The strength of replacement in weak arithmetic

STEPHEN COOK and NEIL THAPEN
Department of Computer Science, University of Toronto

The *replacement* (or *collection* or *choice*) axiom scheme $\mathrm{BB}(\Gamma)$ asserts bounded quantifier exchange as follows: $\forall i < |a|\, \exists x < a\, \phi(i,x) \to \exists w\, \forall i < |a|\, \phi(i,[w]_i)$, for ϕ in the class Γ of formulas. The theory S^1_2 proves the scheme $\mathrm{BB}(\Sigma^b_1)$, and thus in S^1_2 every Σ^b_1 formula is equivalent to a strict Σ^b_1 formula (in which all non-sharply-bounded quantifiers are in front). Here we prove (sometimes subject to an assumption) that certain theories weaker than S^1_2 do not prove either $\mathrm{BB}(\Sigma^b_1)$ or $\mathrm{BB}(\Sigma^b_0)$. We show (unconditionally) that V^0 does not prove $\mathrm{BB}(\Sigma^B_0)$, where V^0 (essentially $\mathrm{I}\Sigma^{1,b}_0$) is the two-sorted theory associated with the complexity class AC^0. We show that PV does not prove $\mathrm{BB}(\Sigma^b_0)$, assuming that integer factoring is not possible in probabilistic polynomial time. Johannsen and Pollett introduced the theory C^0_2 associated with the complexity class TC^0, and later introduced an apparently weaker theory $\Delta^b_1-\mathrm{CR}$ for the same class. We use our methods to show that $\Delta^b_1-\mathrm{CR}$ is indeed weaker than C^0_2, assuming that RSA is secure against probabilistic polynomial time attack.

Our main tool is the KPT witnessing theorem.

Categories and Subject Descriptors: F.4.1 [**Mathematical Logic**]: proof theory, model theory; F.2.2 [**Nonnumerical Algorithms and Problems**]: complexity of proof procedures, computations on discrete structures

General Terms: Theory

Additional Key Words and Phrases: bounded arithmetic, cryptography, PV

1. INTRODUCTION

We are concerned with the strength of various theories of bounded arithmetic associated with the complexity classes P, TC^0, and AC^0. Our goal is to show that some of these theories cannot prove replacement, which is the axiom scheme

$$\forall i < |a|\, \exists x < a\, \phi(i,x) \to \exists w\, \forall i < |a|\, \phi(i,[w]_i), \tag{1}$$

where $\phi(i,x)$ can have other free variables (and $[w]_i$ is defined below). We use $\mathrm{BB}(\Gamma)$ to denote replacement for all formulas ϕ in a class Γ (usually Σ^b_0 or Σ^b_1). Replacement is also sometimes known as "collection" (eg. [Krajíček 1995]) or "choice" (eg. [Zambella 1996]). We begin by briefly describing the main theories of interest.

Authors' addresses: Stephen Cook, Department of Computer Science, University of Toronto, Toronto, Ontario M5S 3G4, Canada, email: sacook@cs.toronto.edu. Neil Thapen, (current address) St Hilda's College, University of Oxford, Cowley Place, Oxford OX4 1DY, UK, email: neil.thapen@st-hildas.ox.ac.uk.

The language of first order arithmetic that we use is

$$\{0,1,+,\cdot,<,|x|,(x)_i,[x]_i,x\#y\}.$$

Here $|x|$ is the length of x in binary notation, $(x)_i$ is the ith bit of x, $[x]_i$ is the ith element of the sequence coded by x, and $x\#y$ is $2^{|x|\cdot|y|}$. All our theories in this language are assumed to include a set of axioms BASIC fixing the algebraic properties of these symbols; see [Buss 1986; Krajíček 1995] for more detail. (These references do not take $[x]_i$ and $(x)_i$ as primitive, but these are simple functions and we can add them, and axioms for them, without changing the power of our theories.)

In the first order setting we will look at $\mathrm{BB}(\Sigma_0^b)$, or "sharply bounded replacement". A sharply bounded or Σ_0^b formula is one in which every quantifier is bounded by a term of the form $|t|$. A Σ_1^b formula is a sharply bounded formula preceded by a mixture of bounded existential and sharply bounded universal quantifiers. A strict Σ_1^b formula is a sharply bounded formula preceded by a block of bounded existential quantifiers.

The strongest theory we look at is S_2^1 [Buss 1986], defined as BASIC together with "length induction", that is the LIND axiom

$$\phi(0)\wedge\forall x<|a|\,(\phi(x)\to\phi(x+1))\to\phi(|a|) \tag{2}$$

for all Σ_1^b formulas ϕ.

S_2^1 proves $\mathrm{BB}(\Sigma_1^b)$, and hence for every S_2^1-formula ϕ there is a strict-Σ_1^b formula ϕ' such that S_2^1 proves $(\phi\leftrightarrow\phi')$. This fact may have influenced Buss's [Buss 1986] original decision not to choose strict Σ_i^b as the standard definition of Σ_i^b. The general definition allows Buss to prove his Thm 2.2 showing that if a theory T^+ extends T by adding Σ_1^b-defined function symbols then Σ_1^b formulas in the extended language are provably equivalent to Σ_1^b formulas in the original language. This result may not hold if Σ_1^b is taken to be strict Σ_1^b and T does not prove replacement. We show here that certain weaker theories (likely) do not prove replacement. For these theories, strict Σ_1^b is a more appropriate definition, and extensions by Σ_1^b-defined functions must be handled with care.

The first order theory we will use most often is PV [Cook 1975] (called PV_1 in [Krajíček 1995] and QPV in [Cook 1998]). This is defined by expanding our language to include a function symbol for every polynomial time algorithm, introduced inductively by Cobham's limited recursion on notation. These are called PV functions, and quantifier free formulas in this language are PV formulas. One way to axiomatize PV is BASIC plus universal axioms defining the new function symbols plus the induction scheme IND

$$\phi(0)\wedge\forall x<a\,(\phi(x)\to\phi(x+1))\to\phi(a)$$

for open formulas $\phi(x)$. However it is an important fact that PV is a universal theory, and can be axiomatized by its universal consequences [Buss 1986; Cook 1998].

PV and S_2^1 are closely linked to the complexity class P. The provably total Σ_1^b (or even strict Σ_1^b) functions in these theories are precisely the polynomial time functions. S_2^1 is Σ_1^b-conservative over PV [Buss 1986], but PV cannot prove the Σ_1^b-LIND axiom scheme (2) for S_2^1 unless the polynomial hierarchy (provably) collapses

[Krajíček et al. 1991; Buss 1995; Zambella 1996].

First order theories are unsuitable for dealing with very weak complexity classes such as AC^0, in which we cannot even define multiplication of strings. In this setting it is more natural to work with a two-sorted or "second order" theory. V^0 is the theory described in the Notes [Cook 2002], page 56. It is based on Σ_0^p-comp [Zambella 1996] and is essentially the same as $\mathrm{I}\Sigma_0^{1,b}$. The two sorts are numbers and strings (finite sets of numbers). There are number axioms giving the basic properties of $0, 1, +, \cdot, \leq$, and two axioms defining the "length" $|X|$ of a finite set X to be 1 plus the largest element in X, or 0 if X is empty. Finally there is the comprehension scheme for Σ_0^B formulas. These are formulas which allow bounded number quantifiers but no string quantifiers, and represent precisely the uniform AC^0 relations on their free string variables.

