

A Dual Semantics for the Data Description Calculus (Extended Abstract)

Yitzhak Mandelbaum^{1,2} Kathleen Fisher¹ David Walker²

¹ AT&T Research

² Princeton University

1 INTRODUCTION

Data description languages are a class of domain specific languages for specifying *ad hoc data formats*, from billing records to TCP packets to scientific data sets to server logs. Examples of such languages include BRO [13], DATASCRIP [1], DEMETER [8], PACKETTYPES [12], PADS/C [3], PADS/ML [11] and XSUGAR [2], among others. All of these languages generate parsers from data descriptions. In addition, and unlike conventional parsing tools such as Lex and Yacc, many also automatically generate auxiliary tools ranging from printers to XML converters to visitor libraries to visualization and editor tools.

In previous work, we developed the *Data Description Calculus* (DDC), a calculus of simple, orthogonal type constructors, designed to capture the core features of many existing type-based data description languages [4, 5]. This calculus had a multi-part denotational semantics that interpreted the type constructors as (1) parsers that transform external bit strings into internal data representations and *parse descriptors* (representations of parser errors), (2) types for the data representations and parse descriptors, and (3) types for the parsers as a whole. We proved that this multi-part semantics was coherent in the sense that the generated parsers always have the expected types and generate representations that satisfy an important *canonical forms* lemma.

The DDC has been very useful already, helping us debug and improve several aspects of PADS/C [3], and serving as a guide for the design of PADS/ML [11]. However, this initial work on the DDC told only a fraction of the semantic story concerning data description languages. As mentioned above, many of these languages not only provide parsers, but also other tools. Amongst the most common auxiliary tools are printers, as reliable communication between programs, either through the file system or over the Web, depends upon both input (parsing) and output (printing).

In this work, we begin to address the limitations of DDC by specifying a printing semantics for the various features of the calculus. We also prove a collection of theorems for the new semantics that serve as duals to our theorems concerning parsing. This new printing semantics has many of the same practical benefits as our older parsing semantics: We can use it as a check against the correctness of our printer implementations and as a guide for the implementation of future data description languages.

$$\begin{array}{lcl}
\text{Kinds } \kappa & ::= & T \mid T \rightarrow \kappa \mid \sigma \rightarrow \kappa \\
\text{Types } \tau & ::= & C(e) \mid \lambda x. \tau \mid \tau e \mid \Sigma x: \tau. \tau \mid \tau + \tau \\
& & \mid \{x: \tau \mid e\} \mid \alpha \mid \lambda \alpha. \tau \mid \tau \tau \mid \dots
\end{array}$$

FIGURE 1. DDC^α syntax, selected constructs

In this extended abstract, we give an brief overview of the calculus, its dual semantics and their properties. A companion technical report contains a complete formal specification [10]. In comparison to our previous work on the DDC at POPL 06 [4], the calculus we present here has been streamlined in several subtle, but useful ways. It has also been improved through the addition of polymorphic types. We call this new polymorphic variant DDC^α. These improvements and extensions, together with proofs, appear in Mandelbaum’s thesis [9] and in a recently submitted journal article [5]. This abstract reviews the DDC^α and extends all the previous work with a printing semantics and appropriate theorems. To be more specific, sections 2 through 3 present the extended DDC^α calculus, focusing on the semantics of polymorphic types for parsing and the key elements of the printing semantics. Then, Section 4 shows that both parsers and printers in the DDC^α are type correct and furthermore that parsers produce pairs of parsed data and parse descriptors in *canonical form*, and that printers, given data in canonical form, print successfully. We briefly discuss related work in Section 5, and conclude in Section 6.

