

Cons-free Programs and Complexity Classes between LOGSPACE and PTIME

Neil D. Jones

Computer Science Department
University of Copenhagen
neil@diku.dk

Siddharth Bhaskar

Computer Science Department
University of Copenhagen
sbhaskar@di.ku.dk

Cynthia Kop

Department of Software Science
Radboud University, Nijmegen
C.Kop@cs.ru.nl

Jakob Grue Simonsen

Computer Science Department
University of Copenhagen
simonsen@di.ku.dk

Programming language concepts are used to give some new perspectives on a long-standing open problem: is $\text{LOGSPACE} = \text{PTIME}$?

Introduction

“ $P = ? NP$ ” is an archetypical question in computational complexity theory, unanswered since its formulation in the 1970s. The question: Is the computational power of polynomially time-bounded programs increased by adding the ability to “guess” (i.e., nondeterminism) ? This is interesting because “polynomial time” is a plausible candidate for “feasibly solvable”.

Perhaps the second most important question is “ $L = ? P$ ”: whether $\text{LOGSPACE} = \text{PTIME}$. Here L is the set of problems solvable by *cursor programs*. These also run in polynomial time, but have no rewritable storage¹. Both questions remain open since Cook and Savitch’s pathbreaking papers in the 1970s [3, 15].

We investigate the question “ $L = ? P$ ” from the viewpoint of *functional programming languages*: a different viewpoint than Turing machines. The link is earlier characterisations of L and P by “cons-free” programs [6, 7, 8]. The net result: a deeper and finer-grained analysis, illuminated by perspectives both from programming languages and complexity theory.

Some new definitions and theorems give fresh perspectives on the question $L = ? P$. We use programs to define and study complexity classes between the two. By [6, 7, 8] cursor programs exactly capture the problem class L ; and cursor programs with *recursive function definitions* exactly capture the problem class P . A drawback though is that recursive cursor programs *can run for exponential time*, even though they exactly capture the *decision problems* that can be solved in polynomial time by Turing machines.

The goal of this paper is to better understand the problems in the interval between classes L and P . Problem class NL is already-studied in this interval, and it is the logspace analog of similar long-standing open problems. Kuroda’s two “LBA problems” posed in 1964 [11]: (1) Is $\text{DSPACE}(n) = ? \text{NSPACE}(n)$ and (2) Is $\text{NSPACE}(n)$ closed under complementation? After both stood unresolved for 23 years, (2) was finally answered “yes” (independently in 1987) by Immerman and by Szelepcsényi [5, 16]: NL and larger nondeterministic space classes (with constructive bounds) are closed under complementation.²

¹One take: a cursor program is a multihead two-way read-only finite automaton. A more classical but equivalent version: a 2-tape Turing machine with n -bit read-only input tape 1, that uses at most $O(\log n)$ bits of storage space on read-write tape 2.

²Kuroda’s other LBA problem $\text{DSPACE}(n) = ? \text{NSPACE}(n)$ is still open, as well as the question $L = ? NL$.

We study the problems solvable by an in-between class CFpoly: recursive cursor programs that *run in polynomial time*. Recursion is in some sense orthogonal to the ability to make nondeterministic choices, i.e., to “guess”. The class CFpoly seems more natural than NL from a programming perspective.

1 Complexity by Turing machines and by programming languages

1.1 Overview

Let $X \subseteq \{0, 1\}^*$ be a set of bit strings. The *decision problem* for X : given $x \in \{0, 1\}^*$, to decide whether or not $x \in X$. We say that X is in PTIME iff it is decidable by a 1-tape deterministic Turing machine that runs within polynomial time $O(n^k)$ [here $n = |x|$ is the length of its input x , and k is a constant independent of x]. Further, X is in LOGSPACE iff it is decidable by a 2-tape Turing machine that uses at most $O(\log n)$ bits of storage space on tape 2, assuming its n -bit input is given on read-only tape 1. Both problem classes are of practical interest as they are decidable by programs with running times bounded by polynomial functions of their input lengths. The essential difference is the amount of allowed storage space. These classes are invariant across a wide range of variations among computation models, and it is easy to see that $\text{LOGSPACE} \subseteq \text{PTIME}$.

However, the question: is $\text{LOGSPACE} \subsetneq \text{PTIME}$? has stood open for many years.

Programs and problem decision. Semantics: a program computes a (partial!) function from bit strings to bits:

$$\llbracket p \rrbracket : \{0, 1\}^* \rightarrow \{0, 1\} \cup \{\perp\}$$

Program semantics is call-by-value; and $\llbracket p \rrbracket(x) = \perp$ means: p does not terminate on input x .

A set $X \subseteq \{0, 1\}^*$ is *decided* by a program p if p terminates on all inputs $x \in \{0, 1\}^*$, and for any x ,

$$\llbracket p \rrbracket(x) = \begin{cases} 1 & \text{if } x \in X \\ 0 & \text{if } x \notin X \end{cases}$$

Complexity by cons-free programs. We use programming languages (several in this paper) to explore the interesting boundary zone between the problem classes LOGSPACE and PTIME. Strong links were established in [7, 8]: Each class was characterised by a small general recursive functional programming language. The one language (called CF for “cons-free”) is limited to programs without data constructors.³ The other, named CFTR, is identical to CF but has restricted control, allowing only *tail recursive* calls to defined functions.⁴ From [8] we know:

Theorem 1 $\text{LOGSPACE} = \{X \subseteq \{0, 1\}^* \mid \text{some CFTR program decides } X\}$

Theorem 2 $\text{PTIME} = \{X \subseteq \{0, 1\}^* \mid \text{some CF program decides } X\}$

A compact notation relating programs and problems

Definition 1 $\{\{L\}\}$ is the set of all problems (i.e., sets $X \subseteq \{0, 1\}^*$) that are decidable by L-programs:

$$\{\{L\}\} \stackrel{\text{def}}{=} \{X \subseteq \{0, 1\}^* \mid \text{some L-program } p \text{ decides } X\}$$

Theorems 1 and 2 can thus be restated as: $\text{LOGSPACE} = \{\{CFTR\}\} \subseteq \{\{CF\}\} = \text{PTIME}$

³“Cons-free” is not to be confused with “context-free”. TR stands for “tail recursive”. Data access is *read-only* in both of our languages, so neither CF nor CFTR is Turing-complete.

⁴These ideas stem from S. A. Cook’s seminal work on complexity classes and pushdown machines [3]. Papers [7, 8] re-express and adapt Cook’s ideas to a cons-free programming context. Paper [8] extends [3], characterising decision powers of higher-order cons-free programs. CFTR is (yet another) version of the “cursor programs” mentioned in the Introduction.

