

The Expressiveness of a Family of Finite Set Languages

Neil Immerman*

Sushant Patnaik*

David Stemple†

Computer and Information Science Department
University of Massachusetts
Amherst, MA 01003

To appear in *Theoretical Computer Science*

Abstract

In this paper we characterise exactly the complexity of a set based database language called *SRL*, which presents a unified framework for queries and updates. By imposing simple syntactic restrictions on it, we are able to express exactly the classes, *P* and *LOGSPACE*. We also discuss the role of ordering in database query languages and show that the *hom* operator of *Machiavelli* language in [OBB89] does not capture all the order-independent properties.

1. Introduction

The expressiveness and complexity of database query and transaction languages are of interest for a number of reasons. Since the size of inputs to expressions in these languages is often very large, controlling the expressiveness of a language can be used to reduce the number of intractable queries posed by naive users, a major clientele of query languages. In addition, powerful optimization techniques are easier to develop

*Research supported by NSF grant CCR-9008416

†Research supported by NSF grants IRI-8606424 and IRI-8822121 and ONR U.R.I. grant N00014-86-K-0764.

and apply to limited languages than to more general languages. It is also often easier to reason formally about limited languages than about more general languages, though it is sometimes hard to isolate the difficulties stemming from the superficial diversity of a language, i. e., a large number of ways of expressing the same computation, from those due to its computational complexity. Our motivation includes the first two reasons, but is also strongly concerned with the third - the tractability of reasoning about finite set computations. While such tractability can be useful in optimizing queries and transactions, it can also be used to assure the quality of systems, for example, in terms of consistency maintenance over transactions.

Here we address the expressiveness of languages for specifying computations over finite sets. The family of languages we consider has very few primitives and its semantics are expressed by a small set of algebraic axioms. It does not start with first order logic and set theory, nor rely on concepts of destructive update or random access memory (or the related concepts of object ids and *ref* types). Unlike many approaches to computing with finite sets, it is designed to be seamlessly combined with other algebraic computational models such as ordinary arithmetic or recursive data type algebras. One of our goals is to be able to reason effectively about the complexity *and other properties* of computations over combined algebras, including finite set algebra.

Our logic base is simply the *if-then-else* function. First order logic is included in our computational model as a result of combining set traversal and the *if-then-else* operator. Our framework includes simple *tuple algebra*, expressing the ability to *construct* typed, fixed arity, non-recursive tuples and to *select* their components. Set traversal is expressed using a single mechanism, the higher order function *set-reduce*, which applies functions as it traverses a set. This is the sole iterative construct, and it can depend on the order of traversal. Its formalization in algebraic axioms makes the order of elements in a set manifest and allows order dependence to be reasoned about using mechanical reasoning capabilities, which are based on Boyer-Moore computational logic ([SS89]). In this way, we can often prove that the result of a computation is order independent even though the ordering is implicitly used as we traverse a set. This allows a new approach to a question raised by Chandra and Harel as to whether there is a language that expresses exactly the polynomial-time, order independent queries. All previous research on polynomial time queries has chosen either to deal with languages that express order dependent queries, or languages in which certain simple order independent queries cannot be expressed.

There have been numerous studies of the expressive power of query languages. For instance, it is well known that first order relational query languages are limited in their expressibility [AU79]. However when augmented with recursion or looping (as an added primitive) they become sufficiently powerful to express exactly the queries

in various complexity classes. Characterizing the expressive power of such languages has been the principal object of study in [Va82], [CH80], [CH82a], [CH82b], [AU79], [Imm82], [AV89]. For example, Immerman and Vardi discovered independently that fixpoint logic plus ordering expresses the set of polynomial time computable queries [Va82], [Imm82].

A common and rather useful way of measuring expressiveness is to use complexity characterizations. One finds classes of queries capturing *LOGSPACE*, *P*, *PSPACE*, *PRIMREC*. Interestingly, most of the classes of queries considered turned out to capture some complexity class. It seems that certain query language comparisons are connected with deep problems of complexity theory. Recently parallel evaluation of recursive queries has also drawn considerable attention in [CK85], [AC89].

In the past the emphasis has been to develop a *natural* set of primitives for a query language so that it can compute all the *computable* queries as in [CH80], [Ch81], [AV88]. Unbounded arity relations or the ability to create new values have been used. For example, Chandra and Harel, in [CH80], define the concept of *computable* queries and present a *complete* database programming language and show that relational algebra augmented with the power of iteration is *complete*.¹ In [HS89b], Hull and Su consider a hierarchy of languages whose complexity is in the super exponential range. However, we are interested in devising a *natural* language whose complexity is clear from the syntax but for *feasible* complexity classes from a database point of view, e.g. *TIME*[n] and *SPACE*[$\log^k n$]. Instead of *computable queries*, we regard primitive recursive queries as the high end of the spectrum. Indeed, all of the interesting complexity classes are contained in *PrimRec*. Our measure of complexity is data-complexity as defined by Vardi [Va82].

The *set-reduce* construct (defined in section 2) can be thought of as a bounded loop primitive. See [SS89] for more details. The *set-reduce* construct resembles the *hom* operator of the database programming language called *Machiavelli* [OBB89]. We define the transaction language, *unrestricted SRL* and show that its corresponding query language captures the primitive recursive properties. We then show that natural restrictions of this language capture *P*, *DSPACE*($\log n$) and *NSPACE*($\log n$).

The expressive power of the bounded loop construct or its variant has been studied before in [AU79], [Va82], [Q89], [CH80], [Imm87], [AV88] but not in this framework. In [Ch81], Chandra raises the question of specifying a set of primitives of the form *forall tuples t in relation R do statement S*, where *S* is restricted so as not to use the order in which the *forall* cycles over all the tuples, such that programs in this style can compute all the computable queries. The *set-reduce* construct provides a partial

¹Here a complete database language is one that can compute every partial recursive function of its database [CH80].

answer.

In [AV88], Abiteboul and Vianu present declarative and procedural update languages and show that they are *complete* in Chandra and Harel's sense. They also define restrictions on their languages and characterize their expressiveness. They define the so-called "*non-deterministic*" updates and show certain languages to be "*non-deterministic*" update complete. Their definition of "*non-deterministic*" updates is actually what we refer to as *order-dependent*. It should not be confused as such with *non-determinism* as referred to in complexity theory. In [Q89], Qian studies the complexity of a bounded looping construct *foreach x in R/p(x) do t(x)* and shows that, under deterministic semantics, her language and a subclass of it have polynomial time and first order expressive power. Her looping construct closely resembles the *set-reduce* operator, however the two corresponding languages differ in their algebras. Her definition of "*non-deterministic*" semantics, identical to Abiteboul and Vianu's, does not lead to a decidable distinction between non-deterministic and deterministic languages. In a recent paper [AK90], Abiteboul and Kanellakis discuss an object oriented database programming language wherein objects are built by applying *set* and *tuple* constructors. They define an algebra for their language which is built from first-order operators augmented with the *powerset* operator. The latter immediately puts the data-complexity of their language in the exponential range. Had they instead defined a bounded iterator operator, as we do, then it would have been possible

***** iteration

to derive sub-languages whose complexities lie between first-order and exponential. In [Gu83], Gurevich characterized the complexity of functions defined using primitive recursion over structures with finite domains. He showed that different versions of the language (with bounded successor) capture functions in polynomial time and logspace.

This paper is organized as follows. Section 2 defines the *set-reduce* language and gives some background. Section 3 presents some tools of descriptive complexity [Imm87] and proves that *SRL* (with set-height at most 1) = *P*. Section 4 describes ways of restricting the complexity of *SRL*. Section 5 shows that *unrestricted SRL* with sets of unbounded width (or, equivalently, *SRL* with invented values) captures the class of primitive recursive functions. Section 6 shows how to deduce the complexity of a *SRL* program from its syntax. Section 7 discusses the role of ordering in database query languages. Section 8 concludes with some comments and open questions.

2. The language

Set-reduce language is a typed language of finite expressions constructed and typed according to the following rules. We assume a set of base types (with equality) that includes the booleans, tuples (i.e., records without explicit attribute names) and sets (without equality). Further, most of the paper can be viewed as assuming, besides the booleans, a single base type with a finite domain.

1. *true* and *false* of type *boolean* are *set-reduce* expressions.
2. *if srebool then sre₁ else sre₂*, where *srebool* is an *set-reduce* expression of type *boolean*, and *sre₁* and *sre₂* are *set-reduce* expressions of the same type, is a *set-reduce* expression of the same type as *sre₁* and *sre₂*.
3. constants of type *T* where *T* includes an equality relation are *set-reduce* expressions of type *T*.
4. [*sre₁*, ..., *sre_n*] of type *tuple(sel₁ : T₁, ..., sel_n : T_n)*,
 *****XXXXXXXXXXXXXXXXX where *T_i* is the type of *sre_i*,
 is a *set-reduce* expression of type *tuple(sel₁ : T₁, ..., sel_n : T_n)*.
5. *sel_i(sre)* where *sre* is of type *tuple(sel₁ : T₁, ..., sel_n : T_n)* is a *set-reduce* expression of type *T_i*.
6. *sre₁ = sre₂*, where *sre₁* and *sre₂* *set-reduce* expressions of the same type, is a *set-reduce* expression of type *boolean*.
7. *emptyset* is a *set-reduce* expression of type *set(alpha)*, where alpha matches any type.
8. *insert(e, s)*, for *s*, a *set-reduce* expression of type *set(T)* and *e*, of type *T*, is a *set-reduce* expression of type *set(T)*.
9. *set-reduce (s, app, acc, base, extra)* is a *set-reduce* expression of type *T'*, where *s*, *base* and *extra* are *set-reduce* expressions of types *set(T)*, *T'* and *extype*, respectively, and *app* and *acc* are formed by appending *lambda(x, y)* to *set-reduce* expressions, in which only *x* and *y* can appear free. The variables *x* and *y* in *app* must appear in places appropriate for *set-reduce* expressions of types *T* and *extype*; and in *acc* in places appropriate for *T* and *T'*, respectively. The typing of lambda expressions follows the normal type inference rules for lambda expression applications.

