A Constructive Proof of Higman's Lemma

Chetan Murthy*
James R. Russell**

TR 89-1049
October 1989

Department of Computer Science
Cornell University
Ithaca, NY 14853-7501

*Supported in part by an NSF graduate fellowship.
**Supported in part by an IBM graduate fellowship.

A Constructive Proof of Higman's Lemma

Chetan Murthy* James R. Russell†
Computer Science Department, Upson Hall
Cornell University
Ithaca, NY 14853

October 26, 1989

Abstract

Higman's Lemma is a special case of the more general Kruskal's tree embedding theorem and the graph minor theorem. Prior to this work, only classical (and impredicative) proofs of the Lemma were known. Recently there has been much interest in developing a constructive proof of the Lemma, primarily via Friedman's A-translation. In this paper we present a direct constructive proof. We achieve this by reducing the problem to a construction of certain sets of *sequential regular expressions.* We then exhibit a well-founded order on such sets, and the Lemma then follows by induction.

1 Introduction and History of Higman's Lemma

The so called "Higman's Lemma" asserts that certain constructions on well quasi ordered sets preserve well quasi orderedness, and was first considered by Higman [Hig52]. It is a special case, and an essential component, of the more general Kruskal's tree embedding theorem and the graph minor theorem. Until recently, only classical (and impredicative) proofs of the Lemma were known.

We originally became interested in the problem of finding a constructive proof of Higman's Lemma via Gabriel Stolzenberg. He drew our attention to the fact that Friedman's A-translation [Fri77] would lead one to believe that Higman's Lemma had a constructive proof, but he could not see such in the classical proof. That is, he could not interpret the classical proof in such a way that the construction (which must be implicit, due to Friedman's result), could be seen. In the course our work on A-translating the classical proof

*Supported in part by an NSF graduate fellowship
†Supported in part by an IBM graduate fellowship

using the NuPRL proof development system [Con86], we discovered a direct constructive proof of the Lemma. Our proof resembles an algorithmic specification, which makes the ordinal of induction quite apparent, and the proof itself easy to follow.

We view our work on Higman's Lemma as a first step towards constructive proofs of Kruskal's tree embedding theorem and the graph minor theorem. Also, since in our proof the ordinal of induction is evident, it may assist an investigation of the minimal theories in which such a proof exists. Finally, the direct constructive proof provides a context in which we can compare and evaluate the result of the mechanical A-traslation of the classical proof.

We now present some essential definitions and the statement of Higman's Lemma.

Definition 1. Given a set Σ, a binary relation $\leq$ is a *well-quasi-order (WQO)* if and only if $\leq$ is a preorder on Σ, and for any infinite sequence $s_1, s_2, s_3, \ldots$ of elements of Σ there are $i < j$ such that $s_i \leq s_j$.

Definition 2. Given strings $u = u_1 u_2 \cdots u_m$ and $v = v_1 v_2 \cdots v_n$ in Σ^*, we say that u can be *embedded* in v (written $u \ll v$) if and only if there is an injective, order-preserving mapping g from $\{1, \ldots, m\}$ to $\{1, \ldots, n\}$ such that for all i in $\{1, \ldots, m\}$, $u_i \leq v_{g(i)}$. Simply put, u can be embedded in v if we can erase some of the v_i and produce a string which is pointwise above u. We call u a *substring* of v, and v a *superstring* of u.

Note that the string embedding $\ll$ is a preorder on Σ^*. Higman's lemma states the following:

If $\leq$ is a WQO on Σ, then $\ll$ is a WQO on Σ^*.

2 A Classical Proof of the Lemma

In this section, we sketch a classical proof of Higman's Lemma. The classical proof of Higman's Lemma is interesting because it contains a *minimal bad sequence* (MBS) argument. The proof proceeds by assuming that Higman's Lemma is false, and arguing to a contradiction, by defining, by transfinite induction, a MBS. This sequence turns out not to be minimal, and we get a contradiction. This MBS argument turns up in Kruskal's Theorem, of which Higman's Lemma is but a small part. The version of the proof we present is due to Gallier [Gal87], and originally presented in this form by Nash-Williams [NW63].

First we state a (classically) trivial lemma.

