Algebrization

Algebrization: A New Barrier in Computational Complexity
[Scott Aaronson, Avi Wigderson; (2008)]

Any proof of $P \neq NP$ will have to overcome two barriers:

- **Relativization**
  T. Baker, J. Gill, and R. Solovay. Relativizations of the $P =?NP$ question. 1975
- **Natural Proofs**
  A. A. Razborov and S. Rudich. Natural proofs. 1997

Relativization was circumvented by **Arithmetization**
C. Lund, L. Fortnow, H. Karloff, and N. Nisan. Algebraic methods for interactive proof systems. 1992
A. Shamir. IP=PSPACE. 1992
Any complexity class separation proved via **diagonalization** is non-naturalizing.

We now have circuit lower bounds that evade both barriers, by combining arithmetization with diagonalization.

- $MA_{EXP} \not\subset P/poly$ [Buhrman, Fortnow, and Thierauf. Nonrelativizing separations. 1998]
- $PP \not\subset SIZE(n^k)$ for every fixed k. [Vinodchandran. A note on the circuit complexity of PP. 2004]
- $PromiseMA \not\subset SIZE(n^k)$ for fixed k. [Santhanam. Circuit lower bounds for Merlin-Arthur classes. 2007]

Could arithmetization and diagonalization already suffice to prove $NEXP \not\subset P/poly$, or $P \neq NP$? Or is there a third barrier, to which even the most recent results are subject? [Santhanam 2007]

Algebrization: A generalization of relativization where the simulating machine gets access not only to an oracle A, but also a low-degree extension $\widetilde{A}$ of A over a finite field or ring.

Algebrization

Definition

The inclusion $C \subseteq D$ **relativizes** if $C^A \subseteq D^A$ for all oracles A.

Definition

Given an oracle $A = \{A_n\}$ with $A_n : \{0,1\}^n \to \{0,1\}$, an **extension** $\widetilde{A}$ of A is a collection of polynomials $\widetilde{A}_n : Z^n \to Z$ satisfying:

- $\widetilde{A}_n(x) = A_n(x)$ for all Boolean $x \in \{0,1\}^n$,
- $deg(\widetilde{A}_n) = O(n)$,
- $size(\widetilde{A}_n(x)) \leq p(size(x))$ for some polynomial p, where

$$size(x) := \sum_{i=1}^{n} \lceil 1 + log_2 |x_i| \rceil$$

Algebrization

Definition

- A complexity class inclusion $C \subseteq D$ **algebrizes** if $C^A \subseteq D^{\widetilde{A}}$ for all oracles A and all (low-degree) extensions $\widetilde{A}$ of A over a finite field or ring.
- A separation $C \not\subset D$ **algebrizes** if $C^{\widetilde{A}} \not\subset D^A$ for all $A, \widetilde{A}$.
- Proving $C \subseteq D$ **requires non-algebrizing techniques** if there exist $A, \widetilde{A}$ such that $C^A \not\subset D^{\widetilde{A}}$.
- Proving $C \not\subset D$ **requires non-algebrizing techniques** if there exist $A, \widetilde{A}$ such that $C^{\widetilde{A}} \subseteq D^A$.

- Almost all known techniques in complexity theory algebrize.
- Any proof of $P \neq NP$ or $P = RP$ or $NEXP \not\subset P/poly$ will require non-algebrizing techniques.

Proving that existing results algebrize

In *Algebraic methods for interactive proof systems* (1992), C. Lund, L. Fortnow, H. Karloff, and N. Nisan show that $coNP \subseteq IP$.

In their protocol Arthur arithmetizes a Boolean formula ϕ to produce a low-degree polynomial $\widetilde{\phi}$. Merlin wants to convince Arthur that $\sum_{x \in \{0,1\}^n} \widetilde{\phi}(x) = 0$.

In the last step Arthur checks that $\widetilde{\phi}(r_1, ..., r_n)$ equals the value claimed by Merlin for some $r_1, ..., r_n$ randomly chosen earlier.

How was the polynomial p produced in the LFKN protocol?

By starting from a Boolean circuit, then multiplying together terms that enforce 'correct propagation' at each gate:

$$\widetilde{A}(x, y)g + (1 - \widetilde{A}(x, y))(1 - g)$$

Arthur and Merlin then reinterpret p not as a Boolean function, but as a polynomial over some larger field.

But what if the circuit contained oracle gates? Then how could Arthur evaluate p over the larger field?
He'd almost need oracle access to a low-degree extension $\widetilde{A}$ of A.

Suppose we want to prove $coNP^{A} \subseteq IP^{\widetilde{A}}$.

Now Arthur's formula φ will in general contain A gates, in addition to the usual AND, OR, and NOT gates.

