(X_{k+1})$.

PROOF. First we prove **b)**. For $k \leq n-1$ and $0 \leq j \leq n$, let $e_{k,j}$ be the horizontal edge on column k with y -coordinate j . Fix $k \leq n-2$. Define the ordered lists (see Figure 4)

$$L_0 = e_{k,0}, e_{k,1}, \dots, e_{k,n}; \quad L_{n+1} = e_{k+1,0}, e_{k+1,1}, \dots, e_{k+1,n}$$

and for $1 \leq j \leq n$:

$$L_j = e_{k+1,0}, \dots, e_{k+1,j-1}, \langle (k+1, j-1), (k+1, j) \rangle, e_{k,j}, \dots, e_{k,n}$$

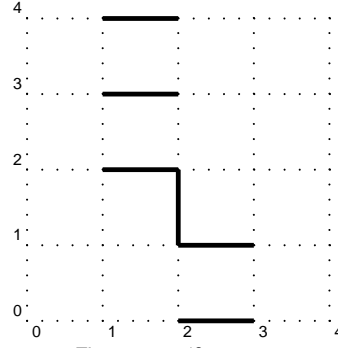
A red edge r is said to be *odd in L_j* if $r \in L_j$, and

$$\text{parity}(\{b : b \text{ is a blue edge that precedes } r \text{ in } L_j\}) = 1$$

(In particular, X_k and X_{k+1} consist of odd edges in L_0 and L_{n+1} , respectively.) For $0 \leq j \leq n+1$, let

$$Y_j = \{r : r \text{ is an odd edge in } L_j\}$$

Thus $Y_0 = X_k$ and $Y_{n+1} = X_{k+1}$.

Fig. 4. L_2 (for $n = 4, k = 1$).

Claim: If $k \neq m - 1$ then

$$\text{parity}(Y_j) = \text{parity}(Y_{j+1})$$

for $j \leq n$.

This is because the symmetric difference of Y_j and Y_{j+1} has either no red edges, or two red edges with the same parity.

Thus by $\Sigma_0^B(\mathcal{L}_{\text{FAC}^0}(2))$ -IND on j we have $\text{parity}(Y_0) = \text{parity}(Y_{n+1})$, and hence $\text{parity}(X_k) = \text{parity}(X_{k+1})$.

The proof of **a**) is similar. The only change here is that $\text{parity}(L_j)$ and $\text{parity}(L_{j+1})$ must differ for exactly one value of j : either $j = y(p_1) - 1$ or $j = y(p_2) - 1$. \square

4. INPUT AS A SEQUENCE OF EDGES

Now suppose that B is a sequence of edges

$$\langle q_0, q_1 \rangle, \langle q_1, q_2 \rangle, \dots, \langle q_{t-2}, q_{t-1} \rangle, \langle q_{t-1}, q_0 \rangle$$

that form a single closed curve (i.e., $t \geq 4$ and q_0, \dots, q_{t-1} are distinct). In this section we will show that the weak base theory V^0 proves two theorems that together imply the Jordan Curve Theorem for grid graphs: The curve B divides the grid into exactly two connected regions. Theorem 4.1 is the analog of Theorem 3.3 (Main Theorem for $V^0(2)$), and states that a sequence of edges forming a path connecting points p_1 and p_2 on different sides of the curve must intersect the curve. Theorem 4.9 states that any point p in the grid off the curve can be connected by a path (in a refined grid) that does not intersect the curve, and leads from p to one of the points p_1 or p_2 .

There is no analog in Section 3 to the last theorem because in that setting it would be false: the definition of a *curve* as a set of edges allows multiple disjoint curves.

4.1. There Are at Least Two Regions

THEOREM 4.1 (MAIN THEOREM FOR V^0). *The theory V^0 proves the following: Let B be a sequence of edges that form a closed curve, and let p_1, p_2 be any two points on different sides of B . Suppose that R is a sequence of edges that connect p_1, p_2 . Then R and B intersect.*

(See Definition 3.1 to explain the notion of points p_1, p_2 being on different sides of a curve.)

We use the fact that the edges B can be directed (i.e., from q_i to q_{i+1}). The Main Theorem follows easily from the Edge Alternation Theorem 4.3, which states that the

horizontal edges on each column m of a closed curve must alternate between pointing right and pointing left.

4.1.1. Alternating edges and proof of the Main Theorem. We start by defining the notion of *alternating sets*, which is fundamental to the proof of the Main Theorem for V^0 . Two sets X and Y of numbers are said to alternate if their elements are interleaved, in the following sense.

Definition 4.2. Two disjoint sets X, Y *alternate* if between every two elements of X there is an element of Y , and between every two elements of Y there is an element of X . These conditions are defined by the following Σ_0^B formulas:

- (i). $\forall x_1, x_2 \in X (x_1 < x_2 \supset \exists y \in Y, x_1 < y < x_2)$,
- (ii). $\forall y_1, y_2 \in Y (y_1 < y_2 \supset \exists x \in X, y_1 < x < y_2)$

The Main Theorem follows easily from the following result.

THEOREM 4.3 (EDGE ALTERNATION THEOREM). (*Provable in V^0*) *Let P be a sequence of edges that form a closed curve. For each column m , let A_m be the set of y -coordinates of left-pointing edges of P on the column, and let B_m be the set of y -coordinates of right-pointing edges of P on the column. Then A_m and B_m alternate.*

The proof of this theorem is rather complicated. Note that a violation of the theorem may occur on any column, so a simple proof by induction on the columns is not possible. Here we will consider segments of the curve that lie entirely (except for its two endpoints) to the left of a vertical line, say, $x = m$ (see Definition 4.5). It suffices to show that for any fixed m the endpoints of all such segments alternate (see the Edge Alternation Lemma 4.6). For this we will need to establish a number of properties of these segments. The theorem will be proved in Subsection 4.1.4, after presenting necessary concepts and lemmas in Subsections 4.1.2 and 4.1.3. First, we show that the Main Theorem can be obtained from the above theorem.

