

A Programming Language for the Inductive Sets, and Applications*

DAVID HAREL^{†,‡}

The Weizmann Institute, Rehovot, Israel

AND

DEXTER KOZEN^{†,§}

Aarhus University, Aarhus, Denmark

We describe a programming language IND that generalizes alternating Turing machines to arbitrary first-order structures. We show that IND programs (resp. everywhere-halting IND programs, loop-free IND programs) accept precisely the inductively definable (resp. hyperelementary, elementary) relations. We give several examples showing how the language provides a robust and computational approach to the theory of first-order inductive definability. We then show: (1) on all acceptable structures (in the sense of Moschovakis (“Elementary Induction on Abstract Structures,” North-Holland, Amsterdam, 1974), r.e. dynamic logic is more expressive than finite-test dynamic logic. This refines a separation result of Meyer and Parikh (“Proc. 12th ACM Sympos. on Theory of Computing,” 1979, pp. 167–175); (2) IND provides a natural query language for the set of fixpoint queries over a relational database, answering a question of Chandra and Harel (*J. Comput. System Sci.* 25, No. 1 (1982), 99–128). © 1984 Academic Press, Inc.

1. INTRODUCTION

In this paper we describe a programming language IND which is very similar to the nondeterministic language introduced in Manna (1970). In

* This paper is a significantly revised version of an unrefereed preliminary report which appeared under the same title in the “Proceedings of the 9th Colloquium on Automata, Languages, and Programming,” Aarhus, Denmark, July 1982, pp. 313–329.

[†] Research supported in part by a Bath–Sheva fellowship.

[‡] Work done in part at IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598.

[§] On leave from IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598.

its most basic form, the language consists of only 3 types of statements:

$$\begin{array}{ll}
 l_1: y \leftarrow \exists & (\text{or } y \leftarrow \forall) \\
 l_2: \text{accept} & (\text{or reject}) \\
 l_3: \text{if } R(\bar{x}) \text{ then goto } l_4
 \end{array}$$

where $R(\bar{x})$ is an atomic first-order formula.

An IND program P can run in any first-order structure of the same similarity type. The input is an initial assignment to variables $\bar{x} = x_1, \dots, x_n$, where \bar{x} contains (at least) all the free variables of the program, and execution starts at the first statement. Statements of the form l_1 assign an arbitrary element of the domain existentially (universally) to variable y , just as in alternating Turing machines (Chandra, Kozen, and Stockmeyer, 1981). A statement of the form l_2 causes immediate acceptance or rejection, and l_3 is an ordinary conditional branch. The definition of acceptance of the input \bar{x} is the same as in alternating Turing machines (Chandra *et al.*, 1981) involving an inductively-defined labeling of the computation tree with either 0 (reject), 1 (accept), or \perp (undefined); see Section 3.

In Section 5 we show:

- (1) IND programs accept precisely the relations definable by elementary (first-order) induction.
- (2) IND programs which halt on all inputs (i.e., either accept or reject) accept precisely the hyperelementary (or inductive, co-inductive) relations.
- (3) Loop-free IND programs accept precisely the elementary (first-order definable) relations.

In countable acceptable structures (see Moschovakis, 1974) such as the natural numbers N , (1) and (2) become

- (1) IND programs accept precisely the Π_1^1 relations.
- (2) IND programs which halt on all inputs accept precisely the Δ_1^1 relations.

IND provides a simple and highly intuitive computational framework for the theory of inductive definability. For example, our proof of (2) involves showing that if both a relation and its complement are accepted by IND programs P_1 and P_2 , then the two programs can be simulated by a third program that always halts, P_3 . P_3 simulates steps of P_1 and P_2 alternately, halting whenever one or the other halts, just as in the usual Turing machine proof that an r.e., co-r.e. set is recursive. Many other elementary results, such as the stage comparison theorem, closure theorem, and

separation theorem (Moschovakis, 1974), have machine-based proofs using IND that recall analogous proofs in recursion theory that use Turing machines. The availability of such a tool is especially important now that concepts central to inductive definability theory have resurfaced in computer science in recent work on program logics (Meyer and Parikh, 1980), and program verification in the presence of fairness or unbounded non-determinism (Apt and Plotkin, 1981; Lehmann, Pnueli, and Stavi, 1981; Grumberg, Francez, Makowsky, and de Roever, 1981).

In Section 6 we use IND to characterize the expressive power of dynamic logic. Meyer and Parikh (1980) have shown that dynamic logic with unrestricted recursively enumerable programs (DL_{re}) is strictly more expressive than many limited versions, such as DL with finite tests (DL_{ft}). The result is proved by transferring the problem to the problem of distinguishing ω^ω and $\omega^\omega \cdot 2$ in fragments of infinitary logic, and does not provide any insight into the inherent computational power of DL. In Section 6 we show that on any acceptable structure, DL_{re} is more expressive than DL_{ft} . Specifically, we show that in any acceptable structure, DL_{re} and DL_{ft} define exactly the IND complexity classes $\tau(\omega_1^{CK})$ and $\tau(\omega)$, respectively, where ω_1^{CK} is the first non-recursive ordinal. In any acceptable structure, $\tau(\omega)$ is the set of first-order definable relations, and on recursive acceptable structures (such as N), $\tau(\omega_1^{CK}) = \Delta_1^1$. The classes $\tau(\omega)$ and $\tau(\omega_1^{CK})$ can be separated on any acceptable structure by a simple diagonalization argument.

It should be emphasized that the expressiveness results of (Meyer and Parikh, 1980) are *schematic*, in the sense that they consider $L_1 \leq L_2$ if there is an interpretation of L_1 in L_2 which holds uniformly over all structures, whereas we will write $L_1 \leq L_2$ if *for each structure* there is an interpretation of L_1 in L_2 . The former gives stronger positive expressibility results, and the latter gives stronger negative expressibility results (such as $DL_{ft} \neq DL_{re}$). Our positive expressibility results (such as $DL_{re} = \tau(\omega_1^{CK})$) are not to be interpreted schematically.

In Section 7 we show the connection between inductive definability and the fixpoint queries of (Aho and Ullman, 1979; Chandra and Harel, 1980) for relational data bases. Coupled with a recent result of Immerman (1982), this observation shows that IND defines exactly the class of fixpoint queries FP, thus answering a question of Chandra and Harel (1980).

2. PROGRAMMING EXAMPLES

The definition of acceptance for IND programs involves labelings of an and/or computation tree generated downwards from the root or start configuration, much as in alternating Turing machines (Chandra *et al.*, 1981).

We defer the formal definition of acceptance until Section 3; in this section we whet the reader's intuition with some sample IND programs.

We first add a few more powerful programming constructs. These constructs are defined in terms of the primitive constructs given in the last section, and do not change the expressive power of the language.

An unconditional jump *goto l* can be obtained by using the test $y = y$ in the conditional. Certain more complicated conditional forms, such as

if $R(\bar{x})$ *then goto* l_1 *else goto* l_2
if $R(\bar{x})$ *then accept else reject*
if $\neg R(\bar{x}) \vee S(\bar{x})$ *then goto* l

are obtained by manipulation of control flow. The assignment statement $y \leftarrow t$ is obtained by

$y \leftarrow \exists$	or	$x \leftarrow \exists$
<i>if</i> $y \neq t$ <i>then reject</i>		<i>if</i> $x \neq t$ <i>then reject</i>
		$y \leftarrow \exists$
		<i>if</i> $y \neq x$ <i>then reject</i>

where x is a new variable, if y occurs in t .

