

Reimplementing Mizar in Rust

Mario Carneiro  

Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

This paper describes a new open-source proof processing tool, `mizar-rs`, a wholesale reimplementation of core parts of the Mizar proof system, written in Rust. In particular, the “checker” and “analyzer” of Mizar are implemented, which together form the trusted core of Mizar. This is to our knowledge the first and only external implementation of these components. Thanks to the loose coupling of Mizar’s passes, it is possible to use the checker as a drop-in replacement for the original, and we have used this to verify the entire MML in 11.8 minutes on 8 cores, a $4.8\times$ speedup over the original Pascal implementation. Since Mizar is not designed to have a small trusted core, checking Mizar proofs entails following Mizar closely, so our ability to detect bugs is limited. Nevertheless, we were able to find multiple memory errors, four soundness bugs in the original (which were not being exploited in MML), in addition to one non-critical bug which was being exploited in 46 different MML articles. We hope to use this checker as a base for proof export tooling, as well as revitalizing development of the language.

2012 ACM Subject Classification Mathematics of computing → Mathematical software performance; Software and its engineering → Formal methods; Social and professional topics → Systems analysis and design

Keywords and phrases Mizar, proof checker, software, Rust

Digital Object Identifier 10.4230/LIPIcs.ITP.2023.10

Related Version *Previous Version:* <https://arxiv.org/abs/2304.08391>

Supplementary Material *Software (Source Code):* <https://github.com/digama0/mizar-rs/tree/itp2023/itp2023>, archived at `swh:1:dir:2522beed4f5fce87ce3193ff0359def1dcff1d7c`

Funding Work supported by the Hoskinson Center for Formal Mathematics at Carnegie Mellon.

Acknowledgements I would like to thank Josef Urban and Jeremy Avigad for their support and encouragement, and Adam Naumowicz and SUM for taking the courage to finally open-source Mizar, without which it would not have been possible to publish this paper and accompanying code.

1 Introduction

The Mizar language [3] is a proof language designed by Andrzej Trybulec in 1973 for writing and checking proofs in a block structured natural deduction style. The Mizar project more broadly has been devoted to the development of the language and tooling, in addition to the Mizar Mathematical Library (MML) [2], a compendium of “articles” on a variety of mathematical topics written in the Mizar language. The MML is one of the largest and oldest formal mathematical libraries in existence, containing (at time of writing) 1434 articles and over 65,000 theorems.

The “Mizar system” is a collection of tools for manipulating Mizar articles, used by authors to develop and check articles for correctness, and maintained by the Association of Mizar Users (SUM). Arguably the most important tool in the toolbox is `verifier`, which reads a Mizar article and checks it for logical correctness. The starting point for this work is the goal of extracting formal proofs from the `verifier` which can be checked by a tool not directly connected to Mizar. This turned out to be quite challenging, and this paper will explain how we achieved a slightly different goal – to build `mizar-rs`, a Rust program which checks proofs in the same manner as the `verifier`, with an eye for proof generation.



© Mario Carneiro;

licensed under Creative Commons License CC-BY 4.0

14th International Conference on Interactive Theorem Proving (ITP 2023).

Editors: Adam Naumowicz and René Thiemann; Article No. 10; pp. 10:1–10:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Why hasn't this been done before?

The first challenge is that there is no “spec” for the Mizar language: a Mizar proof is one that `verifier` accepts. Some of it is “obvious” things which users can figure out by playing with the language, and there are reference works at varying levels of detail [5, 10], but if we want to check the entire MML then approximations aren't going to cut it. There are some writings online which give the general sense of the algorithm, but there is nowhere to look for details except the source code.

This leads to the second challenge, which is that the source code for the Mizar system is not publicly available, or at least was not while this project was under development.¹ It was only made available to members of the SUM. Luckily, it is possible to obtain a membership, but the price of admission is that one has to write a Mizar article first. Thus, for this project the author contributed an article regarding the divergence of the sum of prime reciprocals [4].

The third challenge is that Mizar does not have a “small trusted kernel” in any reasonable sense. The main components that perform trusted activities are the “analyzer” and the “checker”, and most of the code of the verifier is contained in these modules. Moreover, the code is written in Pascal and Polish (two languages that the author is not very good at), and is 50 years old, meaning that there is a huge amount of technical debt in the code.² To get a good sense of what everything is doing and why it is correct (or not), we estimated the best option would simply be to rewrite it and closely follow what the original code is doing. Rust was chosen as the implementation language because it is well suited to writing console applications with a focus on correctness, performance, and refactoring, all of which were important for something that could play the role of `verifier`.

We are not the first to present aspects of Mizar from an outsider's perspective. J. Harrison's Mizar mode in HOL Light [6] is a simulation of large parts of the Mizar frontend, with the MESON prover used in place of the Mizar checker. J. Urban's MPTP [12] project exports the statements of checker goals and sends them to ATPs. And most recently, Kaliszyk and Pał presented a model of Mizar as encoded in Isabelle [7], but stop short of a fully automatic proof export mechanism.

One may reasonably ask why the MPTP export is not already everything you could wish regarding proof export. There are two problems: (1) Because it is using external ATPs rather than Mizar's own code, it does not and cannot reasonably be expected to achieve 100% success on the MML, and even when it works the Mizar checker often outperforms the external ATP, because the MML is significantly overfit to the Mizar checker. (2) Because it instruments the “checker inference” level, we verify “by” steps but not the skeleton steps that glue them together, and some significant type reasoning happens in the analyzer.

In this paper, we will try to give an implementer's perspective of Mizar: how it works under the hood, and some things we discovered while trying to replicate the functionality. The reader should be aware that this paper is focused on internal details and implementation correctness of the Mizar system, not the abstract theory of Mizar (that is, Tarski–Grothendieck set theory and the soft type system). See [5, 7] for information on these aspects of the language.

¹ A happy after-effect of this project is that the Mizar source code itself has been made public! It is now available at <https://github.com/MizarProject/system>.

² Technically the current incarnation of Mizar (“PC Mizar”) only dates to 1986. See [9] for the early history of Mizar.

2 Mizar internals: A (determined) user’s perspective

When one runs the `verifier` on a Mizar file, say `tarski.miz`, one will see something like this:

```
Verifier based on More Strict Mizar Processor, Mizar Ver. 8.1.11 (Linux/FPC)
Copyright (c) 1990-2022 Association of Mizar Users
Processing: mizshare/mml/tarski.miz
```

```
Parser [ 138] 0:00
MSM [ 132] 0:00
Analyzer [ 137] 0:00
Checker [ 137] 0:00
Time of mizarining: 0:00
```

This refers to modules named `Parser`, `MSM`, `Analyzer`, `Checker`, so one can reasonably guess that the code is split up into multiple passes and these are the names of the modules.

- `Parser` reads the text and converts it into an abstract syntax tree.
- `MSM` stands for “More Strict Mizar”, which is a slightly more elaborated version of the syntax tree. Notably, this pass does (some kinds of) name resolution, as well as resolving and auto-binding reserved variables.
- `Analyzer` is responsible for processing the large scale logical structure of a Mizar document. It resolves symbol overloading and checks types, as well as tracking the “thesis” as it is transformed by Mizar proof steps.
- `Checker` is called by the analyzer whenever there are atomic proof obligations. It is sometimes referred to as the “`by`” proof automation because it is called whenever there is a step proved using the `by` keyword.