10. (*srlexp*) is a *set-reduce* expression if *srlexp* is a *set-reduce* expression, and it has the same type as *srlexp*.

This defines a ground language in which no free variables occur. In the sequel, we use expressions in the inductive language, i. e., with free variables, to stand for ground expressions in which the bindings to the free variables are not shown.

Note that equality is axiomatized only for the base types. For other types, it has to be expressed using the *set-reduce* operation. The *extra* template in the *set-reduce* operator is for the purpose of passing extra variables into functions so that all reference is local; the use of extra parameters makes nested variable scoping and currying unnecessary while allowing succinct expression of a wide variety of functions and algorithms [SS89]. The input to any *set-reduce* expression is a structure or database specified by the name(s) of set(s) or relation(s).

The semantics of *set-reduce* expressions is given by the following rules and equations for which there is a straightforward reduction mechanism that is complete for deciding equality of ground terms.

Boolean

$$(if\ true\ then\ e_1\ else\ e_2) = e_1$$

$$(if\ false\ then\ e_1\ else\ e_2) = e_2$$

Other types

The equality of constants of types other than boolean, tuple and set is defined by the types.

Tuples

Tuple construction is a function. For tuples t and t' of type $tuple(sel_1 : T_1, \dots, sel_n : T_n)$,

e_i typed T_i for $i = 1$ to n , $t = [e_1, \dots, e_i, \dots, e_n]$ and $t' = [e'_1, \dots, e'_i, \dots, e'_n]$

$$(t = t') \rightarrow (e_i = e'_i) \text{ for } i = 1 \text{ to } n$$

$$sel_i(t) = e_i \text{ for } i = 1 \text{ to } n$$

Finite sets

The interested reader is referred to [SS89] for a formal specification of the semantics of *emptyset*, *insert*, *choose* and *rest* (the latter two used in the semantics of *set-reduce*). Here, we point out the salient features. We assume that each base type is equipped with an ordering (and hence, every type has an ordering induced by the

latter). For any non-empty set, S , of type T , $choose(S)$ is defined to return the minimal (in the ordering of T) element in S , and $rest(S)$ returns $S/choose(S)$.

$$\begin{aligned} \text{set-reduce}(s, app, acc, base, extra) = \\ \text{if } s = \text{emptyset} \\ \text{then } base \\ \text{else } acc(app(choose(s), extra), \text{set-reduce}(rest(s), app, acc, base, extra)) \end{aligned}$$

The semantics of lambda expressions are given by straightforward reduction rules.

Parentheses

$$(srlexp) = srlexp$$

The above specifies a many-sorted signature and a set of axioms. The axioms for finite sets specify the existence of a total order on the domain type of a finite set type. Any model (algebra) of the specification must supply an order. Although an ordering is assumed to exist on types, an order symbol is not included in the language, its only use is through the *set-reduce* construct. The results may still be order-independent. (In our use of this specification a user of the type does not supply an order. The implementation supplies the order. A user may observe the order, but may not encode any information in it, nor depend on it in any way other than on its existence.)

In the following definition, we consider the expressiveness of *set-reduce language* over structures with a finite domain. We consider functions in our framework to accept structures (databases) as input and return structures as output.

Definition 2.1 The class of *set-reduce functions* is the smallest class of functions computed by such *set-reduce* expressions and closed under *composition* and *set-reduce* operations. We denote it as $\mathcal{F}(u\text{-SRL})$. We denote the class of decision problems (boolean valued functions) expressible in this language as $\mathcal{L}(u\text{-SRL})$.

Note that boolean *and*, *or*, and *not* can easily be defined with the *if-then-else* function. Also, note that we have made available to us an ordering relation (denoted by \leq) on the domain which is the same order in which the elements are scanned by the *choose* mechanism of *set-reduce* ([SS89]). This is quite natural, since any computation must use an ordering. See Section 7 for a discussion of the ordering.

We believe that the *set-reduce* framework can provide a suitable base on which algebraic specifications of computations over databases can be analysed for purposes of assuring correct behavior and achieving optimized implementations. We wish to

limit the expressiveness of *unrestricted SRL* to within *reasonable* complexity classes so as to make the latter task feasible. We impose syntactic restrictions on *unrestricted SRL* and study their effect on its expressive power.

Definition 2.2 We define *set-height*(α) as follows:

$$\begin{aligned} \text{set-height}(\text{base-type}) &= 0 \\ \text{set-height}(\text{set of } \alpha) &= 1 + \text{set-height}(\alpha) \end{aligned}$$

As a first step, we restrict the use of *set types* to those with *set-height* at most 1, but allow arbitrary, though fixed (for any given expression), nesting and width (or arity) of *tuple types*. Let us denote this restricted version of the language as *SRL*.

Definition 2.3 We define the class of decision problems (respectively, functions) expressible in *SRL* as $\mathcal{L}(\text{SRL})$ ($\mathcal{F}(\text{SRL})$).

Functions expressible in *SRL* are similar to Cobham's recursive functions [Co64] and we show that the two classes are indeed equivalent. The restrictions on *set-height* and *tuple-width* are, as shown later, quite crucial to our result. To get started, we use the following fact:

Fact 2.4 ([SS89]) *Finite set functions such as union, intersection, difference, membership; predicates for universal and existential quantification such as forall, forsome; and relational operators such as join, project and select can be expressed in SRL.*

3. Expressiveness of SRL

Our approach to characterizing the expressive complexity of *SRL* follows the conventions of descriptive complexity [Imm87]. We will code all inputs as finite logical structures. The universe of structure is $\{0, 1, \dots, n-1\}$ and is denoted by D . We assume that the ordering on the universe (or, synonymously, domain) is given by the one on natural numbers. A vocabulary $\tau = \langle R_1^{a_1}, R_2^{a_2}, \dots, R_s^{a_s} \rangle$ is a tuple of input relation symbols of fixed arities. Let $STRUCT[\tau]$ denote the set of all finite structures of vocabulary τ . We will think of all complexity theoretic problems as subsets of $STRUCT[\tau]$ for some τ . The advantage of this approach is that when we consider our inputs as first order structures we may write properties of them in variants of first-order logic.

For any vocabulary τ , there is a corresponding first-order language $L(\tau)$ built up from the relation symbols of τ and the logical relation symbols and constant symbols $: =, \leq, 0, n-1$, using logical connectives $: \vee, \wedge, \neg$, variables $: x, y, z, \dots$, and quantifiers $: \forall, \exists$. Let FO be the set of first-order definable problems:

$$FO = \{S | (\exists \tau)(\exists \varphi \in L(\tau)) S \in STRUCT[\tau] \models \varphi\}.$$

Let us recall the definition of *first-order interpretation* [IL89, Imm87]. We assume a bit-encoding for the relations in the definition. For example, $R(x, y)$ over $D = \{0, \dots, n-1\}$ is represented by a sequence of n^2 bits such that $R(x, y)$ is true iff $nx + y$ -th bit is 1.

Definition 3.1 Let $S \subset STRUCT[\sigma]$, $T \subset STRUCT[\tau]$ be two problems. For simplicity, assume that the vocabularies, σ, τ consist of single input relations, $\langle Q^a \rangle, \langle R^b \rangle$, of arity a and b , respectively.

Let $k \geq 1$ be a constant and let $\varphi(x_1, \dots, x_k)$ be a FO formula from $L(\sigma)$. Then φ defines a mapping $m_\varphi : STRUCT[\sigma] \rightarrow STRUCT[\tau]$. Let $A = \langle n, Q^a \rangle \in STRUCT[\sigma]$ be a string of length n^a . Then $m_\varphi(A) = \langle n^k, R^b \rangle$ is a string of length $n^{bk} = n^k$. Thus the bit numbered (in n -ary) $j_1 j_2 \dots j_{ki}$ is 1 iff $A \models \varphi(j_1, j_2, \dots, j_{ki})$. If for all A in $STRUCT[\sigma]$,

$$A \in S \leftrightarrow m_\varphi(A) \in T$$

then m_φ is a k -ary *first-order interpretation* of S to T and we write $S \leq_{fo} T$ if such an interpretation exists.

In other words, any relation of $T \in STRUCT(\tau)$ is defined by a formula in first order logic over the relations of $S \in STRUCT(\sigma)$. We refer the reader to [IL89] for further details and examples of such reductions.

Definition 3.2 A class C is closed under FO interpretations if for any problem $T \subset STRUCT[\tau]$ in C and for any problem $S \subset STRUCT[\sigma]$, $S \leq_{fo} T$ implies that S is in C .

Let P denote the class of decision problems recognizable by deterministic Turing machines in time polynomial in the length of the input. To prove that $\mathcal{L}(SRL)$ contains P we will show that $\mathcal{L}(SRL)$ is closed under FO interpretations and that it contains a problem that is complete for P via FO interpretations.

Proposition 3.3 $\mathcal{L}(SRL)$ is closed under FO interpretations.

Proof: If $\mathcal{L}(SRL)$ is closed under boolean and quantifier operators, and there is a FO-translation from S to T and T is in $\mathcal{L}(SRL)$, then since $\mathcal{L}(SRL)$ is closed under composition, S is in $\mathcal{L}(SRL)$, i.e., there is a boolean-valued function expressible in SRL that is true precisely on S .

