



Automatically Generalizing Theorems Using Typeclasses

Alexander Best

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

August 1, 2021

Automatically Generalizing Theorems Using Typeclasses

Alex J. Best
Department of Mathematics
Vrije Universiteit Amsterdam - Faculty of Science
De Boelelaan 1111
1081 HV Amsterdam
The Netherlands
alex.j.best@gmail.com

Abstract

When producing large formally verified mathematical developments that make use of typeclasses it is easy to introduce overly strong assumptions for theorems and definitions. We consider the problem of recognizing from the elaborated proof terms when typeclass assumptions are stronger than necessary. We introduce a metaprogram for the Lean theorem prover that finds and informs the user about possible generalizations.

1 Motivation

Developing and maintaining large libraries of formalized mathematics in an interactive theorem prover, such as the libraries `mathlib` [mC20] in Lean [dMKA⁺15], `Mathematical Components` [MT21] in Coq [Tea21], or the `Isabelle/HOL` [NPW02] libraries is a time consuming process. In addition to the overall design and organisation of the library, inclusion of useful results and writing of proofs, library maintenance requires maintaining interoperability of theories and deduplication of overlapping or identical results. Contributors to such a library with a specific formalization goal in mind often prove background results in the generality that they are needed. When such work is included into a large library it is desirable that all results be maximally useful for other potential use cases and not overly specific, this often leads to time consuming refactoring efforts before new material can be included.

The goal of this paper is to demonstrate that when typeclasses are used to implement hierarchies of assumptions on formalized objects it is possible to automate the process of generalizing results by relaxing typeclass assumptions. Pons [Pon02] has previously considered a similar problem, and describes a mechanism for generalizing assumptions about functions and types, but without specifically considering hierarchies of typeclasses.

Our focus is mostly on the applications to the formalization of mathematics, it is also conceivable however that relaxing typeclass assumptions can be of use for other users of interactive theorem provers, such as for software verification.

2 Typeclasses

Typeclasses were introduced by Wadler [WB89] as a way to provide polymorphism in strongly typed programming languages. They are used in interactive theorem provers [Wen97] and their libraries of formalized mathematics

to manage levels of mathematical structures on an object [SvdW10, SvdW11]. When a term is used the user does not ordinarily supply typeclass arguments, instead they are inserted automatically by a typeclass system that finds instances to match the needed argument, possibly by chaining together the instances declared in the library. From the perspective of interactive theorem provers, typeclasses reduce the burden on the user of the tool to be explicit about which instances are used and in fact hide them completely from the end user.

Commonly used typeclass hierarchies for mathematical libraries include: algebraic structures, including binary operations and their properties, properties of relations and ordering, topological spaces, (pseudo-)metric spaces, and their properties. Additionally mixtures of the aforementioned hierarchies that interact, such as ordered algebraic structures, topological groups and normed algebraic structures may also be included in the typeclass system, leading to a highly nontrivial hierarchy.

In addition to the situation mentioned in the introduction, where results introduced for one purpose may be generalized to a more widely useful form, generalization can also become possible when new typeclasses are inserted into an existing hierarchy. For instance in the mathlib library the concepts of non-unital and/or non-associative rings were added in March 2021¹, several years after rings were introduced in core Lean. In the interim period several thousand lemmas and theorems were added to mathlib concerning rings. This leaves the difficult library maintenance problem of filling out the library of lemmas for these newly introduced algebraic structures, without duplicating existing work.

3 Examples

We introduce some of the key considerations by looking at some examples in the Lean theorem prover using the mathlib library.

Consider the following lemma:

```
lemma mul_inv {G : Type*} [ordered_comm_group G] (a b : G) : (a * b)-1 = a-1 * b-1 :=
by rw [mul_inv_rev, mul_comm]
```

It is clear that the statement and the proof require a commutative group structure, but that the assumption that G be an *ordered* commutative group plays no role. Despite the shortness of the statement and proof script the generated proof term contains many chains of typeclass instances including those shown in Figure 1.

