Upper Bounding Time-Space Lower Bounds for Satisfiability Algorithms

Sam Buss
(joint work with Ryan Williams)

MALOA conference
CIRM, Luminy
May 1, 2013

This talk addresses complexity questions towards separating logarithmic space (L) from non-deterministic polynomial time (NP).

$$\mathrm{L} \subseteq \mathrm{P} \subseteq \mathrm{NP} \subseteq \mathrm{PSPACE} \subseteq \mathrm{EXPTIME}.$$

Space hierarchy gives: $\mathrm{L} \neq \mathrm{PSPACE}$.
Time hierarchy gives: $\mathrm{P} \neq \mathrm{EXPTIME}$.
No other separations are known.

A series of results, especially since Fortnow [1997], has proved some *lower bounds* for the time complexity of sublinear space algorithms for Satisfiability (SAT) and thus for NP problems.

This talk discusses *upper bounds* on the *lower bounds* that can be obtained by present techniques of “alternation trading”.

Barriers to separating L, P and NP include:

Oracle results: [Baker-Gill-Solovay, 1975] There are oracles collapsing the classes, so any proof of separation must relativize.

Natural proofs: [Razborov-Rudich, 1997] Cryptographic assumptions imply that certain constructive separations are not possible.

Algebrization: [Aaronson-Wigderson, 2008] Proofs must relativize to algebraic extensions of oracles.

Present talk: Bounds on the power of **alternation-trading** proofs for separating L and NP.

Alternation-trading proofs involve iterating the restricted space methods of Nepomnjasci together with simulations: essentially a sophisticated version of diagonalization.

Theme: Better simulation methods give better diagonalization proofs for separating complexity classes.

Satisfiability

Definition (Satisfiability – SAT)

An instance of satisfiability is a set of clauses.
Each clause is a set of literals.
A *literal* is a negated or nonnegated propositional variable.
Satisfiability (SAT) is the problem of deciding if there is a truth assignment that sets at least one literal true in each clause.

Thm: Satisfiability is NP-complete.

Conjecture: Satisfiability is not polynomial time. $(\mathrm{P} \neq \mathrm{NP}.)$

Best lower bounds to date state that SAT is not computable in simultaneous time n^c and space n^ϵ for certain values of c and of $\epsilon > 0$. (But, not all such values!)

Why is Satisfiability important?

1. Satisfiability is NP-complete.

2. Many other NP-complete problems are many-reducible to SAT in quasilinear time, that is, time $n \cdot (\log n)^{O(1)}$.

3. For a given non-deterministic machine M, the question of whether $M(x)$ accepts is reducible to SAT in quasilinear time. [Cook, 1988; ...].

Thus SAT is a “canonical” and natural non-deterministic time problem. Lower bounds on algorithms for SAT will translate into lower bounds for many other problems.

This talk always uses the Random Access Memory (RAM) model for computation.

Theorem (Cook, 1988; ...)

For any $L \in \mathrm{NTIME}(T(n))$ *there is a quasi-linear time, many-one reduction to instances of SAT of size* $T(n)(\log T(n))^c$*. In fact, each symbol of the instance of SAT is computable in polylogarithmic time* $(\log T(n))^c$*.*

Corollary

If $\mathrm{SAT} \in \mathrm{DTIME}(n^c)$*, then* $\mathrm{NTIME}(n^d) \subset \mathrm{DTIME}(n^{c \cdot d + o(1)})$*.*

"DTIME" / "NTIME" = Deterministic/Nondeterministic time.

Definition

Let $c \geq 1$. $\mathrm{DTS}(n^c)$ is the class of problems solvable in simultaneous deterministic time $n^{c+o(1)}$ and space $n^{o(1)}$.

A series of results by Kannan [1984], Fortnow [1997], Lipton-Viglas, van Melkebeek, Williams, and others gives:

Theorem (R. Williams, 2007)

Let $c < 2\cos(\pi/7) \approx 1.8019$. *Then* $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$.