1.2 The cons-free programming language CF

All programs are first-order in this paper,⁵ and deterministic to begin with (until nondeterminism is expressly added in Section 5). First, a brief overview (details in [6, 7, 8]):

Definition 2 Program syntax: *a CF-program is a finite sequence of mutually recursive function definitions. Write CF programs and fragments in teletype, e.g., program p or expression $\text{tail}(f\ x\ y)$. The first definition, of form $f_1\ x = e^{f_1}$, defines the program entry function (with one argument, always named x). A context-free grammar for programs, definitions and expressions e :*

```
p    ::= def | p def      -- Program = sequence of function definitions
def  ::= f x1...xm = e   -- (where m >= 0)
e    ::= True | False | [] | xi | base | if e then e else e | f e1...em
base ::= not e | null e | head e | tail e      -- Base function call
```

Data types: *bits and bit lists. Variable and expression values are elements of one of the value sets $\{0, 1\}$ or list $\{0, 1\}^*$. A string $x \in \{0, 1\}^*$ is a bit list, and list 1011 can be written $[1, 0, 1, 1]$ as data value. $[]$ is the empty list, and $b:bs$ is the result of prepending bit b to list bs . We sometimes identify 0, 1 with False and True, resp. e.g., in the test position e_0 of an expression $\text{if } e_0 \text{ then } e_1 \text{ else } e_2$.*

Expressions: *Expression e may be a constant; a variable; a base function call (not, null, head, or tail); a conditional $\text{if } e_0 \text{ then } e_1 \text{ else } e_2$; or a call $f\ e_1 \dots e_r$ to some program-defined function f . A program contains no CONS function or other constructors (hence CF for cons-free). Thus CF does not have successor, +, * or explicit storage allocators such as malloc in C, :: in ML, or : in Haskell.*

Function definition: *A program-defined definition has form $f\ x_1\ x_2 \dots x_m = e$ with $m \geq 0$. No function may be defined more than once, left-side variables must be distinct, and any variable appearing on the right side (in e) must be one of x_1, \dots, x_m .*

Semantics: *The semantics of CF is given by the inference rules in Figure 1. Given a program p and input x , these rules define a (unique) computation tree that we call $\mathcal{T}^{p,x}$. The inference rules define statements $\llbracket p \rrbracket(x) = v$ (at the root of the computation tree) or $p, \rho \vdash e \rightarrow v$. The p in $p, \rho \vdash e \rightarrow v$ is a CF program. The e is a subexpression of the right side of some function definition. Further, ρ is an environment that binds the current program variables to respective argument values v_1, \dots, v_m ; and v is the result value of e . Base or defined function calls are evaluated using call-by-value.*

The full inference rule set is given in Figure 1. It is essentially the first-order part of Fig. 1 in [8].

Example 1 *The CF program parity decides membership in the set $X = \{x \in \{0, 1\}^* \mid |x| \text{ is even}\}$.*

```
entry x = even x
even z  = if (null z) then True else not(even(tail z))
```

This satisfies $\llbracket \text{parity} \rrbracket(x) = \text{True}$ if $|x|$ is even, else False. For example, $\llbracket \text{parity} \rrbracket([1, 0, 1]) = \text{False}$.

1.3 The computation tree $\mathcal{T}^{p,x}$ and an evaluation order

Evaluation steps: The computation tree $\mathcal{T}^{p,x}$ can be built systematically, applying the rules of Figure 1 bottom-up and left-to-right. Initially we are given a known program p and *input* $\in \{0, 1\}^*$, and the computation goal is $\llbracket p \rrbracket(\text{input}) = ?$, where $?$ is to be filled in with the appropriate value v (if it exists). Intermediate goals are of form $p, \rho \vdash e \rightarrow ?$, where p, ρ and e are known, and $?$ is again to be filled in.

⁵The larger language described in [8] encompasses both first and higher-order programs.

Axioms:

$$\frac{}{p, \rho \vdash x \rightarrow \rho(x)} \quad \frac{}{p, \rho \vdash \text{True} \rightarrow \text{True}} \quad \frac{}{p, \rho \vdash \text{False} \rightarrow \text{False}} \quad \frac{}{p, \rho \vdash [] \rightarrow []}$$

Base functions:

$$\frac{p, \rho \vdash e \rightarrow \text{False}}{p, \rho \vdash \text{not } e \rightarrow \text{True}} \quad \frac{p, \rho \vdash e \rightarrow []}{p, \rho \vdash \text{null } e \rightarrow \text{True}} \quad \frac{p, \rho \vdash e \rightarrow v_1:v_2}{p, \rho \vdash \text{head } e \rightarrow v_1}$$

$$\frac{p, \rho \vdash e \rightarrow \text{True}}{p, \rho \vdash \text{not } e \rightarrow \text{False}} \quad \frac{p, \rho \vdash e \rightarrow v_1:v_2}{p, \rho \vdash \text{null } e \rightarrow \text{False}} \quad \frac{p, \rho \vdash e \rightarrow v_1:v_2}{p, \rho \vdash \text{tail } e \rightarrow v_2}$$

Condition:

$$\frac{p, \rho \vdash e_0 \rightarrow \text{True} \quad p, \rho \vdash e_1 \rightarrow v}{p, \rho \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rightarrow v} \quad \frac{p, \rho \vdash e_0 \rightarrow \text{False} \quad p, \rho \vdash e_2 \rightarrow v}{p, \rho \vdash \text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rightarrow v}$$

Function call:

$$\frac{p, \rho \vdash e_1 \rightarrow w_1 \quad \dots \quad p, \rho \vdash e_m \rightarrow w_m \quad p, [x_1 \mapsto w_1, \dots, x_m \mapsto w_m] \vdash e^f \rightarrow v}{p, \rho \vdash f e_1 \dots e_m \rightarrow v} \text{ if } f x_1 \dots x_m = e^f \in p$$

Program running:

$$\frac{p, [x \mapsto \text{input}] \vdash e^{f_1} \rightarrow v}{\llbracket p \rrbracket(\text{input}) = v} \text{ If } f_1 x = e^{f_1} \text{ is the entry function definition of } p$$

Figure 1: Inference rules for CF

Axioms have no premises at all, and base function calls have exactly one premise, making their tree construction straightforward. The inference rules for conditional expressions have two premises, and a call to an m -ary function has $m + 1$ premises. We choose to evaluate the premises left-to-right. For the case $\text{if } e_0 \text{ then } e_1 \text{ else } e_2 \rightarrow ?$, there are two possibly applicable inference rules. However both possibilities begin by evaluating e_0 : finding v such that $e_0 \rightarrow v$. Once v is known, only one of the two if rules can be applied.