When Arthur arithmetizes φ to produce a low-degree polynomial $\widetilde{\varphi}$, his description of $\widetilde{\varphi}$ will contain terms of the form $A(z_1, ..., z_k)$.

Inputs $z_1, ..., z_k$ can be non-Boolean.

Arthur calls the oracle $\widetilde{A}$ to get $\widetilde{A}(z_1, ..., z_k)$.

Other Results That Algebrize

- $PSPACE^{A[poly]} \subseteq IP^{\widetilde{A}}$
- $NEXP^{A[poly]} \subseteq MIP^{\widetilde{A}}$
- $PP^{\widetilde{A}} \subset P^{A}/poly$
- $NEXP^{\widetilde{A}}[poly] \subset P^{A}/poly$
- $MA_{EXP}^{\widetilde{A}[exp]} \not\subset P^{A}/poly$
- $PP^{\widetilde{A}} \not\subset SIZE^{A}(n)$
- $PromiseMA^{\widetilde{A}} \not\subset SIZE^{A}(n)$
- $\exists$ OWF secure against $P^{\widetilde{A}} \Rightarrow NP^{A} \subseteq CZK^{\widetilde{A}}$

Proving $P \neq NP$ Will Require Non-Algebrizing Techniques

Theorem

There exists an oracle A, and an extension $\widetilde{A}$, such that $NP^{\widetilde{A}} \subseteq P^{A}$

Proof.

Let A be a $PSPACE$-complete language, and let $\widetilde{A}$ be the unique multilinear extension of A. Then $\widetilde{A}$ is also $PSPACE$-complete [Babai, Fortnow, Lund]. Hence $NP^{\widetilde{A}} = P^{A} = PSPACE$. □

Proving P=RP Will Require Non-Algebrizing Techniques

Theorem

There exist $A, \widetilde{A}$ such that $RP^A \not\subset P^{\widetilde{A}}$.

We have to prove algebraic oracle separations.

- Prove a concrete lower bound on the query complexity of some function.
- Use the query complexity lower bound to diagonalize against a class of Turing machines.

The first step requires us to prove lower bounds in a new model of algebraic query complexity.
An algorithm is given oracle access to a Boolean function $A : \{0,1\}^n \to \{0,1\}$. It is trying to answer some question about A, by querying A on various points. The algorithm can query not just A itself, but also an adversarially-chosen low-degree extension $\widetilde{A}$.

Lemma

Let F be a field, and let $Y \subseteq F^n$ be the set of points queried by the algorithm. Then there exists a polynomial $p : F^n \to F$, of degree at most $2n$, such that

- $p(y) = 0$ *for all* $y \in Y$.
- $p(z) = 1$ *for at least* $2^n - |Y|$ *Boolean points* z.
- $p(z) = 0$ *for the remaining Boolean points.*

Proof.

Given a Boolean point z, let δ_z be the unique multilinear polynomial that is 1 at z and 0 at all other Boolean points. Then we can express any multilinear polynomial r as $r(x) = \sum_{z\in\{0,1\}^n} \alpha_z \delta_z(x)$.
Requiring $r(y) = 0$ for all $y \in Y$, yields $|Y|$ linear equations in 2n unknowns. Hence there exists a solution r such that $r(z) \neq 0$ for at least $2n - |Y|$ Boolean points z. We now set

$$p(x) = \sum_{z\in\{0,1\}^n : r(z)\neq 0} \frac{r(x)\delta_z(x)}{r(z)}$$

□

A standard diagonalization argument now yields the separation between P and RP in the case of finite fields.

Other Oracle Results We Can Prove By Building Polynomials

- $\exists A, \widetilde{A} : NP^{A} \not\subset coNP^{\widetilde{A}}$
- $\exists A, \widetilde{A} : NP^{A} \not\subset BPP^{\widetilde{A}}$ (only for finite fields, not integers)
- $\exists A, \widetilde{A} : NEXP^{\widetilde{A}[exp]} \subset P^{A}/poly$
- $\exists A, \widetilde{A} : NP^{\widetilde{A}} \subset SIZE^{A}(n)$

From Algebraic Query Algorithms to Communication Protocols

$A : \{0,1\}^n \to \{0,1\}$

$A_0(Alice)$	$A_1(Bob)$
$A(000) = 1$	$A(100) = 0$
$A(001) = 0$	$A(101) = 0$
$A(010) = 0$	$A(110) = 1$
$A(011) = 1$	$A(111) = 1$

Alice and Bob's Goal: Compute some property of A using minimal communication.
Let $\widetilde{A} : F^n \to F$ be the unique multilinear extension of A over finite field F.

Theorem

If a problem can be solved using T queries to $\widetilde{A}$, then it can also be solved using $O(Tnlog|F|)$ bits of communication between Alice and Bob.