There is a loop-free program to compute any first-order formula. For example, an element x of a Boolean algebra is *atomless* if it satisfies the formula

$$\forall y \leq x \quad (y \neq 0 \supset \exists z \leq y \quad (0 \neq z \wedge z \neq y)).$$

The set of such elements is accepted by the program

$y \leftarrow \forall$
if $\neg y \leq x \vee y = 0$ *then accept*
 $z \leftarrow \exists$
if $\neg z \leq y \vee 0 = z \vee z = y$ *then reject*
accept.

However, IND programs can accept sets that are not first-order definable. For example, l_1 below accepts all pairs (x, y) in the reflexive transitive closure of R , and l_2 accepts y iff R is well-founded below y :

l_1 : if $x = y$ then accept $z \leftarrow \exists$ if $\neg R(x, z)$ then reject $x \leftarrow z$ goto l_1	l_2 : $x \leftarrow \forall$ if $\neg R(x, y)$ then accept $y \leftarrow x$ goto l_2
---	---

The statement $l_5 \vee l_6$ which accepts if either the program starting at l_5 or that starting at l_6 accepts is encoded by

$$y \leftarrow \exists$$

$$\text{if } y = z \text{ then goto } l_5 \text{ else goto } l_6$$

where y is a new variable and z is any other variable (provided the domain contains at least two elements). The statement $l_5 \wedge l_6$ is defined similarly, using \forall instead of \exists . One can also encode the statement $\neg l_5$ which accepts (rejects) iff the computation starting at l_5 rejects (accepts). This is done by taking the whole program P and constructing its dual \bar{P} by interchanging \forall/\exists and *accept/reject* statements. The program \bar{P} then accepts (rejects) exactly when program P rejects (accepts), starting at any point.

The statement $\neg l_5$ is then given by *goto* \bar{l}_5 where \bar{l}_5 is the statement corresponding to l_5 in the dual program. Of course, the statement $\neg \bar{l}_5$ in the dual program is replaced by *goto* l_5 .

These constructs allow us to encode the statement

$$\text{if } \varphi(\bar{x}) \text{ then goto } l_1 \text{ else goto } l_2,$$

where $\varphi(\bar{x})$ is any first-order formula, or for that matter any relation computed by a program that always halts:

$$(m_1 \wedge l_1) \vee (\neg m_1 \wedge l_2)$$

where m_1 is the first statement of a program computing $\varphi(\bar{x})$.

One of our main results is that there is an IND program that accepts any relation definable by first-order induction. For example, the subgroup H of G generated by a, b is the least subset of G such that

$$x \in H \leftrightarrow x = a \vee x = b \vee \exists y \in H \exists z \in H x = y \cdot z \vee x = y^{-1}.$$

Membership in H is computed by the program

```

l1: if  $x = a \vee x = b$  then accept
       $y \leftarrow \exists$ 
      if  $x = y^{-1}$  then goto l3
       $z \leftarrow \exists$ 
      if  $x \neq y \cdot z$  then reject
       $l_2 \wedge l_3$ 
l2:  $x \leftarrow z$ 
      goto l1
l3:  $x \leftarrow y$ 
      goto l1

```

To give an example involving an unbounded alternation of quantifiers, consider a two-person game like chess or go. The set of board positions from which a given player has a forced win is defined inductively from the legal-move and (immediate) win predicates by

$$\text{force}(x) \leftrightarrow \text{win}(x) \vee \exists y(\text{legal-move}(x, y) \wedge \neg \text{win}(y) \\ \wedge \forall z(\text{legal-move}(y, z) \rightarrow \text{force}(z)))$$

and is accepted by the program

```

l1: if win( $x$ ) then accept
       $y \leftarrow \exists$ 
      if  $\neg \text{legal-move}(x, y) \vee \text{win}(y)$  then reject
       $x \leftarrow \forall$ 
      if  $\neg \text{legal-move}(y, x)$  then accept
      goto l1

```

IND programs can run for more than a finite amount of time and still halt. An example of a program that runs in N for time $\omega + 1$ is

```

       $y \leftarrow \forall$ 
l1: if  $y = 0$  then accept
       $y \leftarrow y - 1$ 
      goto l1

```

In N , the running times of IND programs are exactly the recursive ordinals. In fact, we can take the set of computation trees of IND programs as a set of notations for recursive ordinals.

3. SEMANTICS OF ACCEPTANCE

The semantics of acceptance is formally almost identical to that of alternating Turing machines (Chandra *et al.*, 1981). Intuitively, it consists of two stages: (1) *generation* of the computation tree downwards from the root, and (2) *evaluation* of the acceptance function upwards from the leaves to the root. Associated with each node of the computation tree is a unique *configuration* (l, v) , where l is the label of one of the statements in the program and v is a valuation of program variables over the domain of computation A . If a node of the computation tree is labeled c , then its immediate descendants are labeled with all elements of $N(c)$, the *next configurations* of c , which are defined by cases, depending on the statement $\text{stmt}(c)$ labeled by c 's first component:

$$\begin{aligned}
 N(l, v) &= \{(l', v[y \leftarrow a]) \mid a \in A\} && \text{if } \text{stmt}(l, v) \text{ is } y \leftarrow \exists \text{ or } y \leftarrow \forall \\
 &= \emptyset && \text{if } \text{stmt}(l, v) \text{ is } \textit{accept} \text{ or } \textit{reject} \\
 &= \{(m, v)\} && \text{if } \text{stmt}(l, v) \text{ is } \textit{if } R(\bar{x}) \text{ then goto } m \\
 &&& \text{and } A, v \models R(\bar{x}) \\
 &= \{(l', v)\} && \text{if } \text{stmt}(l, v) \text{ is } \textit{if } R(\bar{x}) \text{ then goto } m \\
 &&& \text{and } A, v \models \neg R(\bar{x})
 \end{aligned}$$

where l' denotes the next statement after l in the program (or the first statement, if l was the last). The root of the tree is the start configuration. The acceptance function e^* can be regarded as a labelling of nodes of the computation tree with either 0 (reject), 1 (accept), or \perp (undefined), but formally its domain is the set C of configurations. It is defined as the supremum of a chain of approximating labelings e^α . Intuitively, $e^\alpha(c) = 1$ (resp. 0) if c has been determined to be an accept (resp. reject) configuration by time α ; $e^\alpha(c) = \perp$ if neither has been determined by time α . At time 0, nothing is determined, thus $e^0(c) = \perp$ for all c . At time 1, the leaves of the tree, corresponding to *accept* and *reject* statements, are labeled 1 and 0, respectively, and everything else is labeled \perp . If $\text{stmt}(c) = y \leftarrow \exists$ and $e^\alpha(d) = 1$ for some $d \in N(c)$, then $e^{\alpha+1}(c) = 1$. If $\text{stmt}(c) = y \leftarrow \forall$ and all $d \in N(c)$ are eventually labeled 1, then c becomes labeled 1 upon completion of the labeling of $N(c)$. Note that because of unbounded nondeter-

minism, it may take more than a finite amount of time for a configuration to become labeled 0 or 1.