2 DDC^α SYNTAX

Figure 1 summarizes the syntax of DDC^α. The syntax is parameterized by that of a *host language* – the language in which DDC^α parsers and printers are encoded. For concreteness, the host language of DDC^α is a straightforward extension of F_{ω} with recursion and a variety of useful constants and operators. The simplest description is a base type $C(e)$. The base type’s parameter e is an expression drawn from the host language. The PADS/ML type `Pstring` is an example of such a base type. Structured types include value abstraction $\lambda x. \tau$ and application τe , which allow us to parameterize types by host language values. The dependent sum type, $\Sigma x: \tau. \tau$, describes a pair of values, where the value of the first element of the pair (x) can be referenced when describing the second element. Variation in a data source can be described with the sum type $\tau + \tau$, which deterministically describes a data source that either matches the first type, or fails to match the first branch but does match the second one. We specify semantic constraints over a data source with type $\{x: \tau \mid e\}$, which describes any data that satisfies the description τ and the predicate e . Within host language expression e , the value x is bound to the result of parsing the data according τ . Type abstraction $\lambda \alpha. \tau$ and application $\tau \tau$ allow us to parameterize types by other types.

To specify the well-formedness of types, we use a kinding judgment of the form

$$\boxed{\llbracket \tau \rrbracket_{\text{rep}} = \sigma}$$

$$\begin{aligned}
\llbracket C(e) \rrbracket_{\text{rep}} &= \mathcal{B}_{\text{type}(C)} + \text{none} \\
\llbracket \lambda x. \tau \rrbracket_{\text{rep}} &= \llbracket \tau \rrbracket_{\text{rep}} \\
\llbracket \tau e \rrbracket_{\text{rep}} &= \llbracket \tau \rrbracket_{\text{rep}} \\
\llbracket \Sigma x: \tau_1. \tau_2 \rrbracket_{\text{rep}} &= \llbracket \tau_1 \rrbracket_{\text{rep}} * \llbracket \tau_2 \rrbracket_{\text{rep}} \\
\llbracket \tau_1 + \tau_2 \rrbracket_{\text{rep}} &= \llbracket \tau_1 \rrbracket_{\text{rep}} + \llbracket \tau_2 \rrbracket_{\text{rep}} \\
\llbracket \{x: \tau \mid e\} \rrbracket_{\text{rep}} &= \llbracket \tau \rrbracket_{\text{rep}} + \llbracket \tau \rrbracket_{\text{rep}} \\
\llbracket \alpha \rrbracket_{\text{rep}} &= \alpha_{\text{rep}} \\
\llbracket \lambda \alpha. \tau \rrbracket_{\text{rep}} &= \lambda \alpha_{\text{rep}}. \llbracket \tau \rrbracket_{\text{rep}} \\
\llbracket \tau_1 \tau_2 \rrbracket_{\text{rep}} &= \llbracket \tau_1 \rrbracket_{\text{rep}} \llbracket \tau_2 \rrbracket_{\text{rep}}
\end{aligned}$$

FIGURE 2. Representation type translation, selected constructs

$\Delta; \Gamma \vdash \tau : \kappa$, where Δ maps type variables to kinds and Γ maps host language value variables to host language types (σ). In our original work [4], the kinding rules were somewhat unorthodox, but we have since simplified them. Details appear in the companion technical report [10].

3 DDC^α SEMANTICS

The primitives of DDC^α each have four interpretations: two types in the host language, one for the data representation itself and one for its parse descriptor, and two functions, one for parsing and one for printing. We therefore specify the semantics of DDC^α types using four semantic functions, each of which precisely conveys a particular facet of the meaning of a type. The functions $\llbracket \cdot \rrbracket_{\text{rep}}$ and $\llbracket \cdot \rrbracket_{\text{PD}}$ describe the type of the data's in-memory representation and parse descriptor, respectively. The semantic functions $\llbracket \cdot \rrbracket_{\text{P}}$ and $\llbracket \cdot \rrbracket_{\text{PP}}$ define the parsing and printing functions generated from DDC^α descriptions.

3.1 DDC^α representation and parse-descriptor types

In Figure 2, we present the representation type of selected DDC^α primitives. While the primitives are dependent types, the mapping to the host language erases the dependency because the host language does not have dependent types. This involves erasing all host language expressions that appear in types as well as value abstractions and applications. A type variable α in DDC^α is mapped to a corresponding type variable α_{rep} in F_{ω} . Polymorphic types and their application become F_{ω} type constructors and type application, respectively.