Lemma 1. For any sequence $(a_i)_{i\geq 1}$ over a WQO set Σ, there is a subsequence $a_{\sigma(1)}, a_{\sigma(2)}, a_{\sigma(3)}, \ldots$ such that for all i, $a_{\sigma(i)} \leq a_{\sigma(i+1)}$ (i.e. the sequence $(a_{\sigma(i)})_{i\geq 1}$ is monotonically increasing).

We now prove Higman's Lemma.

Theorem 1 (Higman's Lemma)

$$\forall \langle \Sigma, \leq \rangle . \ (\leq \text{ is a WQO on } \Sigma) \ \Rightarrow \ (\ll \text{ is a WQO on } \Sigma^*).$$

Proof: Assume that $\ll$ is not a WQO on Σ^*. Then there is at least one sequence $s = s_1, s_2, s_3, \ldots$ such that for all i, j with $i < j$, $s_i \not\ll s_j$. We will call such a sequence *bad*, and we will call non-bad sequences *good.* We define, by impredicative quantification across all bad sequences, a minimal bad sequence, $t = t_1, t_2, t_3, \ldots$ as follows:

- Let t_1 be a string of minimal length which starts a bad sequence.
- Given that $t_1, t_2, t_3, \ldots, t_i$ have been defined, let t_{i+1} be an $i+1$-st string of minimal length from all bad sequences which start with $t_1, t_2, t_3, \ldots, t_i$.

It is easily seen that the new sequence t is bad, since any prefix of it is also a prefix of a bad sequence. Note that there can be *no* instances of ϵ (the empty string) in t since ϵ may be embedded into any string. It is also clear that t is a minimal sequence, in the sense that for any other bad sequence, x, there exists a $k \in \mathsf{N}$ such that $x_i = t_i$ for $i \leq k$, and $|t_{k+1}| \leq |x_{k+1}|$.

Since all strings in t are non-null, let $t_i = a_i s_i$ where $a_i \in \Sigma$ is the leftmost symbol of t_i. The elements a_i define an infinite sequence $a = (a_i)_{i\geq 1}$ in Σ and the elements s_i define an infinite sequence $(s_i)_{i\geq 1}$ in Σ^*. By the lemma above, we know there is a monotonically increasing subsequence $a' = (a_{\sigma(i)})_{i\geq 1}$. We claim that the corresponding sequence $s' = (s_{\sigma(i)})_{i\geq 1}$ is good. If not, there are two cases:

1. $\sigma(1) = 1$: The infinite sequence s' is bad, with $|s_1| < |t_1|$, which contradicts the minimality of t.
2. $\sigma(1) > 1$: The infinite sequence $s' = t_1, t_2, t_3, \ldots, t_{\sigma(1)-1}, s_{\sigma(1)}, s_{\sigma(2)}, s_{\sigma(3)}, \ldots$ is also bad, because $t_k = a_k s_k$ for all $k \geq 1$, and $t_i \ll s_{\sigma(j)}$ implies that $t_i \ll t_{\sigma(j)}$ by the definition of $\ll$. But $|s_{\sigma(1)}| < |t_{\sigma(1)}|$, and this contradicts the minimality of t.

Since the sequence $s' = (s_{\sigma(i)})_{i\geq 1}$ is good, there exist positive integers $i < j$ such that $\sigma(i) < \sigma(j)$ and $s_{\sigma(i)} \ll s_{\sigma(j)}$. We know a' is monotonically increasing, so we have

$$t_{\sigma(i)} = a_{\sigma(i)}s_{\sigma(i)} \ll a_{\sigma(j)}s_{\sigma(j)} = t_{\sigma(j)}.$$

Hence, we conclude t is good, which is a contradiction.

∎

The important things to notice about this proof are

1. The use of excluded middle.

2. The use of impredicative quantification across the set of bad sequences.

Of course, our constructive proof cannot use either of these two devices, so it will be a bit more involved.

3 The Constructive Proof

3.1 Requirements and Preliminaries

Throughout this section we will assume Σ is well-quasi-ordered by $\leq$. Because we are working in constructive mathematics, we require certain other assumptions on $\langle\Sigma, \leq\rangle$. To wit, we require:

- $\leq$ must be a partial order on Σ. In the section on Extensions to this result, we discuss how to relax this requirement to allow $\leq$ to be a general preorder.