Formally, let C be the set of configurations, let Ord be the class of ordinals, and let $*$ \notin Ord with $\alpha < *$ for all $\alpha \in \text{Ord}$. Define the sequence $e^\alpha: C \rightarrow (0, 1, \perp)$ inductively by $e^\alpha = \sqcup_{\beta < \alpha} \tau(e^\beta)$, where τ is the \sqsubseteq -monotone map defined by

$$\begin{aligned} \tau(e)(c) &= 1 && \text{if } \text{stmt}(c) = \text{accept} \\ &= 0 && \text{if } \text{stmt}(c) = \text{reject} \\ &= \bigvee_{d \in N(c)} e(d) && \text{if } \text{stmt}(c) = y \leftarrow \exists \\ &= \bigwedge_{d \in N(c)} e(d) && \text{if } \text{stmt}(c) = y \leftarrow \forall \\ &= e(d) && \text{if } \text{stmt}(c) = \text{if } R \text{ then goto } m \text{ and } N(c) = (d). \end{aligned}$$

The meet \bigwedge and join \bigvee are with respect to the ordering $0 < \perp < 1$, and should not be confused with the approximation ordering \sqsubseteq with join \sqcup , defined $\perp \sqsubseteq 0$, $\perp \sqsubseteq 1$ and extended pointwise to labellings. e^* is the \sqsubseteq -least fixpoint of τ .

We say c becomes properly labeled at time α if α is the least ordinal such that $e^\alpha(c) \neq \perp$, and write $o(c) = \alpha$. If no such α exists, we write $o(c) = *$. Thus $e^*(c) = e^{o(c)}(c)$. The running time of P on input \bar{x} is defined to be $o(c)$, where c is the start configuration (I_1, \bar{x}) . We denote this by $\text{TIME}(P, \bar{x})$. The program P is said to accept \bar{x} if $e^*(c) = 1$ and reject \bar{x} if $e^*(c) = 0$, where c is the start configuration of P on \bar{x} ; in either case $\text{TIME}(P, \bar{x}) < *$ and P is said to halt on \bar{x} . The relation accepted by P is the set

$$\{\bar{x} \mid P \text{ accepts } \bar{x}\}.$$

If $\text{TIME}(P, \bar{x}) = *$, then $e^*(c) = \perp$ and P does not halt on \bar{x} . Call a program P β -time bounded if $\text{TIME}(P, \bar{x}) \leq \beta$ for all inputs \bar{x} accepted by P (P need not halt on other inputs). Define the complexity class

$$\tau(\alpha) = \bigcup_{\beta < \alpha} (\text{relations accepted by } \beta\text{-time bounded IND programs}).$$

4. THE SHUFFLE CONSTRUCTION

If A and \bar{A} are both r.e., then A can be proved recursive by constructing a Turing machine T_3 which simulates steps of T_1 and T_2 alternately, where T_1 accepts A and T_2 accepts \bar{A} . In this section we give a similar construction for IND programs. This construction will be applied in the proofs of

Theorems 2 and 3. Theorem 1 corresponds to the stage comparison theorem of inductive definability (Manna, 1970).

Suppose P and Q are programs with *disjoint* sets of variables \bar{x} and \bar{y} , respectively, and statement labels l_1, \dots, l_p and m_1, \dots, m_q , respectively. By adding dummy statements, we can assume without loss of generality that p and q are relatively prime. Now we shuffle the statements of P and Q to get the program PQ with $2pq$ statements labeled by all pairs of the form (\underline{l}_i, m_j) and (l_i, \underline{m}_j) , $1 \leq i \leq p$, $1 \leq j \leq q$, arranged in the order (\underline{l}_1, m_1) , (\underline{l}_2, m_1) , (\underline{l}_2, m_2) , $(\underline{l}_3, m_2), \dots$, $(\underline{l}_p, m_i), (l_1, \underline{m}_i), \dots$, (\underline{l}_p, m_q) , (l_1, \underline{m}_q) . The underline tells which statement of P or Q is the next to be simulated. The statement of PQ labeled by (\underline{l}_i, m_j) is the same as the statement of P labeled by l_i , unless it is a conditional jump

$$l_i: \text{if } R(\bar{x}) \text{ then goto } l_k$$

in which case we take

$$(\underline{l}_i, m_j): \text{if } R(\bar{x}) \text{ then goto } (l_k, \underline{m}_j).$$

A symmetric remark holds for statements labeled (l_i, \underline{m}_j) . Thus PQ simulates steps of P and Q alternately. Since the variables of P and Q are disjoint, these simulations do not interfere with each other. The formal statement of this property involves the relationship between the successor configuration maps N_{PQ} and N_P, N_Q . Observe that there is a natural one-to-one correspondence between the configurations C_{PQ} of PQ and pairs $(c, d) \in C_P \times C_Q \cup C_Q \times C_P$:

$$(\underline{l}_i, m_j, \bar{a}, \bar{b}) \mapsto ((l_i, \bar{a}), (m_j, \bar{b}))$$

$$((l_i, \underline{m}_j), \bar{a}, \bar{b}) \mapsto ((m_j, \bar{b}), (l_i, \bar{a})).$$

The order of the components in $(c, d) \in C_P \times C_Q \cup C_Q \times C_P$ tells which of P, Q is next to be simulated. Hence we will identify elements of C_{PQ} with the corresponding elements of $C_P \times C_Q \cup C_Q \times C_P$. Then by construction of PQ ,

$$N_{PQ}(c, d) = \{(d, c') \mid c' \in N(c)\}$$

where $N(c) = N_P(c)$ if $c \in C_P$, $N_Q(c)$ if $c \in C_Q$.

The following lemma and theorem state that the label assigned by e^* to a particular configuration $(c, d) \in C_P \times C_Q$ of PQ is either the one assigned to c by P or the one assigned to d by Q , depending on which is labeled sooner. Moreover (c, d) is labeled in PQ within at most double the time it takes to label either c or d in P or Q , respectively.