Claim. the evaluation order above for given p and input is general: if $\llbracket p \rrbracket(\text{input}) = v$ is deducible by any finite proof at all, then the evaluation order will terminate with result $\llbracket p \rrbracket(\text{input}) = v$ at the root.

A consequence: tree $\mathcal{T}^{p,x}$ is unique if it exists, so CF is a deterministic language. (Nondeterministic variants of CF will be considered later.)

1.4 The suffix property

CF programs have no constructors. This implies that all values occurring in the tree $\mathcal{T}^{p,x}$ must be boolean values, or suffixes of the input. Expressed more formally:

Definition 3 For program p and input $x \in \{0, 1\}^*$, define its range of variation and its reachable calls by

$$V_x = \{0, 1\} \cup \text{suffixes}(x) \text{ and}$$

$$\text{Reach}_p(x) = \{(f, \rho) \mid e^f \text{ is called at least once with environment } \rho \text{ while computing } \llbracket p \rrbracket(x)\}$$

Lemma 1 The computation tree $\mathcal{T}^{p,x}$ is finite iff p terminates on input x .

Lemma 2 If $\mathcal{T}^{p,x}$ contains $p, \rho \vdash e \rightarrow v$ then $v \in V_x$ and $\rho(z) \in V_x$ for every $z \in \text{domain}(\rho)$.

Proof: a straightforward induction on the depth of computation trees. If the entry function definition is $f\ x = e^f$, then the root $\llbracket p \rrbracket(x) = v$ has a parent of the form $p, \rho_0 \vdash e^f \rightarrow v$ where the initial environment is $\rho_0 = [x \mapsto x]$. Of course x is a suffix of x . At any non-root node in the computation tree, any new value must either be a Boolean constant or $[]$, or constructed by a base function `not`, `head`, `tail`, or `null` from a value already proven to be in V_x .

Consequence: a CF computation has at most polynomially many different function call arguments.

Lemma 3 For any CF program p there is a polynomial $\pi(|x|)$ that bounds $\#Reach_p(x)$ for all $x \in \{0, 1\}^*$.

Proof: By Lemma 2, any argument of any m -ary p function f has at most $\#V_x = 3 + |x|$ possible values, so the number of reachable argument tuples for f can at most be $(3 + |x|)^m$.

Remark: this does not necessarily imply that p 's running time is polynomial in $|x|$. We say more on this in Section 2.

1.5 The tail recursive programming language CFTR

A CFTR program is a CF program p such that every function call occurring on the right side of any defined function is *in tail form*. This is a *syntactic* condition; the semantic purpose is to enable a no-stack CFTR implementation, as outlined in Section 2.1. The operational intention is that if a call $f\ e_1 \dots e_m$ is in tail form then, after the call has been performed, control is immediately returned from the current function: there is “nothing more to do” after the call.

Definition 4 Define function $\alpha : Exp \rightarrow \{X, T, N\}$ as follows, where the set of descriptions is ordered by $X < T < N$. The intention: if expression e contains no function calls then $\alpha(e) = X$, else if all function calls in e are in tail form then $\alpha(e) = T$, else if e contains a function call in one or more non-tail-call positions then $\alpha(e) = N$.

| | | |
|--|------------------------------------|--|
| $\alpha(e)$ | $= X$ | if e is a constant |
| $\alpha(\text{base } e)$ | $= X$ | if $\alpha(e) = X$ |
| $\alpha(\text{base } e)$ | $= N$ | if $\alpha(e) = T$ or $\alpha(e) = N$ |
| $\alpha(f\ e_1 \dots e_m)$ | $= T$ | if $\alpha(e_1) = \dots = \alpha(e_m) = X$ |
| $\alpha(f\ e_1 \dots e_m)$ | $= N$ | otherwise |
| $\alpha(\text{if } e_0 \text{ then } e_1 \text{ else } e_2)$ | $= \max(\alpha(e_1), \alpha(e_2))$ | if $\alpha(e_0) = X$ |
| $\alpha(\text{if } e_0 \text{ then } e_1 \text{ else } e_2)$ | $= N$ | otherwise |

Definition 5 CF program p is in CFTR iff $\alpha(e) \in \{X, T\}$ for all function definitions $f\ x_1 \dots x_m = e$.

Some instances in the program `parity`, to clarify the definition:

1. $\alpha(\text{null } z) = \alpha(\text{tail } z) = X$. Neither calls any program-defined functions.
2. Expression `even x` is a call, and is in tail form: $\alpha(\text{even } x) = T$.
3. The expression `even(tail z)` is in tail form: $\alpha(\text{even}(\text{tail } z)) = T$.
4. The expression `not(even(tail z))` is *not* in tail form: $\alpha(\text{not}(\text{even}(\text{tail } z))) = N$.

By point 4, `parity` is *not* a CFTR program. However the set X it decides *can* be decided by the following alternative program whose recursive calls are in tail form. Thus X is CFTR decidable.

Example 2 *Program parity'*

```

entry' x = f x True
f x y    = if (null x) then y else f (tail x) (not y)

```

A very useful result from [1] is the following: we can assume, without loss of generality, that a CF program contains *no nested function calls* such as $f\ x_1 \dots x_m = \dots g(h(e)) \dots$.

Lemma 4 *For any CF program p there is a CF program p' such that $\llbracket p \rrbracket = \llbracket p' \rrbracket$, and for each function definition $f\ x_1 \dots x_m = e^f$ in p' , either $\alpha(e^f) \leq T$, or $e^f = \text{if } e_0 \text{ then } e_1 \text{ else } e_2$ and $\alpha(e_i) \leq T$ for $i \in \{0, 1, 2\}$.*

An interesting step in the proof of Lemma 4 is to show that a function $f : X \rightarrow Y$ is CF-computable if and only if its graph $G^f = \{(x, y) \in X \times Y \mid f(x) = y\}$ is CF-decidable (see [1] for details). In a CFTR program, no function calls are allowed to occur in the test part of a tail-recursive conditional expression (by the last line of the α definition). Thus Lemma 4 does not rule out a call nested in a conditional's test part, and so does not imply that p' is tail-recursive.

