



Innovations in computational type theory using Nuprl

S.F. Allen^{*}, M. Bickford, R.L. Constable, R. Eaton, C. Kreitz,
L. Lorigo, E. Moran

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

Available online 15 December 2005

Abstract

For twenty years the Nuprl (“new pearl”) system has been used to develop software systems and formal theories of computational mathematics. It has also been used to explore and implement computational type theory (CTT)—a formal theory of computation closely related to Martin-Löf’s intuitionistic type theory (ITT) and to the calculus of inductive constructions (CIC) implemented in the Coq prover.

This article focuses on the theory and practice underpinning our use of Nuprl for much of the last decade. We discuss innovative elements of type theory, including new type constructors such as unions and dependent intersections, our theory of classes, and our theory of event structures.

We also discuss the innovative architecture of Nuprl as a distributed system and as a transactional database of formal mathematics using the notion of abstract object identifiers. The database has led to an independent project called the Formal Digital Library, FDL, now used as a repository for Nuprl results as well as selected results from HOL, MetaPRL, and PVS. We discuss Howe’s set theoretic semantics that is used to relate such disparate theories and systems as those represented by these provers.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Martin-Löf type theory; Dependent intersection types; Union types; Polymorphic subtyping; Logic of events; Formal digital libraries; Computational type theory; Proofs as programs; Program extraction; Tactics

^{*} Corresponding author.

E-mail addresses: sfa@cs.cornell.edu (S.F. Allen), markb@cs.cornell.edu (M. Bickford), rc@cs.cornell.edu (R.L. Constable), eaton@cs.cornell.edu (R. Eaton), kreitz@cs.cornell.edu (C. Kreitz), lolorigo@cs.cornell.edu (L. Lorigo), evan@cs.cornell.edu (E. Moran).

1. Introduction

One of the most profound contributions of computer science to intellectual history is the body of demonstration that computers can automate many intellectual processes (not only calculation), and that by having access to knowledge in digital form they become *indispensable partners in knowledge formation*. The first clear articulation of this idea was in our field, applied mathematical logic and automated reasoning, about fifty years ago.

In its full scope, the idea of automating intellectual processes is profound and applies to knowledge formation in every instance. In its full generality, we might call it computer-mediated knowledge formation, or *digital knowledge* for short. Like logic itself, this is a meta-idea, transforming the very concept of knowledge. It rests on the insights of Euclid, Aristotle, Leibniz, Boole, Kant, Frege, Russell, Brouwer, Cantor, Gödel—giants of human achievement—and legions of logicians, philosophers, mathematicians, and computer scientists. It has transformed the methodology of the physical sciences and created the new methodology of the information sciences of which automated reasoning and formal knowledge management are a key part.

This article is about one ongoing theme in the above story. The PRL research group in applied logic and automated reasoning at Cornell University plays a role in developing computational logics, implementing them, and demonstrating their value in mathematics, theoretical computer science, and software system design and verification. We will highlight our latest contributions in each of these activities and put them into both historical and contemporary perspective, suggesting directions for promising future work. We start with a discussion of the major themes that have influenced our work.

Proofs and rules of thought Aristotle and Euclid gave us the idea of deductive proof as a means to establish truth. The Nuprl (“new pearl”) proof development system is based on a formal account of deduction. Proofs are the main characters and are used not only to establish truth but also to denote evidence, including computational evidence in the form of programs (functional and distributed). The idea of a *proof term* is a key abstraction; it is a meaningful mathematical expression denoting evidence for truth. We advocate using propositions as a means of classifying the evidence; the structure of the proposition determines properties of the evidence [27].

Propositions and logical truth Propositions are abstract objects about which it is sensible to make truth judgments (among others). Some propositions are about abstract objects, some are about physical objects. Logic is concerned with those aspects of truth that are applicable to reasoning in any domain.

A great deal has been written justifying philosophically the laws of logic. The Platonic view is that there are immutable laws in the world of ideas. The Kantian view is that these truths are grounded in properties of the mind. Brouwer focused on the computational powers of the mind (and Chomsky on the brain). It is no surprise that computer scientists are drawn to computational justifications because they provide guidelines for implementing logics and for automating intellectual processes on digital computers.

The PRL group has focused on computational justifications in the spirit of Bishop [20] and Martin-Löf [89–91]. We have applied their ideas in the context of computer science, significantly extending them, and experimenting with them using the Nuprl prover.

Type theories The elementary parts of logic, including Aristotle’s discoveries are relatively unproblematic. For example, the propositional calculus is quite well understood and beautifully presented in classic books such as Smullyan [105], Church [24], and Kleene [69]. It can easily be given several philosophical justifications, including a classical (Boolean) one or the Brouwer/Heyting/Kolmogorov [107] one based on computation for the Constructive (or Intuitionist) fragment. The Gödel dialectica interpretation [107] also applied to the constructive fragment.

Serious and difficult issues are lurking one small step away from elementary logic. Suppose $Prop$ is the domain of discourse for propositional logic; so a tautology such as $P \Rightarrow P$ would be made into a closed proposition by writing $\forall p : Prop.(p \Rightarrow p)$. According to Russell [101], it is not a well defined proposition because giving the meaning of $\forall p : Prop.(p \Rightarrow p)$ requires already having defined what is in $Prop$. He labelled such an act as impredicative concept formation, and created his *theory of types* to control it.

In a contemporary type theory, we would stratify $Prop$ into *levels*, say $Prop_1, Prop_2, \dots$. An expression such as $\forall p : Prop_1.(p \Rightarrow p)$ would belong to $Prop_2$ but not to $Prop_1$. Each $Prop_i$ is regarded as a distinct type. This is a common feature of the type theories implemented in Coq and Nuprl.

Russell’s ramified type theory has led quite directly to those used in several modern theorem provers, e.g., Church’s Simple Theory of Types (STT) a direct descendent of Russell’s *Principia Mathematica* logic (PM), the Calculus of Inductive Constructions (CIC) of Coquand, Huet, Paulin [15,40], Intuitionistic Type Theory (ITT) of Martin-Löf, and Computational Type Theory (CTT) of the Cornell group [29,31] (which was derived from one version of ITT, and will be discussed in detail).

Type theory based provers are natural for computer science because the idea of logical types helped organize the concept of a *data type* as used in programming languages (say Algol68, Pascal, ML, Haskell, Java, C++, etc.). The relationship between logical types and data types is explicitly treated in [29,59,104].

Foundations of mathematics—sets and types *Principia Mathematica* [110] was intended to be a logical foundation for mathematics, but it was more complex and less intuitive than formalizations of set theory in first-order logic. Most mathematicians seem to prefer Zermelo Fraenkel Set Theory (ZF) with the Axiom of Choice (ZFC) as the default formalization. There are other essentially equivalent set theories such as Bernays Gödel (BG).

A distinct approach to foundations is through theories of *logical type*, in which (1) the notions of *function* and *type* are basic, a type being the domain of a function, and (2) relations and properties are functions over their possible arguments, and so the meaning of propositional expressions is treated explicitly in terms of how functions are denoted. Russell’s and Church’s type theories are obviously of this kind.

Among modern type theories used as theories of logical type are CIC, CTT, and ITT, which are adequate foundations (some persons further requiring the law of excluded middle to be added to them to achieve adequacy). Indeed, these three theories have a great deal in

common and are in some sense equivalent and converging. If we call the common theory Computational Types (CT), then CT with Choice, (CTC), is a canonical foundation for mathematics as well. In ZFC, the familiar types such as integers (\mathbb{Z}), reals (\mathbb{R}), functions ($A \rightarrow B$), Cartesian products ($A \times B$) and so forth are defined in terms of sets. In CTC, sets are defined in terms of types.

Foundations of computing Computational type theory (CTT) and the calculus of inductive constructions (CIC) are intended to be foundations for computing. This means that special attention is given in logic design to expressions that can express algorithms and data structures, designing rules adequate for reasoning about them with types. In practice, this means that these theories can provide a semantics for real programming languages and for the logics of programs associated with them. See also [10] for a fine exposition of such methods extending ITT.

For both CTT and CIC, the programming languages native to the theory are simple functional programming languages. In order to define and reason about side effects, modules, and objects requires a great deal of work [41,53,54]. CTT has been extended to reason about events and distributed computing [17]. Here we call it CTT-E, and we describe it briefly later.

Proofs as programs and processes The constructive type theories, CIC, CTT, and ITT, express logic by reducing propositions to types, and so serve as theories of logical type. This reduction is often called the *propositions-as-types principle* [45,61,89]. The usual enunciation of this principle (Martin-Löf’s innovation on it will be described below) is to associate with every proposition P a type, say $[P]$, consisting of objects representing proofs of P . The type $[A \& B]$ for a conjunction could be the Cartesian product $[A] \times [B]$. For an implication $A \Rightarrow B$, the type $[A \Rightarrow B]$ is $[A] \rightarrow [B]$, the type of computable functions from the content of A into the content of B . Likewise, $[\forall x : T. B_x]$ consists of the dependent function space $x : [T] \rightarrow [B_x]$, where $[T]$ is the data type associated with the mathematical type T .¹ This topic is explained well in the literature [12,49,98,106].

When f is a proof of $A \Rightarrow B$, then its content, $[f]$, is an element of $[A] \rightarrow [B]$, that is a computable function. The proof object f is like a program in that its computational content is a computable function. Systems that implement CTT must have operations that *extract* the program from f , that is, operations to construct $[f]$ from f .

As we explain later, for expressions that use the type of events, \mathbb{E} , expressions of the form $\forall e : \mathbb{E}p. \exists e' : \mathbb{E}q. R(e, e')$ describe relations between events at one agent (or location) p and those at another q . A proof that this relation is realizable will denote a computing process that implements it.

Martin-Löf’s innovation on the propositions-as-types principle was to use *computational content* of proofs of P rather than proofs as such. Thus, when explaining the meaning of proposition P as a type, only the computational content of all its possible proofs (i.e., the data available to functions implicit in proofs of implications from P) need be stipu-

¹ This rather direct translation only works when any data computationally derivable from proofs of $t \in T$ are already computationally derivable from t itself. See [108, Chapter 1, Section 3] and [6, Chapter 7] for discussion.

lated. The production of any non-computational epistemic content is deferred to producing an argument that $[P]$ has a member.

Under this version of the principle, one may be given a member e which has type T without being also provided a method for constructing from e alone an argument for this fact. Propositions that are difficult to prove, say having the form $\forall x : \mathbb{N}. f(x) = 0$, can still have trivial computational content $\lambda x.r$ where r is a trivial constant.

Automating reasoning The Nuprl system is a tactic style theorem prover in the spirit of Edinburgh LCF [50]; moreover it introduced the notion of a *tactic-tree* so that users could interact directly with the proof object [37,52]. We take the approach that the basic elements of automated proof development can be seen as providing algorithms for formal metamathematical results, so it is computation in the *metalanguage* [33,71]. There is not space in this article to discuss or illustrate the tactic mechanism, but the interested reader can see thousands of examples of tactic-tree proofs in our on-line library of formal mathematics at <http://www.nuprl.org>. In addition, the library contains many examples of the basic tactics written in classical ML, the language of LCF which we adapted as the main programming language for Nuprl.

The first person to build substantial tactics for Nuprl was Douglas Howe, and he was able to use them to settle an important open problem about the typed lambda calculus, called the *Girard Paradox* [62]. Another person who solved an important open problem in mathematics was Chetan Murthy in generating a constructive proof of Higman's Lemma from its classical proof [95].

Over the years, many users have added to the Nuprl tactic library, sometimes by importing important algorithms from other provers, adapting packages from other provers such as HOL and Isabelle. We benefited a great deal from the use of a version of Nuprl at Edinburgh for many years and have adapted many of the ideas from Alan Bundy's group there, especially their reconstruction of methods from the Boyer and Moore prover Nqthm [22]. Chetan Murthy worked with the Coq group at INRIA and this resulted in sharing other ideas, designs and code (in both directions). The MetaPRL [57,58] group has contributed many ideas and programs as well. Now when we are using Nuprl, we are acutely aware that its capabilities rest on the contributions of dozens of modern researchers and the legacy of the logicians, philosophers, mathematicians and computer scientists we mentioned before. We will discuss later the contribution of fully automatic proving, through JProver, to proof automation in Nuprl, MetaPRL, and Coq.

Formal digital libraries The important role of knowledge, and the information resources from which it arises, is clear from the behavior of the best system designers and programmers. They bring extensive knowledge to bear in their reasoning about designs and code, sharing knowledge and reasoning in a coordinated yet distributed manner. However, the tools used to help them create and check inferences have access to limited fragments of formal proofs. It is a terrible waste of time and huge cost to unnecessarily recreate knowledge. Several researchers have made significant strides to overcome the challenges that sharing formal mathematics create [2,9,23,72,73,111]. The community of Mathematical Knowledge Management (MKM) has emerged with these challenges at heart. The PRL

research group has built a prototype Formal Digital Library in response to a MURI grant [30], and continues to focus efforts in this domain.

2. Computational Type Theory

Computational Type Theory (CTT) is a family of closely related theories designed to provide a foundation for reasoning about computation both in mathematics and computing practice. It is strongly influenced by the computer scientist’s wish to have a universal programming language as an integral part of the theory. There are two implementations of this theory, the initial one in Nuprl and a newer one in MetaPRL. Any implementation of the theory will provide an implementation of this language. This requirement on a foundational theory for computing goes back to early work in computer science on programming logic, including work at Cornell in which PLC programs were the algorithms of the logic, which was called PLCV [35,36,38]. Interestingly, some of the reasoning methods from PLCV were directly recoded in the Nuprl system, the arithmetic decision procedure for example, presented in [35].

Computational Type Theory was also influenced by our experience trying to generalize the type system of PLCV to make it more like Algol68 and to integrate the assertions of the programming logic with the types. Our efforts in this direction were dramatically affected when we read the work of Per Martin-Löf [89]. We incorporated his ideas into our programming logic [39], and then decided with Joe Bates to adopt Martin-Löf’s theory as the core of our type theory. We were deeply impressed by Martin-Löf’s idea that an untyped lambda calculus underlies the semantics of the type theory, and that methods of describing types and values could be rich enough that one must sometimes carry out complex arguments that a purported expression actually succeeds at describing a type or a value of a certain type. That is, the *well-formedness* of expressions for types and typed values is not a mere decidable syntactic condition.

Our technical presentation of types will begin with Martin-Löf’s Intuitionistic Type Theory as it appears in Nuprl. But first we describe some differences in the logic designed for the Nuprl system.

2.1. Differences in logic design

Martin-Löf’s texts, especially [90], which informed our understanding of his methods, formalize arguments about types as natural deduction style proofs under hypotheses of form $x \in A$, for variable x and type expression A . A peculiarity of this style of proof was that any proof introducing such a hypothesis must be built by extending a proof comprising a proof of A being a well-formed type expression. Thus, a common form of a proof that $(\Pi x : A)B(x)$ denotes a type would be a proof that A denotes a type followed immediately by introducing hypothesis $x \in A$ and proving $B(x)$ denotes a type under that hypothesis, then finishing with an inference to the proof goal discharging the hypothesis.

This did not suit the Nuprl design goal of organizing proofs by assertions to be proved, typically proving an assertion by proving independently provable assertions as premises. The meaning of Martin-Löf’s “ $B(x)$ type ($x \in A$)” *presupposes* that A is a type, and under

that presupposition means that $B(x)$ describes a type-valued function in x over A . Hence proving “ $B(x)$ type ($x \in A$)” requires proving A is a type. The alternative logic design choice made for Nuprl was to stipulate (and we use a different syntax here to avoid confusion) that “ $x : A \vdash B(x)$ type” means that *if* A is a type then $B(x)$ describes a type-valued function in x over A . Consequently, the problem of showing that “ $B(x)$ type ($x \in A$)” reduces to independently showing both “ A type” and “ $x : A \vdash B(x)$ type”.