Both the analyzer and the checker are soundness-critical components. Clearly the checker is important because it actually checks proofs – a bug here will mean that an incorrect proof is accepted – but the analyzer is also important because it determines the proof obligations that will be sent to the checker. If the analyzer sends the wrong assertion to the checker or simply doesn’t call the checker at all, then it might assert a theorem that hasn’t been properly checked.

Another visible effect of running `verifier` on a `.miz` file is that a huge number of other files are generated with bespoke extensions, which reveals another implementation quirk of the Mizar system, which is that the four components above are very loosely coupled and communicate only via the file system. The parser reads the `.miz` file and produces a `.wsx` file, which is read by `MSM` to produce `.msx`, which is read by the `transfer2analyzer` module (not mentioned above) to produce `.par`, which is read by `Analyzer` to produce `.xml`, which is read by `Checker` and verified.

The general data flow of a Mizar verification looks like this:

- The first tool to process the article is the “accommodator,” `accom`. This is the only tool which reads data from articles other than the current one; it is responsible for aggregating data from imports and collecting them for the current article.
 - reads: `mml.ini`, `mml.vct`, `mizar.dct`, `art.miz`, and `dep`.{`dco`, `dre`, `dcl`, `dno`, `did`, `drd`, `dpr`, `def`, `the`, `sch`}
 - writes: `art`.{`vc1`, `sg1`, `cho`, `aco`, `atr`, `ere`, `ecl`, `eno`, `nol`, `ano`, `frm`, `prf`, `eid`, `erd`, `epr`, `dfs`, `dfe`, `dfx`, `eth`, `esh`, `fil`, `err`, `log`}

Here `art` is the article being processed and `dep` is an article which is imported by this one. For example `xboole_0.miz` has a `theorems TARSKI` declaration in its `environ` section, so the accommodator will load `tarski.the` to get the theorems from article `tarski`

10:4 Reimplementing Mizar

and put them in `xbool_0.eth`. Similar things are done for constructors, notations, definitions, etc., each of which is in a separate file, which is why so many files are read and written in this step.

- The next pass is the parser, which parses the body of the `art.miz` article and produces a `art.wsx` file (Weakly Strict Mizar) containing the AST, along with `art.frx` containing the formats declared in the article.
- The MSM pass reads `art.wsx` and writes `art.msx`, performing name resolution and filling in the types of `reservation` variables in statements.
- The `transfer2analyzer` pass appears to be for backward compatibility reasons, as it reads `art.msx` and translates it to `art.par`, which is the same thing but in a less extensible format. It also resolves format references.
- The analyzer pass reads `art.{eno, epr, dfs, par}` and writes `art.xml`, which is the article AST again but using fully elaborated and typechecked terms, and with all the statements of checker subgoals explicitly annotated.
- The checker is the last pass. It reads `art.{ref, ere, atr, ecl, dfe, dfx, epr, eid, erd, eth, esh, xml}` and verifies the theorems.

The underlined files are the ones which we needed to parse while implementing the analyzer and checker passes.

Thanks to the efforts of J. Urban in 2004 [11], most of these internal files between the components are in XML format. Those that are not XML are highlighted in **red**. The only non-XML files we need to read are `art.ref` and `art.ere`, and both of these have rather simplistic number-list formats.

If one opens one of these files, one is presented with the next major challenge, which is that terms are pervasively indexed and hence reading the expressions can be quite difficult. Moreover, there are many *distinct* index classes which are differentiated only by the context in which they appear, so without knowing how the program processes the indices or what array is being accessed it is hard as an outsider to follow the references. In the Rust implementation this issue is addressed by using “newtypes” to wrap each integer to help distinguish different indexing sets. There are currently 36 of these newtypes defined: for example there are numbers for functors, selectors, predicates, attributes, formats, notations, functor symbols (not the same!), left bracket symbols, reserved identifiers, etc.

After parsing these, one ends up with an expression such as $\forall_ : M_1, R_4(K_1(B_1, K_2(N_2)), N_1)$, which means something like “for all x of the first type, fourth relation holds of the first function applied to x and the second function applied to 2, and 1”. For many purposes, this is sufficient for debugging, but one tends to go cross-eyed staring at these expressions for too long. Ideally we would be able to reverse this indexification to obtain the much more readable expression $\forall x : \mathbf{Nat}, x \cdot (-2) \leq 1$. However, the code to do this does not exist anywhere in Mizar, because the Mizar checker never prints expressions. The only output of the checker is a list of `(line, col, err_code)` triples, which are conventionally postprocessed by the `errflag` tool to insert markers like the following in the text:

```
for x being Nat holds x = x
proof
  let x be Nat;
end;
:: >, 70
:: >
:: > 70: Something remains to be proved
```

This error message is pointing at the **end** keyword, but notably it does not say *what* remains to be proved, here $x = x$. A debug build of Mizar will actually print out expressions to the `art.inf` file, but they are similar to the $R_4(K_1(\dots))$ style.

Luckily, this issue has been addressed outside the main Mizar codebase: J. Urban’s HTMLizer is an XSLT stylesheet which can transform the XML intermediate files into fully marked up reconstructed Mizar documents, and which we adapted to design the formatter.

With the present version of the formatter, and with appropriate debugging enabled, an input like the following:

```
for x,y being Nat holds x = 1 & y = 2 implies x + y = 3;
```

yields this debugging trace:

```
input: ∃ b0: natural set, b1: natural set st
  (b0 = 1) ∧ (b1 = 2) ∧ ¬((b0 + b1) = 3)
refuting 0 @ TEST:28:33:
  ∃ b0: natural set, b1: natural set st
    (b0 = 1) ∧
      (b0 c= 1) ∧
        (1 c= b0) ∧
          (b1 = 2) ∧
            (b1 c= 2) ∧
              (2 c= b1) ∧
                ¬((b0 + b1) = 3) ∧
                  ((b0 + b1) c= 3) → ¬(3 c= (b0 + b1)))
```

In addition to showing some of the formatting and indentation behavior of the reconstructed expression, this also reveals some aspects of the checker, like how the goal theorem has been negated and the definitional theorem $x = y \leftrightarrow x \subseteq y \wedge y \subseteq x$ has been eagerly applied during preprocessing.

One other feature that is demonstrated here is “negation desugaring”, which requires some more explanation. Internally, Mizar represents all expressions using only \neg , \forall and n -ary \wedge . So $P \rightarrow Q$ is mere syntax for $\neg(P \wedge \neg Q)$ and $\exists x, P(x)$ is actually $\neg\forall x, \neg P(x)$. (Even $P \leftrightarrow Q$ is desugared, to $\neg(P \wedge \neg Q) \wedge \neg(Q \wedge \neg P)$, so too much recursive use of \leftrightarrow can cause a blowup in formula size.) This normalization ensures that different spellings of the same formula are not distinguished, for example if the goal is $P \vee Q$ then one may prove it by **assume not P; thus Q**;. Double negations are also cancelled eagerly. In the formatter, we try to recover a natural-looking form for the expression by pushing negations to the leaves of an expression, and also writing $A \wedge B \rightarrow C \vee D$ if after pushing negations we get a disjunction such that the first few disjuncts have an explicit negation. (Mizar actually carries some annotations on formulas to help reconstruct whether the user wrote **not A or B** or **A implies B**, but we chose not to use this information as it is often not available for expressions deep in the checker so we wanted a heuristic that works well in the absence of annotation.)

