# Cluster Computing and the Power of Edge Recognition 

Lane A. Hemaspaandra*<br>Department of Computer Science<br>University of Rochester<br>Rochester, NY 14627<br>USA

Christopher M. Homan ${ }^{\dagger}$<br>Department of Computer Science<br>Rochester Institute of Technology<br>Rochester, NY 14623<br>USA

Sven Kosub ${ }^{\ddagger}$<br>Institut für Informatik<br>Technische Universität München<br>D-85748 Garching b. München<br>Germany<br>Univ. of Rochester Comp. Sci. Dept. Technical Report TR-2005-878<br>September 16, 2005


#### Abstract

We study the robustness - the invariance under definition changes - of the cluster class CL\#P [HHKW05]. This class contains each \#P function that is computed by a balanced Turing machine whose accepting paths always form a cluster with respect to some length-respecting total order with efficient adjacency checks. The definition of CL\#P is heavily influenced by the defining paper's focus on (global) orders. In contrast, we define a cluster class, CLU\#P, to capture what seems to us a more natural model of cluster computing. We prove that the naturalness is costless: CL\#P $=$ CLU\#P. Then we exploit the more natural, flexible features of CLU\#P to prove new robustness results for CL\#P and to expand what is known about the closure properties of CL\#P.

The complexity of recognizing edges - of an ordered collection of computation paths or of a cluster of accepting computation paths - is central to this study. Most particularly, our proofs exploit the power of unique discovery of edges - the ability of nondeterministic functions to, in certain settings, discover on exactly one (in some cases, on at most one) computation path a critical piece of information regarding edges of orderings or clusters.


[^0]
## 1 Introduction

Cluster computing, in the complexity-theoretic use of the term, was introduced by Kosub in [Kos99], though he notes there that there was earlier work that focused in a rather different sense on cluster-like behavior ([Wag90,VW95] and we mention in passing that the so-called telescoping normal form of the boolean hierarchy $\left[\mathrm{CGH}^{+} 88\right]$ and the parallel census technique of Selman [Sel90, GH00] also provide early examples of the type of behavior Kosub was there observing, namely, settings in which "yes" answers always occur in a contiguous block). In particular, Kosub defined and studied the class c\#P, which is the set of all \#P functions computed by (i.e., given by the number of accepting paths of) lexicographical cluster machines-loosely put, machines such that on each input, all the accepting paths are lexicographically adjacent to each other (they form a contiguous block). He obtained quite comprehensive results, but they depended critically on the very simple structure of lexicographical order, namely, that if one knows the left and right edges of a lexicographical cluster, it is easy to compute the size of the cluster.

Yet the underlying motivating issue - to what extent does requiring that all accepting paths be closely related in some order restrict the ability of nondeterministic Turing machines to compute \#P functions? - certainly is not tied to the artificial simplicity of lexicographical order. Just as self-reducibility [Sch76, MP79] has not only been defined with respect to a focus on the lexicographical order and decreasing chains with respect to length there (as in [BDG95, BDG90]) but also (and most elegantly) has been defined with respect to having polynomially length-bounded decreasing chains within appropriate more general classes of orders (as in [MP79,Sch76]), so also is it natural to study cluster computing with respect to more flexible ordering.

To imagine how to naturally do this, we think of the model underlying $\mathrm{c} \# \mathrm{P}$, which, again, is the class of functions that are the numbers of accepting computation paths of balanced (a Turing machine is balanced if there is some polynomial $p$ such that on each input $x$ it holds that each nondeterministic path has exactly $p(|x|)$ binary nondeterministic guesses) Turing machines in which the accepting paths are always lexicographically adjacent. So the accepting block on a given input is, assuming any paths accept, just a lexicographically contiguous block among the length $p(|x|)$ strings, where one viewsas we do throughout this paper-each accepting path (on a given input) as being named by its nondeterministic guesses. Intuitively speaking, we suggest that it might be very natural to generalize this by keeping essentially the entire setting mentioned above, except on input $x$ viewing the strings at length $p(|x|)$ not as being in lexicographical order, but rather viewing them as follows. For each balanced nondeterministic machine whose number of accepting paths defines a function in our new class, there must be polynomial-time computable functions $b$ (the bottom function), $t$ (the top function), and $\prec$ (the adjacency function) such that: We view $b(x) \in\{0,1\}^{p(|x|)}$ as the least string of length $p(|x|) ; \prec(x, y, z)$
tells whether on input $x$ the string $z \in\{0,1\}^{p(|x|)}$ comes immediately after $y \in\{0,1\}^{p(|x|)}$ in our linear ordering of the length $p(|x|)$ strings; if one using those two functions starts at $b(x)$ and moves through one string after another under the adjacency rule specified by $\prec(x, \cdot, \cdot)$, one goes though each string of length $p(|x|)$ and ends up at $t(x)$; and if there are any accepting paths on input $x$, then all the accepting paths on input $x$ form a cluster-a contiguous block-within this ordering. In particular, regarding the ordering, we allow an arbitrary linear ordering of the length $p(|x|)$ strings subject to it being easy to tell the biggest and smallest elements in our new order, and to recognize adjacency in our new order. Let us call the class thus defined CLU\#P.

Though we suggest that the CLU\#P definition and model are very easy to work with, it is very important to note that a previous paper already defined a generalization of Kosub's notion with exactly the goal of handling more general orderings. In particular, this was done by [HHKW05], resulting in the class CL\#P. CL\#P's definition, however, is heavily influenced by the overall focus of that paper on global orders (rather than input-specific orderings). In particular, that paper requires all inputs to have their computation paths share the same order with respect to defining what it means to be a cluster. For example, if on input $x$ computation paths $y$ and $z$ exist and $y \prec z$ (respectively, $y \nprec z$ ), then for each input $x^{\prime}$ on which those computation paths exist (namely, all strings $x^{\prime}$ on which the nondeterminism polynomial happens to evaluate to the same value on $\left|x^{\prime}\right|$ as it does on $|x|$, and so certainly for all strings $x^{\prime}$ of the same length as $x$ ) it must also hold that $y \prec z$ (respectively, $y \nprec z$ ). Further, the fact that that paper really requires a global-over all of $\Sigma^{*}$ - order forces the ordering for each input $x$ to smoothly link the strings related to computation on input $x$ to the other, utterly irrelevant paths. Although these constraints are arguably reasonable in a paper whose focus is on global, total orders (in the formal sense), we here suggest that if one were to simply take the idea of Kosub and shuffle ${ }^{1}$ the paths that apply to that input, the notion of CLU\#P would seem a more natural approach to and model of doing that.

Fortunately, one does not have to choose between the classes CL\#P and CLU\#P. This is because our main result is that the new class CLU\#P, which was defined to directly capture a natural, local, machine-directed notion of cluster computing, has exactly the same descriptive power as the class CL\#P, which is based on a global shared order: CL\#P = CLU\#P. This result is in Section 3, which also shows another robustness result that will be central to our later study of two other notions-free cluster and circular cluster machines. That other robustness result is essentially that unambiguity of cluster edge recognition is sufficient to ensure that even seemingly more flexible models in fact generate just the CL\#P functions.

[^1]Section 4, partially by using our newfound freedom to study CL\#P by studying CLU\#P, shows a number of closure properties of CL\#P. For example, [HHKW05] proved that if CL\#P is closed under increment then UP $=$ coUP, and we show that the converse holds.

Our model, CLU\#P-type machines, has $\mathrm{FP}_{\mathrm{t}}$ (total, polynomial-time computable) top and bottom elements on each input. Section 5 studies two alternate models. CLU\#P $\mathrm{P}_{\text {free }}$ removes any explicit complexity requirement regarding the top and bottom elements. CLU\# $\mathrm{P}_{\text {circular }}$ requires the ordering on our computation paths to be circular-thus there is no top or bottom element. We prove a number of results about these classes, and most particularly we show (a) that $\mathrm{UP}=\mathrm{coUP}$ is a sufficient condition for $\mathrm{CL} \# \mathrm{P}=\mathrm{CLU} \# \mathrm{P}_{\text {free }}=$ CLU\# $\mathrm{P}_{\text {circular }}$, and (b) that UP $=$ coUP is a necessary condition for $\mathrm{CL} \# \mathrm{P}=\mathrm{CLU} \# \mathrm{P}_{\text {free }}$, and even for $\mathrm{CLU} \# \mathrm{P}_{\text {free }} \subseteq \mathrm{CLU} \# \mathrm{P}_{\text {circular }}$. Result (b) can be viewed as reasonably strong evidence that CLU\# $\mathrm{P}_{\text {free }}$ is a strictly more powerful, flexible class than CL\#P, and can also be viewed as reasonably strong evidence that some CLU\# $\mathrm{P}_{\text {free }}$ functions are not in CLU\# $\mathrm{P}_{\text {circular }}$. So freeing the endpoints from their $\mathrm{FP}_{\mathrm{t}}$ constraint seems to yield a real increase in descriptive power.

The proofs in this paper are thematically linked. Most of them focus on the power of what we will call "unique discovery" of facts about about top and bottom elements and about greatest and least accepting paths-i.e., about "edges." By unique discovery we mean that critical pieces of edge-related information used in our proofs are partial or total UPSV (unambiguous polynomial-time single-valued) functions [GS88,Kos99]. Informally speaking, we mean that our proof strategy will often be:

1. Seek to guess some critical piece of information (such as the right edge of a cluster).
2. If we succeeded on the current path in guessing that information correctly, do FOO and otherwise do BAR,
and, critically, our settings will variously ensure that in step (1) either exactly one or at most one path guesses the critical information, that that path "knows" it has done so (i.e., could write on an output tape the information and set a bit declaring it has successfully obtained the information), and each other path knows that it has not done so.

## 2 Definitions

$\Sigma=\{0,1\}$ will be our alphabet. The boolean relation $\prec_{\text {lex }}$ is defined as: $a \prec_{\text {lex }} b$ is true when $b$ is the lexicographical successor of $a$ and is false otherwise, e.g., $111 \prec_{\text {lex }} 0000$ and $010 \prec_{\text {lex }} 011$, but $00 \varliminf_{\text {lex }} 11$. We use NPTM as a shorthand for "nondeterministic polynomialtime Turing machine." As is common, for a given nondeterministic machine $M$ and a string $x, \operatorname{acc}_{M}(x)$ denotes the set of accepting paths of machine $M$ on input $x$, and $\# a c c_{M}(x)$
is defined as $\left\|\operatorname{acc}_{M}(x)\right\| . \mathrm{FP}_{\mathrm{t}}$ denotes the total, polynomial-time computable functions (usually from $\Sigma^{*}$ to $\Sigma^{*}$ ).

Given any string $x \in \Sigma^{*}$ and any integer $n \leq|x|$, $\operatorname{prefix}(x, n)$ denotes the first $n$ bits of $x$ and $\operatorname{suffix}(x, n)$ denotes the last $n$ bits of $x$. If $n>|x|$ these functions are undefined.