Another purpose of the Nuprl design was to automatically derive witnesses to claims that some type has a member, forestalling the actual construction of the witness as long as practical. This allows the user to carry out much of a typical proof as though it was an ordinary proof of a proposition; so a person might simply prove a theorem of the form “ $\forall x : A. P(x)$ ” without always having to think about what functional expression is implicit in its construction. This provides an elegant solution to the project goal of implicitly extracting programs from proofs. Hence, “ $\vdash T$ ” was stipulated to mean that there is an expression t such that $t \in T$. And similarly, “ $x : A \vdash B(x)$ ” was stipulated to mean that if A is a type then $B(x)$ describes a type-valued function in x over A and there is an expression $t(x)$ that describes a $B(x)$ -valued function in x over A .

By stipulating for each rule of inference how to automatically construct the witness for the conclusion from any witnesses for the premises, one has shown how to automatically construct a witness, often a program, from any complete proof.

This goal of constructing programs automatically from proofs that their specifications can be met was an original design goal for the PRL project [13], and Martin-Löf’s theory provided an elegant basis for it [14], superior to the one suggested in [25].

2.2. Direct computation rules

A striking feature of the computationally type-free basis for the type theory as given in [90] is that *any* expression that evaluates to an expression of a given type is also an expression of that type and denotes the same value in that type. Thus, $(\lambda x.0)(a) = 0 \in \mathbb{N}$ no matter what (variable-free) untyped expression might be used for a . In the formal ITT systems provided by Martin-Löf there is no general rule exploiting this semantic feature. In contrast, such exploitation has become a hallmark of CTT methodology.

As reported in the 1986 book [31], standard Nuprl included *direct computation* rules, which allowed inference by various rewrites within a proof goal by steps of type-free computation. Thus, for example, showing the fact mentioned above reduces by one use of direct computation to showing that $0 = 0 \in \mathbb{N}$. Using these direct computation rules sometimes allows one to assign significant types to programs by induction over possible arguments. For example, a simple recursive unbounded search of a denumerable sequence of numbers for its first root, such as Kleene’s minimization operation, can be typed $\{f : (\mathbb{N} \rightarrow \mathbb{N}) \mid \exists n : \mathbb{N}. f(n) = 0\} \rightarrow \mathbb{N}$ and shown to find the first root of any function in this domain.² Indeed, such a typing argument was one of the motives for including the direct computation rules in the system originally.

² For a CTT proof of similar facts see “Recursive Functions” at www.nuprl.org/NuprlBasics.

Users developed the practice of using the untyped Y -combinator to define functions recursively instead of using various primitive recursion combinators, and then proving they had the appropriate types by induction.³ The original direct-computation rules as presented in [31], justified ultimately by methods in [5], involved unpalatable restrictions on contexts permitting rewrites, and were eventually simplified to allow type-free rewrites freely throughout a proof goal, justified by [63].

2.3. Intuitionistic type theory of Martin-Löf

Types are built from a few primitive types using basic *type constructors*. These provide templates for constructing various kinds of objects. The type constructors of ITT are motivated by mathematical considerations. However, they also correspond to data type constructors used in programming. The notes by C.A.R. Hoare, *Notes on Data Structuring*, also explain the ties between data types and logical types [59].

We start with the basic type constructors of ITT.

2.3.1. Cartesian products

If A and B are types, then so is their *product*, written $A \times B$. There will be many *formation rules* of this form, so we adopt a simple convention for stating them. We write

$$\frac{A \text{ is a Type} \quad B \text{ is a Type}}{A \times B \text{ is a Type}}$$

The elements of this product are pairs, $\langle a, b \rangle$. Specifically if a belongs to A and b belongs to B , then $\langle a, b \rangle$ belongs to $A \times B$. We abbreviate this by writing

$$\frac{a \in A \quad b \in B}{\langle a, b \rangle \in A \times B}$$

In programming languages these types are generalized to n -ary products, say $A_1 \times A_2 \times \dots \times A_n$.

We define equality on the elements by

$$\langle a, b \rangle = \langle c, d \rangle \text{ in } A \times B \quad \text{iff} \quad a = c \text{ in } A \text{ and } b = d \text{ in } B$$

In set theory, equality is uniform and built-in, but in type theory we define equality with each constructor; in ITT these equalities are built-in.

There is essentially only one way to decompose pairs. We say things like, “take the first element of the pair P ”, symbolically we might say *first*(P) or *1of*(P). We can also “take the second element of P ”, *second*(P) or *2of*(P).

A uniform way to access both elements of a pair simultaneously is the *spread* operation. The formula *spread*($P; x, y.t$) expresses that every free occurrence of x in the expression t represents the first element of the pair P and that every free occurrence of y in t represents the second element of P . It is often displayed as *let*(x, y) = P in t .

³ Indeed, the Y -combinator itself, not just its applications, can be assigned appropriate types if one uses the intersection type constructor of Section 2.5.1 below, and reformulates the function space constructor purely in terms of typing its applications to arguments, ignoring convergence. See “Typing Y ” at www.nuprl.org/NuprlBasics.

2.3.2. Function space

We use the words “function space” as well as “function type” for historical reasons. If A and B are types, then $A \rightarrow B$ is the type of *computable functions* from A to B . The formation rule is:

$$\frac{A \text{ is a Type} \quad B \text{ is a Type}}{A \rightarrow B \text{ is a Type}}$$

The informal notation we use for functions comes from mathematics texts, e.g. Bourbaki’s *Algebra* [21]. We write expressions like $x \mapsto b$ or $x \xrightarrow{f} b$; the latter gives a name to the function. For example, $x \mapsto x^2$ is the squaring function on numbers.

If b computes to an element of B when x has value a in A for each a , then we say $(x \mapsto b) \in A \rightarrow B$. Formally we use lambda notation, $\lambda(x.b)$ for $x \mapsto b$. The informal rule for typing a function $\lambda(x.b)$ is to say that $\lambda(x.b) \in A \rightarrow B$ provided that when x is of type A , b is of type B . We can express these *typing judgments* in the form $x : A \vdash b \in B$. The typing rule is then

$$\frac{x : A \vdash b \in B}{\vdash \lambda(x.b) \in A \rightarrow B}$$

If f, g are functions, we define their equality extensionally as

$$f = g \quad \text{iff} \quad f(x) = g(x) \quad \text{for all } x \text{ in } A$$

If f is a function from A to B and $a \in A$, we write $f(a)$ for the value of the function.

2.3.3. Disjoint unions

Forming the union of two sets, say $x \cup y$, is a basic operation in set theory. It is basic in type theory as well, *but* for computational purposes, we want to discriminate based on the type to which an element belongs. To accomplish this we put tags on the elements to keep them disjoint. Here we use *inl* and *inr* as the tags.

$$\frac{A \text{ is a Type} \quad B \text{ is a Type}}{A + B \text{ is a Type}}$$

The membership rules are

$$\frac{a \in A}{\text{inl}(a) \in A + B} \quad \frac{b \in B}{\text{inr}(b) \in A + B}$$

We say that $\text{inl}(a) = \text{inl}(a')$ iff $a = a'$ and likewise for $\text{inr}(b)$.

We can now use a case statement to detect the tags and use expressions like

$$\begin{aligned} &\text{if } x = \text{inl}(z) \quad \text{then } \dots \quad \text{some expression in } z \dots \\ &\text{if } x = \text{inr}(z) \quad \text{then } \dots \quad \text{some expression in } z \dots \end{aligned}$$

in defining other objects. The test for $\text{inl}(z)$ or $\text{inr}(z)$ is computable by an operation called *decide* that discriminates in the type tags. The typing rule and syntax for it are given in terms of a typing judgment of the form $E \vdash t \in T$ where E is a list of declarations of the form $x_1 : A_1, \dots, x_n : A_n$ called a *typing environment*. The A_i are types and x_i are variables

declared to be of type A_i . The rule is

$$\frac{E \vdash d \in A + B \quad E, u : A \vdash t_1 \in T \quad E, v : B \vdash t_2 \in T}{E \vdash \text{decide}(d; u.t_1; v.t_2) \in T}$$

2.3.4. Dependent product

Here is a simple example of the dependent product type constructor. Suppose you are writing an application needing a data type for the date.

$$\text{Month} = \{1, \dots, 12\}, \quad \text{Day} = \{1, \dots, 31\}, \quad \text{and} \quad \text{Date} = \text{Month} \times \text{Day}$$

We would need a way to check for valid dates. The pair, 2, 31 is a perfectly legal member of *Date*, although it is not a valid date. One thing we can do is to define

$$\begin{aligned} \text{Day}(1) &= \{1, \dots, 31\} \\ \text{Day}(2) &= \{1, \dots, 29\} \\ &\vdots \\ \text{Day}(12) &= \{1, \dots, 31\} \end{aligned}$$

and we will now write our data type as $\text{Date} = m : \text{Month} \times \text{Day}(m)$.

We mean by this that the second element of the pair belongs to the type indexed by the first element. Now, 2, 20 is a legal date since $20 \in \text{Day}(2)$, and 2, 31 is illegal because $31 \notin \text{Day}(2)$.

Many programming languages implement this or a similar concept in a limited way. An example is Pascal's *variant records*. While Pascal requires the indexing element to be of scalar type, we will allow it to be of any type.

One view of this is that we are generalizing the product type. It is very similar to $A \times B$. Let us call this type $\text{prod}(A, x.B)$. We can display this as $x : A \times B$. The typing rules are:

$$\frac{E \vdash a \in A \quad E \vdash b \in B[a/x]}{E \vdash \text{pair}(a, b) \in \text{prod}(A, x.B)}$$

$$\frac{E \vdash p \in \text{prod}(A, x.B) \quad E, u : A, v : B[u/x] \vdash t \in T}{E \vdash \text{spread}(p; u, v.t) \in T}$$

Note that we haven't added any elements. We have just added some new typing rules.

2.3.5. Dependent functions

If we allow B to be a family in the type $A \rightarrow B$, we get a new type, denoted by $\text{fun}(A; x.B)$, or $x : A \rightarrow B$, which generalizes the type $A \rightarrow B$. The rules are:

$$\frac{E, y : A \vdash b[y/x] \in B[y/x]}{E \vdash \lambda(x.b) \in \text{fun}(A; x.B)} \text{ new } y \quad \frac{E \vdash f \in \text{fun}(A; x.B) \quad E \vdash a \in A}{E \vdash f(a) \in B[a/x]}$$

Back to our example *Dates*. We see that $m : \text{Month} \rightarrow \text{Day}[m]$ is $\text{fun}(\text{Month}; m.\text{Day})$, where *Day* is a family of twelve types. And $\lambda(x.\text{maxday}[x])$ is a term in it.

2.3.6. Universes

A key aspect of ITT is the ability to denote types and type-valued functions using all the generic methods of the theory, especially those for reasoning about (typed) functions. For example, one might choose to recursively define the types of n -place curried numeric functions, such that $\mathbb{N}^0 \rightarrow \mathbb{N} = \mathbb{N}$ and $\mathbb{N}^{n+1} \rightarrow \mathbb{N} = \mathbb{N} \rightarrow (\mathbb{N}^n \rightarrow \mathbb{N})$. A convenient explanation of how this defines a type would be to employ a generic form for iteration of functions which explains, for any type T , for base case $b \in T$, and any function $F \in T \rightarrow T$, how to define a function $g \in \mathbb{N} \rightarrow T$ such that $g(0) = b$ and $g(n+1) = F(g(n))$.

To apply this generic form of reasoning about functions, we would need to find an appropriate type T such that $\mathbb{N} \in T$ and $\lambda(A.\mathbb{N} \rightarrow A) \in T \rightarrow T$. The obvious first choice would be to introduce a type \mathbb{U} of all types, but there cannot be one in ITT [60].

However, one can introduce a series of types \mathbb{U}_n called *universes* [90], and one can stratify the explanation of how expressions denote types by taking whatever basic type construction methods one might have and iteratively introducing a sequence of type expressions \mathbb{U}_n [5]. The base definition of type expressions is defined simply from the base construction methods. Then one defines the level $k+1$ system from the level k system by adding \mathbb{U}_k as a new constant type expression whose members are the type expressions defined at level k . The union of this sequence of systems then assigns every type expression to a universe, assigning $\mathbb{U}_n \in \mathbb{U}_{n+1}$, and making the universes cumulative, i.e., \mathbb{U}_n is a subtype of \mathbb{U}_{n+1} .

Returning to our example, for each k , adopting \mathbb{U}_k as the type to be used as T in the generic iterated function pattern, we get $\lambda(n.\mathbb{N}^n \rightarrow \mathbb{N}) \in \mathbb{U}_k$. Similarly, when expressing and proving general facts about types, one quantifies over universes rather than a non-existent type of all types, similar to Russell’s ramified logic.

Note that the *level variable* n in \mathbb{U}_n is a *not* a variable over which one quantifies within the logic. Indeed, the most perspicuous formulation of a logic for type theory with universes would, like both Martin-Löf’s and the original Nuprl logic of [31], require this universe level in any particular expression within the logic to be a constant, a numeral. But requiring users to produce these constants proved inconvenient; it too often happened that a lemma would be proved for level k that would later have been more useful had it been formulated at a higher level instead.

As a result, based upon [6], the Nuprl logic was reformulated to use more general level expressions for universes, including *level variables* (as letters), an operator for taking the *maximum* of two or more levels, and an operator for incrementing a level by a constant. A key rule was added that allowed the inference from any sequent to another gotten by uniformly instantiating all the level variables by any level-expressions.⁴

Hence, a proof of “ $\forall A : \mathbb{U}_k. P(A)$ ” takes on the force of quantification over all defined types, and such a lemma could be applied to a type of any level by instantiating k with an appropriate level expression. As a notational matter, when all the level expressions throughout an expression of CTT are simply occurrences of the same level variable, we often write simply “*Type*” instead of \mathbb{U}_i , say.

⁴ Note that, therefore, since one can always change level variables, that there is no inferential significance of the same level variable occurring in different sequents of an inference.

2.3.7. Propositions as types

One of the key distinguishing features of Martin-Löf type theory [89–91] and the Nuprl type theory [31] is that propositions are considered to be types, and a proposition is true iff it is inhabited. The inhabitants of a proposition are mathematical objects that provide evidence for the truth of the proposition.

This approach to propositions is related to the so-called *Curry–Howard isomorphism* [44,61] between propositions and types. But the correspondence is elevated in our type theory to the status of a principle. The Curry-Howard isomorphism explains why the definition is sensible for Heyting’s formalization of Brouwer’s constructive semantics for propositions. We briefly recall these semantics and state the Nuprl definitions for the logical operators. According to Brouwer the logical operators have these constructive meanings:

- \perp is never true (read \perp as “false”).
- $A \& B$ is true (provable) iff we can prove A and we can prove B .
- $A \vee B$ is true iff we can prove either A or B , and we say which is proved.
- $A \Rightarrow B$ is true iff we can effectively transform any proof of A into a proof of B .
- $\neg A$ holds iff $A \Rightarrow \perp$.
- $\exists x : A. B$ is true iff we can construct an element a of type A and a proof of $B[a/x]$.
- $\forall x : A. B$ is true iff we have an effective method to construct a proof of $B[a/x]$ for any a in A .

For an atomic proposition, P , with no logical substructure, we say it is true exactly when we have evidence for it. The evidence can be taken to be atomic as well, e.g. unstructured. So we use the unit element, \bullet , as evidence.