PROOF OF THE MAIN THEOREM 4.1 FROM THE EDGE ALTERNATION THEOREM. The proof is by contradiction. Assume that R does not intersect B . We construct a sequence of edges P from B and R that form a closed curve, but that violate the Edge Alternation Theorem.

Without loss of generality, assume that p_1, p_2 and B, R are as in Figure 3. Also, suppose that the sequence R starts from p_1 and ends in p_2 . We may assume that the edge b_1 is from right to left (otherwise reverse the curve). Assume that the point $\langle x(p_1) + 1, y(p_1) \rangle$ is not on B or R . (This can be achieved by doubling the density of the grid.)

We merge B and R into a sequence of edges as in Figure 5. Let P be the resulting sequence of edges. Then P is a closed curve. However, the edges r_1 and b_1 have the same direction, and thus violate the Edge Alternation Theorem. \square

4.1.2. Bijections between alternating sets. Suppose that X and Y alternate and $f : X \rightarrow Y$ is a bijection from X to Y . Let $x_1, x_2 \in X$, $x_1 < x_2$, and suppose that neither $f(x_1)$ nor $f(x_2)$ lies between x_1 and x_2 . Since the open interval (x_1, x_2) contains more elements of Y than X , it must contain an image $f(z)$ of some $z \in X$ where either $z < x_1$ or $z > x_2$.

The above property can be formalized and proved in the theory VTC^0 , where f is given by its graph: a finite set of ordered pairs. However, it is not provable in V^0 , because it implies the surjective Pigeonhole Principle, which is not provable in V^0 [CN10]. Nevertheless it is provable in V^0 under the assumption that f satisfies the condition that connecting each x to its image $f(x)$ by an arc above the line \mathbb{N} does not

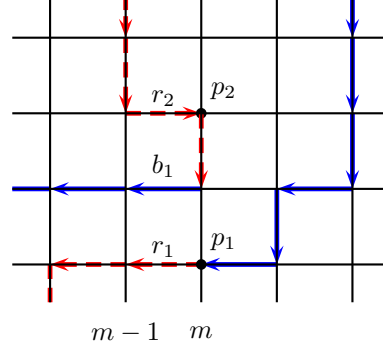
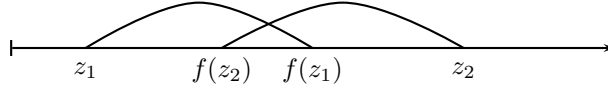


Fig. 5. Merging the red (dashed) path and the blue (undashed) curve.

create any “crossings”, i.e.

$$\begin{aligned} &\text{the sets } \{z_1, f(z_1)\} \text{ and } \{z_2, f(z_2)\} \text{ are not alternating,} \\ &\text{for all } z_1, z_2 \in X, z_1 \neq z_2. \end{aligned} \quad (5)$$

(See Figure 6).

Fig. 6. f violates (5)

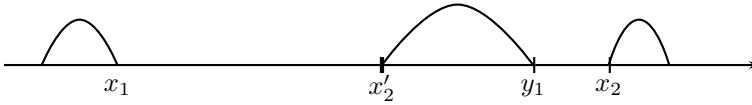
We need the following result to prove the Edge Alternation Theorem.

LEMMA 4.4 (ALTERNATION LEMMA). *(Provable in V^0) Suppose that X and Y alternate and that f (given by a finite set of ordered pairs) is a bijection between X and Y that satisfies (5). Let $x_1, x_2 \in X$ be such that $x_1 < x_2$ and neither $f(x_1)$ nor $f(x_2)$ is in the interval (x_1, x_2) . Then,*

$$\exists z \in X, (z < x_1 \vee z > x_2) \wedge x_1 < f(z) < x_2 \quad (6)$$

PROOF. We prove by contradiction, using the number minimization principle. Let x_1, x_2 be a counter example with the least difference $x_2 - x_1$.

Let $y_1 = \max(\{y \in Y : y < x_2\})$. We have $x_1 < y_1 < x_2$. Let x'_2 be the pre-image of y_1 : $f(x'_2) = y_1$. By our assumption that (6) is false, $x_1 < x'_2 < x_2$. In addition, since $y_1 = \max(\{y \in Y : y < x_2\})$ and X, Y alternate, we have $x_1 < x'_2 < y_1$. (See Figure 7.)

Fig. 7. $f(x_1), f(x_2) \notin (x_1, x_2)$, and $f(x'_2) = y_1$.

Now by (5), for all $z \in X$, $x'_2 < z < y_1$ implies that $x'_2 < f(z) < y_1$. Hence the pair x_1, x'_2 is another counter example, and $x'_2 - x_1 < x_2 - x_1$, contradicts our choice of x_1, x_2 . \square

4.1.3. *Alternating endpoints of curve segments.* For the remainder of Section 4, P denotes a sequence of edges

$$\langle p_0, p_1 \rangle, \langle p_1, p_2 \rangle, \dots, \langle p_{t-2}, p_{t-1} \rangle, \langle p_{t-1}, p_0 \rangle$$

that form a single closed curve (i.e., $t \geq 4$ and p_0, \dots, p_{t-1} are distinct).

For convenience, we assume that P has a point on the first vertical line ($x = 0$) and a point on the last vertical line ($x = n$). To avoid wrapping around the last index, we pick some vertical edge on the line ($x = n$) and define p_0 to be the forward end of this edge. In other words, the edge $\langle p_{t-1}, p_0 \rangle$ lies on the line ($x = n$).

It is easy to prove in V^0 that for every m , $0 \leq m \leq n$, P must have a point on the vertical line ($x = m$). For otherwise there is a largest $m < n$ such that the line ($x = m$) has no point on P , and we obtain a contradiction by considering the edge $\langle p_{i-1}, p_i \rangle$, where i is the smallest number such that $x(p_i) \leq m$.

For $a < b < t$, let $P_{[a,b]}$ be the oriented segment of P that contains the points p_a, p_{a+1}, \dots, p_b , and let $P_{[a,a]} = \{p_a\}$. We are interested in the segments $P_{[a,b]}$ where $x(p_a) = x(p_b)$.

