Formal Reasoning Using Low-Complexity Concepts

Lecture 3: Theories for $\mathbf{AC}^0(2)$ and other classes

Stephen Cook

Fall School, Prague

September, 2008

Recall Uniform Circuit Complexity Classes

Problems are specified by a (uniform) poly-size family $\langle C_n \rangle$ of Boolean circuits.
C_n solves problems with input length n.

$\mathbf{AC^0}$: bounded depth, unbounded fan-in $\wedge, \vee$.
(Log time hierarchy for Alternating TMs)

$\mathbf{AC^0}(2)$: allow unbounded fan-in parity gates.
Cannot count mod 3 [Raz 87],[Smo 87]

$\mathbf{TC^0}$: allow threshold gates.
Contains binary $\times$

Appropriate uniform definitions come from descriptive complexity literature (e.g. [BIS]).

Our characterizations: Each function class has a function which is complete under (so-called) $\mathbf{AC^0}$-reductions.

$parity(X)$ is complete for $\mathbf{AC^0}(2)$
$numones(X)$ is complete for $\mathbf{TC^0}$.

Semantic Definition: We say that F (resp. f) is Σ_0^B-**definable** from $\mathcal{L}$ if its bit graph (resp. graph) is represented by a $\Sigma_0^B(\mathcal{L})$ formula.

If $\mathcal{L}$ consists only of functions in $\mathbf{FAC^0}$, then F is Σ_0^B-definable from $\mathcal{L}$ iff F is in $\mathbf{FAC^0}$. The same for f. (Because $\mathcal{L}_{\mathbf{FAC^0}}$ for $\overline{\mathbf{V^0}}$ represents exactly $\mathbf{FAC^0}$.)

Semantic Definition: We say that a string function F (resp. a number function f) is $\mathbf{AC^0}$-**reducible** to $\mathcal{L}$ if there is a sequence of string functions $F_1, \ldots, F_n$ ($n \geq 0$) such that F_i is Σ_0^B-definable from $\mathcal{L} \cup \{F_1, \ldots, F_{i-1}\}$, for $i = 1, \ldots, n$; and that F (resp. f) is Σ_0^B-definable from $\mathcal{L} \cup \{F_1, \ldots, F_n\}$. [Think circuits]

$\mathbf{FAC^0}$ is closed under $\mathbf{AC^0}$ reducibility.
But in general a function $\mathbf{AC^0}$-reducible to $\{G\}$ may not be Σ_0^B-definable from $\{G\}$.

Notation: The function $Row(i, X) = X^{[i]}$ has bit definition

$$X^{[i]}(j) \leftrightarrow X(i, j) \qquad i.e. \quad \leftrightarrow X(\langle i, j \rangle)$$

Example:

$$parity(X) = \begin{cases} 1 & \text{if } X \text{ has an odd number of elements} \\ 0 & \text{otherwise} \end{cases}$$

Define $F(X)$ and $G(X)$ by
$F(X)(i) \leftrightarrow i < |X| \wedge parity(X^{[i]}) = 1$.
$G(X) = parity(F(X))$.

Then $G(X)$ is $\mathbf{AC^0}$-reducible to $\{parity\}$, but apparently $G(X)$ is not Σ_0^B-definable from $\{parity\}$.

Program for Chapter 9 [Nguyen]
Introduce a (two-sorted) minimal canonical theory **VC** for each major complexity class $\mathbf{C}$, for $\mathbf{AC^0} \subseteq \mathbf{C} \subseteq \mathbf{P}$.

• **VC** has vocabulary $\mathcal{L}_A^2$.

• $\mathbf{V}\mathbf{AC^0} = \mathbf{V^0}$

• For $\mathbf{C} \neq \mathbf{AC^0}$, $\mathbf{VC} = \mathbf{V^0}+\{\text{one axiom}\}$ (finitely axiomatizable)
The axiom asserts the existence of a solution to a complete problem for $\mathbf{C}$ under $\mathbf{AC^0}$ reductions.

• The provably total (i.e. Σ_1^B-definable) functions in **VC** are those in **FC**.

• **VC** has a universal conservative extension $\overline{\mathbf{VC}}$ in the style of $\overline{\mathbf{V^0}}$.