Figure 3 gives the types of the parse descriptors corresponding to selected DDC^α types. All parse descriptors share a common structure, consisting of two components, a header and a body. The header reports on the corresponding representation as a whole. It stores the number of errors encountered during parsing,

$$\boxed{\llbracket \tau \rrbracket_{\text{PD}} = \sigma}$$

$$\begin{aligned} \llbracket C(e) \rrbracket_{\text{PD}} &= \text{pd_hdr} * \text{unit} \\ \llbracket \lambda x. \tau \rrbracket_{\text{PD}} &= \llbracket \tau \rrbracket_{\text{PD}} \\ \llbracket \tau e \rrbracket_{\text{PD}} &= \llbracket \tau \rrbracket_{\text{PD}} \\ \llbracket \Sigma x: \tau_1. \tau_2 \rrbracket_{\text{PD}} &= \text{pd_hdr} * \llbracket \tau_1 \rrbracket_{\text{PD}} * \llbracket \tau_2 \rrbracket_{\text{PD}} \\ \llbracket \tau_1 + \tau_2 \rrbracket_{\text{PD}} &= \text{pd_hdr} * (\llbracket \tau_1 \rrbracket_{\text{PD}} + \llbracket \tau_2 \rrbracket_{\text{PD}}) \\ \llbracket \{x: \tau \mid e\} \rrbracket_{\text{PD}} &= \text{pd_hdr} * \llbracket \tau \rrbracket_{\text{PD}} \\ \llbracket \alpha \rrbracket_{\text{PD}} &= \text{pd_hdr} * \alpha_{\text{PDb}} \\ \llbracket \lambda \alpha. \tau \rrbracket_{\text{PD}} &= \lambda \alpha_{\text{PDb}}. \llbracket \tau \rrbracket_{\text{PD}} \\ \llbracket \tau_1 \tau_2 \rrbracket_{\text{PD}} &= \llbracket \tau_1 \rrbracket_{\text{PD}} \llbracket \tau_2 \rrbracket_{\text{PDb}} \end{aligned}$$

$$\boxed{\llbracket \tau \rrbracket_{\text{PDb}} = \sigma}$$

$$\llbracket \tau \rrbracket_{\text{PDb}} = \sigma \text{ where } \llbracket \tau \rrbracket_{\text{PD}} \equiv \text{pd_hdr} * \sigma$$

FIGURE 3. Parse-descriptor type translation, selected constructs

$$\boxed{\llbracket \tau: \kappa \rrbracket_{\text{PT}} = \sigma}$$

$$\begin{aligned} \llbracket \tau: \text{T} \rrbracket_{\text{PT}} &= \text{bits} * \text{offset} \rightarrow \text{offset} * \llbracket \tau \rrbracket_{\text{rep}} * \llbracket \tau \rrbracket_{\text{PD}} \\ \llbracket \tau: \sigma \rightarrow \kappa \rrbracket_{\text{PT}} &= \sigma \rightarrow \llbracket \tau e: \kappa \rrbracket_{\text{PT}}, \text{ for any } e. \\ \llbracket \tau: \text{T} \rightarrow \kappa \rrbracket_{\text{PT}} &= \forall \alpha_{\text{rep}}. \forall \alpha_{\text{PDb}}. \llbracket \alpha: \text{T} \rrbracket_{\text{PT}} \rightarrow \llbracket \tau \alpha: \kappa \rrbracket_{\text{PT}} \quad (\alpha_{\text{rep}}, \alpha_{\text{PDb}} \notin \text{FTV}(\kappa) \cup \text{FTV}(\tau)) \end{aligned}$$

FIGURE 4. Host language types for parsing functions

an error code indicating the degree of success of the parse—success, success with errors, or failure—and the span of data (location in the source) described by the descriptor. To be precise, the type of the header (`pd_hdr`) is `int * errcode * span`. The body contains parse descriptors for the subcomponents of the representation. For types without subcomponents, we use `unit` as the body type. As with the representation types, dependency is uniformly erased.