In this talk, we review these results and prove their optimality relative to currently known proof techniques.

Nepomnjasci's method

Definition

$$^{b}(\exists n^{c})^{d}\mathrm{DTS}(n^{e})$$

denotes the class of problems taking inputs of length $n^{b+o(1)}$, existentially choosing $n^{c+o(1)}$ bits, keeping in memory a total of $n^{d+o(1)}$ bits (using time $n^{c+o(1)}$) which are passed to a deterministic procedure that uses time $n^{e+o(1)}$ and space $n^{o(1)}$.

Theorem (by method of Nepomnjasci, 1970)

$^{b}DTS(n^{c}) \subseteq {}^{b}(\exists n^{x})^{\max\{b,x\}}(\forall n^{0})^{b}\mathrm{DTS}(n^{c-x})$.

Proof next page....

${}^{b}DTS(n^c) \subseteq {}^{b}(\exists n^x)^x(\forall n^0)^b\mathrm{DTS}(n^{c-x})$, for $x \geq b$

Proof idea: Split the n^c time computation into n^x many blocks. Existentially guess the memory contents (apart from the input) at each block boundary ($n^{x+o(1)}$ bits), then universally choose one block to verify correctness ($O(\log n) = n^{o(1)}$ universal choices), and simulate that block's computation (in n^{c-x} time).

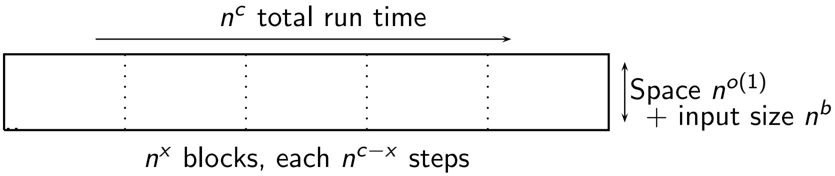

Alternation trading proofs [Williams]

An *alternation trading proof* is a proof that $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$, for some fixed $c \geq 1$. It is a proof by contradiction, based on deducing

$$\mathrm{DTS}(n^a) \subseteq \mathrm{DTS}(n^b)$$

for some $a > b$, from the assumption that $\mathrm{SAT} \in \mathrm{DTS}(n^c)$.

The lines of an alternation trading proof are of the form

$$^1(\exists n^{a_1})^{b_2}(\forall n^{a_2})^{b_3} \cdots {}^{b_k}(Qn^{a_k})^{b_{k+1}}\mathrm{DTS}(n^{a_{k+1}}).$$

There are two kinds of inferences: "speedup" inferences that add quntifiers and reduce run time (based on Nepomnjascii) and "slowdown" inferences that remove a quantifier and increase run time (based on Cook's theorem)....

The rules of inferences for alternation trading proofs are:

Initial speedup: $(x \leq a)$

$$^{1}\mathrm{DTS}(n^{a}) \subseteq {}^{1}(\exists n^{x})^{\max\{x,1\}}(\forall n^{0})^{1}\mathrm{DTS}(n^{a-x}),$$

Speedup: $(0 < x \leq a_{k+1})$

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTS}(n^{a_{k+1}}) \subseteq \cdots {}^{b_k}(\exists n^{\max\{x,a_k\}})^{\max\{x,b_{k+1}\}}(\forall n^{0})^{b_{k+1}}\mathrm{DTS}(n^{a_{k+1}-x}),$$

Slowdown:

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTS}(n^{a_{k+1}}) \subseteq \cdots {}^{b_k}\mathrm{DTS}(n^{\max\{cb_k,ca_k,cb_{k+1},ca_{k+1}\}}).$$

Example: alternation trading proof.