The next Definition is useful in identifying segments of P that are “examined” as we scan the curve from left to right. See Figure 8 for examples.

Definition 4.5. A segment $P_{[a,b]}$ is said to *stick to* the vertical line ($x = m$) if $x(p_a) = x(p_b) = m$, and for $a < c < b$, $x(p_c) \leq m$. A segment $P_{[a,b]}$ that sticks to ($x = m$) is said to be *minimal* if $b - a > 1$, and for $a < c < b$ we have $x(p_c) < m$. Finally, $P_{[a,b]}$ is said to be *entirely on* ($x = m$) if $x(p_c) = m$, for $a \leq c \leq b$.

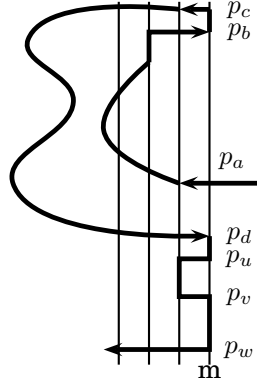


Fig. 8. The segments $P_{[a,b]}, P_{[a,c]}, \dots, P_{[u,w]}, P_{[v,w]}$ all stick to the vertical line ($x = m$). Among these, $P_{[a,b]}, P_{[c,d]}$ and $P_{[u,v]}$ are minimal, while $P_{[b,c]}, P_{[d,u]}$ and $P_{[v,w]}$ are entirely on ($x = m$).

Notice that minimal segments that stick to a vertical line ($x = m$) are disjoint. Also, if $P_{[a,b]}$ is a minimal segment that sticks to ($x = m$), then the first and the last edges of the segments must be horizontal edges in column $m - 1$, i.e., $y(p_a) = y(p_{a+1})$ and $y(p_b) = y(p_{b-1})$. In fact, the left-pointing horizontal edges in column $m - 1$ are precisely those of the form $\langle p_a, p_{a+1} \rangle$ for some minimal segment $P_{[a,b]}$ that sticks to the vertical line ($x = m$), and the right-pointing horizontal edges in column $m - 1$ are precisely those of the form $\langle p_{b-1}, p_b \rangle$ for some such minimal segment $P_{[a,b]}$.

These facts are provable in V^0 , and show that the Edge Alternation Theorem 4.3 is equivalent to the following lemma (see Figure 9). Here (and elsewhere) the assertion

that two sets of points on a vertical line alternate means that the two corresponding sets of y -coordinates alternate.

LEMMA 4.6 (EDGE ALTERNATION LEMMA). *(Provable in V^0) Let $P_{[a_1, b_1]}, \dots, P_{[a_k, b_k]}$ be all minimal segments that stick to the vertical line $(x = m)$. Then the sets $\{p_{a_1}, \dots, p_{a_k}\}$ and $\{p_{b_1}, \dots, p_{b_k}\}$ alternate.*

Note that although in V^0 we can define the set of all segments $P_{[a_i, b_i]}$ in the lemma above, we are not able to define k , the total number of such segments. Thus the index k is used only for readability.

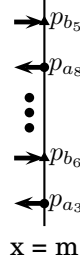


Fig. 9. The end-edges of minimal segments that stick to $(x = m)$ alternate.

Before proving the Edge Alternation Lemma we give two important lemmas needed for the proof. The first of these states that the endpoints of two non-overlapping segments of P that stick to the same vertical line do not alternate on the vertical line.

LEMMA 4.7 (MAIN LEMMA). *(Provable in V^0) Suppose that $a < b < c < d$ and that the segments $P_{[a, b]}$ and $P_{[c, d]}$ both stick to $(x = m)$. Then the sets $\{y(p_a), y(p_b)\}$ and $\{y(p_c), y(p_d)\}$ do not alternate.*

PROOF. We argue in V^0 using induction on m . The base case ($m = 0$) is straightforward: both $P_{[a, b]}$ and $P_{[c, d]}$ must be entirely on $(x = 0)$. For the induction step, suppose that the lemma is true for some $m \geq 0$. We prove it for $m + 1$ by contradiction.

Assume that there are disjoint segments $P_{[a, b]}$ and $P_{[c, d]}$ sticking to $(x = m + 1)$ that violate the lemma. Take such segments with smallest total length $(b - a) + (d - c)$. It is easy to check that both $P_{[a, b]}$ and $P_{[c, d]}$ must be minimal segments.

Now the segments $P_{[a+1, b-1]}$ and $P_{[c+1, d-1]}$ stick to the vertical line $(x = m)$, and their endpoints have the same y -coordinates as the endpoints of $P_{[a, b]}$ and $P_{[c, d]}$. Hence we get a contradiction from the induction hypothesis. \square

From the Main Lemma we can prove an important special case of the Edge Alternation Lemma.

LEMMA 4.8. *(Provable in V^0) Let $P_{[a, b]}$ be a segment that sticks to $(x = m)$, and let $P_{[a_1, b_1]}, \dots, P_{[a_k, b_k]}$ be all minimal subsegments of $P_{[a, b]}$ that stick to $(x = m)$, where $a \leq a_1 < b_1 < \dots < a_k < b_k \leq b$. Then the sets $\{p_{a_1}, \dots, p_{a_k}\}$ and $\{p_{b_1}, \dots, p_{b_k}\}$ alternate.*

PROOF. We show that between any two p_{a_i} 's there is a p_{b_j} . The reverse condition is proved similarly. Thus let $i \neq j$ be such that $y(p_{a_i}) < y(p_{a_j})$. We show that there is some ℓ such that $y(p_{a_i}) < y(p_{b_\ell}) < y(p_{a_j})$.

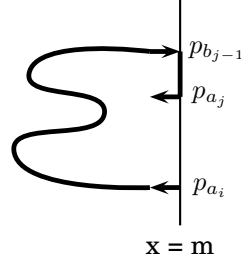


Fig. 10. Proof of Lemma 4.8

Consider the case where $i < j$ (the other case is similar). Then the segment $P_{[b_{j-1}, a_j]}$ is entirely on $(x = m)$. Now if $y(p_{b_{j-1}}) < y(p_{a_j})$, then $y(p_{a_i}) < y(p_{b_{j-1}})$, and we are done. Thus, suppose that $y(p_{b_{j-1}}) > y(p_{a_j})$ (see Figure 10).