LEMMA.

$$\begin{aligned} e^\alpha(c, d) &= e^{\lceil \alpha/2 \rceil}(c) && \text{if } o(c) \leq o(d), \\ &= e^{\lfloor \alpha/2 \rfloor}(d) && \text{if } o(c) > o(d). \end{aligned}$$

Remark on notation. $\lceil \alpha/2 \rceil$ denotes the least ordinal β such that $2\beta \geq \alpha$, and $\lfloor \alpha/2 \rfloor$ denotes the greatest ordinal β such that $2\beta \leq \alpha$. These ordinals exist and are unique; if α is of the form $\lambda + n$ where λ is a limit ordinal and n is finite, then $\lceil \alpha/2 \rceil = \lambda + \lceil n/2 \rceil$ and $\lfloor \alpha/2 \rfloor = \lambda + \lfloor n/2 \rfloor$, where $\lceil n/2 \rceil$ and $\lfloor n/2 \rfloor$ denote the usual ceiling and floor functions on the integers. A simple consequence of this definition is

$$\begin{aligned} \lceil \lambda/2 \rceil &= \lfloor \lambda/2 \rfloor = \lambda && \text{if } \lambda \text{ is a limit ordinal,} \\ \lceil \alpha/2 \rceil &= \lfloor (\alpha + 1)/2 \rfloor, \\ \lfloor \alpha/2 \rfloor + 1 &= \lceil (\alpha + 1)/2 \rceil. \end{aligned}$$

Proof of Lemma. At 0 or limit ordinals,

$$\begin{aligned} e^\lambda(c, d) &= \bigcup_{\alpha < \lambda} e^\alpha(c, d) \\ &= \bigcup_{\alpha < \lambda} e^{\lceil \alpha/2 \rceil}(c) && \text{if } o(c) \leq o(d), \\ &= \bigcup_{\alpha < \lambda} e^{\lfloor \alpha/2 \rfloor}(d) && \text{if } o(c) > o(d), \\ &= e^\lambda(c) && \text{if } o(c) \leq o(d), \\ &= e^\lambda(d) && \text{if } o(c) > o(d). \end{aligned}$$

For successor ordinals $\alpha + 1$, we proceed by cases, depending on the form of $\text{stmt}(c, d) = \text{stmt}(c)$. If $\text{stmt}(c, d) = \text{accept}$, then $o(c) = 1 \leq o(d)$, and

$$e^{\alpha+1}(c, d) = 1 = e^{\lceil (\alpha+1)/2 \rceil}(c).$$

The case $\text{stmt}(c, d) = \text{reject}$ is similar. For $\text{stmt}(c, d) = \text{if...then...}$, let c' be the unique successor of c . Then (d, c') is the unique successor of (c, d) . Moreover, $o(c) = o(c') + 1$ and $e^{\alpha+1}(c) = e^\alpha(c')$. Then

$$\begin{aligned} e^{\alpha+1}(c, d) &= e^\alpha(d, c') \\ &= e^{\lceil \alpha/2 \rceil}(d) && \text{if } o(d) \leq o(c'), \\ &= e^{\lfloor \alpha/2 \rfloor}(c') && \text{if } o(d) > o(c') \end{aligned}$$

by induction hypothesis,

$$\begin{aligned}
e^{\alpha+1}(c, d) &= e^{\perp(\alpha+1)/2\downarrow}(d) && \text{if } o(d) < o(c), \\
&= e^{\perp\alpha/2\downarrow+1}(c) && \text{if } o(d) \geq o(c), \\
&= e^{\perp(\alpha+1)/2\downarrow}(d) && \text{if } o(d) < o(c), \\
&= e^{\lceil(\alpha+1)/2\rceil}(c) && \text{if } o(d) \geq o(c).
\end{aligned}$$

The last and most involved case is $\text{stmt}(c, d) = y \leftarrow \exists$ or $\text{stmt}(c, d) = y \leftarrow \forall$. We prove the case $y \leftarrow \exists$, since the two cases are symmetric.

$$\begin{aligned}
e^{\alpha+1}(c, d) &= \bigvee_{c' \in N(c)} e^{\alpha}(d, c') \\
&= \bigvee_{\substack{c' \in N(c) \\ o(d) \leq o(c')}} e^{\lceil\alpha/2\rceil}(d) \vee \bigvee_{\substack{c' \in N(c) \\ o(d) > o(c')}} e^{\perp\alpha/2\downarrow}(c') \\
&= \sigma \vee \tau, \tag{*}
\end{aligned}$$

where σ (resp. τ) is the left-hand (resp. right-hand) disjunct. Either

(a) $\forall c' \in N(c) \ o(d) > o(c')$, in which case σ vanishes and (*) evaluates to

$$e^{\perp\alpha/2\downarrow+1}(c) = e^{\lceil(\alpha+1)/2\rceil}(c);$$

(b) $\forall c' \in N(c) \ o(d) \leq o(c')$, in which case τ vanishes and (*) evaluates to

$$e^{\lceil\alpha/2\rceil}(d) = e^{\perp(\alpha+1)/2\downarrow}(d);$$

(c) neither (a) nor (b), in which case (*) becomes

$$e^{\perp(\alpha+1)/2\downarrow}(d) \vee \tau.$$

Moreover,

$$\tau \leq \bigvee_{c' \in N(c)} e^{\perp\alpha/2\downarrow}(c') = e^{\lceil(\alpha+1)/2\rceil}(c).$$

Suppose first that $o(c) \leq o(d)$. If $e^{\lceil(\alpha+1)/2\rceil}(c) = 0$, then (a) above holds. If $e^{\lceil(\alpha+1)/2\rceil}(c) = 1$, then $\exists c' \in N(c) \ o(c) = o(c') + 1$ and $e^{\perp\alpha/2\downarrow}(c) = 1$, thus $\tau = 1$ and (*) evaluates to 1. If $e^{\lceil(\alpha+1)/2\rceil}(c) = \perp$, then $e^{\perp(\alpha+1)/2\downarrow}(d) = \perp$ since $o(c) \leq o(d)$, and $\tau \leq \perp$. If (a) holds, there is nothing to prove. (b) can only hold when $o(c) = o(d) = *$, in which case (*) evaluates to $e^{\perp(\alpha+1)/2\downarrow}(d) = e^{\lceil(\alpha+1)/2\rceil}(c) = \perp$. If (c) holds, then (*) evaluates to $\perp \vee \tau = \perp$. Thus

$$e^{\alpha+1}(c, d) = e^{\lceil(\alpha+1)/2\rceil}(c)$$

whenever $o(c) \leq o(d)$.

Now suppose $o(c) > o(d)$. Under no circumstances can (a) occur; if (b) occurs, we are done; therefore assume (c). Then $\tau \neq 1$, otherwise $o(c) \leq o(d)$. Thus $\tau \leq \perp$ and (*) evaluates to

$$e^{L(\alpha+1)/2 \downarrow}(d) \vee \tau.$$

If this is not equal to $e^{L(\alpha+1)/2 \downarrow}(d)$, then $e^{L(\alpha+1)/2 \downarrow}(d) = 0$ and $\tau = \perp$. But this is impossible, because if $e^{\lceil \alpha/2 \rceil}(d) \neq \perp$ and $o(c') < o(d)$ then $e^{L\alpha/2 \downarrow}(c') \neq \perp$. Thus

$$e^{\alpha+1}(c, d) = e^{L(\alpha+1)/2 \downarrow}(d)$$

whenever $o(c) > o(d)$. ■

THEOREM 1.

- (i) $e^*(c, d) = e^*(c)$ if $o(c) \leq o(d)$,
 $= e^*(d)$, if $o(c) > o(d)$.
- (ii) $o(c, d) = \min(2 \cdot o(c) - 1, 2 \cdot o(d))$.

Remark. The expression $2 \cdot o(c) - 1$ in (ii) makes sense, because c can only become first properly labeled at a successor ordinal.