3 The checker

Mizar is broadly based on first order logic, with the non-logical axioms of Tarski–Grothendieck set theory, with a type system layered on top. While types can depend on terms, so one may call it a dependent type theory, this is not a type system in the sense of Martin-Löf type theory: types are essentially just predicates over terms in an untyped base logic, and the language allows typing assertions to be treated as predicates.

3.1 Core syntax

The core syntax of Mizar uses the following grammar:

$a, b, c, d, e, F, G, H, K, U, V, P, R, S, T_M, T_G ::= \text{ident}$			
$t ::= a, b, c, d, e$	variable	$\varphi ::= \top$	true
n	numeral	$\neg\varphi$	negation
$\{F, G, K, U\}(\vec{t}_i)$	function application	$\bigwedge \vec{\varphi}_i$	conjunction
$H(\vec{t}_i) := t'$	local function app.	$\bigwedge_{b=t}^{t'} \varphi$	flex conjunction
the τ	choice	$\forall b : \tau. \varphi$	for all
$\{t \mid \vec{b}_i : \tau_i \mid \varphi\}$	Fraenkel	$\{P, R\}(\vec{t}_i)$	predicate
$\tau ::= \vec{\chi}_i T_M(\vec{t}_i)$	mode	$S(\vec{t}_i) := t'$	local pred.
$\vec{\chi}_i T_G(\vec{t}_i)$	struct type	$t \text{ is } \chi$	attribute
$\chi ::= \pm V(\vec{t}_i)$	attribute	$t : \tau$	qualification

Broadly speaking, terms are first order, meaning applications of function symbols to variables. The exception is the Fraenkel operator $\{t \mid \vec{b}_i : \tau_i \mid \varphi\}$, which should be read as “the set of $t(\vec{b}_i)$ where the $\vec{b}_i : \tau_i$ are such that $\varphi(\vec{b}_i)$ ”, where the variables $\vec{b}_i : \tau_i$ are quantified in both t and φ .

All terms t have a type τ , and there is a robust subtyping system – most terms have *many* types simultaneously. All types are required to be nonempty, which is what justifies the **the** τ constructor for indefinite description. Types are composed of a collection of attributes (a.k.a. clusters) χ applied to a base type $\{T_M, T_G\}(\vec{b}_i)$. Regular types are called “modes”; new modes can be defined by carving out a subset of an existing mode, and modes need not define a set (in particular, **object** is a primitive mode which is the supertype of everything, and does not constitute a set per Russell’s paradox). Structure types are roughly modeled after partial functions from some set of “tags”, but they are introduced axiomatically, similarly to how structure types are treated in a dependent type theory such as Coq or Lean.

► **Remark.** Although one can define a mode for any ZFC set or class, modes and sets are not interchangeable in Mizar because they lie in different syntactic classes. Modes are types, which go in the type argument of quantifiers such as the **Nat** in $\forall x : \mathbf{Nat}. x \geq 0$, while sets are objects, which can be passed to functions, like $|\mathbf{NAT}| = \aleph_0$. In the MML, **NAT** is the set of natural numbers, while **Nat** and **Element of NAT** are modes which describe the type of natural numbers (and in general **Element of A** can be used to treat a set as a type).

Attributes, also known as “adjectives” or “clusters”, are modifiers on types, which may be described as intersection typing in modern terminology, although attributes do not stand alone as types. For example, “**x is non empty finite set**” means $\neg(x \text{ is empty}) \wedge (x \text{ is finite}) \wedge (x : \mathbf{set})$. The Mizar system treats the collection of attributes on a type as an unordered list for equality comparisons.

Formulas are composed mainly from negation, conjunction and for-all, but there are some extra formula constructors that deserve attention:

- Not represented in the grammar is that $\neg\neg\varphi$ is identified with φ , and internally there is an invariant that $\neg\neg$ never appears.
- Similarly, conjunctions are always flattened, so $P \wedge (Q \wedge R) \wedge S$ becomes $P \wedge Q \wedge R \wedge S$.
- Unlike most type theories, typing assertions are reified into an actual formula, so it is possible to say $4/2 : \mathbb{N}$ and $\neg(-1 : \mathbb{N})$.
- There is also a predicate for having an attribute, and $(t : \chi \tau) \leftrightarrow (t \text{ is } \chi) \wedge (t : \tau)$ is provable.

- Flex-conjunction is written with syntax such as $P[1] \ \& \ \dots \ \& \ P[n]$ (with literal “...”), and Mizar knows that this expression is equivalent both to $\forall x : \mathbb{N}. 1 \leq x \leq n \rightarrow P(x)$ as well as to an explicit conjunction (when n is a numeral). For example $P[1] \ \& \ \dots \ \& \ P[3]$ will be expanded to $P[1] \ \& \ P[2] \ \& \ P[3]$ in the checker.

The different letters for functions, predicates, and variables correspond to the different roles that these can play, although in most situations they are treated similarly.

- “Locus variables” (a) are used in function declarations to represent the parameters of a function, mode, cluster, etc. For example a local function might be declared as $K(\vec{a}_i) : \tau := t$ where τ and t are allowed to depend on the \vec{a}_i , and then when typing it we would have that $K(\vec{t}_i)$ has type $\tau[\vec{a}_i \mapsto \vec{t}_i]$.
- “Bound variables” (b) are de Bruijn *levels*, used to represent variables in all binding syntaxes. So for example, $\forall x : \mathbb{N}. \forall y : \mathbb{N}. x \leq x + y$ would be represented as $\forall_1 : \mathbb{N}. \forall_2 : \mathbb{N}. b_0 \leq b_0 + b_1$. NB: there are two conventions for locally nameless variables, called “de Bruijn indices” (variables are numbered from the inside out) and “de Bruijn levels” (variables are numbered from the outside in), and Mizar’s choice of convention is the less common one.
- “Constants” (c) are variables that have been introduced by a **consider**, **given**, or **take** declaration, as well as Skolem constants introduced in the checker when there are existentially quantified assumptions. Constants may or may not be defined to equal some term (for example **take** $x = t$ will introduce a variable x equal to t) and the checker will use this in the equalizer if available.
- “Inference constants” (d) are essentially checker-discovered abbreviations for terms. This is used to allow for subterm sharing, as terms are otherwise completely unshared.
- “Equivalence classes” (e) are used in the equalizer to represent equivalence classes of terms up to provable equality. So for example if the equalizer sees a term $x + y$ it will introduce an equivalence class e_1 for it and keep track of all the things known to be in this class, say $e_1 = x + y = y + x = e_1 + 0$.
- Functors (K) and predicates (R) are the simplest kind of definition, they correspond to function and relation symbols in traditional FOL.
- Scheme functors (F) and scheme predicates (P) are the higher-order analogue of constants. They are only valid inside scheme definitions, which declare these at the start and then use them in the statement and proof.
- Local functors (H) and local predicates (S) are function declarations within a local scope, declared with the **deffunc** or **defpred** keywords. Internally every predicate carries the result of substituting its arguments, so if we declare **deffunc** $H_1(x, y) := x + y$ then an expression like $2 \cdot H_1(x - 1, y)$ is really stored as $2 \cdot (H_1(x - 1, y) := x - 1 + y)$. Some parts of the system treat such an expression like a function application, while others treat it like an abbreviation for a term.
- An aggregate functor (G) is Mizar terminology for a structure constructor $\{\text{foo} := x, \text{bar} := y\} : \text{MyStruct}$, and the converse operator is a selector (U), which is the projection $t.\text{foo}$ for a field. (The Mizar spelling for these operators is **MyStruct**(# x , y #) for the constructor and **the foo of** t for the projection.)