An example that suggests that *some* call nesting is essential: the *MCV* program of Section 2.3 has function calls inside *if* tests. It is well-known that the function that *MCV* computes can be computed tail-recursively if and only if $\text{LOGSPACE} = \text{PTIME}$ [14].

2 About time and space measures

Theorems 1 and 2 relate computations by two rather different computation models: cons-free programs and Turing machines. We now focus on time and space relations between the two. We define the time and space used to run given CF program p on input $x \in \{0, 1\}^*$:

Definition 6

- $\text{time}_p(x)$ is the number of evaluation steps to run program p on x (as in Section 1.3).
- $\text{space}_p(x)$ is the number of bits required to run p on input x .

Both $\text{time}_p(x)$ and $\text{space}_p(x)$ are non-negative integers (or \perp for nontermination). We call $\text{time}_p(x)$ the “native time” for CF or CFTR (in contrast to Turing machine time). For the example, $\text{time}_{\text{parity}}(x) = O(|x|)$. It is reasonably simple to define running time as $\text{time}_p(x) = |\mathcal{T}^{p,x}|$, i.e., the number of nodes in the computation tree $\mathcal{T}^{p,x}$.

A full definition of $\text{space}_p(x)$ requires a lower-level execution model than the abstract semantics above. (Section 2.1 will show a bit more detail.) Consider the *parity* example. The length of the call stack needed to implement the program is linear in the depth of the proof tree, and so $O(|x|)$. Moreover, each non-root node has an environment ρ that binds x or z to its value. By Lemma 2 any such value is a suffix of the input. Any suffix can be represented in $O(\log |x|)$ bits, so the total space that *parity* uses is

$$\text{space}_{\text{parity}}(x) = O(|x| \log |x|)$$

2.1 Space usage: a lower-level operational semantics, and the tail call optimisation

To more clearly define space usage and present the tail call optimisation, we use a finer, more operational level of detail based on traditional implementations⁶. Programs are executed sequentially as in the proof-tree-building algorithm seen earlier. The expression on the right side of a function definition is evaluated

⁶Readers may think of machine-code run-time states, e.g., value stacks, stack frames, . . . ; or of tapes and state transition relations. Such implementation paraphernalia are well-known both to compiler writers and programming language theorists.

one step at a time as in Section 1.3; but with an extra data structure, a stack σ used to record some of the environments seen before. The stack's topmost activation record corresponds closely to the current environment ρ .

Each defined function $f \ x_1 \dots x_m = e$ has an activation record format that contains the values v_1, \dots, v_m of the arguments with which f has been called (plus shorter-lived “temporary results,” i.e., as-yet-unused subexpression values). On any call to f , a new activation record is pushed onto the stack σ . This activation record will be popped when control returns from the call.

Definition 7 *The call history for CF program p and input $x \in \{0, 1\}^*$ is a sequence (finite or infinite) of configurations, each of form $(f, \rho) \in \text{Reach}_p(x)$ where f is a defined function name and ρ is an environment. The first configuration always has form $(p - \text{entry}, [z \mapsto x])$.*

How the syntactic CFTR condition in Definition 5 allows more efficient implementation:

The tail call optimisation

Claim: every CFTR program can be implemented *using at most one stack frame* at a time.

Suppose CF program p has a definition $f \ x_1 \dots x_r = \dots (g \ e_1 \dots e_s) \dots$ where the g call is a tail call. If so, then after the call $(g \ e_1 \dots e_s)$ has been started there can be no future references to the current values of x_1, \dots, x_r . Informally, there is “nothing more to do” after a tail call has been completed. (One can review the parity’ and parity examples in this light.)

The tail call optimisation: The new activation record for g *overwrites* f 's current activation record. In a CFTR program *every* call will be a tail call. Thus there is never more than one frame at a time in the runtime stack. One stack frame requires $O(\log |x|)$ bits. For example, this gives parity’ a major improvement over the space used by parity: $\text{space}_{\text{parity}'}(x) = O(\log |x|)$ bits.

A relaxation of the tail call restriction. Sketch (details in [1]): It is sufficient that the functions defined in p can be given a partial order $<$ such that no call is to another that is earlier in the order, and recursive calls must be tail calls. Under this lighter restriction there may be more than one frame at a time in the runtime stack; but these frames are properly ordered by $<$. Consequently every restricted program can be implemented within a *constant depth of stack frames*.

2.2 CFTR programs run in polynomial time; but CF can take exponential time

Theorem 3 *Any terminating CFTR program runs in (native) time polynomial in $|x|$.*

Proof: in the call history of p on input x , all runtime configurations must be distinct (else the program is in an infinite loop). There is a constant bound on the number of stack frames. Thus, by Lemma 3 there are only polynomially many distinct runtime configurations.

Theorem 4 *A terminating run of a CF program can take time exponential in its input length.*

First, a trivial example with non-polynomial running time (from [8]):

Example 3 `f x = if (null x) then True else
if f(tail x) then f(tail x) else False`

It is easy to see that $\text{time}_q(x) = 2^{O(|x|)}$ due to the repeated calls to `tail x`. Exponential time is not necessary, though, since the computed function satisfies $f(x) = \text{True}$.

2.3 A more fundamental example: a CF-decidable problem that is PTIME-complete

MCV is the *Monotone Circuit Value Problem*. Technically it is the set of all *straight-line Boolean programs* whose last computed value is True, e.g.,

$$x_0 := \text{False}; x_1 := \text{True}; x_2 := x_1 \vee x_0; x_3 := x_2 \wedge x_0; x_4 := x_3 \vee x_2; x_5 := x_4 \vee x_3$$

MCV is a well-known “hardest” problem for PTIME [14]. In complexity theory, completeness means two things: (1) *MCV* \in PTIME; and (2) if $X \subseteq \{0, 1\}^*$ is any problem in PTIME, then X is logspace-reducible⁷ to *MCV*. By transitivity of logspace-reduction, *MCV* \in LOGSPACE if and only if LOGSPACE = PTIME.