For each polynomial $p$ and each NPTM $M, M$ will be said to be $p$-balanced (see [Kos99]) exactly if for each input $x$ the set of nondeterministic guesses along the computation paths of $M$ is precisely $\{0,1\}^{p(|x|)}$. That is, $M$ on input $x$ has exactly $2^{p(|x|)}$ computation paths, one corresponding to each possible guess of $p(|x|)$ bits. Note that we do not require that each step of the machine involves a nondeterministic guess.

We turn immediately to defining the central class of this paper, CLU\#P. In defining CLU\#P, we seek to keep Kosub's notion of a cluster as a block of adjacent paths, but we allow that adjacency to be with respect to a "shuffing" of the paths, rather than to have to be with respect to lexicographical order. However, the shuffle must be simple enough that in polynomial time we can get the first and last paths' names, and also in polynomial time we can, given paths $q$ and $r$, determine whether the path immediately greater than (i.e., right-adjacent to) $q$ is $r$. And a function $f$ belongs to CLU\#P if the function gives the number of accepting paths of an (appropriately balanced) Turing machine whose accepting paths always form a cluster of this sort. Although the formal definition is a bit intimidating, we stress that it is merely rigorously capturing this intuitively simple notion.

Definition 2.1 A (total) function $f: \Sigma^{*} \rightarrow \mathbb{N}$ belongs to CLU\#P if ( $\exists$ polynomial $p$ ) $(\exists$ p-balanced NPTM $M)\left(\exists b, t \in \mathrm{FP}_{\mathrm{t}}\right)(\exists$ 3-argument, polynomial-time computable predicate $\prec)(\forall x)\left(\exists\right.$ bijection $h_{x}$ from $\Sigma^{p(|x|)}$ to $\left.\Sigma^{p(|x|)}\right)$ such that:

1. $|b(x)|=|t(x)|=p(|x|)$.
2. $h(b(x))=0^{p(|x|)} \wedge h(t(x))=1^{p(|x|)}$.
3. $\left(\forall y, z \in\{0,1\}^{p(|x|)}\right)\left[\prec(x, y, z) \Longleftrightarrow h_{x}(y) \prec_{\operatorname{lex}} h_{x}(z)\right]$.
4. All accepting paths are clustered with respect to $\prec(x, \cdot, \cdot)$. That is, if $f(x) \neq 0$ then $\left(\exists \ell, u \in\{0,1\}^{p(|x|)}\right)\left[a c c_{M}(x)=\left\{w \in\{0,1\}^{p(|x|)} \mid h_{x}(\ell) \leq_{\text {lex }} h_{x}(w) \leq_{\operatorname{lex}} h_{x}(u)\right\}\right]$.
5. $f(x)=\# \operatorname{acc}_{M}(x)$.

As mentioned in the introduction, even for two same-length strings $x$ and $y$, it is completely possible that $\prec(x, \cdot, \cdot)$ and $\prec(y, \cdot, \cdot)$ will differ dramatically. That is, CLU\#P focuses heavily on reordering the paths related to the given input, and just those paths, and indeed may do so in a way that can vary based on the input. (Though formally speaking the definition above requires $\prec$ to be defined on all input triples, it is easy to see from the above definition that on input $x$ all that matters is what $\prec(x, \cdot, \cdot)$ does when its second two arguments are distinct strings in $\{0,1\}^{p(|x|)}$. For all other inputs, we can typically just ignore $\prec$ 's output or view it as being false.)

We now turn to the definition of CL\#P [HHKW05]. That definition requires the entire universe of paths - over all inputs to a machine - to be embedded in a single, shared order. As noted earlier, this limits one in two ways: the obvious constraint that one must embed paths over different inputs into the same order (and so when inputs have the same length, their paths must be identically shuffled) and a more subtle side-constraint that even though all computation paths of a machine on a given input are of the same length, in this setting the adjacency test must work even between that length and other lengths, i.e., all of $\Sigma^{*}$ must be woven into a single, giant order with the right feasibility properties.

To support the definition of CL\#P, we briefly define some related notions (see [HHKW05], from which we take these definitions essentially word for word, for consistency), namely, "length-respecting total order $A$ " and " $A$-cluster." A binary relation $A \subseteq \Sigma^{*} \times \Sigma^{*}$ is a partial order if it is reflexive, antisymmetric (i.e., $\left(\forall x, y \in \Sigma^{*}\right)[x \neq y \Longrightarrow$ $((x, y) \notin A \vee(y, x) \notin A)])$, and transitive. A partial order $A$ is a total order if, for all $x, y \in \Sigma^{*},(x, y) \in A$ or $(y, x) \in A$. We write $x \prec_{A} y$ if $x<_{A} y$ and there is no $z$ such that $x<_{A} z<_{A} y$. If $x \prec_{A} y$, we say that $x$ is left-adjacent to $y$ or, equivalently, $y$ is right-adjacent to $x$. Let $M$ be NPTM that is $p$-balanced for some polynomial $p$. Let $y$ and $z$ encode computation paths of $M$ on $x$. By the above assumption that $M$ is balanced, $|y|=|z|$. Fix a total order $A$ on $\Sigma^{*}$. We say that $y \sim_{A, M, x} z$ if and only if (a) $y \prec_{A} z$ or $z \prec_{A} y$, and (b) $M$ on $x$ accepts on path $y$ if and only if $M$ on $x$ accepts on path $z$. Let $\equiv_{A, M, x}$ be the equivalence closure (i.e., the reflexive-symmetric-transitive closure) of $\sim_{A, M, x}$. Then the relation $\equiv_{A, M, x}$ is an equivalence relation and thus induces a partitioning of the computation tree of $M$ on $x$. An $A$-cluster is an equivalence class whose representatives are accepting paths. Additionally, we consider $\emptyset$ to be a valid $A$-cluster. An order $A$ on $\Sigma^{*}$ is said to be length-respecting if and only if, for all $x, y,|x|<|y|$ implies $x<_{A} y$.

Definition 2.2 ([HHKW05]) A function $f$ belongs to the class CL\#P if there exist a polynomial p, a p-balanced NPTM M, and a length-respecting total order A with efficient adjacency checks such that, for all $x$, the following conditions hold:

1. The set of all accepting paths of $M$ on $x$ is an $A$-cluster.
2. $f(x)=\# a c c_{M}(x)$.

We now define the classes CLU\# $\mathrm{P}_{\text {free }}$ and CLU\# $\mathrm{P}_{\text {circular }}$. Their definitions are similar to that of CLU\#P. However, CLU\# $\mathrm{P}_{\text {free }}$ removes the constraint that top- and bottom-finding must be polynomial-time computable, though $\prec$ will implicitly create top and bottom elements. $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ makes the order be a circular order, thus removing any notion of "top" and "bottom."

Definition 2.3 A (total) function $f: \Sigma^{*} \rightarrow \mathbb{N}$ belongs to CLU\# $\mathrm{P}_{\text {free }}$ if ( $\exists$ polynomial $p$ ) $(\exists$ p-balanced NPTM M) $(\exists$ 3-argument, polynomial-time computable predicate $\prec)(\forall x)$
$\left(\exists\right.$ bijection $h_{x}$ from $\Sigma^{p(|x|)}$ to $\left.\Sigma^{p(|x|)}\right)$ such that:

1. $\left(\forall y, z \in\{0,1\}^{p(|x|)}\right)\left[\prec(x, y, z) \Longleftrightarrow h_{x}(y) \prec_{\operatorname{lex}} h_{x}(z)\right]$.
2. All accepting paths are clustered with respect to $\prec(x, \cdot, \cdot)$. That is, if $f(x) \neq 0$ then $\left(\exists \ell, u \in\{0,1\}^{p(|x|)}\right)\left[\operatorname{acc}_{M}(x)=\left\{w \in\{0,1\}^{p(|x|)} \mid h_{x}(\ell) \leq_{\text {lex }} h_{x}(w) \leq_{\text {lex }} h_{x}(u)\right\}\right]$.
3. $f(x)=\# \operatorname{acc}_{M}(x)$.

Definition 2.4 (total) function $f: \Sigma^{*} \rightarrow \mathbb{N}$ belongs to $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ if $(\exists$ polynomial $p$ ) $(\exists$ p-balanced NPTM $M)(\exists$ 3-argument, polynomial-time computable predicate $\prec)(\forall x)$ $\left(\exists\right.$ bijection $h_{x}$ from $\Sigma^{p(|x|)}$ to $\left.\Sigma^{p(|x|)}\right)$ such that:

1. $\left(\forall y, z \in\{0,1\}^{p(|x|)}\right)$

$$
\left[\prec(x, y, z) \Longleftrightarrow\left(h_{x}(y) \prec_{\operatorname{lex}} h_{x}(z) \vee\left(h_{x}(y)=1^{p(|x|)} \wedge h_{x}(z)=0^{p(|x|)}\right)\right)\right] .
$$

2. All accepting paths are clustered with respect to $\prec(x, \cdot, \cdot)$. That is, if $f(x) \neq 0$ then $\left(\exists \ell, u \in\{0,1\}^{p(|x|)}\right)\left[\operatorname{acc}_{M}(x)=\left\{w \in\{0,1\}^{p(|x|)} \mid h_{x}(\ell) \leq_{\operatorname{lex}} h_{x}(w) \leq_{\operatorname{lex}} h_{x}(u)\right\}\right]$.
3. $f(x)=\# a c c_{M}(x)$.

The reader may reasonably worry that our definition of $C L U \# \mathrm{P}_{\text {circular }}$ is cheating. In particular, one may worry that Definition 2.4's part 1 has the adjacency definition go "around the corner" (that is, it adjacency-links $0^{p(|x|)}$ and $1^{p(|x|)}$ in the under-the-image-of- $h$ space), but that Definition 2.4's part 2 doesn't similarly allow the accepting paths to go "around the corner," and that this is a somewhat strange and striking asymmetry of approach between those two aspects of the definition. However, note that in the definition of a $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ function we can without loss of generality require that the preimage of the bijection $h_{x}$ has the property that $h_{x}^{-1}\left(0^{p(|x|)}\right)$ is an accepting path of the machine (on that input) if any accepting paths exist (on that input). That is, the first condition in Definition 2.4 is invariant under cyclic shifts of the numbering $h_{x}$ of the elements in $\{0,1\}^{p(|x|)}$. So the above-mentioned worry about the definition turns out, upon some thought, not to be a worry at all. Indeed, later in the paper this observation will be a useful feature, namely, in the proof of Proposition 5.6.

Note that it follows immediately from the definitions that CLU\#P $\subseteq C L U \# \mathrm{P}_{\text {free }}$ and $\mathrm{CLU} \# \mathrm{P} \subseteq \mathrm{CLU} \# \mathrm{P}_{\text {circular }}$.