3.2 Structure of the checker

The checker is called whenever there is a proof such as “ $2 + 2 = 4;$ ” or “ $x \leq y$ **by** $A1, \text{Thm2};$ ”: basically any time there is a **by** in the proof text, as well as when propositions just end in a semicolon, this being the nullary case of **by**. It consists of three

major components, although there is another piece that is used even before the checker is properly called on a theorem statement:

0. **Attribute inference** (“rounding-up”) is used on every type $\vec{\chi}^- \tau$ to prove that it is equivalent to $\vec{\chi}^+ \tau$ where $\vec{\chi}^+$ is a superset of $\vec{\chi}^-$. Internally we actually keep *both* versions of the attributes (the “lower cluster” $\vec{\chi}^-$ being the one provided by the user and the “upper cluster” $\vec{\chi}^+$ being what we can infer by applying all inference rules), because it is useful to be able to prove that two clusters are equal if $\vec{\chi}_1^- \subseteq \vec{\chi}_2^+$ and $\vec{\chi}_2^- \subseteq \vec{\chi}_1^+$.
1. The **pre-checker** takes the list of assumptions, together with the negated conjecture, and performs a number of normalizations on them, skolemizing existentials, removing vacuous quantifiers, and expanding some definitions. Then it converts the whole formula into *disjunctive* normal form (DNF) and tries to refute each clause.
2. The **equalizer** does most of the “theory reasoning”. It replaces each term in the clause with an equivalence class, adding equalities for inference constants and defined constants, as well as registered equalities and **symmetry** declarations, as well as any equalities in the provided clause. It also evaluates the numeric value (for $2 + 2 = 4$ proofs) and polynomial value (for $(x + 1)^2 = x^2 + 2x + 1$ proofs) of equality classes and uses them to union classes together. These classes also have many *types* since many terms are going into the classes, and all of the attributes of these types are mixed together into a “supercluster” and rounded-up some more, potentially leading to a contradiction. For example if we know that $x = y$ and x is positive and y is negative then the supercluster for $\{x, y\}$ becomes **positive negative** which is inconsistent.
3. The last step is the **unifier**, which handles instantiation of quantifiers. This is relatively simplistic and non-recursive: for each assumption $\forall \vec{b}. P(\vec{b})$ it will instantiate $P(\vec{v})$ (where \vec{v} are metavariables) and then construct the possible assignments (as a DNF) to the variables that would make $P(\vec{v})$ inconsistent with some other assumption. It tries this for each forall individually, and if it fails, it tries taking forall assumptions in pairs and unifying them. If that still fails then it gives up – it does not attempt a complete proof method.

Finally, there is one more component which is largely separate from the “**by**” automation:

4. The **schematizer** is the automation that is called for justifications starting with “**from**”. These are scheme instantiations. In this case it is very explicitly given the list of hypotheses in the right order, so the only thing it needs to do is to determine an assignment of the scheme variables (F, P) to regular functors and predicates (K, R) or local functors and predicates (H, S). Users are often required to introduce local predicates in order to apply a scheme. Nullary scheme functors (a.k.a constant symbols) can be unified with arbitrary terms, however.

Although we cannot go into full detail on the algorithms here, in the following sections we will go into some of the highlights, with an emphasis on what it takes to audit the code for logical soundness, through the lens of root-causing some soundness bugs.

3.3 Requirements

One aspect of the Mizar system that is of particular interest is the concept of “requirements”, which are definitions that the checker has direct knowledge of. For example, the grammar given above does not make any special reference to equality: it is simply one of the possible relation symbols $R_i(x, y)$, and the relation number for equality can vary from one article to the next depending on how the accommodator decides to order the imported relation symbols. Nevertheless, the checker clearly needs to reason about equality to construct equality equivalence classes.

To resolve this, there is a fixed list of “built-in” notions, and one of the files produced by the accommodator (*art.ere*) specifies what the relation/functor/mode/etc. number of each requirement is. Importantly, if a constructor is identified in this way as a requirement, this not only allows the checker to recognize and produce expressions like $x = y$, it is *also an assertion that this relation behaves as expected*. If a requirement is given a weird definition, for example if we were to open `xcmplx_0.miz` and change the definition of $x + y$ to mean subtraction instead, we would be able to prove false in a downstream theorem which enables the requirement for $+$, because we would still be able to prove $2 + 2 = 4$ by evaluation.

At the surface syntax level, requirements are enabled in groups, using the `requirements` directive in the import section. The requirements are, in rough dependency order:

- `HIDDEN` (introduced after the `HIDDEN` article) is a requirement that is automatically enabled for every Mizar file. It introduces the modes `object` and `set`, as well as $x = y$ and $x \in y$. (The article `HIDDEN` itself is somewhat magical, and cannot be processed normally because every file takes an implicit dependency on `HIDDEN`.)
- `BOOLE` (introduced after `XBOOLE_0`) introduces the adjective `x is empty`, as well as set operators \emptyset , $A \cup B$, $A \cap B$, $A \setminus B$, $A \oplus B$, and A meets B (i.e. $A \cap B \neq \emptyset$). These operators have a few extra properties such as $A \cup \emptyset = A$ that are used in the equalizer.
- `SUBSET` (introduced after `SUBSET_1`) introduces the mode `Element of A`, along with $\mathcal{P}(A)$, $A \subseteq B$, and the mode `Subset of A`. The checker knows about how these notions relate to each other, for example if $x \in A$ then $x : \text{Element}(A)$. (Because types have to be nonempty, $x : \text{Element}(A)$ is actually equivalent to $x \in A \vee (A = \emptyset \wedge x = \emptyset)$. So the reverse implication used by the checker is $(x : \text{Element}(A)) \wedge \neg(A \text{ is empty}) \rightarrow x \in A$.)
- `NUMERALS` (introduced after `ORDINAL1`) introduces `succ(x)`, `x is natural`, the set \mathbb{N} (spelled `NAT` or `omega`), `0`, and `x is zero`. Note that `0` is not considered a numeral in the sense of section 3.1, it is a functor symbol $K_i()$. Numbers other than `0` can be written even before the `ORDINAL1` article, but they are uninterpreted sets; after this requirement is added the system will give numerals like `37` the type `Element of NAT` instead of `set`.
- `REAL` (introduced after `XXREAL_0`) introduces $x \leq y$, `x is positive`, `x is negative`, along with some basic implications of these notions.
 - `NUMERALS + REAL` enables the use of flex-conjunctions $\bigwedge_{i=a}^b \varphi(i)$, since these expand to the expression $\forall i : \mathbb{N}. a \leq i \leq b \rightarrow \varphi(i)$ which requires \mathbb{N} and \leq to write down.
- `ARITHM` (introduced after `XCMPLX_0`) introduces algebraic operators on the complexes: $x + y$, $x \cdot y$, $-x$, x^{-1} , $x - y$, x/y , i , and `x is complex`. This also enables the ability to do complex rational arithmetic on numerals, as well as polynomial normalization.