From the Main Lemma for the segments $P_{[a_i, b_{j-1}]}$ and $P_{[a_j, b_j]}$ it follows that $y(p_{a_i}) < y(p_{b_j}) < y(p_{b_{j-1}})$. Since $P_{[b_{j-1}, a_j]}$ is entirely on $(x = m)$, it must be the case that $y(p_{a_i}) < y(p_{b_j}) < y(p_{a_j})$. \square

4.1.4. Proof of the Edge Alternation Theorem. To prove Theorem 4.3 it suffices to prove the Edge Alternation Lemma 4.6. The proof relies on Lemma 4.8, the Main Lemma, and the Alternation Lemma 4.4.

PROOF OF LEMMA 4.6. We argue in V^0 and use downward induction on m . The base case, $m = n$, follows from Lemma 4.8, where the segment $P_{[a, b]}$ has $a = 0$ and $b = t - 1$. (Recall our numbering convention that the edge $\langle p_{t-1}, p_0 \rangle$ lies on the vertical line $(x = n)$.)

For the induction step, suppose that the conclusion is true for $m + 1$, we prove it for m by contradiction.

Let $\{P_{[a'_1, b'_1]}, \dots, P_{[a'_k, b'_k]}\}$ be the definable set of all minimal segments that stick to the line $(x = m + 1)$. (k is not definable in V^0 , we use it only for readability.)

Notation Let $a_\ell = (a'_\ell + 1)$, $b_\ell = (b'_\ell - 1)$ and $A = \{y(p_{a_\ell})\}$, $B = \{y(p_{b_\ell})\}$.

Then, since

$$y(p_{a_\ell}) = y(p_{a'_\ell}) \quad \text{and} \quad y(p_{b_\ell}) = y(p_{b'_\ell}),$$

it follows from the induction hypothesis that A and B alternate. (Note that each $P_{[a_\ell, b_\ell]}$ sticks to $(x = m)$, but might not be minimal.)

Now suppose that there are horizontal P -edges e_1 and e_2 on column $m - 1$ that violate the lemma, with $y(e_1) < y(e_2)$. Thus both e_1 and e_2 point in the same direction, and there is no horizontal P -edge e on column $(m - 1)$ with $y(e_1) < y(e) < y(e_2)$. We may assume that both e_1 and e_2 point to the left. The case in which they both point to the right can be argued by symmetry (or we could strengthen the induction hypothesis to apply to both of the curves P and the reverse of P).

Let the right endpoints of e_1 and e_2 be p_c and p_d , respectively. Thus $x(p_c) = x(p_d) = m$ and $y(p_c) < y(p_d)$.

Let $P_{[a_i, b_i]}$ be the segment of P containing p_c , and let $P_{[a_j, b_j]}$ be the segment of P containing p_d . Note that the segments $P_{[a_i, b_i]}$ and $P_{[a_j, b_j]}$ stick to $(x = m)$, but they are not necessarily minimal. It follows from Lemma 4.8 that $i \neq j$.

We may assume that p_{a_j} lies above p_c . This is because if p_{a_j} lies below p_c , then we claim that p_{a_i} lies below p_d (since otherwise the segments $P_{[a_i, c]}$ and $P_{[a_j, d]}$ would

violate the Main Lemma). Thus the case p_{a_j} lies below p_c would follow by the case we consider, by interchanging the roles of a_i, c with a_j, d , and inverting the graph.

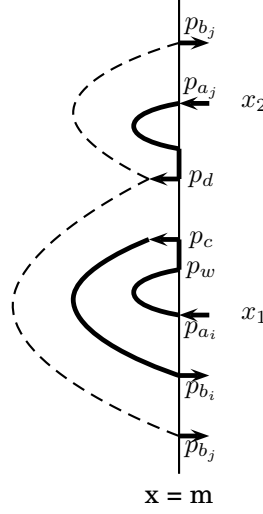


Fig. 11. Case I: $y(p_{a_i}) < y(p_d)$

Case I: $y(p_{a_i}) < y(p_d)$ (See Figure 11)

We apply the Alternation Lemma 4.4 for the alternating sets A and B with the bijection $f(y(p_{a_\ell})) = y(p_{b_\ell})$ and $x_1 = y(p_{a_i})$ and $x_2 = y(p_{a_j})$. Note that f satisfies the non-arc-crossing condition (5) by the Main Lemma.

We claim that both $f(x_1)$ and $f(x_2)$ are outside the interval $[x_1, x_2]$. We show this for $f(x_1)$; the argument for $f(x_2)$ is similar. Thus we are to show that the point p_{b_i} does not lie on the vertical line ($x = m$) between the points p_{a_i} and p_{a_j} .

First we show p_{b_i} does not lie between p_{a_i} and p_c . This is obvious if the segment $P_{[a_i, c]}$ lies entirely on ($x = m$). Otherwise let $w < c$ be such that the segment $P_{[w, c]}$ lies entirely on $x = m$. (Note that $y(p_{a_i}) < y(p_w) < y(p_c)$, because there is no horizontal edge in column $m - 1$ between p_c and p_d .) Then p_{b_i} does not lie between p_{a_i} and p_w by the Main Lemma applied to the segments $P_{[a_i, w]}$ and $P_{[c, b_i]}$.

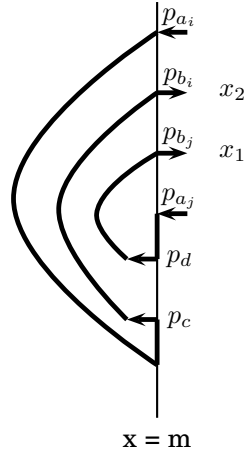
Next, note that p_{b_i} does not lie between p_c and p_d , because there is no horizontal edge in column $m - 1$ between these two points. Finally we claim that p_{b_i} does not lie between p_d and p_{a_j} . This is obvious if $a_j = d$, and otherwise use the Main Lemma applied to the segments $P_{[a_j, d]}$ and $P_{[a_i, b_i]}$.