Finally, let us state the definitions of the function classes $\mathrm{UPSV}_{\mathrm{t}}$ and $\mathrm{UPSV}_{\mathrm{p}}$ [GS88, Kos99], which are the (respectively total and partial) unambiguous versions of the central, single-valued nondeterministic function classes $\mathrm{NPSV}_{\mathrm{t}}$ and $\mathrm{NPSV}_{\mathrm{p}}$ [BLS84, BLS85, Sel94]. When speaking of nondeterministic machines as computing (possibly partial) functions from $\Sigma^{*}$ to $\Sigma^{*}$, we view each path as having no output if the path is a rejecting path, and if a path is an accepting path then it is viewed as outputting whatever string $s \in \Sigma^{*}$ is on
the output tape (along that path) when that path halts. A (potentially partial) function $f: \Sigma^{*} \rightarrow \Sigma^{*}$ belongs to $\operatorname{UPSV}_{\mathrm{p}}$ if there is an NPTM $M$ that (a) on each input has at most one accepting path, (b) on each input $x$ on which $M$ has exactly one accepting path, $f(x)$ is the output on that path, and (c) on each input $x$ on which $M$ has no accepting paths, $f(x)$ is undefined (i.e., $\operatorname{domain}(f)=\{x \mid M(x)$ has no accepting paths $\}$ ). A function $f: \Sigma^{*} \rightarrow \Sigma^{*}$ belongs to UPSV $_{\mathrm{t}}$ if $f$ belongs to $\mathrm{UPSV}_{\mathrm{p}}$ and $f$ is total. $\mathrm{UPSV}_{\mathrm{p}}$ and $\mathrm{UPSV}_{\mathrm{t}}$ functions capture the flavor of "unique discovery," and will (often implicitly and sometimes explicitly) be central in our proofs.

## 3 Robustness of CLU\#P

In this section, we study the robustness of CLU\#P. CLU\#P on its surface might seem to be far more flexible than CL\#P, given that unlike CL\#P it is not chained by the requirement of a global order and the related need to have same-length strings' paths coexist in the same order and to link consistently between lengths. ${ }^{2}$ Nonetheless, we now prove that these two classes are equal: $\mathrm{CL} \# \mathrm{P}=\mathrm{CLU} \# \mathrm{P}$.

Briefly put, to show that CLU\#P $\subseteq \mathrm{CL} \# \mathrm{P}$ we tie together the exponential number of orderings (over all inputs sharing the same path length). To show that CL\#P $\subseteq C L U \# P$, we uniquely discover the top and the bottom elements and then embed into a broader search space a clone of the action of our CL\#P machine on the current input.

Theorem 3.1 CLU\#P = CL\#P.
We first prove a simple lemma. We do so in part because it will be helpful in the proof (though one could work around it if needed), and mostly because the proof provides a simple initial example of how to prove things about cluster classes.

We say a polynomial $p$ is monotonic exactly if, for all natural numbers $n, p(n)<p(n+1)$.
Lemma 3.2 If $f \in \mathrm{CLU} \mathrm{\# P}$, then $f \in \mathrm{CLU} \mathrm{\# P}$ via some integer-coefficient nondeterminism polynomial (in the sense of Definition 2.1) that is monotonic.

Proof Let $f \in \mathrm{CLU} \# \mathrm{P}$. Let $p, M, \prec, t$, and $b$ capture $f$ in the sense of Definition 2.1. If $p$ is monotonic, then we are done. Note that asymptotically the polynomial $p$ cannot

[^2]

Figure 1: Figure for the proof of Lemma 3.2. Key: R denotes paths that certainly are rejecting paths.
diverge to negative infinity and can never take on a negative value. So there will exist a monotonic, integer-coefficient polynomial $p^{\prime}$ such that $(\forall n \in \mathbb{N})\left[0 \leq p(n)<p^{\prime}(n)\right] . f$ will be computed in the sense of Definition 2.1 by $p^{\prime}$-balanced NPTM $M^{\prime}, b^{\prime}$, $t^{\prime}$, and $\prec^{\prime}$, each defined as follows.
$b^{\prime}(x)=0^{p^{\prime}(|x|)-p(|x|)} b(x)$ and $t^{\prime}(x)=1^{p^{\prime}(|x|)}$. NPTM $M^{\prime}$ works as follows: On input $x, M^{\prime}$ guesses $p^{\prime}(|x|)-p(|x|)$ bits, call them $\alpha$, and then guesses $p(|x|)$ bits, call them $\beta$. If $\alpha \in 0^{*}$ and $M(x)$ accepts along computation path $\beta$ (recall that we speak of paths by naming their nondeterministic guess bits), then we accept, and otherwise we reject. Define the predicate $\prec^{\prime}$ as follows: $\prec^{\prime}(x, y, z)$ will evaluate to true exactly if $|y|=|z|=p^{\prime}(|x|)$ and either

1. prefix $\left(y, p^{\prime}(|x|)-p(|x|)\right)=\operatorname{prefix}\left(z, p^{\prime}(|x|)-p(|x|)\right)=0^{p^{\prime}(|x|)-p(|x|)}$ and $\prec(x$, suffix $(y, p(|x|))$, suffix $(z, p(|x|))$, or
2. $\operatorname{prefix}\left(y, p^{\prime}(|x|)-p(|x|)\right) \neq 0^{p^{\prime}(|x|)-p(|x|)}$ and $y \prec_{\operatorname{lex}} z$, or
3. $y=0^{p^{\prime}(|x|)-p(|x|)} t(x)$ and $z=0^{p^{\prime}(|x|)-p(|x|)-1} 10^{p(|x|)}$.

That is, we guess dummy bits, simulate the underlying machine $M$ in the leftmost subtree, and weave all the paths naturally together by inheriting the adjacency operator for the leftmost subtree, and for the rest we follow lexicographical order (and are careful at the boundary about the connection between the leftmost subtree and the rest).

Figure 1a shows this proof pictorially, with paths shown (left to right) in lexicographical order. Figure 1b pictures the same construction, but in the way we will use from now on in all our figures. Namely, Figure 1b shows paths, left to right, not in lexicographical order, but rather "pre-permuted" into our order. (So in part b of the figure, the "inherited order" is a guide to how the figure has been pre-permuted. But that just is an issue of our illustration. Far more critical is to keep in mind that what really is inherited in that segment is the order itself-which paths should be considered adjacent to which. Also, to be clear, the curved arrow on the bottom of Figure 1a is denoting a single adjacency - that the path $0^{p^{\prime}(|x|)-p(|x|)} t(x)$ is left-adjacent to the path $0^{p^{\prime}(|x|)-p(|x|)-1} 10^{p(|x|)}$. But the curved arrow on the bottom of Figure 1b indicates that in the pre-permuted picture shown the adjacencies sweep one at a time from $0^{p^{\prime}(|x|)-p(|x|)} b(x)$ up to $0^{p^{\prime}(|x|)-p(|x|)} t(x)$ in the order inherited from the underlying order we are building upon.)
Proof of Theorem 3.1 For CLU\#P $\subseteq$ CL\#P, let $f \in \operatorname{CLU} \# \mathrm{P}$. Let $p, M, \prec, t$, and $b$ capture $f$ in the sense of Definition 2.1. By Lemma 3.2, we may assume without loss of generality that $p$ is monotonic. We show that $f \in \mathrm{CL} \# \mathrm{P}$ by constructing $p^{\prime}, M^{\prime}$, and length-respecting total order $A^{\prime}$ that capture $f$ in the sense of Definition 2.2. Define $p^{\prime}$ on input $n \in \mathbb{N}$ as $p(n)+n$.
$M^{\prime}$ works as follows: $M^{\prime}$ on input $x \in \Sigma^{*}$ nondeterministically guesses $|x|$ bits $y$. If $y=x$ then $M^{\prime}$ simulates $M$ on input $x$, where each nondeterministic branch accepts iff the corresponding branch in $M$ accepts. If $y \neq x$ then $M^{\prime}$ nondeterministically guesses $p(|x|)$ bits but then ignores them and rejects. Clearly, each nondeterministic branch of $M^{\prime}$ uses exactly $p^{\prime}(|x|)$ guess bits.


Figure 2: Figure for the CLU\#P $\subseteq \mathrm{CL} \# \mathrm{P}$ part of the proof of Theorem 3.1. Key: R denotes paths that certainly are rejecting paths and $x^{\prime}$ (respectively, $x^{\prime \prime}$ ) is the string lexicographically preceding (respectively, succeeding) $x$.
$A^{\prime}$ is the same as the lexicographical order, except for those strings that, for some $n \in \mathbb{N}$, are of length $p^{\prime}(n)$. For those strings, $0^{n} b\left(0^{n}\right)$ comes first and $1^{n} t\left(1^{n}\right)$ comes last. For all $x, y \in \Sigma^{n+p(n)}, x \prec_{A} y$ iff

1. $\operatorname{prefix}(x, n)=\operatorname{prefix}(y, n)$ and $\prec(\operatorname{prefix}(x, n)$, suffix $(x, p(n))$, suffix(y,p(n))), or
2. $\operatorname{prefix}(x, n) \prec_{\operatorname{lex}} \operatorname{prefix}(y, n)$ and $\operatorname{suffix}(x, p(n))=t(\operatorname{prefix}(x, n))$ and $\operatorname{suffix}(y, p(n))=$ $b($ prefix $(y, n))$.

Figure 2 shows pictorially this part of the proof.
For $\mathrm{CL} \# \mathrm{P} \subseteq \mathrm{CLU} \# \mathrm{P}$, let $f \in \mathrm{CL} \# \mathrm{P}$. Let $p, M$, and $A$ capture $f$ in the sense of Definition 2.2.

Let $U$ be an NPTM that works as follows: On input $x \in \Sigma^{*}, U$ nondeterministically guesses strings $u, y, z \in \Sigma^{*}$, where $|y|=|z|=p(|x|)$ and $|u|=2 p(|x|)$. If $p(|x|)=0$ then accept. If $p(|x|)>0$, $\operatorname{prefix}(u, p(|x|)-1) \prec_{A} y$, and $z \prec_{A} \operatorname{suffix}(u, p(|x|)+1)$ then accept. Otherwise reject. Clearly, $U$ on any input has exactly one accepting path (recall that $A$ is length-respecting), is balanced for polynomial $4 p$, and we can in polynomial time determine, for any $x \in \Sigma^{*}$ and $y \in \operatorname{acc}_{U}(x)$, what the least and greatest (with respect to $A$ ) strings of length $p(|x|)$ are.

We show $f \in \mathrm{CLU} \mathrm{\# P}$ by providing $p^{\prime}, M^{\prime}, \prec^{\prime}, t^{\prime}$, and $b^{\prime}$, defined below, that capture $f$ in the sense of Definition 2.1. Let $p^{\prime}=5 p$. For any $x \in \Sigma^{*}$, let $b^{\prime}(x)=0^{p^{\prime}(|x|)}$ and $t^{\prime}(x)=1^{p^{\prime}(|x|)}$.
$M^{\prime}$ works as follows: On input $x \in \Sigma^{*}, M^{\prime}$ nondeterministically guesses $w, y \in \Sigma^{*}$, where


Figure 3: Figure for CL\#P $\subseteq$ CLU\#P part of the proof of Theorem 3.1. Key: R denotes paths that certainly are rejecting paths.
$|w|=4 p(|x|)$ and $|y|=p(|x|)$. If $w \in \operatorname{acc}_{U}(x)$ and $y \in a c c_{M}(x)$, then accept. Otherwise, reject.

