

The Opacity of Backbones

Lane A. Hemaspaandra

Department of Computer Science
University of Rochester
Rochester, NY 14627, USA

David E. Narváez

College of Computing and Information Sciences
Rochester Institute of Technology
Rochester, NY 14623, USA

Abstract

A backbone of a boolean formula F is a collection S of its variables for which there is a unique partial assignment a_S such that $F[a_S]$ is satisfiable (Monasson et al. 1999; Williams, Gomes, and Selman 2003). This paper studies the nontransparency of backbones. We show that, under the widely believed assumption that integer factoring is hard, there exist sets of boolean formulas that have obvious, non-trivial backbones yet finding the values, a_S , of those backbones is intractable. We also show that, under the same assumption, there exist sets of boolean formulas that obviously have large backbones yet producing such a backbone S is intractable. Further, we show that if integer factoring is not merely worst-case hard but is frequently hard, as is widely believed, then the frequency of hardness in our two results is not too much less than that frequency.

1 Introduction

An important concept in the study of the SAT problem is the notion of backbones. The term was first used by Monasson et al. (1999), and the following formal definition was provided by Williams, Gomes, and Selman (2003).

Definition 1. *Let F be a boolean formula. A collection S of the variables of F is said to be a backbone if there is a unique partial assignment a_S such that $F[a_S]$ is satisfiable.*

In that definition, a_S assigns a value (true or false) to each variable in S , and $F[a_S]$ is a shorthand meaning F except with each variable in S assigned the value specified for it in a_S . A backbone S is *nontrivial* if $S \neq \emptyset$. The *size* of a backbone S is the number of variables in S . For a backbone S (for formula F), we say that a_S is the *value* of the backbone S .

For example, every satisfiable formula has the trivial backbone $S = \emptyset$. The formula $x_1 \wedge \overline{x_2}$ has four backbones, \emptyset , $\{x_1\}$, $\{x_2\}$, and $\{x_1, x_2\}$, with respectively the values (listing values as bit-vectors giving the assignments in the lexicographical order of the names of the variables in S) ϵ , 1, 0, and 10. The formula $x_1 \vee \overline{x_2}$ has no nontrivial backbones. (Every formula that has a backbone will have a maximum backbone—a backbone that every other backbone is a

subset of. Backbone variables have been called “frozen variables,” because each of them is the same over all satisfying assignments.)

As Williams, Gomes, and Selman (2003) note, “backbone variables are useful in studying the properties of the solution space of a... problem.”¹

And that surely is so. But it is natural to hope to go beyond that and suspect that if formulas have backbones, we can use those to help SAT solvers. After all, if one is seeking to get one’s hands on a satisfying assignment of an F that has a backbone, one need but substitute in the value of the backbone to have put all its variables to bed as to one’s search, and thus to “only” have all the other variables to worry about.

The goal of the present paper is to understand, at least in a theoretical sense, the difficulty of—the potential obstacles to doing—what we just suggested. We will argue that even for cases when one can quickly (i.e., in polynomial time) recognize that a formula has at least one nontrivial backbone, it can be intractable to find one such backbone. And we will argue that even for cases when one can quickly (i.e., in polynomial time) find a large, nontrivial backbone, it can be intractable to find the value of that backbone. In particular, we will show that if integer factoring is hard, then both the just-made claims hold. Integer factoring is widely believed to be hard; indeed, if it were in polynomial time, RSA (the Rivest-Shamir-Adleman cryptosystem) itself would fall.

In fact, integer factoring is even believed to be hard on average. And we will be inspired by that to go beyond the strength of the results mentioned above. Regarding our results mentioned above, one might worry that the “intractability” might be very infrequent, i.e., merely a rare, worst-case behavior. But we will argue that if integer factoring—or indeed any problem in $\text{NP} \cap \text{coNP}$ —is frequently hard, then the bad behavior types we mention above happen “almost” as often: If the frequency of hardness of integer factoring

¹We mention in passing that backbones also are very interesting for problems even harder than SAT, such as model counting (see Gomes, Sabharwal, and Selman, 2009)—i.e., #SAT (Valiant 1979): counting the number of satisfying assignments of a propositional boolean formula—especially given the currently huge run-time differences between SAT-solvers and model counters. After all, a backbone for a formula helpfully pinpoints into a particular subspace all the solutions that one is seeking to count.

is $d(n)$ for strings up to length n , then for some $\epsilon > 0$ the frequency of hardness of our problems is $d(n^\epsilon)$.

None of this means that backbones are not an excellent, important concept. Rather, this is saying—proving, in fact, assuming that integer factoring is as hard as is generally believed—that although the definition of backbone is merely about a backbone existing, one needs to be aware that going from a backbone existing to finding a backbone, and going from having a backbone to knowing its value, can be computationally expensive challenges.