Proof.

$$\begin{aligned} \text{(i)} \quad e^*(c, d) &= \bigsqcup_{\alpha} e^{\alpha}(c, d) \\ &= \bigsqcup_{\alpha} e^{\lceil \alpha/2 \rceil}(c) && \text{if } o(c) \leq o(d), \\ &= \bigsqcup_{\alpha} e^{L\alpha/2 \downarrow}(d) && \text{if } o(c) > o(d), \\ &= e^*(c) && \text{if } o(c) \leq o(d), \\ &= e^*(d) && \text{if } o(c) > o(d). \end{aligned}$$

- (ii) Let $\mu\alpha \cdot R(\alpha)$ denote the least α such that $R(\alpha)$.

$$\begin{aligned} o(c, d) &= \mu\alpha \cdot e^{\alpha}(c, d) \neq \perp \\ &= \mu\alpha \cdot e^{\lceil \alpha/2 \rceil}(c) \neq \perp && \text{if } o(c) \leq o(d), \\ &= \mu\alpha \cdot e^{L\alpha/2 \downarrow}(d) \neq \perp && \text{if } o(c) > o(d), \\ &= 2 \cdot o(c) - 1 && \text{if } o(c) \leq o(d), \\ &= 2 \cdot o(d) && \text{if } o(c) > o(d), \\ &= \min(2 \cdot o(c) - 1, 2 \cdot o(d)). \quad \blacksquare \end{aligned}$$

If the variables of P and Q are *not* disjoint, define the shuffle PQ as follows: rename the variables of P to get a program P' having no variables in common with Q . Let x_1, \dots, x_k be the variables common to P and Q and let y_1, \dots, y_k be their replacements in P' . Define PQ to be the program which assigns $y_i \leftarrow x_i$, $1 \leq i \leq k$, then runs $P'Q$.

COROLLARY 1. *Let P, Q be two programs with a common set \bar{x} of input variables. Then PQ accepts (rejects) \bar{x} iff either*

- (i) $\text{TIME}(P, \bar{x}) \leq \text{TIME}(Q, \bar{x})$ and P accepts (rejects) \bar{x} ; or
- (ii) $\text{TIME}(P, \bar{x}) > \text{TIME}(Q, \bar{x})$ and Q accepts (rejects) \bar{x} .

5. MAIN RESULTS

THEOREM 2. (i) *IND programs accept precisely the relations definable by first-order induction;*

(ii) *IND programs which halt on all inputs accept precisely the hyper-elementary (or inductive, coinductive) relations;*

(iii) *loop-free IND programs accept precisely the first-order definable relations.*

Proof. (i) Every inductively definable relation is given by a first-order formula $\varphi(S, \bar{x})$ with free variables $\bar{x} = x_1, \dots, x_n$, an n -ary predicate symbol S occurring only positively (i.e., under an even number of negations) in φ , and some constants $\bar{a} = a_1, \dots, a_m$, $m \leq n$. The *fixpoint* defined by φ is the least S^* such that

$$S^* = \{\bar{d} \mid \models \varphi(S^*, \bar{d})\}$$

where \models denotes truth in the structure under consideration. The inductive relation defined by φ, \bar{a} is the $(n-m)$ -ary relation $S^*(a_1, \dots, a_m, x_{m+1}, \dots, x_n)$.

Given such an inductive definition φ, \bar{a} , a program to accept all (x_{m+1}, \dots, x_n) satisfying $S^*(a_1, \dots, a_m, x_{m+1}, \dots, x_n)$ works as follows: First it assigns a_i to x_i , $1 \leq i \leq m$, and then enters a loop labeled l_1 which determines whether $S^*(\bar{x})$. Within the loop it decomposes $\varphi(S, \bar{x})$, using $y \leftarrow \forall$ and $y \leftarrow \exists$ to eliminate quantifiers, $l_1 \wedge l_2$, $l_1 \vee l_2$, $\neg l_1$ to eliminate logical connectives, and conditionals for atomic formulas; this leaves only occurrences of $S(\bar{y})$, which are handled by assigning \bar{y} to \bar{x} followed by an unconditional jump back to l_1 .

Conversely, if P is any IND program with n statements l_1, \dots, l_n and program variables \bar{x} , let a, b be distinct elements of the domain (which we

assume is nontrivial), and let y_1, \dots, y_n be new variables not occurring in P . We will use y_1, \dots, y_n, a, b to simulate a "program counter." Let \bar{a}_i denote the sequence of a s and b s of length n such that there is an a in the i th position, and the rest are b s, and let $\bar{y} = \bar{a}_i$ abbreviate the relation

$$y_i = a \wedge \bigwedge_{j \neq i} y_j = b.$$

The relation $\bar{y} = \bar{a}_i$ will simulate that P is about to execute l_i .

We now construct a first-order formula $\varphi(S, \bar{x}, \bar{y})$ with $S = S(\bar{x}, \bar{y})$ occurring only positively in φ , such that S^* describes exactly the nodes of the computation tree that are labeled 1. Let

$$\begin{aligned} \varphi_i &= \text{true} \\ &= \text{false} \\ &= \exists x_j S(\bar{x}, \bar{a}_{i+1}) \\ &= \forall x_j S(\bar{x}, \bar{a}_{i+1}) \\ &= (R(\bar{x}) \wedge S(\bar{x}, \bar{a}_k)) \vee (\neg R(\bar{x}) \wedge S(\bar{x}, \bar{a}_{i+1})) \end{aligned}$$

if l_i is a statement of the form

$$\begin{aligned} l_i &: \text{accept} \\ l_i &: \text{reject} \\ l_i &: x_j \leftarrow \exists \\ l_i &: x_j \leftarrow \forall \\ l_i &: \text{if } R(\bar{x}) \text{ then } l_k, \end{aligned}$$

respectively. Let

$$\varphi(S, \bar{x}, \bar{y}) = \bigvee_i \bar{y} = \bar{a}_i \wedge \varphi_i.$$

The relation accepted by P is then $S^*(\bar{x}, \bar{a}_1)$, where l_1 is the first statement of P .

(ii) Any program P accepting S which halts on all inputs has a dual \bar{P} which also halts on all inputs, and accepts the complement of S . Thus by (i), S is both inductive and coinductive. Conversely, suppose the set S is both inductive and coinductive. By (i), there are programs P and Q accepting S and \bar{S} , respectively. Modify P and Q so that they never reject, by replacing all statements l : *reject* with l : *goto l*. By Corollary 1, the shuffle $P\bar{Q}$ accepts S and rejects \bar{S} .

(iii) It has already been argued in Section 2 that every first-order definable relation is computed by a loop-free program. The converse is obtained by observing that every loop-free program is equivalent to one with only forward jumps; a formula is now easy to construct. ■

Observe from the proof of Theorem 2(i) that there is a strong connection between the running times of IND programs and the ordinals at which inductive definitions close (see Moschovakis, 1974). The *closure ordinal* κ^A of a structure A is defined in *ibid.* as the supremum of closure ordinals of all possible inductive definitions. By the proof of Theorem 2(i) we see that (for infinite structures) this is just the supremum of running times of IND programs in A . The following theorem relates these concepts to the complexity classes $\tau(\alpha)$ defined in Section 3.

THEOREM 3. (i) $\tau(\omega) = \{\text{first-order definable relations}\}$.