Theories VC for various classes C

Recall $\mathbf{VC} = \mathbf{V^0} + \mathbf{Axiom_C}$
where $\mathbf{Axiom_C} = (\mathbf{Complete_C}$ has a solution)

class	theory	$\mathbf{Complete_C}$
$\mathbf{AC^0}$	$\mathbf{V^0}$	none
$\mathbf{AC^0}(2)$	$\mathbf{V^0}(2)$	$parity(X)$
$\mathbf{TC^0}$	$\mathbf{VTC}^0$	$numones(X)$
$\mathbf{NC}^1$	$\mathbf{VNC}^1$	tree-$MCVP$
$\mathbf{L}$	$\mathbf{VL}$	$UniConn(z, a, E)$
$\mathbf{NL}$	$\mathbf{VNL}$	$Conn(z, a, E)$
$\mathbf{P}$	$\mathbf{VP}$	$MCVP$

Robustness Theorems

$\mathbf{VTC}^0 \simeq \Delta_1^B$-**CR** [JP] (proved in [Nguyen])
$\mathbf{VNC}^1 \simeq \mathbf{AID}$ [Arai] (proved in [CM])
$\mathbf{VNC}^1 \simeq ALV \simeq ALV'$ [Clote] (proved by [Nguyen])
$\mathbf{VL} = \Sigma_0^B$-Rec [Zam97]
$\mathbf{VNL} = \mathbf{V\text{-}Krom}$ [Kolokolova]

The Theory $\mathbf{V^0}(2)$

Uniform $\mathbf{FAC^0}(2)$ is the closure of $\{parity\}$ under $\mathbf{AC^0}$-reductions.

We will use $Parity(x, Y)$ to axiomatize $\mathbf{V^0}(2)$. Here $Parity(x, Y) = Z$ iff $|Z| \leq x + 1$ and for $i \leq x$, $Z(i)$ holds iff there is an odd number of ones in $Y(0)Y(1)\ldots Y(i-1)$. Thus

$$Z = Parity(x, Y) \leftrightarrow |Z| \leq x + 1 \wedge \delta_{parity}(x, Y, Z)$$

where δ_{parity} is a Σ_0^B-formula:

$$\delta_{parity}(x, Y, Z) \equiv \neg Z(0) \wedge \forall i < x(Z(i+1) \leftrightarrow (Z(i) \oplus Y(i)))$$

Definition: The theory $\mathbf{V^0}(2)$ has vocabulary $\mathcal{L}_A^2$ and axioms those of $\mathbf{V^0}$ and

$$\exists Z \leq x + 1 \ \delta_{parity}(x, Y, Z)$$

Obviously $Parity(x, Y)$ is Σ_1^B-definable in $\mathbf{V^0}(2)$.

Theorem: $\mathbf{V^0}(2)$ captures $\mathbf{AC^0}(2)$. That is, the Σ_1^B-definable functions comprise exactly $\mathbf{FAC^0}(2)$.

This is proved by defining an extension $\overline{\mathbf{V^0}(2)}$ of $\mathbf{V^0}(2)$ analogous to $\overline{\mathbf{V^0}}$ as an extension of $\mathbf{V^0}$.

Thus $\overline{\mathbf{V^0}(2)}$ is a universal theory which is a conservative extension of $\mathbf{V^0}(2)$ and whose string function symbols range over the string functions in $\mathbf{FAC^0}(2)$. $\overline{\mathbf{V^0}(2)} \vdash \Sigma_0^B(\mathcal{L})$-**IND**.

NOTE: In general the number functions of **FC** all have the form $|F(\vec{x}, \vec{X})|$ for some string function in the class.

$\overline{\mathbf{V^0}(2)}$ has a vocabulary $\mathcal{L}_{\mathbf{FAC^0}(2)}$ which includes $Parity$, but we need an open axiom defining $Parity$. We omit the quantifier $\forall i < x$ and use

$$\begin{aligned} Z = Parity(x, Y) \ \supset \ & [|Z| \leq x + 1 \wedge \neg Z(0) \wedge \\ & i < x \supset (Z(i+1) \leftrightarrow (Z(i) \oplus Y(i)))] \end{aligned}$$