2 Results

Section 2.1 will formulate our results without focusing on density. Then in Section 2.2 we will discuss how the frequency of hardness of sets of the type we have discussed is related to that of the sets in $\text{NP} \cap \text{coNP}$ having the highest frequencies of hardness.

The present section focuses only on presenting the results and what they mean. We will provide proofs in Section 3.

2.1 Basic Results

We first look at whether there can be simple sets of formulas for which one can easily compute/obtain a nontrivial backbone, yet one cannot easily find the value of that backbone.

Our basic result on this is stated below as Theorem 1. In this and most of our results, we state as our hypothesis not that “integer factoring cannot be done in polynomial time,” but rather that “ $\text{P} \neq \text{NP} \cap \text{coNP}$.” This in fact makes our claims stronger ones than if they had as their hypotheses “integer factoring cannot be done in polynomial time,” since it is well-known (because the decision version of integer factorization is itself in $\text{NP} \cap \text{coNP}$) that “integer factoring cannot be done in polynomial time” implies “ $\text{P} \neq \text{NP} \cap \text{coNP}$.” SAT will, as usual, denote the set of satisfiable (propositional) boolean formulas. (We do not assume that SAT by definition is restricted to CNF formulas.)

Theorem 1. *If $\text{P} \neq \text{NP} \cap \text{coNP}$, then there exists a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas such that:*

1. *There is a polynomial-time computable function f such that $(\forall F \in A)[f(F)$ outputs a nontrivial backbone of F].*
2. *There does not exist any polynomial-time computable function g such that $g(F)$ computes the value of backbone $f(F)$.*

Theorem 1 remains true even if one restricts the backbones found by f to be of size 1. We state that, in a slightly more general form, as follows.

Theorem 2. *Let $k \in \{1, 2, 3, \dots\}$. If $\text{P} \neq \text{NP} \cap \text{coNP}$, then there exists a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas such that:*

1. *There is a polynomial-time computable function f such that $(\forall F \in A)[f(F)$ outputs a size- k backbone of F].*
2. *There does not exist any polynomial-time computable function g such that $g(F)$ computes the value of backbone $f(F)$.*

Now let us turn to the question of whether, when it is obvious that there is at least one nontrivial backbone, it can be

hard to efficiently produce a nontrivial backbone. The following theorem shows that, if integer factoring is hard, the answer is yes.

Theorem 3. *If $\text{P} \neq \text{NP} \cap \text{coNP}$, then there exists a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas (each having at least one variable) such that:*

1. *Each formula $F \in A$ has a backbone whose size is at least 49% of F 's total number of variables.*
2. *There does not exist any polynomial-time computable function g such that, on each $F \in A$, $g(F)$ outputs a backbone whose size is at least 49%—or even at least 2%—of F 's variables.*

The 49% and 2% above are not at all magic, but are just for concreteness. It is easy to see that our proof that establishes the above theorem is in fact tacitly establishing the following more general claim (note: the smaller the ϵ the stronger the claim, and so the fact that below we speak only of $\epsilon \leq 1$ is not a weakness). The above theorem is the $\epsilon = 1$ case.

Theorem 4. *For each fixed ϵ , $0 < \epsilon \leq 1$, the following claim holds. If $\text{P} \neq \text{NP} \cap \text{coNP}$, then there exists a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas (each having at least one variable) such that:*

1. *Each formula $F \in A$ has a backbone whose size is at least $(50 - \epsilon)\%$ of F 's total number of variables.*
2. *There does not exist any polynomial-time computable function g such that, on each $F \in A$, $g(F)$ outputs a backbone whose size is at least $(2\epsilon)\%$ of F 's variables.*

2.2 Frequency of Hardness

A practical person might worry about results of the previous section in the following way. (Here, $|F|$ will denote the number of bits in the representation of F .) “Just because something is hard, doesn’t mean it is hard *often*.” For example, consider Theorem 3. Perhaps there is a polynomial-time function g' that, though it on infinitely many $F \in A$ fails to compute the value of the backbone $f(F)$, has the property that for each $F \in A$ for which it fails it then is correct on the (in lexicographical order) next $2^{2^{2^{|F|}}}$ elements of A . In this case, the theorem is indeed true, but it is a worst-case extreme that doesn’t recognize that in reality the errors may be few and far—very, very far—between.”

In this section, we address that reasonable worry. We show that if *even one problem* in $\text{NP} \cap \text{coNP}$ is frequently hard, then the sets in our previous sections can be made “almost” as frequently hard, in a sense of “almost” that we will make formal and specific. Since it is generally believed—for example due to the generally believed typical-case hardness of integer factoring—that there are sets in $\text{NP} \cap \text{coNP}$ that are quite frequently hard, it follows that the $2^{2^{2^{|F|}}}$ behavior our practical skeptic was speculating about cannot happen. Or at least, if that behavior did happen, then that would imply that every single problem in $\text{NP} \cap \text{coNP}$ has polynomial-time heuristic algorithms that make extraordinarily few errors.