Predicate $\prec^{\prime}$ is defined as follows. For $x, u, u^{\prime} \in \Sigma^{*}, \prec^{\prime}\left(x, u, u^{\prime}\right)$ evaluates to true exactly if, for some $v, w, y, z \in \Sigma^{*}$, where $|v|=|w|=4 p(|x|)$ and $|y|=|z|=p(|x|), u=v y, u^{\prime}=w z$, it holds that:

1. $\{v, w\} \cap \operatorname{acc}_{U}(x)=\emptyset$ and $v y \prec_{\text {lex }} w z$, or
2. $\{v, w\} \subseteq \operatorname{acc}_{U}(x)(\Rightarrow v=w)$ and $y \prec_{A} z$, or
3. $v \prec_{\text {lex }} w$ and $w \in \operatorname{acc}_{U}(x)$ and $y=1^{p(|x|)}$ and $z$ is the minimum (with respect to $A$ ) of the set of all computational paths of $M$ on $x$, or
4. $v \prec_{\text {lex }} w$ and $v \in \operatorname{acc}_{U}(x)$ and $z=0^{p(|x|)}$ and $y$ is the maximum (with respect to $A$ ) of the set of all computational paths of $M$ on $x$.

Figure 3 shows pictorially this part of the proof.
We now derive a robustness result that might seem a bit less natural than Theorem 3.1. However, this robustness result provides a critical tool for proving natural results and gives substantial insight into what suffices to make cluster computation simple.

To state the result, we must define notions of the greatest element and the least element of an accepting path cluster. The slight unnaturalness occurs in the circular model, in particular in the case when all paths are accepting paths since in that case, even though there is no natural choice of greatest and least accepting paths, our definition makes a choice.

Note that we are speaking not about top and bottom notions among all paths of a given length, but rather are seeking the greatest and least accepting paths with respect to a given input and the ordering implicit in $\prec$.

Let $p, M$, and $\prec$ be a nondeterminism polynomial, machine, and adjacency predicate in either the $\mathrm{CLU} \# \mathrm{P}_{\text {free }}$ model or the CLU\# $\mathrm{P}_{\text {circular }}$ model. We define two partial functions ( $p$ is implicit in $M$, but for uniformity and clarity in settings like this we include $p$ throughout the paper) greatest $_{p, M, \prec}$ and least $_{p, M, \prec}$ as follows. Let $M$ compute the function $f$, i.e., on input $x, f(x)=\# a c c_{M}(x)$.

If $f(x)=0$, then (in both the free and the circular models) greatest $_{p, M, \prec}(x)$ and least $_{p, M, \prec}(x)$ are undefined. In the free model, if $f(x) \neq 0$, then greatest $_{p, M, \prec}(x)$ is the unique length $p(|x|)$ string $z$ that is an accepting path to which no length $p(|x|)$ accepting path is right-adjacent (i.e., the unique string $z$ of length $p(|x|)$ such that $z$ is an accepting path of $M$ on input $x$ and yet $\left.\left(\forall w \in\{0,1\}^{p(|x|)}\right)\left[w \in \operatorname{acc}_{M}(x) \Longrightarrow \neg \prec(x, z, w)\right]\right)$. Similarly, in the free model, if $f(x) \neq 0$, then least $t_{p, M, \prec}(x)$ is the unique length $p(|x|)$ string $z$ that is an accepting path to which no length $p(|x|)$ accepting path is left-adjacent.

The above definitions will not work in the circular model if $M$, on input $x$, accepts on all paths, as there the definition would give nothing, but for our proofs we cannot allow that to happen. It actually is fine to break this impasse by saying that when that happens, the greatest function takes on the value of any of $M$ 's accepting paths such that the least function has as its value the path right-adjacent to that path. However, for clarity and specificity, we sacrifice a bit of flexibility and choose one particular impassebreaking splitting point as follows for the circular model. If $f(x)=0$, greatest $_{p, M, \prec}(x)$ and least ${ }_{p, M, \prec}(x)$ still are undefined. If $f(x)=2^{p(|x|)}$ then greatest $_{p, M, \prec}(x)=1^{p(|x|)}$. If $f(x)=2^{p(|x|)}$ and $p(|x|)=0$ then least $_{p, M, \prec}(x)=\epsilon$. If $f(x)=2^{p(|x|)}$ and $p(|x|) \neq 0$ then least $t_{p, M, \prec}(x)$ equals the unique length $p(|x|)$ string $z$ satisfying $\prec\left(x, 1^{p(|x|)}, z\right)$. In the circular model, if $0 \leq f(x)<2^{p(|x|)}$, greatest and least are defined exactly as in the free model.

Theorem 3.3 1. Let $f$ be computed by $p, M$, $\prec$ in the free model. If greatest ${ }_{p, M, \prec} \in$ $\mathrm{UPSV}_{\mathrm{p}}$ and least ${ }_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$, then $f \in \mathrm{CLU} \# \mathrm{P}$.
2. Let $f$ be computed by $p, M$, $\prec$ in the circular model. If greatest $t_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$ and least $_{p, M, \prec} \in \operatorname{UPSV}_{\mathrm{p}}$, then $f \in \mathrm{CLU} \# \mathrm{P}$.

That is, unique discovery of boundaries is sufficient in both the free and the circular models to remove any power beyond that of CLU\#P.

Briefly summarized, our proof will seek to uniquely discover the greatest and least accepting paths, and on the (at most one) block that discovers them will simulate the original machine except with each path sheathed in three dummy rejecting paths. Blocks that fail to make the unique discovery will follow lexicographical order, and the unique


Figure 4: First figure for the proof of Theorem 3.3. Key: R denotes paths that certainly are rejecting paths. The string $y$ is the accepting path of $U(x)$ concatenated with the accepting path of $U^{\prime}(x)$, if such paths exist. The string $y^{\prime}$ (respectively, $y^{\prime \prime}$ ) is the string lexicographically preceding (respectively, succeeding) $y$.
discovery block (which will exist exactly when $f(x)>0$ ) will adopt a somewhat complex order that allows us to indeed be CLU\#P-like.
Proof of Theorem 3.3 The proofs for items 1 and 2 are essentially the same. Let $f$ be computed by $p, M$, and $\prec$ in either the free or the circular model. Suppose for polynomials $q$ and $q^{\prime}$ that greatest $_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$ via $q$-balanced NPTM $U$ and least $_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$ via $q^{\prime}$-balanced NPTM $U^{\prime}$. We now define $p^{\prime}, M^{\prime}, \prec^{\prime}, b^{\prime}$, and $t^{\prime}$ that capture $f$ in the sense of Definition 2.1. Let $p^{\prime}(n)=q(n)+q^{\prime}(n)+p(n)+2$.
$M^{\prime}$ works as follows. On input $x \in \Sigma^{*}, M^{\prime}$ nondeterministically guesses strings $y, z, u, v \in \Sigma^{*}$, where $|y|=q(|x|),|z|=q^{\prime}(|x|),|u|=p(|x|)$, and $|v|=2$. If $y \in \operatorname{acc}_{U}(x)$, $z \in \operatorname{acc}_{U^{\prime}}(x), u \in \operatorname{acc}_{M}(x)$, and $v=01$ then accept. Otherwise reject.

We define $b^{\prime}$ on input $x \in \Sigma^{*}$ as

$$
b^{\prime}(x)= \begin{cases}0^{q(|x|)+q^{\prime}(|x|)} l e a s t_{p, M, \prec}(x) 01 & \text { if } 0^{q(|x|)} \in a c c_{U}(x) \text { and } 0^{q^{\prime}(|x|)} \in \operatorname{acc}_{U^{\prime}}(x) \\ 0^{p^{\prime}(|x|)} & \text { otherwise. }\end{cases}
$$

On input $x \in \Sigma^{*}, t^{\prime}$ is defined as $1^{p^{\prime}(|x|)}$.
We define $\prec^{\prime}$, on inputs $x, w, w^{\prime} \in \Sigma^{*}$ as follows. Predicate $\prec^{\prime}\left(x, w, w^{\prime}\right)$ evaluates to true exactly if there exist strings $y, y^{\prime}, u, u^{\prime}, v, v^{\prime} \in \Sigma^{*}$ such that $|y|=\left|y^{\prime}\right|=q(|x|)+q^{\prime}(|x|)$, $|u|=\left|u^{\prime}\right|=p(|x|),|v|=\left|v^{\prime}\right|=2, w=y u v, w^{\prime}=y^{\prime} u^{\prime} v^{\prime}$, and it holds that:

1. $\left\{w, w^{\prime}\right\} \cap a c c_{M^{\prime}}(x)=\emptyset$ and $w \prec_{\operatorname{lex}} w^{\prime}$, or


Figure 5: Second figure for the proof of Theorem 3.3. Key: R denotes paths that certainly are rejecting paths and A denotes accepting paths. String $y$ is as in Figure 4.
2. $\left\{w, w^{\prime}\right\} \subseteq \operatorname{acc}_{M^{\prime}}(x)$ and $\prec\left(x, u, u^{\prime}\right)$, or
3. $w \in \operatorname{acc}_{M^{\prime}}(x)$ and $u=$ greatest $_{p, M, \prec}(x)$ and prefix $\left(w, q(|x|)+q^{\prime}(|x|)\right)=$ $\operatorname{prefix}\left(w^{\prime}, q(|x|)+q^{\prime}(|x|)\right)$ and $\operatorname{suffix}\left(w^{\prime}, p(n)+2\right) \in 0^{*}$, or
4. $\left.\left(\exists w^{\prime \prime} \in \operatorname{acc}_{M^{\prime}}(x)\right)\left[w \prec_{\operatorname{lex}} w^{\prime \prime} \prec_{\operatorname{lex}} w^{\prime}\right)\right]$, or
5. $w^{\prime} \in \operatorname{acc}_{M^{\prime}}(x)$ and $u^{\prime} \in$ least $_{p, M, \prec}(x)$ and $y \prec_{\text {lex }} y^{\prime}$ and $\operatorname{suffix}(w, p(n)+2) \in 1^{p(|x|)+2}$.

Figures 4 and 5 show how this construction works.
We will employ Theorem 3.3 in Section 5, but let us note now that focusing on the boundaries of the accepting block is enough to speak to issues regarding the complexity of the top and bottom functions.

Corollary 3.4 If in Definition 2.1 " $\exists b, t \in \mathrm{FP}_{\mathrm{t}}$ " is replaced with " $\exists b, t \in \operatorname{UPSV}_{\mathrm{t}}$," the class defined by the new definition remains precisely CLU\#P.

Proof Let $p, M, \prec, b$, and $t$ satisfy the definitions of CLU\#P except altered as noted in the statement of this corollary.

Notice that on a given input $x$, a computation path $\rho$ of length $p(|x|)$ is the value of least $_{p, M, \prec}$ if $\rho$ is an accepting path and either (a) $x$ is in the domain of (the $\mathrm{UPSV}_{\mathrm{t}}$ function) $b$ and $b(x)=\rho$, or (b) there is a rejecting computation path that is left-adjacent to $\rho$ (and
note that if such a path exists it is unique). Thus, keeping in mind that if there are no accepting paths at length $p(|x|)$ on input length the test just described will not select any
 accepting, and then doing the above test and outputting the path if the test succeeds, we have shown that least ${ }_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$. By a similar argument, greatest $t_{p, M, \prec} \in \mathrm{UPSV}_{\mathrm{p}}$. So by Theorem 3.3 the function computed by $p, M$, and $\prec$ is in CLU\#P.