Closure under boolean operations follows from the definition of SRL . Closure under quantification is implicit in 2.4. Thus, for example to see that $\mathcal{L}(SRL)$ is closed under universal quantification, let $\varphi_1(\bar{R}, y)$ be a boolean function $\in SRL$, where y is a variable ranging over domain D and \bar{R} are relation names, and let $\varphi(\bar{R})$ be the first-order logic formula, $\forall y \varphi_1(\bar{R}, y)$. Then, φ can be expressed in SRL as :

$$\varphi(\bar{R}) = \text{set-reduce}(D, \lambda(d, e)\varphi_1(e, d), \wedge, \text{true}, \bar{R})$$

The existential quantifier case is handled similarly. ■

Definition 3.4 Let an alternating graph $G = (V, E, A)$ be a directed graph whose vertices are labeled universal or existential. Let $APATH(x, y)$ be the smallest relation on vertices of G such that

1. $APATH(x, x)$,
2. If x is existential (i.e. $\neg A(x)$) and for some edge (x, z) $APATH(z, y)$ holds then $APATH(x, y)$,
3. If x is universal (i.e. $A(x)$) there is at least one edge leaving x and for all edges (x, z) $APATH(z, y)$ holds then $APATH(x, y)$.

Let $AGAP = \{G | APATH(V_0, V_{max})\}$.

The \leq predicate, included in the base functions of SRL is crucial to the following result.

Fact 3.5 ([Imm87]) *AGAP is complete for P under first-order reductions.*

Consider the following monotone operator Γ [Imm87]:-

$$\Gamma(R)[x, y] \equiv (x = y) \vee [(\exists z)(E(x, z) \wedge R(z, y)) \wedge (A(x) \rightarrow ((\forall z)E(x, z) \rightarrow R(z, y)))]$$

It is easy to see that $LFP(\Gamma) = APATH$. We show that it is possible to express $AGAP$ as a function in SRL in the following lemma:

Lemma 3.6 *APATH is expressible in SRL.*

Proof: We shall specify the types in our *SRL* function for *APATH* only at the beginning and then use variables without mentioning types to enhance the readability. We shall use Fact 2.4 extensively.

Let *NODES* of type $set(V)$ and *EDGES* of type $set([from, to : V, label : \{AND, OR\}])$ be the input. Thus, the set of *AND* and *OR* labeled vertices can be obtained as follows:

$$ANDS = project(select(EDGES, \lambda(x)x.label = AND), from)$$

$$ORS = project(select(EDGES, \lambda(x)x.label = OR), from).$$

We can write Γ in *SRL* easily and then simulate the least fixed point operator on Γ which is of arity 2, by writing a loop which runs n^2 times.

$$\begin{aligned} \Gamma(x, y, R) = & (x = y) \vee (forsome(NODES, \lambda(z)(member([z, y], R) \wedge \\ & \hspace{10em} member([x, z], EDGES))) \\ & \wedge (\neg(member(x, ANDS)) \vee \\ & forall(NODES, \lambda(z)(\neg(member([x, z], EDGES)) \vee \\ & \hspace{10em} member([z, y], R)))))) \end{aligned}$$

$$\begin{aligned} \Gamma(R) = & set-reduce(NODES, \\ & \lambda(d_1, S)(set-reduce(NODES, \lambda(d_2, e)([e, d_2]), \\ & \hspace{10em} \lambda(t, T)(if \neg(member([t.1, t.2], T)) \wedge \Gamma(t.1, t.2, T) \\ & \hspace{10em} then insert([t.1, t.2], T) \\ & \hspace{10em} else T), \\ & \hspace{10em} S, \\ & \hspace{10em} d_1)), \\ & union, \\ & R) \end{aligned}$$

$LFP_\Gamma = ITERATE()$ where

$$\begin{aligned} ITERATE() = & set-reduce(NODES, identity, \\ & \lambda(d, Z)(set-reduce(NODES, identity, \lambda(d, X)\Gamma(X), Z)), \\ & \{\}) \end{aligned}$$

■

Corollary 3.7 $P \subseteq \mathcal{L}(SRL)$.

Proof: Since *AGAP* is complete for P under *FO* reductions (by Fact 3.5), and it is expressible in *SRL* (by Lemma 3.6), and $\mathcal{L}(SRL)$ is closed under *FO* reductions (by Proposition 3.3), it follows that $P \subseteq \mathcal{L}(SRL)$. ■

Since we have defined *SRL* so that *set-height* is at most 1 and *tuple nesting and width* are constant, we have that

Proposition 3.8 Let l be the tuple nesting and w be the tuple width of a tuple type α . Let S be of type set of α and let n be the number of elements in the input domain D . Then, $|S| \in O(n^w)$.

Proof: The maximum size of any set S that can be formed is equal to the number of possible tuples of width w and nesting l which is easily seen to be $w^{l-1}n^w \in O(n^w)$. ■

It now follows that

Lemma 3.9 $\mathcal{L}(SRL) \subseteq P$.

Proof: We define depth, d , of a *set-reduce function*, recursively as follows.

$$\begin{aligned} \text{depth}(\text{base functions}) &= 0 \\ \text{depth}(\text{set-reduce}(S, \text{appf}, \text{accf}, b, e)) &= 1 + \max(\text{depth}(S), \text{depth}(\text{appf}), \text{depth}(\text{accf}), \\ &\quad \text{depth}(b), \text{depth}(e)) \end{aligned}$$

We show by induction on d that each function F in *SRL* can be computed in time polynomial in n and therefore produces sets of polynomial size.

Base case: $d = 0$. The base functions can clearly be computed in P . Only *insert* increases the set-size by 1.

Inductive Step: Any function in *SRL* is of the form

$$F(S, e) = \text{set-reduce}(S, \text{appf}, \text{accf}, b, e)$$

By the inductive hypothesis, *accf*, *appf*, *base*, e and S can be computed in time $\leq n^k$, for some constant k . Thus, we have $|S|$ applications of *appf*, *accf* on inputs of size at most n^w by the proposition above. Total time to compute this recursion is

= the time to compute $accf, appf$ $|S| \leq n^w$ times on input of size n^w
 + time to compute $base$ + the time to compute e
 $\leq 2(n^w)(n^w)^k + n^k + n^k$
 $\leq n^{k'}$, for some constant $k' = w(k + 1) + 1$.

■

Theorem 3.10 $P = \mathcal{L}(SRL)$.

Proof: It follows from the Corollary 3.7 and Lemma 3.9 above.

■

Remarks:

- It is possible to show that $DTIME(n^k) \subseteq SRL$ by directly simulating the Turing machine computation. Refer to section 6 where we give tighter bounds on the complexity of an SRL expression from its syntax.
- Let FP denote the class of functions computable in polynomial time. Then, it follows from the previous theorem that

Corollary 3.11 $\mathcal{F}(SRL) = FP$.

- Restricting to a single usage of *set-reduce* does not help to restrict the complexity. It still remains sufficiently powerful to express $AGAP$ and hence the whole of P .
- The restriction on *set-height* is crucial as the following example shows. With *set-height* 2, it is possible to express a function in SRL that constructs a set of size exponential in the size of the input set.

Example 3.12 Consider the following function, $powerset(S)$, which given a set S constructs the power set $P(S)$ of S . For example, $powerset(\{1, 2\})$ returns a set of sets, $\{\{\}, \{1\}, \{2\}, \{1, 2\}\}$.

$powerset(S) = set-reduce(S, identity, sift, \{\{\}\})$

Sift takes an element x and a set of sets T and calls *finset* to insert x in each one of the elements of T and returns T unioned with all these new sets containing x .

$$sift(x, T) = set-reduce(T, \lambda(y, e)([y, e]), finsert, \{\}, x)$$

Finsert takes as arguments x , a two tuple of a set and an element, and a set of sets, T and returns $T \cup \{x.1\} \cup \{x.2 \text{ inserted in } x.1\}$.

$$finsert(x, T) = insert(x.1, insert(insert(x.2, x.1), T))$$

Similarly, it can be shown that such a situation exists if we do not restrict the *tuple nesting*. In this case, we are able to represent a list, e.g. $\langle 1, 2, 3 \rangle$ as $[1, [2, [3, [-, -]]]]$. Then, T in the program above is typed as a set of tuples of width 2 and arbitrary nesting, and we redefine *finsert* as follows:

$$finsert(x, T) = insert(x.1, insert([x.2, x.1], T))$$

$$powerset(S) = set-reduce(S, identity, sift, \{[-, -]\})$$

In [SS89], a *list-reduce* construct is defined which is exactly the same as *set-reduce* except that the object we recurse over is a *list*, and *not* a *set*. The difference is that the items appear in a specific order in the list. Clearly any function realized using *set-reduce* can be implemented using *list-reduce* by simply substituting the former by the latter construct. Define *list-height* analogous to *set-height*. Let us denote *LRL* as the corresponding language with *list-reduce* replacing *set-reduce* as the recursion operator and with *list-height* ≤ 1 and $\mathcal{F}(LRL)$ as the functions that can be expressed in this language. As observed above, $\mathcal{F}(SRL) \subseteq \mathcal{F}(LRL)$. But $\mathcal{F}(LRL) \not\subseteq FP$. This can be seen from the following function which is not in P , but is readily seen to be in *LRL* viz. $F(\langle 1 \rangle, \langle 1, 2, \dots, n \rangle) = \langle 1, 1, \dots, (2^n \text{ times}), \dots, 1 \rangle$. Note that lists can be of arbitrary length in this language. In fact, we will see that $\mathcal{F}(LRL)$ exactly equals the class of *primitive recursive functions*.