Definition: The vocabulary $\mathcal{L}_{\mathbf{FAC^0}(2)}$ for $\overline{\mathbf{V^0}(2)}$ is the smallest set containing $\mathcal{L}_{\mathbf{FAC^0}}$ and $Parity$ and for each open formula $\varphi(z, \vec{x}, \vec{X})$ over $\mathcal{L}_{\mathbf{FAC^0}(2)}$ and $\mathcal{L}_A^2$-term $t(\vec{x}, \vec{X})$, there is a string function $F_{\varphi(z),t}$ in $\mathcal{L}_{\mathbf{FAC^0}(2)}$ with defining axiom

$$F_{\varphi(z),t}(\vec{x}, \vec{X})(z) \leftrightarrow z < t(\vec{x}, \vec{X}) \wedge \varphi(z, \vec{x}, \vec{X})$$

Definition: $\overline{\mathbf{V^0}(2)}$ is the universal theory over $\mathcal{L}_{\mathbf{FAC^0}(2)}$ whose axioms consist of the axioms for $\overline{\mathbf{V^0}}$, the open defining axiom for $Parity$, and the above defining axioms for $F_{\varphi(z),t}$.

Claim 1: $\overline{\mathbf{V^0}(2)}$ extends $\mathbf{V^0}(2)$. (Easy)

Claim 2: The function symbols in $\mathcal{L}_{\mathbf{FAC^0}(2)}$ represent precisely the functions in $\mathbf{FAC^0}(2)$. (Easy, by the way we defined $\mathbf{AC^0}(2)$).

Claim 3: $\overline{\mathbf{V^0}(2)}$ is conservative over $\mathbf{V^0}(2)$. (Hard – the open $\mathcal{L}_{\mathbf{FAC^0}(2)}$-formulas are not equivalent to $\Sigma_0^B(Parity)$ formulas. [Witness]

Solution: Introduce notion of 'Aggregate of a function'.

Surjective Pigeonhole Principle SPHP

As for PHP, the string variable X represents a mapping from pigeons to holes: $X(i,j)$ means pigeon i maps to hole j.

$$SPHP(n,X) \equiv \neg(total(n,X) \wedge INJ(n,X) \wedge SUR(n,X))$$

where $total(n,X) \equiv$
$\forall i \leq n \exists j < n X(i,j)$

$INJ(n,X) \equiv$
$\forall i,j \leq n \forall k < n (X(i,k) \wedge X(j,k) \supset i = j)$

$SUR(n,X) \equiv$
$\forall k < n \exists i \leq n X(i,k)$

Theorem: $\mathbf{V^0}(2) \vdash$ SPHP.

Conjecture: $\mathbf{V^0}(2) \nvdash$ PHP.

Discrete Jordan Curve Theorem
[Nguyen/Cook LICS 07]

Original statement: A simple closed curve divides the plane into exactly two connected components.

(Hales gave a computer-verified proof involving 44,000 proof steps. His proof started with a discrete version. Warmup for Kepler Conjecture.)

Discrete Setting: The curve consists of edges connecting grid points in the plane.

Case I: The curve is given as a set of edges such that every grid point has degree 0 or 2.

(Then there may be more than 2 connected components.)

Theorem: $\mathbf{V^0}(2)$ proves the following:
If B is a set of edges forming a curve and p_1, p_2 are two points on different sides of B, and R is a set of edges that connects p_1 and p_2, then B and R intersect.

Jordan Curve Cont'd

Theorem: [Buss] $\mathbf{V^0}$ cannot prove the previous version of JCT.

Case II: The curve is given as a sequence of edges.

Theorem: $\mathbf{V^0}$ proves that a curve given by a sequence of edges divides the plane into exactly two connected components.

Lemma: (Provable in $\mathbf{V^0}$) For each column in the planar grid, the edges of a closed curve alternate in direction.

(The proof is difficult in $\mathbf{V^0}$, since no counting is allowed, even mod 2.)

Note that these statements of JCT (and connectivity) are formalized with Σ_0^B-formulas.

MAJOR OPEN QUESTION: Are there Σ_0^B-theorems which separate $\mathbf{V^0}(2)$ and $\mathbf{VTC}^0$?

(We know there are Σ_1^B-theorems which separate $\mathbf{V^0}(2)$ and $\mathbf{VTC}^0$, since $\mathbf{AC^0}(2) \neq \mathbf{TC^0}$, by [Raz/Smo].)