We express Brouwer’s semantics by these definitions of the logical operators:

Definition

$$A \& B \equiv A \times B$$

$$A \vee B \equiv A + B$$

$$A \Rightarrow B \equiv A \rightarrow B$$

$$\exists x : A. B \equiv x : A \times B$$

$$\forall x : A. B \equiv x : A \rightarrow B$$

$$\perp \equiv \text{void}$$

$$\top \equiv \mathbf{1}$$

We can also define other logical operators that express shades of meaning not commonly expressed. For example, $\cap x : A. B$ (Section 2.5.1) makes sense as a kind of universal quantifier.

We use $Prop_i$ as a synonym for $Type_i$.

2.4. Early extension to ITT

2.4.1. Subset and quotient types

One of the first extensions we made to ITT type theory has been one of the most useful and pervasive, a type we called the *subset type* [26]. The idea seems obvious from set theory, but its computational interpretation is subtle. Given a type A and a propositional function on A , say $B : A \rightarrow Prop_i$, we form the type of those elements of A satisfying B , $\{x : A \mid B(x)\}$. For example, if \mathbb{Z} is the type of integer, then the natural numbers can be defined as $\{z : \mathbb{Z} \mid z \geq 0\}$. This type is different from $z : \mathbb{Z} \times z \geq 0$ because the elements of $\{z : \mathbb{Z} \mid z \geq 0\}$ are simply numbers whereas the elements of the product are pairs $\langle z, p \rangle$ where p is the content of the proposition $z \geq 0$. The subset type is very useful in a computational theory because it allows us to separate data from its logical properties.

Another key extension we introduced at the same time is the *quotient type*. Given a type A , it comes with an equality relation, $=_A$. We can form a new type by changing the equality. If E is an equivalence relation on A , then $A//E$ is a new type, the quotient of A by E . The elements of $A//E$ are the elements of A , but the equality on $A//E$ is E . Good examples are the congruences of integers. That is, if $E_n(x, y)$ is $x = y \bmod n$, then $\mathbb{Z}//E_n$ are the integers modulo n , often denoted \mathbb{Z}_n . In \mathbb{Z}_2 , $0 = 2 = 4 = 6 = \dots$.

2.4.2. Inductive types

For recursive types ITT uses W-types (well-founded trees, Brouwer ordinals). An alternative approach was used for the Nuprl system by which recursive types are defined as minimal fixed points of suitable operators on types [93]. For example, a list-like type could be defined as $(\mu X. Unit + (A \times X))$, and Martin-Löf's type $(Wx \in A)B(x)$ could be defined as $(\mu W.x : A \times (B(x) \rightarrow W))$.

We say that $t_1 = t_2 \in \mu X.F(X)$ iff $t_1 = t_2 \in F(\mu X.F(X))$.

To construct elements of $\mu X.F(X)$, we basically just *unwind* the definition. That is,

$$\text{if } t \in F(\mu X.F(X)) \text{ then } t \in \mu X.F(X).$$

There is a natural notion of *subtyping* in this theory. We say that $A \sqsubseteq B$ iff $a = a' \in A$ implies that $a = a' \in B$. Semantically, a recursive type defined by $(\mu X.F(X))$ is meant to be the \sqsubseteq -minimal transitive symmetric relation on terms that is closed under F .

Induction over this type can be expressed as a rule, similar to rules in [58,93],

$$\frac{H, x : (\mu Z.T), J, u : \mathbb{U}_i, u \sqsubseteq (\mu Z.T), w : (x : u \rightarrow G[x/x]), z : T[u/Z] \vdash G[z/x]}{H, x : (\mu Z.T), J \vdash G[x/x]}$$

where $A \sqsubseteq B$ is simply an operator notationally defined as $a : A \rightarrow (a \in B)$, which expresses subtyping, since there is a member iff equal expressions of type A are equal expressions of type B .

The power of recursive types comes from the fact that we can define total computable functions over them and we can prove properties of elements recursively. (These are the same idea.) Recursive definitions are given by this term, $\mu\text{-ind}(a; f, z.b)$ called a *recursor* or *recursive-form*. It obeys the computation rule

$$\mu\text{-ind}(a; f, z.b) \text{ evaluates in one step to } b[a/z, (\lambda y.\mu\text{-ind}(y; f, z.b))/f].$$

It is rather clear what type $(\mu X.F(X))$ is when it is a type at all. It may be more controversial when $(\mu X.F(X))$ actually succeeds at describing a type. One widely acceptable criterion would be when $F(X)$ describes a \sqsubseteq -monotonic operation on types, which would be handily formulated by a rule

$$\frac{H, Z : \mathbb{U}_i \vdash T \in \mathbb{U}_i \quad H, X_1 : \mathbb{U}_i, X_2 : \mathbb{U}_i, (X_1 \sqsubseteq X_2) \vdash T[X_1/Z] \sqsubseteq T[X_2/Z]}{H \vdash (\mu Z.T) \in \mathbb{U}_i}$$

Some constructivists, however, might prefer a rule that imposes other constraints (perhaps intensional ones, involving positivity, say). One advantage of introducing particular kinds of recursive types, such as W-types and lists, instead of using this generic recursive type constructor is that the conditions for legitimately describing types would be quite narrow, and so more widely accepted.

2.5. Class extensions

As we expanded our efforts from verifying moderately large programs (say up to one thousand lines [1]) to verifying systems, we were faced with formalizing some of the basic concepts from object-oriented programming. Jason Hickey took on this task at the start of his thesis work, and along with Karl Crary, Alexey Nogin, and Aleksey Kopylov, we made a great deal of progress in building a *theory of classes* in CTT. These ideas are reported in several publications, among them [32,42,43,58,75]. The MetaPRL prover was built using object-oriented concepts, even in the tactic mechanism.

The key concepts turned out to be the subtyping relation and a generalization of the intersection type to dependent intersections—an idea that has had a profound impact on both the theory and its applications to designing systems.

2.5.1. Intersection types and the top type

Given two types A and B , it makes sense to consider the elements they have in common; we call that type the *intersection* of A and B , written $A \cap B$. We require that $(a = b$ in $A \cap B)$ iff $a = b$ in A and $a = b$ in B . To illustrate, it is clear that $\text{void} \cap A$ is void for any type A and $A \cap A$ is A .

It might be a surprise that $(1 \rightarrow 1) \cap (\text{void} \rightarrow \text{void})$ is not void but contains the identity function, $\lambda(x.x)$. This is because the base objects of Nuprl are *untyped*, so $\lambda(x.x)$ is polymorphic, belonging to all types $A \rightarrow A$. It is clear that $\{x : A \mid P(x)\} \cap \{x : A \mid Q(x)\}$ is $\{x : A \mid P(x) \& Q(x)\}$.

The intersection type is defined for a family of types as well. Let $B(x)$ be a family of types for $x \in A$, then $\cap x : A.B(x)$ is their intersection; and $b \in \cap x : A.B(x)$ iff $b \in B(a)$ for all $a \in A$. also $b = b'$ in $\cap x : A.B(x)$ iff $b = b'$ in $B(a)$ for all $a \in A$. Notice that when A is empty, $\cap x : A.B(x)$ is not empty but has exactly one element. This element is denoted by any closed expression of the theory, e.g. $\text{void} = 1$ in $\cap x : \text{void}.B(x)$. Such equalities follow from the understanding that under the assumption that $x \in \text{void}$, we can infer anything.

Types such as $\cap x : \text{void}.B(x)$ are maximal in the sense that every type is a subtype and every closed term denotes the single member of the type. We pick one such type expression and call it the *top type*, denoted Top . For example, $\cap x : \text{void}.x$ can be Top .

Type equality follows the pattern for structural equality, namely

$$(\cap x : A.B(x) = \cap x : A'.B'(x)) \quad \text{iff} \quad A = A' \text{ and } B(a) = B'(a)$$

Intersection types can express conditional typehood, which is useful in typing partial functions. If B does not depend on x , then $\cap x : A.B$ expresses the notion that B is a type provided that A is inhabited, as pointed out by Backhouse et al. [10]. We write this as B given A , following an analysis of this idea by Stuart Allen. Allen calls these *guarded types*.

Guarded types can specify the conditions under which $A \Rightarrow B$ is a type in Nuprl. Namely, if A is nonempty, the B must be a type. So the typing is

$$\lambda(A, B.A \rightarrow B) \in A : \text{Type} \rightarrow B : (\cap x : A.\text{Type}) \rightarrow \text{Type}$$

Guarded types are also useful in expressing one-one correspondences between types. For example we know that mappings from pairs p in $x : A \times B(x)$ into $C(p)$ are isomorphic to curried mappings of the type $x : A \rightarrow y : B(x) \rightarrow C(\langle x, y \rangle)$. Similarly, mappings $x : \{x : A \mid B(x)\} \rightarrow C(x)$ are isomorphic to those in $x : A \rightarrow \cap y : B(x).C(x)$.

2.5.2. Subtyping

As mentioned before, the natural notion of subtyping is to define $A \sqsubseteq B$ iff $a = a' \in A$ implies that $a = a' \in B$. For example, $\{x : A \mid P(x)\} \sqsubseteq A$ for any predicate $P(x)$. We clearly have $\text{void} \sqsubseteq A$ for any type A and $A \sqsubseteq \text{Top}$ for any type A , and $A \cap B \sqsubseteq A$ for any types A and B . It is easy to see that the following relationships hold.

If $A \sqsubseteq A'$ and $B \sqsubseteq B'$ then

1. $A \times B \sqsubseteq A' \times B'$
2. $A + B \sqsubseteq A' + B'$
3. $A' \rightarrow B \sqsubseteq A \rightarrow B'$.

The relation in 3 holds because functions in type theory are polymorphic. Given $f \in A' \rightarrow B$, we see that it belongs to $A \rightarrow B'$ as well because on inputs restricted to A it will produce results in B . In set theory this relationship would not hold, as we see in examples from the discussion of classes.

2.5.3. Records

We follow the idea from algebra and programming languages that classes are like algebraic structures, sets with operators. But classes are defined using signatures, and these are large types. Without the signature definition, as a large type, we could not have the definition of a class as a small type. The precise type we start with is a record type, these will also be classes.

Let F be a discrete type, the field names, and define a signature over F as a function from F into types, say $\text{Sig} \in F \rightarrow \text{Types}$. A *record type on discrete type F with signature $\text{Sig} \in F \rightarrow \text{Types}$* is the dependent type $i : F \rightarrow \text{Sig}(i)$. If F is a finite type, say $\{x_1, \dots, x_n\}$, then we write the record as $\{x_1 : \text{Sig}(x_1); \dots; x_n : \text{Sig}(x_n)\}$. If we are given types T_1, \dots, T_n we can also write simply $\{x_1 : T_1; \dots; x_n : T_n\}$ and this implies that we built the function assigning type T_i to name x_i .

Notice that if F_1 and F_2 are discrete types such that $F_1 \sqsubseteq F_2$ and $Sig_1 : F_1 \rightarrow \text{Type}$, $Sig_2 : F_2 \rightarrow \text{Type}$ and $Sig_1(x) \sqsubseteq Sig_2(x)$ for $x \in F_1$ then

$$i : F_2 \rightarrow Sig_2(i) \sqsubseteq j : F_1 \rightarrow Sig_2(j)$$

An important example of these relationships will be used as a running illustration. We consider it here briefly. Given a type M as the carrier of algebraic structure, we get the structure of a monoid over M , $Monoid_M$, is $\{op : M \times M \rightarrow M; id : M\}$. We can extend this to a group structure, $Group_M$, as $\{op : M \times M \rightarrow M; op : M; inv : M \rightarrow M\}$. Notice that $Group_M \sqsubseteq Monoid_M$.

2.5.4. Uniform records

Extending records is a basic operation, as in extending a Monoid structure to a Group structure. This can be done in a uniform way if we agree on a common set of names to be used as names of the components (“fields” as they are often called). Let us take $Label$ as an infinite discrete type. Signatures Sig over $Label$ are functions in $i : Label \rightarrow \text{Type}$, and records with this signature are $i : Label \rightarrow Sig(i)$.

We generally consider records in which only finitely many labels are names of significant components. If we map the insignificant components to the type Top , then intersection of arbitrary uniform records makes sense and is defined as

$$(i : Label \rightarrow Sig_1(i)) \cap (j : Label \rightarrow Sig_2(j)) = i : Label \rightarrow Sig_1(i) \cap Sig_2(i)$$

Record intersection is an operation that extends records as we see if we assume that op , id and inv are elements of $Label$ and assume that $MonSig_M(i) = Top$ if $i \neq op$ and $i \neq id$. Generally when we write $\{x_1 : Sig(x_1); \dots; x_n : Sig(x_n)\}$ we will assume that $Sig(i) = Top$ for $i \neq x_j$. So now consider $\{op : M \times M \rightarrow M; id : M\} \cap \{inv : M \rightarrow M\}$.

The result is $Group_M$, $\{op : M \times M \rightarrow M; id : M; inv : M \rightarrow M\}$ because in the structure $\{inv : M \rightarrow M\}$, the components for op and id are Top and in Monoid the component for inv is Top . Thus for example, the type at op in the intersected structure is

$$op : (M \times M \rightarrow M) \cap Top \quad \text{and} \quad (M \times M \rightarrow M) \cap Top = M \times M \rightarrow M$$

When two records share a common field, the intersected record has the intersection of the types, that is $\{x : T\} \cap \{x : S\} = \{x : T \cap S\}$. So generally,

$$\{x : T\} \cap \{y : S\} = \begin{cases} \{x : T, y : S\} & \text{if } x \neq y \\ \{x : T \cap S\} & \text{if } x = y \end{cases}$$

2.5.5. Dependent records

Records can be seen as (Cartesian) products with labels for the components. Can we define the record analogue of dependent products? The right notation would be $\{x_1 : T_1; x_2 : T_2(x_1); \dots; x_n : T_n(x_1, \dots, x_{n-1})\}$ where $T_i(x_1, \dots, x_{i-1})$ is a type depending on x_1, \dots, x_{i-1} . Can we define this structure in terms of a function $Sig : \{x_1, \dots, x_n\} \rightarrow \text{Type}$? We see that $Sig(x_1)$ is just a type, but $Sig(x_2)$ uses a function $F \in Sig(x_1) \rightarrow \text{Type}$ and it is F applied to x_1 . We might think of writing $F(x_1)$, but x_1 is a name and F wants an element of $Sig(x_1)$. Basically we need to talk about an arbitrary element of the type we are

building. That is if $r \in i : \{x_1, \dots, X_n\} \rightarrow \text{Sig}(i)$, then $\text{Sig}(x_2)$ is $F(r(x_1))$. So to define Sig we need r , and to type r we need Sig .

We see that an element of a type such as $\{x_1 : T_1; \dots; x_n : T_n(x_1, \dots, x_{n-1})\}$ has the property that its type depends on its “previous values”. That is, the type of $r(x_1)$ is $\text{Sig}(x_1)$ but the type of $r(x_2)$ is $\text{Sig}(x_2)(r(x_1))$. We can encode this type using the *very-dependent function type* denoted $\{f \mid x : A \rightarrow B[f, x]\}$. We give Jason Hickey’s original definition below. Now we use it to define the dependent records. We take A to be the finite type $\{0 \dots (n-1)\}$. Let

$$T_f(n) \equiv f \rightarrow i \in \{0 \dots (n-1)\} \left[\begin{array}{l} \text{case}(i) \text{ of} \\ 0 \quad \Rightarrow T_0 \\ 1 \quad \Rightarrow T_1(f(0)) \\ \quad \vdots \\ n-1 \Rightarrow T_{n-1}(f(0))(f(1)) \dots (f(n-2)) \end{array} \right]$$

If we have a function F that returns the values T_i for $0 \leq i < n$ given the index i , we can give this type a more concise form. Let $\text{fix}(f.b)$ define the fixed point of the program b . That is, $\text{fix}(f.b) \equiv Y\lambda f.b$. Given functions F and f , and an index i we can define a function napply that applies F to $f(0), \dots, f(n-1)$:

$$\text{napply} \equiv \lambda f. \lambda i. \text{fix}(g. \lambda G. \lambda j. \text{if } j = i \text{ then } G \text{ else } g(G(f(j)))(j+1))$$

The type for F is then $T_F(n) \equiv \{F \mid i : \{0 \dots (n-1)\} \rightarrow \text{napply}(F)(i)(\lambda x.x)(0)\}$.