If we add to V^0 a function $X \cdot Y$ for string multiplication, we get a theory equivalent to the first order theory $\Sigma_0^b - \mathrm{LIND}$. The number sort would correspond to sharply bounded numbers and the string sort to "large" numbers; the Σ_0^B induction available in V^0 would correspond to $\Sigma_0^b - \mathrm{LIND}$.

With this correspondence (known as RSUV isomorphism [Takeuti 1993; Razborov 1993]) in mind, we consider V^0 and the first order fragments of S_2^1 as fitting naturally into one hierarchy of theories of bounded arithmetic. The only differences between the two approaches will be in the notation for strings and sequences. $(z)_i = 1$ in the first order setting corresponds to $Z(i)$ or $i \in Z$ in the second order setting; $[z]_i$ corresponds to $Z^{[i]}$ (see next paragraph).

In second order bounded arithmetic the replacement scheme (1) becomes

$$\forall i < n\, \exists X < n\, \phi(i, X) \rightarrow \exists W\, \forall i < n\, \phi(i, W^{[i]}).$$

Here $\exists X < n\, \phi$ stands for $\exists X(|X| < n \wedge \phi)$ and $W^{[i]}(u)$ is formally $W(\langle i, u\rangle)$ where $\langle i, u\rangle$ is a standard pairing function (so $W^{[i]}$ is row i in the two-dimensional bit array W).

Our main results are that V^0 does not prove Σ_0^B replacement (unconditionally) and that, unless integer factoring is possible in probabilistic polynomial time, PV does not prove Σ_0^b replacement. (As mentioned above, S_2^1 does prove Σ_0^b replacement.)

We summarize our results with a picture of the structure of theories between S_2^1 and V_0. An arrow on the diagram represents inclusion. To the right of an arrow we give a sufficient condition for the two theories to be distinct. A bold arrow indicates that this condition is true, and that the theories in fact are distinct. To the left of an arrow we show the conservativity between the two theories.

We will begin with the bottom of the diagram. We have already talked about V^0 and PV. $\Delta_1^b - \mathrm{CR}$ was introduced in [Johannsen and Pollett 2000] to correspond to the complexity class TC^0 of constant-depth circuits with threshold gates. The Σ_1^b functions provably total in $\Delta_1^b - \mathrm{CR}$ are precisely the uniform TC^0 functions. The theory is defined as the closure of the BASIC axioms and the LIND axioms for open formulas under the normal rules of logical deduction together with the Δ_1^b-comprehension rule: if we can *prove* that a Σ_1^b formula $\phi(x)$ is equivalent to a Π_1^b formula $\psi(x)$, then we are allowed to introduce comprehension for ϕ,

$$\exists w\, \forall i < |a|\,, (w)_i = 1 \leftrightarrow \phi(i).$$

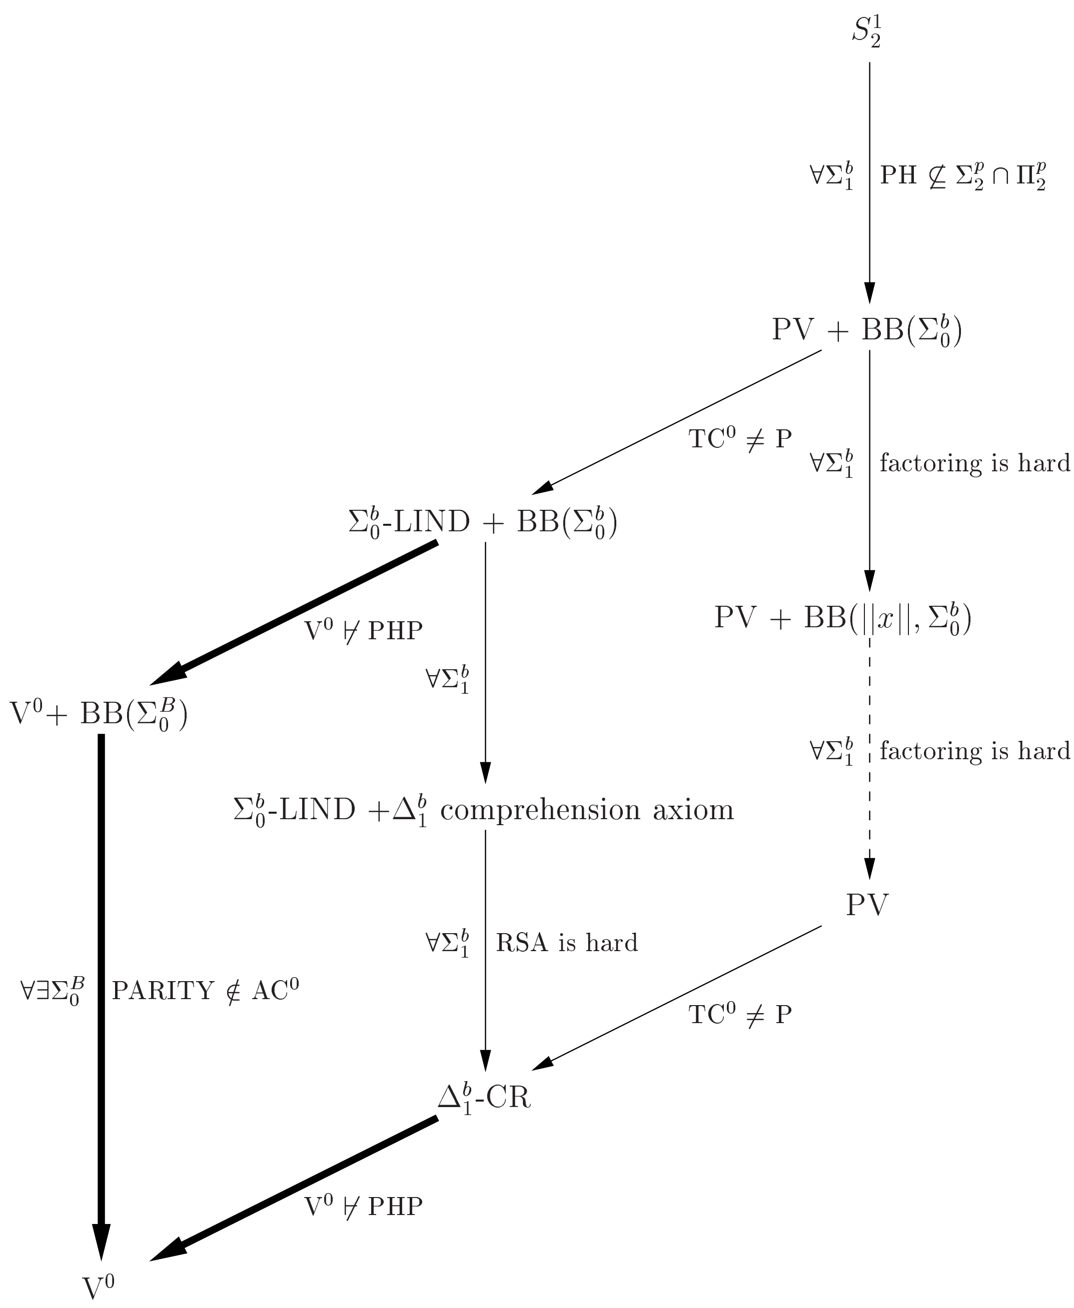
S^1_2
$\forall\Sigma^b_1$
$\mathrm{PH} \not\subseteq \Sigma^p_2 \cap \Pi^p_2$
$\mathrm{PV} + \mathrm{BB}(\Sigma^b_0)$
$\mathrm{TC}^0 \neq \mathrm{P}$
$\forall\Sigma^b_1$
factoring is hard
Σ^b_0-LIND + $\mathrm{BB}(\Sigma^b_0)$
$\mathrm{V}^0 \not\vdash \mathrm{PHP}$
$\mathrm{PV} + \mathrm{BB}(||x||, \Sigma^b_0)$
$\forall\Sigma^b_1$
$\mathrm{V}^0 + \mathrm{BB}(\Sigma^B_0)$
$\forall\Sigma^b_1$
factoring is hard
Σ^b_0-LIND $+\Delta^b_1$ comprehension axiom
PV
$\forall\Sigma^b_1$
RSA is hard
$\forall\exists\Sigma^B_0$
$\mathrm{PARITY} \notin \mathrm{AC}^0$
$\mathrm{TC}^0 \neq \mathrm{P}$
Δ^b_1-CR
$\mathrm{V}^0 \not\vdash \mathrm{PHP}$
V^0

