Artifact Report: Trocq: Proof Transfer for Free, With or Without Univalence

Cyril Cohen[1], Enzo Crance[2,3], and Assia Mahboubi[2,4(✉)]

[1] Université Côte d'Azur, Inria, Biot, France
Cyril.Cohen@inria.fr
[2] Nantes Université, École Centrale Nantes, CNRS, INRIA, LS2N, UMR 6004, Nantes, France
{Enzo.Crance,Assia.Mahboubi}@inria.fr
[3] Mitsubishi Electric R&D Centre Europe, Rennes, France
[4] Vrije Universiteit Amsterdam, Amsterdam, The Netherlands

Trocq [5] is both the name of a calculus, describing a parametricity framework, and of a Coq plugin [6] that provides tactics for performing representation changes in goals, as well as vernacular commands for specifying the expected translations. More precisely, from an initial goal of type `G`, the `trocq` tactic simultaneously computes using the Trocq calculus [5] a translation `G'` and a justification `w : G' -> G`. If successful, the user is thus left with proving `G'`.

The plugin orchestrates this double synthesis, by assembling existing building blocks known to the tactic, in the course of a linear traversal of the input term `G`. These building blocks are of two natures. First, the actual rules of the parametricity framework [5] govern the synthesis rule attached to each term construction of CC_ω. The other nature of building blocks is the collection of registered pairs of user-defined constants. These pairs come equipped with a witness of their relatedness at some level, a data registered via the `Trocq Use` command. When the linear traversal of the input term hits a constant, it queries the database registering these user-defined relations, looking for the corresponding constant and witness to be used in the synthesis.

The Trocq plugin is implemented in Elpi [9]: a dialect of λProlog which can be used as a meta-language for Coq, through the Coq-Elpi [20] plugin. The latter encodes Coq terms in higher-order abstract syntax (HOAS) which provides native support for bound variables, complemented by a comprehensive API (typechecking, elaboration, interacting with the global environment, etc).

1 Example

Let us translate the induction principle associated with type `nat`, the unary representation of ℕ, to type `N`, the binary one. Types `nat` and `N` are equivalent and we use the `Trocq Use` command to register such pairs of related types:

```
Definition RN : (N <=> nat)%P := Iso.toParamSym N.of_nat_iso.
Trocq Use RN. (* registering a pair of related types *)
```

Proof `RN` coerces to a relation of type `N -> nat -> Type`, and we also register proofs that it relates the respective zero and successor constants of these types:

S. Weirich (Ed.): ESOP 2024, LNCS 14576, pp. 269–274, 2024.
https://doi.org/10.1007/978-3-031-57262-3_11

```
Definition RN0 : RN 0%N 0%nat. Proof. done. Qed.
Definition RNS m n : RN m n -> RN (N.succ m) (S n). Proof. by case. Qed.
Trocq Use RN0. Trocq Use RNS. (* registering related constants *)
```

We can now use tactic `trocq` to prove a useful induction principle on type `N`:

```
Lemma N_Srec : forall (P : N -> Type), P 0%N ->
    (forall n, P n -> P (N.succ n)) -> forall n, P n.
Proof. trocq. (* replaces N by nat in the goal *) exact nat_rect. Qed.
```

Inspecting the proof term actually reveals that univalence was not needed in the proof of `N_Srec`. The `example` directory of the artifact provides more examples, for weaker relations than equivalences, and beyond representation independence.

2 Architecture of the plugin

A Trocq parametricity sequent $\Delta \vdash_t M \ @ \ A \sim M' \ \because \ M_R$ expresses that *terms M and M' are related at type A with witness M_R* in context Δ. Unlike standard, unequivocal parametricity translations, each construct of CC_ω gives rise to a *family* of possible synthesis rules, indexed by annotations on M and A.

Encoding CC_ω^+ . To implement the annotation calculus CC_ω^+ , we just annotate Coq's sort `Type` with a pair (n, m) using *convertible* synonyms `(PType n m)`, where `PType := fun (_ _ : label) => Type`. The two thrown-away arguments code for the annotation. In the course of the synthesis, arguments of certain occurrences of `PType` are left as holes and filled by a constraint solving algorithm.