- $\leq$ come equipped with a well-founded induction scheme over non-increasing sequences of elements from Σ. The fact that $\leq$ is a WQO on Σ tells us that non-increasing sequences of elements from Σ may not be infinite. We require more, to wit, that we may use well-founded induction upon these sequences, under the prefix ordering. Classically, this is easily gotten from the WQO-ness of $\leq$, but constructively we must have this as an assumption. Of course, this assumption also tells us that $\leq$ is a well-founded order, with an induction scheme. We denote this ordering by $\sqsubset_{seq}$, and $A \sqsubset_{seq} B$ exactly when A and B are both non-increasing sequences over Σ, and A is a proper extension of B. After a moment's reflection, it should be obvious to the reader that the constructive equivalent of the classical notion of WQO is a well-founded order on sequences under the prefix ordering. That is, the classical notion of WQO entails that there are no infinite decreasing or incomparable chains. This tells us classically that the prefix ordering on sequences of elements of Σ is well-founded, and so we assume this constructively. The proof we present will basically prove that the prefix ordering on sequences of elements from $\Sigma^*$is well-founded also, from which the Higman's Lemma result will fall out.

- $\leq$ is decidable. This requirement is obvious.

We also require that the proof system be powerful enough to:

- encode and manipulate the decidable fragment of regular expression theory.
- manipulate regular expressions as structured data types in the theory.
- do well-founded induction on a (constructively) well-founded set which will be given later.

Finally, in the proof we will make considerable use of a class of regular expressions we will call *sequential* regular expressions. Sequential r.e.'s are (possibly empty) concatenations of r.e.'s from the following two categories:

- Expressions of the form $(b - A)$, where $b \in \Sigma$, and A is a finite non-increasing sequence, say of length k, over Σ. $x \in (b - A) \Leftrightarrow x \leq b \mathbf{and} \forall i \in \{1, \ldots, k\}. a_i \not\leq x$. That is, x is beneath b, but not in the upward closure of A. We will also write these expressions as $(b - A)^k$, which is understood to be equivalent for all purposes to k concatenations of $(b - A)$. We call these *constant* expressions.

- Expressions of the form $(\Sigma - A)^*$, where $A = a_1, a_2, a_3, \ldots, a_l$, is a non-increasing sequence, and we think of this expression as denoting the set

$$\{w \in \Sigma^* | \text{no symbol of } w \text{ is in the upward closure of } A\}.$$

  We call these *starred* expressions.

An example of a sequential r.e. for the set $\Sigma = \{a, b, c, d\}$ is $(\Sigma - (b))^* bd(\Sigma - (c, d, b))^* a$.

In the case where $\Sigma = \{1, \ldots, m\}$, we can specialize sequential r.e.'s and, for instance, instead of $\{1, 2, 3, 4\} - (2, 4)$, write $\{1, 3\}$. Of course, in this case, the partial order $\leq$ is just the equality on integers.

We need to use this kind of r.e. so that we can "subtract" a string x, and all strings which it can be embedded into, from an r.e. σ and end up with a finite set of "simpler" r.e.'s as a result. We will write $s \in \sigma$ to mean that the string s matches the r.e. σ.

3.2 An Overview of the Constructive Proof with Examples

In this section we will give a general overview of the constructive proof and present some examples to motivate the detailed version which follows. The object of our proof is to show that any sequence of strings in Σ^* is good – that is, for $s = s_1, s_2, s_3, \ldots$ there are $i < j$ with $s_i \ll s_j$. The key to this is the observation that for a prefix $s_1, s_2, \ldots, s_n$ of s we can use a set of sequential r.e.'s to represent all the strings that could follow s_n without containing any $s_i, 1 \leq i \leq n$, as a substring. If for some prefix of s the corresponding set of r.e.'s is empty, then any following string must contain one in the prefix as a substring and the sequence s must be good. As we consider successively longer prefixes of s, we can see intuitively that there are "fewer" possible following strings, and that the corresponding sets of r.e.'s must get "smaller" (closer to empty). In essence, what our constructive proof does is formalize the construction of the sets of sequential r.e.'s, and inductively show that starting with any prefix the corresponding sets must eventually be reduced to empty, and hence any sequence s is good.

To give a general feel for this construction, we present some examples. For a prefix $s_1, s_2, \ldots, s_{k-1}$, we will call the corresponding set of r.e.'s E_k.

Example 1. Suppose $\Sigma = \{1, 2\}$ and the sequence s under consideration begins $1, 22222, 2222, 22, \ldots$. We construct the E_i as follows:

1) $E_1 = \{\Sigma^*\} = \{(1+2)^*\}$.