Following is the core of a Haskell program to decide membership in *MCV*, together with a sample run. We chose a compact Boolean program representation, so the program above becomes a list:

```
program = [ OR 5 4 3, OR 4 3 2, AND 3 2 0, OR 2 1 0 ]
```

Haskell encoding: Assignment $x_i := x_j \wedge x_k$ is represented by omitting $:=$, and coding \wedge as AND in prefix position (and similarly for \vee). The program is presented in reverse order; and variables X0, X1 always have predefined values False, True resp., so their assignments are omitted in the representation.

```
type Program      = [Instruction]
data Instruction = AND Int Int Int | OR Int Int Int

mcv :: Program      -> Bool
vv  :: (Int,Program) -> Bool

mcv ((OR lhs x y):s) = if vv(x,s) then True  else vv(y,s)
mcv ((AND lhs x y):s) = if vv(x,s) then vv(y,s) else False

vv(0,s) = False -- vv(v,s) = value of variable v at program suffix s
vv(1,s) = True  --
vv(v,s) = case s of
  ((AND lhs x y):s') -> if v==lhs then mcv(s) else vv(v,s')
  ((OR lhs x y):s')  -> if v==lhs then mcv(s) else vv(v,s')
```

This works by recursive descent, beginning at the end of the program; and uses no storage at all (beyond the implicit implementation stack). Following is a trace of nontrivial calls to *vv*, ended by the final result of running *p*. Note that variable X2 is evaluated repeatedly.

```
program = [X5:=X4 OR X3, X4:=X3 OR X2, X3:=X2 AND X0, X2:=X1 OR X0]

vv(X4, [ X4 := X3 OR X2, X3 := X2 AND X0, X2 := X1 OR X0 ])
vv(X3, [ X3 := X2 AND X0, X2 := X1 OR X0 ])
vv(X2, [ X2 := X1 OR X0 ])
vv(X2, [ X3 := X2 AND X0, X2 := X1 OR X0 ])
vv(X2, [ X2 := X1 OR X0 ])
```

True

⁷Details may be found in [6], Chapters 25, 26.

Re-expressing the Boolean program as a bit string for CF input

Of course CF programs do not have data values of type `Int`. However, any straight-line Boolean program can be coded by a CF bit string. Example: the Boolean program above can be coded as a 44-bit string⁸:

```
[1,1,1,0,  1,0,1, 0, 1,0,0, 0,1,1,
      1,0,0, 0, 0,1,1, 0,1,0,
      0,1,1, 1, 0,1,0, 0,0,0,
      0,1,0, 0, 0,0,1, 0,0,0]
```

```
-- x5 := OR  x4  x3
-- x4 := OR  x3  x2
-- x3 := AND x2  x0
-- x2 := OR  x1  x0
```

Bag of tricks: The program has 5 variables. The index i of any variable x_i can be coded by a 3-bit binary sequence we denote by \hat{i} , e.g., x_3 is coded by $\hat{3} = 0, 1, 1$. This has length 3 since $3 = \lceil \log_2 5 \rceil$.

The beginning `1, 1, 1, 0` of the program code gives (in unary) the code block length, 3 for this program. A Boolean assignment $x_i := x_j \text{ op } x_k$ is coded as $\hat{i} \text{ op } \hat{j} \hat{k}$, where the Boolean operators \vee, \wedge are coded as 0, 1 respectively.

Using this encoding, it is not hard to transform the Haskell-like program to use if-then-else statements rather than case; and then implement it in the true CF language.

A dramatic contrast in running times

- The CF program just sketched has repeated (and nested) function calls with the same argument values. For example, Boolean variable `X2` in program `p` is evaluated 3 times in the trace above. More generally, the running time bound is exponential:

$$time_{mcv}(x) = 2^{O(|x|)}$$

- On the other hand, one can show that *MCV can be decided by a Turing machine* in polynomial time: Execute the Boolean instructions from the beginning, store the values of left side variables when computed, and refer to stored variable values as needed.

This major difference is due to data storage: the Turing machine tape can save already-computed values, to be looked up as needed. A CF program has no such storage, and must resort to recomputation⁹.

Is this contrast necessary? It seems strange that the cost of strengthening tail recursive programs (in CFTR) by adding recursion (to get CF) is to raise run times from polynomial to exponential. The next sections, abbreviated from [8], show that the problem is a general one of the relation between LOGSPACE and PTIME. Thus it is not peculiar to the MCV problem, nor to the way we have programmed its solution.

3 How LOGSPACE = {{CFTR}} and PTIME = {{CF}} are proven

3.1 Proof sketch of Theorem 1, that LOGSPACE = {{CFTR}}

- For \supseteq , we simulate a CFTR program `p` by a LOGSPACE Turing machine (informal). Given an input $x = a_1 \dots a_n \in \{0, 1\}^*$, by Lemma 2 any reachable configuration (f, ρ) satisfies $0 \leq |\rho(x_i)| \leq \max(n, 1)$ for $i = 1, \dots, m$. Each v_i can be coded in $O(\log n)$ bits. Now `f` and the number of `f`'s arguments are independent of x , so an entire configuration (f, ρ) can be coded into $O(\log n)$ bits.

⁸The `--` parts delimit comments, and are not part of the Boolean program's bit-string encoding.

⁹Duplicating CF variables does not suffice, since the number of variables is independent of the length of the input data.

The remaining task is to show that the operational semantics of running p on x can be simulated by a LOGSPACE Turing machine. The key: because p is tail recursive, there is *no nesting* of calls to program-defined functions. The construction can be described by cases:

1. Expressions without calls to any defined function: Suppose p 's current environment ρ is represented on the Turing machine work tape. Simulating this evaluation is straightforward.
 2. Evaluation of an expression $f e_1 \dots e_m$: Since p is tail recursive, none of e_1, \dots, e_m contains a call, and there is “nothing more to do” after the call $f e_1 \dots e_m$ has been simulated. Assume inductively that $e_1 \dots e_m$ have all been evaluated. Given the definition $f x_1 \dots x_m = e$ of f , the remaining steps¹⁰ are to collect the values v_1, \dots, v_m construct a new environment $\rho' = [x_1 \mapsto v_1, \dots, x_m \mapsto v_m]$, and replace the current ρ by ρ' .
- For \subseteq , we simulate a LOGSPACE Turing machine by a CFTR program. This can be done using Lemma 2: represent a Tape 2 value by one or more suffixes of x . (A more general result is shown in [8] by “counting modules”.)

Remark: given an input x , the *number of times* that a call (f, ρ) appears in the call history of $\llbracket p \rrbracket(x)$ may be much larger than the number of such calls *with distinct argument values*, even exponentially larger.

3.2 Proof sketch of Theorem 2, that PTIME = {{CF}}

- For \subseteq , we simulate a PTIME Turing machine Z by a CF program. (this is the core of *Proposition 6.5* from [8]). Assume Turing machine $Z = (Q, \{0, 1\}, \delta, q_0)$ is given¹¹ and that it runs in time at most n^k for inputs x of length n . Consider an input $x = a_1 a_2 \dots a_n \in \{0, 1\}^*$.