$\Delta_1^b - \mathrm{CR}$ proves induction for sharply bounded formulas, so we can think of V^0 as a subtheory of it. In fact [Nguyen 2004] defines an extension VTC^0 of V^0 by adding an axiom for the function $\mathrm{NUMONES}(X)$ (which counts the number of 1's in the string X) and proves VTC^0 is RSUV isomorphic to $\Delta_1^b - \mathrm{CR}$. But VTC^0 proves the pigeonhole principle, as represented by a Σ_0^B formula $\mathrm{PHP}(X,n)$ [Nguyen 2004], and V^0 does not [Cook 2002]. Hence $\Delta_1^b - \mathrm{CR}$ is strictly stronger than V^0.

The Δ_1^b-comprehension rule is a derived rule of PV. This is because by results in [Buss 1986] if a formula ϕ is provably Δ_1^b in PV, then PV proves that the characteristic function of ϕ is computable in polynomial time, and hence that comprehension holds for ϕ. Thus PV is an extension of $\Delta_1^b - \mathrm{CR}$.

PV is separated from $\Delta_1^b - \mathrm{CR}$ by the circuit value principle, which says that "for all circuits C and all inputs $\bar{x}$, there exists a computation of C on $\bar{x}$". This is provable in PV, but under the assumption that P does not equal uniform TC^0 it is not provable in $\Delta_1^b - \mathrm{CR}$.

Turning now to the top of the diagram, [Buss 1986] proves the $\forall\Sigma_1^b$-conservativity of S_2^1 over PV. If $\mathrm{PV}+\mathrm{BB}(\Sigma_0^b)$ proves S_2^1, then $\mathrm{PV} \vdash S_2^1$ [Zambella 1996] and hence the bounded arithmetic hierarchy collapses to PV and the polynomial hierarchy PH collapses to $\Sigma_2^p \cap \Pi_2^p$ [Zambella 1996; Buss 1995].

The $\forall\exists\Sigma_0^B$-conservativity of $V^0+\mathrm{BB}(\Sigma_0^B)$ over V^0 is from Zambella [1996]. $\Sigma_0^b - \mathrm{LIND}+\mathrm{BB}(\Sigma_0^b)$ was introduced by Johannsen and Pollett [1998](where they call it C_2^0), and proved to be $\forall\Sigma_1^b$ conservative over $\Delta_1^b - \mathrm{CR}$ in [Johannsen and Pollett 2000]. From these conservativity results it follows that $V^0+\mathrm{BB}(\Sigma_0^B)$ does not prove the pigeonhole principle and $\Delta_1^b - \mathrm{CR} + \mathrm{BB}(\Sigma_0^b)$ does not prove the circuit value principle (unless P equals uniform TC^0), which gives us the separations between the three theories with replacement.

In the body of the paper we show the separations between the theories with and without various kinds of replacement, using a similar argument in all cases.

In section 2 we describe how our general argument goes. In section 3 we use it together with the fact that parity is not computable in nonuniform AC^0 to separate V^0 from $V^0 + \mathrm{BB}(\Sigma_0^b)$.

In section 4 we show that if PV proves Σ_0^b-replacement, then factoring is possible in probabilistic polynomial time. (This strengthens a result in [Thapen 2002] where the weaker conclusion "RSA is insecure" was proved.) We observe that this is true even if we look at weak versions of Σ_0^b-replacement, where we code very short sequences of witnesses; for example $\mathrm{BB}(\Sigma_0^b, ||x||)$ in the diagram is the scheme of replacement for sequences of double-log length:

$$\forall i < ||a||\, \exists y < a\, \phi(i,y) \to \exists w\, \forall i < ||a||\, \phi(i,[w]_i).$$

The dotted line in the diagram represents the fact that if factoring is hard, then all the theories $\mathrm{BB}(\Sigma_0^b, |x|)$, $\mathrm{BB}(\Sigma_0^b, ||x||)$, $\mathrm{BB}(\Sigma_0^b, |||x|||)$, $\ldots$ are distinct (in fact we show something slightly stronger than this). By a similar argument, all these theories are distinct over V^0 (in place of PV), without any assumptions, but for the sake of tidiness we have not put this on the diagram.

The theory of strong Δ_1^b comprehension is like $\Delta_1^b - \mathrm{CR}$, except that rather than having a rule that if a formula is provably Δ_1^b then comprehension holds for it, we have the "Δ_1^b comprehension axiom scheme"

$$\forall x\,(\phi(x) \leftrightarrow \neg\psi(x)) \rightarrow \exists w\,\forall i\!<\!|a|\,(\phi(i) \leftrightarrow (w)_i = 1) \tag{3}$$

where $\phi, \psi \in \Sigma_1^b$ (and may contain other parameters); so comprehension holds for ϕ in a structure, if ϕ is Δ_1^b in that structure. The question is raised in [Johannsen and Pollett 2000], whether this theory is strictly stronger than $\Delta_1^b - \mathrm{CR}$. We show that it is, under a cryptographic assumption. We consider a principle not shown on the diagram, which we call "unique replacement". We show that if RSA is secure against probabilistic polynomial time attack then PV does not prove unique replacement, and that it follows that PV, and hence $\Delta_1^b - \mathrm{CR}$, does not prove the Δ_1^b comprehension axiom scheme.

We have not looked for a separation between this last theory and $\Sigma_0^b - \mathrm{LIND} + \mathrm{BB}(\Sigma_0^b)$.

A preliminary version of this paper appears in [Cook and Thapen 2004].

2. WITNESSING WITH AN INTERACTIVE COMPUTATION

First we recall a standard lemma.