The proof of Lemma 3.9 goes through as long as Proposition 3.8 is not violated. So what are the operators that can be included in *set-reduce language* such that it still remains within PT Clearly integers and bounded operators on them such as *addition mod x* and *multiplication mod x*, where x is an exponential function of n (the input size), can be added to *set-reduce language* without blowing up its complexity, since the size of such objects is bounded by $\log x$, which is polynomial in n . Let \mathbf{N} denote the set of natural numbers. Let *succ* denote the successor operator on the naturals. It is shown later that allowing objects of type N or integers, with the usual *succ* operator on such types, in *set-reduce language* increases the complexity to that of *primitive recursive functions*. In particular, if we allow addition or multiplication on integers and allow the type, *set of integer*, then we can express the class of primitive recursive functions in this language.

However, if we do not permit the use of the type, *set of N* or *set of integer*, then its complexity is still within P . For example, we can safely add integer types along

with addition (+) on integers in *set-reduce language* while still remaining within P provided we do not allow the type, *set of integer*. We can also add the operation of multiplication (*) on integers, if in addition to the previous restriction, we do not allow the accumulator function, *accf*, to use it. Clearly, addition and multiplication are in P and hence, their result is polynomial sized.

Let x and y be of some ordinal type α , like N . Let $\langle op \rangle$ be some operator such that if $x \langle op \rangle y$ is repeated n times recursively then size of the result is polynomial in the sizes of x and y , and in n . Let $a = x \langle op \rangle y$ and let $|x|$ denote the size or length of the binary representation of x . *SRL* is closed with respect to such an operator $\langle op \rangle$ defined on α , provided the type - *set of α* - is not permitted. An example would be some $\langle op \rangle$ such that $|x \langle op \rangle y|$ is an additive function of $|x|$ and $|y|$. But if $|x \langle op \rangle y|$ were a multiplicative function of $|x|$ and $|y|$, then we have to impose one further restriction - we do not allow *accf* in the set-reduce construct to use $\langle op \rangle$ unless one of the operands is a fixed constant. If op is + then $|a| \leq \max\{|x|, |y|\} + 1$, whereas if op is * then $|a| \leq 2 * \max\{|x|, |y|\}$. Thus, if one allows *accf* to use multiplication, then it is easy to compute x^{2^n} (which cannot be done in P) in *set-reduce language* by repeated squaring. However, we can allow *multiplication by a constant* inside *accf*, while still remaining within P , since in this case, the size of the result of n repeated multiplications by a *constant* is clearly polynomially (in fact, linearly) bounded. It can be easily observed that

Proposition 3.13 *Addition of other operators and functions to SRL will not take us out of P provided that the size of the sets we can build, using those operators, remains polynomially bounded.*

4. Restricted versions of SRL

Let L (NL) denote the class of problems recognized by deterministic (non-deterministic) Turing machines using space no more than logarithmic in the input size. It is well known that $L \subseteq NL \subseteq NC^2$. The question arises as to whether there exist any syntactic restrictions on *SRL* that in an elegant and natural way capture L and NL . Characterizing L and NL as some form of *SRL* would be interesting since problems in these classes are also efficiently parallelizable.

One way of doing this follows easily from the results in [Imm87]. We adopt the same notations. Let $\varphi(\bar{x}, \bar{x}')$ be any *FO* formula. Define $TC[\lambda\bar{x}, \bar{x}'\varphi]$ as the reflexive, transitive closure of the relation φ . Let $(FO + TC)$ be the set of properties expressible using first order logic plus the operator TC . The following characterization is well known:

Fact 4.1 ([Imm87, Imm88]) $NL = (FO + TC)$.

We define a new operator called TC , in SRL as follows. Let the set of vertices be D . $TC(\varphi)$ is computed as follows in SRL :

Let $EDGEp([x, y]) = \varphi(x, y)$, and $EDGES = select(join(D, D), \lambda([x, y])EDGEp([x, y]))$

The function $bothsides(v, EDGES)$ returns the pairs of vertices that are at distance 2 from each other, given the adjacencies specified by the current value of relation $EDGES$:

$$\{[x, y] \mid [x, v] \in EDGES \wedge [v, y] \in EDGES\}$$

$$bothsides(v, EDGES) = join(D, D, \lambda(t1, t2)((t1.to = v) \vee (t2.from = v)), \lambda(t1, t2)([t1.from, t2.to]))$$

Add simply updates $EDGES$ after every iteration.

$$add(v, E) = union(E, bothsides(v, E))$$

The transitive closure is obtained by simply iterating $bothsides$ $|D|$ times.

$$TC(EDGES) = set-reduce(project(EDGES, to), identity, \lambda(x, Y)(add(x, Y)), EDGES)$$

Let $SRFO + TC$ be the class of problems expressible in a subset of SRL that has only the following functions available: *forsome*, *forall*, \neg , \vee , \wedge , \leq , TC . As an immediate corollary to the preceding fact, we have that

Corollary 4.2 $SRFO + TC = NL$.

Proof: Clearly every property expressible in $FO + TC$ can be expressed in $SRFO + TC$ and vice versa. ■

Given a first order relation $\varphi(\bar{x}, \bar{x}')$, let

$$\varphi_d(\bar{x}, \bar{x}') \equiv \varphi(\bar{x}, \bar{x}') \wedge [(\forall \bar{z}) \neg \varphi(\bar{x}, \bar{z}) \vee (\bar{x}' = \bar{z})].$$

That is, $\varphi_d(\bar{x}, \bar{x}')$ is true just if \bar{x}' is the unique descendant of \bar{x} . Define $DTC(\varphi) \equiv TC(\varphi_d)$. Let $(FO + DT C)$ be the set of properties expressible using first order logic plus the operator $DT C$. Then, analogous to the NL case, it comes as no surprise that,

Fact 4.3 ([Imm87]) $L = (FO + DT C)$.

$DTC(\varphi)$ can be computed in SRL as follows:

$$\begin{aligned}\varphi_d(\bar{x}, \bar{y}) &= \varphi(\bar{x}, \bar{y}) \wedge \text{forall}(D, \lambda(z, e)(p(z, e)), [\bar{x}, \bar{y}]) \\ \text{where } p(\bar{z}, e) &= \neg(\varphi(e.1, \bar{z})) \vee (\text{equal}(e.2, \bar{z})) \\ DTC(\varphi) &= TC(\varphi_d).\end{aligned}$$

Let $SRFO + DTC$ be the class of problems expressible in a subset of SRL that has only the following functions available: *forsome*, *forall*, \neg , \vee , \wedge , \leq , DTC . Thus, we have the following easy corollary from Fact 4.3 that

Corollary 4.4 $L = SRFO + DTC$.

Another, perhaps more natural, way of characterizing L is achieved by considering the following restriction on SRL : we restrict the function *acc* in our *set-reduce* template to return just a tuple of bounded width (and set-height zero). Let us denote this version of SRL as $BASRL$ and the set of decision problems expressible in this version of SRL as $\mathcal{L}(BASRL)$. Then, we can show that the class L is exactly equal to $\mathcal{L}(BASRL)$ as follows. The proof is similar in form to that of P equals $\mathcal{L}(SRL)$. We need the following definitions. We treat the elements of D as numbers. This will enable us to perform arithmetic functions conveniently.

Let $BIT(i, x)$ denote the value of the i^{th} bit in the binary representation of x . In the context of SRL , since it only deals with sets and not numbers, we have to impart a meaningful interpretation to this operator. Assume the active domain of any SRL program is D and let $|D|$ denote the size of D and let $n = |D|$.

Note that we have a total order \leq on D which is the order in which the elements of D are scanned by *set-reduce*. We can either assume that it is available to us as a set of pairs say, $S = \{(a, b) | a \leq b\}$, or we can compute the successor or predecessor of an element with respect to \leq whenever we need it. Each element has a unique position in this ordering. Let d_1, d_2 be any elements in D , let i_1, i_2 be the ranks (positions) of d_1, d_2 in that total order. Then, $BIT(d_1, d_2) \equiv BIT(i_1, i_2)$. In a similar vein we define addition, multiplication, exponentiation. Let $d_1, d_2 \in D$ and let i_1, i_2 be their respective ranks in the ordering \leq . Then $d_1 + d_2$ is defined to be $d_3 \in D$ such that if i_3 is the rank of d_3 in \leq then $i_3 = i_1 + i_2$. Multiplication and exponentiation are likewise defined.

Proposition 4.5 *Addition, multiplication, exponentiation are expressible in $BASRL$.*

Proof: We show how to add 1 to a as follows. Starting with $[false, false, a]$, we iterate over D , changing the first *false* to *true* when we find a , changing the second *false* to *true* on the next step when we remember $a + 1$ and then, finally returning the triple, $[true, true, a + 1]$.

$$\begin{aligned} \text{increment}(a) = \text{set-reduce} & (D, \text{identity}, \\ & \lambda(d, X)(\text{if } \neg(X.1) \wedge (d = X.3) \\ & \quad \text{then } [true, false, X.3] \\ & \quad \text{else if } \neg(X.2) \wedge (X.1) \text{ then } [X.1, true, d] \\ & \quad \text{else } X), \\ & [false, false, a]) \end{aligned}$$

Similarly one can define $\text{decrement}(A)$. We have to take care of the boundary cases, $\text{increment}(n)$ and $\text{decrement}(0)$, appropriately. Then,

$$\begin{aligned} \text{ADD}(a, b) = \text{set-reduce}(D, \text{identity}, \\ & \lambda(d, X)(\text{if } \neg(X.1 = n) \wedge \neg(X.2 = 0) \\ & \quad \text{then } [\text{increment}(X.1).3, \text{decrement}(X.2).3] \\ & \quad \text{else if } (X.2 = 0) \text{ then } X \\ & \quad \text{else } [0, \text{decrement}(X.2).3]), \\ & [a, b]) \end{aligned}$$