2) $s_1 = 2$, so E_2 should represent all strings which do not contain "1" – that is, strings in 2^*; any other string will be a superstring of s_1. Thus, we set $E_2 = \{2^*\}$.

3) $s_2 = 22222$, so E_3 should represent strings which don't contain a "1", and don't have "22222" as a substring. These are the strings "2222", "222", "22", "2", and ϵ, so we set $E_3 = \{2222, 222, 22, 2, \epsilon\}$.

4) $s_3 = 2222$, so E_4 is further restricted to $\{222, 22, 2, \epsilon\}$.

At this point we see that there are at most four more strings in s before we find one containing a previous string. Hence, s must be good.

Example 2. This is a more complicated example. $\Sigma = \{1, 2\}$ as before, and we will just show the strings s_i, and the regular expression sets E_i:

1) $E_1 = \{(1+2)^*\}$.

2) $s_1 = 12$; $E_2 = \{2^*1^*\}$.

3) $s_2 = 221111$; $E_3 = \{221^*, 21^*, 1^*, 2^*111, 2^*11, 2^*1, 2^*\}$.

4) $s_3 = 22111$; $E_4 = \{221^*, 21^*, 1^*, 2111, 111, 2^*11, 2^*1, 2^*\}$.

It is not as obvious as in example 1, but the sets E_i are "decreasing" in this example as well (the sense in which they are decreasing will be made precise in section 4.4). Although the size of the sets may not decrease at every step, their complexity, as measured by the number of $+$'s and $*$'s appearing in the expressions, does. Similarly, while we cannot say that at this point we have limited the number of possible following strings (as in example 1), we have limited the number of *forms* these strings may take. And after a moment's reflection, the reader will realize that after one instance of each of the forms in E_4 we will have a finite list limiting the number of possible following strings.

In the rest of this section we formalize the concepts mentioned above and present our proof.

3.3 The Construction of Replacement Sets

A key aspect of our constructive proof is the idea of replacing a regular expression σ with a finite set of regular expressions that collectively match a smaller set of strings than σ. Specifically, given σ and a string $s \in \sigma$, we need the ability to construct a finite set of r.e.'s that collectively match any string *not containing s as a substring* that σ also matches. We describe such a construction in this subsection.

Definition 3. Given a string s, we define s° to be the set of all strings that s can be embedded into. That is, $s^\circ \stackrel{\text{def}}{=} \{x \in \Sigma^* | s \ll x\}$.

Claim 1. Given a sequential r.e. σ and a string $w \in \sigma$, we can construct a finite set of sequential r.e.'s $\Theta(\sigma, w)$ with the property that if $x \in \sigma$ and $w \not\ll x$ then there is an r.e. $\theta \in \Theta(\sigma, w)$ with $x \in \theta$.

Proof: We will describe $\Theta(\sigma, w)$ and show that it has the desired property. The construction is divided into three cases by the structure of σ.

Case 1 — σ is a constant expression: Say $\sigma = (b - A)$ with $w \in \sigma$. We know immediately that w is a single symbol, $w \leq b$, and that w is not in the upward closure of A. We wish to define an r.e. θ such that $x \in (b - A)$ and $w \not\leq x$ implies $x \in \theta$. This is easy:

let $\theta = (b - (A|w))$, where $(A|w)$ is the sequence A with w appended to the end. In this case the set $\Theta(\sigma, w)$ is the singleton $\{\theta\}$, and clearly has the desired property.

Case 2 — σ is a starred expression: Say $\sigma = (\Sigma - A)^*$ and $w \in \sigma$. We want to construct a set of sequential r.e.'s $\Theta(\sigma, w)$ so that $w \not\leq x \Rightarrow (\exists\theta \in \Theta(\sigma, w).x \in \theta)$. What this means is that any string of which w is not a substring will be accepted by some r.e. in $\Theta(\sigma, w)$. In the following, we will write $A|a$ to mean the sequence A followed by the symbol a.

We build the required set of regular expressions $\Theta(\sigma, w)$ in two steps. First, let $w = p_1p_2\cdots p_l$ with $p_i \in \Sigma$, and build the following r.e.:

$$(\Sigma - (A|p_1))^*(p_1 + \epsilon)(\Sigma - (A|p_2))^*(p_2 + \epsilon)\cdots(p_{l-1} + \epsilon)(\Sigma - (A|p_l))^*.$$

Second, distribute the "+" operators over the concatenations, to yield a disjunction of sequential r.e.'s. These sequential r.e.'s comprise $\Theta(\sigma, w)$. It is important to note that if w is the empty string, then $\Theta(\sigma, w) = \emptyset$.