## 4 Closure Properties of CLU\#P

Arithmetic closure properties are not the focus of this paper. However, in this section we briefly study some as an example of the power of unique discovery of boundaries and to take advantage of the fact that Theorem 3.1 allows us to prove closure properties of CL\#P via the easier to work with model of CLU\#P. In particular, we show that an implication of [HHKW05] is in fact a complete characterization.
Theorem 4.1 ([HHKW05]) If CL\#P (equivalently in light of Theorem 3.1, CLU\#P) is closed under increment (i.e., $f \in \mathrm{CL} \# \mathrm{P} \Longrightarrow(\lambda x . f(x)+1) \in \mathrm{CL} \# \mathrm{P}$ ), then $\mathrm{UP}=\mathrm{coUP}$.

We prove that the converse holds and in fact prove that $\mathrm{UP}=$ coUP characterizes a number of closures of CLU\#P. We say a function is natural-number-valued if it maps from $\Sigma^{*}$ to $\mathbb{N}$. All CLU\#P functions are natural-number-valued.

Theorem 4.2 The following statements are equivalent:

1. $\mathrm{UP}=\mathrm{coUP}$.
2. $\mathrm{CLU} \# \mathrm{P}$ is closed under increment.
3. $\mathrm{CLU} \# \mathrm{P}$ is closed under addition of natural-number-valued $\mathrm{FP}_{\mathrm{t}}$ functions.
4. CLU\#P is closed under addition of natural-number-valued $\mathrm{UPSV}_{\mathrm{t}}$ functions.
5. CLU\#P is closed under addition.

Proof Clearly, all natural-number-valued UPSV $_{\mathrm{t}}$ functions are CLU\#P functions, so $5 \Rightarrow$ 4. All $\mathrm{FP}_{\mathrm{t}}$ functions are $\mathrm{UPSV}_{\mathrm{t}}$ so $4 \Rightarrow 3$. Clearly, $3 \Rightarrow 2$. By Theorem $4.1,2 \Rightarrow 1$.

To prove $1 \Rightarrow 5$, suppose UP $=\operatorname{coUP}$. Let $f \in \mathrm{CLU} \mathrm{\# P}$ via $M_{f}, p_{f}, b_{f}, t_{f}$ and $\prec_{f}$ and let $g \in \mathrm{CLU} \mathrm{\# P}$ via $M_{g}, p_{g}, b_{g}, t_{g}$ and $\prec_{g}$. Let $i \in\{f, g\}$. Let $V_{i}$ be an NPTM that, on input $x \in \Sigma^{*}$, nondeterministically guesses $y, z, u, v \in \Sigma^{*}$, where $|y|=|z|=|u|=|v|=p_{i}(|x|)$, and accepts (on the current path) if and only if $\{z, u\} \subseteq \operatorname{acc}_{M_{i}}(x)$ and both of the following hold:

1. $\left(y \notin \operatorname{acc}_{M_{i}}(x) \wedge \prec_{i}(x, y, z)\right)$ or $\left(z=b_{i}(x) \wedge y \in 0^{*}\right)$, and
2. $\left(v \notin \operatorname{acc}_{M_{i}}(x) \wedge \prec_{i}(x, u, v)\right)$ or $\left(u=t_{i}(x) \wedge v \in 0^{*}\right)$.

Clearly, $V_{i}$ has on any input at most one accepting path, and if $V_{i}$ accepts on input $x$ then the string $z$ (respectively, $u$ ) guessed by the accepting path is least $p_{i, M_{i}, \prec_{i}}(x)$ (respectively, greatest $\left.p_{p_{i}, M_{i}, \prec_{i}}(x)\right)$. Assuming UP $=\operatorname{coUP}, \overline{L\left(V_{i}\right)} \in \mathrm{UP}$. Suppose for some polynomial $q_{i}^{\prime}$ that $V_{i}^{\prime}$ is a $q_{i}^{\prime}$-balanced NPTM that decides $\overline{L\left(V_{i}\right)}$ and has on any input at most one accepting path. Let $U_{i}$ be a Turing machine that, on input $x \in \Sigma^{*}$, nondeterministically guesses a string $z$ of length $q_{i}(|x|)$, where $q_{i}$ is a polynomial such that $(\forall n \in \mathbb{N})\left[q_{i}(n)>\right.$ $\left.4 p_{i}(n)+q_{i}^{\prime}(n)\right]$, and accepts (on the current path) if and only if $\left(\exists y \in \Sigma^{*}\right)\left[\left(y \in a c c_{V_{i}}(x) \vee y \in\right.\right.$ $\left.\left.\left.\operatorname{acc}_{V_{i}^{\prime}}(x)\right) \wedge z \in y 0^{*}\right)\right]$. Note that on any input $U_{i}$ has at most one accepting path and that we can in polynomial time determine, for any $x \in \Sigma^{*}$ and $y \in \operatorname{acc}_{U_{i}}(x)$, whether $i(x)>0$ and, if so, what greatest $_{p_{i}, M_{i}, \prec_{i}}(x)$ and least $p_{i, M_{i}, 久_{i}}(x)$ are.

Let $U$ be an NPTM that, on input $x \in \Sigma^{*}$, guesses $y \in \Sigma^{*}$ of length $q_{f}(|x|)+q_{g}(|x|)$ and accepts (on the current path) exactly if $\left(\operatorname{prefix}\left(y, q_{f}(|x|)\right) \in \operatorname{acc}_{U_{f}}(x) \wedge \operatorname{suffix}\left(y, q_{g}(|x|)\right) \in\right.$ $\left.\operatorname{acc}_{U_{g}}(x)\right)$. Clearly, $U$ has on any input exactly one accepting path, and we can, in polynomial time, determine, for any $x \in \Sigma^{*}, y \in \operatorname{acc}_{U}(x)$, and $i \in\{f, g\}$ whether $i(x)>0$ and, if so, what least $_{p_{i}, M_{i}, \prec_{i}}(x)$ and greatest $_{p_{i}, M_{i}, 久_{i}}(x)$ are.

We now construct $M, p, b, t$, and $\prec$ that capture the function $f+g$ in the sense of Definition 2.1. Let $p(n)=q_{f}(n)+q_{g}(n)+q(n)+3$, where $q$ is a polynomial such that $(\forall n \in \mathbb{N})\left[q(n)>p_{f}(n)+p_{q}(n)\right]$.
$M$ on input $x \in \Sigma^{*}$ works as follows: $M$ nondeterministically guesses $y \in \Sigma^{*}$, where $|y|=p(|x|)$, and accepts (on the current path) exactly if $\operatorname{prefix}\left(y, q_{f}(|x|)+q_{g}(|x|)\right) \in a c c_{U}(x)$ and:

1. For some $z \in \operatorname{acc}_{M_{f}}(x)$, suffix $(y, q(|x|)+3) \in 0 z 0^{*} 1$, or
2. for some $z \in \operatorname{acc} M_{M_{g}}(x)$, suffix $(y, q(|x|)+3) \in 1 z 0^{*} 1$.

Clearly, $(\forall x)\left[\# a c c_{M}(x)=f(x)+g(x)\right]$. Note that our accepting paths always have at least two "extra" bits at the end of the path. This allows us to sheath, in a manner similar to what we did in the proof of Theorem 3.3, each accepting path in at least three dummy rejecting paths.

Before defining $\prec$, we will first define greatest $_{p, M, \prec}$ and least $_{p, M, \prec}$. For $x \in \Sigma^{*}$, let $u \in \operatorname{acc}_{U}(x)$ and let
greatest $_{p, M, \prec}(x)= \begin{cases}\text { u1 }^{\text {greatest }_{p_{g}, M_{g}, \prec_{g}}(x) 0^{q(|x|)-p_{g}(x)+1} 1} & \text { if } \text { greatest }_{p_{g}, M_{g}, \prec_{g}}(x) \text { is defined, } \\ \text { ugreatest }_{p_{f}, M_{f}, \prec_{f}}(x) 0^{q(|x|)-p_{f}(x)+1} 1 & \text { otherwise, if } \text { greatest }_{p_{f}, M_{f}, \prec_{f}}(x) \\ & \text { is defined, } \\ \text { undefined } & \text { otherwise. }\end{cases}$

$\xrightarrow{\text { [lexicographical order] }}$


Figure 6: First figure for the proof of Theorem 4.2. Key: R denotes paths that certainly are rejecting paths. The string $y$ is the accepting path of $U(x)$. The string $y^{\prime}$ (respectively, $y^{\prime \prime}$ ) is the string lexicographically preceding (respectively, succeeding) $y$.
and let

$$
\text { least }_{p, M, \prec}(x)= \begin{cases}\text { u0least } p_{f}, M_{f}, \prec_{f}(x) 0^{q(|x|)-p_{f}(x)+1} 1 & \text { if } \text { least }_{p_{f}, M_{f}, \prec_{f}}(x) \text { is defined, } \\ \text { ulleast } p_{g}, M_{g}, \prec_{g} \\ & (x) 0^{q(|x|)-p_{g}(x)+1} 1 \\ \text { otherwise, if least } t_{p_{g}, M_{g}, \prec_{g}}(x) \text { is } \\ \text { undefined } & \text { defined, } \\ \text { otherwise. }\end{cases}
$$

Define $\prec$ on $x, y, z \in \Sigma^{*}$ as follows (we claim that our definition will result in greatest $_{p, M, \prec}$ and least $t_{p, M, \prec}$ as defined above): $\prec(x, y, z)$ evaluates to true exactly if $|y|=|z|=p(|x|)$ and:

1. $\{y, z\} \cap a c c_{M}(x)=\emptyset$ and $y \prec_{\text {lex }} z$, or
2. $\{y, z\} \subseteq \operatorname{acc}_{M}(x)$ and, for $w \in \operatorname{acc}_{U}(x)$ :
(a) for $u, v \in a c c_{M_{f}}, y \in w 0 u 0^{*} 1$ and $z \in w 0 v 0^{*} 1$ and $\prec_{f}(x, u, v)$, or
(b) for $u, v \in a c c_{M_{g}}, y \in w 1 u 0^{*} 1$ and $z \in w 1 v 0^{*} 1$ and $\prec_{g}(x, u, v)$, or
(c) $f(x)>0$ and $g(x)>0$ and $y \in$ wogreatest $_{p_{f}, M_{f}, \prec_{f}}(x) 0^{*} 1$ and $z \in$ $w 1$ least $_{p_{g}, M_{g}, \prec_{g}}(x) 0^{*} 1$,
or
3. $\{y, z\} \cap \operatorname{acc}_{M}(x)=\emptyset$ and, for some $w \in \operatorname{acc}_{M}(x), y \prec_{\operatorname{lex}} w \prec_{\operatorname{lex}} z$, or
4. $z \in \operatorname{acc}_{M}(x)$ and $z=$ least $_{p, M, \prec}(x)$ and $\operatorname{prefix}\left(y, q_{f}(|x|)+q_{g}(|x|)\right) \prec_{\text {lex }} \operatorname{prefix}\left(z, q_{f}(|x|)+\right.$ $\left.q_{g}(|x|)\right)$ and $\operatorname{suffix}(y, q(|x|)+3) \in 1^{*}$, or


Figure 7: Second figure for the proof of Theorem 4.2. Key: R denotes paths that certainly are rejecting paths and A denotes accepting paths. String $y$ is as in Figure 6.
5. $y \in \operatorname{acc}_{M}(x)$ and $y=$ greatest $_{p, M, \prec}(x)$ and $\operatorname{prefix}\left(z, q_{f}(|x|)+q_{g}(|x|)\right) \in \operatorname{acc}_{U}(x)$ and $\operatorname{suffix}(z, q(|x|)+3) \in 0^{*}$.
On input $x \in \Sigma^{*}, b$ is defined as

$$
b(x)= \begin{cases}\text { least }_{p, M, \prec}(x) & \text { if } 0^{q_{f}(|x|)+q_{g}(|x|)} \in a c c_{U}(x) \\ & \text { and } \text { least }_{p, M, \prec}(x) \text { is defined }, \\ 0^{p(|x|)} & \text { otherwise },\end{cases}
$$

and $t(x)$ is defined as $1^{p(|x|)}$. Figures 6 and 7 show how this construction works.
Theorem 4.2 may be viewed as evidence that CLU\#P lacks various closure properties, e.g., closure under increment. In contrast, the following result provides a closure property, proper decrement, that CLU\#P possesses unconditionally.