Note that ADD returns a tuple, $[a + b, 0]$, while increment and decrement return a tuple of the form $([true, true, a'])$ and operators $.1$, $.2$ and $.3$ return the first, second and third component of the tuple, respectively. Multiplication is expressed as follows:

$$\begin{aligned} \text{MULT}(a, b) = \text{set-reduce}(D, \lambda(s, \text{extra})(\text{extra}), \\ & \lambda(e, X)(\text{if } (X.2 = 0) \text{ then } X \\ & \quad \text{else } [\text{ADD}(e, X.1).1, \text{decrement}(X.2).3]), \\ & [0, b], \\ & a) \end{aligned}$$

Note that we use 0 , n to simply mean the first and last elements respectively in \leq . It is readily checked that $x = 0$ or, $x = n$ can be expressed in BASRL by seeing whether x is the first or, last element of the ordering. Exponentiation can now be expressed as shown below:

$$\begin{aligned} \text{EXP}(a, b) = \text{set-reduce}(D, \lambda(s, x)(x), \\ & \lambda(x, T)(\text{if } (T.2 = 0) \text{ then } T \\ & \quad \text{else } [\text{MULT}(x, T.1).1, \text{decrement}(T.2).2]), \\ & [1, b], \\ & a) \end{aligned}$$

■

Lemma 4.6 *BIT is expressible in BASRL.*

Proof: We shall use the proposition above. First we show how to divide by 2 in *BASRL*:

$$\begin{aligned}
 SHIFT(a) = \text{set-reduce}(D, \text{identity}, \\
 \lambda(x, e)(& \text{if } \neg(e.1) \wedge ((ADD(x, x).1) = e.2) \\
 & \text{then } [true, x, false] \\
 & \text{else if } (\text{increment}(ADD(x, x).1).3 = e.2) \\
 & \text{then } [true, x, true] \\
 & \text{else } e), \\
 & [false, a, false])
 \end{aligned}$$

Note that we have also defined a predicate, *PARITY* of a number as *true iff number is odd*:

$$PARITY(x) = SHIFT(x).3$$

Finally, *BIT*(i, a), i.e. the i^{th} bit of a is given by the parity of a divided by 2^i as follows:

$$\begin{aligned}
 REM(i, a) = \text{set-reduce}(D, \text{identity}, \\
 \lambda(s, X)(& \text{if } \neg(X.1 = 0) \\
 & \text{then } [\text{decrement}(X.1).3, SHIFT(X.2).2] \\
 & \text{else } X), \\
 & [i, a])
 \end{aligned}$$

$$BIT(i, a) = PARITY(REM(i, a).2)$$

If a is a bounded width tuple of elements from D , then, *BIT* is interpreted with respect to the ordering on the tuple induced by \leq . It is a straightforward but tedious exercise to extend the proof to this case. ■

Corollary 4.7 *$\mathcal{L}(BASRL)$ is closed with respect to FO interpretations that also use BIT.*

Proof: Let $struct(\sigma)$ and $struct(\tau)$ be some vocabularies. Let $A \subset struct(\sigma)$ and $B \subset struct(\tau)$ be two problems. Given that $A \in \mathcal{L}(BASRL)$ and $B \leq_{fo+bit} A$ we have to show that $B \in \mathcal{L}(BASRL)$. It follows immediately from 3.3 that $\mathcal{L}(BASRL)$ is closed with respect to quantification and boolean operations, since the *set-reduce* functions defined in that proof satisfy the definition of *BASRL*. Closure under *BIT* operation i.e. for any function f , f expressible in *BASRL* \rightarrow $BIT(f, i)$ expressible in *BASRL*, follows from Lemma 4.6 above. Note that f either returns a singleton element from the active domain or it returns a bounded width tuple of elements. ■

Definition 4.8 Let π_1, π_2 denote two permutations on $[n] = 1, 2, \dots, n$. Let composition, $*$, denote the following operation:

$$\pi_1 * \pi_2(i) = \pi_2(\pi_1(i)), 1 \leq i \leq n$$

Let S_n denote the group of permutations on $[n]$ under composition. Let IM_{S_n} denote the following iterated multiplication problem: given permutations $\pi_1, \dots, \pi_n \in S_n$ as input, compute their composition, i.e. $\pi_1 * \pi_2 * \dots * \pi_n$.

The following theorem indicates the usefulness of IM_{S_n} .

Fact 4.9 ([CM87, IL89]) *IM_{S_n}* is complete for *L* under *FO* reductions with *BIT*.

We show how to express IM_{S_n} in *BASRL*.

Lemma 4.10 *IM_{S_n}* is expressible in *BASRL*.

Proof: We shall express IM_{S_n} as a predicate in the sense that given the input as stated earlier and also two other inputs viz. numbers i and j , our *BASRL* program will return true iff the iterated product permutation maps i to j . The input will be coded as follows: each permutation would be represented by tuples of the type, $[i, [j, k]]$, which means that the i^{th} permutation maps j to k . Thus, the input, say I , is a set of such tuples. Note that since i, j, k are represented by sets of respective cardinalities the input is of *set-height* 2. It is easy to write a program in *BASRL* to check that the permutation group is indeed S_n , where n is the number of elements (permutations) being multiplied. Also, n can be regarded as a constant available to us since one can always define it in *FO* as follows:

$$\exists x \forall y (y \leq x).$$

Then, the following program expresses IM_{S_n} . As before, we do not specify the types in the following to make it easy to read.

$$\begin{aligned}
IP(I, i) = & \text{set-reduce}(I, \text{identity}, \\
& \lambda(x\text{tuple}, \text{pair})(\text{set-reduce}(I, \text{identity}, \\
& \quad \lambda(x, p)(\text{if } (x.1 = p.1) \wedge (x.2.1 = p.2) \wedge \neg(p.1 = n) \\
& \quad \text{then } [\text{increment}(p.1).2, x.2.2] \\
& \quad \text{else } p), \\
& \quad \text{pair}) \\
& [1, i])
\end{aligned}$$

$$IM(I, i, j) = \text{if } IP(I, i).2 = j \text{ then true else false}$$

Note that the accumulator function returns a bounded tuple in the above program. ■

Corollary 4.11 $L \subseteq \mathcal{L}(BASRL)$.

Proof: Since IM_{S_n} is complete for L under *FO interpretations* that include *BIT* (by Fact 4.9), and it is expressible in *BASRL* (by Lemma 4.10), and $\mathcal{L}(BASRL)$ is closed under these reductions (by Corollary 4.7), the result follows. ■

Lemma 4.12 $\mathcal{L}(BASRL) \subseteq L$.

Proof: It suffices to show that a logspace deterministic Turing machine can simulate any *BASRL* program. Since the accumulator function only returns a bounded width tuple, we can just write the tuple on $O(\log n)$ bits of worktape. It is easy to see that the scan done by the set-reduce can be simulated by just scanning the input with the read-only head and an index tape that uses at most $\log n$ bits. Now all that remains is to show the closure under bounded number of compositions. This follows from the well known fact that logspace computable 0-1 functions are closed under compositions. ■

Finally, we have that,

Theorem 4.13 $L = \mathcal{L}(BASRL)$.

Proof: It follows from the previous lemma and the corollary. ■

Remarks. *BASRL* programs can be evaluated efficiently in parallel (since, $L \subseteq IND(\log n) \subseteq NC^2$).

5. Expressiveness of unrestricted SRL

Set-reduce language, without any restrictions, can become intractable, as we observed in the remarks following Theorem 3.10. Let \mathbf{N} be the set of natural numbers. In this section, we consider the complexity of functions, from $\mathbf{N} \rightarrow \mathbf{N}$, expressible in the *set-reduce language*, when extended with an unbounded successor function. We define $SRL + new$ as SRL augmented with another operator, *new*, that gives the language in effect the ability to construct a new element. In particular, let $new(D)$ return an element $\notin D$, where D is any set. Note that this is equivalent to having an unbounded successor operator. At first glance it may seem that this version of SRL is not that different from SRL . However, we show that these versions express all the primitive recursive functions. Observe that SRL contains a bounded successor function whereas $SRL + new$ contains an unbounded successor function.

Let $PrimRec$ denote the class of primitive recursive functions. Recall the definition of $PrimRec$ [DW].

Definition 5.1 Let $g : \mathbf{N} \rightarrow \mathbf{N}$, $h : \mathbf{N} \times \mathbf{N} \rightarrow \mathbf{N}$. Then, $f : \mathbf{N} \times \mathbf{N} \rightarrow \mathbf{N}$ is defined by primitive recursion from g, h if

$$\begin{aligned} f(0, t) &= g(t) \\ f(s + 1, t) &= h(s, g(s, t)) \end{aligned}$$

Let the *initial* functions be given as:

$$\begin{aligned} succ(i) &= i + 1 \\ n(i) &= 0 \\ p_k^n([i_1, \dots, i_n]) &= i_k \end{aligned}$$

A function is *primitive-recursive* if it is obtained from the *initial functions* by a bounded number of compositions and primitive-recursions.

Note that functions in $SRL + new$ give mappings between sets. However, we can consider them as functions from \mathbf{N} to \mathbf{N} , since finite ordered sets can be Gödel numbered in a standard way. For example, in our notation, a mapping between natural numbers and sets is given by:

$$0 = \emptyset, 1 = \{d_1\}, 2 = \{d_1, d_2\}, \dots, n + 1 = n \cup \{new(n)\} \dots$$

In the following, $\mathcal{F}(SRL + new)$ denotes the functions from \mathbf{N} to \mathbf{N} that can be expressed in the corresponding language.

Theorem 5.2 $PrimRec = \mathcal{F}(SR + new)$.

Proof: The initial functions are easily expressible in $SR + new$. Eg.,

$$proj_k(t) = t.k$$

$$(succ(S) = insert(new(S), S)).$$

In fact, this is the only usage of new .

(i): $PrimRec \subseteq \mathcal{F}(SRL + new)$:

It follows easily from the following

Proposition 5.3 $\mathcal{F}(SRL + new)$ is closed with respect to primitive recursion.