3.3.1 Soundness considerations of the requirements

There is a slight mismatch between what the user has to provide in order to enable a requirement and what the checker gets to assume when a requirement is enabled which causes a challenge for proof export or other external soundness verification. As the list above might indicate, generally checker modules corresponding to a requirement are enabled as soon as all of the *constructors* involved in stating them are available; for example we can see this with the `NUMERALS + REAL` prerequisite for flex-conjunctions. (More precisely, flex-conjunctions are enabled exactly when \mathbb{N} and \leq become available.) However, the checker needs more than that to justify the manipulations it does with them.

For example the checker exploits the fact that $\bigwedge_{i=1}^3 \varphi(i)$ is equivalent both to $\forall i : \mathbb{N}. 1 \leq i \leq 3 \rightarrow \varphi(i)$ and to $\varphi(1) \wedge \varphi(2) \wedge \varphi(3)$, and so this amounts to an assertion that $a \leq i \leq b \leftrightarrow i = a \vee i = a + 1 \vee \dots \vee i = b$ is provable when a and b are numerals such that $a \leq b$. The reverse implication follows from numerical evaluation, and the forward implication is a metatheorem that can be proven by induction, assuming the existence of lemma L saying $a \leq i \rightarrow i = a \vee \text{succ}(a) \leq i$:

10:10 Reimplementing Mizar

► **Theorem 1.** *If a and n are numerals and $i : \mathbb{N}$, then*

$$i : \mathbb{N} \vdash a \leq i \leq \text{succ}^n(a) \rightarrow i = a \vee i = \text{succ}(a) \vee \dots \vee i = \text{succ}^n(a)$$

is provable.

Proof. By induction on n . Applying the induction hypothesis with $\text{succ}(a)$ and $n - 1$, we get

$$i : \mathbb{N} \vdash \text{succ}(a) \leq i \leq \text{succ}^n(a) \rightarrow i = \text{succ}(a) \vee \dots \vee i = \text{succ}^n(a)$$

so it suffices to show $a \leq i \rightarrow i = a \vee \text{succ}(a) \leq i$, and we appeal to lemma L. ◀

So ideally, when introducing \leq or the **REAL** requirement, one would be required to supply a proof of lemma L somehow to justify that the checker will be making use of this fact in the following article. Unfortunately, there is no actual place to inject this theorem into the system, because there is no concrete syntax for introducing requirements. That is, even taking all the `.miz` files in the MML together there is nothing that would indicate that `XXREAL_0` is the article which allows the **REAL** requirement to be used.³

The way this actually works is that when the accommodator sees a **requirements REAL**; directive, it reads the (hand-written) `real.dre` XML file, which explicitly names the constructor number for \leq and the fact that it is in article `XXREAL_0`. We would like to propose that this file also contains *justifications* for involved constants so that the theorems aren't smuggled in without proof. (Or even better, requirement declarations become a part of the language proper, so that they can get **correctness** proof blocks like any other justified property.)

The case of **NUMERALS + REAL** enabling flex-conjunctions is especially interesting because neither of these requirements depends on the other, so neither `numerals.dre` nor `real.dre` can state the compatibility theorem $a \leq i \rightarrow i = a \vee \text{succ}(a) \leq i$ between them. (This doesn't require a major reorganization to fix, since of course the article that introduces **REAL**, `xxreal_0.miz`, references the article which introduces **NUMERALS**, `ordinal1.miz`.)

4 The analyzer

The analyzer plays an interesting role in Mizar. This is an essentially completely separate inference system from the checker, which has much stricter rules about equality of expressions, and it is what forms the “glue” between different lines of proof. It is a large module only because the language of Mizar is quite expansive, with 109 keywords, including:

- Different kinds of definitions for modes, functors, predicates; redefinitions;
- Definitions by case analysis;
- “Notations” (**synonym** and **antonym** declarations);
- “Properties” like **commutativity**, **projectivity** or **involutiveness**;
- Cluster registrations (existential, functor, and conditional);
 - Existential clusters assert that $\overrightarrow{\chi} \tau$ is nonempty and hence a legal type
 - Functor clusters assert that t is $\overrightarrow{\chi}$ for some term t
 - Conditional clusters assert that $\overrightarrow{\chi}$ implies $\overrightarrow{\chi}'$

³ Note that is also a Mizar article called `real.miz`, which contains some of the lemmas that are auto-proved by the **REAL** requirement, but there is no formal relation between the article and the requirement, and it is only an incomplete approximation to the lemmas required to justify the requirement, not formally checked. This should be easy to fix, and will more or less fall out of any attempt at proof export.

- “Reductions”, equalities the checker automatically uses for simplification;
 - “Identifications”, equalities the checker automatically uses for congruence closure;
 - Local declarations;
 - Schemes;
 - Reservations (variables with types declared in advance);
 - Propositions and theorems;
- and this is not an exhaustive list.

This also only covers top-level items. Inside a **proof** block there is a different (overlapping but largely disjoint) set of legal items, called skeleton steps, which correspond to natural deduction rules. Inside a proof there is a variable that holds the current “thesis”, the goal to prove, and the **thesis** keyword resolves to it, unless one is in a **now** block, where the thesis is not available and is reconstructed from the skeleton steps.

- **let x be T**; is the forall introduction rule: it transforms the thesis from $\forall x : T. \varphi(x)$ to $\varphi(x)$ and introduces a constant $x : T$.
- **assume A**; is the implication introduction rule: it transforms the thesis from $A \rightarrow B$ to B and pushes a proposition A , which can be referred to using **then**.
- **thus A**; is the conjunction introduction rule: it transforms the thesis from $A \wedge B$ to B , and gives A as a goal to the checker.
- **take t**; is the existential introduction rule: it transforms the thesis from $\exists x. \varphi(x)$ to $\varphi(t)$. There is also **take x = t**; which is the same but introduces x as an abbreviation for t ; this version can also be used in a **now** block.
- **consider x being T such that A**; is existential elimination: it introduces $x : T$ and a proposition $A(x)$ that can be labeled, and gives $\exists x : T. A(x)$ as a goal to the checker.
- **given x being T such that A**; is a combination of implication introduction and existential elimination: it transforms $(\exists x : T. A(x)) \rightarrow B$ to B and introduces $x : T$ and a proposition $A(x)$ that can be labeled.
- **reconsider x = t as T**; introduces $x : T$ as a new local constant known to be equal to t , and gives $(t : T)$ as a proof obligation to the checker. (This is mainly used when $t : T$ is not already obvious to the type system, and allows $t : T$ to be proved by the user.)
- **per cases**; is disjunction elimination, and it has two variations:
 - followed by a sequence of **suppose A_i**; ... **end**; blocks, it gives $\bigvee_i A_i$ as a goal to the checker and makes A_i available in each block, leaving the thesis unchanged;
 - followed by a sequence of **case A_i**; ... **end**; blocks, it gives $\bigvee_i A_i$ as a goal to the checker, and if the thesis is $\bigvee_i (A_i \wedge B_i)$ then B_i becomes the thesis in each block.