Theorem 4.3 CLU\#P is closed under proper decrement (i.e., $f \in \mathrm{CLU} \# \mathrm{P} \Longrightarrow$ $(\lambda x$. $\max \{0, f(x)-1\}) \in \mathrm{CLU} \# \mathrm{P})$.

Proof Let $f \in \mathrm{CLU} \# \mathrm{P}$ via NPTM $M$, nondeterminism polynomial $p$, predicate $\prec$, and functions $b, t \in \mathrm{FP}_{\mathrm{t}}$. Define $U$ to be a machine that, on input $x \in \Sigma^{*}$, guesses a string $y \in \Sigma^{*}$ of length $2 p(|x|)$ and accepts (on the current path) if and only if $\operatorname{prefix}(y, p(|x|)) \in \operatorname{acc}_{M}(x)$ and

1. suffix $(y, p(|x|)) \notin \operatorname{acc}_{M}(x)$ and $\prec(x$, prefix $(y, p(|x|))$, suffix $(y, p(|x|)))$, or
2. $\operatorname{prefix}(y, p(|x|))=t(x)$ and suffix $(y, p(|x|))=0^{p(|x|)}$.

Clearly, $U$ can have at most one accepting path. (Notice that $U$ can be viewed as a balanced UPSV $_{p}$ machine that computes greatest $\left._{p, M, \prec}.\right)$ Now consider an NPTM $N$ that, on input $x \in \Sigma^{*}$, guesses a string $y$, where $|y|=3 p(|x|)$, and accepts (on the current path) if and only if $\operatorname{prefix}(y, 2 p(|x|)) \in \operatorname{acc}_{U}(x) \wedge \operatorname{suffix}(y, p(|x|)) \in \operatorname{acc}_{M}(x) \wedge \operatorname{prefix}(y, p(|x|)) \neq$ $\operatorname{suffix}(y, p(|x|))$. It is clear that $N$ would have as many accepting paths as $M$ has except regarding $N$ 's analog of the greatest accepting path of $M$ (if $M$ has accepting paths at all), which is rejected by $N$. Thus, for all $x \in \Sigma^{*}, \# a c c_{N}(x)=\max \left\{0, \# a c c_{M}(x)-1\right\}$. It remains to show how to define an appropriate predicate $\prec^{\prime}$ : For $x, y, z \in \Sigma^{*}, \prec^{\prime}(x, y, z)$ evaluates to true if and only if $|y|=|z|=3 p(|x|)$ and

1. $\operatorname{prefix}(y, 2 p(|x|))=\operatorname{prefix}(z, 2 p(|x|))$ and $\prec(x, \operatorname{suffix}(y, p(|x|))$, suffix $(z, p(|x|)))$, or
2. $\operatorname{prefix}(y, 2 p(|x|)) \prec_{\operatorname{lex}} \operatorname{prefix}(z, 2 p(|x|))$ and $\operatorname{suffix}(y, p(|x|))=t(x)$ and $\operatorname{suffix}(z, p(|x|))=$ $b(x)$.

Furthermore, for all $x \in \Sigma^{*}$, define $b^{\prime}(x)=0^{2 p(|x|)} b(x)$ and $t^{\prime}(x)=1^{2 p(|x|)} t(x)$. So NPTM $N$, nondeterminism polynomial $3 p$, predicate $\prec^{\prime}$, and $\mathrm{FP}_{\mathrm{t}}$ functions $b^{\prime}, t^{\prime}$ witness the fact that $\# a c c_{N} \in \mathrm{CLU} \# \mathrm{P}$.

## 5 Free Cluster and Circular Cluster Computation

We defined CLU\# $\mathrm{P}_{\text {free }}$ and CLU\# $\mathrm{P}_{\text {circular }}$ in Section 2. Are these seemingly more flexible models truly more powerful than CLU\#P? We have not been able to prove that, though we will later, as Theorems 5.2 and 5.5 , show that unless they are more powerful, certain collapses and closures hold. On the other hand, we now prove that UP $=$ coUP is sufficient to reduce the power of these two seemingly more flexible classes to that of CLU\#P.

Theorem 5.1 UP $=\operatorname{coUP} \Longrightarrow \mathrm{CLU} \# \mathrm{P}=\mathrm{CLU} \# \mathrm{P}_{\text {free }}=\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$.
Proof Let us first consider the CLU\# $\mathrm{P}_{\text {free }}$ model. Suppose we are given some function $f$ in $\mathrm{CLU} \# \mathrm{P}_{\text {free }}$, along with $p, M$, and $\prec$ modeling it. Note that (on input $x$ ) recognizing length $p(|x|)$ elements as being not the top element is a UP test. ${ }^{3}$ Namely, if the length of our paths is zero, then the path is the top element exactly if it is $\epsilon$ (i.e., it is the path containing no guesses). And otherwise our arbitrary, given length $p(|x|)$ path is not the top element exactly if there exists a length $p(|x|)$ element that is right-adjacent to it with respect to $\prec(x, \cdot, \cdot)$. So being the top element is a coUP test (again, relative to $x$; see Footnote 3). And note that for each $x$ there is only one top element at length $p(|x|)$. So, if $\mathrm{UP}=\mathrm{coUP}$, we have a $\mathrm{UPSV}_{\mathrm{t}}$ function that finds the top element on input $x$ : Guess each

[^3]length $p(|x|)$ path, and then simulate the UP test for the coUP question of whether it is the top element. Analogously, if $\mathrm{UP}=\mathrm{coUP}$, we have a $\mathrm{UPSV}_{\mathrm{t}}$ function that finds the bottom element. By Corollary 3.4 we are done.

We now turn to the $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ model. Suppose we are given some function $f$ in CLU\# $\mathrm{P}_{\text {circular }}$, along with the $p, M, \prec$ modeling it. We ignore the case $p(|x|)=0$ as that is an uninteresting and easy to handle special case.

Notice that it is a UP test to determine (on input $x$ ) whether there exists any right edge to the accepting paths of length $p(|x|)$, i.e., whether there exists any accepting path such that its (unique) right-adjacent path is a rejecting path. So we now describe, under the assumption that $\mathrm{UP}=\mathrm{coUP}$, a $\mathrm{UPSV}_{\mathrm{p}}$ function computing greatest. The lack of a right edge is a coUP test, and UP $=$ coUP. And as just mentioned the existence of a right edge is a UP test.

So our $\mathrm{UPSV}_{\mathrm{t}}$-type machine guesses both that there is and that there is not a right edge, and under each such guess it runs the appropriate UP test. If there is a right edge, some single path finds that that is the case (and it is easy to ensure that that path in fact even has the name of the right edge, as that is why that test is in UP in the first place). Output the name of our right edge. If there is no right edge, again some single path discovers that that is the case. Since $p(|x|) \neq 0$ and the path has found that there is no right edge, and our order is circular, there are only two subcases: All length $p(|x|)$ paths reject or all length $p(|x|)$ paths accept. So on the current path our machine does the following to distinguish these two subcases. Our machine tests whether $0^{p(|x|)}$ is an accepting path. If $0^{p(|x|)}$ is not an accepting path, then there are no accepting paths at the current length, and so our machine rejects on the current path. If $0^{p(|x|)}$ is an accepting path, then there are only accepting paths at the current length, and so our machine in order to obey the definition of greatest outputs $1^{p(|x|)}$.

We have shown that in the circular case the partial function greatest is in $\mathrm{UPSV}_{\mathrm{p}}$ if $\mathrm{UP}=\mathrm{coUP}$. The argument that in the circular case the partial function least is in $\mathrm{UPSV}_{\mathrm{p}}$ if $\mathrm{UP}=\mathrm{coUP}$ is quite similar, except it focuses on the left edge of the accepting block. Also, it has one minor twist. In the final subcase $(p(|x|) \neq 0$, we have found that there is no left edge, and $0^{p(|x|)}$ is an accepting path), we must output the unique length $p(|x|)$ string that is right-adjacent to $1^{p(|x|)}$. However, doing this still within a $\mathrm{UPSV}_{\mathrm{p}}$ function is easy, since there indeed is (recall, $p(|x|) \neq 0$ ) exactly one such string and so we can guess and recognize it.

So, under the assumption UP $=$ coUP, in the circular model greatest and least are both UPSV $_{\mathrm{p}}$ functions, and so by Theorem $3.3 f \in \mathrm{CLU} \# \mathrm{P}$.

We now show that if cluster machines with free boundaries no more powerful than the regular or circular models then UP $=$ coUP.

Theorem 5.2 The following are equivalent.

1. $\mathrm{UP}=\mathrm{coUP}$.
2. $\mathrm{CLU}_{\mathrm{\#}}^{\text {free }}$ $=\mathrm{CLU} \# \mathrm{P}$.
3. $\mathrm{CLU} \# \mathrm{P}_{\text {free }} \subseteq \mathrm{CLU} \# \mathrm{P}_{\text {circular }}$.