Synthesis. The logic programming paradigm on which Elpi is based, is ideal to implement algorithms expressed as inference rules, as each rule can be associated to an instance of a predicate. The linear traversal of the input term at the core of the Trocq plugin is operated by the predicate `param`, of arity 4, where `param X T X' XR` stands for the parametricity sequent $\Delta \vdash_t x \ @ \ T \sim x' \ \because \ x_R$ for a certain context Δ. In this sequent, x and T are input values (initially, the source goal and the annotated sort $\square^{(0,1)}$), and the synthesized term x' and witness x_R are outputs. Each construct of CC_ω leads to *one* instance of the predicate. As an example, let us inspect the instance of the `param` predicate for dependent products, which implements the rule TrocqPi of the Trocq calculus. For the sake of readability, we removed lines related to logs, pretty-printing, and fresh universe instance generation. The head of the predicate is:

```
param (prod N A B) (app [pglobal (const PType) _, M1, M2]) Prod' ProdR :-
  param.db.ptype PType, !,
  cstr.univ-link C M1 M2,
```

which matches an input term $\Pi x : A. B$ and our Coq encoding of its annotated type $\square^{(M_1, M_2)}$. Then, following the hypotheses in the inference rule, the predicate

computes the prescribed annotation $(C_A, C_B) = \mathcal{D}_{\Pi}(C)$, and does two recursive calls on A and B with classes C_A and C_B:

```
cstr.dep-pi C CA CB,
cstr.univ-link CA M1A M2A,
param A (app [pglobal (const PType) _, M1A, M2A]) A' AR,
cstr.univ-link CB M1B M2B,
TB = app [pglobal (const PType) _, M1B, M2B],
@annot-pi-decl N A a\ pi a' aR\ param.store a A a' aR =>
  param (B a) TB (B' a') (BR a a' aR),
```

The last step (omitted in the above snippet) is to build the output proof $p^C_{\Pi}\ A_R\ B_R$. As the axioms (univalence, functional extensionality) that might be involved in some proofs are not assumed globally, they are used as an additional argument albeit only in the building blocks that require them. Therefore, we check whether the requested rule requires the addition of an axiom to the list of arguments (in the case of the dependent product, function extensionality). If this axiom is not present in the context at the time of calling this part of the code, the tactic rightfully fails, because the translation is impossible.

Exploiting symmetries. Trocq provides several distinct rules per language construct (such as Π) and per relation structure among the 36 items in the hierarchy: for a same construct, these rules differ by the annotations required on the input of the rule, and by the structure of the relation relating the input term and the synthesized one. For each such rule, a Coq function provides the corresponding rule building block. Making the most of symmetries, the 495 rule building blocks are generated by meta-programming from only 9 manually defined ones.

Handling of constants. Finally, the traversal of the input term collects *constraints* on the annotations, as multiple valid solutions might exists: for instance, an implication might be obtained from weakening an equivalence. The algorithm strives to minimize the requirements on the user-defined building blocks, which also amounts to minimizing the dependency on axioms. This inference procedure is formalized as a finite domain constraint solving problem, and implemented using *Constraint Handling Rules* (CHR) language [10], as available in Elpi.

3 Related work

In the context of type theory, Barthe and Pons [3] already noticed that the computational content of type isomorphisms can serve proof transfer. The first implementation report of a tool based on this idea appeared soon after [16]. Implemented in a meta-language and based on proof rewriting, this heuristic translation produced a candidate proof term from an existing proof term, with no formal guarantee, not even that of being well-typed. Generalized rewriting [17], which generalizes setoid rewriting to preorders, is also a variant of proof transfer, albeit within the same type. As such, it allows in particular rewriting under

binders. The restriction to homogeneous relations however excludes more general instances of proof transfer, *e.g.*, , datatype representation change and quasi-PERs (QPER, or zig-zag complete relations) [13], essentially heterogeneous.

The other proof transfer methods we are aware of all address the case of heterogeneous relations. Incidentally, they can thus also be used for the homogeneous case, and thus for generalized rewriting, although this special case is seldom emphasized. The Coq Effective Algebra Library (CoqEAL) [8,7] and the Isabelle/HOL transfer package [14,11,12,15], pioneered the use of parametricity-based methods for proof transfer, motivated by the refinement of proof-oriented data-structures to computation-oriented counterparts. Together with a subsequent generalization of the CoqEAL approach [21], these tools address the case of a transfer between a subtype of a certain type A and a quotient of a certain type B, *i.e.*, the case of a trivial QPER in which the zig-zag morphism is a surjection from A to B.

Modern approaches to proof transfer rely on univalence, either as an axiom, in the case of univalent parametricity [19] or as a computing primitive [2]. Key ingredients of univalent parametricity were already present in earlier seemingly unpublished work [1], implemented using an ancestor of the MetaCoq library [18].

The columns of Table 1 lists these tools in chronological order, and indicates when the features listed as lines are available (✓), not available (✗) or only partially available (✐). Transfer along *heterogeneous relations*, and while the oldest tool operates via a monolithic translation of an input proof term, others rather prove an *internal* implication lemma. *Anticipation* [19] refers to the need to define a dedicated structure for the signature to be transported. *Binders* (∀) can prevent transfer, as well as *dependent types*, which require univalence.