Proof: Let f be the function obtained from g, h by primitive recursion as defined above. Given that g, h are expressible in $SRL + new$, we show how to compute f in $SRL + new$.

$$f(S, T) = set-reduce(S, identity, \lambda(x, T')(hf(x, T')), [g(T), \{\}])$$

where $hf(x, T') = [h(T'.2, T'.1), insert(x, T'.2)]$ ■

(ii): $\mathcal{F}(SRL + new) \subseteq PrimRec$:

The proof in this direction assumes the following encoding of the sets as numbers: given an ordered domain $D = \{d_0 \leq d_2 \leq \dots d_n \leq \dots\}$, we encode any (ordered) $S \subseteq D$ as the number corresponding to the binary string $\langle s_n, \dots, s_k, \dots, s_0 \rangle$ such that s_i is 1 iff $d_i \in S$. Note that k gives the largest element s_k such that s_k is in S , and that any finite subset gives a finite number with infinitely many leading zeroes (which are ignored in our framework). Also, note that d_i corresponds to the number 2^i in our encoding, and hence, $\text{Log}(d_i) = i$.

The initial functions in $SRL + new$, with the possible exception of $insert$ and new , are clearly primitive recursive. To show that $insert$ is primitive recursive, we define the following functions. Let $\text{Bit}(n, i)$ denote the i -th bit of n , $\text{Div}(n, j)$ denote $\lfloor n/2^j \rfloor$, $\text{Log}(n) = \text{maximum } k \text{ such that } \text{Bit}(n, k) = 1$, $\text{Exp}(n, i) = n^i$, $\text{Rlog}(n) = \text{minimum } k \text{ such that } \text{Bit}(n, k) \text{ is } 1$, $\text{Mod}(n, j)$ denote $n \bmod 2^j$, and $\text{Cond}(b, i, j) = i$, if b is 1 and j , otherwise. Let $\langle n \rangle$ denote the binary representation of the number n . Noting that we have a successor operation, it can be shown that

Fact 5.4 $Bit, Div, Mod, Log, Rlog, Cond$ are primitive recursive.

Using this fact, we show how to express $insert(x, S)$: we set the bit corresponding to $x = d_i$, in this case, the i -th bit, to 1. Let $i = Log(x)$. Then,

$$new(S) \equiv Exp(2, Log(S) + 1)$$

$$insert(x, S) \equiv Cond(Bit(i, S), S, Div(S, i - 1) + 1 + Mod(S, i - 1))$$

Note that all the variables are treated as numbers in binary notation.

We have to show how to simulate the *set-reduce* operator using primitive recursion, and that completes the proof. Note that the order, \leq , in which *set-reduce* scans a set is given by the base functions, *choose* and *rest*. We observe that *choose*(S) is given by 2^i , where i is the position of the first non-zero least significant bit in $\langle S \rangle$, and *rest*(S) by shifting $\langle S \rangle$ to the right by $i + 1$ bits.

$$choose(S) \equiv Exp(2, Rlog(S))$$

$$rest(S) \equiv Div(S, Rlog(S) + 1)$$

Let *accf*, *appf*, *basef* be primitive recursive functions. Then any *set-reduce* expression in $SRL + new$, e.g.

$$f(S, y) = set-reduce(S, appf, accf, basef, y)$$

is equivalent to the following primitive recursive function:

$$\begin{aligned} f(0, y) &= basef(y) \\ f(S, y) &= accf(appf(choose(S), y), f(rest(S), y)) \end{aligned}$$

■

Remarks:

- Let *cons* be the list append operator. It can be shown in a manner similar to the proof above that

Corollary 5.5 $\mathcal{F}(LRL) = PrimRec = \mathcal{F}(SRL + cons)$.

- Note the crucial use of the types, \mathbf{N} and *set of N*, in $SR + new$ in the context of the comments preceding 3.13. Thus, we see that merely throwing in *new* operator increases the complexity of SRL all the way to *PrimRec*.

6. Complexity of SRL from its syntax

Given a program in *set-reduce language*, and the results in this paper, a scan of its syntax allows us to make certain conclusions regarding its complexity. If the user has sets of *set-height* greater than 1 in the program, then its complexity may be exponential. On the other hand, if sets of *set-height* at most 1 are used, then its complexity is polynomial in the size of the input sets. If in addition, the accumulator functions, $accf : \{\alpha, \gamma\} \rightarrow \beta$, (for some types α, γ, β) in his *set-reduce* constructs are such that β for any $accf$ is never of type *set*, then we are certain that the function expressed by the program is in L (or logspace). Any usage of the type *set* of some *unbounded type* in the program would possibly make the function it is computing very hard to optimize, but on the other hand using objects of unbounded type without using *set* of such objects makes it less difficult.

Let a be the maximum width of a *SRL* expression, i.e. the maximum arity of tuples used in a non-input set. Let d be the depth (defined in Lemma 3.6) of the expression. Let T_{ins} be the time complexity of an *insert* operation. Let n be the size of the input. Keeping in mind that the sets dealt with are of size polynomial in n , T_{ins} could be $O(1)$ (implemented by hashing), $O(\log n)$ (implemented by some balanced data structure) or at worst n^a , the maximum size of any set in the expression. Let $DTIME[n]$ denote the class of decision problems recognizable by deterministic Turing machines in time linear in the length of the input. Let $DTIME(f(n))$ denote the class of problems recognized by deterministic Turing machines in time bounded by $O(f(n))$. Then, we can easily bound the time complexity as follows.

Proposition 6.1 *Any SRL expression with width a and depth d is in $DTIME(n^{ad}T_{ins})$.*

Proof: By induction on the depth d .

$d = 0$: The base function *insert* takes time T_{ins} .

Any *set-reduce* over a set, say R , of depth d , where the *accf* and *app* functions are themselves of depth $d - 1$, takes time

$$\begin{aligned} &\leq |R| (\max\{\text{time of the } accf \text{ or } app\}) \\ &\leq n^a n^{a(d-1)} T_{ins} \quad \text{by the ind. hyp.} \\ &= O(n^{ad} T_{ins}) \end{aligned}$$

■

The bound leaves much room for improvement. In actually analysing a particular *SRL* expression, one usually can do much better, since then one can get rid of the overestimated n^a term that appears in the proposition above. Is $DTIME(n)$ expressible by a *SRL* expression with width 1 and depth 1Γ ? Apparently not. We show in the following that $DTIME(n)$ can be expressed by a *SRL* expression of width 2

and depth 3. However, the expression we obtain can actually be evaluated in time $O(n^2T_{ins})$ which is much better than the bound $O(n^6T_{ins})$ given by 6.1 above.

Proposition 6.2 *DTIME(n) is expressible by a SRL expression of width 2 and depth 3.*

Proof: We show how to simulate the computation of a *DTIME(n)* Turing machine by an *SRL* expression. Let σ be the alphabet, x_1, \dots, x_n be the input where $x_1, \dots, x_n \in \sigma$, and n be the input size. Let S denote the input as a set of pairs viz. $\{[1, x_1], [2, x_2], \dots, [n, x_n]\}$. Let us denote the work tape W as another set of pairs. It is easy to write a *SRL* expression, call it *create-tape*, that initializes W with blanks i.e. $\{[1, -], [2, -], \dots, [n, -]\}$. Let us denote the input tape head and work tape head positions by two variables, say P_1, P_2 . Now we can just use a *set-reduce* over S , thereby iterating n times, and in each iteration the *accf*, in this case, $F1$, updates W, P_1, P_2 according to the Turing machine program.

$$\text{Simulate}() = \text{set-reduce}(S, \text{identity}, \lambda(s, T)F1(T), [W, P_1, P_2])$$

$$F1(T) = \text{set-reduce}(S, \text{identity}, \lambda(s, X)F2(s, X), T)$$

Note that X is a 3-tuple of Worktape, position of input head, and position of work-tape head. What $F2(s, X)$ does is retrieve the contents of the tapes under the two heads and modify the work-tape contents and update the head positions according to the transition table of the Turing machine. We merely sketch an informal outline leaving the interested reader to express the functions in *SRL*.

$$\begin{aligned} F2(s, X) = & \text{if } (s.1 = X.2) \\ & \text{then set-reduce}(X.1, \lambda(t, ex)[t, ex], \lambda(tE, Y)\text{update}(tE, Y), \\ & \quad [\{\}, 0, 0], [X.2, X.3, s.2]) \\ & \text{else } X \end{aligned}$$

$$\begin{aligned} \text{update}(tE, R) = & (\text{if } (tE.1.1 = tE.2.2) \\ & \text{then use TM transition table and } tE.1.2 \\ & \quad (\text{work tape content}) \text{ and } tE.2.3(\text{input cell}) \text{ to make} \\ & \quad \text{a move i.e. change work head position } tE.2.2 \\ & \quad \text{and input head position } tE.2.1 \text{ accordingly and} \\ & \quad \text{return } [\text{insert}([tE.1.1, tE.1.2'], R.1), tE.2.1', tE.2.2'] \\ & \text{else } [\text{insert}(tE.1, R.1), R.2, R.3] \end{aligned}$$

Note that W is a set of pairs and hence the width is 2. $F1$ is of depth 2, since it uses one *set-reduce* over S to get the input tape content and another over W to get at the work-tape content. W is set up initially by a depth 1 *set-reduce* function called *create-tape*. The total depth equals 3. ■

Note that we use the *increment* function implicitly in updating the head positions and that the *set-reduce* over W is repeated only once for one full scan of S , and *increment* is also done only once for one full scan of W . An analysis of the time complexity of the expression, `Simulate()`, reveals that the two *set-reduce*'s in $F1$ together take $O(nT_{ins})$ time and since it is iterated over n times, the total complexity is n^2T_{ins} .