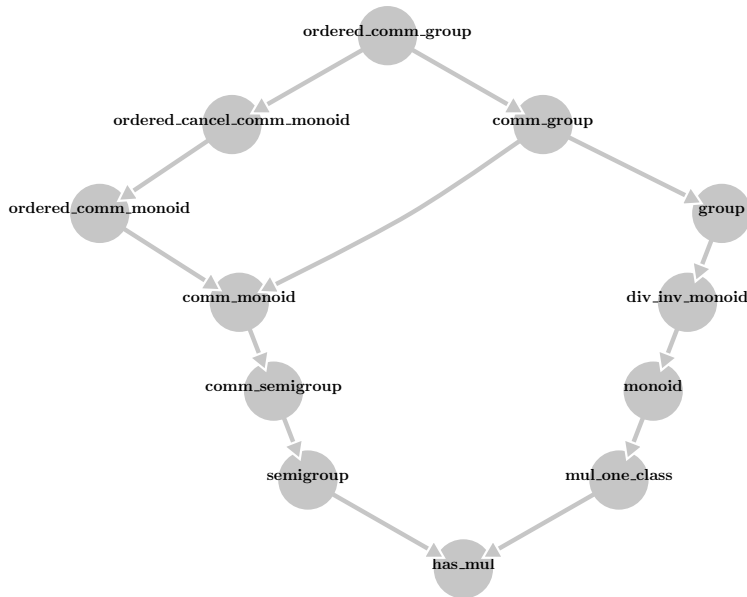


Figure 1: Typeclass chains in the proof of `mul_inv`.

There are in addition other instances between these classes, not made use of in the original proof, but available to the typeclass system.

¹See <https://github.com/leanprover-community/mathlib/pull/6786>.

The only typeclass parameters actually required to satisfy the assumptions of the applied lemmas `mul_inv_rev` and `mul_comm` are `group G` and `comm_semigroup G`. We cannot replace the assumption `ordered_comm_group G` with this pair of assumptions however as they have overlapping fields, they both should refer to the same multiplication operation. This is visible in the figure as there is an instance chain from each of these to `has_mul`. Instead a meet of these two typeclasses in the typeclass graph should be chosen as the generalization, in this case it is `comm_group`.

The following theorem states that given a ring homomorphism between two fields and a natural number p , one of the fields has characteristic p if and only if the other has characteristic p (including $p = 0$):

```
lemma ring_hom.char_p_iff_char_p {K L : Type*} [field K] [field L]
(f : K →+* L) (p : ℕ) : char_p K p ↔ char_p L p :=
begin
  split;
  { introI _c, constructor, intro n,
    rw [← @char_p.cast_eq_zero_iff _ _ p _c n, ← f.injective.eq_iff, f.map_nat_cast,
      f.map_zero] }
end
```

We see that the proof script splits the iff statement into each direction, but both directions are proved by the same tactic block. It is non-trivial to determine just by reading the proof given what the weakest assumptions possible are, and it is not immediately clear from the statement either. While it is of course possible to work out the correct generality by hand by inspecting the assumptions needed for each lemma applied, we hope this provides a realistic example where the “right” answer is not immediately clear.

In fact the weakest possible typeclass assumptions for the proof of `ring_hom.char_p_iff_char_p` are precisely the assumptions needed to apply `ring_hom.injective` to perform the rewrite `← f.injective.eq_iff`. These are that K should be a division ring, and L should be a nontrivial semiring. This example highlights another important point; simply taking the meet in the typeclass graph of the required classes is not always optimal. The meet of the classes `nontrivial` and `semiring` in the present version of `mathlib` is `domain`, once again an overly strong assumption, as there is no reason not to simply assume `[nontrivial K] [semiring K]`. We should allow one assumption to be replaced by several distinct assumptions if there is no conflict in doing so. The typeclasses `nontrivial` and `semiring` both provide instances of the `nonempty` typeclass, which contains no data, so it is a `Prop` and hence a subsingleton in Lean. Thus despite the fact that they do both provide instances of the same class, the instances will be equal and there is no issue allowing both assumptions separately.