NONUNIFORM VERSION: for propositional proof systems: Does $\mathbf{AC^0}(2)$-Frege p-simulate $\mathbf{TC^0}$-Frege?

Candidates: PHP, and also Counting mod 3.

Replacement Axiom Scheme

Definition: $\mathbf{g}\Sigma_1^B(\mathcal{L})$ is the closure of $\Sigma_0^B(\mathcal{L})$ under $\wedge, \vee, \forall x \le t, \exists x \le t$, and $\exists X \le t$.

Thus $\Sigma_1^B(\mathcal{L}) \subset \mathbf{g}\Sigma_1^B(\mathcal{L})$, but
$\forall x \le t \exists Y \le t \varphi$ is in $\mathbf{g}\Sigma_1^B$ but not in Σ_1^B (where φ is in Σ_0^B).

$\mathbf{g}\Sigma_1^B$ is more analogous to Buss's Σ_1^b.

Recall $Z^{[i]}(j) \leftrightarrow Z(i,j)$ ($Z^{[i]}$ is row i of Z)

Definition of Replacement Axiom:
Φ-**REPL** is the set of all formulas

$$(\forall i \le b \exists X \le c\varphi(i,X)) \supset \exists Z \le \langle b,c\rangle \forall i \le b(|Z^{[i]}| \le c \wedge \varphi(i, Z^{[i]}))$$

where φ is in Φ.

Note the LHS is a logical consequence of the RHS.

Lemma A: If a theory $\mathcal{T}$ over $\mathcal{L}$ proves $\Sigma_0^B(\mathcal{L})$-**REPL**, then for each $\mathbf{g}\Sigma_1^B(\mathcal{L})$-formula φ there is a $\Sigma_1^B(\mathcal{L})$-formula φ' such that $\mathcal{T} \vdash \varphi \leftrightarrow \varphi'$.

Proof: Structural induction on φ.

Theorem [Buss 86] $\mathbf{V}^1 \vdash \Sigma_0^B$-**REPL**.

Proof: Note that $\mathbf{V}^1 \vdash \Sigma_1^B$-**IND**.
Assume $\forall i \leq b \exists X \leq c \varphi(i, X)$.
Prove by induction on $n \leq b$

$$\exists Z \leq \langle n, c\rangle \forall i \leq n(|Z^{[i]}| \leq c \wedge \varphi(i, Z^{[i]}))$$

QED

Thus by Lemma A, for $\mathbf{V}^1$ the distinction between Σ_1^B and $\mathbf{g}\Sigma_1^B$ is not important.

But Σ_0^B-**REPL** would be very helpful in showing $\overline{\mathbf{V}^0}(2)$ is conservative over $\mathbf{V^0}(2)$, and $\mathbf{V^0}$ does not prove Σ_0^B-**REPL**.

To show $\mathbf{V^0}$ does not prove Σ_0^B-**REPL** we need the following form of the Herbrand Theorem:

KPT Witnessing Theorem: Let $\mathcal{T}$ be a universal two-sorted theory, let φ be an open $\mathcal{L}$-formula and suppose

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

Then there exists a finite sequence $T_1, \ldots, T_k$ of string terms such that

$$\begin{aligned}\mathcal{T} \vdash & \varphi(X, T_1(X), Z_1) \vee \varphi(X, T_2(X, Z_1), Z_2) \\ & \vee \ldots \vee \varphi(X, T_k(X, Z_1, \ldots, Z_{k-1}), Z_k).\end{aligned}$$

where the notation $T_i(X, Z_1, \ldots, Z_{i-1})$ means that only the displayed variables may occur in T_i.

Proof: Let $B, C_1, C_2, ...$ be a list of new string constants, and let $U_1, U_2, ...$ be an enumeration of all terms built from symbols of $\mathcal{L}$ together with $B, C_1, C_2, ...$, where the only new constants in U_k are among $\{B, C_1, ..., C_{k-1}\}$. It suffices to show that

$$\mathcal{T} \cup \{\neg\varphi(B, U_1, C_1), \neg\varphi(B, U_2, C_2), \ldots, \neg\varphi(B, U_k, C_k)\}$$

is unsatisfiable for some k.