Note that no one currently knows for sure how frequently-hard problems in $\text{NP} \cap \text{coNP}$ can be. But our results are showing that, whatever that frequency is, sets of the sort we've been constructing are hard "almost" as frequently. This can at first seem a bit of a strange notion to get one's head around, especially as complexity theory often doesn't pay much attention to frequency-of-hardness issues (though such issues in complexity theory can be traced back at least as far as the work of Schöning, 1986). But this actually is analogous to something every computer science researcher knows well, namely, NP-completeness. No one today knows for sure whether any NP problems are not in P. But despite that the longstanding NP-completeness framework lets one right now, today, prove clearly for specific problems that *if* any NP problem is not in P then that specific problem is not in P. The results of this section are about an analogous type of argument, except regarding frequency of hardness.

We now give our frequency-of-hardness version of Theorem 1. A claim is said to hold *for almost every* n if there exists an n_0 beyond which the claims always holds, i.e., the claim fails at most at a finite number of values of n . (In the theorems of this section, n 's universe is the natural numbers, $\{0, 1, 2, \dots\}$.)

Theorem 5. *If h is any nondecreasing function and for some $B \in \text{NP} \cap \text{coNP}$ it holds that each polynomial-time algorithm, viewed as a heuristic algorithm for testing membership in B , for almost every n (respectively, for infinitely many n) errs on at least $h(n)$ of the strings whose length is at most n , then there exist an $\epsilon > 0$ and a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas such that:*

1. *There is a polynomial-time computable function f such that $(\forall F \in A)[f(F)$ outputs a nontrivial backbone of F].*
2. *Each polynomial-time computable function g will err (i.e., will fail to compute the value of backbone $f(F)$), for almost every n (respectively, for infinitely many n), on at least $h(n^\epsilon)$ of the strings in A of length at most n .*

The precisely analogous result holds for Theorem 2. The analogous result also holds for Theorem 3, and we state that as the following theorem.

Theorem 6. *If h is any nondecreasing function and for some $B \in \text{NP} \cap \text{coNP}$ it holds that each polynomial-time algorithm, viewed as a heuristic algorithm for testing membership in B , for almost every n (respectively, for infinitely many n) errs on at least $h(n)$ of the strings whose length is at most n , then there exist an $\epsilon > 0$ and a set $A \in \text{P}$, $A \subseteq \text{SAT}$, of boolean formulas such that:*

1. *Each formula $F \in A$ has a backbone whose size is at least 49% of F 's total number of variables.*
2. *Each polynomial-time computable function g will err (i.e., will fail to compute a set of size at least 2% of F 's variables that is a backbone of F), for almost every n (respectively, for infinitely many n), on at least $h(n^\epsilon)$ of the strings in A of length at most n .*

What the above theorems say, looking at the contrapositives to the above results, is that if any of our above cases have polynomial-time heuristic algorithms that don't make errors too frequently, then *every single set* in $\text{NP} \cap \text{coNP}$

(even those related to integer factoring) has polynomial-time heuristic algorithms that don't make errors too frequently.

To make the meaning of the above results clearer, and to be completely open with our readers, it is important to have a frank discussion about the effect of the " ϵ " in the above results. Let us do this in two steps. First, we give as concrete examples two central types of growth rates that fall between polynomial and exponential. And second, we discuss how innocuous or noninnocuous the " ϵ " above is.

As to our examples, suppose that for some fixed $c > 0$ a particular function $h(n)$ satisfies $h(n) = 2^{\Omega((\log n)^c)}$. Note that for each fixed $\epsilon > 0$, it holds that the function $h'(n)$ defined as $h(n^\epsilon)$ itself satisfies the same bound, $h'(n) = 2^{\Omega((\log n)^\epsilon)}$. (Of course, the constant implicit in the " Ω " potentially has become smaller in the latter case.) Similarly, suppose that for some fixed $c > 0$ a particular function $h(n)$ satisfies $h(n) \geq 2^{n^c}$. Then for each fixed $\epsilon > 0$ it will hold that there is a value $c' > 0$, namely, $c' = \epsilon c$, such that $h(n^\epsilon) \geq 2^{n^{c'}}$.

The above at a casual glance might suggest that the weakening of the frequency claims between the most frequently hard problems in $\text{NP} \cap \text{coNP}$ and our problems is a "mere" changing of a constant. In some sense it is, but constants that are standing on the shoulders of exponents have more of a kick than constants sitting on the ground floor. And so as a practical matter, the difference in the actual numbers when one substitutes in for them can be large. On the other hand, polynomial-time reductions sit at the heart of computer science's formalization of its problems, and density distortions from n to n^ϵ based on the stretching of reductions are simply inherent in the standard approaches of theory, since those are the distortions one gets due to polynomial-time reductions being able to stretch their inputs to length $n^{1/\epsilon}$, i.e., polynomially. For example, it is well known that if B is an NP-complete set, then for every $\epsilon > 0$ it holds that B is polynomial-time isomorphic (which is theoretical computer science's strongest standard notion of them being "essentially the same problem") to some set B' that contains at most 2^{n^ϵ} strings at each length n .