4 Algorithm

We now describe in more detail an algorithm to find possible typeclass generalizations, given a proof term.

One important thing to note here is that we also handle typeclasses with more than one argument that are partially applied. For instance it is desirable to be able to generalize `group G` to `has_pow G ℤ`. This is distinct from `monoid G` generalizing to `has_pow G ℕ`, so we treat these partially applied typeclasses as the basic objects of interest, we call them *bound classes*.

The first step is to traverse the entire environment and build a directed graph of bound classes with instances between them. We also precompute the transitive closure and a topological sort of this graph.

Algorithm 1: Generalising typeclasses assumptions in a declaration.

```
foreach typeclass argument c do
  S ← [];
  foreach maximal chain of instances ending in c in the statement or proof do
    | add the head bound class of the chain to S;
  end
  T ← the set of strongly connected components for the subgraph of bound classes reachable from S
    (ignoring subsingletons);
  o ← [];
  foreach u in T do
    | add a meet of u to o;
  end
  return o
end
```

5 Implementation

We have implemented the above ideas for the Lean theorem prover as a metaprogram written in Lean itself². This means that no external tools beyond Lean itself are needed for users to integrate this tool into their Lean developments. We may also make use of the linter framework [vDEL20] which provides a convenient means for users to run different checks on the current file, and for continuous integration tools to regularly lint an entire library on multiple cores. This introduces a dependency on mathlib itself, but this tool can be imported and applied to any Lean development, and the import removed when generalization is complete.

We note that it is important to precompute and cache useful information about the typeclass graph, such as a topological sort of the graph, and its transitive closure. With the current size of mathlib this precomputation is quite reasonable, even with non-optimal algorithms.

6 Summary of results

By running the metaprogram described above on a current version of the mathlib library³, which contains around 80,000 declarations we can gauge the effectiveness of this approach. Note that the process of generalizing typeclass assumptions using such a tool is naturally an iterative process, by generalizing one theorem other theorems that made use of the original theorem may also become generalizable. This makes it most natural to use such a tool interactively in a file when formalizing some theory. Nevertheless, for ease of measurement we simply apply a single pass in our test, thus there are likely many more interesting or useful typeclass relaxations that can be discovered by such a tool if used iteratively.

In total the current implementation produces 2877 results, some of these are false positives, or at least not particularly useful generalizations. To get a sense of what generalizations were possible in a library like mathlib, we list in Table 1 some of the most common generalizations found by this pass.

Table 1: Generalizations found in the Lean mathlib.

Original typeclass	Replacement typeclasses (number of times replaced)
<code>comm_ring</code>	<code>comm_semiring</code> (42), <code>ring</code> (40), <code>semiring</code> (27), <code>has_zero</code> (8)
<code>add_comm_group</code>	<code>add_comm_monoid</code> (96), <code>add_group</code> (5), <code>sub_neg_monoid</code> (5), <code>{has_add, has_neg, has_zero}</code> (3)
<code>semiring</code>	<code>non_assoc_semiring</code> (53), <code>non_unital_non_assoc_semiring</code> (23), <code>{add_comm_semigroup, has_one}</code> (13), <code>{add_comm_monoid, has_mul}</code> (8), <code>has_mul</code> (7), <code>has_zero</code> (5)
<code>field</code>	<code>division_ring</code> (23), <code>comm_ring</code> (12), <code>integral_domain</code> (12), <code>semiring</code> (7), <code>domain</code> (4), <code>ring</code> (4), <code>has_inv</code> (3)
<code>ring</code>	<code>semiring</code> (55), <code>non_assoc_semiring</code> (8), <code>{add_group, has_mul}</code> (4), <code>{has_add, has_neg, mul_zero_one_class}</code> (4)
<code>preorder</code>	<code>has_lt</code> (36), <code>has_le</code> (31), <code>{has_le, is_refl}</code> (3), <code>{has_lt, is_asymm}</code> (3), <code>{has_lt, is_irrefl, is_trans}</code> (3)
<code>comm_semiring</code>	<code>semiring</code> (26), <code>monoid</code> (10), <code>add_zero_class</code> (4), <code>{has_mul, is_associative, is_commutative}</code> (4), <code>has_pow</code> (4), <code>mul_one_class</code> (4)
<code>normed_space</code>	<code>module</code> (23), <code>semi_normed_space</code> (38)
<code>add_monoid</code>	<code>add_zero_class</code> (51), <code>{has_add, has_zero}</code> (5)
<code>monoid</code>	<code>mul_one_class</code> (34), <code>has_mul</code> (11), <code>{has_mul, has_one}</code> (5), <code>has_pow</code> (5)
<code>module</code>	<code>has_scalar</code> (20), <code>distrib_mul_action</code> (8), <code>mul_action</code> (6)
<code>normed_group</code>	<code>semi_normed_group</code> (36), <code>has_norm</code> (10), <code>has_nnorm</code> (3)
<code>integral_domain</code>	<code>comm_ring</code> (15), <code>domain</code> (8), <code>{comm_ring, no_zero_divisors}</code> (7), <code>comm_monoid</code> (3), <code>{has_mul, has_zero, no_zero_divisors}</code> (2), <code>{no_zero_divisors, semiring}</code> (2)
<code>partial_order</code>	<code>preorder</code> (33)