Let k be a constant. The proof above can easily be generalized to show that

Corollary 6.3 *DTIME(n^k) is expressible by a SRL expression of width $k + 1$ and depth $k + 3$.*

Remarks. The SRL expression obtained above can be evaluated in time $O(n^{2k}T_{ins})$.

Let SR_h denote the class of problems expressible by a version of *set-reduce language* that has its *set-height* h and *tuple-width* $= O(1)$. Let n denote the input size. Let $2^{i\#n}$ denote a stack of i 2's, i.e. $2^{0\#n} = n^{O(1)}$, $2^{i+1\#n} = 2^{2^{i\#n}}$.

Then, following the preceding proof, it can be shown that

Corollary 6.4 *For $h = 1, 2, \dots$ $SRL_h = DTIME(2^{h\#n})$.*

Remarks. This hierarchy is mentioned here for the sake of completeness. It is quite similar in notion to the results of [HS88], [AV88] and others.

7. The Rôle of Ordering

A set stored by a computer has its members in some order. Simply put, any object is a sequence of bits, thus falling in place in lexicographical order. This allows any database system to search through a set in lexicographical order à la *set-reduce*; and, also to compute information that may depend on the somewhat arbitrary ordering that ensues. For example, one may compute the order dependent boolean query:

$$\text{Purple}(\text{First}(S))$$

namely that the element that happens to be first in the arbitrary ordering of the set S satisfies the predicate `Purple(\cdot)`.

It is neither surprising, nor especially dangerous that programs that search through a set in a given order may compute some information that depends on that order. If the order is truly independent of any information we wish to be computing and if our programs are correct, then the answers will be independent of the ordering. Furthermore, most sets of data have at least one natural ordering which can be used instead of the arbitrary ordering, for example one can print the elements of a set of employees in order of their names, or, date of hire, etc.

Still, if we are not certain that our programs are correct, then it would be nice to know whether the answers we get depend on the arbitrary ordering of elements within a set. Furthermore, one can imagine difficulties when long queries are suspended and then resume, or when different parts of them are carried out at different sites of a distributed database. In particular, these separated processes may be using different, arbitrary ordering of the same set in which case, just combining their computations without taking note of their dependence on the ordering, could lead to error.

In any case, there is general sentiment in the theoretical database community that ordering is dangerous and that order dependent queries should be avoided. In fact, in the influential paper [CH82a], Chandra and Harel define a *query* to be an order-independent query and they ask the question:

Question 7.1 *Is there a natural language that expresses exactly the set of polynomial-time computable, order-independent queries?*

One can make this question more precise by removing the undefined term “natural” and instead ask:

Question 7.2 *Is there a recursively enumerable set of programs that compute exactly the set of polynomial-time computable, order-independent queries over relational databases?*

The above two questions remain open in spite of many years of intensive study. See [IL90] for a history of this subject. Here we give an overview of what is known about Questions 7.1 and 7.2.

In a preliminary version of the paper [CH82a], Chandra and Harel defined fixed point logic, FP, which is an extension of first-order logic to include applications of the fixed point operator, thus allowing the inductive definition of new relations. In symbols: $FP = (FO(\text{wo}\leq) + LFP)$. Chandra and Harel conjectured that there was a hierarchy of queries in FP consisting of successive applications of LFP and first-order operations. In response, Immerman showed that Chandra and Harel’s conjecture was false:

Fact 7.3 ([Imm82, Imm87]) *Every query in FP is expressible the form*

$$\text{LFP}(\varphi(R)[\bar{t}])$$

where \bar{t} is a tuple of terms and φ is a quantifier-free formula containing no occurrences of LFP.

Perhaps more interesting is the fact that if a total ordering of the universe is present, then the queries expressible in (FO + LFP) are exactly those computable in polynomial time.

Fact 7.4 ([Imm82, Va82])

$$(\text{FO} + \text{LFP}) = \text{P}$$

Fact 7.4 fails badly if we remove the ordering. For example, it is easy to show that without an ordering we cannot count. In fact, if EVEN represents the query that is true if the size of the universe is even, then:

Fact 7.5 ([CH82a]) *EVEN is not expressible in (FO(wo \leq)) + LFP).*

Indeed, before 1989, examples involving the counting of large, unstructured sets were the only problems known to be in order-independent P but not in (FO(wo \leq) + LFP). In 1982, Immerman [Imm82] considered the language (FO(wo \leq)+LFP+count) in which structures are two-sorted, with an unordered domain $D = \{d_0, d_1, \dots, d_{n-1}\}$ and a separate number domain: $N = \{0, 1, \dots, n-1\}$ with the database predicates defined on D and the standard ordering defined on N . The two sorts are combined via counting quantifiers:

$$(\exists^i x)\varphi(x)$$

meaning that there exist at least i elements x such that $\varphi(x)$. Here i is a number variable and x is a domain variable.

For quite a while, it was an open question whether the language (FO(wo \leq)+LFP+count) is equal to order independent P. A positive answer would have provided a nice solution to Question 7.1.

Instead, in [CFI89] it was proved that that (FO(wo \leq) + LFP + count) is strictly contained in order-independent P. See Theorem 7.8 below for an explanation and slight generalization of this result.

See Figure 7.6 for the relationships between the polynomial-time query classes we have been discussing.

$$\begin{aligned}
(\text{FO}(\text{wo}\leq) + \text{LFP}) &\subset (\text{FO}(\text{wo}\leq) + \text{LFP} + \text{count}) \\
&\subset (\text{order-independent P}) \\
&\subset (\text{FO} + \text{LFP}) = \text{P}
\end{aligned}$$

Figure 7.6: Some polynomial-time query classes.(The relation “ \subset ” denotes proper containment.)

Another approach to capturing order-independent queries is worthy of mention here. In [OBB89] the language Machiavelli is defined. It contains an operator called *hom* which is similar to set-reduce. In the following definition, *op* is any previously defined binary operation.

$$\begin{aligned}
\text{hom}(f, \text{op}, z, \{\}) &= z \\
\text{hom}(f, \text{op}, z, \{x_1, x_2, \dots, x_n\}) &= \text{op}(f(x_1), \dots, \text{op}(f(x_n, z)) \dots)
\end{aligned}$$

It is not hard to see that in the presence of an ordering, and with set-height restricted to at most one, the languages SRL and a similar *hom*-based language, which we will refer to as HL, have equivalent expressive power. However, in [OBB89], an instance of *hom* is called *proper* if the corresponding *op* is commutative and associative. It follows that an application of proper *hom* does not derive any information from the ordering in which a set is presented. Thus the language “proper HL” is order-independent and would seem to be a candidate for order-independent P.

One obstacle to this is easily overcome: when *op* is associative, the application of *hom* may be drawn as a binary tree of height $\log n$, and thus evaluated in parallel time $O[\log n]$ times the parallel time to perform a single *op*. It follows that “proper Machiavelli” is contained in the class NC consisting of those problems computable in parallel time $(\log n)^{O[1]}$ using polynomially many processors. NC is believed to be strictly contained in P [C85].

We can alleviate this problem by allowing “proper HL” to iterate an operation polynomially many times. One way to do this is to consider the language similar to $(\text{FO}(\text{wo}\leq) + \text{LFP} + \text{count})$ which has a number domain, N , separate from the database domain. One can then safely allow arbitrary applications of *hom* over the number domain. Define $(\text{FO}(\text{wo}\leq) + N + \text{hom})$ to be this class. Then we have the following proposition which says that “proper HL together with a polynomial iteration operation” is at least as expressive as $(\text{FO}(\text{wo}\leq) + \text{LFP} + \text{count})$. As of this writing, we do not know whether or not this inclusion is proper:

Proposition 7.7

$$(\text{FO}(\text{wo}\leq) + \text{LFP} + \text{count}) \subseteq (\text{FO}(\text{wo}\leq) + N + \text{hom})$$

Proof The above discussion explained why $(\text{FO}(\text{wo}\leq) + N + \text{hom})$ contains $(\text{FO}(\text{wo}\leq) + N + \text{LFP})$. Thus it suffices to show how to count using proper hom. This is easy. Let $f : D \rightarrow N$ be the function that takes everything in the database domain to the number 1. Then we can count a set $S \subseteq D$ using hom as follows:

$$\text{count}(S) = \text{hom}(f, +, 0, S)$$

■

We next show that the lower bound from [CFI89] **does** apply to the language $(\text{FO}(\text{wo}\leq) + N + \text{hom})$. It also applies to the language $(\text{FO}(\text{wo}\leq) + \text{count} + \text{while})$.²

Theorem 7.8 *The set (order-independent P) is not contained in $(\text{FO}(\text{wo}\leq) + N + \text{hom} + \text{while})$.*

Proof The paper [CFI89] constructs a sequence of structures $G_n, H_n, n = 1, 2, \dots$. These structures contain $O[n]$ domain elements. G_n and H_n may be distinguished in linear time if we have access to any ordering on their domains. By contrast, G_n and H_n agree on all sentences in $(\text{FO}(\text{wo}\leq) + \text{count})$ containing at most n distinct variables. (If the simple, polynomial-time order-independent property that characterizes G_n were expressible in $(\text{FO}(\text{wo}\leq) + \text{LFP} + \text{count})$ or in $(\text{FO}(\text{wo}\leq) + \text{count} + \text{while})$ then it would follow that a first-order sentence with a *bounded* number of variables would distinguish the graphs G_n and H_n . This is true because the operators LFP and ‘while’ are simply “formula iterators” and do not increase the number of distinct variables in the formula.)