Simply put, the "almost" in our "almost as frequent" claims is the natural, strong claim, judged by the amounts of slack that in theoretical computer are considered innocuous. And the results do give insight into how much the density does or does not change, e.g., the first above example shows that quasi-polynomial lower bounds on error frequency remain quasi-polynomial lower bounds on error frequency. However, on the other hand, there is a weakening, and even though it is in a "constant," that constant is in an exponent and so can alter the numerical frequency quite a bit.²

²In reality, how nastily small will the " ϵ " be? From looking inside the proofs of the results and thinking hard about the lengths of the formulas involved in the proofs (and resulting from the "Galil" version of Cook-Karp-Levin's Theorem that we will be discussing in the next section), one can see that ϵ 's value is primarily controlled by the running time of the NP machines for the NP sets B and \bar{B} from the theorems of this section. If those machines run in time around n^k , then ϵ will vary, viewed as a function of k , roughly

3 Proofs

We now provide the proofs or proof sketches for our results. We aim more for communication than for extremely detailed rigor. However, readers not interested in proofs may wish to skip this section.

3.1 Proofs for Section 2.1

We will prove all three of the theorems of Section 2.1 hand-in-hand, in a rather narrative fashion, as they share a framework. Each of that section's theorems starts with the assumption that $P \neq NP \cap coNP$. So let B be some set instantiating that, i.e., $B \in (NP \cap coNP) - P$. As all students learn when learning that SAT is NP-complete, we can efficiently transform the question of whether a machine accepts a particular string into a question about whether a certain boolean formula is satisfiable (Cook 1971; Karp 1972; Levin 1975). The original work that did that did not require (and did not *need* to require) that the thus-created boolean formula transparently revealed what machine and input had been the input to the transformation. But it was soon noted that one can ensure that the formula mapped to transparently reveals the machine and input that were the input to the transformation; see Galil (1974).

Galil's insight can be summarized in the following strengthened version of the standard claim regarding the so-called Cook-Karp-Levin Reduction. Let N_1, N_2, \dots be a fixed, standard enumeration of clocked, polynomial-time Turing machines, and w.l.o.g. assume that N_i runs within time $n^i + i$ on inputs of length n , and that N_i and i are polynomially related in size and easily obtained from each other. There is a function $r_{Galil-Cook}$ (for conciseness, we are writing Galil-Cook rather than Galil-Cook/Karp/Levin, although this version is closer to the setting of Karp and Levin than to that of Cook, since Cook used Turing reductions rather than many-one reductions) such that

1. For each N_i and x : $x \in L(N_i)$ if and only if $r_{Galil-Cook}(N_i, x) \in SAT$.
2. There is a polynomial p such that $r_{Galil-Cook}(N_i, x)$ runs within time polynomial (in particular, with p being the polynomial) in $|N_i|$ and $|x|^i + i$.
3. There is a polynomial-time function s such that for each N_i and x , $s(r_{Galil-Cook}(N_i, x))$ outputs the pair (N_i, x) .

We will be using two separate applications of the r function in our construction. But we need those two applications to be variable-disjoint. We will need this as otherwise we'd have interference with some of our claims about sizes of backbones and which variables are fixed and how many variables we have. These are requirements not present in any previous work that used the r function of Galil-Cook. We also will want to be able to have some literal names (in particular, “z”-using literal names of the form $z_\ell, z'_\ell, \bar{z}_\ell$, or \bar{z}'_ℓ , for all ℓ) available to us that we know are not part of the output of any application of the Galil-Cook r function; we need them as our construction involves not just two applications of the r function but also some additional variables. We

as (some constant times) the inverse of k .

can accomplish all the special requirements just mentioned as follows. We will, w.l.o.g., assume that in the output of the Galil-Cook function $r_{Galil-Cook}(N_i, x)$, every variable is of the form x_j (the x there is not a generic example of a letter, but really means the letter “x” just a “z” earlier really means the letter “z”), where j itself, when viewed as a pair of integers via the standard fixed correspondence between \mathbb{Z}^+ and $\mathbb{Z}^+ \times \mathbb{Z}^+$, has N_i as its first component or actually, to be completely precise, the natural number corresponding to N_i in the standard fixed correspondence between positive integers and strings. Though not all implementations of the Galil-Cook r function need have this property (and in fact, none has previously satisfied it as far as we know), we claim that one can implement a legal Galil-Cook r function in such a way that it has this property yet still has the property that this r function will have a polynomial-time inversion function s satisfying the behavior for s mentioned above. (For those wanting more information on how such a function $r_{Galil-Cook}(N_i, x)$ can be implemented that has all the properties claimed above, and indeed seeing such an implementation is important to best understanding what is going on here, we have made available a detailed construction we have built that accomplishes this, in the full technical report version of this paper, which is available online as Hemaspaandra and Narváez, 2016; the construction involves exploiting encoding headroom that one can in some sense piggyback on top of an arbitrary implementation of the Cook-Karp-Levin Reduction.)