Idea: in the space-time diagram of Z 's computation on x (e.g., Fig. 3 in [8]), data propagation is local; the information (control state q , symbol a , whether or not it is scanned) at time t and tape position i is a mathematical function of the same information at time $t - 1$ and tape positions $i - 1, i, i + 1$. This connection can be used to compute the contents of any tape cell at any time. Repeating for $t = 0, 1, \dots, n^k$ steps gives the final result of the computation (x is accepted or not accepted).

The simulating CF program computes functions

$$\begin{aligned} \text{state}(t) &= \text{the state } q \text{ that } Z \text{ is in at time } t \\ \text{position}(t) &= \text{the scanning position } i \in \{0, 1, 2, \dots, n^k\} \text{ at time } t \\ \text{symbol}(t, i) &= \text{the tape symbol } a \text{ found at time } t \text{ and scanning position } i \end{aligned}$$

Initially, at time $t = 0$ we have total state $(q_0, 0, \tau_0)$ where $\tau_0(i) = a_i$ for $1 \leq i \leq n$, else $\tau_0(i) = B$. Now suppose $t > 0$ and the total state at time $t - 1$ is (q, i, τ) , and that $\delta(q, a) = (q', a', d)$. Then the total state at time t will be $(q', i + d, \tau')$ where $\tau'(i) = a'$ and $\tau'(j) = \tau(j)$ if $j \neq i$.

Construction of a CF program z with definitions of *state*, *position* and *symbol* is straightforward. Arguments $t, i \in \{0, 1, 2, \dots, n^k\}$ can be uniquely encoded as tuples of suffixes of $x =$

¹⁰These steps are done using Turing machine representations of the environments as coded on Tape 2.

¹¹As usual Q is a finite set with initial state q_0 , and transition function of type $\delta : Q \times \{0, 1, B\} \rightarrow Q \times \{0, 1, B\} \times \{0, 1, -1\}$. A *total state* is a triple (q, i, τ) where $q \in Q, i \geq 0$ and the tape is $\tau : \mathbb{N} \rightarrow \{0, 1, B\}$. Scan positions are counted $i = 0, 1, \dots$ from the tape left end. Transition $\delta(q, a) = (q', a', d)$ means: if Z is in state q and $a = \tau(i)$ is the scanned tape symbol, then change the state to q' , write a' in place of the scanned symbol, and move the read head from its current position i to position $i + d$.

$a_1 a_2 \dots a_n \in \{0, 1\}^*$. It is not hard to see that such an encoding is possible; [8] works out details for an imperative version of the Turing machine (cf. Fig. 5).

- For \supseteq , we simulate an arbitrary CF program p by a PTIME Turing machine (call it Z_p). We describe the Z_p computation informally. Given an input $x \in \{0, 1\}^*$, Z_p will systematically build a *cache* containing $Reach_p(x)$. By Lemma 3 its size is polynomially bounded.

The cache (call it c) at any time contains a set of triples (f, ρ, v') where the simulator has completed evaluation of function f on argument values in environment ρ , and this f call has returned value v .

Concretely, c is built “breadth-first”: when a p call $f e_1 \dots e_m$ is encountered (initially c is empty):

- First, arguments e_1, \dots, e_m are evaluated to values v_1, \dots, v_m ; and these are collected into an environment ρ .
- Second, cache c is searched. If it contains (f, ρ, v) , then simulation continues using the value v as the value of $f e_1 \dots e_m$.
- If c contains no such triple, then p must contain a function definition $f x_1 \dots x_m = e^f$. Then expression e^f is evaluated in environment ρ to yield some value v . After finishing, add the triple (f, ρ, v) to c , and return the value v .

An observation: *The CF program of Section 3.2 (simulating a polynomial time Turing machine) has exponential runtime.*

Paradoxically, as observed in Section 2.2: even though CF exactly characterises PTIME, its programs do not run in polynomial time. The polynomial versus exponential contrast between the running times of CFTR and CF programs is interesting since both program classes are natural; and the *decision powers* of CFTR and CF are exactly the complexity classes LOGSPACE and PTIME. Alas, we have found no CF algorithm to simulate polynomial-time Turing machines. We explain how this happens in more detail.

Definition 8 *Call overlap occurs if a CF program can call the same function more than once with the same tuple of argument values.*

How can this happen? By Lemma 3, only polynomially many argument tuples are distinct. Consequently, superpolynomial running time implies that some simulation functions may be called repeatedly with the same argument value tuples, even though only polynomially many of them are distinct.

This can be seen in the proof that $PTIME \supseteq \{\{CF\}\}$: the value of $symbol(t+1, i)$ may depend on the values of $symbol(t, i-1)$, and $symbol(t, i)$, and $symbol(t, i+1)$. In turn, the value of $symbol(t, i)$ may depend on the values of $symbol(t-1, i-1)$, and $symbol(t-1, i)$, and $symbol(t-1, i+1)$.

Net effect: a symbol a_i on the final tape (with $t = n^k$) may depend many distinct function call sequences that “bottom out” at $t = 0$. The number of call sequences may be exponential in n .

Lemma 5 *Call overlap will occur if the length of the call history for $\llbracket p \rrbracket(x)$ is greater than $\#Reach_p(x)$.*

Unfortunately, it is hard to see *which* calls will overlap (it seems impossible without storage). Furthermore, the “caching” technique used to prove Theorem 2 cannot be used because, in contrast with Turing machines, CF programs do not have a memory that can accumulate previously computed values.

4 Closer to the boundary between CFTR and CF

Viewed *extensionally*, the difference (if any) between CFTR and CF corresponds to the difference (if any) between LOGSPACE and PTIME. Viewed *intensionally*, there seem to be significant differences, e.g., in algorithm expressivity as well as in running time. A relevant early result:

A **program transformation by Greibach** shows that programs whose calls (to defined functions) are all linear¹² can be made tail recursive [4]. Using this transformation, a CF program p whose calls are tail calls or linear calls may be transformed into one containing only tail calls (and no CONS or other constructors), so it is in CFTR. Thus p decides a problem in LOGSPACE.

There is a price to be paid, though: in general, the transformed program may run *polynomially slower than the original*, due to re-computation. For instance, the Greibach method can transform our first “parity” program into CFTR form. The cost: to raise time complexity from linear to quadratic.

Following is a new tool to investigate the problem. It is analogous to the set of programs for polynomial-time Turing machines, but adapted to the CF world. By Theorem 4, $\text{CFpoly} \subsetneq \text{CF}$.