(ii) $\tau(\kappa^A) = \{\text{hyper-elementary relations}\}$.

Remark. Part (ii) is related to the *closure theorem* of inductive definability theory (Moschovakis, 1974, p. 33).

Proof. (i) Clearly, any loop-free program can run for only finitely many steps, independent of the input. Conversely, any c -time bounded program, $c < \omega$, can be made to halt on all inputs by shuffling it with a "clock," i.e., a program that on all inputs runs for $c + 1$ steps, then rejects. The resulting program now has a finite, uniform time bound d , independent of the input. But any such program is equivalent to a loop-free program obtained by unwinding the loops $d + 1$ times. The result follows from Theorem 2(iii).

(ii) (\supseteq) Let P be a program that halts on all inputs. Let P_1 be P modified so as never to reject, as in the proof of Theorem 2(ii), and let P_2 be \bar{P} modified so as never to reject, where \bar{P} is the dual of P . Then $P_1 P_2$ accepts all inputs and $\forall \bar{x}, \text{TIME}(P, \bar{x}) \leq \text{TIME}(P_1 P_2, \bar{x})$. Let P_3 be the program which chooses the input universally by executing $y \leftarrow \forall$ for all input variables, then executes $P_1 P_2$. Then $\text{TIME}(P_3, \bar{x})$ is a constant β independent of the input, and $\forall \bar{x} \text{TIME}(P, \bar{x}) \leq \beta < \kappa^A$.

(\subseteq) Let P be β -time bounded, $\beta < \kappa^A$. If we can construct an α -clock, $\beta < \alpha$, then it can be shuffled with P to give a program accepting the same set as P , but always halting. The result then follows from Theorem 2(ii). Since $\beta < \kappa^A$, there exists a program Q which runs for time $\alpha > \beta$ on some input \bar{x} . Let Q_1 assign \bar{x} to all input variables, then run Q . Q_1 halts on all inputs in time exactly $\alpha + c$ and either accepts all inputs or rejects all inputs, so either Q_1 or \bar{Q}_1 gives an appropriate clock. ■

6. AN APPLICATION TO DYNAMIC LOGIC

The programming language IND originally arose in our attempt to clarify a result of Meyer and Parikh (1980) on the relative expressibility of four variants of first-order dynamic logic (DL), namely DL_{reg} , DL_{cf} , DL_{ft} , and DL_{re} . Programs of DL_{re} are all r.e. sets of sequences of assignments $x := t$ and tests $\varphi(\bar{x})$, called *seqs*, where t is a term and φ a formula of DL_{re} . One obtains DL_{reg} , DL_{cf} , and DL_{ft} by allowing, respectively, only regular expressions or flowchart programs (so that the set of seqs is regular), recursion schemes (so that the set of seqs is context free), or r.e. sets of seqs, but each with at most finitely many distinct tests.

Meyer and Parikh prove that DL_{reg} is strictly less expressive than DL_{re} (in symbols, $DL_{reg} < DL_{re}$) by the following sequence:

$$DL_{reg} \leqslant_{(a)} DL_{cf} \leqslant_{(b)} DL_{ft} \leqslant_{(c)} L_{ba} <_{(d)} L_{\omega_1\omega}^{CK} \equiv_{(e)} DL_{re} \quad (**)$$

where $L_{\omega_1\omega}^{CK}$ is infinitary first-order logic with r.e. disjunctions, and L_{ba} is the same language restricted to bounded quantifier alternation. The bulk of the proof is devoted to (d), which uses an Ehrenfeucht–Frassé argument to show that L_{ba} cannot distinguish between the ordinals ω^ω and $\omega^\omega \cdot 2$, while $L_{\omega_1\omega}^{CK}$ can define any recursive ordinal up to isomorphism.

In this part of their paper, all resemblance to dynamic logic has been lost. This was taken as evidence in support of the stand that dynamic logic should enjoy less autonomy, and one should do all one's work in infinitary logic (Meyer and Tiuryn, 1981). We disagree, for the simple reason that $L_{\omega_1\omega}^{CK} \equiv L_{ba}$ in virtually every structure arising in computer science (e.g., the natural numbers N , any recursively defined data type, or any structure whose elements are all named by closed terms). This is because every $L_{\omega_1\omega}^{CK}$ formula is equivalent to a quantifier-free formula, by replacing $\exists x\varphi(x)$ with $\bigvee_t \varphi(t)$, where the join is over the set of closed terms. One's intuition is still that $DL_{reg} < DL_{re}$, even restricted to such structures. Our results of this section show that $DL_{ft} < DL_{re}$ on any *acceptable* structure (Moschovakis, 1974) (or *arithmetic universe* (Harel, 1979)). These structures contain a first-order definable copy of the natural numbers and first-order predicates for coding and decoding sequences of elements into single elements. This allows assigning codes or Gödel numbers to programs and formulas so that they can be decoded and manipulated by other programs and formulas. The proof reveals the computational power of the various versions of DL in terms of the complexity classes $\tau(\alpha)$.

THEOREM 4. *On any acceptable structure,*

$$(i) \quad DL_{reg} \equiv DL_{cf} \equiv DL_{ft} \equiv \tau(\omega)$$

- (ii) $DL_{re} \equiv \tau(\omega_1^{CK})$
- (iii) $\tau(\omega_1^{CK}) - \tau(\omega) \neq \emptyset$

where ω_1^{CK} is the first nonrecursive ordinal.

For example, on N , whose closure ordinal is ω_1^{CK} , $DL_{re} \equiv \Delta_1^1$ and $DL_{ft} \equiv$ (first-order definable relations). This follows from Theorems 3, 4, and Kleene's Theorem ($\Delta_1^1 =$ hyper elementary on N).

Proof. (i), (ii), (\equiv) This direction does not need the assumption of acceptability. The case (i) follows from Theorem 3(i) and the fact that DL contains first-order logic. Similarly, on any structure \mathcal{A} , $\tau(\omega_1^{CK}) \leq DL_{re}$, since if P is any IND program and α any recursive ordinal, there is a DL_{re} -formula φ_α such that

$$\mathcal{A} \models \varphi_\alpha(c) \quad \text{iff} \quad e^\alpha(c) = 1$$

for any configuration c , defined recursively by

$$\begin{aligned} \varphi_0 &\leftrightarrow \text{false} \\ \varphi_{\alpha+1}(c) &\leftrightarrow \text{stmt}(c) = \text{accept} \\ &\vee \text{stmt}(c) = y \leftarrow \exists \wedge \exists d \in N(c) \varphi_\alpha(d) \\ &\vee \text{stmt}(c) = y \leftarrow \forall \wedge \forall d \in N(c) \varphi_\alpha(d) \\ &\vee \text{stmt}(c) = \text{if...then...} \wedge \exists d \in N(c) \varphi_\alpha(d) \\ \varphi_\lambda &\leftrightarrow \langle \{ \varphi_\alpha? \mid \alpha < \lambda \} \rangle \text{ true, } \lambda \text{ a limit ordinal.} \end{aligned}$$

The crucial point of this definition is that it is *effective*, in the sense that there is a recursive function r such that $r("a") = "\varphi_\alpha"$, where " a " and " φ " denote codes (fixed in advance) for recursive ordinals and DL_{re} formulas. This fact is needed in the definition of φ to insure that the set $\{ \varphi_\alpha? \mid \alpha < \lambda \}$ is r.e., so that φ_λ will be a DL_{re} formula. Now, if $A \in \tau(\omega_1^{CK})$, then A is accepted by an IND program P which is α -time bounded for some recursive ordinal α ; thus for any input $\bar{a} \in \mathcal{A}^k$,

$$P \text{ accepts } \bar{a} \leftrightarrow e^\alpha(l_1, \bar{a}) = 1 \leftrightarrow \mathcal{A} \models \varphi_\alpha(l_1, \bar{a}).$$

Then $\varphi_\alpha(l_1, \bar{x})$ is a DL_{re} formula defining A .