Like other types, DDC^α type variables α are translated into a pair of header and a body. The body has abstract type α_{PDb} . This translation makes it possible for polymorphic parsing code to examine the header of a PD, even though it does not know the DDC^α type it is parsing. DDC^α abstractions are translated into F_ω type constructors that abstract the body of the PD (as opposed to the entire PD) and DDC^α applications are translated into F_ω type applications where the argument type is the PD body type.

$$\boxed{\llbracket \tau : \kappa \rrbracket_{\text{PPT}} = \sigma}$$

$$\begin{aligned} \llbracket \tau : \text{T} \rrbracket_{\text{PPT}} &= \llbracket \tau \rrbracket_{\text{rep}} * \llbracket \tau \rrbracket_{\text{PD}} \rightarrow \text{bits} \\ \llbracket \tau : \sigma \rightarrow \kappa \rrbracket_{\text{PPT}} &= \sigma \rightarrow \llbracket \tau e : \kappa \rrbracket_{\text{PPT}}, \text{ for any } e. \\ \llbracket \tau : \text{T} \rightarrow \kappa \rrbracket_{\text{PPT}} &= \forall \alpha_{\text{rep}}. \forall \alpha_{\text{PDb}}. \llbracket \alpha : \text{T} \rrbracket_{\text{PPT}} \rightarrow \llbracket \tau \alpha : \kappa \rrbracket_{\text{PPT}} \quad (\alpha_{\text{rep}}, \alpha_{\text{PDb}} \notin \text{FTV}(\kappa) \cup \text{FTV}(\tau)) \end{aligned}$$

FIGURE 5. Host language types for printing functions

3.2 DDC^α parsing semantics.

The parsing semantics of a type τ with kind T is a function that transforms some amount of input into a pair of a representation and a parse descriptor, the types of which are determined by τ . The parsing semantics for types with higher kind are functions that construct parsers, or functions that construct functions that construct parsers, *etc.* Figure 4 specifies the host-language types of the functions generated from well-kinded DDC^α types.

For each (unparameterized) type, the input to the corresponding parser is a bit string to parse and an offset at which to begin parsing. The output is a new offset, a representation of the parsed data, and a parse descriptor.

For any type, there are three steps to parsing: parse the subcomponents of the type (if any), assemble the resultant representation, and tabulate meta-data based on subcomponent meta-data (if any). For the sake of clarity, we have factored the latter two steps into separate representation and PD constructor functions which we define for each type. For example, the representation and PD constructors for the dependent sums are R_{Σ} and P_{Σ} , respectively. We have also factored out some commonly occurring code into auxiliary functions. These constructors and functions appear in the companion technical report [10].

The PD constructors determine the error code and calculate the error count. There are three possible error codes: `ok`, `err`, and `fail`, corresponding to the three possible results of a parse: it can succeed, parsing the data without errors; it can succeed, but discover errors in the process; or, it can find an unrecoverable error and fail. The error count is determined by subcomponent error counts and any errors associated directly with the type itself.

Figure 6 specifies the parsing semantics of a selected portion of DDC^α. We explain the interpretations of select types, from which the interpretation of the remaining types may be understood. The full semantics appears in the technical report [10]. A dependent sum parses the data according to the first type, binding the resulting representation and PD to x before parsing the remaining data according to the second type. It then bundles the results using the dependent sum constructor functions.

A type variable translates to an expression variable whose name corresponds directly to the name of the type variable. These expression variables are bound in