Let $1 < c < \sqrt{2}$. Then, if $\mathrm{SAT} \in \mathrm{DTS}(n^c)$,

$$\begin{aligned} \mathrm{DTS}(n^2) &\subseteq (\forall n^1)^1(\exists n^0)^1\mathrm{DTS}(n^1) \\ &\subseteq (\forall n^1)^1\mathrm{DTS}(n^c) \\ &\subseteq \mathrm{DTS}(n^{c^2}). \end{aligned}$$

which is a contradiction. Proof uses a speedup-slowdown-slowdown pattern, also denoted **100**.

This proves:

Theorem (Lipton-Viglas, 1999)

$\mathrm{SAT} \notin \mathrm{DTS}(n^{\sqrt{2}})$.

Better results can be found with more alternations.

Theorem (Fortnow, van Melkebeek, et. al)

$\mathrm{SAT} \notin \mathrm{DTS}(n^c)$, *where* $c < \phi \approx 1.618$, *the golden ratio.*

The optimal refutation with seven inferences derives:

Theorem (Williams)

$\mathrm{SAT} \notin \mathrm{DTS}(n^{1.6})$.

This proof was found with a Maple-based linear programming algorithm. It uses a pattern of inferences: **1100100**, where "**1**" denotes a speedup and "**0**" denotes a slowdown.

Theorem (Williams)

Let $c < 2\cos(\pi/7) \approx 1.801$. *Then* $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$.

This used proofs of the following $\mathbf{1}/\mathbf{0}$ patterns:

$$\mathbf{1}^n(\mathbf{10})^*(\mathbf{0}(\mathbf{10})^*)^n.$$

These were gleaned from patterns found with Maple experiments, and conjectured by Williams to be the best possible refutations.

In this talk, we show how to prove these conjectures, at least in the framework of currently known rules for alternation trading proofs.

Remark: If $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$ for all c, then NP $\notin$ L ("L"=logspace), something thought to be hard to prove.

L $\subseteq$ NL $\subseteq$ P $\subseteq$ NP $\subseteq$ PSPACE.

Main Theorem I

There are alternation trading proofs of $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$ for exactly the values $c < 2\cos(\pi/7)$.

Reduced alternation trading proofs

Two simplifications for a 'reduced" system:
1. Replace the superscripts "1" with "0".
2. Get rid of half the exponents! Replace each quantifier "$(Qn^{a_i})^{b_i}$" with just "Q^{b_i}".

The intuition is:

Firstly, that the values "1" can be made infinitesimal by making a_i's and b_i's large. Then the "1"s can be replaced by zeros.

Secondly, the a_i's are always dominated by the b_i's and thus are never important.

The simplified rules for alternation proofs become:

Initialization: $^{0}\mathrm{DTS}(n^a) \vdash {}^{0}\exists^{0}\mathrm{DTS}(n^a).$

Speedup: $(0 < x \leq a)$

$$\cdots {}^{b_k}\exists^{b_{k+1}}\mathrm{DTS}(n^a) \vdash \cdots {}^{b_k}\exists^{\max\{x,b_{k+1}\}}\forall^{b_{k+1}}\mathrm{DTS}(n^{a-x}),$$

Slowdown: $\cdots {}^{b_k}\exists^{b_{k+1}}\mathrm{DTS}(n^a) \vdash \cdots {}^{b_k}\mathrm{DTS}(n^{\max\{cb_k,cb_{k+1},ca\}}).$

Theorem

The reduced system has a refutation iff the original system has a refutation.

Approximate inference

Defn: Given Ξ and Ξ':

$$\begin{aligned} \Xi &= {}^{0}\exists^{b_2}\forall^{b_3} \cdots {}^{b_k} Q^{b_{k+1}} \mathrm{DTS}(n^a) \\ \Xi' &= {}^{0}\exists^{b'_2}\forall^{b'_3} \cdots {}^{b'_k} Q^{b'_{k+1}} \mathrm{DTS}(n^{a'}). \end{aligned}$$

$\Xi \le \Xi'$ means $a \le a'$ and each $b_i \le b'_i$.