This establishes the hypotheses for the Alternation Lemma. By that Lemma it follows that there must be some p_{a_ℓ} outside the vertical interval between p_{a_i} and p_{a_j} such that p_{b_ℓ} lies in that interval. But this is impossible, by applying the Main Lemma as above. This contradiction shows that Case I is impossible.

Case II: $y(p_{a_i}) > y(p_d)$ (See Figure 12)

In this case we must have $y(p_{a_i}) > y(p_{a_j})$, by the Main Lemma applied to the segments $P_{[a_i, c]}$ and $P_{[a_j, d]}$. In fact, by repeated use of the Main Lemma we can show

$$y(p_{a_j}) < y(p_{b_j}) < y(p_{b_i}) < y(p_{a_i})$$

Fig. 12. Case II: $y(p_{a_i}) > y(p_d)$

We get a contradiction by applying the Alternation Lemma, this time using the inverse bijection $f^{-1} : B \rightarrow A$, with $x_1 = y(p_{b_j})$ and $x_2 = y(p_{b_i})$. \square

4.2. There Are at Most Two Regions

Here we formalize and prove the idea that if P is a sequence of edges that form a closed curve, and p_1 and p_2 are points on opposite sides of P , then any point in the plane off P can be connected to either p_1 or p_2 by a path that does not intersect P . However this path must use points in a refined grid, in order not to get trapped in a region such as that depicted in Figure 13. Thus we triple the density of the points by tripling n to $3n$, and replace each edge in P by a triple of edges. We also assume that originally the curve P has no point on the border of the grid. (This assumption is different from our convention stated in Section 4.1.3.)

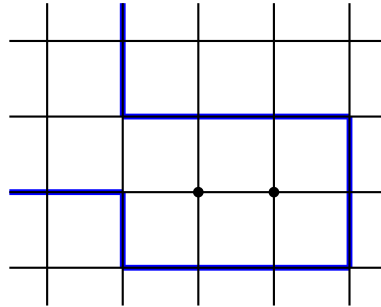


Fig. 13. An “unwanted” region with two points.

Let P' denote the resulting set of edges. Note that the new grid has size $(3n) \times (3n)$.

THEOREM 4.9. *The theory V^0 proves the following: Let P be a sequence of edges that form a closed curve, and suppose that P has no point on the border of the grid. Let P' be the corresponding sequence of edges in the $(3n) \times (3n)$ grid, as above. Let p_1, p_2*

be any two points on different sides of P' (Definition 3.2). Then any point p (on the new grid) can be connected to either p_1 or p_2 by a sequence of edges that does not intersect P' .

PROOF. Since edges in P' are directed it makes sense to speak of edges a distance 1 to the left of P' and a distance 1 to the right of P' . Thus, taking care when P' turns corners, it is straightforward to define (using Σ_0^B -COMP) two sequences Q_1, Q_2 of edges on either side of P' , i.e., both Q_1 and Q_2 have distance 1 (on the new grid) to P' . Then p_1 and p_2 must lie on Q_1 or Q_2 . By the Main Theorem for V^0 , p_1 and p_2 cannot be on the same Q_i . So assume w.l.o.g. that p_1 is on Q_1 and p_2 is on Q_2 .

We describe informally a procedure that gives a sequence of edges connecting any point p to p_1 or p_2 . First we compute (using the number minimization principle) the Manhattan distances $d(p, Q_1)$ and $d(p, Q_2)$ respectively from p to Q_1 and Q_2 . Suppose w.o.l.g. that

$$d(p, Q_1) \leq d(p, Q_2)$$

Let q be a point on Q_1 so that $d(p, q) = d(p, Q_1)$. Then any shortest sequence of edges that connect p and q does not intersect P' . Concatenate one such sequence and the sequence of edges on Q_1 that connect q and p_1 , we have a sequence of edges that connects p and p_1 without intersecting P' . \square

5. EQUIVALENCE TO THE ST-CONNECTIVITY PRINCIPLE

The st-connectivity principle states that it is not possible to have a red path and a blue path which connect diagonally opposite corners of the grid graph unless the paths intersect. Here we show that over V^0 this principle is equivalent to the discrete Jordan Curve Theorem. As a result, the set-of-edges version of this principle is provable in $V^0(2)$, and the sequence-of-edges version is provable in V^0 . First we consider the set-of-edges setting.

THEOREM 5.1. *The theory V^0 proves that the following are equivalent:*

- (a) *Suppose that B is a set of edges forming a curve, p_1 and p_2 are two points on different sides of B , and that R is a set of edges that connects p_1 and p_2 . Then B and R intersect.*
- (b) *Suppose that B is a set of edges that connects $\langle 0, n \rangle$ and $\langle n, 0 \rangle$, and R is a set of edges that connects $\langle 0, 0 \rangle$ and $\langle n, n \rangle$. Then B and R intersect.*

PROOF. First we show that (a) implies (b). Let B and R be sets as in (b).

We extend the grid to size $(n+2) \times (n+2)$ as in Figure 14. Although the y -coordinates now range from -1 to $n+1$, this can be easily fixed and we omit the details.