LEMMA 2.1. *Over* BASIC, *Σ_0^b-replacement is equivalent to strict Σ_1^b-replacement. Hence over* PV, *Σ_0^b-replacement is equivalent to replacement for* PV *formulas, since* PV *proves that every* PV *formula is equivalent to a strict Σ_1^b formula.*

Similarly over V^0, Σ_0^B-replacement is equivalent to Σ_1^B-replacement, where a Σ_1^B formula is a Σ_0^B formula preceded by a block of bounded existential string quantifiers. □

Our main tool in this paper is the KPT witnessing theorem. We state it here for PV and polynomial time, although it holds in a much more general form.

THEOREM 2.2. *[Krajíček et al. 1991] Let ϕ be a* PV *formula and suppose* $\mathrm{PV} \vdash \forall x\,\exists y\,\forall z\,\phi(x,y,z)$. *Then there exists a finite sequence $f_1, \ldots, f_k$ of* PV *function symbols such that*

$$\begin{aligned}\mathrm{PV} \vdash \forall x\,\forall \bar{z}, \phi(x, f_1(x), z_1) &\vee \phi(x, f_2(x, z_1), z_2) \\ &\vee \ldots \vee \phi(x, f_k(x, z_1, \ldots, z_{k-1}), z_k).\end{aligned}$$

PROOF. Let $b, c_1, c_2, ...$ be a list of new constants, and let $t_1, t_2, ...$ be an enumeration of all terms built from symbols of PV together with $b, c_1, c_2, ...$, where the only new constants in t_k are among $\{b, c_1, ..., c_{k-1}\}$. It suffices to show that

$$\mathrm{PV} \cup \{\neg\phi(b, t_1, c_1), \neg\phi(b, t_2, c_2), \ldots, \neg\phi(b, t_k, c_k)\}$$

is unsatisfiable for some k.

Suppose otherwise. Then by compactness

$$\mathrm{PV} \cup \{\neg\phi(b, t_1, c_1), \neg\phi(b, t_2, c_2), ...\} \tag{4}$$

has a model M. Since PV is universal, the substructure M' consisting of the denotations of the terms $t_1, t_2, ...$ is also a model for (4). It is easy to see that

$$M' \models \mathrm{PV} + \forall y \exists z \neg\phi(b, y, z)$$

and hence $\mathrm{PV} \nvdash \forall x \exists y \forall z \phi(x, y, z)$. □

Now choose a function f which can be computed in polynomial time but which is hard to invert (in a more general setting, we would choose a function which is in the complexity class corresponding to the theory we are looking at, but whose inverse probably is not). Suppose PV proves the following instance of replacement (which has a and y as parameters, and $m = |a|$):

$$\forall i < m\, \exists u < a\, f(u) = [y]_i \rightarrow \exists w\, \forall j < m\, f([w]_j) = [y]_j.$$

We can rewrite this as

$$\exists i < m\, \exists w\, \forall u < a\,,\ f(u) = [y]_i \rightarrow \forall j < m\, f([w]_j) = [y]_j.$$

Applying our witnessing theorem, we get $k \in N$ and functions $g_1, \ldots, g_k$ and $h_1, \ldots h_k$ (which have a as a suppressed argument), such that

$$\begin{aligned}
\mathrm{PV} \vdash \forall \bar{z} < a\,,\ & (f(z_1) = [y]_{g_1(y)} \rightarrow \forall j < m\, f([h_1(y)]_j) = [y]_j) \\
& \vee (f(z_2) = [y]_{g_2(y,z_1)} \rightarrow \forall j < m\, f([h_2(y,z_1)]_j) = [y]_j) \\
& \vee \ldots \\
& \vee (f(z_k) = [y]_{g_k(y,z_1,\ldots,z_{k-1})} \rightarrow \\
& \qquad \forall j < m\, f([h_k(y,z_1,\ldots,z_{k-1})]_j) = [y]_j)
\end{aligned}$$

This allows us to write down an algorithm which, given an input y (considered as a sequence $[y]_0, \ldots, [y]_{m-1}$), will ask for a pre-image of f on at most k elements of y. With this information it will output a number w coding a sequence of pre-images of all m elements of y.

The algorithm is as follows. Let $w = h_1(y)$. If $\forall j < m\, f([w]_j) = [y]_j$ then output w and halt. Otherwise calculate $g_1(y)$ and ask for a pre-image of $[y]_{g_1(y)}$; store the answer as z_1. Then let $w = h_2(y, z_1)$. If $\forall j < m\, f([w]_j) = [y]_j$ then output w and halt. Otherwise calculate $g_2(y, z_1)$ and ask for a pre-image of $[y]_{g_2(y,z_1)}$; store the answer as z_2, and so on. By our assumption the algorithm will run for at most k steps of this form before it outputs a suitable w.

Now fix a such that $|a| = m > k$, and choose a sequence $[x]_0, \ldots, [x]_{m-1}$ of numbers less than a. Let y encode the pointwise image of x under f. Run the algorithm above, and reply to queries with elements of x. We will end up with w encoding a sequence of pre-images of y, which will clash in some way with our assumption that f is hard to invert. If f is an injection, w will be the same as x; we use this in section 3. If f is not an injection and x was chosen at random, then w is probably different from x; we use this in sections 4 and 5. [1]

[1] In this paper we only consider worst-cast complexity. Russell Impagliazzo has pointed out that if we consider average-case complexity, we can use our algorithm to show that no one-way permutations exist (under our assumption about replacement). Suppose f is a polynomial time permutation that maps m-bit strings to m-bit strings. We will show that f is not one-way, by showing that it is not hard to invert in the average case. Let v be a random string, which we want to find a pre-image of. Choose strings $u_1, \ldots, u_m$ at random and let $v_1, \ldots, v_m$ be their images under f. Insert v into this sequence of images at a random place to get a sequence of $m + 1$ strings uniformly distributed amongst all such sequences (since f is a permutation), and give this sequence to our algorithm. It will ask for k pre-images and with high probability we will be able to give correct answers, using the u_is. Then the algorithm will output pre-images for every string

The important properties of PV used in the argument above are that it is universal and can define functions by cases (needed for the KPT witnessing theorem) and that it can manipulate sequences. We show now how to make V^0 into a universal theory in which we can carry out the same argument.

We start by referring to [Cook 2002], pp 66–73. A relation $R(\bar{x}, \bar{Y})$ is in (uniform) AC^0 iff it is defined by some Σ_0^B formula $A(\bar{x}, \bar{Y})$. A number function $f : \mathbb{N}^k \times (\{0,1\}^*)^\ell \longrightarrow \mathbb{N}$ is an AC^0 function iff there is an AC^0 relation R and a polynomial p such that

$$f(\bar{x}, \bar{Y}) = \min z < p(\bar{x}, |\bar{Y}|)\ R(z, \bar{x}, \bar{Y}) \tag{5}$$

A string function $F(\bar{x}, \bar{Y})$ is an AC^0 function iff $|F(\bar{x}, \bar{Y})| \leq p(\bar{x}, |\bar{Y}|)$ for some polynomial p, and the bit graph

$$B_F(i, \bar{x}, \bar{Y}) \equiv F(\bar{x}, \bar{Y})(i)$$

is an AC^0 relation.