The *weakening rule* allows inferring Ξ' from Ξ; deduction with weakening is denoted $\Xi \;\overset{w}{\vdash}\; \Xi'$. The weakening rule does not add any power to the proof system.

Defn: $(\Xi + \epsilon)$ is obtained from Ξ by increasing a and each b_i by ϵ.

Definition (Approximate inference, $\Vdash$)

$\Xi \Vdash \Lambda$ if and only if for all $\epsilon > 0$ there exists a $\delta > 0$ such that

$$(\Xi + \delta) \;\overset{w}{\vdash}\; (\Lambda + \epsilon).$$

Achievability

Definition

Let $\mu \geq 1$ and $0 < \nu$. The pair $\langle \mu, \nu \rangle$ is *c-achievable* provided that, for all values a, b and d satisfying $c\mu b = \nu d$,

$$^{a}\exists^{b}\mathrm{DTS}(n^{d}) \Vdash {}^{a}\exists^{\mu b}\mathrm{DTS}(n^{\nu d}).$$

Theorem

If $\langle \mu, \nu \rangle$ is c-achievable for $\nu < 1/c$, then $\mathrm{SAT} \notin \mathrm{DTS}(n^{c})$.

Pf :

$$
\begin{aligned}
{}^{0}\mathrm{DTS}(n^{1}) &\vdash {}^{0}\exists^{0}\mathrm{DTS}(n^{1}) && \text{Initialization}\\
&\vdash^{w} {}^{0}\exists^{\nu/(c\mu)}\mathrm{DTS}(n^{1}) && \text{Weakening}\\
&\Vdash {}^{0}\exists^{\nu/c}\mathrm{DTS}(n^{\nu}) && \text{By a } \langle \mu, \nu \rangle \text{ step}\\
&\vdash {}^{0}\mathrm{DTS}(n^{c\nu}) && \text{Slowdown}
\end{aligned}
$$

Note $c\nu < 1$. (Converse to proof holds too.)

Theorem

$\langle 1,\ c{-}1\rangle$ is c-achievable with $\mathbf{(10)}^*$ derivations

Pf. Let $\Xi = {}^a\exists^b\mathrm{DTS}(n^d)$, with $cb \le d$. Then

$$\Xi \vdash {}^a\exists^b\forall^b\mathrm{DTS}(n^{d-b}) \vdash {}^a\exists^b\mathrm{DTS}(n^{\max\{cb,c(d-b)\}}) = {}^a\exists^b\mathrm{DTS}(n^{d'}).$$

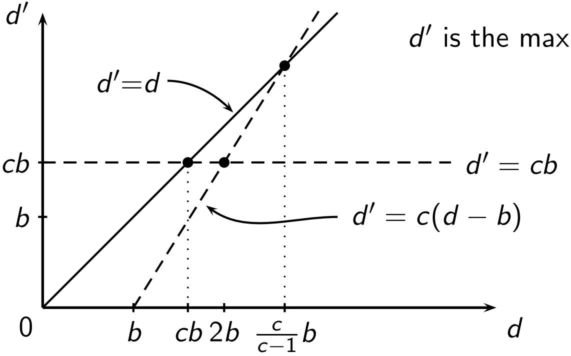

d' is the max of the dashed lines

"q.e.d."

Composition of c-achievable pairs

Theorem

Let $\langle \mu_1, \nu_1 \rangle$ and $\langle \mu_2, \nu_2 \rangle$ be c-achievable, with $c\nu_1\mu_2 \geq \mu_1$. Then $\langle \mu, \nu \rangle$ is c-achievable, where

$$\mu = c\nu_1\mu_2 \quad \text{and} \quad \nu = \frac{c\mu_1\nu_1\nu_2}{\mu_1 + \nu_1\nu_2}.$$

Pf idea: Use a speedup, followed by a $\langle \mu_2, \nu_2 \rangle$ step, then a slowdown, and finally a $\langle \mu_1, \nu_1 \rangle$ step. If $c\nu_1\mu_2 < \mu_1$, then theorem holds with $\mu = \max\{c\nu_1\mu_2, \mu_1\}$ instead.