(i), (ii), (\subseteq) We describe first an IND program to decide the satisfiability of DL_{re} formulas in \mathcal{A} , consisting of a main program $SATIS(p, x)$ and subroutine $COMPUTE(q, x, y)$. $SATIS("a", "\bar{a}')$ will determine if $\mathcal{A}, \bar{a} \models a$ and $COMPUTE("a", "\bar{a}, "\bar{b}')$ will determine if state \bar{a} goes to state \bar{b} under program π , where " a ," " b " are codes of

sequences \bar{a} , \bar{b} of elements of \mathcal{A} , and “ φ ” and “ π ” are codes of a DL_{re} formula φ and a DL_{re} program π .

Initially, SATIS (“ φ ,” “ \bar{a} ”) assigns a code for the list of free variables of φ (available from “ φ ”) to a variable v , and assigns “ \bar{a} ” to w . This models the assignment of the values \bar{a} to the list of variables in v , in the same order. It now proceeds by cases, depending on the form of φ . If $\varphi = \psi \wedge \sigma$, it uses the program construct \wedge defined in Section 2 to check both ψ and σ , and similarly for \vee , \neg . If $\varphi = \exists y\psi$, it executes $z \leftarrow \exists$; then, if the DL_{re} variable y is in the list v , it modifies the corresponding value in the list w to the value of z ; otherwise, it appends the name of y to the list v and the value of z to the list w . If $\varphi = \forall y\psi$ it does the same, using $z \leftarrow \forall$ instead of $z \leftarrow \exists$. If $\varphi = \langle \pi \rangle \psi$, since

$$\mathcal{A}, \bar{a} \models \langle \pi \rangle \psi \quad \text{iff} \quad \exists \bar{b} \quad \bar{a} \text{ goes to } \bar{b} \text{ under } \pi \text{ and } \mathcal{A}, \bar{b} \models \psi,$$

SATIS executes $z \leftarrow \exists$ and interprets the result as a code “ \bar{b} .” It then calls SATIS(“ ψ ,” “ \bar{b} ”) and COMPUTE(“ π ,” “ \bar{a} ,” “ \bar{b} ”) in parallel, using \wedge .

Finally, if $\varphi = R(\bar{x})$ where $R(\bar{x})$ is atomic, it picks out the current values in the list w corresponding to DL_{re} variables \bar{x} and assigns them to IND variables \bar{y} , then executes

if $R(\bar{y})$ then accept else reject.

COMPUTE(“ π ,” “ \bar{a} ,” “ \bar{b} ”) determines whether state \bar{a} goes to state \bar{b} under DL_{re} program π . Recall that a program π consists of an r.e. set of seqs; each seq is a finite sequence $s_0; \dots; s_{k-1}$ for some k ; and each s_i is either an assignment $y := t$ or a test $\psi?$. The code “ π ” gives a Gödel number for the set of seqs, and \bar{a} goes to \bar{b} under π iff \bar{a} goes to \bar{b} under some seq of π . COMPUTE chooses a seq existentially using $i \leftarrow \exists$, and then tries to determine if the i th seq of π , say $\text{seq}_i = s_0; \dots; s_{k-1}$, takes \bar{a} to \bar{b} . It could do this by starting from $\bar{a}_0 = \bar{a}$, deterministically applying s_0, s_1, \dots, s_{k-1} in succession to get a sequence $\bar{a}_1, \dots, \bar{a}_k$ of intermediate states, and accepting if $\bar{a}_k = \bar{b}$. However, for a later application, it will be better to keep COMPUTE loop-free. Thus, the program instead guesses a code for the entire sequence $\bar{a}_0, \bar{a}_1, \dots, \bar{a}_k$ with a single $z \leftarrow \exists$, and then determines whether s_j takes \bar{a}_j to \bar{a}_{j+1} , $0 \leq j < k$, by executing $j \leftarrow \forall$; if j is not the code for a natural number $< k$ then accept; check if s_j takes \bar{a}_j to \bar{a}_{j+1} . For s_j of the form $y := t$, the check is straightforward. For s_j of the form $\psi?$, the program checks whether $\bar{a}_j = \bar{a}_{j+1}$, then whether $\mathcal{A}, \bar{a}_j \models \psi$ by a recursive call to SATIS.

Holding “ φ ” fixed, SATIS(“ φ ,” “ \bar{a} ”) accepts (the codes of) the set defined by φ . The theorem is now proved by analyzing the time complexity of SATIS and COMPUTE on fixed φ . All encoding and decoding operations can be done without loops, since they are first-order definable. The choice

of seq_i in COMPUTE can be done without a loop since π is r.e. and thus first-order definable. We were careful to avoid loops in the processing of a seq in COMPUTE. Thus each iteration of SATIS and COMPUTE takes constant time before it recurs on a subformula; therefore there is a constant c such that

$$\forall \bar{a} \text{ TIME}(\text{SATIS}(\varphi, \bar{a})) \leq c \cdot h(\varphi)$$

where $h(\varphi)$ is the *height* of φ , defined by

$$\begin{aligned} h(\varphi) &= 1, \quad \varphi \text{ atomic,} \\ h(\varphi \vee \psi) &= h(\varphi \wedge \psi) = \max\{h(\varphi), h(\psi)\} + 1, \\ h(\forall x \varphi) &= h(\exists x \varphi) = h(\neg \varphi) = h(\varphi) + 1 \\ h(\langle \pi \rangle \varphi) &= \max\{h(\pi), h(\varphi)\} + 1, \\ h(\pi) &= \sup\{h(\sigma) \mid \sigma \text{ a seq of } \pi\} + 1, \\ h(\sigma) &= \sup\{h(\varphi) \mid \varphi? \text{ a test of } \sigma\} + 1, \quad \sigma \text{ a seq.} \end{aligned}$$

Thus it remains to show that

$$\begin{aligned} h(\varphi) &< \omega_1^{CK}, & \varphi \text{ in } \text{DL}_{\text{re}}, \\ h(\varphi) &< \omega, & \varphi \text{ in } \text{DL}_{\text{ft}}. \end{aligned}$$

The former follows from the fact that there is a recursive code “ φ ” for each φ in DL_{re} , and h is effective with respect to this code. The latter follows from the fact that the suprema in the definition of $h(\pi)$ and $h(\sigma)$ are finite, since there are only finitely many tests.