Additionally, the formulas don’t have to exactly match what the skeleton steps say. For example one can start a proof of $\forall x : \mathbb{N}. \varphi(x)$ using **let x be set**; because **set** is a supertype of **Nat**. Definitional unfolding can also be forced by a skeleton step, for example if the thesis is $A \subseteq B$ then **let x be set**; is a legal step provided the right **definitions** directive is supplied.

The other major role of the analyzer is to elaborate types, terms, and formulas from their input form to the core grammar shown in section 3.1. There are two things that make this challenging:

- Mizar heavily uses overloading, where the same function symbol can have several definitions (and **redefinitions**, which are definitions which use the same base term but can have different input and output types). These are resolved by declaration order (last declaration wins) and typing. Types are propagated exclusively from the inside out, using this type-based overload resolution, although you can use **e qua A** as a type ascription to influence the selection.

- Many declarations have “invisible arguments”, also known as “implicit arguments” in the literature. These are filled in by a straightforward first-order unification process.⁴

5 Mizar soundness bugs

One of the fortuitous side effects of going over each line of code and rewriting it to something morally equivalent in a different language is that one can find a lot of bugs. Bugs can happen even when one takes great efforts to avoid them [1, 8], but external review can definitely help. Mizar is a large project with a long history and a small team, whose source code was not publicly accessible, with many soundness-critical parts, which is pretty much a worst case scenario for finding soundness bugs. This `mizar-rs` project has been tremendously successful in ferreting out these bugs, with no less than four proofs of false that will be given below. These errors have been reported to the Mizar developers, and a patched version is available⁵.

While it is unfortunate that the software wasn’t perfect to start with, this is evidence for the usefulness of external checkers, and it is a way for us to improve the original Mizar. We hope that by telling the story of how these bugs work we can give some sense of some of the issues that can arise when doing proof checking, as well as some more internal details whose importance may not have been obvious.

5.1 Exhibit 1: polynomial arithmetic overflow

This is the largest of the contradiction proofs, and we will show only the main part of it.⁶ The only non-MML notion used is the adjective **a is x-ordered** defined as $x < a$.

theorem contradiction

proof

```

consider x being 1-ordered Nat such that not contradiction;
  1 is 0-ordered; then
    A1: x * x is 1-ordered; then
      consider x1 being 0-ordered Nat such that B1: x1 = x * x;
      A2: 1 < x1 by A1,B1; then
        x1 * x1 is x1-ordered by XREAL_1:155; then :: 0 < a ∧ 1 < b → a < a · b
          consider x2 being x1-ordered Nat such that B2: x2 = x1 * x1;
          ...
          consider x31 being x1-ordered Nat such that B31: x31 = x30 * x30;
          consider x32 being x1-ordered Nat such that B32: x32 = x31 * x31;
          C: x32 * x1 = x1 by
            B1, B2, B3, B4, B5, B6, B7, B8, B9, B10, B11, B12, B13, B14, B15, B16,
            B17, B18, B19, B20, B21, B22, B23, B24, B25, B26, B27, B28, B29, B30, B31, B32;
          x32 is x1-ordered implies x1 < x32;
          then 0 < x1 & 1 < x32 by A2, XREAL_0:2; :: a ≤ b ∧ b ≤ c → a ≤ c
          hence contradiction by C, XREAL_1:155; :: 0 < a ∧ 1 < b → a < a · b
end;

```

⁴ Unfortunately, the unification process does not have completely unique solutions, because attributes are unordered, and this can lead to overfitting in the MML to the Mizar attribute ordering. For example, `rng(f) ⊆ B` returns the range of a function $f : A \rightarrow B$ as an element of `Subset of B` (and `rng` has hidden arguments `rngA,B(f)` inferred from the types). But the function type is split into separate attributes as `A-defined B-valued Function`, so if `f` is `B-valued C-valued Function` then it is up to a variety of implementation details whether you get the type of `rng f` as `Subset of B` or `Subset of C`. This happens in practice for the empty function, which is registered as `NAT-valued`, `RAT-valued`, and a few others.

⁵ <https://github.com/digama0/mizar-system>

⁶ The full proof can be found at

<https://github.com/digama0/mizar-rs/blob/itp2023/itp2023/false1/false1.miz>.

As mentioned earlier, there is a module in the equalizer for polynomial evaluation. This means that each expression which is a complex number will also be expressed as a polynomial in terms of basic variables, and two expressions which compute to equal polynomials will be equated. This is how things like $-(a + b) = -a + -b$ are proved automatically by the checker.

These polynomials are of the form $\sum_i c_i \prod_j x_{i,j}^{n_{i,j}}$, and the starting point for this proof was the discovery that the $n_{i,j}$ are represented as signed 32-bit integers and overflow is not checked. Turning this into an exploit is surprisingly difficult however, because one cannot simply write down $x^{2^{32}}$ because the power function is not one of the requirements (see section 3.3). The best thing we have for creating larger polynomials is multiplication, but we can get to $x^{2^{32}}$ with 32 steps of repeated squaring, which is what the large block of **consider** statements is doing.

So the strategy is to construct $x_1 = x^2$, $x_2 = x^4$, all the way up to $x_{30} = x^{2^{30}}$. After this things start to get weird: $x_{31} = x^{-2^{31}}$ because of signed overflow, and $x_{32} = x^0$. The system does not recognize $x_{32} = 1$ however, because it maintains an invariant of monomial powers being nonzero, and single powers are also handled specially, so we have to go up to $x_{32} \cdot x_1 = x^2$. Since x_1 is also x^2 , the system will equate $x_{32} \cdot x_1 = x_1$ which is the key step **C**. After this, we simply need to separately prove that since $x_{32} \cdot x_1$ is really $x^{2^{32}+2}$, it is strictly larger than $x_1 = x^2$ as long as we choose $x > 1$.

To prove this last fact we use the attribute inference mechanism to prove that all of the intermediates are strictly greater than x_1 , since $x_1 < a$ implies $x_1 < a \cdot a$ as long as $x_1 > 1$. Hence $1 < x_1 < x_{32}$ so $x_1 < x_{32} \cdot x_1 = x_1$, which is a contradiction.

5.2 Exhibit 2: negation in the schematizer

This one requires absolutely no imports; we include the complete proof below including the import section.

```

environ begin

scheme Foo{P[set,set]}: P[1,1] implies P[1,1]
proof thus thesis; end;

theorem contradiction
proof
  1 = 1 implies 1 <> 1 from Foo;
  hence thesis;
end;

```

This defines a very trivial scheme which just says that $P(1,1)$ implies $P(1,1)$. The interesting part is the instantiation of this scheme in the main proof. In the schematizer, we do not negate the thesis and try to prove false, we just directly match the thesis against the goal. It does not attempt any fancy higher-order unification: if it needs to unify $P(\vec{t}_i) \stackrel{?}{=} \pm R(\vec{t}'_i)$ it will just assign $P := \pm R$ and proceed with $t_i \stackrel{?}{=} t'_i$ for each i . Or at least it should do that, but there is a bug wherein it instead assigns $P := R$ and ignores the \pm part, so we can trick it into unifying $P(x,y) := (x = y)$ and then think that $P(x,y) \stackrel{?}{=} (x \neq y)$ is still true. So we use $1 = 1$ to prove $1 \neq 1$ and then prove a contradiction.