The technique used to prove Theorem 5.2 (one that we will later use to prove Theorem 5.7) is, roughly speaking, to impose over an unambiguous or "nearly unambiguous" NPTM an order that skips over any accepting paths but then, at the end, sticks the skipped paths to the top of the order. Because the free model imposes no restrictions on the last element, this is relatively easy to do. But this makes accepting paths relatively easy to locate in the other, "non-free" models, and that will allow us to get coUP $\subseteq$ UP.
Proof of $\mathbf{5 . 2}$ By Theorem 5.1, $1 \Rightarrow 2$. Because CLU\#P $\subseteq$ CLU\# $\mathrm{P}_{\text {circular }}, 2 \Rightarrow 3$. For $3 \Rightarrow 1$, let $L \in$ coUP. For some polynomial $p$, let $U^{\prime}$ be a $p$-balanced NPTM such that $L\left(U^{\prime}\right)=\bar{L}$ and that, on any input, has at most one accepting path. We may assume without loss of generality that $p>0$ and that, on any $x \in \Sigma^{*}, U^{\prime}$ never accepts on path $1^{p(|x|)}$. Let $U$ be an NPTM that is the same as $U^{\prime}$ except that, on any $x \in \Sigma^{*}, U$ on input $x$ always accepts on path $1^{p(|x|)}$. Thus $x \in L \Rightarrow U$ on input $x$ has exactly one accepting path and $x \notin L \Rightarrow U$ on input $x$ has exactly two accepting paths. Define predicate $\prec$ as follows. For $x, y, z \in \Sigma^{*}, \prec(x, y, z)$ evaluates to true exactly if $|y|=|z|=p(|x|)$ and:

1. $\{y, z\} \cap \operatorname{acc}_{U^{\prime}}(x)=\emptyset$ and either
(a) $y \prec_{\text {lex }} z$, or
(b) $\left(\exists w \in \operatorname{acc}_{U^{\prime}}(x)\right)\left[y \prec_{\operatorname{lex}} w \prec_{\operatorname{lex}} z\right]$,

## or

2. $z \neq y=1^{p(|x|)}$ and $z \in \operatorname{acc}_{U^{\prime}}(x)$.

Thus, $p, U$ and $\prec$ capture $\# a c c_{U}$ in the sense of Definition 2.3. (Note, however, that in defining $\prec$, which sets the adjacencies in $U$, we very deliberately refer to accepting paths of $U^{\prime}$ rather than $U$, since this somewhat simplifies the definition of $\prec$.)

Suppose that $\# a c c_{U} \in \operatorname{CLU} \# \mathrm{P}_{\text {circular }}$ via $p^{\prime}, M^{\prime}$, and $\prec^{\prime}$. Then the NPTM $N$, described below, witnesses that $L \in \mathrm{UP}$. On input $x \in \Sigma^{*}, N$ nondeterministically guesses $y, z, w \in \Sigma^{*}$ where $|y|=|z|=|w|=p^{\prime}(|x|)$ and accepts (on the current path) exactly if $z \in \operatorname{acc}_{M^{\prime}}(x)$, $\{y, w\} \cap \operatorname{acc}_{M^{\prime}}(x)=\emptyset, \prec^{\prime}(x, y, z)$, and $\prec^{\prime}(x, z, w)$.

The most pressing open question posed by Theorems 5.1 and 5.2 and indeed by this paper is whether UP $=$ coUP is a necessary condition for CLU\#P $=C L U \# P_{\text {circular }}$.

Finally, we present three results that show that the free and circular classes are in some ways relatively close to CLU\#P.

Let $0-1-\mathrm{F}$ denote all $0-1$-valued total functions, i.e., total functions $f$ mapping from $\Sigma^{*}$ to $\{0,1\}$.

Theorem 5.3 CLU\# $\mathrm{P}_{\text {free }} \cap 0-1-\mathrm{F}=\mathrm{CLU} \# \mathrm{P} \cap 0-1-\mathrm{F}$.
Proof Let $f \in \mathrm{CLU} \# \mathrm{P}_{\text {free }} \cap 0-1-\mathrm{F}$ via some NPTM $M$ and nondeterminism polynomial $p$. Then, for all $x \in \Sigma^{*}, M$ has, at most, one accepting path. $M$ and $p$ together with $\prec, b$, and $t$ defined below witness that $f \in \mathrm{CLU} \# \mathrm{P} \cap 0-1-\mathrm{F}$. For all $x \in \Sigma^{*}$, define $b(x)$ as $0^{p(|x|)}$ and $t(x)$ as $1^{p(|x|)}$. For all $x, y, z \in \Sigma^{*}, \prec(x, y, z)$ evaluates to true exactly if $|y|=|z|=p(|x|)$ and $y \prec_{\text {lex }} z$.

Whether Theorem 5.3 holds for $0-1-2$-valued functions is open. (If we knew that the accepting-path cluster could without loss of generality be assumed never to extend to the top or bottom element, then 0-1-2-valued functions, $\mathcal{O}(1)$-valued functions, and much more would work in Theorem 5.3, via Theorem 3.3. However, the desired "without loss of generality" is not currently known to hold.)

Somewhat related to Theorem 5.3 is the following result, which shows that the sole obstacle to achieving CLU\#P $=C L U \# \mathrm{P}_{\text {circular }}$ is the possibility that the $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ machine (i.e., $M$ of $p, M, \prec$ ) has all paths accept on a hard-to-predict set of inputs. In particular, define $\widetilde{C L U \#} \mathrm{P}_{\text {circular }}$ to be the class of all $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ functions $f$ whose membership in CLU\# $\mathrm{P}_{\text {circular }}$ is instantiated by some $p, M$, and $\prec$ (in the sense of Definition 2.4) such that $\left\{x \mid f(x)=2^{p(|x|)}\right\} \in \mathrm{P}$.

Theorem 5.4 $\widetilde{\text { CLU\# }} \mathrm{P}_{\text {circular }}=\mathrm{CLU} \# \mathrm{P}$.
Proof The $\subseteq$ direction follows easily from Theorem 3.3, since by hypothesis we can test in P whether $f(x)=2^{p(|x|)}$ and with the correct answer in hand can produce in each case in a $\mathrm{UPSV}_{\mathrm{p}}$ fashion the accepting-block edge functions least and greatest. For the $\supseteq$ direction, let $p, M, \prec, t$, and $b$ capture $f$ in the sense of Definition 2.1. Clearly, $\left\{x \mid f(x)=2^{p(|x|)}\right\} \in \mathrm{P}$, via a deterministic polynomial-time Turing machine that simulates $t$ and $b$. Let $\prec^{\prime}$ be the same as $\prec$, except that, for any $x \in \Sigma^{*}, \prec^{\prime}(x, t(x), b(x))$ evaluates to true. Then $p, M$, and $\prec^{\prime}$ are a $\widehat{\text { CLU\# }} \mathrm{P}_{\text {circular }}$ model for $f$.

Considering Theorem 3.3 and the above proof closely, it is not hard to see that one can prove the $\subseteq$ direction above even if in the definition of $\widetilde{\text { CLU\# }} \mathrm{P}_{\text {circular }}$ one were to replace $"\left\{x \mid f(x)=2^{p(|x|)}\right\} \in \mathrm{P}$ " with " $\left\{x \mid f(x)=2^{p(|x|)}\right\} \in \mathrm{UP} \cap$ coUP." And the $\supseteq$ direction of course holds with this adjustment. Thus, Theorem 5.4 remains true even under that less restrictive alternate definition.

Though Theorem 5.4 shows that CLU\# $\mathrm{P}_{\text {circular }}$ and CLU\#P have only one obstacle blocking their equality, the following result provides a bit of evidence against their equality. It shows that from equality there follows a consequence that we do not see how to establish without the assumption of equality. (We say that a function $g: \Sigma^{*} \rightarrow \mathbb{N}$ is strictly positive if $\left(\forall x \in \Sigma^{*}\right)[g(x)>0]$.)

Theorem 5.5 If CLU\#P $=\operatorname{CLU} \# \mathrm{P}_{\text {circular }}$, then for every $\mathrm{CLU} \# \mathrm{P}$ function $f$ there is a strictly positive $\mathrm{FP}_{\mathrm{t}}$ function $g$ such that $f+g$ is in CLU\#P.

The proof of this theorem relies on the following "complementarity" property of the circular model.

Proposition 5.6 Let $f$ be computed by $M, p$, $\prec$ in the circular model. Then the function $\bar{f}(x)=2^{p(|x|)}-f(x)$ can be computed in the circular model.

Proof Let $M^{\prime}$ be the Turing machine that is exactly the same as $M$ except that, for each $x \in \Sigma^{*}$ and each path $\rho$ of $M(x)$, if $M(x)$ halts and rejects on path $\rho$ then $M^{\prime}(x)$ halts and accepts on its analogous path, and if $M(x)$ halts and accepts on path $\rho$ then $M^{\prime}(x)$ halts and rejects on its analogous path. We claim that $M^{\prime}, p$, and $\prec$ compute $\bar{f}$ in the circular model. To see this, choose $x \in \Sigma^{*}$. Suppose that $0<f(x)<2^{p(|x|)}$. Recall, as per the paragraph immediately following Definition 2.4, that we are allowed to adjust the bijection $h_{x}$ in such a way that $h_{x}^{-1}\left(0^{p(|x|)}\right)$ is an accepting path of $M(x)$ and $h_{x}^{-1}\left(1^{p(|x|)}\right)$ is a rejecting path. Thus, we may assume that the first element under the image of $h_{x}$ is the least accepting path of $M$. Then, obviously, the last element under the image of $h_{x}$ is the greatest accepting path of $M^{\prime}$ on input $x$. If $f(x) \in\left\{0,2^{p(|x|)}\right\}$ then $h_{x}$ witnessing that the computational paths of $M$ on input $x$ are clustered in the circular model will also witness that the computational paths of $M^{\prime}$ on input $x$ are clustered in the circular model. Since we do not change $p$ or $\prec$ ( or $h_{x}$ ), we easily see that $M^{\prime}$ computes exactly $\bar{f}(x)=2^{p(|x|)}-f(x)$ in the circular model.
Proof of Theorem 5.5 Let $f$ be a CLU\#P function via NPTM $M$, nondeterminism polynomial $p$, predicate $\prec, \mathrm{FP}_{\mathrm{t}}$ functions $b$ and $t$. Without loss of generality, we may assume that for all $x \in \Sigma^{*}, f(x)<2^{p(|x|)} .{ }^{4}$ Clearly, $f$ is in CLU\#P $\mathrm{P}_{\text {circular }}$ via the same NPTM $M$, the same nondeterminism polynomial $p$, and the new predicate $\prec^{\prime}$ that is the same as $\prec$ except that, for any $x \in \Sigma^{*}, \prec^{\prime}(x, t(x), b(x))$ evaluates to true. Thus, by Proposition 5.6 the function $\bar{f}$, defined for all $x \in \Sigma^{*}$ as $\bar{f}(x)=2^{p(|x|)}-f(x)$, belongs to CLU\# $\mathrm{P}_{\text {circular }}$. Define the function $f^{\prime}$ on each $x \in \Sigma^{*}$ as $f^{\prime}(x)=\max \{0, \bar{f}(x)-1\}(=\bar{f}(x)-1$, because $\bar{f}$ is strictly positive). From the assumption CLU\#P $=\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ and Theorem 4.3, we conclude that $f^{\prime} \in \mathrm{CLU} \# \mathrm{P}$, say, using the nondeterminism polynomial $q$. Here, we may assume that for all $x \in \Sigma^{*}, q(|x|) \geq p(|x|)$. Again, we have $f^{\prime} \in \mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ and this can be witnessed using the same polynomial $q$. Now, let $\overline{f^{\prime}}$ be the function defined for all $x \in \Sigma^{*}$ as $\overline{f^{\prime}}(x)=2^{q(|x|)}-f^{\prime}(x)$. Then, by Proposition 5.6, $\overline{f^{\prime}}$ is a CLU\# $\mathrm{P}_{\text {circular }}$ function

[^4]and thus, once more by our assumption, a CLU\#P function. Define an $\mathrm{FP}_{\mathrm{t}}$ function $g$, for each $x \in \Sigma^{*}$, as $g(x)=2^{q(|x|)}-2^{p(|x|)}+1$. It follows that for all $x \in \Sigma^{*}, g(x) \geq 1$ and $\overline{f^{\prime}}(x)=2^{q(|x|)}-\left(2^{p(|x|)}-f(x)-1\right)=f(x)+g(x)$. So $f+g \in$ CLU\#P.