We denote by $V^0(\mathrm{FAC}^0)$ a conservative extension of V^0 obtained by adding a set FAC^0 of function symbols with universal defining axioms for all AC^0 functions, based on the above characterizations. FAC^0 is essentially $\mathcal{R} - def$ in [Zambella 1996].) This can be done in such a way that $V^0(\mathrm{FAC}^0)$ is a universal theory. In particular, the Σ_0^B comprehension axioms follow since for every Σ_0^B formula ϕ there is a FAC^0 string function whose range is the set of strings asserted to exist by the the comprehension axiom for ϕ. Further, from (5) it is clear that for every Σ_0^B formula ϕ there is a quantifier-free formula ϕ' in the language of $V^0(\mathrm{FAC}^0)$ such that

$$V^0(\mathrm{FAC}^0) \vdash (\phi \leftrightarrow \phi')$$

From these remarks, it is clear that the usual proof of the KPT witnessing theorem can be adapted to show the following:

Theorem 2.3. *Let $\phi(X, Y, Z)$ be a Σ_0^B formula such that $\forall X \exists Y \forall Z \phi(X, Y, Z)$ is provable in V^0. Then there are FAC^0 functions $F_1, ..., F_k$ such that*

$$\begin{aligned} V^0(\mathrm{FAC}^0) \vdash \forall X \forall \bar{Z},\ & \phi(X, F_1(X), Z_1) \vee \phi(X, F_2(X, Z_1), Z_2) \\ & \vee \ldots \vee \phi(X, F_k(X, Z_1, ..., Z_{k-1}), Z_k). \end{aligned}$$

Using this we can show that if V^0 proves Σ_0^B-replacement, then for any AC^0 function F there exists $k \in N$ and a uniform AC^0 algorithm that will find a pre-image under F of any sequence $Y^{[0]}, \ldots, Y^{[m-1]}$ of strings by asking at most k queries of the form "what is a pre-image of $Y^{[i]}$?"

3. REPLACEMENT IN $\mathbf{V^0}$ AND PARITY

Let $PARITY$ be the set of all strings over $\{0,1\}$ with an odd number of 1s. By a (nonuniform) AC^0 circuit family we mean a polynomial size bounded depth family $\langle C_n : n \in \mathbb{N} \rangle$ of Boolean circuits over $\wedge, \vee, \neg$ such that C_n has n inputs and one output. Ajtai's theorem [Ajtai 1983; Furst et al. 1984] states that no such circuit family accepts $PARITY$.

in the sequence, including v.

We show that if V^0 proves the Σ_0^B replacement scheme, then (using KPT witnessing) there exists a (uniform) randomized AC^0 algorithm for $PARITY$. This algorithm shows the existence of a (uniform) AC^0 circuit family such that each circuit has a vector $\bar{r}$ of random input bits in addition to the standard input bits, and with probability $p > 2/3$ the circuit correctly determines whether the standard input is in $PARITY$ and with probability $1-p$ the circuit produces an output indicating failure. From this a standard argument shows the existence of a nonuniform AC^0 circuit family for parity, violating the above theorem.

Let PAR be the function that maps a binary string of length m to its parity vector. That is, $PAR(m, Y) = X$ if $|X| < m$ and, for each $i < m$, $X(i)$ is the parity of the string $Y(0) \dots Y(i)$. In what follows we take m to be a parameter, assume Y is an m-bit string, and suppress the argument m from $PAR(m, Y)$.

Plainly $PAR(Y)$ cannot be computed in AC^0. However its inverse, which we will call $UNPAR$, is in uniform AC^0: the ith bit of $UNPAR(X)$ is given by the Σ_0^B formula $(i = 0 \wedge X(i)) \vee (i > 0 \wedge X(i-1) \oplus X(i))$. Here $UNPAR$ has an argument m, which we suppress.

Notice also that for all m-bit strings A, B, C, writing $\oplus$ for bitwise XOR, if $A = B \oplus C$ then $PAR(A) = PAR(B) \oplus PAR(C)$.

Theorem 3.1. *V^0 does not prove* $\mathrm{BB}(\Sigma_0^B)$.

Proof. Suppose $V^0 \vdash \mathrm{BB}(\Sigma_0^B)$. Then applying the argument of section 2 to the function $UNPAR$, for some fixed k there is a uniform AC^0 algorithm which, for any sequence $Y^{[0]}, \dots, Y^{[m-1]}$ of binary strings of length m makes k queries of the form "what is $PAR(Y^{[i]})$?" and outputs the sequence of parity vectors of Y.

We will show how to use this algorithm to compute the parity of a single string in uniform randomized AC^0. Suppose $m \geq 3k$ and let I be the input string of length m which we want to compute the parity of.

Choose m strings $U_0, \dots, U_{m-1}$ in $\{0,1\}^m$ at random, and for each i compute $V_i = UNPAR(U_i)$. Choose a number r, $0 \leq r < m$, uniformly at random. Define the string Y (thought of as an $m \times m$ binary matrix) by the condition

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

Since for each m the function $UNPAR$ defines a bijection from the set $\{0,1\}^m$ to itself, and since for each I with $|I| < m$ the map $X \mapsto I \oplus X$ also defines a bijection from that set to itself, it follows that the string Y defined above, interpreted as an $m \times m$ bit matrix, is uniformly distributed over all such matrices.

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

Since at most k different values of i are compared to r and since for each input I each pair (Y, r) is equally likely to have been chosen, it follows that the computation will be aborted with probability at most $k/m \leq 1/3$.

Hence with probability at least $2/3$ the algorithm is not aborted, we are able to answer all the queries correctly, and we obtain W such that $W^{[r]} = PAR(Y^{[r]}) =$

$PAR(I \oplus V_r)$. But $I = V_r \oplus (I \oplus V_r)$ and hence

$$\begin{aligned} PAR(I) &= PAR(V_r) \oplus PAR(I \oplus V_r) \\ &= U_r \oplus W^{[r]} \end{aligned}$$

We use this to compute $PAR(I)$ and use bit $m-1$ of $PAR(I)$ to determine whether $I \in PARITY$.

For each input I the algorithm succeeds with probability at least $2/3$, where the probability is taken over its random input bits.

Since no such AC^0 algorithm exists, it follows that V^0 does not prove the Σ_0^B replacement scheme. □

4. REPLACEMENT IN PV AND FACTORING

We adapt the proof [Rabin 1979] that cracking Rabin's cryptosystem based on squaring modulo n is as hard as factoring.

Let n be the product of distinct odd primes p and q. Suppose $0 < x_1 < n$ and $\gcd(x_1, n) = 1$. Let $c = x_1^2$. Then c has precisely four square roots x_1, x_2, x_3, x_4 modulo n. This can be seen as follows: let $x_p = (x_1 \bmod p)$ and $x_q = (x_1 \bmod q)$. By the Chinese remainder theorem there are uniquely determined numbers x_1, x_2, x_3, x_4 with $0 < x_i < n$ such that

$$\begin{array}{ll} x_1 \equiv x_p \pmod p & x_1 \equiv x_q \pmod q \\ x_2 \equiv x_p \pmod p & x_2 \equiv -x_q \pmod q \\ x_3 \equiv -x_p \pmod p & x_3 \equiv x_q \pmod q \\ x_4 \equiv -x_p \pmod p & x_4 \equiv -x_q \pmod q \end{array}$$

Now $x_1 - x_2 \equiv 0 \pmod p$ and $x_1 - x_2 \equiv 2x_q \not\equiv 0 \pmod q$, so $\gcd(x_1 - x_2, n) = p$. So from x_1 and x_2 we can recover p, and similarly from x_1 and x_3 we can recover q.