	[16]	[17]	[7]	[14]	[21]	[19]	[2]	[4]	Trocq
Heterogen. rel.	✓	✗	✓	✓	✓	✓	✓	✓	✓
Internal	✗	✓	✓	✓	✓	✓	✓	✓	✓
No anticipation	✓	✓	✓	✓	✓	✓	✗	✓	✓
Under ∀	✓	✓	✗	✓	✓	✓	✓	✓	✓
Dep. types	✓	✗	✗	✗	✗	✓	✓	✗	✓
Univalence-free	✓	✓	✓	✓	✓	✗	✗	✓	✓
Subrelations	✗	✓	✗	✗	✗	✗	✗	✗	✐
QERs	✗	✐	✐	✐	✐	✗	✓	✗	✐
Subtyping	✗	✗	✐	✐	✐	✗	✗	✐	✐
	Coq	Coq	Coq	Isabelle/HOL	Coq	HoTT	CubicalAgda	Coq	Coq or HoTT

Table 1. Comparison of proof transfer automation devices

Acknowledgments. The authors would like to thank Enrico Tassi and the anonymous reviewers. This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 101001995).

References

1. Anand, A., Morrisett, G.: Revisiting parametricity: Inductives and uniformity of propositions (2017), http://arxiv.org/abs/1705.01163
2. Angiuli, C., Cavallo, E., Mörtberg, A., Zeuner, M.: Internalizing representation independence with univalence. Proc. ACM Program. Lang. **5**(POPL), 1–30 (2021)
3. Barthe, G., Pons, O.: Type isomorphisms and proof reuse in dependent type theory. In: FoSSaCS. LNCS, vol. 2030, pp. 57–71. Springer (2001)
4. Blot, V., Cousineau, D., Crance, E., de Prisque, L.D., Keller, C., Mahboubi, A., Vial, P.: Compositional pre-processing for automated reasoning in dependent type theory. In: CPP. pp. 63–77. ACM (2023)
5. Cohen, C., Crance, E., Mahboubi, A.: Trocq: proof transfer for free, with or without univalence. In: Weirich, S. (ed.) Programming Languages and Systems. LNCS, vol. 14576, pp. 239–268. Springer, Cham (2024). https://doi.org/10.1007/978-3-031-57262-3_10
6. Cohen, C., Crance, E., Mahboubi, A.: coq-community/trocq: Trocq 0.1.5 (Jan 2024). https://doi.org/10.5281/zenodo.10563382
7. Cohen, C., Dénès, M., Mörtberg, A.: Refinements for free! In: CPP. LNCS, vol. 8307, pp. 147–162. Springer (2013)
8. Dénès, M., Mörtberg, A., Siles, V.: A refinement-based approach to computational algebra in coq. In: ITP. LNCS, vol. 7406, pp. 83–98. Springer (2012)
9. Dunchev, C., Guidi, F., Coen, C.S., Tassi, E.: ELPI: fast, embeddable, λProlog interpreter. In: LPAR. LNCS, vol. 9450, pp. 460–468. Springer (2015)
10. Frühwirth, T., Raiser, F.: Constraint handling rules: Compilation, execution, and analysis (2011)
11. Haftmann, F., Krauss, A., Kuncar, O., Nipkow, T.: Data refinement in Isabelle/HOL. In: ITP. LNCS, vol. 7998, pp. 100–115. Springer (2013)
12. Huffman, B., Kuncar, O.: Lifting and transfer: A modular design for quotients in Isabelle/HOL. In: CPP. LNCS, vol. 8307, pp. 131–146. Springer (2013)
13. Krishnaswami, N.R., Dreyer, D.: Internalizing relational parametricity in the extensional calculus of constructions. In: CSL. LIPIcs, vol. 23, pp. 432–451. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2013)
14. Lammich, P.: Automatic data refinement. In: ITP. LNCS, vol. 7998, pp. 84–99. Springer (2013)
15. Lammich, P., Lochbihler, A.: Automatic refinement to efficient data structures: A comparison of two approaches. J. Autom. Reason. **63**(1), 53–94 (2019)
16. Magaud, N.: Changing data representation within the Coq System. In: TPHOLs. LNCS, vol. 2758, pp. 87–102. Springer (2003)
17. Sozeau, M.: A new look at generalized rewriting in type theory. J. Formaliz. Reason. **2**(1), 41–62 (2009)
18. Sozeau, M., Anand, A., Boulier, S., Cohen, C., Forster, Y., Kunze, F., Malecha, G., Tabareau, N., Winterhalter, T.: The metacoq project. J. Autom. Reason. **64**(5), 947–999 (2020)
19. Tabareau, N., Tanter, É., Sozeau, M.: The marriage of univalence and parametricity. Journal of the ACM (JACM) **68**(1), 1–44 (2021)
20. Tassi, E.: Elpi: an extension language for Coq (Metaprogramming Coq in the Elpi λprolog). In: CoqPL (January 2018), https://hal.inria.fr/hal-01637063
21. Zimmermann, T., Herbelin, H.: Automatic and transparent transfer of theorems along isomorphisms in the coq proof assistant. In: CICM (Work in Progress). pp. 50–62 (2015)