Then we turn B into a closed curve B' by adding the following blue edges that connect $\langle 0, n \rangle$ and $\langle n, 0 \rangle$:

$$\begin{aligned} &(\langle 0, n \rangle, \langle 0, n+1 \rangle), \\ &(\langle i, n+1 \rangle, \langle i+1, n+1 \rangle) \text{ for } 0 \leq i \leq n+1, \\ &(\langle n+2, j+1 \rangle, \langle n+2, j \rangle) \text{ for } 0 \leq j \leq n, \\ &(\langle n+1, 0 \rangle, \langle n, 0 \rangle), (\langle n+2, 0 \rangle, \langle n+1, 0 \rangle) \end{aligned}$$

Similarly we turn R into a red path R' that connects $p_1 = \langle n+1, -1 \rangle$ to $p_2 = \langle n+1, 1 \rangle$ by adding the following red edges:

$$(\langle 0, -1 \rangle, \langle 0, 0 \rangle), \text{ and } (\langle i+1, -1 \rangle, \langle i, -1 \rangle) \text{ for } 0 \leq i \leq n$$

and

$$(\langle n, n \rangle, \langle n+1, n \rangle), \text{ and } (\langle n+1, i+1 \rangle, \langle n+1, i \rangle) \text{ for } 1 \leq i \leq n-1$$

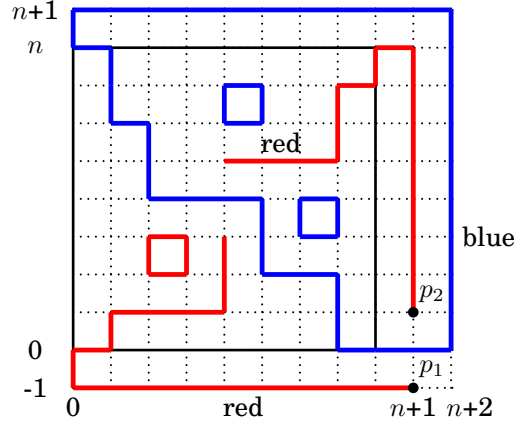


Fig. 14. Reduction from st-connectivity to discrete Jordan Curve Theorem

By (a) B' and R' intersect. The newly added paths are outside the original grid and clearly do not intersect each other, so it follows that B and R intersect.

Now we prove (a) from (b). Basically we have to turn a red path connecting two points p_1, p_2 into a red path that connects two opposite corners of the grid, and a blue curve into a blue path that connects the other two corners. This turns out to be nontrivial; the following construction is by Neil Thapen.

Suppose that we remove p from the blue curve. Then we have a blue path that connects the two neighbors p_3 and p_4 of p . Using the st-connectivity principle we need to show that this blue path intersects the red path. Intuitively, what we want is a “continuous” mapping that sends p_1, p_2 to two opposite corners of the grid, and p_3, p_4 to the other two. On the real plane \mathbb{R}^2 this suggests a mapping like $(x, y) \mapsto (\frac{x}{x^2+y^2}, \frac{y}{x^2+y^2})$. The mapping given below is based on reflection and is simpler than this suggestion, but the underlying idea is the same.¹

Let B and R be sets as in (a). Suppose for a contradiction that B and R do not intersect. By extending the grid if necessary, we can assume that the midpoint p of p_1 and p_2 (as in Figure 2) is the center of the $n \times n$ grid (see Figure 15).

Using B and R our goal is to construct sets B' and R' that form nonintersecting paths which connect opposite corners of a $2n \times 2n$ grid. This violates (b) and we are done.

We will informally describe the sets B' and R' ; formal definitions are straightforward and are left to the reader. Consider the four triangular quarters of the original grid which are determined by the two diagonals. Take the image of each triangle by reflection through its grid edge base. The results, together with the original grid, form a $\sqrt{2}n \times \sqrt{2}n$ square whose four vertices q_0, q_1, q_2, q_3 are reflection images of the center p through the edges of the original grid (see Figure 16). The $2n \times 2n$ grid is determined by the appropriate vertical and horizontal lines that go through q_0, q_1, q_2, q_3 . (So q_0, q_1, q_2, q_3 will be the midpoints of the edges of the $2n \times 2n$ grid.)

The image of the red path R are disconnected segments that lie inside the square $q_0q_1q_2q_3$ but outside the original $n \times n$ grid. It is easy to add vertical and horizontal lines to connect these segments. For example, consider a point r where R cuts a diagonal as

¹The referee points out another way of viewing this reduction as follows. View the grid as lying on a sphere. Now puncture the sphere at p , and consider the embedding of this punctured sphere into the plane. This embedding has the same spirit as the mapping described above.

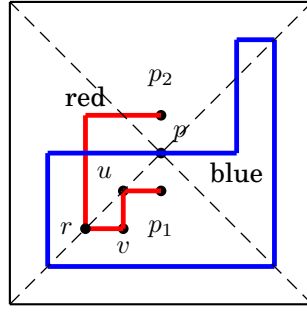
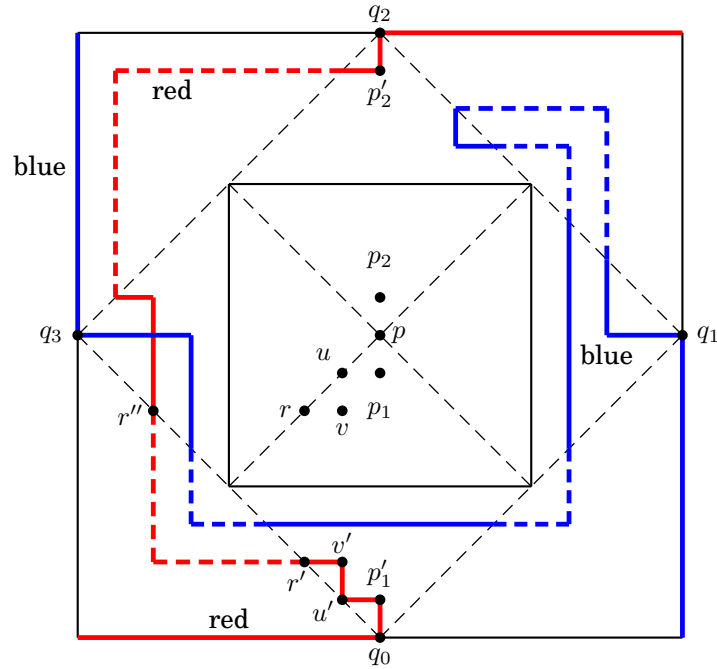
Fig. 15. Sets B and R as in discrete Jordan Curve Theorem

Fig. 16. Reduction from discrete Jordan Curve Theorem to st-connectivity

in Figure 15. Its two images r', r'' can be connected by two dashed red lines as drawn in in Figure 16. As a result, we obtain a red path that connects the images p_1', p_2' of p_1, p_2 . This red path is in turn easily extended to a path that connects the lower-left and upper-right corners of the $2n \times 2n$ grid as shown in Figure 16.