Hence if we have one square root of c, and are then given a square root at random, we can factor n with probability $\frac{1}{2}$.

THEOREM 4.1. *If* PV *proves replacement for sharply bounded formulas, then factoring (of products of two odd primes) is possible in probabilistic polynomial time.*

PROOF. We will use our standard argument, taking squaring modulo n as our function f (so f has n as a parameter).

If PV proves $\mathrm{BB}(\Sigma_0^b)$ then there is polynomial time algorithm which, for some fixed $k \in \mathbb{N}$, given any sequence $y_0, \ldots, y_{m-1}$ of squares (modulo n), makes at most k queries of the form "what is the square root of y_i?" and, if these are answered correctly, outputs square roots of all the y_is.

Now suppose n is large enough that $m = |n| > k$. Choose numbers $x_0, \ldots, x_{m-1}$ uniformly at random with $0 < x_i < n$. We may assume that $\gcd(x_i, n) = 1$ for all i, since otherwise we can immediately find a factor of n.

For each i let $y_i = (x_i^2 \bmod n)$. Let y code the sequence $y_0, \ldots y_{m-1}$, so $[y]_i = y_i$. Notice that each x_i is distributed uniformly amongst the four square roots of $[y]_i$.

Run our algorithm, and to each query "what is the square root of $[y]_i$?", answer with x_i. We will get as output w coding a sequence $[w]_0, \ldots, [w]_{m-1}$ of square roots of $[y]_0, \ldots, [y]_{m-1}$.

If we think of n as fixed, the value of w depends only on the inputs given to the algorithm, namely y and the k many numbers x_i that we gave as replies. Let i be some index for which x_i was not used. Then x_i is distributed at random among the square roots of $[y]_i$, and $[w]_i$ is a square root of $[y]_i$ that was chosen without using any information about which square root x_i is. Hence $\gcd(x_i - [w]_i, n)$ is a factor of n with probability $\frac{1}{2}$. $\square$

Notice that the only property of the function $|\ |$ we used was that we could find some n with $|n| > k$. So any nondecreasing, not eventually constant function would do in the place of $|\ |$. Hence if PV only proves replacement for very short sequences, that is still enough to give us factoring.

In fact under the assumption that factoring is hard we can show that these replacement schemes form a hierarchy. For any α with one argument, let $\mathrm{BB}(\alpha, \mathrm{PV})$ be the axiom scheme:

$$\forall i < \alpha(b)\, \exists y < b\, \phi(i, y) \rightarrow \exists w\, \forall i < \alpha(b)\, \phi(i, [w]_i)$$

for all PV formulas ϕ. We will assume that our base theory proves that $\alpha(x) < |x|$ and that α is increasing.

We need a generalization of a result of Zambella, lemma 3.3 of [Zambella 1996]. The lemma there is presented for a two-sorted system similar to V^0 and with $|x|$ rather than $\alpha(x)$.

An $\exists^b\mathrm{PV}$ formula is a PV formula preceded by a bounded existential quantifier; modulo PV this is the same as a strict Σ^b_1 formula.

Lemma 4.2. *Any model $N \models \mathrm{PV}$ has an $\exists^b\mathrm{PV}$-elementary extension to a model $M \models \mathrm{PV} + \mathrm{BB}(\alpha, \mathrm{PV})$ such that every element of M is of the form $f(a, \bar{b})$ for some $f \in \mathrm{PV}$, $a \in N$ and $\bar{b} \subseteq \alpha(M)$, where $\alpha(M) = \{x \in M : x < \alpha(y),\ \text{some } y \in M\}$. Informally, M is formed from N by only adding new "α-small" elements and closing under* PV *functions.*

Proof. Let L be the language of PV with the addition of a name for every element of N, and let T be the universal theory of N in this language, so every model of T will be an $\exists$-elementary, and hence $\exists^b\mathrm{PV}$-elementary, extension of N. Enumerate as $(t_1, \phi_1(x, y)), (t_2, \phi_2(x, y)), \ldots$ all pairs consisting of closed terms in L and binary PV formulas with parameters from L. We will use this to construct a chain $T = T_0 \subseteq T_1 \subseteq T_2 \subseteq \ldots$ of theories.

Suppose that T_i has been constructed and is a consistent, universal theory. If $T_i \vdash \forall x < \alpha(t_{i+1})\, \exists y\, \phi_{i+1}(x, y)$ then put $T_{i+1} = T_i$. Otherwise introduce a new constant symbol c and put

$$T_{i+1} = T_i \cup \{c < \alpha(t_{i+1})\} \cup \{\forall y\, \neg\phi_{i+1}(c, y)\}.$$

Note that T_{i+1} is consistent and universal.

Let T^* be the union of this chain of theories, and let L^* be L together with all the new constant symbols that were added in the construction of T^*. Enumerate all pairs of closed terms and binary formulas in L^*, and repeat the above construction to get a theory T^{**} and a language L^{**}. Repeat this step ω times, and let T^+ be the union of the theories and L^+ its language.

T^+ is consistent and universal, so there is a model $M \vDash T^+$ each element of which is named by some closed L^+-term. $M \vDash T$, so M is an $\exists^b\mathrm{PV}$-elementary extension of N. Also, each time a new constant c was introduced to L^+, $c < \alpha(t)$ was introduced to T^+ for some term t. So M is the closure of elements of N and new "α-small" elements, as required.

To show that M is a model of $\mathrm{BB}(\alpha, \mathrm{PV})$, suppose that a is an element of M and $\phi(x, y)$ is a PV formula with parameters from M, and

$$M \vDash \forall x < \alpha(a)\, \exists y\, \phi(x, y).$$

Then by the construction of M, we may assume that a is named by some closed L^+ term t and that $\phi(x, y)$ is a parameter-free L^+ formula; and by the construction of T^+ we must have that $T^+ \vdash \forall x < \alpha(t)\, \exists y\, \phi(x, y)$, since T^+ either proves this or its negation. But T^+ is a universal theory, so by using Herbrand's theorem and the properties of PV we can find a PV function symbol f (with parameters) such that $T^+ \vdash \forall x < \alpha(t)\, \phi(x, f(x))$. Now by the comprehension available in PV, we can find some $w \in M$ such that $M \vDash \forall x < \alpha(t)\, \phi(x, [w]_x)$, as required. □

We can now adapt the proof of the KPT witnessing theorem to get the following:

THEOREM 4.3. *Suppose*

$$\mathrm{PV} + \mathrm{BB}(\alpha, \mathrm{PV}) \vdash \forall x\, \exists y\, \forall z\, \phi(x, y, z)$$

for an $\exists^b\mathrm{PV}$ *formula* ϕ. *Then there exist* $k \in \mathbb{N}$, *a term* $s(x, \bar{z})$ *and functions* $f_1, \ldots, f_k$ *such that*

$$\begin{aligned} \mathrm{PV} \vdash \forall x\, \forall \bar{z}, &\exists i < \alpha(s)^k\, \phi(x, [f_1(x)]_i, [z_1]_i) \\ &\vee \exists i < \alpha(s)^k\, \phi(x, [f_2(x, z_1)]_i, [z_2]_i) \\ &\vee \ldots \vee \exists i < \alpha(s)^k\, \phi(x, [f_k(x, z_1, \ldots, z_{k-1})]_i, [z_k]_i) \end{aligned}$$