Proof.

Given any point $y \in F^n$, we can write $\widetilde{A}(y)$ as a linear combination of the values taken by A on the Boolean cube

$$\widetilde{A}(y) = \sum_{x \in \{0,1\}^n} A(x)\delta_x(y) = \widetilde{A_0}(y) + \widetilde{A_1}(y)$$

Let $y_1 \in F^n$ be the first point queried. Then Alice computes the partial sum $\widetilde{A}_0(y_1) = \sum_{x \in \{0,1\}^{n-1}} \delta_{0x}(y)A(0x)$ and sends $(y_1, \widetilde{A}_0(y_1))$ to Bob. Bob computes $\widetilde{A}_1(y_1) = \sum_{x \in \{0,1\}^{n-1}} \delta_{1x}(y)A(1x)$ from which he learns $\widetilde{A}_1(y_1) = \widetilde{A}_0(y_1) + \widetilde{A}_1(y_1)$. Bob computes $\widetilde{A}_1(y_2)$ and sends $(y_2, \widetilde{A}_1(y_2))$, and so on for T rounds.

Each message uses $O(nlog|F|)$ bits, from which it follows that the total communication cost is $O(Tnlog|F|)$.

□

- This argument works in the randomized world, the nondeterministic world, the quantum world.
- Any communication complexity lower bound leads to an algebraic query complexity lower bound.
- It yields multilinear extensions instead of multiquadratic ones.
- It works just as easily over the integers as over finite fields.
- The lower bounds one gets from communication complexity are more contrived.

Some Applications to Communication Complexity

Theorem

Let Alice and Bob hold 3SAT instances φ_A, φ_B respectively of size N. Suppose there is no IP-protocol with communication cost $O(polylogN)$, by which Merlin can convince Alice and Bob that φA and φB have a common satisfying assignment. Then $NL \neq NP$.

Theorem

Let φ be a 3SAT instance of size N. Suppose there is no bounded-error randomized verifier that decides whether φ is satisfiable by

- *making $O(polylogN)$ queries to a binary encoding of φ, and*
- *exchanging $O(polylogN)$ bits with a competing yes-prover and no-prover, both of whom know φ and can exchange private messages not seen by the other prover. Then $P \neq NP$.*

Separations in Communication Complexity Imply Algebraic Oracle Separations

$\Omega(2^n)$ randomized lower bound for Disjointness [KS 1987] [Razborov 1990]	$\exists A, \widetilde{A} : NP^A \not\subset BPP^{\widetilde{A}}$
$\Omega(2^{\frac{n}{2}})$ quantum lower bound for Disjointness [Razborov 2002]	$\exists A, \widetilde{A} : NP^A \not\subset BQP^{\widetilde{A}}$
$\Omega(2^{\frac{n}{2}})$ lower bound on MA-protocols for Disjointness [Klauck 2003]	$\exists A, \widetilde{A} : coNP^A \not\subset MA^{\widetilde{A}}$
Exponential separation between classical and quantum communicat. complexities [Raz 1999]	$\exists A, \widetilde{A} : BQP^A \not\subset BPP^{\widetilde{A}}$
Exponential separation between MA and QMA communication complexities [Raz-Shpilka 2004]	$\exists A, \widetilde{A} : QMA^A \not\subset MA^{\widetilde{A}}$

Conclusions

Arithmetization had a great run. It led to IP=PSPACE, the PCP Theorem, non-relativizing circuit lower bounds.
Yet we showed it''s fundamentally unable to resolve barrier problems like P vs. NP, or even P vs. BPP or NEXP vs. P/poly.

Why? It 'doesn't pry open the black-box wide enough'.
I.e. it uses a polynomial-size Boolean circuit to produce a low-degree polynomial, which it then evaluates as a black box. It doesn't exploit the small size of the circuit in any 'deeper' way.

To reach this conclusion, we introduced a new model of algebraic query complexity, which has independent applications (e.g. to communication complexity) and lots of nooks and crannies to explore in its own right.

Algebrization provides nearly the precise limit on the non-relativizing techniques of the last two decades. We speculate that going beyond this limit will require fundamentally new methods.

Open Problems

- Develop non-algebrizing techniques.
- Do there exist $A, \widetilde{A}$ such that $coNP^{A} \not\subset AM^{\widetilde{A}}$?
- Improve $PSPACE^{A[poly]} \subset IP^{\widetilde{A}}$ to $PSPACE^{\widetilde{A}[poly]} = IP^{\widetilde{A}}$.
- The power of 'double algebrization'.
- Integer queries of unbounded size.
- Algebraic query lower bounds ⇒ communication lower bounds?
- Generalize to arbitrary error-correcting codes (not just low-degree extensions)?
- Test if a low-degree extension came from a small circuit?