Similarly, the image of B can be turned into a blue path B' connecting the upper-left and lower-right corners of the $2n \times 2n$ grid. Given that B and R do not intersect, it can be verified that B' and R' do not intersect, and this completes our proof. \square

The next theorem is for the sequence-of-edges setting.

THEOREM 5.2. *The theory V^0 proves that the following are equivalent:*

- (a) *Suppose that B is a sequence of edges forming a curve, p_1 and p_2 are two points on different sides of B , and that R is a sequence of edges that connects p_1 and p_2 . Then B and R intersect.*

(b) Suppose that B is a sequence of edges that connects $\langle 0, n \rangle$ and $\langle n, 0 \rangle$, and R is a sequence of edges that connects $\langle 0, 0 \rangle$ and $\langle n, n \rangle$. Then B and R intersect.

PROOF. This theorem is proved similarly to the previous theorem. However, here the reductions have to output sequences of edges, as opposed to just sets of edges. In other words, given j , we need to specify the j -th edge on the paths/curves produced by the reductions.

For the direction (a) \implies (b) we can essentially use the same reduction given in the proof of Theorem 5.1 (see Figure 14). Given sequences B and R for the blue path from $\langle 0, n \rangle$ to $\langle n, 0 \rangle$ and the red path from $\langle 0, 0 \rangle$ to $\langle n, n \rangle$, it is straightforward to define the new sequences of edges for the curve B' and path R' described in the first part of the previous proof.

The proof of (b) \implies (a) is a bit more involved than before. Consider the reduction depicted in Figure 16 and let R' be the red path from p'_1 to p'_2 . To specify the sequence of red edges on R' an immediate problem is to compute its length, and this requires computing the total length of all dashed red lines. In general, such computation is not in AC^0 and hence not formalizable in V^0 .

To get around this problem, the idea is to refine the grid so as to make R' exactly $16n^2$ times longer than R , the original red path from p_1 to p_2 . (Similarly for the new blue path B' that connects q_1 and q_3 .) Thus, let path R be the sequence e_0, e_1, \dots, e_k . We will refine the grid so that each edge e_i gives rise to precisely $16n^2$ red edges

$$e'_{16n^2i}, e'_{16n^2i+1}, \dots, e'_{16n^2(i+1)-1}$$

on R' . As a result, for any j the j -th edge e'_j will be easily specified by looking at $e_{\lfloor j/16n^2 \rfloor}$.

We will distinguish between two kinds of undashed edges on R' . The first kind, called “outward edges”, consists of those that are followed by (two) dashed lines, for example (v', r') in Figure 16. All other undashed edges on R' are called “inward edges”. We will turn every inward edge into a path of length $16n^2$, and every outward edge, together with the dashed lines immediately following it, into a path of length $16n^2$.

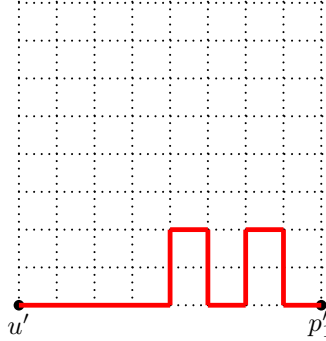


Fig. 17. Inward edge (p'_1, u') is turned into a path of length $16n^2$ (here $n = 1$)

To this end we will refine the grid $8n$ times (thus each unit square becomes an $8n \times 8n$ square). As a result, each grid edge becomes a path of length $8n$. For each inward edge we further lengthen this path by making its first half travel inside one quarter of the $8n \times 8n$ square (say to its right). For example, in Figure 17 we travel north then south $2n$ times, traversing $(4n - 2)$ edges each time, so in effect we add to the red path $4n$

vertical segments of length $(4n - 2)$ each. Consequently we obtain a path of length $16n^2$.

Now consider an outward edge and the two dashed lines that immediately follow it. Without loss of generality consider the edge (v', r') and the path (r', r'') in Figure 16. From the coordinates of r' we can compute their total length, which is of the form $(2\ell + 1)$ for some $1 \leq \ell < n$. After refinement these become a path of length $(2\ell + 1)8n$. To increase the length of this path to $16n^2$ we increase the length of the $8n$ -edge path (v', r') by $4n(4n - 4\ell - 2)$ by making its first half travel inside one quarter of the $8n \times 8n$ square to the right of (v', r') as above. Here we also go north then south $2n$ times, but now each north-south path is of length $4n - 4\ell - 2$.

Similarly we turn B' into a path of length exactly $16n^2$ times the length of B . It can be seen that the new segments that we add for each original undashed edge take up only one quarter of the $8n \times 8n$ square to its right, therefore they do not create intersection. By (b) B' and R' intersect, it follows that B and R intersect. \square

6. PROPOSITIONAL PROOFS

Buss [Bus06] defines the $STCONN(n)$ tautologies to formalize the st-connectivity principle (see the previous section), where the blue path and red path are given as sets of edges. Thus there are propositional variables e_b and e_r for each horizontal and vertical edge slot e in the $n \times n$ grid, where e_b asserts that edge e is a blue edge and e_r asserts that e is a red edge. $STCONN(n)$ is the negation of a CNF formula whose clauses assert that the four corners each have degree one, the upper left and lower right corners each touch blue edges, the other two corners each touch red edges, every other node has degree zero or two and cannot touch both a blue and red edge.

Propositional proofs in AC^0 -Frege-systems (also called constant-depth Frege systems [Kra95]) allow formulas with unbounded AND and OR gates, as long as the total depth of the formula does not exceed a constant d , which is a parameter of the system. Every true Σ_0^B formula φ translates into a polynomial size family of constant depth propositional tautologies which have polynomial size AC^0 -Frege-proofs if φ is provable in V^0 (see [CN10]).