Theorem

The constructions above "subsume" all alternation trading proofs. There is an alternation trading proof of $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$ iff an c-achievable pair with $\nu < 1/c$ can be constructed using the previous two theorems.

Understanding what is achievable

The expressions for μ and ν can be rewritten as:

$$\frac{1}{\mu} = \frac{1}{R}\left(\frac{1}{\mu_2}\right) \quad \text{and} \quad \frac{1}{\nu} = \frac{1}{T} - \frac{1}{R}\left(\frac{1}{T} - \frac{1}{\nu_2}\right).$$

where $\frac{1}{R} = \frac{1}{c\nu_1}$ and $\frac{1}{T} = \frac{\nu_1}{(c(\nu_1 - 1)\mu_1}$. Without loss of generality $\nu_1 > 1/c$ (otherwise we are done), and thus $\frac{1}{R} < 1$.

We think of $\langle \mu_1, \nu_1 \rangle$ as transforming $\langle \mu_2, \nu_2 \rangle$ to yield $\langle \mu, \nu \rangle$, and write this as

$$\langle \mu_1, \nu_1 \rangle : \langle \mu_2, \nu_2 \rangle \mapsto \langle \mu, \nu \rangle$$

This transformation makes μ_2 increase geometrically, and makes ν_2 contract inverse-geometrically towards T.

Define $\langle \mu_i, \nu_i \rangle$ by:

$$\begin{aligned} \langle \mu_0, \nu_0 \rangle &= \langle 1, c{-}1 \rangle, \\ \langle \mu_0, \nu_0 \rangle &: \langle \mu_i, \nu_i \rangle \mapsto \langle \mu_{i+1}, \nu_{i+1} \rangle. \end{aligned}$$

If

$$T_0 = \frac{(c\nu_0 - 1)\mu_0}{\nu_0} = \frac{c(c-1)-1}{c-1} < 1/c,$$

then some $\nu_i < 1/c$. This will give an alternation trading proof of $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$. For $1 \le c \le 2$, this is equivalent to

$$\boxed{c^3 - c^2 - 2c + 1 < 0,}$$

i.e., $c < 2\cos(\pi/7)$.

This gives the desired alternation trading proof that $\mathrm{SAT} \notin \mathrm{DTS}(n^{2\cos(\pi/7)})$. [Williams]

The next theorem states $c = 2\cos(\pi/7)$ is the best possible. A key point is that the attraction points "T" only increase.

Lemma

If $\langle \mu_1, \nu_1 \rangle : \langle \mu_2, \nu_2 \rangle \mapsto \langle \mu, \nu \rangle$ *and if* $T_1 \geq 1/c$*, then* $T \geq T_2$*.*

Theorem

There are alternation trading proofs of $\mathrm{SAT} \notin \mathrm{DTS}(n^c)$ *for exactly the values* $c < 2\cos(\pi/7)$*.*

Time-Space Tradeoff Lower Bounds

Definition

$\mathrm{DTISP}(n^c, n^\epsilon)$ is the class of problems decidable in deterministic time $n^{c+o(1)}$ and space $n^{\epsilon+o(1)}$.

The notion of alternation trading proofs can be expanded to give proofs that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$ for various values $1 \leq c < 2\cos(\pi/7)$ and $0 < \epsilon < 1$.

This is done by giving alteration trading proofs of

$$\mathrm{DTISP}(n^{\alpha c}, n^{\alpha\epsilon}) \subseteq \mathrm{DTISP}(n^{\beta c}, n^{\beta\epsilon})$$

for some $\alpha > \beta > 0$.

Rules of inference for DTISP

Initial speedup: $(e < x \leq a)$

$$^1\mathrm{DTISP}(n^a, n^e) \subseteq {}^1(\exists n^x)^{\max\{x,1\}}(\forall n^0)^{\max\{e,1\}}\mathrm{DTISP}(n^{a-x+e}, n^e)$$

Invoked only with $a = c \cdot e/\epsilon$.