If we parameterize these types over n , we have the types

$$S_f \equiv n : \mathbb{N} \times T_f(n)$$

$$S_F \equiv n : \mathbb{N} \times T_F(n)$$

The type S_F specifies *descriptions* of dependent products of arbitrary finite arity, and the type S_f specifies the corresponding dependent product type. The inhabitants of S_f are pairs $\langle n, f \rangle$ where n is the arity of the product, and f is a function with domain $\{0 \dots (n-1)\}$ that serves as the projection function.

This finite arity dependent product type will be useful to define the type *Signature*. It can also be used to define the type of sequents in a constructive logic: the hypotheses of the sequents are defined as a dependent product of finite arity, and the goal is a type that may depend on all hypotheses.

2.6. Later extensions

2.6.1. Dependent intersection

The general subtyping of record types is captured exactly by the *contravariance* in function space subtyping: $A \sqsubseteq A'$ and $B \sqsubseteq B'$ implies $(A' \rightarrow B) \sqsubseteq (A \rightarrow B')$. This property makes it possible to express records as intersections of singletons, as illustrated in Section 2.5.4.

To express *dependent records* (see Section 2.5.5) in a similar fashion, a type constructor stronger than simple intersection is needed. Kopylov [74] elaborated an elegant definition

Table 1
Inference rules for dependent intersection

$\frac{\Gamma \vdash A \text{Type} \quad \Gamma; x : A \vdash B[x] \text{Type}}{\Gamma \vdash (x : A \cap B[x]) \text{Type}}$	(TypeFormation)
$\frac{\Gamma \vdash A = A' \quad \Gamma; x : A \vdash B[x] = B'[x]}{\Gamma \vdash (x : A \cap B[x]) = (x : A' \cap B'[x])}$	(TypeEquality)
$\frac{\Gamma \vdash a \in A \quad \Gamma \vdash a \in B[a] \quad \Gamma \vdash x : A \cap B[x] \text{Type}}{\Gamma \vdash a \in (x : A \cap B[x])}$	(Introduction)
$\frac{\Gamma \vdash a = a' \in A \quad \Gamma \vdash a = a' \in B[a] \quad \Gamma \vdash x : A \cap B[x] \text{Type}}{\Gamma \vdash a = a' \in (x : A \cap B[x])}$	(Equality)
$\frac{\Gamma; u : (x : A \cap B[x]); \Delta[u]; x : A; y : B[x] \vdash C[x, y]}{\Gamma; u : (x : A \cap B[x]); \Delta[u] \vdash C[u, u]}$	(Elimination)

of dependent records in terms of a new type constructor that he discovered: *dependent intersections*. Let A be a type and $B[x]$ be a type for all x of the type A . We define *dependent intersection* $x : A \cap B[x]$ as the type containing all elements a from A such that a is also in $B[a]$. This is a type that we might write as $\{x : A \mid x \in B_x\}$, but $x \in B_x$ is not a well-formed proposition of Nuprl unless it is true.

As an example, let $A = \mathbb{Z}$ and $B[x] = \{y : \mathbb{Z} \mid y > 2x\}$ (i.e., $B[x]$ is a type of all integers y , such that $y > 2x$). Then $x : A \cap B[x]$ is a set of all integers, such that $x > 2x$.

The ordinary binary intersection is just a special case of a dependent intersection with a constant second argument: $A \cap B = x : A \cap B$.

The dependent intersection is not the same as the intersection of a family of types $\bigcap_{x:A} B[x]$. The latter refers to an intersection of types $B[x]$ for all x in A . The difference between these two type constructors is similar to the one between dependent products $x : A \times B[x] = \Sigma_{x:A} B[x]$ and the product of a family of types $\prod_{x:A} B[x] = x : A \rightarrow B[x]$.

Two elements, a and a' , are equal in the dependent intersection $x : A \cap B[x]$ when they are equal both in A and $B[a]$. The rules for the new type appear in Table 1.

2.6.2. Dependent record types as dependent intersections

Now we will give a method for constructing record types and their elements using dependent intersection. This type is defined by other means in *Nuprl's Class Theory and its Applications* [32], and there are other quite different accounts [16,100].

Records Elements of record types are defined as in [32]. They map labels to the corresponding fields. We can build up these functions component by component starting with an empty record. We write the extension (or field update) as $r.x := a$; think of assigning a to $r.x$, where x is a label.

$$\{\} \equiv \lambda(l.l) \quad (\text{the empty record})$$

We could pick any function as a definition of an empty record.

$$(r.x := a) \equiv \lambda(l. \text{if } l = x \text{ then } a \text{ else } r \ l) \quad (\text{field extension})$$

$$r.x \equiv r(x) \quad (\text{field extraction})$$

Now we use a common notation for the elements of record types.

$$\{x_1 = a_1; \dots; x_n = a_n\} \equiv \{l.x_1 := a_1.x_2 := a_2. \dots .x_n := a_n$$

These definitions provide that

$$\begin{aligned} \{x_1 = a_1; \dots; x_n = a_n\} = \lambda l. & \text{if } l = x_1 \text{ then } a_1 \text{ else} \\ & \vdots \\ & \text{if } l = x_n \text{ then } a_n \end{aligned}$$

Record types We will show how to build finite record types by intersection. First we define a singleton record.

$$\{x : A\} \equiv \{x\} \rightarrow A \quad (\text{single-field record type})$$

where $\{x\} \equiv \{l : \text{Label} \mid l = x \in \text{Label}\}$ is a singleton set.

$$\{R_1; R_2\} \equiv R_1 \cap R_2 \quad (\text{independent concatenation of record types})$$

Thus $\{n : \text{Nat}; x : \text{Real}\}$ is $\{n : \text{Nat}\} \cap \{x : \text{Real}\}$.

Next we include dependency.

$$\{self : R_1; R_2[self]\} \equiv self : R_1 \cap R_2[self] \quad (\text{dependent concatenation})$$

Here *self* is a variable bound in R_2 . We will usually use the name “self” for this variable and use the shortening $\{R_1; R_2[self]\}$ for this type. Further, we will omit “self.” in the body of R_2 , e.g. we will write just x for $self.x$, when such notation does not lead to misunderstanding.⁵

We assume that this concatenation is a left associative operation, and we will omit inner braces. For example, we will write $\{x : A; y : B[self]; z : C[self]\}$ instead of $\{\{\{x : A\}; \{y : B[self]\}\}; \{z : C[self]\}\}$.

Note that in this expression there are two distinct bound variables *self*. The first one is bound in B , and refers to the record itself as a record of the type $\{x : A\}$. The second *self* is bound in C ; it also refers to the same record, but it has type $\{x : A; y : B[self]\}$.

The definition of independent concatenation is just a special case of dependent concatenation, when R_2 does not depend on *self*. We can also define *right associating records*, a less natural idea, but nevertheless expressible:

$$\{s : x : A; R[s]\} \equiv self : \{x : A\} \cap R[self.x]$$

Here x is a variable bound in R that represents a field x . Note that we can α -convert the variable s , but not the label x , e.g., $\{s : x : A; R[s]\} = \{y : x : A; R[y]\}$, but $\{s : x : A; R[s]\} \neq \{y : z : A; R[y]\}$. We will usually use the same name for labels and corresponding bound variables. This connection is right associative, e.g., $\{s : x : A; t : y : B[s]; z : C[s, t]\}$ stands for $\{s : x : A; \{t : y : B[s]\}; \{z : C[s, t]\}\}$.

⁵ Note that Allen’s notational devices allow Nuprl to display these terms as written here [7].

These dependent records can be used to capture theories built in a context. For example, here is how a *monoid* could be defined:

$$\{M : \text{Type}_i;$$

$$op : M \times M \rightarrow M;$$

$$id : M;$$

$$assocax : \forall x, y, z : M. op(x, op(y, z)) = op(op(x, y), z);$$

$$idax : \forall x : M. (op(x, id) = x \ \& \ op(id, x) = x)\}.$$

Here we use the fact that propositions are types—another key Automath contribution.

2.7. Unions, singletons, and mainstream sets

Given the utility and elegance of the intersection operator, and given the fact that those two properties derive in large part from \cap 's position as the infimum of the subtyping order, it is natural to wonder what we could accomplish if we also had the dual. The union type operator has, in fact, recently been introduced into the CTT family [58,76]. Guided and inspired not only by the supremum operation of order theory but also by mainstream set theory's union, the present implementation is a good compromise in the face of the infeasibility of getting the best of both those two worlds.

Part of the compromise is uneven compatibility with set-theoretic expectations, as exemplified by the fact that a union $A \cup B$ can have fewer distinct members than do its branches A and B . That fascinating information-hiding phenomenon is characteristic of general unions, but we will not pursue it further since the rest of this paper only requires restricted and unsurprising applications of \cup . Specifically, each collection of types that we will union together will be pairwise disjoint, and that special case of the operator is reliably like the familiar ZFC union, both in its behavior and in its versatility.

As hoped, these two new ways of combining types, intersection and union, do greatly enrich the foundational system, and the community is still in the early stages of exploiting them. Kopylov, for instance, is using unions and intersections to describe objects [76]. Another ongoing effort is the use of the enhanced language to re-explain types in a way that creates a pervasive, bi-directional connection between CTT and ZFC.

In the limited space remaining, a preview of the latter project is presented. Our focus here is on one central aspect of the re-explanation: the fact that types are constructed from the bottom up, using unions and intersections as mortar and using singletons as stones. To see what is meant by that, consider the following unusual rendition of $\mathbb{N}_2 \rightarrow \mathbb{N}_2$.⁶ It has a superficial similarity to the analogous ZFC function space, $2 \rightarrow 2$, since it is the union of four (disjoint) things. If we look more closely, we see that the similarity to the native ZFC function space $2 \rightarrow 2$ goes much further than that; the two match all the way down to the

⁶ Throughout this section, we work modulo extensional type equality.

leaves.

$$\cup \left[\begin{array}{l} \cap \left[\begin{array}{l} \{0\} \rightarrow \{0\} \\ \{1\} \rightarrow \{1\} \end{array} \right. \\ \cap \left[\begin{array}{l} \{0\} \rightarrow \{0\} \\ \{1\} \rightarrow \{0\} \end{array} \right. \\ \cap \left[\begin{array}{l} \{0\} \rightarrow \{1\} \\ \{1\} \rightarrow \{0\} \end{array} \right. \\ \cap \left[\begin{array}{l} \{0\} \rightarrow \{1\} \\ \{1\} \rightarrow \{1\} \end{array} \right. \end{array} \right. \left. \left. \begin{array}{l} \left\{ \begin{array}{l} (0, 0) \\ (1, 1) \end{array} \right\} \\ \left\{ \begin{array}{l} (0, 0) \\ (1, 0) \end{array} \right\} \\ \left\{ \begin{array}{l} (0, 1) \\ (1, 0) \end{array} \right\} \\ \left\{ \begin{array}{l} (0, 1) \\ (1, 1) \end{array} \right\} \end{array} \right.$$

Thinking of them in this way, we could say that they're structurally identical, although the one on the left is admittedly much more heavily decorated.

To see why the left-hand type expression and $\mathbb{N}_2 \rightarrow \mathbb{N}_2$ are extensionally equal, consider the union's third branch. That intersection's top half, $\{0\} \rightarrow \{1\}$, says that 0 maps to 1, while its bottom half, $\{1\} \rightarrow \{0\}$, says that 1 maps to 0. It is thus a brute-force description of negation. The identity function and the two constant functions are visible in the remaining branches.

Doug Howe, in [64,66,67], essentially showed that the connection illustrated above can be developed into a comprehensive system, one that encompasses not only finite function spaces like $\mathbb{N}_2 \rightarrow \mathbb{N}_2$ but also infinite function spaces, higher-order function spaces, dependent function spaces, dependent products, universes, W types, special cases of the intersection, union, and recursion operators, and even a variant of the quotient operator. The result underlies a novel set-theoretic semantics of a new, classical member of the CTT family.

The semantic framework discussed here is an extension and re-examination of his work, taking different approaches in many areas (in particular, by focusing on singletons instead of oracles) and aiming to ultimately explain a richer collection of types. The core ideas and methods are still the same, though. To get a glimpse of the mechanics, let's first note that the metatheory is $ZFC + \exists$ infinitely many inaccessible, abbreviated $ZFCI^\omega$, and that the meaning of a type T with κ equivalence classes is a set $\llbracket T \rrbracket$ with exactly κ elements.

Although entirely conventional up to this point, perfectly in line with [109, Chapter 9] and [46], if we attempt to continue with a straightforward mapping we immediately encounter a snag. Namely, the set $\llbracket \{0\} \rightarrow \{0\} \rrbracket$ must contain exactly one element, but the type $\{0\} \rightarrow \{0\}$ contains terms such as $\lambda x.0$ and $\lambda x.x$ whose meanings must be distinct. It turns out that this counterexample is typical in the sense that the challenges that we face all boil down to difficulties with singletons. The essential problem is $\llbracket S \rrbracket$ when S is a singleton, and the counterexample shows that this can be perplexing even at the level of the oldest and simplest types.

The solution is to abandon, at least initially, the idea that $\llbracket S \rrbracket$ will be a set of things that inhabit S . Instead, $\llbracket S \rrbracket$ is no more (and no less) than a description of how S can be built from other singletons. In a bit more detail, the unique member of $\llbracket S \rrbracket$ is a well-founded tree, the encoding of which is not particularly important. The tree's rank can be (and typically is) greater than several inaccessible cardinals. At each leaf, there is a code for either the

singleton Top or a concrete singleton like $\{1\}$. At each internal node, there is a code for a type operator such as $\cdot \rightarrow \cdot$, $\cdot \times \cdot$, \cap , or $//_{\text{Unit}}$.

These trees, these blueprints for constructing singletons, have a number of useful and intriguing properties. First and foremost is the comprehensive, pervasive match that often exists between the blueprint and its native ZFCI^ω counterpart. This is true not only for function spaces, as illustrated with $\mathbb{N}_2 \rightarrow \mathbb{N}_2$ earlier, but also for products, disjoint unions, and so on, and it is the basis of Howe’s justification of bi-directional sharing between STT and his classical CTT.

This singleton-based rendition of type theory is also remarkable for the fact that types are constructed bottom-up, building singletons from singletons via a fairly limited and orthogonal collection of type operators. It seems plausible that, with sufficient effort, this approach could be developed into a full-fledged, alternative formulation of a rich successor to Howe’s classical CTT. Due to the percolation upward from mainstream set theory, we might aspire to a foundation that is compact, stable, and directly connected to ZFC.

