

The Abella Interactive Theorem Prover (System Description)

Andrew Gacek

Department of Computer Science and Engineering, University of Minnesota
200 Union Street SE, Minneapolis, MN 55455, USA

1 Introduction

Abella [3] is an interactive system for reasoning about aspects of object languages that have been formally presented through recursive rules based on syntactic structure. Abella utilizes a two-level logic approach to specification and reasoning. One level is defined by a specification logic which supports a transparent encoding of structural semantics rules and also enables their execution. The second level, called the reasoning logic, embeds the specification logic and allows the development of proofs of properties about specifications. An important characteristic of both logics is that they exploit the λ -tree syntax approach to treating binding in object languages. Amongst other things, Abella has been used to prove normalizability properties of the λ -calculus, cut admissibility for a sequent calculus and type uniqueness and subject reduction properties. This paper discusses the logical foundations of Abella, outlines the style of theorem proving that it supports and finally describes some of its recent applications.

2 The Logic Underlying Abella

Abella is based on \mathcal{G} , an intuitionistic, predicative, higher-order logic with fixed-point definitions for atomic predicates and with natural number induction [4].

Representing binding. \mathcal{G} uses the λ -tree syntax approach to representing syntactic structures [7], which allows object level binding to be represented using meta-level abstraction. Thus common notions related to binding such as α -equivalence and capture-avoiding substitution are built into the logic, and the encodings of object languages do not need to implement such features.

To reason over λ -tree syntax, \mathcal{G} uses the ∇ quantifier which represents a notion of generic judgment [9]. A formula $\nabla x.F$ is true if F is true for each x in a generic way, *i.e.*, when nothing is assumed about any x . This is a stronger statement than $\forall x.F$ which says that F is true for all values for x but allows this to be shown in different ways for different values.

For the logic \mathcal{G} , we assume the following two properties of ∇ :

$$\nabla x.\nabla y.F \ x \ y \equiv \nabla y.\nabla x.F \ x \ y \quad \nabla x.F \equiv F \text{ if } x \text{ not free in } F$$

A natural proof-theoretic treatment for this quantifier is to use *nominal constants* to instantiate ∇ -bound variables [16]. Specifically, the proof rules for ∇ are

$$\frac{\Gamma, B[a/x] \vdash C}{\Gamma, \nabla x. B \vdash C} \nabla\mathcal{L} \qquad \frac{\Gamma \vdash C[a/x]}{\Gamma \vdash \nabla x. C} \nabla\mathcal{R}$$

where a is a nominal constant which does not appear in the formula underneath the ∇ quantifier. Due to the equivalence of permuting ∇ quantifiers, nominal constants must be treated as permutable, which is captured by the initial rule.

$$\frac{\pi. B = B'}{\Gamma, B \vdash B'} id_\pi$$

Here π is a permutation of nominal constants.

Definitions. The logic \mathcal{G} supports fixed-point definitions of atomic predicates. These definitions are specified as clauses of the form $\forall \bar{x}. (\nabla \bar{z}. H) \triangleq B$ where the head H is an atomic predicate. This notion of definition is extended from previous notions (*e.g.*, see [9]) by admitting the ∇ -quantifier in the head. Roughly, when such a definition is used, in ways to be explained soon, these ∇ -quantified variables become instantiated with nominal constants from the term on which the definition is used. The instantiations for the universal variables \bar{x} may contain any nominal constants not assigned to the variables \bar{z} . Thus ∇ quantification in the head of a definition allows us to restrict certain pieces of syntax to be nominal constants and to state dependency information for those nominal constants.

Two examples hint at the expressiveness of our extended form of definitions. First, we can define a predicate *name* E which holds only when E is a nominal constant. Second, we can define a predicate *fresh* $X E$ which holds only when X is a nominal constant which does not occur in E .

$$(\nabla x. \text{name } x) \triangleq \top \qquad \forall E. (\nabla x. \text{fresh } x E) \triangleq \top$$

Note that the order of quantification in *fresh* enforces the freshness condition.

Definitions can be used in both a positive and negative fashion. Positively, definitions are used to derive an atomic judgment, *i.e.*, to show a predicate holds on particular values. This use corresponds to unfolding a definition and is similar to back-chaining. Negatively, an atomic judgment can be decomposed in a case analysis-like way based on a closed-world reading of definitions. In this case, the atomic judgment is unified with the head of each definitional clause, where eigenvariables are treated as instantiatable. Also, both the positive and negative uses of definitions consider permutations of nominal constants in order to allow the ∇ -bound variables \bar{z} to range over any nominal constants. A precise presentation of these rules, which is provided in Gacek *et al.* [4], essentially amounts to introduction rules for atomic judgments on the right and left sides of sequents in a sequent calculus based presentation of the logic.