²the implementation is available at <https://github.com/alexjbest/lean-generalisation>

³commit 29b63a7e91d079b159dfc2cf0fb4d2a1ce1c409b, not including core Lean library

Here several generalizations such as changing a field typeclass to a (commutative) ring give rather large relaxations of structure, and potentially introduce a large number of new useful lemmas in the library.

The replacement of `integral_domain` with `comm_ring` and `no_zero_divisors` appears mathematically trivial, but in fact results in the removal of the assumption that the ring in question be nontrivial (which is part of `mathlib`'s definition of an integral domain). Replacements such as this are mathematically not significant and it is not the case that this generalization would be useful to an end user directly. However, making this generalization relieves the user of such a theorem from the burden of adding a non-trivial assumption which can then proliferate through the library. These small assumptions that nonetheless need verifying whenever a theorem is applied can slow down formalization efforts and make future formalization unduly time consuming and tedious compared with the actual mathematics being formalized.

It is important that such a tool be usable interactively when working on formalizing a mathematical theory. Running our implementation on the `mathlib` file `topology/ordered/basic.lean` (the second longest in `mathlib` by number of lines) takes between 1 and 2 minutes on a modern laptop, depending on if the typeclass graph is already computed and cached. This checks 475 declarations and reports 183 results currently.

7 Limitations and Further Work

By only inspecting the existing proof of a theorem naturally not all possible generalizations of theorems can be found. It is quite easy to write proofs that require stronger assumptions than necessary, especially when using powerful automation. Trying to remove this limitation completely shifts the problem to one of automatic theorem proving. However, an intermediate problem where some automation may be possible is to find theorems for which the same proof script proves a generalized theorem. This is especially likely with small proofs that make heavy use of automation to begin with. Here a brute force strategy of weakening typeclasses in theorem statements and re-running the same proof script seems far too slow to be used on the scale of a large library in a prover such as Lean, however similar techniques have been used in the Mizar system [Ala13]. Nevertheless, a tool to automate this process could still be useful for small new developments.

Currently the system as implemented in Lean will provide to the user a list of possible typeclass generalizations by printing the name of the declaration and describing the argument that can be generalized. The structure of Lean files allows for variables to be defined within a `section` these assumptions can then hold for all declarations in large chunks of the file. This decoupling of the location of the assumptions used for each theorem and the theorems themselves sometimes makes it tedious to rearrange the file to actually implement the suggestions. Therefore better tooling to reorganise Lean files in a content-aware manner will improve the usability of the tool described in this paper.

A mild extension of this work would be to consider not just assumptions, but also concrete types in a theorem statement. For instance theorems proven about explicitly constructed types such as the natural, rational or real numbers often make use of (ordered) algebraic properties of these types, which are filled in via typeclasses. Recognizing these typeclass chains in the same way as described here could allow parts of libraries concerning these concrete types to be generalized to include any object satisfying some typeclass hypotheses.