$$\boxed{\llbracket \tau \rrbracket_P = e}$$

$$\begin{aligned} \llbracket C(e) \rrbracket_P &= \lambda(B, \omega). \mathcal{B}_{\text{imp}}(C) (e) (B, \omega) \\ \llbracket \lambda x. \tau \rrbracket_P &= \lambda x. \llbracket \tau \rrbracket_P \\ \llbracket \tau e \rrbracket_P &= \llbracket \tau \rrbracket_P e \\ \llbracket \Sigma x: \tau. \tau' \rrbracket_P &= \\ &\lambda(B, \omega). \\ &\quad \text{let } (\omega', r, p) = \llbracket \tau \rrbracket_P (B, \omega) \text{ in} \\ &\quad \text{let } x = (r, p) \text{ in} \\ &\quad \text{let } (\omega'', r', p') = \llbracket \tau' \rrbracket_P (B, \omega') \text{ in} \\ &\quad (\omega'', R_\Sigma(r, r'), P_\Sigma(p, p')) \\ \llbracket \tau + \tau' \rrbracket_P &= \\ &\lambda(B, \omega). \\ &\quad \text{let } (\omega', r, p) = \llbracket \tau \rrbracket_P (B, \omega) \text{ in} \\ &\quad \text{if } \text{isOk}(p) \text{ then } (\omega', R_{+\text{left}}(r), P_{+\text{left}}(p)) \\ &\quad \text{else let } (\omega', r, p) = \llbracket \tau' \rrbracket_P (B, \omega) \text{ in} \\ &\quad (\omega', R_{+\text{right}}(r), P_{+\text{right}}(p)) \\ \llbracket \{x: \tau \mid e\} \rrbracket_P &= \\ &\lambda(B, \omega). \\ &\quad \text{let } (\omega', r, p) = \llbracket \tau \rrbracket_P (B, \omega) \text{ in} \\ &\quad \text{let } x = (r, p) \text{ in} \\ &\quad \text{let } c = e \text{ in} \\ &\quad (\omega', R_{\text{con}}(c, r), P_{\text{con}}(c, p)) \\ \llbracket \alpha \rrbracket_P &= \text{parse}_\alpha \\ \llbracket \lambda \alpha. \tau \rrbracket_P &= \Lambda \alpha_{\text{rep}}. \Lambda \alpha_{\text{pDb}}. \lambda \text{parse}_\alpha. \llbracket \tau \rrbracket_P \\ \llbracket \tau_1 \tau_2 \rrbracket_P &= \llbracket \tau_1 \rrbracket_P \llbracket \tau_2 \rrbracket_{\text{rep}} \llbracket \tau_2 \rrbracket_{\text{pDb}} \llbracket \tau_2 \rrbracket_P \end{aligned}$$

FIGURE 6. DDC^α parsing semantics, selected constructs

$$\boxed{\llbracket \tau \rrbracket_{PP} = e}$$

$$\begin{aligned} \llbracket C(e) \rrbracket_{PP} &= \lambda(r, \text{pd}). \mathcal{B}_{\text{pp}}(C) (e) (r, \text{pd}) \\ \llbracket \lambda x. \tau \rrbracket_{PP} &= \lambda x. \llbracket \tau \rrbracket_{PP} \\ \llbracket \tau e \rrbracket_{PP} &= \llbracket \tau \rrbracket_{PP} e \\ \llbracket \Sigma x: \tau_1. \tau_2 \rrbracket_{PP} &= \\ &\lambda(r, \text{pd}). \\ &\quad \text{let } x = (r.1, \text{pd}.2.1) \text{ in} \\ &\quad \text{let } \text{bs}_1 = \llbracket \tau_1 \rrbracket_{PP} x \text{ in} \\ &\quad \text{let } \text{bs}_2 = \llbracket \tau_2 \rrbracket_{PP} (r.2, \text{pd}.2.2) \text{ in} \\ &\quad \text{bs}_1 @ \text{bs}_2 \\ \llbracket \tau_1 + \tau_2 \rrbracket_{PP} &= \\ &\lambda(r, \text{pd}). \\ &\quad \text{case } (r, \text{pd}.2) \text{ of} \\ &\quad \mid (\text{inl } r_1, \text{inl } p_1) \Rightarrow \llbracket \tau_1 \rrbracket_{PP} (r_1, p_1) \\ &\quad \mid (\text{inr } r_2, \text{inr } p_2) \Rightarrow \llbracket \tau_2 \rrbracket_{PP} (r_2, p_2) \\ &\quad \mid _ \Rightarrow \text{badInput}() \\ \llbracket \{x: \tau \mid e\} \rrbracket_{PP} &= \\ &\lambda(r, \text{pd}). \\ &\quad \text{case } (r, \text{pd}.2) \text{ of} \\ &\quad \mid (\text{inl } r_1, p_1) \Rightarrow \llbracket \tau \rrbracket_{PP} (r_1, p_1) \\ &\quad \mid (\text{inr } r_2, p_2) \Rightarrow \llbracket \tau \rrbracket_{PP} (r_2, p_2) \\ \llbracket \alpha \rrbracket_{PP} &= \text{print}_\alpha \\ \llbracket \lambda \alpha. \tau \rrbracket_{PP} &= \Lambda \alpha_{\text{rep}}. \Lambda \alpha_{\text{pDb}}. \lambda \text{print}_\alpha. \llbracket \tau \rrbracket_{PP} \\ \llbracket \tau_1 \tau_2 \rrbracket_{PP} &= \llbracket \tau_1 \rrbracket_{PP} \llbracket \tau_2 \rrbracket_{\text{rep}} \llbracket \tau_2 \rrbracket_{\text{pDb}} \llbracket \tau_2 \rrbracket_{PP} \end{aligned}$$