Induction. \mathcal{G} supports induction over natural numbers. By augmenting the predicates being defined with a natural number argument, this induction can serve as a method of proof based on the length of a bottom-up evaluation of a definition.

$$\begin{aligned} & \forall m, n, a, b [of\ m\ (arr\ a\ b) \wedge of\ n\ a \supset of\ (app\ m\ n)\ b] \\ & \forall r, a, b [\forall x [of\ x\ a \supset of\ (r\ x)\ b] \supset of\ (abs\ a\ r)\ (arr\ a\ b)] \end{aligned}$$

Fig. 1. Second-order hereditary Harrop formulas for typing

3 The Structure of Abella

The architecture of Abella has two distinguishing characteristics. First, Abella is oriented towards the use of a specific (executable) specification logic whose proof-theoretic structure is encoded via definitions in \mathcal{G} . Second, Abella provides tactics for proof construction that embody special knowledge of the specification logic. We discuss these aspects and their impact in more detail below.

3.1 Specification Logic

It is possible to encode object language descriptions directly in definitions in \mathcal{G} , but there are two disadvantages to doing so: the resulting definitions may not be executable and there are common patterns in specifications with λ -tree syntax which we would like to take advantage of. We address these issues by selecting a specification logic which has the features that the \mathcal{G} lacks, and embedding the evaluation rules of this specification logic instead into \mathcal{G} . Object languages are then encoded through descriptions in the specification logic [6].

The specification logic of Abella is second-order hereditary Harrop formulas [8] with support for λ -tree syntax. This allows a transparent encoding of structural operational semantics rules which operate on objects with binding. For example, consider the simply-typed λ -calculus where types are either a base type i or arrow types constructed with arr . Terms are encoded with the constructors app and abs . The constructor abs takes two arguments: the type of the variable being abstracted and the body of the function. Rather than having a constructor for variables, the body argument to abs is an abstraction in our specification logic, thus object level binding is represented by the specification logic binding. For example, the term $(\lambda f : i \rightarrow i. (\lambda x : i. (f\ x)))$ is encoded as

$$abs\ (arr\ i\ i)\ (\lambda f. abs\ i\ (\lambda x. app\ f\ x)).$$

In the latter term, λ denotes an abstraction in the specification logic. Given this representation, the typing judgment $of\ m\ t$ is defined in Figure 1. Note that these rules do not maintain an explicit context for typing assumptions, instead using a hypothetical judgment to represent assumptions. Also, there is no side-condition in the rule for typing abstractions to ensure the variable x does not yet occur in the typing context, since instead of using a particular x for recording a typing assumption, we quantify over all x .

Our specification of typing assignment is executable. More generally, the Abella specification logic is a subset of the language λ Prolog [11] which can be compiled and executed efficiently [12]. This enables the animation of specifications, which is convenient for assessing specifications before attempting to

prove properties over them. This also allows specifications to be used as testing oracles when developing full implementations.

The evaluation rules of our specification logic can be encoded as a definition in \mathcal{G} . A particular specification is then encoded in a separate definition which is used by the definition of evaluation in order to realize back-chaining over specification clauses. Reasoning over a specification is realized by reasoning over its evaluation via the definition of the specification logic. Abella takes this further and is customized towards the specification logic. For example, the context of hypothetical judgments in our specification logic admits weakening, contraction, and permutation, all of which are provable in \mathcal{G} . Abella automatically uses this meta-level property of the specification logic during reasoning. Details on the benefits of this approach to reasoning are available in Gacek *et al.* [5].

3.2 Tactics

The user constructs proofs in Abella by applying tactics which correspond to high-level reasoning steps. The collection of tactics can be grouped into those that generically orchestrate the rules of \mathcal{G} and those that correspond to meta-properties of the specification logic. We discuss these classes in more detail below.

Generic tactics. The majority of tactics in Abella correspond directly to inference rules in \mathcal{G} . The most common tactics from this group are the ones which perform induction, introduce variables and hypotheses, conduct case analysis, apply lemmas, and build results from hypotheses. In the examples suite distributed with Abella, these five tactics make up more than 90% of all tactic usages. The remaining generic tactics are for tasks such as splitting a goal of the form $G_1 \wedge G_2$ into two separate goals for G_1 and G_2 , or for instantiating the quantifier in a goal of the form $\exists x.G$.