Finally, we have the following result, which shows that if it does hold that CLU\#P $=$ CLU\# $\mathrm{P}_{\text {free }}$ or $\mathrm{CLU} \# \mathrm{P}=\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$, then we probably can expect that CLU\#P will at least need to in some cases use more nondeterminism than the other two classes.

## Theorem 5.7

1. If for each $p, M$, and $\prec$ that instantiate a $\mathrm{CLU} \# \mathrm{P}_{\text {circular }}$ function that function is also instantiated by a CLU\#P machine having nondeterminism exactly $p$, then $\mathrm{P}=\mathrm{UP}$.
2. If for each $p, M$, and $\prec$ that instantiate $a \mathrm{CLU}_{\mathrm{P}} \mathrm{P}_{\text {free }}$ function that function is also instantiated by a CLU\#P machine having nondeterminism exactly $p$, then $\mathrm{P}=\mathrm{UP}$.
3. If each $\mathrm{CLU}_{\mathrm{H}} \mathrm{P}_{\text {free }}$ machine has $\mathrm{UPSV}_{\mathrm{t}}$ functions $t$ and $b$, then $\mathrm{UP}=\mathrm{coUP}$.
4. If each $\mathrm{CLU} \# \mathrm{P}_{\text {free }}$ machine has $\mathrm{FP}_{\mathrm{t}}$ functions $t$ and $b$, then $\mathrm{P}=\mathrm{UP}$.

Proof Let $L \in$ UP and let $U$ be an NPTM such that $L(U)=L$ and that $U$ has on any input at most one accepting path. Assume without loss of generality that $U$ is balanced via polynomial $p$ where $p>0$ and that $U$ never accepts on path $1^{p(|x|)}$. Let $U^{\prime}$ be an NPTM that is the same as $U$ except that the accepting and rejecting states are switched.

For item 1, define $\prec^{\prime \prime}$ as follows: For all $x, y, z \in \Sigma^{*}, \prec^{\prime \prime}(x, y, z)$ evaluates to true exactly if $|y|=|z|$ and either $y \prec_{\text {lex }} z$ or $\left(y \in 1^{*}\right.$ and $\left.z \in 0^{*}\right)$. Clearly, $U^{\prime}, p$, and $\prec^{\prime \prime}$ captures \#acc $U_{U^{\prime}}$ in the sense of Definition 2.4.

For items 2, 3, and 4, define $\prec$ as follows: For all $x, y, z \in \Sigma^{*}, \prec(x, y, z)$ evaluates to true exactly if:

1. $\{y, z\} \subseteq \operatorname{acc}_{U^{\prime}}(x)$ and either
(a) $y \prec_{\text {lex }} z$, or
(b) $\left(\exists w \notin \operatorname{acc}_{U^{\prime}}(x)\right)\left[y \prec_{\operatorname{lex}} w \prec_{\operatorname{lex}} z\right]$,
or
2. $y=1^{p(|x|)}$ and $z \notin \operatorname{acc}_{U^{\prime}}(x)$.

Clearly, $p, U^{\prime}$, and $\prec$ captures $\# a c c_{U^{\prime}}$ in the sense of Definition 2.3.
Since we are only considering cluster machines, all the accepting paths must be contiguous with respect to $\prec$, and this forces any rejecting path to occur as the bottom or the top path. Under the assumptions of item 3 , then, $\bar{L} \in \mathrm{UP}$ by a machine that simulates in sequence UPSV $_{\mathrm{t}}$ machines for the $t$ and $b$ functions of $p, U^{\prime}$, and $\prec$. Under the assumptions
of item $4, L \in \mathrm{P}$ via a deterministic polynomial-time Turing machine that simulates in sequence the $t$ and $b$ functions of $p, U^{\prime}$, and $\prec$.

For items 2 and 1 , suppose that $\# a c c_{U^{\prime}} \in \mathrm{CL} \# \mathrm{P}$ via some $M, p^{\prime}, b^{\prime}, t^{\prime}, \prec^{\prime}$ and suppose that $p=p^{\prime}$. Then $L \in P$ via a deterministic, polynomial-time Turing machine that, on input $x \in \Sigma^{*}$ accepts if and only if $\left\{b^{\prime}(x), t^{\prime}(x)\right\} \nsubseteq a c c_{M}(x)$. Note that because $M$ has the same amount of nondeterminism as $U^{\prime}$ it follows that $M$ has, at most, one rejecting path (which must occur at the top or bottom of the order) and this path occurs if and only if $x \in L$.

Acknowledgments We are deeply grateful to Klaus W. Wagner for hosting the visit to Würzburg during which much of this work was done, for proposing the class CLU\# $\mathrm{P}_{\text {free }}$, and for many other valuable suggestions and insights.

## References

[BDG90] J. Balcázar, J. Díaz, and J. Gabarró. Structural Complexity II. EATCS Monographs in Theoretical Computer Science. Springer-Verlag, 1990.
[BDG95] J. Balcázar, J. Díaz, and J. Gabarró. Structural Complexity I. EATCS Texts in Theoretical Computer Science. Springer-Verlag, 2nd edition, 1995.
[BLS84] R. Book, T. Long, and A. Selman. Quantitative relativizations of complexity classes. SIAM Journal on Computing, 13(3):461-487, 1984.
[BLS85] R. Book, T. Long, and A. Selman. Qualitative relativizations of complexity classes. Journal of Computer and System Sciences, 30(3):395-413, 1985.
$\left[\mathrm{CGH}^{+} 88\right]$ J. Cai, T. Gundermann, J. Hartmanis, L. Hemachandra, V. Sewelson, K. Wagner, and G. Wechsung. The boolean hierarchy I: Structural properties. SIAM Journal on Computing, 17(6):1232-1252, 1988.
[GH00] C. Glaßer and L. Hemaspaandra. A moment of perfect clarity I: The parallel census technique. SIGACT News, 31(3):37-42, 2000.
[GS88] J. Grollmann and A. Selman. Complexity measures for public-key cryptosystems. SIAM Journal on Computing, 17(2):309-335, 1988.
[HHKW05] L. Hemaspaandra, C. Homan, S. Kosub, and K. Wagner. The complexity of computing the size of an interval. Technical Report TR-856, Department of Computer Science, University of Rochester, Rochester, NY, February 2005. Revised, March 2005. This is an expanded version of [HKW01].
[HKW01] L. Hemaspaandra, S. Kosub, and K. Wagner. The complexity of computing the size of an interval. In Proceedings of the 28th International Colloquium on Automata, Languages, and Programming, pages 1040-1051. Springer-Verlag Lecture Notes in Computer Science \#2076, July 2001.
[Kos99] S. Kosub. A note on unambiguous function classes. Information Processing Letters, 72(5-6):197-203, 1999.
[MP79] A. Meyer and M. Paterson. With what frequency are apparently intractable problems difficult? Technical Report MIT/LCS/TM-126, Laboratory for Computer Science, MIT, Cambridge, MA, 1979.
[Sch76] C. Schnorr. Optimal algorithms for self-reducible problems. In Proceedings of the 3rd International Colloquium on Automata, Languages, and Programming, pages 322-337. Edinburgh University Press, July 1976.
[Sel90] A. Selman. A note on adaptive vs. nonadaptive reductions to NP. Technical Report 90-20, Department of Computer Science, State University of New York at Buffalo, Buffalo, NY, September 1990.
[Sel94] A. Selman. A taxonomy of complexity classes of functions. Journal of Computer and System Sciences, 48(2):357-381, 1994.
[VW95] H. Vollmer and K. Wagner. Complexity classes of optimization functions. Information and Computation, 120(2):198-219, 1995.
[Wag90] K. Wagner. Bounded query classes. SIAM Journal on Computing, 19(5):833846, 1990.


[^0]:    *Supported in part by grant NSF-CCF-0426761. Work done in part while visiting Julius-MaximiliansUniversität Würzburg. URL: www.cs.rochester.edu/u/lane.
    ${ }^{\dagger}$ URL: www.cs.rit.edu/~ ${ }^{\text {cmh. }}$
    ${ }^{\ddagger}$ URL: www14.in.tum.de/personen/kosub.

[^1]:    ${ }^{1}$ Throughout this paper, we use "shuffle" in its common-language sense of permuting a single collection, rather than in the very different sense in which the term is sometimes used in theoretical computer science, namely, taking two separate lists and interleaving them.

[^2]:    ${ }^{2}$ One might note that, on the other hand, CL\#P lacks the $\mathrm{FP}_{\mathrm{t}}$ constraints (on the top and bottom elements among the computation paths) that CLU\#P obeys, and in that way at least potentially might seem to have some flexibility that CLU\#P might lack. However, though CL\#P does not explicitly speak of top and bottom functions at each length, it is not hard to see that it has top and bottom functions (mapping from each $x$-or even from $0^{|x|}$ _ to the top and bottom elements at length $\left.p(|x|)\right)$ that are computable in $\mathrm{UPSV}_{\mathrm{t}}$. We will show later in this section that CLU\#P remains unchanged if one allows its top and bottom functions to be drawn not just from $\mathrm{FP}_{\mathrm{t}}$ but even from $\mathrm{UPSV}_{\mathrm{t}}$. Thus, CLU\#P is not at a disadvantage on this issue.

[^3]:    ${ }^{3}$ In this proof we always are speaking relative to $x$. so this statement actually means that $\{\langle x, y\rangle \mid 0 \neq$ $\left.|y|=p(|x|) \wedge\left(\exists z \in\{0,1\}^{p(|x|)}\right)[\prec(x, y, z)]\right\}$.

[^4]:    ${ }^{4}$ Otherwise consider an NPTM $N$ that on input $x \in \Sigma^{*}$, guesses a string $u \in\{0,1\}^{p(|x|)+1}$, simulates $M$ along path suffix $(u, p(|x|))$ if the first bit of $u$ is zero, and otherwise rejects along the path $u$. For all $x, y, z \in$ $\Sigma^{*}, a, b \in \Sigma$, the corresponding predicate $\prec^{\prime}(x, a y, b z)$ evaluates to true if and only if $[a=b \wedge \prec(x, y, z)$ ] or $[a=0 \wedge y=t(x) \wedge b=1 \wedge z=b(x)]$. Bottom and top functions $b^{\prime}$ and $t^{\prime}$ are given by $b^{\prime}(x)=0 b(x)$ and $t^{\prime}(x)=1 t(x)$.