Now, we show that over the structures G_n, H_n applications of hom give us no new expressive power. This is because G_n and H_n are almost ordered. That is, there is a first-order, quasi-total ordering on the vertices. The vertices are partitioned into color classes of size at most 4 and the color classes are totally ordered. Thus we

²In [Va82], Vardi defined the language $(\text{FO} + \text{while})$, i.e. first-order logic together with an unbounded iteration operator, and showed that its expressive power is equal to PSPACE. (See also [Imm82b] for an equivalent formulation of an unbounded iterator applied to FO giving PSPACE.) See also [AV91] for a surprising new result: $(\text{FO}(\text{wo}\leq) + \text{while}) = (\text{FO}(\text{wo}\leq) + \text{LFP})$ if and only if $\text{P} = \text{PSPACE}$.

can compute hom of a set by walking through the color classes occurring in the set, applying the operator by hand to at most four elements in each class. ■

One of us (Immerman) has studied the issue of ordering because of its intimate connection with his study of descriptive and computational complexity [IL90]. Another of us (Stemple) has developed a theory of finite sets because of their importance in database transactions [SS89]. It is an unaesthetic aspect of any such theory to date, that in order to develop a theory of unordered finite sets that is rich enough to describe computation, one seems to need an ordering on these sets.

It seems to us unacceptable to use impoverished query and transaction languages in order to have the aesthetically desirable characteristic of order-independence. Our view is that one should use a language that we know includes all the feasible queries, i.e. P. But, that one should use a theorem prover such as Sheard's extended Boyer-Moore theorem prover [SS89] to prove that our queries and transactions are correct. Correctness here would mean that the queries and transactions do what we want them to do. In particular, they preserve the database integrity constraints, and, when desired, they compute only order-independent properties. Thus we can add to Figure 7.6 the class (proved order-independent P) of those queries in SRL, or equivalently in (FO + LFP) that our theorem-prover has shown to be order-independent.

8. Conclusions

The inference mechanism in [SS89] on finite set terms with variables proves only properties that are true in all models. It can be used to prove that a *set-reduce* expression is independent of order, though it is of necessity incomplete with respect to this problem. (Any algebra meeting the specification is powerful enough to express P problems, and thus the order independence of an arbitrary expression in the language cannot be decided.) Likewise we can prove that some expressions depend on the order. The language of expressions that do not vary with order would have the properties of a specification with an initial algebra, but this language is not recursively enumerable. However, it may be that the set of expressions that their prover can prove are order independent includes all polynomial-time computable, order independent queries. We are investigating this possibility.

8.1. Open Problems

1. Problems Related to Ordering:

- (a) Settle Question 7.1. In particular, prove or disprove the conjecture that the subset of SRL that can be proved order-independent using Sheard's Boyer-Moore theorem prover is exactly order-independent P.
 - (b) Settle variants of Question 7.1 for smaller complexity classes, e.g. L, NL, NC. Note that for complexity classes NP and above, the question is easily settled because an ordering can simply be existentially quantified and thus no ordering need be provided.
2. Our results show that there is a clear demarcation between SRL which expresses the polynomial-time computable queries and unrestricted SRL which computes all primitive recursive queries. Thus, it is very desirable to improve our characterization of this demarcation line. We would like to be able to say in a very general way, "Yes, *these* sorts of operations and functionalities can all safely be added, without taking us out of P. On the other hand, any of *those* will bring us all the way up to Primitive Recursive complexity."
 3. 6.1 shows that to a certain extent the time complexity of an SRL expression can be read off its face. However, we suspect that the complexity bounds we give here can be improved.
 4. The classical complexity classes L, NL, P give an interesting basis for comparing the expressibility of query and transaction languages. On the other hand, these are clearly not precisely the complexity classes that are appropriate for studying the true costs of queries and transactions for modern database systems. We are in the process of taking a step in this direction by defining and studying complexity classes more appropriate for database systems. In particular the cost of disk I/O's is given its due place, and incremental complexity is emphasized: we consider the complexity of processing a long sequence of transactions on-line. Much more work is needed in this direction.

References

- [AU79] A. Aho, J. Ullman: Universality of data retrieval languages. *Proceedings of Sixth ACM Symposium on POPL*, Jan. 1979, 110-117.
- [AK90] S. Abiteboul, P. Kanellakis: Database theory column: Query languages for complex object databases. *SIGACT News 21 No. 3*, Summer 1990, 9-18.
- [AK89] S. Abiteboul, P. Kanellakis: Object identity as a query language primitive. *Rapports de Recherche No. 1022, INRIA*, Apr 1989.
- [AB88] S. Abiteboul, C. Beeri: On the power of languages for the manipulation of complex objects. *Rapports de Recherche no. 846, INRIA*, May 1988.

- [AC89] F. Arafati, S. Cosmadakis: Expressiveness of restricted recursive queries, *Proceedings of 21st ACM STOC*, 1989, 113-126.
- [AV88] S. Abiteboul, V. Vianu: Procedural and declarative database update languages. *Proceedings of the Seventh ACM SIGACT-SIGMOD-SIGART Symposium on PODS*, 1988, 240-250.
- [AV89] S. Abiteboul, V. Vianu: Fixpoint extensions of first-order logic and datalog-like languages. *Proceedings of LICS*, 1989, 2-11.
- [AV91] S. Abiteboul, V. Vianu: Generic computation and its complexity. To appear in *32nd IEEE Symposium on FOCS*, 1991.
- [BIS88] D. Barrington, N. Immerman, H. Straubing: On uniformity within NC^1 . *Journal of Computer Systems and Science* 41, No. 3, 1990, 274-306.
- [BM] R.S.Boyer, J.S.Moore:*A Computational Logic*, Academic Press, New York, 1979.
- [CFI89] J. Cai, M. Furer, N. Immerman: An optimal lower bound on the number of variables for graph identification. *Proceedings of 30th IEEE Symposium on FOCS*, 1989, 612-617.
- [CH80] A. Chandra, D. Harel: Computable queries for relational databases. *Journal of Computer Systems and Science* 21, 1980, 156-178.
- [CH82a] A. Chandra, D. Harel: Structure and complexity of relational queries. *Journal of Computer Systems and Science* 25, 1982, 99-128.
- [CH82b] A. Chandra, D. Harel: Horn clauses and fixpoint query hierarchy. *Proceedings of 14th ACM STOC*, May 1982, 158-163.
- [Ch81] A. Chandra: Programming primitives for database languages. *Proceedings of ACM Symposium on POPL*, 1981, 50-62.
- [CSV84] A. Chandra, L. Stockmeyer, U. Vishkin: Constant-depth reducibility, *SIAM Journal of Computing* 13, May 1984, 423-439.
- [Co64] A. Cobham: The intrinsic computational difficulty of functions. *Proceedings of the 1964 Congress for Logic, Philosophy and Methodology of Science*, North Holland, 24-30.
- [C85] S. Cook: A taxonomy of problems with fast parallel algorithms. *Information and Control* 64, (1985), 2-22.
- [CM87] S. Cook, P. McKenzie: Problems complete for deterministic logspace. *Journal of Algorithms* 8, 1987, 385-394.
- [CK85] S. Cosmadakis, P. Kanellakis: Parallel evaluation of recursive rule queries. *Proceedings of 5th ACM Symposium on PODS*, 1986, 280-290.
- [DW] M. Davis, S. Weyukar: Computability, Complexity and Languages. *Academic Press*, 1983.
- [Gu83] Y. Gurevich: Algebras of Feasible Functions. *Proceedings of 24th IEEE Symposium on Foundations of Computer Science*, October 1983, 210-214.

- [HS89a] R. Hull, J. Su: Untyped sets, invention and computable queries. *Proceedings of 8th ACM Symposium on PODS*, 1989, 347-360.
- [HS89b] R. Hull, J. Su: On bulk data type constructors and manipulation primitives - a framework for analyzing expressive power and complexity. *Proceedings of 2nd International Workshop on Database Programming Languages*, June 1989, 396-410.
- [HS88] R. Hull, J. Su: On the expressive power of database queries with intermediate types. *Proceedings of 7th ACM Symposium on PODS*, Mar. 1988, 39-51.
- [Imm87] N. Immerman: Languages that capture complexity classes. *SIAM Journal on Computing* 16 No. 4, Aug. 1987, 760-778.
- [Imm82] N. Immerman: Relational queries computable in polynomial time. *Proceedings of the 14th ACM STOC*, May 1982, 147-152.
- [Imm82b] N. Immerman: Upper and lower bounds for first order expressibility. *Journal of Computer and System Sciences* 25, 1982, 76-98.
- [Imm88] N. Immermann: Nondeterministic space is closed under complementation. *SIAM J. Comput.* 17, No. 5, (1988), 935-938.
- [IL89] N. Immerman, S. Landau: The complexity of iterated multiplication. *Proceedings of 4th Structure in Complexity Theory Conference*, 1989, 104-111.
- [IL90] N. Immerman, E. Lander: Describing graphs: a first-order approach to graph canonization. *Complexity Theory Retrospective*, A. Selman, ed., Springer Verlag (1990).
- [K88] P. Kanellakis: Elements of relational database theory. *Tech. Report CS-88-09, Dept. Of Computer Science, Brown University*, Apr. 1988.
- [OBB89] A. Ohori, P. Buneman, V. Breazu-Tannen: Database programming in Machiavelli - a polymorphic language with static type interference. *Proceedings of the ACM SIGMOD*, June 1989, 46-57.
- [Q89] X.Qian: The expressive power of the bounded iteration construct. *Proceedings of the 2nd International Workshop on Database Programming Languages*, 1990.
- [SS89] T. Sheard, D. Stemple: Automatic verification of database transaction safety. *ACM transactions on Database Systems* 14, No. 3, Sep. 1989, 322-368.
- [Va82] M.Y. Vardi: The complexity of relational query languages. *Proceedings of 14th ACM STOC*, May 1982, 137-146.
- [VS90] J. Vitter, E. Shriver: Optimal disk I/O with parallel block transfer. *Proceedings of the 22nd ACM STOC*, May 1990, 159-169.