Proof of KPT continued:

Suppose otherwise. Then by compactness

$$\mathcal{T} \cup \{\neg\varphi(B, U_1, C_1), \neg\varphi(B, U_2, C_2), ...\} \qquad (1)$$

has a model $\mathcal{M}$. Since $\mathcal{T}$ is universal, the substructure $\mathcal{M}'$ consisting of the denotations of the terms $U_1, U_2, ...$ is also a model for (1). It is easy to see that

$$\mathcal{M}' \models \mathcal{T} + \forall Y \exists Z \neg\varphi(B, Y, Z)$$

and hence $\mathcal{T} \nvdash \forall X \exists Y \forall Z \varphi(X, Y, Z)$.

Theorem: [CoTh] $\overline{\mathbf{V^0}}$ (and hence $\mathbf{V^0}$) does not prove Σ_0^B-**REPL**.

Proof idea: Suppose for all n F is a bijection from strings of length n to itself, and $F \in \mathbf{FAC^0}$ but $F^{-1} \notin \mathbf{FAC^0}$. Suppose $\overline{\mathbf{V^0}}$ proves the following instance of **REPL**:

$$(\forall i < t\, \exists Z < u F(Z) = W^{[i]}) \supset \exists Y \forall j < u F(Y^{[j]}) = W^{[j]} \tag{2}$$

We can rewrite this as

$$\exists i < t\, \exists Y \forall Z < u \Big(F(Z) = W^{[i]} \supset \forall j < t\, F(Y^{[j]}) = W^{[j]} \Big)$$

Applying KPT we get a positive integer k and functions $g_1, \ldots, g_k, H_1, \ldots, H_k$ such that T proves

$$\begin{aligned}
&(F(Z_1) = W^{[g_1(W)]} \supset \forall j < t\, F(H_1(W)^{[j]}) = W^{[j]}) \\
&\vee (F(Z_2) = W^{[g_2(W,Z_1)]} \supset \forall j < t\, F(H_2(W,Z_1)^{[j]}) = W^{[j]}) \\
&\vee \ldots \\
&\vee (F(Z_k) = W^{[g_k(W,Z_1,\ldots,Z_{k-1})]} \supset \\
&\forall j < t\, F(H_k(W,Z_1,\ldots,Z_{k-1})^{[j]}) = W^{[j]})
\end{aligned}$$

This allows the "student", given an input W (considered as a sequence $W^{[0]}, \ldots, W^{[t-1]}$), to compute Y coding a sequence of pre-images of F of all t elements of W, by asking the "teacher" for pre-images of at most k elements of W.

Let $PAR(m, X)(i)$ be the parity of number of 1's in $X(0)...X(i)$, $i < m$. Fix m: PAR is a bijection from strings for length $\leq m$ to itself. $UNPAR \in \mathbf{FAC^0}$ but $PAR \notin \mathbf{FAC^0}$.
Let m be a parameter, set $F = UNPAR$.

Suppose $m \geq 2k$. We will show how to use this algorithm to compute the parity of a single string I, $|I| \leq m$, in uniform randomized $\mathbf{AC^0}$.

Choose m strings $U_1, \ldots, U_m$ of m bits each uniformly at random, and for each i compute $V_i = UNPAR(U_i)$. Choose a number r, $1 \leq r \leq m$, uniformly at random. For $1 \leq i \leq m$ define $W^{[i]}$ by the condition

$$W^{[i]} = \begin{cases} V_i & \text{if } i \neq r \\ I \oplus V_r & \text{if } i = r. \end{cases}$$

Now run our student-teacher $\mathbf{AC^0}$ algorithm on W. If the student asks "what is $PAR(W^{[i]})$?" for $i \neq r$, reply with U_i (or $W^{[0]}$ if $i = 0$) (which is the correct answer). If the algorithm queries "what is $PAR(Y^{[r]})$?", then abort the computation.

Since $PAR(W^{[i]})$ is queried for at most k different values of i and since for each input I each pair (W, r) is equally likely to have been chosen, it follows that the computation will be aborted with probability at most $k/m \leq 1/2$.

Hence with probability at least 1/2 the algorithm is not aborted, we are able to answer all the queries correctly, and we can recover $PAR(I)$ from Y.

This algorithm can be converted to a nonuniform $\mathbf{AC^0}$ circuit family for PARITY, which is impossible.