As an example, if $\Sigma = \{1, 2, 3, 4\}$, $\sigma = (\Sigma - (2))$, and $w = 1441$, then we get the r.e.

$$(3+4)^*(1+\epsilon)(1+3)^*(4+\epsilon)(1+3)^*(4+\epsilon)(3+4)^*$$

which, after distributing $+$, becomes the set of r.e.'s

$$\begin{aligned}
\Theta(\sigma, w) \quad = \quad & \{(3+4)^*(1+3)^*(1+3)^*(3+4)^*, \\
& (3+4)^*1(1+3)^*(1+3)^*(3+4)^*, \\
& (3+4)^*(1+3)^*4(1+3)^*(3+4)^*, \\
& (3+4)^*1(1+3)^*4(1+3)^*(3+4)^*, \\
& (3+4)^*(1+3)^*(1+3)^*4(3+4)^*, \\
& (3+4)^*1(1+3)^*(1+3)^*4(3+4)^*, \\
& (3+4)^*(1+3)^*4(1+3)^*4(3+4)^*, \\
& (3+4)^*1(1+3)^*4(1+3)^*4(3+4)^*\}.
\end{aligned}$$

Of course, some of these r.e.'s accept the same language as others, but that's not important because there is still only a finite list of them. It is clear from the construction that if $x \not\leq w$ then there exists $\theta \in \Theta(\sigma, w)$ with $x \in \theta$. The proof is simple, but tedious.

Case 3 — σ is a sequential r.e.: Say $\sigma = a_1a_2\cdots a_n$ where each a_i is either a constant expression or a starred expression, and $w \in \sigma$. Then we must be able to break w into n (possibly empty) pieces which match respectively the expressions a_i. That is, there exist (possibly empty) strings $w_1, w_2, \ldots, w_n$ such that $w = w_1 w_2 \cdots w_n$ and $w_i \in a_i$ for all $i \in \{1, \ldots, n\}$.

Now we define $\Theta(\sigma, w)$ by

$$\Theta(\sigma, w) \stackrel{\text{def}}{=} \bigcup_{i=1}^{n} \{a_1 \cdots a_{i-1} \theta a_{i+1} \cdots a_n \mid \theta \in \Theta(a_i, w_i)\}.$$

In other words, we build the set of all sequential r.e.'s derived from σ by replacing a subexpression a_i with an r.e. in $\Theta(a_i, w_i)$. As in the previous case, note that if $w = \epsilon$, then $\Theta(\sigma, w) = \emptyset$.

For any $x \in \sigma$, we can break x into pieces, $x = x_1 x_2 \cdots x_n$, with $x_i \in a_i$ for all i. If $w \not\preccurlyeq x$ then we know $w_j \not\preccurlyeq x_j$ for some $j \in \{1, \ldots, n\}$. From the previous cases we know that there is $\theta \in \Theta(a_j, w_j)$ with $x_j \in \theta$, hence we have

$$x \in a_1 \cdots a_{i-1} \theta a_{i+1} \cdots a_n.$$

∎

The idea behind this construction is that for a string to *not* be a superstring of w, it can only contain a *proper* substring of w. So what we do is write down r.e.'s which each accept classes of strings containing *different* substrings of w.

Lemma 2. Given a set of finite strings $S \subseteq \Sigma^*$, and a finite set of r.e.'s E such that for every $s \in S$ there is $\sigma \in E$ with $s \in \sigma$, then for any such s, σ we have

$$\forall s' \in S \setminus s^\circ.\ \exists \sigma' \in (E \setminus \{\sigma\}) \cup \Theta(\sigma, s).\ s' \in \sigma'.$$

In other words, if we remove the r.e. σ from E and replace it with the set $\Theta(\sigma, s)$, then the resulting set of r.e.'s represents (at least) all the strings in S not containing s.

Proof: $s' \in S \setminus s^\circ \subseteq S$, so there is $\tau \in E$ with $s' \in \tau$. If $\tau \neq \sigma$, then $\tau \in (E \setminus \{\sigma\}) \cup \Theta(\sigma, s)$. If $\tau = \sigma$, then since $s \not\preccurlyeq s'$ we conclude by the previous claim that there is $\sigma' \in \Theta(\sigma, s)$ with $s \in \sigma'$.