Speedup: $(e < x \leq a_{k+1}.)$

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTISP}(n^{a_{k+1}}, n^e)$$
$$\subseteq \quad \cdots {}^{b_k}(\exists n^{\max\{x,a_k\}})^{\max\{x,b_{k+1}\}}(\forall n^0)^{\max\{b_{k+1},e\}}\mathrm{DTISP}(n^{a_{k+1}-x+e}, n^e),$$

Slowdown: Let $a = \max\{b_k, a_k, b_{k+1}, a_{k+1}\}$.

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTISP}(n^{a_{k+1}}, n^e) \subseteq \cdots {}^{b_k}\mathrm{DTISP}(n^{ca}, n^{\epsilon a}).$$

Based on extension of the theory of acheivable pairs to "acheivable triples", and on a computer-based search, aided by theorems about pruning the searches:

Theorem [Buss-Williams] The following pairs are the optimal values c and ϵ for which there are alternating trading proofs that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$.

ϵ	c
0.001	1.80083
0.01	1.79092
0.1	1.69618
0.25	1.55242
0.5	1.34070
0.75	1.15765
0.9	1.06011
0.99	1.00583
0.999	1.00058

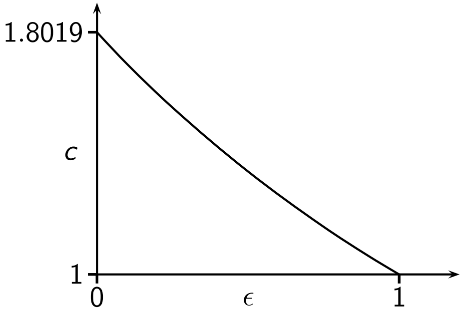

These values for c and ϵ are better than prior known lower bounds.

Open problems

- Find a closed form solution for the optimal $\mathrm{DTISP}(n^c, n^\epsilon)$ proofs. Even, find a simple characterization of how to construct the optimal proofs without resorting to a brute-force (pruned) search.
- There are many other flavors of alternation trading proofs, for instance for nondeterministic algorithms for tautologies. One could try giving proofs that the known alternation trading proofs are optimal.
- Most interesting: Try to find *new* principles that go beyond the presently known speedup and slowdown inferences, to give improved lower bound proofs.

Thank you!

Time-Space Tradeoff Lower Bounds

Definition

$\mathrm{DTISP}(n^c, n^\epsilon)$ is the class of problems decidable in deterministic time $n^{c+o(1)}$ and space $n^{\epsilon+o(1)}$.

We expand the notion of alternation trading proofs to give proofs that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$ for various values $1 \le c < 2\cos(\pi/7)$ and $0 < \epsilon < 1$.

This will be done by giving alteration trading proofs of

$$\mathrm{DTISP}(n^{\alpha c}, n^{\alpha\epsilon}) \subseteq \mathrm{DTISP}(n^{\beta c}, n^{\beta\epsilon})$$

for some $\alpha > \beta > 0$.

Rules of inference for DTISP

Initial speedup: $(e < x \leq a)$

$$^{1}\mathrm{DTISP}(n^a, n^e) \subseteq {}^{1}(\exists n^x)^{\max\{x,1\}}(\forall n^0)^{\max\{e,1\}}\mathrm{DTISP}(n^{a-x+e}, n^e)$$

Invoked only with $a = c \cdot e/\epsilon$.