(we include the exponent k *here because the range of* α *might not be closed under multiplication).*

PROOF. Enumerate all pairs of PV functions as $(s_1, f_1), (s_2, f_2), \ldots$ with infinite repetitions in such a way that for each k both s_k and f_k take k or fewer arguments. Assume that the conclusion of the theorem is false, and let T be the theory

$$\begin{aligned} \mathrm{PV} + \{&\forall i < \alpha(s_1(b, c_1))^1\, \neg\phi(b, [f_1(b)]_i, [c_1]_i), \\ &\forall i < \alpha(s_2(b, c_1, c_2))^2\, \neg\phi(b, [f_2(b, c_1)]_i, [c_2]_i), \ldots\} \end{aligned}$$

where b and $c_1, c_2, \ldots$ are new constant symbols. Then T is finitely satisfiable (we can take the term s in the statement of the theorem as the sum of our finite set of terms $s_1, \ldots, s_k$).

Let N be a model of T, and let $N' \subseteq N$ be the substructure consisting of all the elements named by terms. Since T is universal, $N' \models T$. Let M be the extension of N given by lemma 4.2 to a model of $\mathrm{BB}(\alpha, \mathrm{PV})$. By $\exists^b\mathrm{PV}$ elementariness, M is also a model of T.

Now let a be any element of M. By the construction of M, for some $\bar{d} \subseteq \alpha(M)$, some $e \in N'$ and some PV function g we have $a = g(\bar{d}, e)$. Furthermore by the construction of N' we know that $\bar{d} < \alpha(h_1(b, c_1, \ldots, c_k))$ and $e = h_2(b, c_1, \ldots, c_k)$ for some k and some PV functions h_1 and h_2.

In this paragraph we identify a number $i < \alpha(h_1(b, \bar{c}))^k$ with the sequence $\bar{i} = i_1 \ldots i_k$ of numbers less than $\alpha(h_1(b, \bar{c}))$ that it codes. We can find $l >$

k such that f_l is the PV function symbol that takes as input $b, c_1, \ldots, c_l$ and outputs (as a single number) the sequence $w_1 \ldots w_{\alpha(h_1(b,c_1,\ldots,c_k))^k}$ where $w_i = g(\bar{i}, h_2(b, c_1, \ldots, c_k))$. Then $a = [f_l(b, c_1, \ldots, c_l)]_d$ and since $M \models T$ we have $M \models \neg\phi(b, a, [c_{l+1}]_d)$. Here a was chosen arbitrarily, so we have shown that $M \models \mathrm{PV} + \mathrm{BB}(\alpha, \mathrm{PV}) + \neg\forall x\, \exists y\, \forall z\, \phi(x, y, z)$. □

Corollary 4.4. *Suppose that factoring is not possible in probabilistic polynomial time. Then* $\mathrm{BB}(\alpha, \mathrm{PV})$ *is not provable in* $\mathrm{PV} + \mathrm{BB}(\gamma, \mathrm{PV})$*, for terms* α, γ *where* $\alpha(x), \gamma(x) < |x|$ *and* α *grows faster than any polynomial in* γ*.*

Proof. Our standard argument is that if replacement is provable in PV, then there is a polynomial time interactive algorithm that queries k square roots and outputs $|n|$ square roots, for some fixed $k \in \mathbb{N}$.

By theorem 4.3 we can show, by a similar argument, that if $\mathrm{PV} + \mathrm{BB}(\gamma, \mathrm{PV}) \vdash \mathrm{BB}(\alpha, \mathrm{PV})$ then we have a polynomial time interactive algorithm that queries $k\gamma(n)^k$ square roots modulo n and outputs $\alpha(n)$ square roots, for some fixed $k \in \mathbb{N}$.

So if n is sufficiently large that $\alpha(n) > k\gamma(n)^k$, we can use the argument of theorem 4.1 to factor n. □

This gives a hierarchy of theories

$$\mathrm{PV} + \mathrm{BB}(|x|, \mathrm{PV}) \supset \mathrm{PV} + \mathrm{BB}(||x||, \mathrm{PV}) \supset \ldots$$

The same argument goes through in V^0. One way to see this is to notice that the important difference between PV and V^0 is that the PV functions are closed under polynomial time iteration, and no such iteration is used in the proof here. So we have the unconditional separation result

Theorem 4.5. $\mathrm{BB}(\alpha, \Sigma_0^B)$ *is not provable in* $V^0 + \mathrm{BB}(\gamma, \Sigma_0^B)$*, for terms* α, γ *where* $\alpha(n), \gamma(n) < n$ *and* α *grows faster than any polynomial in* γ*.*

Proof. If the theorem is false, then there is $k \in \mathbb{N}$ and an interactive algorithm that, given $\alpha(n)$ many vectors $v_1, \ldots, v_{\alpha(n)}$, each of length n, will make $k\gamma(n)^k$ queries of the form "what is the parity vector of v_i?" and then output the parity vectors of all the v_is. So if $\alpha(n) \geq 3k\gamma(n)^k$, then by adapting the argument of section 3 we get a probabilistic uniform AC^0 algorithm which computes parity. □

5. UNIQUE REPLACEMENT IN PV AND RSA

We define "unique replacement" to be the scheme

$$\forall i < |a|\, \exists! x < b\, \phi(i, x) \rightarrow \exists w\, \forall i < |a|\, \phi(i, [w]_i).$$

Theorem 5.1. *If* PV *proves unique replacement for sharply bounded formulas, then the injective* WPHP *for* PV *formulas can be witnessed in probabilistic polynomial time (and hence in particular we can crack RSA [Krajíček and Pudlák 1998]).*

Proof. (Simplified from the model-theoretic proof in [Thapen 2002].) First notice that it is sufficient to show that PV does not prove unique replacement for some PV formula ϕ. For suppose that ϕ is decided by the polynomial time machine with code e, and that for some fixed i there is a unique x such that $\phi(i, x)$. Then there is a unique pair (z, x) such that z is an accepting computation of the machine

e on input (i,x), and the property of being an accepting computation is sharply bounded.

In the rest of this proof x and y will code sequences of $|n|$ numbers each of size $< n^{|n|}$ and with elements $[x]_i, [y]_i$, and z will code a sequence of $|n|$ numbers each of size $< n$ and with elements $\langle z\rangle_i$.

Suppose that h is a PV function from $n^{|n|}$ to n. Note that from any PV function $g : 2n \to n$ we can derive such a function h with the property that a witness to WPHP for h yields in polynomial time a witness to WPHP for g ([Paris et al. 1988], or see [Thapen 2002] for an explicit polynomial time construction).

Choose $x < n^{|n|^2}$ at random and let $z < n^{|n|}$ be such that $\langle z\rangle_0 = h([x]_0)$, ..., $\langle z\rangle_{|n|-1} = h([x]_{|n|-1})$.