We now can specify the sets A needed by the theorems of Section 2.1. Recall we have (thanks to the assumptions of the theorems) fixed a set $B \in (NP \cap coNP) - P$. $B \in NP$ so let i be a positive integer such that N_i is a machine from the abovementioned standard enumeration such that $L(N_i) = B$. $\bar{B} \in NP$ so let j be a positive integer such that N_j is a machine from the abovementioned standard enumeration, such that $L(N_j) = \bar{B}$. Fix any positive integer k . Then for the case of that fixed value k , the set A of Theorem 2 is as follows: $A_{3,k} = \{(z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (r_{Galil-Cook}(N_i, x))) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (r_{Galil-Cook}(N_j, x))) \mid x \in \Sigma^*\}$. One must keep in mind in what follows that, as per the previous paragraph, $r_{Galil-Cook}$ never outputs literals with names involving subscripted z s or z' s and the outputs of $r_{Galil-Cook}(N_i, x)$ and $r_{Galil-Cook}(N_j, x)$ share no variable names (since $i \neq j$).

Let us argue that $A_{3,k}$ indeed satisfies the requirements of the A for the “ k ” case of Theorem 2.

$A \in P$: Given a string y whose membership in A we are testing, we make sure y syntactically matches the form of the elements of A (i.e., elements of $A_{3,k}$). If it does, we then check that its k matches our k , and we use s to get decoded pairs (i', x') and (j'', x'') from the places in our parsing of y where we have formulas—call them F_{left} and F_{right} —that we are hoping are the outputs of the r function. That is, if our input parses as $(z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (F_{left})) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (F_{right}))$, then if $s(F_{left})$ gives $(N_{i'}, x')$ our decoded pair is (i', x') , and F_{right} is handled analogously. We also check to make sure that $x' = x'', i = i'$, and $j = j''$. If anything mentioned so far fails, then $y \notin A$. Otherwise,

we check to make sure that $r_{\text{Galil-Cook}}(N_i, x') = F_{\text{left}}$ and $r_{\text{Galil-Cook}}(N_j, x') = F_{\text{left}}$, and reject if either equality fails to hold. (Those checks are *not* superfluous. s by definition has to correctly invert on strings that are the true outputs of $r_{\text{Galil-Cook}}$, but we did not assume that s might not output sneaky garbage when given other input values, and since F_{left} and F_{right} are coming from our arbitrary input y , they could be anything. However, the check we just made defangs the danger just mentioned.) If we have reached this point, we indeed have determined that $y \in A$, and for each $y \in A$ we will successfully reach this point.

$A \subseteq \text{SAT}$: For each x , either $x \in B$ or $x \notin B$. In the former case ($x \in B$), $r_{\text{Galil-Cook}}(N_i, x) \in \text{SAT}$ and so the left disjunct of $(z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (r_{\text{Galil-Cook}}(N_i, x))) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (r_{\text{Galil-Cook}}(N_j, x)))$ can be made true using that satisfying assignment and setting each z_ℓ to true. On the other hand, if $x \notin B$, then $r_{\text{Galil-Cook}}(N_j, x) \in \text{SAT}$ and so the whole formula can be made true using that satisfying assignment and setting each z_ℓ to false.

There is a polynomial-time computable function f such that $(\forall F \in A)[f(F)$ outputs a nontrivial backbone of F]: On input $F \in A$, f will simply output $\{z_1, z_2, \dots, z_k\}$, which is a nontrivial backbone of F . Why is it a nontrivial backbone? If the x embedded in F satisfies $x \in B$, then not only does $r_{\text{Galil-Cook}}(N_i, x) \in \text{SAT}$ hold, but also $r_{\text{Galil-Cook}}(N_j, x) \notin \text{SAT}$ must hold (since otherwise we would have $x \notin B \wedge x \in B$, an impossibility). So if the x embedded in F satisfies $x \in B$, then there are satisfying assignments of $(z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (r_{\text{Galil-Cook}}(N_i, x))) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (r_{\text{Galil-Cook}}(N_j, x)))$, and every one of them has each z_ℓ set to true. Similarly, if the x embedded in F satisfies $x \notin B$, then our long formula has satisfying assignments, and every one of them has each z_ℓ set to false. Thus $\{z_1, z_2, \dots, z_k\}$ indeed is a size- k backbone.