The propositional proof system $AC^0(2)$ -Frege (resp. TC^0 -Frege) is an extension of AC^0 -Frege which allows parity gates $\oplus(x_1, \dots, x_n)$ (resp. threshold gates $T_k(x_1, \dots, x_n)$) and has suitable axioms defining these gates. (The formula T_k is true when at least k of the inputs are true.) There are propositional translation results as above, where Σ_0^B -theorems of the theory $V^0(2)$ (resp. VTC^0) translate into polynomial size $AC^0(2)$ -Frege-proofs (resp. TC^0 -Frege-proofs) (see [CN10]).

It is shown in [Bus06] that the tautologies $STCONN(n)$ have polynomial size TC^0 -Frege proofs. The following stronger statement follows immediately from Theorem 5.1 and the translation theorem for $VZ(2)$.

THEOREM 6.1. *$STCONN(n)$ has polynomial size $AC^0(2)$ -Frege proofs.*

The Main Theorem for V^0 , which states the st-connectivity principle when paths are given as sequences of edges, translates into a family $STSEQ(n)$ of tautologies. Here the propositional variables have the form $b_{e,i}$ and $r_{e,i}$, which assert that edge e is the i -th edge in the blue (resp. red) path, $1 \leq i \leq n^2$. From the Main Theorem and the translation theorem for V^0 we obtain

THEOREM 6.2. *$STSEQ(n)$ has polynomial size AC^0 -Frege proofs.*

7. SOME OPEN PROBLEMS

As many problems in propositional proof complexity remain open, here we are not able to prove the “Reverse” part of Bounded Reverse Mathematics for our results. For example, it is open whether $\Sigma_0^B(\mathbf{V}^0(2))$, the theory axiomatized by the Σ_0^B consequences of $\mathbf{V}^0(2)$, can be axiomatized by \mathbf{V}^0 together with the Jordan Curve Theorem. In fact even the question of whether $\Sigma_0^B(\mathbf{V}^0(2))$ is finitely axiomatizable is open.

In [Bus06] Buss shows that

$$\text{Mod}_2 \preceq_{\text{cd}\mathcal{F}} \text{STCONN}$$

where Mod_2 denotes the counting modulo 2 principle, and $Q \preceq_{\text{cd}\mathcal{F}} R$ means that the principle Q has polynomial-size constant-depth proofs in a Frege proof system augmented with all instances of principle R . Buss’s proof shows that

$$\mathbf{V}^0 \vdash \text{STCONN} \supset \text{Mod}_2$$

for an appropriate first-order formalization of Mod_2 . At the time of writing this paper, it is still open whether the reverse direction can also be proved in \mathbf{V}^0 . If so, it would follow that $\mathbf{V}^0(2)$ could be replaced by $\mathbf{V}^0 + \text{Mod}_2$ in the statement of Theorem 3.3.

It is also open for primes $p > 2$ whether any of the implications

$$\text{STCONN} \supset \text{Mod}_p, \quad \text{Mod}_p \supset \text{STCONN}$$

is provable in \mathbf{V}^0 .

In this paper we have not considered the case in which the inputs are given as a set of *directed* edges. Possibly in this case the Jordan Curve Theorem can be proved in a subtheory of $\mathbf{V}^0(2)$. It is also interesting to consider cases where the input setting is mixed, e.g., the curve is a set of edges while the path is a sequence of edges.

ACKNOWLEDGMENT

We would like to thank Neil Thapen for the discussions related to Section 5, and the anonymous referees for constructive comments,

REFERENCES

- Miklós Ajtai. The Complexity of the Pigeonhole Principle. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 346–355, 1988.
- Samuel Buss. Polynomial-size Frege and Resolution Proofs of st-Connectivity and Hex Tautologies. *Theoretical Computer Science*, 357:35–52, 2006.
- Stephen Cook and Phuong Nguyen. *Logical Foundations of Proof Complexity*. ASL Perspectives in Logic Series. Cambridge University Press, 2010.
- Stephen Cook. Theories for Complexity Classes and Their Propositional Translations. In Jan Krajíček, editor, *Complexity of computations and proofs*, pages 175–227. Quaderni di Matematica, 2005.
- Stephen Cook and Charles Rackoff. Unpublished research notes, 3 June, 1997.
- Thomas Hales. A Verified Proof of the Jordan Curve Theorem. Seminar Talk, Department of Mathematics, University of Toronto, 8 Dec, 2005.
- Thomas Hales. The Jordan Curve Theorem, formally and informally. *American Mathematical Monthly*, 114(10):882–894, 2007.
- Neil Immerman. *Descriptive Complexity*. Springer, 1999.
- Jan Krajíček. *Bounded Arithmetic, Propositional Logic, and Complexity Theory*. Cambridge University Press, 1995.
- Phuong Nguyen and Stephen Cook. Theory for TC^0 and Other Small Complexity Classes. *Logical Methods in Computer Science*, 2(1), 2005.
- Phuong Nguyen and Stephen Cook. The Complexity of Proving Discrete Jordan Curve Theorem. In *Proc. 22nd IEEE Symposium on Logic in Computer Science*, pages 245–254, 2007.
- Phuong Nguyen. *Bounded Reverse Mathematics*. PhD thesis, University of Toronto, 2008. <http://www.cs.toronto.edu/~pnguyen/>.

- Alexander A. Razborov. Bounded Arithmetic and Lower Bounds in Boolean Complexity. In P. Clote and J. B. Remmel, editors, *Feasible Mathematics II*, pages 344–386. Birkhäuser, 1995.
- Stephen Simpson. *Subsystems of Second Order Arithmetic*. Springer, 1999.
- Carsten Thomassen. The Jordan-Schonflies Theorem and the Classification of Surfaces. *Amer. Math. Monthly*, 99(2):116–131, 1992.
- Alasdair Urquhart. Hex example. Email correspondence, Toronto theory group, April, 2001.
- Domenico Zambella. Notes on Polynomially Bounded Arithmetic. *Journal of Symbolic Logic*, 61(3):942–966, 1996.

Received February 2010; revised December 2010; accepted January 2011