Speedup: $(e < x \leq a_{k+1}.)$

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTISP}(n^{a_{k+1}}, n^e)$$
$$\subseteq \cdots {}^{b_k}(\exists n^{\max\{x,a_k\}})^{\max\{x,b_{k+1}\}}(\forall n^0)^{\max\{b_{k+1},e\}}\mathrm{DTISP}(n^{a_{k+1}-x+e}, n^e),$$

Slowdown: Let $a = \max\{b_k, a_k, b_{k+1}, a_{k+1}\}$.

$$\cdots {}^{b_k}(\exists n^{a_k})^{b_{k+1}}\mathrm{DTISP}(n^{a_{k+1}}, n^e) \subseteq \cdots {}^{b_k}\mathrm{DTISP}(n^{ca}, n^{\epsilon a}).$$

An equivalent simplified reduced system

Line now have the form

$$^{0}\exists^{b_1}\forall^{b_2}\exists^{b_3}\cdots{}^{b_k}Q^{b_{k+1}}\mathrm{DTISP}(n^a, n^e)$$

where $a \geq e$ and each $b_i \geq e$.

Initialization:

$$^{0}\mathrm{DTISP}^*(n^a, n^e) \vdash {}^{0}\exists^{e}\mathrm{DTISP}^*(n^a, n^e).$$

Speedup: $(e < x \leq a.)$

$$\cdots{}^{b_k}\exists^{b_{k+1}}\mathrm{DTISP}(n^a, n^e) \vdash \cdots{}^{b_k}\exists^{\max\{x,b_{k+1}\}}\forall^{b_{k+1}}\mathrm{DTISP}(n^{a-x+e}, n^e),$$

Slowdown: Let $a' = \max\{b_k, b_{k+1}, a\}$.

$$\cdots{}^{b_k}\exists^{b_{k+1}}\mathrm{DTISP}(n^a, n^e) \vdash \cdots{}^{b_k}\mathrm{DTISP}^*(n^{ca'}, n^{\epsilon a'}).$$

(c, ϵ)-Achievable Triples

Approximate inference, $\Vdash$, is defined similarly for DTISP as for DTS.

Achievable *triples* are defined as:

Definition

Let $\mu \geq 1$ and $0 < \nu < 1$ and $1 \leq \ell \in \mathbb{N}$. Then $\langle \mu, \nu, \ell \rangle$ is (c, ϵ)-achievable provided that, when b, d and e satisfy $(c+\epsilon)\mu b = \nu(d + \ell e)$ and $e \leq b \leq d$,

$$^{a}\exists^{b}\mathrm{DTISP}(n^{d}, n^{e}) \ \Vdash \ ^{a}\exists^{\mu b}\mathrm{DTISP}^{*}(n^{c\mu b}, n^{\epsilon \mu b}).$$

Definition

$$R = \rho(\mu, \nu, \ell) = \frac{c(c + \ell\epsilon)\nu}{c + \epsilon}.$$

Definition

If there is a (c, ϵ)-achievable triple with $c\mu\epsilon < 1$ and $R < 1$, then there is a alternation trading proof that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$.

Before (for DTS bounds), R was just $c\nu$. Now, unfortunately, R depends on ℓ, and ℓ will be increasing as (c, ϵ)-achievable triples are formed.

Theorem

$\langle 1, c{+}\epsilon{-}1, 1\rangle$ is (c,ϵ)-achievable with **(10)*** *derivations.*

Theorem

Let $\langle \mu_1, \nu_1, \ell_1\rangle$ and $\langle \mu_2, \nu_2, \ell_2\rangle$ be (c,ϵ)-achievable. Define

$$\begin{aligned} \mu &= \frac{c(c+\ell_1\epsilon)}{c+\epsilon}\nu_1\mu_2 \\ \nu &= \frac{c(c+\epsilon)(c+\ell_1\epsilon)\mu_1\nu_1\nu_2}{(c+\epsilon)^2\mu_1 + c(c+\ell_1\epsilon)\nu_1\nu_2} \\ \ell &= \ell_2 + 1. \end{aligned}$$

Suppose that $\mu \geq \mu_1$. Then $\langle \mu, \nu, \ell\rangle$ is (c,ϵ)-achievable.

(If $\mu < \mu_1$, use $\mu = \mu_1$ instead.)