There does not exist any polynomial-time computable function g such that $g(F)$ computes the value of backbone $f(F)$: Suppose by way of contradiction that such a polynomial-time computable function g does exist. Then we would have that $B \in \text{P}$, by the following algorithm. Let f be the function constructed in the previous paragraph, i.e., the one that outputs $\{z_1, z_2, \dots, z_k\}$ when $F \in A$. Given x , in polynomial time— g and f are polynomial-time computable, and although r in general is not since its running time's polynomial degree varies with its first argument and so is not uniformly polynomial, r here is used only for the first-component values N_i and N_j and under that restriction it indeed is polynomial-time computable—compute $g(f((z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (r_{\text{Galil-Cook}}(N_i, x))) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (r_{\text{Galil-Cook}}(N_j, x)))))$. This must either tell us that the z_ℓ s are true in all satisfying assignments, which tells us that it is the left disjunct that is satisfiable and thus $x \in B$, or it will tell us that the z_ℓ s are false in all satisfying assignments, from which we similarly can correctly conclude that $x \notin B$. So $B \in \text{P}$, yet we chose B so as to satisfy $B \in (\text{NP} \cap \text{coNP}) - \text{P}$. Thus our assumption that such a g exists is contradicted.

That ends our proof of Theorem 2—and so implicitly also of Theorem 1, since Theorem 1 follows immediately from Theorem 2.

Having seen the above proof, the reader will not need a detailed treatment of the proof of Theorem 3. Rather, we describe how to convert the above construction into one that proves Theorem 3. Recall that for the “ k ” case of Theorem 2 our set A was $\{(z_1 \wedge z_2 \wedge \dots \wedge z_k \wedge (r_{\text{Galil-Cook}}(N_i, x))) \vee (\bar{z}_1 \wedge \bar{z}_2 \wedge \dots \wedge \bar{z}_k \wedge (r_{\text{Galil-Cook}}(N_j, x))) \mid x \in \Sigma^*\}$.

For Theorem 3, let us use almost the same set. Except we will make two types of changes. First, in the above, replace the two occurrences of k each with the smallest positive integer m' satisfying

$$\frac{\text{numvars}(r_{\text{Galil-Cook}}(N_i, x)) + \text{numvars}(r_{\text{Galil-Cook}}(N_j, x)) + 2m'}{m'} \geq \frac{49}{100},$$

where numvars counts the number of variables in a formula, e.g., $\text{numvars}(\bar{x}_1 \wedge x_2 \wedge \bar{x}_2) = 2$, due to the variables x_1 and x_2 . Let m henceforward denote that value, i.e., the smallest (positive integer) m' that satisfies the above equation. Second, in the right disjunct, change each \bar{z}_ℓ to z'_ℓ .

Note that if $x \in B$, then $\{z_1, z_2, \dots, z_m\}$ is a backbone whose value is the assignment of true to each variable, and that contains at least 49% of the variables in the formula that x put into A . Similarly, if $x \notin B$, then $\{z'_1, z'_2, \dots, z'_m\}$ is a backbone whose value is the assignment of false to each variable, and that contains at least 49% of the variables in the formula that x put into A . It also is straightforward to see that our thus-created set A belong to P and satisfies $A \subseteq \text{SAT}$.

So the only condition of Theorem 3 that we still need to show holds is the claim that, for the just-described A , there does not exist any polynomial-time computable function g such that, on each $F \in A$, $g(F)$ outputs a backbone whose size is at least 2% of F 's variables. Suppose by way of contradiction that such a function g does exist. We claim that would yield a polynomial-time algorithm for B , contradicting the assumption that $B \notin \text{P}$. Let us give such a polynomial-time algorithm. To test whether $x \in B$, in polynomial time we create the formula in A that is put there by x , and we run our postulated polynomial-time g on that formula, and thus we get a backbone, call it S , that contains at least 2% of F 's variables. Note that we ourselves do not get to choose which large backbone g outputs, so we must be careful as to what we assume about the output backbone. We in particular certainly cannot assume that g happens to always output either $\{z_1, z_2, \dots, z_m\}$ or $\{z'_1, z'_2, \dots, z'_m\}$. But we don't need it to. Note that the two backbones just mentioned are variable-disjoint, and each contains 49% of F 's variables.

Now, there are two cases. One case is that S contains at least one variable of the form z_ℓ or z'_ℓ . In that case we are done. If it contains at least one variable of the form z_ℓ then $x \in B$. Why? If $x \in B$, then the left-hand disjunct of the formula x puts into A is satisfiable and the right-hand disjunct is not. From the form of the formula, it is clear that each z_ℓ is always true in each satisfying assignment in this case, yet that for each z'_ℓ there are satisfying assignments where z'_ℓ is true and there are satisfying assignments where z'_ℓ is false. So if $x \in B$, no z'_ℓ can belong to any backbone.