FIGURE 7. DDC^α printing semantics, selected constructs

the interpretations of type abstractions. We interpret type abstraction as a function over other parsing functions. Because those parsing functions can have arbitrary DDC^α types (of kind \top), the interpretation must be a polymorphic function, parameterized by the representation and PD-body type of the DDC^α type parameter. For clarity, we present this type parameterization explicitly. Type application $\tau_1 \tau_2$ simply becomes the application of the interpretation of τ_1 to the representation-type, PD-type, and parsing interpretations of τ_2 .

3.3 DDC^α printing semantics

The definition of the printing semantics for a DDC^α description uses a similar set of concepts as the parsing semantics. To begin, the semantic function $\llbracket \tau: \kappa \rrbracket_{\text{PPT}} = \sigma$ gives the host language type σ for the printer generated from type τ with kind κ . As shown in Figure 5, the printing semantics for descriptions with higher kind are

functions that construct printers, while the printing semantics for descriptions with base kind are simple first-order functions that map a representation and a parse descriptor into a string of bits.

Figure 7 presents the printing semantics of selected DDC^α constructs. Base types $C(e)$ are printed in various ways according to the definition \mathcal{B}_{pp} , which is a parameter to the semantics. The base type printer \mathcal{B}_{pp} accepts the parse descriptor as a parameter, and in the case of an error, prints nothing. Dependent sums print one component and then the next in order. An ordinary sum prints the underlying tagged value. Notice that the structure of the parse descriptor and the representation should be isomorphic – both should be left injections or both should be right injections. Any pair of structures generated by the parser are guaranteed to satisfy this invariant. If the pair do not match, then the programmer is using the printer incorrectly. In this case, the printer calls an unspecified error routine named `badInput()`.

The semantics of printing parameterized types follows similar lines to the semantics of parsing this construct. In particular, whenever a type parameter is introduced in the syntax of a description, a corresponding value parameter with printer function type is introduced in the generated printer code. We give the value parameter the name $print_\alpha$.

4 METATHEORY

To validate our semantic definitions, we have proven two key metatheoretic results, one concerning *type correctness* and the other concerning *canonical forms*.

4.1 Type Correctness.

Demonstrating that generated parsers and printers are well formed and have the expected types is nontrivial primarily because the generated code expects parse descriptors to have a particular shape, and it is not completely obvious they do in the presence of polymorphism. We discuss these and other issues in more detail in the companion technical report [10]; below we present the final theorem on type correctness.

Theorem 1 (Type Correctness of Closed Types)

- If $\vdash \tau : \kappa$ then $\vdash \llbracket \tau \rrbracket_P : \llbracket \tau : \kappa \rrbracket_{PT}$.
- If $\vdash \tau : \kappa$ then $\vdash \llbracket \tau \rrbracket_{PP} : \llbracket \tau : \kappa \rrbracket_{PPT}$.

4.2 Canonical Forms for Parsed Data.

DDC^α parsers generate pairs of representations and parse descriptors designed to satisfy a number of invariants. Of greatest importance is the fact that when the parse descriptor says there are no errors in a particular substructure, the programmer can count on the representation satisfying all of the syntactic and semantic

Normal Types	$v ::= C(e) \mid \lambda x.\tau \mid \Sigma x:\tau.\tau \mid \tau + \tau \mid \{x:\tau \mid e\} \mid \lambda\alpha.\tau$
Types	$\tau ::= v \mid \tau e \mid \tau\tau \mid \alpha$

FIGURE 8. DDC^α normal types, selected constructs

constraints expressed by the DDC^α type description. When a parse descriptor and representation satisfy these invariants, we say the pair of data structures is in *canonical form*. While generated parsers produce canonical outputs, generated printers expect canonical inputs.