3.4 A Well-Founded Ordering

In this section we will construct a well-founded ordering (hereinafter WFO) on finite sets of sequential r.e.'s and show that the replacement of an r.e. by a set $\Theta(\sigma, s)$ as in lemma 2 reduces the original set in this order. As usual, we proceed by constructing the WFO. We begin with a general lemma.

Lemma 3. If $\sqsubset$ is a WFO over a set Γ, then $\sqsubset^M$ is a WFO over finite multisets (sets with repetitions) of elements of Γ, where $\sqsubset^M$ is defined as follows: $A \sqsubset^M B$ iff there exist U, V, W such that

$$A = U \uplus V$$
$$B = U \uplus W$$
$$\forall v \in V \; \exists w \in W. \; v \sqsubset w$$

(we use the symbol $\uplus$ to denote multiset union, in which multiple copies of the same element are not identified). That is, we obtain A from B by removing elements from B, and replacing them with finitely many elements smaller in the $\sqsubset$ order.

Proof: Omitted.

We will now proceed to construct a number of WFO's, building up to the one we want.

By the assumption in section 4.1 we have a WFO, $\sqsubset_{seq}$, on finite, non-increasing sequences over Σ.

We can put a WFO on starred expressions, by mapping $(\Sigma - A)$ to A. In this ordering, if a is not in the upward closure of A, $(\Sigma - (A|a))$ is beneath $(\Sigma - A)$. We denote this ordering by $\sqsubset_*$, and $(\Sigma - A)^* \sqsubset_* (\Sigma - B)^*$ iff $A \sqsubset_{seq} B$.

We can put a WFO on constant expressions likewise, and we can then put a WFO on the set of constant *and* starred expressions by placing all constant expressions *below* all starred expressions. We denote the order on constant expressions by $\sqsubset_{const}$ and the order on the union of both classes by $\sqsubset_{exp}$.

We can put a well-founded order on finite multisets of starred and constant r.e.'s, where A is beneath B exactly when we can transform B to A by removing elements from B and replacing them with a finite number of elements which are lesser in the order $\sqsubset_{exp}$. The preceding lemma justifies this step, and we call this ordering $\sqsubset_{setexp}$.

We can put a WFO on sequential r.e.'s, in which we map each sequential r.e. to a finite multiset of it's component starred and constant expressions, and use the previous ordering on these constructed sets. We denote this order by $\sqsubset_{re}$, and $\alpha \sqsubset_{re} \beta$ when $\alpha = a_1 \cdots a_k$, $\beta = b_1 \cdots b_l$, and $\biguplus_{i=1}^{k}\{a_i\} \sqsubset_{setexp} \biguplus_{i=1}^{l}\{b_i\}$.

Finally, we can put a WFO on finite sets of sequential r.e.'s by direct application of the above lemma, and we denote this ordering by $\sqsubset_{setre}$.

Lemma 4. All of the orders defined above are WFO's.

Proof: Straightforward from the definitions.

Lemma 5. Given an r.e. σ and a string $w \in \sigma$, then for all $\theta \in \Theta(\sigma, w)$ we have $\theta \sqsubset_{re} \sigma$.

Proof: Direct from the definitions.

Lemma 6. Given a set of r.e.'s E, an r.e. $\sigma \in E$ and a string $s \in \sigma$, we have

$$(E \setminus \{\sigma\}) \cup \Theta(\sigma, w) \ \sqsubset_{setre} \ E.$$

Proof: This follows from the definition of $\sqsubset_{setre}$ and the previous lemma, from which we know every r.e. in $\Theta(\sigma, w)$ is beneath σ.

3.5 The Proof Itself

In this section we give the constructive proof of Higman's Lemma. Given a set of r.e.'s E, we will write $\mathcal{L}(E)$ for the set $\{w \in \Sigma^* | \exists \sigma \in E \text{ with } w \in \sigma\}$.

Theorem 2. Let $s = s_1, s_2, \ldots$ be a sequence of strings in Σ^*. Given any set of r.e.'s E, then for any integer $k \geq 1$ the following holds:

$$\Sigma^* \setminus (s_1{}^\circ \cup s_2{}^\circ \cup \cdots s_{k-1}{}^\circ) \subseteq \mathcal{L}(E) \ \Rightarrow \ \exists j \geq k.\ \exists i < j.\ s_i \ll s_j.$$

Proof: By induction on the set E.