By analogous reasoning, if S contains at least one variable of the form z'_ℓ then $x \notin B$. (It follows from this and

the above that S cannot possibly contain at least one variable that is a subscripted z and at least one variable that is a subscripted z' , since then x would have to simultaneously belong and not belong to B .)

The final case to consider is the one in which S does not contain at least one variable of the form z_ℓ or z'_ℓ . We argue that this case cannot happen. If this were to happen, then every variable of F other than the variables $\{z_1, z_2, \dots, z_m, z'_1, z'_2, \dots, z'_m\}$ must be part of the backbone, since S must involve 2% of the variables and $\{z_1, z_2, \dots, z_m, z'_1, z'_2, \dots, z'_m\}$ comprise 98% of the variables. But that is impossible. We know that the variables used in $r_{\text{Galil-Cook}}(N_i, x)$ and $r_{\text{Galil-Cook}}(N_j, x)$ are disjoint. So the variables in the one of those two that is not the one that is satisfiable can and do take on any value in some satisfying assignment, and so cannot be part of any backbone. (The only remaining worry is the case where one of $r_{\text{Galil-Cook}}(N_i, x)$ or $r_{\text{Galil-Cook}}(N_j, x)$ contains no variables. However, the empty formula is by convention considered illegal, in cases such as here where the formulas are not considered to be trapped into DNF or CNF. There is a special convention regarding empty DNF and CNF formulas, but that is not relevant here.)

We have thus concluded the proof of Theorem 3.

3.2 Proof Sketches for Section 2.2

We will treat this section briefly and informally, as we will argue that its claims can be seen as following from the previous section's proofs.

The crucial thing to note is that the mapping from strings x (as to whether they belong to B) into the string that x puts into A is (a) polynomial-time computable (and so the one string that x puts into A is at most polynomially longer than x), and (b) one-to-one.

So, any collection of m instances up to a given length n that fool a particular polynomial-time algorithm for B are associated with at least m distinct instances in A all of length at most n^q (where the polynomial bound on the length of the formula that x puts into A is that it is of length at most n^{q^3}). So if one had an algorithm for the “ A ” set such that the algorithm had at most m' errors on the strings up to length n^q , it would certainly imply an algorithm for B that up to length $n^{1/q}$ made at most m' errors. Namely, one's heuristic of that form for B would be to take x , map it to the string it put into A , and then run the heuristic for A on that string.

The above discussion establishes what the results in Section 2.2 are asserting.

³We have for simplicity in this brief analysis left out any lower-order terms and the leading-term constant, but that is legal except at $n \in \{0, 1\}$ —since starting with $n = 2$ we can boost q if needed—and no finite set of values, such as $\{0, 1\}$ can cause problems to our theorem, as it is about the “infinitely-often” and “almost-everywhere” cases. However, such boosting does potentially interfere with the inverse-of- k relation mentioned in Footnote 2, and so if we wanted to maintain that, we would in this argument instead use a lowest-degree-possible monotonic polynomial bounding the growth rate.

4 Related Work

Our results can be viewed as part of a line of work that is so underpopulated as to barely merit being called a line of work, at least regarding its connections to AI. The true inspiration for this work was a paper of Alan Demers and Allan Borodin (1976) from the 1970s that never appeared in any form other than as a technical report. Though quite technical, that paper in effect showed sufficient conditions for creating simple sets of satisfiable formulas such that it was unclear why they were satisfiable.

Even in the theoretical computer science world, where Borodin and Demers's work is set, the work has been very rarely used. In particular, it has been used to get characterizations regarding unambiguous computation (Hartmanis and Hemaspaandra 1988), and Rothe and his collaborators have used it in various contexts to study the complexity of certificates (Hemaspaandra, Rothe, and Wechsung 1997; Rothe 1999), see also Fenner et al. (2003) and Valiant (1976).

There has been just one paper that previously has sought to bring the focus of this line to a topic of interest in AI. Although it appeared in a theoretical computer science venue, the work of Hemaspaandra, Hemaspaandra, and Menton (2013) shows that some problems from computational social choice theory, a subarea of multiagent systems, have the property that if $P \neq NP \cap \text{coNP}$ then their search versions are not polynomial-time Turing reducible to their decision problems—a rare behavior among the most familiar seemingly hard sets in computer science, since so-called self-reducibility (Meyer and Paterson 1979) is known to preclude that possibility for most standard NP-complete problems. The key issue that 2013 paper left open is whether the type of techniques it used, descended from Borodin and Demers (1976), might be relevant anywhere else in AI, or whether its results were a one-shot oddity. The present paper in effect is arguing that the former is the case. Backbones are a topic important in AI and relevant to SAT solvers, and this paper shows that the inspiration of the line of work initiated by Borodin and Demers (1976) can be used to establish the opacity of backbones.