5.3 Exhibit 3: flex-and unfolding

10:14 Reimplementing Mizar

```
theorem contradiction
proof
  (1 = 1 or 1 = 1) & ... & (2 = 1 or 1 = 1);
  hence thesis;
end;
```

This one is bewilderingly short. We start by asserting a (true) flex-and statement $\bigwedge_{i=1}^2 (i = 1 \vee 1 = 1)$. We prove this by using the forall expansion of the flex-and, $\forall i : \mathbb{N}. 1 \leq i \leq 2 \rightarrow i = 1 \vee 1 = 1$, which is of course true because the $1 = 1$ disjunct is provable.

The second part is simply **hence contradiction**, so we are calling the checker to *disprove* the same statement. So we assume $\bigwedge_{i=1}^2 (i = 1 \vee 1 = 1)$ and one of the things the pre-checker does here is to expand out the conjunction, and we would expect it to produce $(1 = 1 \vee 1 = 1) \wedge (2 = 1 \vee 1 = 1)$ from which no contradiction can be found. However, what it actually produces is $1 = 1 \wedge 2 = 1$, which can be disproved.

To see why this happens, we have to look more specifically at the forall-expansion without the negation sugar employed thus far. What Mizar actually sees for the expansion of the flex-and expression is $\forall i : \mathbb{N}. \neg(1 \leq i \wedge i \leq 2 \wedge (i \neq 1 \wedge 1 \neq 1))$, except that as mentioned previously conjunctions are always flattened into their parents. The code for expanding flex-and expects the expansion body to be the third conjunct past the forall and negation, but after flattening the third conjunct is actually $i \neq 1$ and there is an unexpected fourth conjunct $1 \neq 1$ that is forgotten.

5.4 Exhibit 4: flex-and substitution

This is the most worrisome of the bugs that have been presented, because it is not simply a bad line of code but rather an issue with the algorithm itself. As a result, this one survived the translation to Rust and was discovered to affect both versions, and moreover it has still not been fixed in `mizar-rs`, because the fix could not be rolled out without breaking the MML. (There is an “unsound flag” which controls whether to do the thing that is unsound or the thing that is sound but fails to validate the MML.)

```
theorem contradiction
proof
  A: now
    let n be Nat;
    assume n > 3 & (1 + 1 <> 3 & ... & n <> 3);
    hence contradiction;
  end;
  ex n being Nat st n > 3 & (1 + 1 <> 3 & ... & n <> 3)
proof
  take 2 + 2;
  thus 2 + 2 > 3 & (1 + 1 <> 3 & ... & 2 + 2 <> 3);
end;
hence thesis by A;
end;
```

To explain what is happening here, we have to talk about another aspect of flex-and expressions which was not mentioned in section 3.1, which is that a flex-and expression $\bigwedge_{i=a}^b \varphi(i)$ is actually stored as five pieces of information: the bounds a and b , the expansion

$\forall i : \mathbb{N}. a \leq i \leq b \rightarrow \varphi(i)$, from which $\varphi(i)$ can be reconstructed (if one is careful – see section 5.3), and the bounding expressions $\varphi(a)$ and $\varphi(b)$. As the reader may have noticed, the syntax of Mizar very much prefers to only discuss the bounding expressions, since flex-and expressions in the concrete syntax are written as $\varphi(a) \wedge \dots \wedge \varphi(b)$ with literal dots, and no place to supply a , b or $\varphi(i)$.

When one writes an expression like $P \wedge \dots \wedge Q$, Mizar essentially diff's the two expressions to determine what is changing and what the values are on each side. So if one writes $1 + 1 \neq 3 \wedge \dots \wedge n \neq 3$ then Mizar infers that the desired expression is $\bigwedge_{i=1+1}^n i \neq 3$.

The problem is that this operation is not stable under substitution. If we take that expression and substitute $n := 2 + 2$, then we end up with $\bigwedge_{i=1+1}^{2+2} i \neq 3$, but if we were to write out $1 + 1 \neq 3 \wedge \dots \wedge 2 + 2 \neq 3$ ourselves we would end up with $\bigwedge_{i=1}^2 i + i \neq 3$ instead. This is somewhat annoying, especially if one is trying to write an exploit example, but it's not obviously a soundness bug yet because internally we are still storing the expansion which has all the details – $\bigwedge_{i=1+1}^{2+2} i \neq 3$ and $\bigwedge_{i=1}^2 i + i \neq 3$ are different expressions.

However, the equality check between two flex-and expressions is simply $\varphi(a) = \varphi'(a)$ and $\varphi(b) = \varphi'(b)$! The example above is crafted to demonstrate that this is not sound in general, since $\bigwedge_{i=1+1}^{2+2} i \neq 3$ is false but $\bigwedge_{i=1}^2 i + i \neq 3$ is true.

In the first part of the proof we show that $\forall n. \neg(n > 3 \wedge \bigwedge_{i=1+1}^n i \neq 3)$, which is true. The interesting part is the second half, where we take $n := 2 + 2$. At this point the thesis is $2 + 2 > 3 \wedge \bigwedge_{i=1+1}^{2+2} i \neq 3$, but we cannot write this expression. What we write instead is $2 + 2 > 3 \wedge (1 + 1 \neq 3 \wedge \dots \wedge 2 + 2 \neq 3)$, which as mentioned will be interpreted as $2 + 2 > 3 \wedge \bigwedge_{i=1}^2 i + i \neq 3$. This should be rejected for not matching the thesis, but it has the same endpoints so it is accepted, and the checker is able to prove it by case analysis. If we were to write **thus thesis** to let Mizar insert the unmentionable proposition, then the analyzer will accept it as being the right thesis but the checker will not be able to prove it.

The reason this is a hard bug to squash is because the MML actually relies on it; there are several proofs in `aofa_100.miz` that break when strict flex-and checking is enabled. All of the examples found seem to be variations on $(\bigwedge_{i=a+c}^{b+c} \varphi(i)) \leftrightarrow (\bigwedge_{i=a}^b \varphi(i + c))$, which is true, so there is some hope that a more subtle check can be used to fix the soundness issue without making the prover much worse.

5.5 Honorable mention: attributes that don't exist

This is not an exploitable bug to my knowledge, but it is notable for being widespread, and it is difficult for `mizar-rs` to support without resulting in very weird behavior. Consider the following Mizar article:

```

environ
  vocabularies ZFMISC_1, SUBSET_1;
  notations ZFMISC_1, SUBSET_1;
  constructors TARSKI, SUBSET_1;
  requirements SUBSET; ::, BOOLE;
begin
for x, B being set, A being Element of bool B st x in A holds x in B;

```

This checks as one would expect, and it is in fact a true statement (and it is not at all contrived). However, the reasoning that gets the checker to accept the proof is... somewhat suspicious. It goes as follows:

1. Suppose $A : \text{Element}(\mathcal{P}(B))$, $x \in A$ and $x \notin B$.
2. Because $A : \text{Element}(\mathcal{P}(B))$ and $x \in A$, it follows that B is not empty and $x : \text{Element}(B)$.
3. Because $x \notin B$, B is not empty, and $x : \text{Element}(B)$, contradiction.

10:16 Reimplementing Mizar

■ **Table 1** Comparison of the original (Pascal) implementation of Mizar with `mizar-rs` (Rust). Tests were performed on an 8 core 11th Gen Intel Core i7-1165G7 @ 2.80GHz, with 8 threads.