Specification logic tactics. Since our specification logic is encoded in \mathcal{G} , we can formally prove meta-level properties for it. Once such properties are proved, their use in proofs can be built into tactics. Two important properties that Abella uses in this way are instantiation and cut admissibility. In detail, negative uses of the specification logic \forall quantifier are represented in \mathcal{G} as nominal constants (*i.e.*, the ∇ quantifier), and the instantiation tactic allows such nominal constants to be instantiated with specific terms. The cut tactic allows hypothetical judgments to be relieved by showing that they are themselves provable.

4 Implementation

Abella is implemented in OCaml. The most sophisticated component of this implementation is higher-order unification which is a fundamental part of the logic \mathcal{G} . It underlies how case analysis is performed, and in the implementation, unification is used to decide when tactics apply and to determine their result. Thus an efficient implementation of higher-order unification is central to an efficient prover. For this, Abella uses the the higher-order pattern unification

package of Nadathur and Linnell [10]. We have also extended this package to deal with the particular features and consequences of reasoning in \mathcal{G} .

Treatment of nominal constants. As their name suggests, nominal constants can be treated very similarly to constants for most of the unification algorithm, but there are two key differences. First, while traditional constants can appear in the instantiation of variables, nominal constants cannot appear in the instantiation of variables. Thus dependency information on nominal constants is tracked via explicit raising of variables. Second, nominal constants can be permuted when determining unifiability. However, even in our most sophisticated examples the number of nominal constants appearing at the same time has been at most two. Thus, naive approaches to handling permutability of nominal constants have sufficed and there has been little need to develop sophisticated algorithms.

Simple extensions. The treatment of case analysis via unification for eigenvariables creates unification problems which fall outside of the higher-order pattern unification fragment, yet still have most general unifiers. For example, consider the clause for β -contraction in the λ -calculus:

$$\text{step } (\text{app } (\text{abs } R) M) (R M).$$

Case analysis on a hypotheses of the form $\text{step } A B$ will result in the attempt to solve the unification problem $B = R M$ where B , R , and M are all instantiatable. This is outside of the higher-order pattern unification fragment since R is applied to an instantiatable variable, but there is a clear most general unifier. When nominal constants are present, this situation is slightly more complicated with unification problems such as $B x = R M x$ or $B x = R (M x)$, where x is a nominal constant. The result is the same, however, that a most general unifier exists and is easy to find.

5 Examples

This section briefly describes sample reasoning tasks we have conducted in Abella. The detailed proofs are available in the distribution of Abella [3].

Results from the λ -calculus. Over untyped λ -terms, we have shown the equivalence of big-step and small-step evaluation, preservation of typing for both forms of evaluation, and determinacy for both forms of evaluation. We have shown that the λ -terms can be disjointly partitioned into normal and non-normal forms. Over simply-typed λ -terms, we have shown that typing assignments are unique.

Cut admissibility. We have shown that the cut rule is admissible for a sequent calculus with implication and conjunction. The representation of sequents in our specification logic used hypothetical judgments to represent hypotheses in the sequent. This allowed the cut admissibility proof to take advantage of Abella's built-in treatment of meta-properties of the specification logic.

The POPLmark challenge. The POPLmark challenge [1] is a selection of problems which highlight the traditional difficulties in reasoning over systems which manipulate objects with binding. The particular tasks of the challenge involve reasoning about evaluation, typing, and subtyping for $F_{<}$, a λ -calculus with bounded subtype polymorphism. We have solved parts 1a and 2a of this challenge using Abella, which represent the fundamental reasoning tasks involving objects with binding.

Proving normalizability à la Tait. We have shown that all closed terms in the call-by-value, simply-typed λ -calculus are normalizable using the logical relations argument in the style of Tait [14]. Fundamental in this proof was the encoding of arbitrary cascading substitutions which allows one to consider all closed instantiations for an open λ -term. Encoding and reasoning over this form of substitution makes essential use of the extended form of definitions in \mathcal{G} .

6 Future and Related Work

Induction and coinduction. The logic \mathcal{G} currently supports induction on natural numbers. Similar logics have been extended to support structural induction and coinduction on definitions [15]. Already, the implementation of Abella has support for these features. A paper which describes the extended logic supporting these features is in preparation.

User programmability. Tactics-based theorem provers often support *tacticals* which allow users to compose tactics in useful ways. Some systems even go beyond this and offer a full programming language for creating custom tactics. We would like to extend Abella with such features.

Proof search. Many proofs in Abella follow a straightforward pattern of essentially induction, case analysis, and building from hypotheses. We would like to extend Abella to perform these types of proofs automatically. Recent results on focusing in similar logics may offer some insight into a disciplined approach to automated proof search [2].