2.8. A logic of events

As a result of working with system designers on the specification, design, and verification of distributed systems [19,56,85], we have developed a theory of event structures [17] that is the basis for extending CTT and the extending the method of correct-by-construction programming. We present distributed systems from the point of program synthesis that we helped create in the 80’s. Logical synthesis of programs in constructive logic uses the *propositions-as-types principle*: a specification is a proposition and a constructive proof of the proposition yields a member of the corresponding type [14]. Because we can adequately characterize a sequential computation as a function from its input parameters to its output value, a computational theory of function types is a sufficient basis for logical synthesis of sequential programs. Specifications are propositions that correspond to function types and programs are members of these function types.

From Lamport’s work [83,84], we know that to characterize a distributed computation we need a fundamental model of *messages*, i.e., information flow between agents. This naturally introduces abstract agents and temporal concepts into our logic since a message is sent from some agent at some time and is received by another agent at some, causally later, time. These aspects of distributed computation are not naturally modeled as functions and, hence, a distributed program is not merely a member of some function type. Instead, distributed algorithms are specified by describing the desired set S of possible runs, and a distributed program p satisfies the specification if all of its runs are in the set S . Thus a specification is a proposition about runs and a program is something that constrains the set of possible runs.

A *system* is a set of runs. A protocol p generates a system, namely the set of all runs consistent with p . Again, following Lamport, we characterize a run of a program as a set of events. Each event is associated with a unique agent. When we want to speak more abstractly about *events happening* before the agent is fully defined, we talk about the events at a *location*. A location is an abstraction of the notion of an *agent* or a *process*. Formally, events are elements of a type E , and there is a function loc of type $E \rightarrow \text{Loc}$, where Loc is some set of locations. For each $i \in \text{Loc}$, the set of events e such that $\text{loc}(e) = i$ is totally

ordered. Intuitively, this set of events is the *history* of events at location i . If $\text{first}(e)$ holds, then e is the first event in the history associated with $\text{loc}(e)$; if not, then e has a predecessor $\text{pred}(e)$. Events are further partitioned by their *kinds*. The $\text{kind}(e)$ of event e is either $\text{rcv}(l)$ —a receive event on *link* l , or else $\text{local}(a)$ —a local event of kind a .⁷ Every receive event has a *sender event* $\text{sender}(e)$. The sender event of a message is the event when the message was sent.

In this setting we can define Lamport’s [83] *causal order* on events as the transitive closure of the sender-receiver and predecessor relations. The local state of a process is represented as the values of a collection of state variables. Formally, state variables are just identifiers that are assigned a value at each event.⁸

There are binary functions *when* and *after* that describe the values of state variables before and after an event takes place. We typically write these functions using infix notation. Thus, if $\text{loc}(e) = i$ then $(x \text{ when } e)$ describes the value of the state variable x at location i just before e , and $(x \text{ after } e)$ describes its value after e . Note that state variables are local variables and two locations may have state variables with the same name.

Every event e also has a *value* $\text{val}(e)$. The value of a receive event is the message that is received, and the value of a local event represents a value (satisfying some constraints) chosen (non-deterministically) when the local event is generated. For example, if whenever a local event of kind a occurs an integer value is chosen and a message with twice that value is sent on link l , then events with $\text{kind}(e) = \text{rcv}(l)$ and $\text{kind}(\text{sender}(e)) = \text{local}(a)$ will have $\text{val}(e) = 2 * \text{val}(\text{sender}(e))$.

Our six axioms of event structures say that an event is the sender of only a finite number of messages; the predecessor function is one-to-one; causal order is well-founded; the local predecessor of an event has the same location; the sender of an event has the location of the source of the link on which the message was received; and state variables change only at events, so that: $(x \text{ after } \text{pred}(e)) = (x \text{ when } e)$.

Working in a powerful higher-order logic, we are free to define further derived concepts on top of the event structure model. An event structure is a rich enough structure that we can define various “history” operators that list or count previous events having certain properties. Because we can define operators like these we do not need to add “history variables” to the states in order to write specifications and prove them. For example, the basic history operator lists all the locally prior events at a location.

$$\text{before}(e) = \text{if } \text{first}(e) \text{ then nil else } \text{pred}(e) :: \text{before}(\text{pred}(e))$$

We can also define useful notations for concepts such as “event e changes state variable x ”

$$x \Delta e \equiv (x \text{ after } e \neq x \text{ when } e)$$

This language is a theoretical contribution that enables natural declarative descriptions of distributed systems, and it is a basis for formal verification. Now we need to examine the issue of code synthesis.

⁷ Receive events are further partitioned by a *tag* so that we can restrict the kinds of events sending messages on a given link with a given tag without restricting other uses of that link. To simplify the discussion in this paper, we have suppressed all mention of these tags.

⁸ State variables are typed, but to simplify our discussion we suppress all type declarations.

2.8.1. Worlds and executable code

In order to relate the event structure abstraction to executable code, we define a general model of the execution of a distributed system. We call instances of this general model *worlds*. Any concurrent programming language can be given a semantics in terms of its possible worlds, and every world gives rise to an event structure. We are now using one concurrent programming language that we call *message automata* (and their coding as Java programs). These automata are closely related to the IO-automata defined by Lynch [88] but have a built-in syntax for messages and a built-in mechanism, the frame conditions, to allow simple extension. We will summarize below how the semantics of XYZaa is defined in terms of possible worlds.

Definition of world A world is an idealized trace of the execution of a distributed system. It has locations, and links from a graph $\langle Loc, Lnk, src, dst : Lnk \rightarrow Loc \rangle$. In contrast to an event structure, in which time is only a partial order, a world has a notion of global, discrete time, modeled as the natural numbers \mathbb{N} . This global time is not expressible in the logic of events and is only used to construct these models. By observing the system at every location i and every time t , we have a *state* $s(i, t)$, an *action* $a(i, t)$, and a list of *messages* $m(i, t)$.

The state $s(i, t)$ is the state of the part of the system at location i at time t . The state at location i is a general record (a dependent function from names to values).

The action $a(i, t)$ is the action that was chosen by the system to be executed next at location i and time t . It may be a null action, indicating that no action was taken at $\langle i, t \rangle$. Other actions will be local actions and the action of receiving a message. Every action has a *kind* of one of these forms (null, local, or receive), and it also has a *value* whose type depends on the kind and location of the action.

The messages $m(i, t)$ are the list of messages sent from location i at time t . As in event structure, a message consists of a link, a tag, and a value.

Event system of a world If w is a world, then we can construct an event structure from w . The basic construction is to let the events E be the points $\langle i, t \rangle$ in spacetime such that the action $a(i, t)$ in world w is not null. So E is the set of points where an event occurred. We then define all the operations $<_{loc}$, $<$, *first*, *pred*, *when*, *after*, etc., to construct the event structure $Ev(w)$ of world w . *It is then a theorem that $Ev(w)$ satisfies all the axioms for event structures.*

Note, our logic is not about worlds. They are part of the meta-reasoning about execution. We reason about processes which we turn to next.

2.8.2. Message automata

Event structures and worlds are infinite objects, but they arise from the behaviors of *finite* distributed programs. We call our representations of these finite programs *message automata*. A message automaton is a finite collection of declarations and clauses each of which has an assigned location. The *declarations* provide the names and types of state variables, local actions, input messages, and output messages. The *clauses* **initialize** variables, *state preconditions* on local actions, give the **effects** of actions on state variables, and give the *output* messages **sent** by actions. In addition, two further kinds of clauses, called

frame conditions, restrict the kinds of actions that affect **only** a given state variable or send *only* on a given link.⁹ For example, the following message automaton was automatically assembled from a proof of an algorithm for leader election in a ring.

```

state me :  $\mathbb{N}$ ; initially uid(i)
state done :  $\mathbb{B}$ ; initially false
state x :  $\mathbb{B}$ ; initially false
action vote; precondition  $\neg$ done
  effect done := true
  sends [msg(out(i), vote, me)]
action rcvin(i)(vote)(v) :  $\mathbb{N}$ ;
  sends if  $v > me$  then [msg(out(i), vote, v)] else []
  effect x := if  $me = v$  then true else x
action leader; precondition  $x = true$ 
only rcvin(i)(vote) affects x
only vote affects done
only {vote, rcvin(i)(vote)} sends out(i), vote

```

The semantics of a message automaton D is the set of possible worlds w that are consistent with it. To be consistent, w must be a fair-fifo world and respect the meanings of the clauses of the message automata at each location. We have a formal definition $Consistent(D, w)$ of this relation, and this defines the semantics of distributed systems of message automata.

Rules for message automata A message automaton constrains the possible worlds that can be executions of the system. We state these constraints as rules on the event structures that come from the possible worlds. A *rule* of the form $M : \psi$ means that:

$$\exists w : Consistent(M, w). \forall D : D_{sys}. \forall w : World. \\ Consistent(D, w) \wedge M \subseteq D \Rightarrow Ev(w) \models \psi$$

It says that there is at least one event structure of a possible world for M and the event structure of any possible world of any distributed system containing M will satisfy ψ . Since there are only six kinds of clauses in message automata, *we have a set of only six basic rules and one composition rule*. For example, the rules for initialization clauses and effect clauses are

$$@i \text{ state } x : T; E_{\text{initially}} x = v : \\ \forall e @ i. \text{first}(e) \Rightarrow x \text{ when } e = v$$

⁹ Actually, send on a given link with a given tag.

@i state $x : T1$; action $k : T2$;
 $k(v)$ effect $x := f(s, v)$:
 $\forall e @ i. \text{kind}(e) = k \Rightarrow x \text{ after } e = f(s \text{ when } e, \text{val}(e))$

2.8.3. Implementation of the logic of events and extraction of process code

Our current Logical Programming Environments supports proof and program development by top down refinement of goals into subgoals and bottom up synthesis of programs from fragments of code derived from proofs of subgoals. We are extending this mechanism, called *program extraction*, to the synthesis of *processes* to support *process extraction*.

Our library of declarative knowledge about distributed systems will contain many theorems that state that some property ϕ of event structures is *realizable* (which we write as $\models \phi$.) A property ϕ is realizable if and only if there is a distributed system D that is both *feasible*—which implies that there is at least one world and, hence, at least one event structure, consistent with D —and realizes the property; every event structure consistent with D satisfies the property.

The basic rules for message automata provide initial knowledge of this kind—all the properties of single clauses of message automata are realizable. We add to our knowledge by proving that more properties are realizable. In these proofs, the system will automatically make use of the knowledge already in the library when we reach a subgoal that is known to be realizable.

To make this automated support possible, some new features, which we believe are unique to the Nuprl system, have been added to the system. In order to motivate a discussion of these features, let us examine in detail the steps a user takes in proving a new realizability result.

Suppose that we want to prove that ϕ is realizable, and we start a proof of the top-level goal $\models \phi$. From the form of the goal, the proof system knows that we must produce a feasible distributed system D that realizes ϕ so it adds a new abstraction $D(x, \dots, z)$ to the library (where x, \dots, z are any parameters mentioned in ϕ). The new abstraction has no definition initially—that will be filled in automatically as the proof proceeds. This initial step leads to a goal where from the hypothesis that an event structure es is consistent with $D(x, \dots, z)$ we must show the conclusion that $\phi(es)$, i.e., that es satisfies ϕ .

Now, suppose that we can prove a lemma stating that in any event structure, es , $\psi_1(es) \wedge \psi_2(es) \Rightarrow \phi(es)$ (the proof of this might be done automatically by a good decision procedure for event structures or interactively in the standard way). In this case, the user can refine the initial goal $\phi(es)$ by asserting the two subgoals $\psi_1(es)$ and $\psi_2(es)$ (and then finishing the proof of $\phi(es)$ using the lemma).

If ψ_1 is already known to be realizable, then there is a lemma $\models \psi_1$ in the library and, since the proofs are constructive, there is a realizer A_1 for ψ_1 . Thus to prove $\psi_1(es)$, it is enough to show that es is consistent with A_1 , and since this follows from the fact that es is consistent with $D(x, \dots, z)$ and that $A_1 \subset D(x, \dots, z)$, the system will automatically refine the goal $\psi_1(es)$ to $A_1 \subset D(x, \dots, z)$. If ψ_2 is also known to be realizable with realizer A_2 then the system produces the subgoal $A_2 \subset D(x, \dots, z)$, and if not, the user uses other lemmas about event structures to refine this goal further.

Whenever the proof reaches a point where the only remaining subgoals are that $D(x, \dots, z)$ is feasible or have the form $A_i \subset D(x, \dots, z)$, then the proof can be completed automatically by defining $D(x, \dots, z)$ to be the join of all the A_i . In this case, all the subgoals of the form $A_i \subset D(x, \dots, z)$ are automatically proved, and only the feasibility proof remains. Since each of the realizers A_i is feasible, the feasibility of their join follows automatically from the pairwise compatibility of the A_i and the system will prove the pairwise compatibility of the realizers A_i automatically if they are indeed compatible, in which case the proof of the realizability of ϕ is complete and its realizer has been constructed and stored in the library.

How might realizer A_1 and A_2 be incompatible? For instance, A_1 might contain a clause that initializes a state variable x to true while A_2 contains a clause that initializes the same variable to false. Or, A_1 might declare that an action of kind k_1 has an effect on x while A_2 declares that only actions of kinds k_2 or k_3 may affect x .

If the variable x occurs explicitly in the top goal ϕ then the user has simply made incompatible design choices in his attempt to realize ϕ and must change the proof. However, if the variable x is not explicitly mentioned in ϕ then it is the case that x can be renamed to y without affecting ϕ . It is often the case that x can be renamed independently in the proofs of the subgoals ψ_1 and ψ_2 (say to y and z) and hence the realizers $A_1(y)$ and $A_2(z)$ will no longer be incompatible.

Incompatibilities such as these can arise when names for variables, local actions, links, locations, or message tags that may be chosen arbitrarily and independently, happen to clash. Managing all of these names is tedious and error prone, so we have added automatic support for managing these names.

By adding some restrictions to the definition mechanism, we are able to ensure that the names inherent in any term are always visible as explicit parameters. We have also defined the semantics of Nuprl in such a way that the *permutation rule* is valid. The permutation rule says that if proposition $\phi(x, y, \dots, z)$ is true, where x, y, \dots, z are the names mentioned in ϕ , then proposition $\phi(x', y', \dots, z')$ is true, where x', y', \dots, z' is the image of x, y, \dots, z under a permutation of all names.

Using the permutation rule, our automated proof assistant will always permute any names that occur in realizers brought in automatically as described above, so that any names that were chosen arbitrarily in the proof of the realizability lemma but were not explicit in that lemma will be renamed to fresh names that will not clash with any names chosen so far. This strategy is supported by the primitive rules of the logic and it guarantees that any incompatibility in the realizer built by the system was inherent in the user's design choices and was not just an unfortunate accident in the choice of names.

3. System architecture

The architecture of the Nuprl [3,78,99] system is the product of many evolutions aimed at providing a theorem proving environment as rich and robust as its type theory. The resulting implementation composes a set of communicated processes, distributing tasks amongst the varied components. The independent communicating processes are centered around a common knowledge base, called the *library*. The library contains definitions,

theorems, inference rules, and meta-level code (e.g. tactics), and serves as a transaction broker for the other processes. Those processes include inference engines (*refiners*), user interfaces (*editors*), rewrite engines (*evaluators*), and *translators*. Translators between the formal knowledge stored in the library and, for instance, programming languages like Java or Ocaml [79,80] allow the formal reasoning tools to supplement real-world software from various domains and thus provide a *logical programming environment* for the respective languages.

Multiple instances and kinds of processes can connect to the library, and can do so at any time during the user's session. This allows for multiple refiners in a session so that proofs can be refined in parallel, for example. At its core, the Nuprl system is based upon a transactional library, connected components, or clients, and library-client communications.