A core design principle of Lean 4, the next iteration of the Lean theorem prover [MU21], is that the majority of the system be written in Lean itself. This allows for extensibility of the Lean system by users, rather than having to build different versions of Lean for user extensions. Using this it may be possible to integrate the minimisation of typeclass assumption into Lean's own typeclass system, providing a more accurate and efficient tool.

8 Conclusion

Recognizing too strong typeclass assumptions automatically is possible, and can be done efficiently. Such tools can save maintainers of mathematical libraries time by providing more general results for free from an existing library, and provide users of formal proof systems with interesting information about what generality theorems hold in.

8.1 Acknowledgements

I would like to thank Floris van Doorn for helpful discussions on an earlier version of this work, and the referees for their useful comments and suggestions. This work was supported by the Hariri Institute for Computing,

the Simons Collaboration on Arithmetic Geometry, Number Theory, and Computation, via Simons Foundation grant #550023, and NWO Vidi grant 639.032.613.

References

- [Ala13] Jesse Alama. Eliciting Implicit Assumptions of Mizar Proofs by Property Omission. *Journal of Automated Reasoning*, 50(2):123–133, February 2013.
- [dMKA⁺15] Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. The Lean Theorem Prover (System Description). In Amy P. Felty and Aart Middeldorp, editors, *Automated Deduction - CADE-25*, volume 9195, pages 378–388, Cham, 2015. Springer International Publishing. Series Title: Lecture Notes in Computer Science.
- [mC20] The mathlib Community. The lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020*, pages 367–381, New York, NY, USA, January 2020. Association for Computing Machinery.
- [MT21] Assia Mahboubi and Enrico Tassi. *Mathematical Components*. Zenodo, January 2021. Version Number: 1.0.1.
- [MU21] Leonardo de Moura and Sebastian Ullrich. The Lean 4 Theorem Prover and Programming Language. In André Platzer and Geoff Sutcliffe, editors, *Automated Deduction – CADE 28*, Lecture Notes in Computer Science, pages 625–635, Cham, 2021. Springer International Publishing.
- [NPW02] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. *Isabelle/HOL: a proof assistant for higher-order logic*, volume 2283. Springer Science & Business Media, 2002.
- [Pon02] Olivier Pons. Generalization in Type Theory Based Proof Assistants. In Paul Callaghan, Zhaohui Luo, James McKinna, Robert Pollack, and Robert Pollack, editors, *Types for Proofs and Programs*, Lecture Notes in Computer Science, pages 217–232, Berlin, Heidelberg, 2002. Springer.
- [SvdW10] Bas Spitters and Eelis van der Weegen. Developing the Algebraic Hierarchy with Type Classes in Coq. In Matt Kaufmann and Lawrence C. Paulson, editors, *Interactive Theorem Proving*, Lecture Notes in Computer Science, pages 490–493, Berlin, Heidelberg, 2010. Springer.
- [SvdW11] Bas Spitters and Eelis van der Weegen. Type Classes for Mathematics in Type Theory. *Mathematical Structures in Computer Science*, 21(4):795–825, August 2011. arXiv: 1102.1323.
- [Tea21] The Coq Development Team. The Coq Proof Assistant, January 2021.
- [vDEL20] Floris van Doorn, Gabriel Ebner, and Robert Y. Lewis. Maintaining a Library of Formal Mathematics. In Christoph Benzmüller and Bruce Miller, editors, *Intelligent Computer Mathematics*, Lecture Notes in Computer Science, pages 251–267, Cham, 2020. Springer International Publishing.
- [WB89] P. Wadler and S. Blott. How to make ad-hoc polymorphism less ad hoc. In *Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages, POPL '89*, pages 60–76, New York, NY, USA, January 1989. Association for Computing Machinery.
- [Wen97] Markus Wenzel. Type classes and overloading in higher-order logic. In Elsa L. Gunter and Amy Felty, editors, *Theorem Proving in Higher Order Logics*, Lecture Notes in Computer Science, pages 307–322, Berlin, Heidelberg, 1997. Springer.