Definition 9 *The programming language CFpoly has the same semantics as CF; but CFpoly’s programs are restricted to be those CF programs that terminate in polynomial time.*

Immediate:

$$\text{LOGSPACE} = \{\{\text{CFTR}\}\} \subseteq \{\{\text{CFpoly}\}\} \subseteq \{\{\text{CF}\}\} = \text{PTIME}$$

By Theorem 3, every terminating CFTR program is in CFpoly. One can see Greibach’s result as transforming a subset of CFpoly into CFTR.

Lemma 6 *If a terminating CF program does not have call overlap, then it is in CFpoly.*

The reason: by Lemma 3, such a CF program must run in polynomial time.

Remark: On the other hand, it may have non-tail calls, and so not be in CFTR.

Lemma 7 *Problem class $\{\{\text{CFpoly}\}\}$ is closed under \cup, \cap and complement.*

5 Nondeterminism and cons-free programs

Nondeterministic programs allow expression evaluation and program runs to be relations (\rightsquigarrow) rather than functions. A nondeterministic version of Figure 1 would use $\llbracket p \rrbracket(x) \rightsquigarrow v$ and $p, \rho \vdash e \rightsquigarrow v$ in place of $\llbracket p \rrbracket(x) \rightarrow v$ and $p, \rho \vdash e \rightarrow v$. Alas, the algorithm of Section 1.3 cannot be used if p is nondeterministic.

Definition 10 *A set $X \subseteq \{0, 1\}^*$ is decided by an NCF program p if for all inputs $x \in \{0, 1\}^*$,*

$$x \in X \text{ if and only if } p \text{ has a computation } \llbracket p \rrbracket(x) \rightsquigarrow \text{True}$$

5.1 Nondeterministic versions of the CF program classes

Remark: the reasoning used in Theorem 3 clearly extends to show that $\text{NLOGSPACE} = \{\{\text{NCFTR}\}\}$. The following is particularly interesting since the question $\text{LOGSPACE} =? \text{NLOGSPACE}$ is still open.

Theorem 5 $\{\{\text{CF}\}\} = \{\{\text{NCF}\}\}$

¹²A call is linear if it is not contained in a call to any defined function. An example with a linear call that is not a tail call: $\text{not}(\text{even}(\text{tail } z))$ in Section 1.5.

This was proven by Bonfante [2] by a technique that stems back to in Cook [3]. The implication is that $\{\{CF\}\}$ is closed under nondeterminism, since both problem classes are equal to PTIME. (This does *not* imply PTIME = NPTIME, since it is not hard to see that any NCF program can be simulated by a deterministic polynomial-time Turing machine: Lemma 3 holds for NCF as well as for CF.)

However it is not known whether $\{\{CFpoly\}\}$ is closed under nondeterminism, since Bonfante and Cook’s reasoning does not seem to apply. Why? The memoisation used in [2] yields a *Turing machine* polynomial time algorithm. Consequently, the problem is decidable by some CF program; but we know no way to reduce its time usage from exponential to polynomial.

A “devil’s advocate” remark: by Theorem 2 the classes $\{\{CFpoly\}\}$ and NLOGSPACE are both between LOGSPACE and PTIME:

$$\text{LOGSPACE} = \{\{CFTR\}\} \subseteq \text{NLOGSPACE} \subseteq \{\{CF\}\} = \text{PTIME}$$

So why bother with yet another class? One answer: CFpoly seems more natural than NCFTR from a programming viewpoint; and intuitively $\{\{CFpoly\}\}$ seems to be a larger extension of LOGSPACE than NLOGSPACE or the program class NCFTR. However we have no proof that $\text{NLOGSPACE} \neq \{\{CFpoly\}\}$, and no solid reason to think that either class contains the other.

5.2 Simulating CFpoly by a nondeterministic algorithm

Theorem 6 $\{\{CFPoly\}\} \subseteq \text{NSPACE}(\log^2 n)$.

Proof: If $p \in \text{CFPoly}$, there is a polynomial π such that for any input $x \in \{0, 1\}^*$ one can decide, in time $\pi(|x|)$, whether or not $\llbracket p \rrbracket(x) \rightarrow v$. The question: can this be done in significantly less space? We answer “yes”. The proof uses a space-economical nondeterministic algorithm to find v at the root $\llbracket p \rrbracket(x) = v$ of tree $\mathcal{T}^{p,x}$. First, observe that any reachable statement $p, \rho \vdash e \rightarrow w$ can be represented in $O(\log n)$ bits.

- To evaluate p on input x , we nondeterministically *guess* a value v such that $\llbracket p \rrbracket(x) = v$. The remaining task is to *confirm* that this statement is true. If we cannot do that, the whole algorithm has failed.
- To *confirm* a statement such as $p, \rho \vdash f e \rightarrow v$, we must confirm the existence of an evaluation tree of the form:

$$\frac{\frac{\dots}{p, \rho \vdash e \rightarrow w} \quad \frac{\dots}{p, [x_1 \mapsto w] \vdash e^f \rightarrow v}}{p, \rho \vdash f e \rightarrow v}$$

To do this, we now *guess* $w \in V_x$, and now need to confirm two statements. We also *guess* which of these two statements has the **shortest** evaluation tree.

For example, suppose we guess that this is the case for $p, [x_1 \mapsto w] \vdash e^f \rightarrow v$. Then we do a recursive call to confirm this statement (which temporarily stores the current state on the algorithm’s stack). Afterwards, we *tail-recursively* confirm the other statement, $p, \rho \vdash e \rightarrow w$. Since this is a tail call, the algorithm’s stack size does not increase.

- This extends naturally to multi-argument function calls.
- To confirm a statement $p, \rho \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rightarrow v$, we *guess* whether e_1 reduces to True or False; for example, we guess that it reduces to False. Then it suffices to confirm that an evaluation tree of the following form exists:

$$\frac{\frac{\dots}{p, \rho \vdash e_1 \rightarrow \text{False}} \quad \frac{\dots}{p, \rho \vdash e_3 \rightarrow v}}{p, \rho \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rightarrow v}$$

For this, we again have to confirm two statements. As before, we *guess* which of the two statements has the shorter evaluation tree, evaluate that statement first, and then evaluate the other statement tail-recursively.

If all guessed values are correct, then this algorithm returns the correct result. Furthermore, if we always guessed correctly which statement has the shortest evaluation tree, it does so with an overall stack depth that is never more than $\log(|\mathcal{T}^{p,x}|)$. The reason is that every time we do a non-tail recursive call, this shortest subtree has size at most half of the previous subtree's size. Since a statement $p, \rho \vdash e \rightarrow w$ can be represented in $O(\log n)$ bits we have the desired result.