(iii) This is a straightforward diagonalization. Construct an IND program P which, on input “ $\varphi(x)$,” $\varphi(x)$ a first-order formula with one free variable x , accepts iff $\neg \varphi(\ulcorner \varphi \urcorner)$. P runs for time $c \cdot h(\varphi) < \omega$, so the set it accepts is in $\tau(\omega + 1) \subseteq \tau(\omega_1^{CK})$, and not in $\tau(\omega) = \{\text{first-order-definable sets}\}$ for obvious reasons. ■

7. IND AS A DATA BASE QUERY LANGUAGE

There has been much recent work in the theory of relational data bases. In the relational model, a data base is a collection of finite tables (Codd, 1970) and can be viewed simply as a finite first-order structure $B = (D, R_1, \dots, R_k)$. Queries are (partial) functions from data bases to relations, and a query language is a set of formal expressions defining such functions; (Chandra and Harel, 1980).

Codd (1972) introduced the languages of the relational algebra and calculus, which are equivalent in expressive power. The latter is essentially the first-order language of similarity type $(=, R_1, \dots, R_k)$. In (Aho and Ullman, 1979) it was pointed out that many useful queries definable naturally by least fixpoints of first-order formulas, such as the transitive closure of a binary relation, are not first-order definable, and it was suggested therein that the first-order language of Codd (1972) be augmented with an appropriate least fixpoint operator. In such a language the transitive closure of R would be the least fixpoint S of $S = R \cup R \circ S$, where \circ is relational composition. A formal version of such an extension was subsequently supplied in (Chandra and Harel, 1982a), where fixpoint operators were allowed to alternate with any number of first-order constructs. A hierarchy of height ω^2 of sets of queries is defined in *ibid.*, in which those queries at level $\omega \cdot i$ are obtained by applying a least fixpoint operator to queries at lower levels. The set of queries constituting the entire hierarchy is termed FP, for *fixpoint queries*.

It is shown in Chandra and Harel (1982a) that FP is a very restricted subset of the set of all computable queries (Chandra and Harel, 1980); in particular, all queries in FP are polynomial-time computable. There is also a close correspondence with the queries definable by Kowalski's logic programs; see Kowalski (1974), Gallaire and Minker (1978), and Chandra and Harel (1982b). However, it was left as an open problem in Chandra and Harel (1982a) whether there is a natural computational query language for defining the fixpoint queries.

At this point, one observes that the least fixpoint operator as defined in Aho and Ullman (1979) and Chandra and Harel (1982a) corresponds exactly to an inductive definition as defined in Moschovakis (1974), so that a single fixpoint operator applied to a first-order formula corresponds to a first-order inductive definition. Recently, however, Immerman (1982) has shown that the hierarchy of fixpoint queries in fact collapses down to level ω when finite structures are considered. In other words, all queries in FP are definable by a single application of a fixed-point operator to a first-order formula. Hence, we obtain

LEMMA. *A relational function on finite structures is in FP iff it is uniformly first-order inductively definable (i.e., there is a single first-order inductive definition which, given the input structure, defines the output relation).*

THEOREM. *IND defines precisely the fixed-point queries on relational data bases.*

We might remark that the $x \leftarrow \forall$ statement of IND and the parallel method of execution implied by its semantics reminds one of the "for all

tuples t in relation R ” construct used in some real query language with parallel execution semantics; see (Aho and Ullman, 1979, Sect. 7). It remains to be seen whether a rigorous definition of the semantics of such a language, together with the dual “for some tuple t in R ,” yields a language equivalent to IND.

RECEIVED: January 17, 1983; ACCEPTED: February 28, 1985

REFERENCES

- AHO, A. V., AND ULLMAN J. D., Universality of data retrieval languages. *Proc. 6th ACM Symp. on Principles of Programming Languages*, Jan. 1979, 110–117.
- APT, K., AND PLOTKIN G. A Cook’s Tour of Countable Nondeterminism, ICALP ’81.
- BARWISE, J. *Admissible Sets & Structures*. Springer-Verlag 1975.
- CHANDRA, A. K., AND HAREL, D. Computable queries for relational data bases. *JCSS* **21**; 2, (1980), 156–178.
- CHANDRA, A. K., AND HAREL, D. Structure and Complexity of Relational Queries, *JCSS* **25**; 1 (1982a), 99–128.
- CHANDRA, A. K., AND HAREL, D. Horn clauses and the fixpoint query hierarchy. *SIGACT-SIGMOD Symp. on Principles of Data Base Systems*, March, 1982b.
- CHANDRA, A. K., KOZEN, D., AND STOCKMEYER, L. Alternation, *J. ACM*, (1981), 114–133.
- CODD, E. F. A relational model for large shared data bases. *CACM* **13**; 6, (1970), 377–387.
- CODD, E. F. Relational completeness of data base sublanguages. In *Data Base Systems* (Rustin, ed.), Prentice Hall, 1972.
- COOK, S. A. Soundness and Completeness of an Axiom System for Program Verification, *SIAM J. on Computing* **7**; 1, (1978).
- GALLAIRE, H., AND MINKER J. (eds.), *Logic and Data Bases*, Plenum, New York (1978).
- GRUMBERG, O., FRANCEZ, N., MAKOWSKY, J., AND DE ROEVER, W. A proof Rule for fair Termination of Guarded Commands. In *Algorithmic Languages* (de Bakker, van Vliet, eds.) North Holland (1981), 399–416.
- HAREL, D. *First-order Dynamic Logic*, Lecture Notes in Computer Science 68, Springer-Verlag 1979.
- IMMERMAN, N. Relational queries computable in polynomial time. *14th ACM Symp. on Theory of Computing*, May 1982.
- KOWALSKI, R. A. Predicate logic as a programming language. *Proc. IFIP74*, North-Holland 1974, 556–574.
- LEHMANN, D., PNUELI, A., AND STAVI, J. Impartiality, Justice and Fairness: The Ethics of Concurrent Termination, ICALP ’81.
- MANNA, Z. The Correctness of Nondeterministic Programs, *Artificial Intelligence* **1** (1970), 1–26.
- MEYER, A. R., AND TIURYN, J. A Note of Equivalence Among Logics of Programs, *Proc. Workshop on Logics of Programs 1981*, Lecture Notes in Computer Science 131, Springer-Verlag, 282–299.
- MEYER, A. R., AND PARIKH, R. Definability in Dynamic Logic, *Proc. 12th ACM Symp. on Theory of Computing* (1980), 1–8.
- MEYER, A. R., AND WINKELMANN, K. On the Expressive Power of Dynamic Logic, *Proc. 12th ACM Symp. on Theory of Computing* (1979), 167–175.

- MIRKOWSKA, G. On Formalized Systems of Algorithmic Logic, *Bull. Acad. Pol. Sci., Ser. Math. Astr. Phys.* **22** (1974), 421–428.
- MOSCHOVAKIS, Y. N. *Elementary Induction on Abstract Structures*, North-Holland, 1974.
- PRATT, V. Semantical Considerations on Floyd-Hoare Logic. *Proc. 17th IEEE Symp. on Found. of Comp. Science* (1976), 109–121.