3.1. Transactional library

The ways in which knowledge is stored, structured, changed, and in general managed, stem from desires to provide users with a reliable robust and sustainable system.

All content in the library is stored in the form of library objects. The library maintains a mapping of *abstract object identifiers* to *object content*. Abstract object identifiers prevent name collision problems, since their values cannot be imposed on the library. Instead, unique object identifiers are created by the library while mnemonic names chosen by users are included as simply part of an object's content, and thus need not be unique. When an object refers to another object, it contains that object's identifier. The mapping of object identifiers to object content can be viewed as a graph where an identifier occurring in some object constitutes an edge from this object to the identified object.

Apart from the object map and embedded abstract links to objects, the library does not impose any predefined organization. The absence of a predefined library organization is a prerequisite for integrating formal knowledge from other systems besides Nuprl without requiring these systems to change their representation structure. It is expected that contributors define their own organization of content using the simple but powerful primitives provided in the library. Translators of formal knowledge relating content of disparate systems would require specialized methods of accessing each systems content.

So that users need not worry about losing unsaved work, the library uses a transactional database to make objects persistent. All changes to object contents are committed to the database when a transaction ends, ensuring that knowledge does not get lost in case of a system failure.

In the intended use of the library to support mathematical knowledge management, it is expected that clients will be reading content repeatedly. To facilitate this use, the library implements a *distributed lookup table* to allow clients to maintain local projections of the library content. A projection is a mapping of a subset of library content. Clients can tailor the content in their local tables by defining their own projections. The library maintains the consistency of the projection by broadcasting updates after modification of objects within the subset of the projection.

The library database is constructive in the sense that it never overwrites content. Instead the object identifier is rebound to newly constructed object content. The chain of bindings is preserved so that old versions may be accessed. This allows recovery of user data corrupted

or destroyed erroneously. The original content is preserved until it is deleted during garbage collection of the database. Garbage collection must be explicitly invoked and retains a parameterizable number of generations of the library.

The ACID model is an old and important concept that sets prerequisites for the design of a reliable database. Drawing from this standard, we designed the knowledge base to be sure to provide the four ACID database properties.

Atomicity Each object entry in the library table has three attributes. A *sticky bit* marks the object as a root object, which prevents it from being deleted during garbage collection. A *liveness bit* indicates whether an object can be included in a projection. The *content* is the term data associated with the object.

All persistent modifications to the library are performed via six fundamental *directives*: *Bind* and *unbind* modify the binding of object identifiers to content. *Allow* and *disallow* modify the sticky bit. *Activate* and *deactivate* modify the liveness bit.

There are many functions to manipulate properties and projections of content, but such modifications must be made persistent by unbinding the old content from the object identifier and binding the new content to it.

Having a small symmetric set of basic commands provides for easy implementation of atomicity. As modifications are performed within a transaction the library remembers which command was called (for an unbind it also remembers content prior to unbind). If a transaction fails, then modifications can be undone simply by performing the inverse operation in reverse order. If a transaction succeeds then the committed commands are written to a log.

Each deactivate or activate generates a broadcast of a conditional update. Each undo or commit will then generate a broadcast to either retract or confirm the update. Broadcasts are filtered and routed to clients depending on the projections they request. It is required that an object be inactive at bind or unbind. Thus the following sequence of actions is required to modify an object: deactivate, which generates a conditional delete broadcast; unbind old content; bind new content; activate, which generates a new conditional insert broadcast; and finally commit will generate appropriate confirm broadcasts for the delete and insert updates.

Consistency By extracting access list updates during the scope of a request, the *access list* is used to record dependencies between objects. For example, by extracting the accesses recorded during a refinement the lemma dependencies of the inference created can be recorded. Dependency accounting data are stored with the object content. If clients are robust in reporting such dependencies, dependency accounting can be a powerful tool to convert ephemeral runtime dependencies into persistent links between objects.

Dependency links and object content time stamps provide the ability to detect stale dependencies. If a referenced object has a time stamp later than the referencing object, then some doubt is cast on the validity of the referencing object. Similar tests could be used to cause failure if such conditions are detected at modification of the referenced object. This would prevent stale dependencies, but it also requires more forethought when attempting modification of referenced objects. In the current implementation we have opted to allow inconsistent states, but have provided utilities for detecting such states and refreshing af-

ected objects to resolve them. The concept of *certificates* (see Section 4.3.2) addresses similar concerns in a more rigorous manner.

Isolation Serializability is provided by locking objects. Current implementation uses only write locks and a transaction will fail if it attempts to gain a lock on an object held by another transaction. Read locks are not used but instead Nuprl supports multiple views. Modifications of objects by concurrent transactions are not visible to other transactions. At commit time, modified content is stamped. Each time an object is accessed the transaction remembers the stamp at access time. While Nuprl can record library lookup of objects, clients must also report lookups in distributed tables for the access list to be complete. After a transaction is complete, it can determine if any objects accessed were modified by a concurrent transaction which has since committed. Thus dirty reads are allowed, but can be detected. Given the nature of the expected content it is expected that transactions will read many objects and write few.

Durability The database provides durability for a library. A *library instance* is created from a log in the database. A log is a journal of committed directives. A user opens a library by choosing a log from the database, and a library process can then construct a library instance from the log. The content terms are stored separately from the log and thus each log entry consists of an indication of the operation, the object identifier, and possibly a pointer to the content. The log entries are represented as terms and the database is used to store logs as a sequence of terms. When a log is closed, an abbreviated copy is made containing only the entries needed to produce the library instance. If a library process fails, a new library instance can be produced from the raw log. In this manner durability is provided. When a new library instance is created, a new log is created as well. If the instance was created from a properly closed log, then the first entry will be a persistent pointer to the old log. Otherwise, the new log is initialized with the state of the library after the old log has been completely read.

3.2. Editing and refining proofs

Besides the library, the Nuprl system includes as native clients an editor and a Nuprl refiner for editing and refining proofs, respectively. While one process of each is the most standard configuration, multiple editor and refiner processes can be connected to a single library process. For example, multiple editors can allow for cooperative work on a shared library, and multiple refiners allow for parallel refinement

The Nuprl editor is a structure editor suited to manipulating terms. It provides many features to facilitate browsing and modification of library content and offers special support for editing terms, definitions, theorems, programs, proof tactics, etc. As a structure editor it enables the user to modify terms in a structured fashion by marking subterms and editing slots in the displayed term. Users may also follow hyperlinks between object identifiers contained in a term and the object they refer to.

The Nuprl refiner implements the formal system described in Section 2. It contains a *rule compiler* that translates rule objects contained in the library into inference mecha-

nisms, which makes it possible to extend the logic in the future without re-implementing the refiner. The refiner can also extract programs from proofs and evaluate them.

JProver [81,82,103], a complete theorem prover for first-order intuitionistic and classical logic, is included as an extension to the Nuprl refiner. To ensure validity of a proof built from refiners of disparate logics, JProver is used as a heuristic, and the proof terms are translated and rerun in the Nuprl system. The communication between Nuprl and JProver utilizes the MathBus design [92].

3.3. Communication between the components

The advantages of an open distributed system are vast, providing us not only with the robustness and flexibility we aimed for, but also with an excellent division of labor, particularly useful for targeting mathematical knowledge management. Still, managing the communication between the library and its clients required a significant evolution of earlier versions of the Nuprl system.

In our open architecture, refiner and editor processes will be repeatedly reading content. To facilitate this, the library has robust methods to push object content projections to clients. In our implementation, clients may subscribe to various projections and kinds of content. The library distributes lookup tables to the clients at the time of connection, and sends updates appropriately as changes are made. This minimizes load on the library and maintains the correct state for all processes.

All interprocess communication is done via the exchange of terms. Each process maintains a list of links to connected processes. Currently, term transfer is done over TCP sockets. The terms are encoded into a byte stream on output and then decoded back to a term on input.

Work is accomplished in Nuprl by evaluating requests. Requests must be evaluated in the scope of an *environment*. An environment is characterized by its subscriptions and what requests it can evaluate.

Just as an environment will process requests, it can make requests of other environments. Requests are sent by marshalling the request as a term and then routing the request to a connected environment.

Nuprl also supports a fixed API with predefined commands. The fixed API is suitable for clients which do not wish to support the distributed table features. Such clients could be used for reading Nuprl content or for batch imports of content into Nuprl. There are provisions in the API for a client to provide limited services to Nuprl.

To tightly integrate a system as a Nuprl client, several protocols need to be supported. The message protocol specifies the form and valid sequences of all messages. This protocol includes fields for relating messages to transaction scope and describes forms for delimiting the extent of transactions. The distributed table protocol specifies the form of broadcast messages used to update the distributed tables. The evaluation protocol specifies the form of messages used to make requests, return responses and report distributed table lookups. A client is not required to implement an evaluator in the library's native language (ML). Instead, the library can be extended to build client-specific commands to access client services provided the commands are expressible as terms.

4. A formal digital library

It is very gratifying to researchers in automated reasoning that a great deal of formal mathematics has now been created using interactive theorem provers. There are probably as many as 50,000 formally proved theorems of interest among the files of a dozen major provers.

In the mid 90's, Doug Howe and his collaborators at Bell Labs who were using HOL and Nuprl on various projects decided that it was more efficient to carry out the basic theoretical and practical work needed to share mathematics across these provers than it was to rebuild the results separately. They worked to this end, showing how to embed HOL in Classical Nuprl [64,65] and using the combined library in their work [47,48]. We built on this work to also include PVS results in a library of formal mathematics [3]. Out of that work and our work on the Nuprl System, emerged our project on a Formal Digital Library (FDL). The FDL allows for multiple logics, and the integration of multiple support tools and services. We discuss some of our results and contributions in this section. This is an on-going effort of our research group [30], and it ties us into the Mathematical Knowledge Management community.

4.1. FDL design goals

Mathematical Knowledge Management is concerned with providing access to mathematical knowledge in order to facilitate the creation or discovery of mathematics, facilitate mathematics education, and make feasible intelligent information retrieval in large bodies of digital mathematics. To better serve these goals, a digital library of formalized mathematical and algorithmic knowledge should meet the following objectives.

Connectivity The FDL must be able to connect to multiple clients (proof tools, users, etc.) independently, asynchronously, and in parallel.

Usability Clients of the FDL must be able to browse library contents, search for information by a variety of search criteria, and contribute new knowledge to the library.

Interoperability The FDL shall support the cooperation of proof systems in the development of formal algorithmic knowledge. Different proof systems will be based on different formal theories and on different internal representations of knowledge. The representation of knowledge in the FDL has to be generic, so that it can be translated into a large variety of formats when providing knowledge to clients or receiving formal knowledge from them.

Accountability The FDL needs to be able to account for the integrity of the formalized knowledge it contains. As it supports interoperability between very different proof tools, there cannot be an “absolute” notion of correctness. Instead, the FDL has to provide justifications for the validity of proofs, which will depend upon what rules and axioms are admitted and on the reliability of the inference engines employed. Furthermore, these justi-

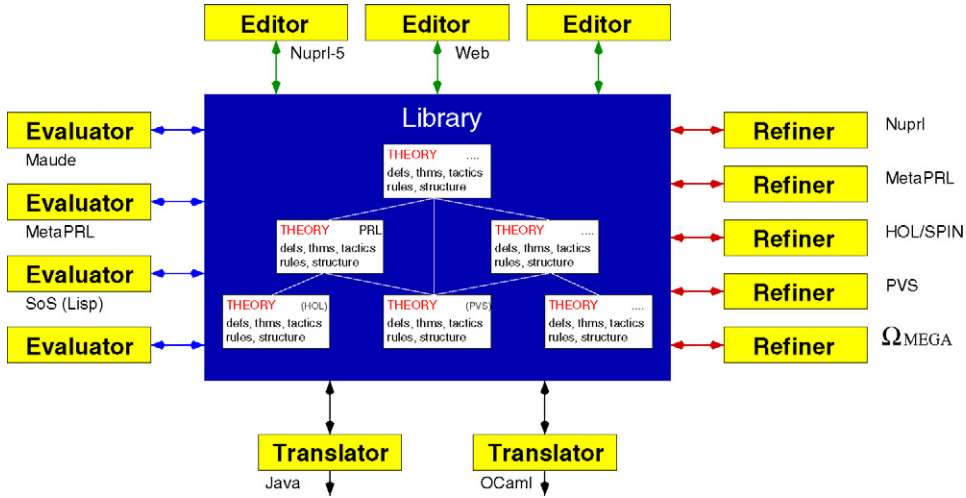


Fig. 1. FDL distributed open architecture.

fications must be exposed to determine the extent to which one may rely upon the provided knowledge. We call these justifications *certificates*.

Information preservation The FDL has to guarantee that existing knowledge and justifications cannot be destroyed or corrupted by clients or system crashes.

Archiving The FDL has to support the management of knowledge on a large scale such as merging separate developments of large theories and performing context-specific tasks. This requires the use of abstract references to knowledge objects, or *abstract identifiers*, as traditional naming schemes do not scale.

We have extended the Nuprl library described in Section 3.1 to a prototype FDL that interfaces with Nuprl refiners and editors and other tools and systems that support formalized mathematical reasoning and knowledge. The FDL system is depicted in Fig. 1. Multiple logics are accommodated via translators, which can map subsets of one logic onto another [4,48,64,65]. Throughout this ongoing project, we have studied concepts, core functionalities and features of formal digital libraries, which we will discuss in the rest of this section. The latest version of the Nuprl library was in fact designed for the intended use of mathematical knowledge management. Constructing the FDL independently from the Nuprl Refiner and Editor as a fully autonomous system, and integrating many of the concepts discussed below provided us with an excellent prototype upon which to innovate further to meet the needs of mathematical knowledge management. Furthermore, it allowed us to enhance the Nuprl system, because we can connect Nuprl refiners to it to get the benefits of the additional FDL functionalities. While concepts including abstract identifiers (Section 4.2.1) and general accounting mechanisms have been fully implemented into this prototype, closed maps (Section 4.2.2) and certificates have not yet been made available to the users, while their design is complete.

4.2. Organizing concepts

4.2.1. Abstract identifiers

The usual method of interaction with the FDL is to build and develop a *client work space*, i.e., a collection of named object contents that provide a specific view of the data and can be tailored to the specific needs and permissions of a client.

In a work space, *abstract object identifiers* [8] are linked to concrete names chosen by a user. This allows the user to organize objects in folders, to use the same name in different folders, and to establish “private” links between objects. The work space may also restrict a client’s access to certain library objects. Most importantly, however, it protects internal identifiers and object contents from being modified without going through the FDL, helps preventing name collisions, and makes proof mechanisms independent of particular naming schemes.

The *library manager* provides clients with utilities for building, storing, and sharing collections of *session objects*. It maintains the work spaces and thus enforces a discipline for building named collections, thus preserving the *coherency* of the collections.

4.2.2. Closed maps

Work spaces are represented by *maps* from a finite set of names to library objects. These maps have to be *closed* in the sense that the objects they refer to do not contain any references to objects that have no name in the map. Thus the basic model of interacting with the library is to maintain a *current closed map* as a part of state that is updated repeatedly as one works.

In general, a *closed map* is a function of type $D \rightarrow \text{Term}(D)$, where D is a finite discrete type of indices and $\text{Term}(D)$ is the type of terms whose subterms only contain abstract identifiers in D . Usually we identify objects in a closed map with their index (or *name*).

In practice the class D will be varied continually. For example, extending a closed map requires selecting a larger index class. Deleting members of a closed map requires a smaller index class. In both cases, we have to make sure that the resulting map remains closed.