For each DDC^α type, its canonical forms are defined via two (mutually recursive) relations. The first relation, $\text{Canon}_v(r, p)$, defines the canonical form of a representation r and a parse descriptor p at normal type v . *Normal types* are those closed types with base kind \top that are defined in Figure 8. Types with higher kind such as abstractions are not described by this relation as they cannot directly produce representations and PDs. Another relation, $\text{Canon}^*_\tau(r, p)$ (formal definition omitted) normalizes τ , thereby eliminating outermost type and value applications. For brevity, we write $p.h.nerr$ as $p.nerr$ and use pos to denote the function that returns zero when passed zero and one when passed another natural number.

Definition 2 (Canonical Forms (selected constructs))

$\text{Canon}_v(r, p)$ iff exactly one of the following is true:

- $v = C(e)$ and $r = \text{inl } c$ and $p.nerr = 0$.
- $v = C(e)$ and $r = \text{inr none}$ and $p.nerr = 1$.
- $v = \Sigma x:\tau_1.\tau_2$ and $r = (r_1, r_2)$ and $p = (h, (p_1, p_2))$ and $h.nerr = \text{pos}(p_1.nerr) + \text{pos}(p_2.nerr)$, $\text{Canon}^*_{\tau_1}(r_1, p_1)$ and $\text{Canon}^*_{\tau_2[(r,p)/x]}(r_2, p_2)$.
- $v = \tau_1 + \tau_2$ and $r = \text{inl } r'$ and $p = (h, \text{inl } p')$ and $h.nerr = \text{pos}(p'.nerr)$ and $\text{Canon}^*_{\tau_1}(r', p')$.
- $v = \tau_1 + \tau_2$ and $r = \text{inr } r'$ and $p = (h, \text{inr } p')$ and $h.nerr = \text{pos}(p'.nerr)$ and $\text{Canon}^*_{\tau_2}(r', p')$.
- $v = \{x:\tau' \mid e\}$, $r = \text{inl } r'$ and $p = (h, p')$, and $h.nerr = \text{pos}(p'.nerr)$ and $\text{Canon}^*_{\tau'}(r', p')$ and $e[(r', p')/x] \rightarrow^* \text{true}$.
- $v = \{x:\tau' \mid e\}$, $r = \text{inr } r'$ and $p = (h, p')$, and $h.nerr = 1 + \text{pos}(p'.nerr)$ and $\text{Canon}^*_{\tau'}(r', p')$ and $e[(r', p')/x] \rightarrow^* \text{false}$.

The first part of Theorem 3 states that parsers for well-formed types (of base kind) produce a canonical pair of representation and parse descriptor if they produce anything at all. Conversely, the second part states that, given a canonical representation and parse descriptor, the printer for well-formed types (of base kind) will not “go wrong” by calling the `badInput()` function.

Theorem 3 (Parsing to/Printing from Canonical Forms)

- If $\vdash \tau : \mathbb{T}$ and $\llbracket \tau \rrbracket_P (B, \omega) \rightarrow^* (\omega', r, p)$ then $\text{Canon}^*_\tau(r, p)$.
- If $\vdash \tau : \mathbb{T}$, $\text{Canon}^*_\tau(r, p)$ and $\llbracket \tau \rrbracket_{PP} (r, p) \rightarrow^* e$ then $e \neq \text{badInput}()$.

5 RELATED WORK

There are other formalisms for defining parsers, most famously, regular expressions and context-free grammars. In terms of recognition power, our type theory contains dependency and hence easily express languages that are not context-free. Perhaps more importantly though, unlike standard theories of context-free grammars, we do not treat our type theory merely as a recognizer for a collection of strings. Our type-based descriptions define external data formats, rich invariants on the internal parsed data structures and parser and printer functions to relate them. This multi-part interpretation of types lies at the heart of tools such as PADS.