Assume $E = \emptyset$. If for some k

$$\Sigma^* \setminus (s_1{}^\circ \cup s_2{}^\circ \cup \cdots s_{k-1}{}^\circ) = \emptyset,$$

then clearly $s_k \in s_i{}^\circ$ for some $i \leq k-1$, hence $s_i \ll s_k$.

Now assume $\emptyset \sqsubset_{setre} E$. Suppose for some $k \geq 1$ we have

$$\Sigma^* \setminus (s_1{}^\circ \cup s_2{}^\circ \cup \cdots s_{k-1}{}^\circ) \subseteq \mathcal{L}(E).$$

If $s_k \notin \mathcal{L}(E)$, then $s_k \in s_i{}^\circ$ for some $i \leq k-1$ and $s_i \ll s_k$. If $s_k \in \sigma$ for some $\sigma \in E$, then let $E' = (E \setminus \{\sigma\}) \cup \Theta(\sigma, s_k)$. Then by lemma 6 we know $E' \sqsubset_{setre} E$, and by lemma 2 we know that

$$\Sigma^* \setminus (s_1{}^\circ \cup s_2{}^\circ \cup \cdots s_k{}^\circ) \subseteq \mathcal{L}(E').$$

Hence, by the induction hypothesis, we conclude that there is $j \geq k+1$ and $i < j$ with $s_i \ll s_j$.

∎

Corollary 1 (Higman's Lemma) Any sequence of strings in Σ^* is good – that is, $\ll$ is a WQO on Σ^*.

Proof: For a sequence $s_1, s_2, \ldots$ of strings in Σ^*, we apply the theorem with $E = \{\Sigma^*\}$ and $k = 1$, and conclude that there are $i < j$ with $s_i \ll s_j$. Hence $\ll$ is a WQO on Σ^*.

4 Extensions

In this proof we assumed that $\leq$ was a decidable partial order. if instead we let $\leq$ be a decidable preorder, and in our construction let our equality on elements of Σ be $a \equiv b$ iff $a \leq b$ **and** $b \leq a$, then we can use this proof for arbitrary preorders, too. We still require a well-founded induction scheme, though, and in this case the scheme must require that the predicate which is being proven by induction respect the equivalence relation induced by the preorder. Also, when, in the proof we say $a < b$, meaning $a \leq b$ **and** $a \neq b$, we must now use $a \leq b$ **and** $b \not\leq a$. This guarantees that we can perform our induction on the preorder.

5 Conclusions and Directions for Further Work

We have presented a direct, constructive proof of Higman's Lemma for arbitrary preorders, where a well-founded induction scheme which respects the equality induced by the preorder is provided. Further research on the complexity of the algorithm in terms of the input sequence, and the applicability of this approach to Kruskal's Theorem, are in progress. We do not see much chance of applying the construction here presented in anything like a direct manner, but we do hope that the insight gained here will allow us to make progress toward a direct proof of Kruskal's Theorem.

Note

While privately circulating this proof for comments, we learned that Friedman has recently proven Higman's Lemma in a classical system, *without* impredicativity. Also, we have recently become aware that Richman and Stolzenberg [RS89] have independently developed a constructive proof of Higman's Lemma which is very similar to ours.

References

[Con86] Robert L. Constable, et al. *Implementing Mathematics with the Nuprl Proof Development System.* Prentice-Hall, Englewood Cliffs, New Jersey, 1986.

[Fri77] Harvey Friedman. Classically and intuitionistically provable recursive functions. Technical report, Ohio State University, September 1977.

[Gal87] J. H. Gallier. What's so special about Kruskal's Theorem and the ordinal Γ_0. Technical Report MS-CIS-87-27, University of Pennsylvania, Philadelphia, PA, April 1987.

[Hig52] G. Higman. Ordering by divisibility in abstract algebras. In *Proc. London Math. Soc.*, volume 2, pages 236–366, 1952.

[NW63] C. St. T A. Nash-Williams. On well-quasi-ordering finite trees. In *Proc. Cambridge Phil. Soc.*, volume 59, pages 833–835, 1963.

[RS89] F. Richman and G. Stolzenberg. Well quasi-ordered sets yellow, 1989. Unpublished paper.