If the restriction of a closed map $m \in D \rightarrow \text{Term}(D)$ to a subclass $X \subseteq D$ is itself a closed map, then we call it a *submap* of m . Similarly a *supermap* of m is a closed extension of m to a class $Y \supseteq D$. Two closed maps m and m' are *equivalent*, if they are simply renamings of each other.

Closed maps are essential for defining the notion of *dependency*. Objects depend on others if they directly or indirectly refer to them. An expression $t \in \text{Term}(D)$ *refers* directly to an object (index) $x \in D$ if x occurs within a subterm of t . The notion of dependency is the key to defining *correctness*. While it is possible to define useful notions of correctness with respect to state, the enduring ones can only be formulated in terms of closed maps: the correctness of an object should only depend on the correctness of the object it refers to but not on library objects that are not within the current closed map.

The library is a repository not of closed maps per se, but is rather a repository of data and instructions for building closed maps modulo choice of abstract identifiers. In a session the current closed map is initialized from the library, transformed through a sequence of operations, and then stored back into the library for later retrieval. Several operations can

be defined on closed maps including operations on two maps such as *Merging*, in which objects can be identified, or operations that build a new map by shrinking or expanding a current map, or a subset of a current map, utilizing dependency information.

4.3. Accounting mechanisms

One of the central aspects of a formal digital library is to account for the integrity of its contents and to support arguments for claims of the following form:

Because the library contains a proof of theorem T that refers to a given collection of proof rules and proof engines, theorem T is true if those rules are valid and those engines run correctly.

Accounting mechanisms determine how to execute *inferences* as specified by a proof tactic depending on the actual contents of the library and produce *certificates*, which attest that certain actions were taken at a certain time to account for the validity of an object. Accounting mechanisms are also needed to determine whether a proof built from a collection of certified inferences is acceptable for a given purpose. This would be trivial if the FDL were restricted to a single uniform logic and to a single inference engine. But inferences that may employ a variety of logics and inference engines cannot be simply combined. Instead, certain stipulations limiting proofs to a given set of rules, axioms, and other objects on which these may depend must be expressed and checked.

To account for the validity of library objects, the knowledge base supports *dependency tracking*. For this purpose a variety of information is stored together with an object, e.g. the logical rules and theorems on which it depends, the exact version of the refiners that were used to prove it, timestamps, etc. This information will help a user to validate theorems that rely on knowledge created by several systems, provided that the conditions for *hybrid validity* with respect to the underlying logics are well understood and stored in the library. For instance, a theorem referring to lemmata from Nuprl (constructive type theory) and HOL (classical higher order logic) would be marked as constructively valid if the HOL theorems involve only decidable predicates.

The library only provides methods to store and retrieve dependency data recorded for an object. Tools for examining the relationships between objects from disparate domains must be provided by developers of the systems being related.

Key units of the FDL accounting mechanisms include accounting at the level *inference*, and maintenance of *certificates*.

4.3.1. Inferences

One of the most fundamental mechanisms to account for the validity of library contents is the application of logical inferences from a finite number of premises to a conclusion. Inferences are represented by trees of *inference steps*, which in turn are represented as library objects. A library process checks or generates an inference step by applying *inference engines*, which create proofs in some formal calculus according to user specified methods.

The fact that an inference step has been verified by a given inference engine is represented by an external certificate that refers to the inference step. There may be multiple certificates for the same inference, certifying that the inference has been checked by different inference engines. Depending on the contents of the available certificates, the inference may be considered valid or not in a specific context.

Inference engines support the development of new formal knowledge by providing mechanisms for interactive, tactical, and fully automated reasoning in a specific formal language. As the formal digital library supports almost any formal language, it can be connected to a variety of inference engines that will provide justifications for its formal content.

4.3.2. Certificates

Certificates are the basis for logical accounting. They attest that certain library actions were taken at a certain time to validate the contents of, or identity between, objects. A certificate will be realized as an object, which can then be referenced and accessed like other objects save for certain constraints. A certificate cannot be created or modified except by the library process following a procedure specific to the kind of certificate in question.

Although certificate contents are expected to be rather compact, largely consisting of object references, they will often be expensive to establish. By realizing certificates as objects the library can build certificates that depend on others whose correctness is independently established. Thus one process of certification can contribute to many other certifications without having to be redone.

The paradigmatic certificates are those created to validate proofs. An *inference step certificate* attests to the fact that a specified inference engine accepted that a certain inference step is valid. It is built by applying the engine to the inference, and includes references to the inference step as well as to the instructions for building or deploying the inference engine. A proof is a rooted dag of inference steps. A *proof certificate* is created only when there is an inference certificate for the root inference, and there are already proof certificates for all the proofs of the premises of the root inference.

A certificate may fall into doubt when any object it refers to is modified and needs to be reconsidered and modified in this case. Certificates may also be reconsidered by explicit demand.

As certificates only attest to the fact that the objects they refer to satisfy certain policies for creating these certificates, they provide justifications that are more general than formal proofs in a specific target logic. Nevertheless, they are equally rigorous in the sense that they provide the exact reasons that were used for declaring an object valid.

This aspect is particularly interesting when one considers the possibility that certain inference engines may not be generally accepted. People who do trust a certain inference engine will accept the certificates produced by it while others will insist that the same inferences have to be checked by inference engines they trust. The connection between the FDL and JProver accounts for these two levels of trust: one may either trust that matrix proofs produced by JProver [81] are valid, or one may require that the algorithm for translating matrix proofs into sequent proofs [82] be executed and that the results will be checked with a proof checker for the intuitionistic sequent calculus.

4.4. Information science considerations

With the capabilities to merge vast collections of formal mathematics, to account for their correctness and origins, and to use these libraries in the further developments of formal mathematics, we see both a need and an opportunity to utilize and adapt information science techniques.

4.4.1. Publishing and reading

One of the key services that a library must provide is an interface that makes formal algorithmic knowledge accessible to users. Efforts such as MoWGLI and the HELM project [55,94] attest to this importance. As we envision a variety of users who would benefit from being able to inspect the contents of a formal digital library, we have developed a *publication mechanism* that enables external users to access the logical library and to browse its contents without having to run a local copy of it.

In [97] we made a first step towards publishing our formal mathematics on the web. Our interface supports viewing formal contents at all levels of precision—from the rough sketch of a theory or proof down to the level of the underlying logic. We expect that the ability to unveil formal details on demand will have a significant educational value for teaching and understanding mathematical and algorithmic concepts.

We have built a large utility for converting related collections of objects to HTML, along with automatically generated structure-revealing documents and auxiliary annotations. When the utility is run, the various browsable pages and some print forms are generated, missing links are reported diagnostically, and a decision is rendered about whether the result is suitable for posting. The structure of the generated web material is thoroughly explained on the several documentation pages to which one may link from every page so created.

Currently, the publication mechanisms are capable of handling formal content created by Nuprl. Automatically showing new kinds of content in this way, such as pages about PVS libraries, requires appropriate modification and specialization. Not only might the presentation of each object need adjustment, but the relations to be revealed between documents must be determined and accommodated.

4.4.2. A whole greater than its sum

From an information science perspective, the FDL is a dynamic information network, with library objects as its vertices and relationships between them as its arcs, not unlike web pages and hypertext. There can coexist multiple kinds of arcs, defined at varying levels of detail, such as object to object or theory to theory, as well as multiple kinds of objects. The collection of these objects into such a network facilitates a deeper and richer understanding than instances of these objects, or even instances of these objects together with meta-properties. We have already seen this in numerous accounts for the web [70, 102].

One meaningful and straightforward relationship in the FDL network is the logical dependency relationship. If we limit the links to logical dependencies, for example, then, unlike with hypertext on the web, we gain a well defined linking mechanism. A popular link analysis method to reveal relevancy in web search, HITS [70], was applied and

adapted to the FDL dependency network in [87]. The algorithm finds popular nodes, which is useful for ranking query-based search results, and also finds communities, or clusters of related objects, which may be useful for automated organization of the library. Looking at the graph as a whole, it appears that our FDL mathematics collections are scale-free [11]. A scale-free network includes “very connected” nodes, which greatly shorten the likely distance between two nodes from that of a randomly connected graph with the same number of nodes and arcs. A characteristic depth and breadth of the graph was also revealed. Initial comparisons were made between the FDL network and the dependency network of a collection of Coq proofs, yielding a mix of similar and dissimilar patterns. Relationships to other objects, and the context within which a math object resides, can bring added support to the task of information retrieval in the math domain.

5. Conclusion

The Nuprl system, its Computational Type Theory, and the design of a formal digital library are the result of more than 20 years of research and practical experiments with mathematical proof assistants. In addition to aspects of theoretical developments and issues of implementing state-of-the-art proof environments there have been numerous applications ranging from mathematics and formal courseware to hardware and software design, verification, and optimization [1,18,19,28,34,56,62,64,68,77,79,80,85,86,96]. Much of our current research supports the interaction between Nuprl and other mathematical assistants, proof engines, and publication mechanisms—both from the technological and from the logical perspective—which will lead to a greater acceptance of these systems in the practice of mathematics and software development.

Acknowledgements

This material is based upon work supported by the National Science Foundation under Grants No. 0204193 (Proof Automation in Constructive Type Theory) and 0208536 (Innovative Programming Technology for Embedded Systems). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- [1] M. Aagaard, M. Leiser, Verifying a logic synthesis tool in Nuprl, in: G. Bochmann, D. Probst (Eds.), Proceedings of Workshop on Computer-Aided Verification, in: Lecture Notes in Computer Science, vol. 663, Springer-Verlag, Berlin, 1993, pp. 72–83.
- [2] J. Abbott, A. van Leeuwen, A. Strotmann, Openmath: Communicating mathematical information between co-operating agents in a knowledge network, *J. Intelligent Syst.* (1998).
- [3] S. Allen, R. Constable, R. Eaton, C. Kreitz, L. Lorigo, The Nuprl open logical environment, in: D. McAllester (Ed.), Proceedings of 17th Conference on Automated Deduction, in: Lecture Notes in Artificial Intelligence, vol. 1831, Springer-Verlag, Berlin, 2000, pp. 170–176.

- [4] S. Allen, M. Bickford, R. Constable, R. Eaton, C. Kreitz, A Nuprl-PVS connection: Integrating libraries of formal mathematics, Technical Report, Cornell University. Department of Computer Science, 2003.
- [5] S.F. Allen, A non-type-theoretic definition of Martin-Löf's types, in: [51], pp. 215–224.
- [6] S.F. Allen, A non-type-theoretic semantics for type-theoretic language, PhD thesis, Cornell University, 1987.
- [7] S.F. Allen, From dy/dx to $[]_p$: a matter of notation, in: Proceedings of the Conference on User Interfaces for Theorem Provers, Eindhoven, The Netherlands, 1998.
- [8] S.F. Allen, Abstract identifiers, intertextual reference and a computational basis for recordkeeping, First Monday 9 (2) (2004).
- [9] A. Asperti, L. Padovani, C. Sacerdoti Coen, F. Guidi, I. Schena, Mathematical knowledge management in HELM, Ann. Math. Artificial Intelligence 38 (1–3) (2003) 27–46.
- [10] R.C. Backhouse, P. Chisholm, G. Malcolm, Do-it-yourself type theory (part II), EATCS Bull. 35 (1988) 205–245.
- [11] A.L. Barabasi, R. Albert, Emergence of scaling in random networks, pp. 509–512.
- [12] H.P. Barendregt, Handbook of Logic in Computer Science, vol. 2, Oxford University Press, Oxford, 1992, Chapter: Lambda Calculi with Types, pp. 118–310.
- [13] J.L. Bates, A logic for correct program development, PhD thesis, Cornell University, 1979.
- [14] J.L. Bates, R.L. Constable, Proofs as programs, ACM Trans. Programming Lang. Syst. 7 (1) (1985) 53–71.
- [15] Y. Bertot, P. Castéran, Interactive Theorem Proving and Program Development; Coq'Art: The Calculus of Inductive Constructions, Texts in Theoretical Computer Science, Springer-Verlag, 2004.
- [16] G. Betarte, A. Tasistro, Extension of Martin-Löf's type theory with record types and subtyping, in: Twenty-Five Years of Constructive Type Theory, Oxford University Press, Oxford, 1998, pp. 21–39.
- [17] M. Bickford, R.L. Constable, A logic of events, Technical Report TR2003-1893, Cornell University, 2003.
- [18] M. Bickford, C. Kreitz, R. van Renesse, X. Liu, Proving hybrid protocols correct, in: R. Boulton, P. Jackson (Eds.), Proceedings of 14th International Conference on Theorem Proving in Higher Order Logics, in: Lecture Notes in Computer Science, vol. 2152, Springer-Verlag, Berlin, 2001, pp. 105–120.
- [19] K. Birman, R. Constable, M. Hayden, J. Hickey, C. Kreitz, R. van Renesse, O. Rodeh, W. Vogels, The Horus and Ensemble projects: Accomplishments and limitations, in: DARPA Information Survivability Conference and Exposition, IEEE Computer Society Press, 2000, pp. 149–160.
- [20] E. Bishop, Foundations of Constructive Analysis, McGraw-Hill, New York, 1967.
- [21] N. Bourbaki, Elements of Mathematics, Algebra, vol. 1, Addison-Wesley, Reading, MA, 1968.
- [22] R.S. Boyer, J.S. Moore, A Computational Logic, Academic Press, New York, 1979.
- [23] D. Carlisle, P. Ion, R. Minder, N. Poppelier, Mathematical markup language (MathML) 2.0 w3c recommendation, February 2001, www.w3.org/TR/2001/REC-MathML2-20010221/.
- [24] A. Church, Introduction to Mathematical Logic, vol. I, Princeton University Press, 1956.
- [25] R.L. Constable, Constructive mathematics and automatic program writers, in: Proceedings of the IFIP Congress, North-Holland, Amsterdam, 1971, pp. 229–233.
- [26] R.L. Constable, Mathematics as programming, in: Proceedings of the Workshop on Programming and Logics, in: Lecture Notes in Computer Science, vol. 164, Springer-Verlag, Berlin, 1983, pp. 116–128.
- [27] R.L. Constable, The semantics of evidence, Technical Report TR85-684, Cornell University, 1985.
- [28] R.L. Constable, Creating and evaluating interactive formal courseware for mathematics and computing, in: M.F. Iskander, M.J. Gonzalez, G.L. Engel, C.K. Rushforth, M.A. Yoder, R.W. Grow, C.H. Durney (Eds.), Frontiers in Education, IEEE, 1996.
- [29] R.L. Constable, Types in logic, mathematics and programming, in: S.R. Buss (Ed.), Handbook of Proof Theory, Elsevier, Amsterdam, 1998, pp. 683–786, Chapter X.
- [30] R.L. Constable, S. Allen, M. Bickford, J. Caldwell, J. Hickey, C. Kreitz, Steps Toward a World Wide Digital Library of Formal Algorithmic Knowledge, vol. 1, 2003.
- [31] R.L. Constable, et al., Implementing Mathematics with the Nuprl Proof Development System, Prentice-Hall, Englewood Cliffs, NJ, 1986.
- [32] R.L. Constable, J. Hickey, Nuprl's class theory and its applications, in: F.L. Bauer, R. Steinbrueggen (Eds.), Foundations of Secure Computation, in: NATO ASI Series, Series F: Computer & System Sciences, IOS Press, Amsterdam, 2000, pp. 91–116.
- [33] R.L. Constable, D.J. Howe, Implementing metamathematics as an approach to automatic theorem proving, in: R.B. Banerji (Ed.), Formal Techniques in Artificial Intelligence: A Source Book, Elsevier Science, Amsterdam, 1990, pp. 45–76.