It is important to acknowledge that our proofs regarding Section 2.1 are drawing on elements of the insights of Borodin and Demers (1976), although in ways unanticipated by that paper. And the addition of density transfer arguments to the world of Borodin-Demers arguments is due to Hemaspaandra, Hemaspaandra, and Menton (2013), and we are benefiting from that argument.

5 Conclusions

We argued, under assumptions widely believed to be true such as the hardness of integer factoring, that knowing a large backbone exists doesn't mean one can efficiently find a large backbone, and finding a nontrivial backbone doesn't mean one can efficiently find its value. Further, we showed that one can ensure that these effects are not very infrequent, but rather that they can be made to happen with “almost” the same density of occurrence as the error rates of the most densely hard sets in $NP \cap \text{coNP}$.

Most of our results relied on the assumption that $P \neq$

$NP \cap coNP$, which as noted above is likely true, since if it is false then integer factoring is in P and the RSA encryption scheme falls. It would be interesting, however, to see whether one can get any (likely weaker) backbone-opacity results under the weaker assumption that $P \neq NP$. We in fact have done so, but we consider those results unsatisfying, and they in effect rely on a particular feature of the Williams, Gomes, and Selman (2003) definition of backbones, namely, that unsatisfiable formulas have no backbones. That feature has no effect on the results of this paper, since in all our theorems we produced sets A whose elements all are satisfiable formulas.

Acknowledgments

We are deeply grateful to the anonymous AAAI-17 reviewers for very helpful comments and suggestions.

References

- Borodin, A., and Demers, A. 1976. Some comments on functional self-reducibility and the NP hierarchy. Technical Report TR 76-284, Department of Computer Science, Cornell University, Ithaca, NY.
- Cook, S. 1971. The complexity of theorem-proving procedures. In *Proceedings of the 3rd ACM Symposium on Theory of Computing*, 151–158. ACM Press.
- Fenner, S.; Fortnow, L.; Naik, A.; and Rogers, J. 2003. Inverting onto functions. *Information and Computation* 186(1):90–103.
- Galil, Z. 1974. On some direct encodings of nondeterministic Turing machines operating in polynomial time into P-complete problems. *SIGACT News* 6(1):19–24.
- Gomes, C.; Sabharwal, A.; and Selman, B. 2009. Model counting. In Biere, A.; Heule, M.; van Maaren, H.; and Walsh, T., eds., *Handbook of Satisfiability*. IOS Press. 633–654.
- Hartmanis, J., and Hemachandra, L. 1988. Complexity classes without machines: On complete languages for UP. *Theoretical Computer Science* 58(1–3):129–142.
- Hemaspaandra, L., and Narváez, D. 2016. The opacity of backbones. Technical Report arXiv:1606.03634v2 [cs.AI], Computing Research Repository, arXiv.org/corr/.
- Hemaspaandra, E.; Hemaspaandra, L.; and Menton, C. 2013. Search versus decision for election manipulation problems. In *Proceedings of the 30th Annual Symposium on Theoretical Aspects of Computer Science*, 377–388. Leibniz International Proceedings in Informatics (LIPIcs).
- Hemaspaandra, L.; Rothe, J.; and Wechsung, G. 1997. Easy sets and hard certificate schemes. *Acta Informatica* 34(11):859–879.
- Karp, R. 1972. Reducibilities among combinatorial problems. In Miller, R., and Thatcher, J., eds., *Complexity of Computer Computations*, 85–103.
- Levin, L. 1975. Universal sequential search problems. *Problems of Information Transmission* 9(3):265–266. March 1975 translation into English of Russian article originally published in 1973.
- Meyer, A., and Paterson, M. 1979. With what frequency are apparently intractable problems difficult? Technical Report MIT/LCS/TM-126, Laboratory for Computer Science, MIT, Cambridge, MA.
- Monasson, R.; Zecchina, R.; Kirkpatrick, S.; Selman, B.; and Troyansky, L. 1999. Determining computational complexity from characteristic ‘phase transitions’. *Nature* 400:133–137.
- Rothe, J. 1999. Complexity of certificates, heuristics, and counting types, with applications to cryptography and circuit theory. Habilitation thesis, Friedrich-Schiller-Universität Jena, Institut für Informatik, Jena, Germany.
- Schöning, U. 1986. Complete sets and closeness to complexity classes. *Mathematical Systems Theory* 19(1):29–42.
- Valiant, L. 1976. The relative complexity of checking and evaluating. *Information Processing Letters* 5(1):20–23.
- Valiant, L. 1979. The complexity of computing the permanent. *Theoretical Computer Science* 8(2):189–201.
- Williams, R.; Gomes, C.; and Selman, B. 2003. Backdoors to typical case complexity. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence*, 1173–1178. Morgan Kaufmann.