Related work. A closely related system is Twelf [13] which is based on a dependently typed λ -calculus for specification. Controlling for dependent types, the most significant difference is that our meta-logic is significantly richer than the one in Twelf. Also related is the Nominal package [17] for Isabelle/HOL which allows for reasoning over α -equivalence classes. This approach leverages on existing theorem proving work, but does not address the full problem of reasoning with binding. In particular, all work related to substitution is left to the user. A more detailed comparison with these works is available in Gacek *et al.* [5].

Acknowledgements

I am grateful to David Baelde, Dale Miller, Gopalan Nadathur, Randy Pollack, and Alwen Tiu for their input and feedback on the development of Abella.

Anonymous reviewers provided helpful comments on an earlier version of this paper. This work has been supported by the NSF Grant CCR-0429572 and by a grant from Boston Scientific. Opinions, findings, and conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Aydemir, B.E., Bohannon, A., Fairbairn, M., Foster, J.N., Pierce, B.C., Sewell, P., Vytiniotis, D., Washburn, G., Weirich, S., Zdancewic, S.: Mechanized metatheory for the masses: The PoplMark challenge. In: Hurd, J., Melham, T. (eds.) TPHOLs 2005. LNCS, vol. 3603, pp. 50–65. Springer, Heidelberg (2005)
2. Baelde, D., Miller, D.: Least and greatest fixed points in linear logic. In: Dershowitz, N., Voronkov, A. (eds.) LPAR 2007. LNCS (LNAI), vol. 4790, pp. 92–106. Springer, Heidelberg (2007)
3. Gacek, A.: The Abella system. Available in source code (2008), <http://abella.cs.umn.edu/>
4. Gacek, A., Miller, D., Nadathur, G.: Combining generic judgments with recursive definitions. In: Pfenning, F. (ed.) 23th Symp. on Logic in Computer Science. IEEE Computer Society Press, Los Alamitos (to appear, 2008)
5. Gacek, A., Miller, D., Nadathur, G.: Reasoning in Abella about structural operational semantics specifications. In: LFMTTP 2008 (to appear, 2008), <http://arxiv.org/abs/0804.3914>
6. McDowell, R., Miller, D.: Reasoning with higher-order abstract syntax in a logical framework. *ACM Trans. on Computational Logic* 3(1), 80–136 (2002)
7. Miller, D.: Abstract syntax for variable binders: An overview. In: Palamidessi, C., Moniz Pereira, L., Lloyd, J.W., Dahl, V., Furbach, U., Kerber, M., Lau, K.-K., Sagiv, Y., Stuckey, P.J. (eds.) CL 2000. LNCS (LNAI), vol. 1861, pp. 239–253. Springer, Heidelberg (2000)
8. Miller, D., Nadathur, G., Pfenning, F., Scedrov, A.: Uniform proofs as a foundation for logic programming. *Annals of Pure and Applied Logic* 51, 125–157 (1991)
9. Miller, D., Tiu, A.: A proof theory for generic judgments. *ACM Trans. on Computational Logic* 6(4), 749–783 (2005)
10. Nadathur, G., Linnell, N.: Practical higher-order pattern unification with on-the-fly raising. In: Gabbriellini, M., Gupta, G. (eds.) ICLP 2005. LNCS, vol. 3668, pp. 371–386. Springer, Heidelberg (2005)
11. Nadathur, G., Miller, D.: An Overview of λ Prolog. In: Fifth International Logic Programming Conference, Seattle, August 1988, pp. 810–827. MIT Press, Cambridge (1988)
12. Nadathur, G., Mitchell, D.J.: System description: Teyjus—A compiler and abstract machine based implementation of Lambda Prolog. In: Ganzinger, H. (ed.) Proceedings of the 16th International Conference on Automated Deduction, Trento, Italy, July 1999, pp. 287–291. Springer, Heidelberg (1999)
13. Pfenning, F., Schürmann, C.: System description: Twelf — A meta-logical framework for deductive systems. In: Ganzinger, H. (ed.) 16th Conference on Automated Deduction. LNCS (LNAI), pp. 202–206. Springer, Heidelberg (1999)
14. Tait, W.W.: Intensional interpretations of functionals of finite type I. *J. of Symbolic Logic* 32(2), 198–212 (1967)

15. Tiu, A.: A Logical Framework for Reasoning about Logical Specifications. PhD thesis, Pennsylvania State University (May 2004)
16. Tiu, A.: A logic for reasoning about generic judgments. In: Momigliano, A., Pientka, B. (eds.) International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice (LFMTP 2006) (2006)
17. Urban, C., Tasson, C.: Nominal techniques in Isabelle/HOL. In: Nieuwenhuis, R. (ed.) CADE 2005. LNCS (LNAI), vol. 3632, pp. 38–53. Springer, Heidelberg (2005)