	PC Mizar	<code>mizar-rs</code>	ratio
lines, checker	30 458	14 474	2.10
lines, checker + analyzer	56 926	20 096	2.83
compiling MML, checker only, real time	57.37 min	11.33 min	5.06
compiling MML, checker only, CPU time	417.57 min	73.70 min	5.67
compiling MML, checker + analyzer, real time	71.38 min	11.71 min	6.10
compiling MML, checker + analyzer, CPU time	490.55 min	81.93 min	5.99

Because the example has been minimized, there is really not much going on in the proof because these are all essentially primitive inferences. The problem here is “ B is not empty”, because the `BOOLE` requirement which supplies the `B is not empty` predicate is not available (note the environment section). The constructor for `empty` is actually available in the environment because it has been brought in indirectly via the `constructors SUBSET_1` directive, but without the requirement the checker just sees it as a normal attribute.

So what, then, is the checker doing? This seems to be a case of multiple bugs cancelling each others’ effects. Requirements are internally represented by an integer index, where zero means that the requirement is not available. Normally any handling of a requirement involves a check that the requirement is nonzero first, but we forget that in step 2, and as a result we end up adding attribute 0 to B , which is meaningless. The second bug is in step 3, where we again forget to ask whether the requirement index is nonzero, and so we find attribute 0 and interpret it (correctly) as meaning that B is not empty.

This would just be a curious bug, but for the fact that it is exploited all over the place because when creating articles it is standard to minimize the environment section by removing anything that keeps the proof valid, and this bug allows one to remove the `BOOLE` requirement without breaking the proof. We had to patch 46 articles that all had this same issue. In all cases we only need to add the `BOOLE` requirement to fix the issue.

6 Results

The whole project is 20 096 lines of code (see table 2), or 14 474 if we restrict attention to those files used in the checker, which is 1/2 to 1/3 of the equivalent code in Pascal (see table 1). We credit this mainly to the language itself – Rust is able to express many complex patterns that would be significantly more verbose to write in Pascal, and PC Mizar is also written in a fairly OOP-heavy style that necessitates a lot of boilerplate. Furthermore, Pascal is manually memory-managed while Rust uses “smart-pointer” style automatic memory management, so all the destructors simply don’t need to be written which decreases verbosity and eliminates many memory bugs.

The performance improvements are most striking: we measured a 5-6 \times reduction in processing time to check the MML. This is most likely a combination of characteristics of the LLVM compiler pipeline, along with many small algorithmic improvements and removing redundant work. (The reduction is even larger when adding the other phases of the Mizar system – accom, parser, MSM – so that we can skip the costly I/O and serialization steps of section 2.)

■ **Table 2** Line counts for the files in `mizar-rs`. Files in red are only required for the analyzer, not the checker.

file	lines
<code>analyze.rs</code>	3471
<code>main.rs</code>	2747
<code>equate.rs</code>	2617
<code>types.rs</code>	2114
<code>parser/msm.rs</code>	1489
<code>unify.rs</code>	1346
<code>parser/mod.rs</code>	1280
<code>checker.rs</code>	1083
<code>reader.rs</code>	890
<code>parser/article.rs</code>	759
<code>ast.rs</code>	662
<code>equate/polynomial.rs</code>	586
<code>format.rs</code>	581
<code>bignum.rs</code>	372
<code>util.rs</code>	99
total	20096

7 Conclusion & Future work

We have implemented a drop-in replacement for the `verifier -c` checker of the Mizar system that is able to check the entire MML, which gets a significant performance improvement. Furthermore, we have improved Mizar by uncovering and reporting some soundness bugs.

While this implementation is explicitly trying not to diverge from the Mizar language as defined by the PC Mizar implementation and the MML, there are many possible areas where improvements are possible to remove unintentional or undesirable restrictions. For example, anyone who has played with Mizar will have undoubtedly noticed that all the article names are 8 letters or less, for reasons baked deeply into PC Mizar. Should `mizar-rs` become the official Mizar implementation, it would be possible to lift this restriction without much difficulty.

This project was also originally started to get proof export from Mizar, and to that end replacing one large trusted tool with another one does not seem like much of an improvement. But (bugs notwithstanding) we saw nothing while auditing the checker that is not proof-checkable or unjustified, and remain confident that proof export is possible.

References

- 1 Mark Miles Adams. Proof auditing formalised mathematics. *Journal of Formalized Reasoning*, 9(1):3–32, 2016.
- 2 Grzegorz Bancerek, Czesław Byliński, Adam Grabowski, Artur Korniłowicz, Roman Matuszewski, Adam Naumowicz, and Karol Pąk. The role of the Mizar Mathematical Library for interactive proof development in Mizar. *Journal of Automated Reasoning*, 61:9–32, 2018.
- 3 Grzegorz Bancerek, Czesław Byliński, Adam Grabowski, Artur Korniłowicz, Roman Matuszewski, Adam Naumowicz, Karol Pąk, and Josef Urban. Mizar: State-of-the-art and beyond. In *Intelligent Computer Mathematics: International Conference, CICM 2015, Washington, DC, USA, July 13–17, 2015, Proceedings.*, pages 261–279. Springer, 2015.

10:18 Reimplementing Mizar

- 4 Mario Carneiro. The Divergence of the Sum of Prime Reciprocals. *Formalized Mathematics*, 30(3):209–210, 2022. doi:doi:10.2478/forma-2022-0015.
- 5 Adam Grabowski, Artur Kornilowicz, and Adam Naumowicz. Mizar in a nutshell. *Journal of Formalized Reasoning*, 3(2):153–245, 2010.
- 6 John Harrison. A Mizar Mode for HOL. In Joakim von Wright, Jim Grundy, and John Harrison, editors, *Theorem Proving in Higher Order Logics: 9th International Conference, TPHOLS'96*, volume 1125 of *Lecture Notes in Computer Science*, pages 203–220, Turku, Finland, 1996. Springer-Verlag.
- 7 Cezary Kaliszyk and Karol Pąk. Semantics of Mizar as an Isabelle object logic. *Journal of Automated Reasoning*, 63:557–595, 2019.
- 8 Ondřej Kunčar and Andrei Popescu. A consistent foundation for Isabelle/HOL. In *Interactive Theorem Proving: 6th International Conference, ITP 2015, Nanjing, China, August 24-27, 2015, Proceedings 6*, pages 234–252. Springer, 2015.
- 9 Roman Matuszewski and Piotr Rudnicki. Mizar: the first 30 years. *Mechanized mathematics and its applications*, 4(1):3–24, 2005.
- 10 Michal Muzalewski. *An outline of PC Mizar*. Fondation Philippe le Hodey, 1993.
- 11 Josef Urban. XML-izing Mizar: Making Semantic Processing and Presentation of MML Easy. In Michael Kohlhase, editor, *Mathematical Knowledge Management, 4th International Conference, MKM 2005, Bremen, Germany, July 15-17, 2005, Revised Selected Papers*, volume 3863 of *Lecture Notes in Computer Science*, pages 346–360. Springer, 2005. doi:10.1007/11618027_23.
- 12 Josef Urban. MPTP 0.2: Design, implementation, and initial experiments. *Journal of Automated Reasoning*, 37:21–43, 2006.