Assume that PV proves the following instance of unique replacement:

$$\begin{aligned}
&\exists i < |n|\, \forall u < n^{|n|}\, h(u) \neq \langle z\rangle_i \\
&\vee \exists i < |n|\, \exists u_1 < u_2 < n^{|n|}\, h(u_1) = h(u_2) \\
&\vee \exists y < n^{|n|^2}\, \forall i < |n|\, h([y]_i) = \langle z\rangle_i .
\end{aligned}$$

Then by our witnessing theorem, for some k (independent of n) there is a deterministic interactive computation which takes n and z as its initial input. Then for k steps it gives us an index $i < |n|$ and expects an input $y < n^{|n|}$; if we can guarantee that for each such step we have $h(y) = \langle z\rangle_i$, then the computation outputs either u_1 and u_2 mapping to the same thing, in which case we are done (and this case is the only one that is different from normal replacement), or $y < n^{|n|^2}$ satisfying $\forall i < |n|\, h([y]_i) = \langle z\rangle_i$.

Run the computation, and to each index i queried respond with $[x]_i$. The computation must output some y satisfying $\forall i < |n|\, h([y]_i) = \langle z\rangle_i$. Now the computation is deterministic, and if we think of n as fixed, there were $n^{|n|(k+1)}$ possible different inputs to the machine: namely $n^{|n|}$ different possibilities for z and $(n^{|n|})^k$ different possibilities for the k responses $[x]_i$. Hence there are at most $n^{|n|(k+1)}$ possible outputs y. However x was originally chosen at random from $n^{|n|^2}$ possibilities. So if $k < n-1$ then with high probability x is not a possible output of the machine, so $x \neq y$ and for some $i < |n|$ we have $[x]_i \neq [y]_i$ but $h([x]_i) = \langle z\rangle_i = h([y]_i)$. $\square$

Notice that part of this argument can be formalized in PV, to show that if PV proves unique replacement, then PV proves that the surjective WPHP for PV functions implies the injective WPHP for PV functions. In the proof above randomness was used to find some x outside the range of a given polynomial time algorithm; in the formal PV proof we would use the surjective WPHP to provide such an x.

Corollary 5.2. *Suppose* PV *proves the Δ^b_1 comprehension axiom scheme (3). Then* PV *proves unique replacement for* PV *formulas and by theorem 5.1 we can crack RSA.*

Proof. Let $\phi(i,x)$ be any PV formula (with parameters) and suppose that the hypothesis of the theorem holds. Let $M \models$ PV, $a, b \in M$ and suppose $M \models \forall i <$

$|b|\,\exists! x < a\, \phi(i,x)$. Then

$$M \models \forall i < |b|\, \forall j < |a|,$$
$$\exists x < a\, (\phi(i,x) \wedge x_j = 1) \leftrightarrow \forall x < a\, (\phi(i,x) \rightarrow x_j = 1).$$

Over PV, ϕ is equivalent to both a Σ_1^b and a Π_1^b formula, so we can apply comprehension and get some w such that

$$M \models \forall i < |b|\, \forall j < |a|,\, ([w]_i)_j = 1 \leftrightarrow \exists x < a\, (\phi(i,x) \wedge x_j = 1).$$

Here we assume without loss of generality that a is a power of 2, so that we can switch easily between thinking of w as a binary sequence of length $|b||a|$ and as a sequence of $|b|$ many binary numbers $[w]_1 \ldots [w]_{|b|}$, each of length $|a|$. We also use the fact that in PV the formula $\phi(i,x)$ can be written in both a strict Σ_1^b and a strict Π_1^b way, which we need to apply comprehension.

Now pick any $i < |b|$. There is some unique $x \in M$ such that $\phi(i,x)$; and by the construction of w, for each $j < |a|$ we know $([w]_i)_j = 1$ if and only if $x_j = 1$. Hence $[w]_i = x$.

So $M \models \forall i < |b|\, \phi(i, [w]_i)$. □

REFERENCES

AJTAI, M. 1983. Σ_1^1-formulae on finite structures. *Annals of Pure and Applied Logic 24*, 1–48.

BUSS, S. 1986. *Bounded Arithmetic*. Bibliopolis.

BUSS, S. 1995. Relating the bounded arithmetic and polynomial time hierarchies. *Annals of Pure and Applied Logic 75,* 1–2, 67–77.

COOK, S. 1975. Feasibly constructive proofs and the propositional calculus. *Proceedings of the 7th Annual ACM Symposium on Theory of computing*, 83–97.

COOK, S. 1998. Relating the provable collapse of P to NC^1 and the power of logical theories. *DIMACS Series in Discrete Mathematics and Theoretical Computer Science 39*, 73–91.

COOK, S. 2002. *CSC 2429 course notes: Proof Complexity and Bounded Arithmetic*. Available from the web at `www.cs.toronto.edu/~sacook/csc2429h/`.

COOK, S. AND THAPEN, N. 2004. The strength of replacement in weak arithmetic. *Proceedings of the Nineteenth Annual IEEE Symposium on Logic in Computer Science*.

FURST, M., SAXE, J. B., AND SIPSER, M. 1984. Parity, circuits and the polynomial-time hierarchy. *Math. Systems Theory 17*, 13–27.

JOHANNSEN, J. AND POLLETT, C. 1998. On proofs about threshold circuits and counting hierarchies (extended abstract). In *Proc. 13th IEEE Symposium on Logic in Computer Science*. 444–452.

JOHANNSEN, J. AND POLLETT, C. 2000. On the Δ_1^b-bit-comprehension rule. In *Logic Colloquium 98*, S. Buss, P. Hájek, and P. Pudlák, Eds. ASL Lecture Notes in Logic. 262–279.

KRAJÍČEK, J. 1995. *Bounded Arithmetic, Propositional Logic and Computational Complexity*. Cambridge University Press.

KRAJÍČEK, J. AND PUDLÁK, P. 1998. Some consequences of cryptographical conjectures for S_2^1 and EF. *Information and Computation 140,* 1, 82–89.

KRAJÍČEK, J., PUDLÁK, P., AND TAKEUTI, G. 1991. Bounded arithmetic and the polynomial hierarchy. *Annals of Pure and Applied Logic 52*, 143–153.

NGUYEN, P. 2004. *VTC^0: A Second-Order Theory for TC^0*. MSc Thesis, Department of Computer Science, University of Toronto.

PARIS, J., WILKIE, A., AND WOODS, A. 1988. Provability of the pigeonhole principle and the existence of infinitely many primes. *Journal of Symbolic Logic 53,* 4, 1235–1244.

RABIN, M. 1979. Digitalized signatures and public-key functions as intractable as factorization. Tech. Rep. MIT/LCS/TR-212, MIT Laboratory for Computer Science.

Razborov, A. A. 1993. An equivalence between second order bounded domain bounded arithmetic and first order bounded arithmetic. In *Arithmetic, Proof Theory and Computational Complexity*, P. Clote and J. Krajicek, Eds. Oxford University Press, 247–77.

Takeuti, G. 1993. RSUV isomorphism. In *Arithmetic, Proof Theory and Computational Complexity*, P. Clote and J. Krajicek, Eds. Oxford University Press, 364–86.

Thapen, N. 2002. A model-theoretic characterization of the weak pigeonhole principle. *Annals of Pure and Applied Logic 118*, 175–195.

Zambella, D. 1996. Notes on polynomially bounded arithmetic. *Journal of Symbolic Logic 61,* 3, 942–966.

Received September 2004; revised January 2005; accepted January 2005