As mentioned in the introduction, there are many domain-specific specification languages that allow users to write down descriptions of data and generate tools from them [1, 8, 12, 3, 11, 2]. Many of these projects contain great ideas, but we would like to highlight just two of them: DEMETER [8] and XSUGAR [2]. DEMETER generates a number of different kinds of visitor patterns, including visitors for printing, from high-level type-based descriptions called *class dictionaries*. XSUGAR is a specialized tool for converting back and forth from ad hoc data to XML. A static analysis guarantees that XSUGAR’s generated transforms are inverses of one another, something we have not yet proven for DDC^α . On the other hand, neither DEMETER nor XSUGAR supports dependent types or polymorphism.

There is also a clear relation between our work and the parser/printer combinator libraries developed for many functional programming languages. One difference between the DDC^α and, for instance, Haskell combinator libraries [7, 14], is that a single DDC^α description specifies multiple programs (*i.e.*, a parser and a printer) whereas a program built from Haskell combinators will generally only specify one thing – either a parser or a printer.

The DDC^α and its semantics are also closely tied to *type-directed* programming techniques (See, for example, Hinze’s work [6]). Indeed, the DDC^α parsers and printers really are just specially crafted (dependently) type-directed programs.

6 CONCLUSIONS

The DDC^α , a dependent data description calculus, now has both parsing and printing semantics. Moreover, these semantics have been proven to satisfy important type correctness and canonical forms theorems. In the future, we hope to provide semantics for other tools generated by data description languages in general, and the PADS family of languages in particular. The ease with which we were able to augment our existing semantic framework with a printing semantics suggests this goal is within our reach.

Acknowledgments: This material is based upon work supported by the National Science Foundation under Grant Nos. 0615062 and 0612147. 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 NSF.

REFERENCES

- [1] G. Back. DataScript - A specification and scripting language for binary data. In *Generative Programming and Component Engineering*, volume 2487, pages 66–77. Lecture Notes in Computer Science, 2002.
- [2] C. Brabrand, A. Møller, and M. I. Schwartzbach. Dual syntax for XML languages. In *Proc. 10th International Workshop on Database Programming Languages, DBPL '05*, volume 3774 of LNCS, pages 27–41. Springer-Verlag, August 2005.
- [3] K. Fisher and R. Gruber. PADS: A domain specific language for processing ad hoc data. In *ACM Conference on Programming Language Design and Implementation*, pages 295–304. ACM Press, June 2005.
- [4] K. Fisher, Y. Mandelbaum, and D. Walker. The next 700 data description languages. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 2 – 15, Jan. 2006.
- [5] K. Fisher, Y. Mandelbaum, and D. Walker. The next 700 data description languages. *Submitted for journal publication.*, Dec. 2006. See www.padsproj.org/doc.html.
- [6] R. Hinze. A new approach to generic functional programming. In *Proceedings of the 27th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 119–132, New York, NY, USA, 2000. ACM Press.
- [7] G. Hutton and E. Meijer. Monadic parsing in Haskell. *Journal of Functional Programming*, 8(4):437–444, July 1998.
- [8] K. Lieberherr. Object-oriented programming with class dictionaries. *Lisp and Symbolic Computation*, 1:185–212, 1988.
- [9] Y. Mandelbaum. *The Theory and Practice of Data Description*. PhD thesis, Princeton University, Sept. 2006.
- [10] Y. Mandelbaum, K. Fisher, D. Walker, M. Fernandez, and A. Gleyzer. PADS/ML: A functional data description language. Technical Report TR-761-06, Princeton University, July 2006.
- [11] Y. Mandelbaum, K. Fisher, D. Walker, M. Fernandez, and A. Gleyzer. PADS/ML: A functional data description language. In *ACM Symposium on Principles of Programming Languages*. ACM Press, Jan. 2007.
- [12] P. McCann and S. Chandra. PacketTypes: Abstract specificationa of network protocol messages. In *ACM Conference of Special Interest Group on Data Communications*, pages 321–333. ACM Press, August 2000.
- [13] V. Paxson. A system for detecting network intruders in real-time. In *Computer Networks*, Dec. 1999.
- [14] P. Wadler. A prettier printer. *The Fun of Programming, Cornerstones of Computing*, pages 223–244, 2003.