- [34] R.L. Constable, P.B. Jackson, P. Naumov, J. Uribe, Constructively formalizing automata, in: *Proof, Language and Interaction: Essays in Honour of Robin Milner*, MIT Press, Cambridge, MA, 2000, pp. 213–238.
- [35] R.L. Constable, S. Johnson, C. Eichenlaub, *Introduction to the PL/CV2 Programming Logic*, Lecture Notes in Computer Science, vol. 135, Springer-Verlag, Berlin, 1982.
- [36] R.L. Constable, S.D. Johnson, *A PL/CV précis*, in: *Principles of Programming Languages*, ACM, New York, 1979, pp. 7–20.
- [37] R.L. Constable, T. Knoblock, J.L. Bates, Writing programs that construct proofs, *J. Automat. Reason.* 1 (3) (1984) 285–326.
- [38] R.L. Constable, M.J. O'Donnell, *A Programming Logic*, Winthrop, MA, 1978.
- [39] R.L. Constable, D.R. Zlatin, The type theory of PL/CV3, *ACM Trans. Programming Languages and Systems* 6 (1) (1984) 94–117.
- [40] C. Cornes, J. Courant, J.-C. Filliâtre, G. Huet, P. Manoury, C. Paulin-Mohring, C. Muñoz, C. Murthy, C. Parent, A. Saïbi, B. Werner, *The Coq Proof Assistant reference manual*, Technical Report, INRIA, 1995.
- [41] K. Cray, Foundations for the implementation of higher-order subtyping, in: *1997 ACM SIGPLAN International Conference on Functional Programming*, 1997.
- [42] K. Cray, Simple, efficient object encoding using intersection types, Technical Report TR98-1675, Department of Computer Science, Cornell University, April 1998.
- [43] K. Cray, *Type-theoretic methodology for practical programming languages*, PhD thesis, Cornell University, Ithaca, NY, August 1998.
- [44] H.B. Curry, R. Feys, W. Craig, *Combinatory Logic*, vol. I, *Studies in Logic and the Foundations of Mathematics*, North-Holland, Amsterdam, 1958.
- [45] N.G. de Bruijn, The mathematical language Automath: its usage and some of its extensions, in: J.P. Seldin, J.R. Hindley (Eds.), *Symposium on Automatic Demonstration*, in: *Lecture Notes in Mathematics*, vol. 125, Springer-Verlag, Berlin, 1970, pp. 29–61.
- [46] P. Dybjer, Inductive sets and families in Martin-Löf's type theory and their set-theoretic semantics, in: G. Huet, G. Plotkin (Eds.), *Logical Frameworks*, Cambridge University Press, Cambridge, 1991, pp. 280–306.
- [47] A. Felty, F. Stomp, A correctness proof of a cache coherence protocol, in: *Proceedings of 11th Annual Conference on Computer Assurance*, 1996, pp. 128–141.
- [48] A.P. Felty, D.J. Howe, Hybrid interactive theorem proving using Nuprl and HOL, in: W. McCune (Ed.), *Proceedings of the 14th International Conference on Automated Deduction*, in: *Lecture Notes in Artificial Intelligence*, vol. 1249, Springer-Verlag, Berlin, 1997, pp. 351–365.
- [49] J.-Y. Girard, P. Taylor, Y. Lafont, *Proofs and Types*, Cambridge Tracts in Computer Science, vol. 7, Cambridge University Press, Cambridge, 1989.
- [50] M. Gordon, R. Milner, C. Wadsworth, *Edinburgh LCF: A Mechanized Logic of Computation*, Lecture Notes in Computer Science, vol. 78, Springer-Verlag, Berlin, 1979.
- [51] D. Gries (Ed.), *Proceedings of the 2nd IEEE Symposium on Logic in Computer Science*, IEEE Computer Society Press, 1987.
- [52] T.G. Griffin, *Notational definition and top-down refinement for interactive proof development systems*, PhD thesis, Cornell University, 1988.
- [53] R. Harper, M. Lillibridge, A type-theoretic approach to higher-order modules with sharing, in: *Proceedings of 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Portland, OR, January 1994, pp. 123–137.
- [54] R. Harper, J.C. Mitchell, On the type structure of standard ML, *ACM Trans. Programming Lang. Syst.* 15 (2) (1993) 211–252.
- [55] HELM, Hypertextual electronic library of mathematics, <http://www.cs.unibo.it/~asperti/HELM>.
- [56] J. Hickey, N. Lynch, R. van Renesse, Specifications and proofs for Ensemble layers, in: R. Cleaveland (Ed.), *Proceedings of 5th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, in: *Lecture Notes in Computer Science*, vol. 1579, Springer-Verlag, Berlin, 1999, pp. 119–133.
- [57] J. Hickey, A. Nogin, R.L. Constable, B.E. Aydemir, E. Barzilay, Y. Bryukhov, R. Eaton, A. Granicz, A. Kopylov, C. Kreitz, V.N. Krupski, L. Lorigo, S. Schmitt, C. Witty, X. Yu, *MetaPRL—A modular logical environment*, in: D. Basin, B. Wolff (Eds.), *Proceedings of the 16th International Conference on Theorem Proving in Higher Order Logics*, in: *Lecture Notes in Computer Science*, vol. 2758, Springer-Verlag, Berlin, 2003, pp. 287–303.

- [58] J.J. Hickey, The MetaPRL logical programming environment, PhD thesis, Cornell University, January 2001.
- [59] C.A.R. Hoare, Notes on data structuring, in: *Structured Programming*, Academic Press, New York, 1972.
- [60] J.G. Hook, D.J. Howe, Impredicative strong existential equivalent to type:type, Technical Report TR 86-760, Cornell University, Department of Computer Science, June 1986.
- [61] W. Howard, The formulas-as-types notion of construction, in: *To H.B. Curry: Essays on Combinatory Logic, Lambda-Calculus and Formalism*, Academic Press, New York, 1980, pp. 479–490.
- [62] D.J. Howe, The computational behaviour of Girard’s paradox, in: D. Gries (Ed.), *LICS-87—Second Annual Symposium on Logic in Computer Science*, IEEE Computer Society Press, 1987, pp. 205–214.
- [63] D.J. Howe, Equality in lazy computation systems, in: *Proceedings of 4th IEEE Symposium on Logic in Computer Science*, IEEE Computer Society Press, pp. 198–203.
- [64] D.J. Howe, Importing mathematics from HOL into Nuprl, in: J. von Wright, J. Grundy, J. Harrison (Eds.), *Theorem Proving in Higher Order Logics*, in: *Lecture Notes in Computer Science*, vol. 1125, Springer-Verlag, Berlin, 1996, pp. 267–282.
- [65] D.J. Howe, Semantic foundations for embedding HOL in Nuprl, in: M. Wirsing, M. Nivat (Eds.), *Algebraic Methodology and Software Technology*, in: *Lecture Notes in Computer Science*, vol. 1101, Springer-Verlag, Berlin, 1996, pp. 85–101.
- [66] D.J. Howe, A classical set-theoretic model of polymorphic extensional type theory, Unpublished manuscript, 1997, available at <ftp://ftp.research.bell-labs.com/dist/howe/set-nuprl.ps.gz>.
- [67] D.J. Howe, Toward sharing libraries of mathematics between theorem provers, in: D.M. Gabbay, M. de Rijke (Eds.), *Frontiers of Combining Systems 2*, in: *Studies in Logic and Computation*, vol. 7, Research Studies Press/Wiley, 2000, pp. 161–176.
- [68] P.B. Jackson, Exploring abstract algebra in constructive type theory, in: A. Bundy (Ed.), *Proceedings of the 12th Conference on Automated Deduction*, in: *Lecture Notes in Artificial Intelligence*, vol. 814, Springer-Verlag, Berlin, 1994.
- [69] S.C. Kleene, *Introduction to Metamathematics*, D. Van Nostrand, Princeton, 1952.
- [70] J. Kleinberg, Authoritative sources in a hyperlinked environment, *Proceedings of the Ninth ACM-SIAM Symposium on Discrete Algorithms*, 1998.
- [71] T. Knoblock, *Metamathematical extensibility in type theory*, PhD thesis, Cornell University, 1987.
- [72] M. Kohlhase, OmDoc: An infrastructure for openmath content dictionary information, in: *Bulletin of the ACM Special Interest Group for Algorithmic Mathematics SIGSAM*, 2000.
- [73] M. Kohlhase, OmDoc: Towards an Internet standard for the administration, distribution and teaching of mathematical knowledge, in: *Proceedings of Artificial Intelligence and Symbolic Computation*, in: *Lecture Notes in Computer Science*, vol. 1930, Springer-Verlag, Berlin, 2000.
- [74] A. Kopylov, Dependent intersection: A new way of defining records in type theory, Department of Computer Science TR2000-1809, Cornell University, 2000.
- [75] A. Kopylov, Dependent intersection: A new way of defining records in type theory, in: *Proceedings of 8th IEEE Symposium on Logic in Computer Science*, 2003, pp. 86–95.
- [76] A.P. Kopylov, Type theoretical foundations for data structures, classes, and objects, PhD thesis, Cornell University, January 2004.
- [77] C. Kreitz, Automated fast-track reconfiguration of group communication systems, in: R. Cleaveland (Ed.), *Proceedings of 5th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, in: *Lecture Notes in Computer Science*, vol. 1579, Springer-Verlag, Berlin, 1999, pp. 104–118.
- [78] C. Kreitz, The Nuprl Proof Development System, version 5: Reference Manual and User’s Guide, Cornell University, Department of Computer Science, 2002.
- [79] C. Kreitz, Building reliable, high-performance networks with the Nuprl proof development system, *J. Funct. Programming* 14 (1) (2004) 21–68.
- [80] C. Kreitz, M. Hayden, J. Hickey, A proof environment for the development of group communication systems, in: C. Kirchner, H. Kirchner (Eds.), *Proceedings of 15th Conference on Automated Deduction*, in: *Lecture Notes in Artificial Intelligence*, vol. 1421, Springer-Verlag, Berlin, 1998, pp. 317–331.
- [81] C. Kreitz, J. Otten, Connection-based theorem proving in classical and non-classical logics, *J. Universal Comput. Sci.* 5 (3) (1999) 88–112.
- [82] C. Kreitz, S. Schmitt, A uniform procedure for converting matrix proofs into sequent-style systems, *J. Inform. Comput.* 162 (1–2) (2000) 226–254.

- [83] L. Lamport, Time, clocks and the ordering of events in a distributed system, *Comm. ACM* 21 (7) (1978) 558–565.
- [84] L. Lamport, The mutual exclusion problem, part ii: statement and solutions, *J. ACM* 33 (2) (1986) 327–348.
- [85] X. Liu, C. Kreitz, R. van Renesse, J. Hickey, M. Hayden, K. Birman, R. Constable, Building reliable, high-performance communication systems from components, in: *Proceedings of 17th ACM Symposium on Operating Systems Principles (SOSP’99)*, *Oper. Syst. Rev.* 34 (1999) 80–92.
- [86] X. Liu, R. van Renesse, M. Bickford, C. Kreitz, R. Constable, Protocol switching: Exploiting meta-properties, in: L. Rodrigues, M. Raynal (Eds.), *International Workshop on Applied Reliable Group Communication (WARGC 2001)*, IEEE Computer Society Press, 2001, pp. 37–42.
- [87] L. Lorigo, J. Kleinberg, R. Eaton, R. Constable, A graph-based approach towards discerning inherent structures in a digital library of formal mathematics, in: *Lecture Notes in Computer Science*, Springer-Verlag, Berlin, 2004, pp. 220–235.
- [88] N. Lynch, *Distributed Algorithms*, Morgan Kaufmann, San Mateo, CA, 1996.
- [89] P. Martin-Löf, An intuitionistic theory of types: Predicative part, in: *Logic Colloquium ’73*, North-Holland, Amsterdam, 1973, pp. 73–118.
- [90] P. Martin-Löf, Constructive mathematics and computer programming, in: *Proceedings of the Sixth International Congress for Logic, Methodology, and Philosophy of Science*, North-Holland, Amsterdam, 1982, pp. 153–175.
- [91] P. Martin-Löf, *Intuitionistic Type Theory*, *Studies in Proof Theory*, Lecture Notes, vol. 1, Bibliopolis, Napoli, 1984.
- [92] The MathBus Term Structure, www.nuprl.org/mathbus/mathbusTOC.htm.
- [93] P.F. Mender, Inductive definition in type theory, PhD thesis, Cornell University, 1988.
- [94] MoWGLI: Mathematic on the web: Get it by logics and interfaces, <http://mowgli.cs.unibo.it>.
- [95] C. Murthy, A computational analysis of Girard’s translation and LC, in: *Seventh Symposium on Logic in Computer Science*, 1992, pp. 90–101.
- [96] C.R. Murthy, J.R. Russell, A constructive proof of Higman’s lemma, in: J.C. Mitchell (Ed.), *Fifth Annual Symposium on Logic in Computer Science*, 1990, IEEE Computer Society Press, 1990, pp. 257–267.
- [97] P. Naumov, Publishing formal mathematics on the web, Technical Report TR98-1689, Cornell University, Department of Computer Science, 1998.
- [98] B. Nordström, K. Petersson, J.M. Smith, *Programming in Martin-Löf’s Type Theory*, Oxford Sciences Publication, Oxford, 1990.
- [99] Nuprl home page, <http://www.nuprl.org>.
- [100] R. Pollack, Dependently typed records for representing mathematical structure, in: J. Harrison, M. Aagaard (Eds.), *Proceedings of 13th International Conference on Theorem Proving in Higher Order Logics*, in: *Lecture Notes in Computer Science*, vol. 1869, Springer-Verlag, Berlin, 2000, pp. 461–478.
- [101] B. Russell, Mathematical logic as based on a theory of types, *Am. J. Math.* 30 (1908) 222–262.
- [102] L. Page, S. Brin, The anatomy of a large-scale hypertextual (web) search engine, in: *Seventh International World Wide Web Conference*, 1998.
- [103] S. Schmitt, L. Lorigo, C. Kreitz, A. Nogin, JProver: Integrating connection-based theorem proving into interactive proof assistants, in: R. Gore, A. Leitsch, T. Nipkow (Eds.), *International Joint Conference on Automated Reasoning*, in: *Lecture Notes in Artificial Intelligence*, vol. 2083, Springer-Verlag, Berlin, 2001, pp. 421–426.
- [104] D. Scott, Data types as lattices, *SIAM J. Comput.* 5 (1976) 522–587.
- [105] R.M. Smullyan, *First-Order Logic*, Springer-Verlag, Berlin, 1968.
- [106] S. Thompson, *Type Theory and Functional Programming*, Addison-Wesley, Reading, MA, 1991.
- [107] A. Sjerp Troelstra, *Metamathematical Investigation of Intuitionistic Mathematics*, *Lecture Notes in Mathematics*, vol. 344, Springer-Verlag, Berlin, 1973.
- [108] A.S. Troelstra, D. van Dalen, *Constructivism in Mathematics, An Introduction*, vols. I, II, North-Holland, Amsterdam, 1988.
- [109] University of Cambridge, Computer Laboratory, *The HOL System: DESCRIPTION*, March 1994.
- [110] A.N. Whitehead, B. Russell, *Principia Mathematica*, vols. 1, 2, 3, second ed., Cambridge University Press, 1925–1927.
- [111] J. Zimmer, M. Kohlhasse, The Mathweb Software Bus for distributed mathematical reasoning, in: H. Ganzinger (Ed.), *Proceedings of 18th International Conference on Automated Deduction*, in: *Lecture Notes in Artificial Intelligence*, vol. 1632, Springer-Verlag, Berlin, 2002, pp. 217–221.