Theorem:

The above two constructions "subsume" all possible alternation trading proofs. This there is an alternation trading proof that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$ iff there is a (c,ϵ)-achievable triple with $R<1$. [Recall $R = c(c+\ell\epsilon)\nu/(c+\epsilon)$.]

The second theorem can be rewritten as:

$$\frac{1}{\mu} = \frac{1}{R_1}\cdot\frac{1}{\mu_2} \quad \text{and} \quad \frac{1}{\nu} = \frac{1}{T_1} - \frac{1}{R_1}\left(\frac{1}{T_1} - \frac{1}{\nu_2}\right).$$

where $T_1 = (c+\epsilon)\mu_2(1-1/R_1)$.

Alternate expression for $1/\nu$

$$\frac{1}{\nu} = \frac{1}{(c+\epsilon)\mu_1} + \frac{1}{R_1\nu_2}.$$

Unfortunately, we are unable to give a closed form analysis for the best time-space tradeoff bounds. Instead, we resorted to an exhaustive computer search for all (c,ϵ)-achievable triples.

Our initial searches generated huge search domains and frequently failed to find optimal c/ϵ pairs, even after discarding "subsumed" (c,ϵ)-achievable triples.

Defn: A pair of (c,ϵ)-achievable triples τ_1' and τ_1'' *multi-subsume* a triple τ_1 provided $\ell_1', \ell'' \le \ell_1$ and their "alternate expression" lines dominate that of τ_1. That is, iff, for every ν_2, at least one τ_1' or τ_1'' gives a better ν value than τ_1 does.

Theorem

Any multi-subsumed triple may be pruned from the search space without loss of generality.

With this pruning of multi-subsumed triples, we are completely successful in our computer searches in finding the optimal pairs c and ϵ for which there are alternating trading proofs that $\mathrm{SAT} \notin \mathrm{DTISP}(n^c, n^\epsilon)$.

ϵ	c
0.001	1.80083
0.01	1.79092
0.1	1.69618
0.25	1.55242
0.5	1.34070
0.75	1.15765
0.9	1.06011
0.99	1.00583
0.999	1.00058

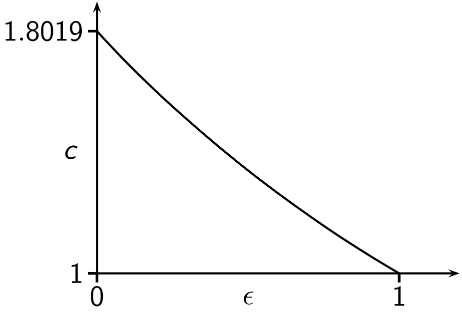

These values for c and ϵ are better than prior known values.

ϵ	c	Number of Rounds	Number of Triples	Has Refutation
0.001	1.80084	7	167	No
	1.80083	11	455	Yes
0.01	1.79093	20	764	No
	1.79092	11	278	Yes
0.1	1.69619	248	3633	No
	1.69618	26	435	Yes
0.25	1.55242	249	2932	No
	1.55242	33	297	Yes
0.5	1.34071	203	1533	No
	1.34070	44	406	Yes
0.75	1.15766	155	1379	No
	1.15765	27	167	Yes
0.9	1.06012	146	454	No
	1.06011	19	88	Yes
0.99	1.00584	99	260	No
	1.00583	7	20	Yes
0.999	1.00059	3	3	No
	1.00058	24	10	Yes

Open problems

- Find a closed form solution for the optimal $\mathrm{DTISP}(n^c, n^\epsilon)$ proofs. Even, find a simple characterization of how to construct the optimal proofs without resorting to a brute-force (pruned) search.
- There are many other flavors of alternation trading proofs, for instance for nondeterministic algorithms for tautologies. One could try giving proofs that the known alternation trading proofs are optimal.
- Most interesting: Try to find *new* principles that go beyond the presently known speedup and slowdown inferences, to give improved lower bound proofs.