The idea of applying tail-recursion to the largest subtree to save space was also used to implement Quicksort [9]. However our problem is much more general; and we must use nondeterminism because it cannot be known in advance which subtree will be the largest.

We expect that this construction can be extended to show $\{\{\text{NCFpoly}\}\} \subseteq \text{NSPACE}(\log^2 n)$ as well.

5.3 Closure under complementation

The following was shown by a subtle nondeterministic state-enumeration and state-counting algorithm. It was devised independently by Immerman and Szelepcsényi, and solved Kuroda's second and long-standing open question [11, 5, 16]. It is still open whether $\text{LOGSPACE} \subsetneq \text{NLOGSPACE}$.

Theorem 7 *NLOGSPACE is closed under complement.*

6 Conclusions, future work

We have probed the question $\text{LOGSPACE} =? \text{PTIME}$ from a programming language viewpoint. A starting point was that the "cons-free" programming language CF exactly captures the Turing complexity class PTIME; while its cons-free tail recursive subset CFTR exactly captures LOGSPACE. Section 3 recapitulates the reasoning used in [8].

In more detail: all CFTR programs run in polynomial time; but on the other hand, some CF programs run for exponentially many steps. Further the sets decided by the two program classes have seemingly different closure properties: the questions $\text{LOGSPACE} =? \text{NLOGSPACE}$ and $\text{PTIME} =? \text{NPTIME}$ and even $\text{LOGSPACE} =? \text{PTIME}$ have been open for decades. Given this, it seems almost paradoxical that $\{\{\text{CF}\}\} = \{\{\text{NCF}\}\}$ (from [2]).

Trying to understand these differences made it natural to consider CFpoly - the class of polynomially time-bounded CF programs since they have feasible running times (even though some CF programs may have superpolynomial behavior). One test of CFpoly was to see whether it contained any PTIME-complete problems. As a case study we wrote (Section 2.3) a CF-program for MCV (the monotone circuit value) problem to clarify where non-tail-recursion was necessary, and where superpolynomial runtimes came into the picture. (One key was nested recursion in function calls, clearly visible in function `mcv` in the MCV code.) Another test was to see whether $\{\{\text{CFpoly}\}\}$ is perhaps a smaller complexity class than PTIME. Theorem 6 leads in this direction, with an interesting proof construction and the surprising upper bound $\text{NSPACE}(\log^2 n)$.

Many questions for CFpoly are still to be investigated. One is to see whether the Immerman-Szelepcsényi algorithm can be adapted to NCFpoly.

References

- [1] Siddharth Bhaskar & Jakob G. Simonsen (2020): *Implicit Complexity via Controlled Construction and Destruction*. Technical Report, Department of Computer Science, University of Copenhagen.
- [2] Guillaume Bonfante (2006): *Some Programming Languages for Logspace and Ptime*. In Michael Johnson & Varmo Vene, editors: *Algebraic Methodology and Software Technology, 11th International Conference, AMAST 2006, Kuressaare, Estonia, July 5-8, 2006, Proceedings, Lecture Notes in Computer Science 4019*, Springer, pp. 66–80, doi:10.1007/11784180_8.
- [3] Stephen A. Cook (1971): *Characterizations of Pushdown Machines in Terms of Time-Bounded Computers*. *J. ACM* 18(1), pp. 4–18, doi:10.1145/321623.321625.
- [4] Sheila A. Greibach (1975): *Theory of Program Structures: Schemes, Semantics, Verification*. *Lecture Notes in Computer Science* 36, Springer, doi:10.1007/BFb0023017.
- [5] Neil Immerman (1988): *Nondeterministic Space is Closed Under Complementation*. *SIAM J. Comput.* 17(5), pp. 935–938, doi:10.1137/0217058.
- [6] Neil D. Jones (1997): *Computability and complexity - from a programming perspective*. Foundations of computing series, MIT Press, doi:10.7551/mitpress/2003.001.0001.
- [7] Neil D. Jones (1999): *LOGSPACE and PTIME Characterized by Programming Languages*. *Theor. Comput. Sci.* 228(1-2), pp. 151–174, doi:10.1016/S0304-3975(98)00357-0.
- [8] Neil D. Jones (2001): *The expressive power of higher-order types or, life without CONS*. *J. Funct. Program.* 11(1), pp. 55–94, doi:10.1017/S0956796800003889. Available at <http://journals.cambridge.org/action/displayAbstract?aid=68581>.
- [9] Donald E. Knuth (1973): *The Art of Computer Programming, Volume III: Sorting and Searching*. Addison-Wesley.
- [10] Cynthia Kop & Jakob Grue Simonsen (2017): *The Power of Non-determinism in Higher-Order Implicit Complexity - Characterising Complexity Classes Using Non-deterministic Cons-Free Programming*. In Hongseok Yang, editor: *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Lecture Notes in Computer Science 10201*, Springer, pp. 668–695, doi:10.1007/978-3-662-54434-1_25.
- [11] Sige-Yuki Kuroda (1964): *Classes of languages and linear-bounded automata*. *Information and Control* 7(2), pp. 207–223, doi:10.1016/S0019-9958(64)90120-2.
- [12] John McCarthy (1960): *Recursive Functions of Symbolic Expressions and Their Computation by Machine, Part I*. *Commun. ACM* 3(4), pp. 184–195, doi:10.1145/367177.367199.
- [13] Yiannis Moschovakis (2019): *Abstract Recursion and Intrinsic Complexity*. Cambridge University Press (Lecture Notes in Logic), doi:10.1017/9781108234238.
- [14] Christos H. Papadimitriou (1994): *Computational complexity*. Addison-Wesley.
- [15] Walter J. Savitch (1970): *Relationships Between Nondeterministic and Deterministic Tape Complexities*. *J. Comput. Syst. Sci.* 4(2), pp. 177–192, doi:10.1016/S0022-0000(70)80006-X.
- [16] Róbert Szelepcsényi (1988): *The Method of Forced Enumeration for Nondeterministic Automata*. *Acta Inf.* 26(3), pp. 279–284, doi:10.1007/BF00299636.
- [17] Alan M. Turing (1936-7): *On Computable Numbers with an Application to the Entscheidungsproblem*. *Proceedings of the London Mathematical Society* 42(2), pp. 230–265, doi:10.1112/plms/s2-42.1.230.