

# Proof Patterns for Formal Methods

Leo Freitas and Iain Whiteside

School of Computing Science, Newcastle University, U.K.  
{\*.\*}@newcastle.ac.uk

**Abstract.** *Design patterns* represent a highly successful technique in software engineering, giving a reusable ‘best practice’ solution to commonly occurring problems in software design. Taking inspiration from this approach, this paper introduces *proof patterns*, which aim to provide a common vocabulary for solving formal methods proof obligations by capturing and describing solutions to common patterns of proof.

**Keywords:** proof pattern, formal verification, proof obligations

## 1 Introduction

A key advantage of formal specification of software systems using mathematical models is a precise characterisation of the requirements that is amenable to scrutiny through deductive reasoning methods. Popular *formal methods* languages like Z, B, and VDM [1,25,36], share a similar methodology: a software system is modelled at an abstract level and refined stepwise to a more concrete representation by infusing the model with design decisions and concrete datatypes. In each level of ‘refinement’ a system state is usually described as an abstract datatype or record with an *invariant* attached. Operations of the system are described with preconditions and postconditions that may modify the system state. The correctness guarantees provided by ‘formal methods’ stem from the *proof obligations (POs)* that are generated for any given model. Within formal methods, POs tend to have a predictable shape. Furthermore, the process of model design tends to exhibit predictable shapes. This repetition in the phrasing of theorems and in the solution to particular modelling problems suggests the possibility of repeated proofs. This notion of repetitive proof has been corroborated in practice from our personal experiences [11,15,16].

One of the main challenges for the widespread adoption of formal methods in industry is the expense, both in time and human expertise, of solving these proof obligations. For example, industrial partners using B claim each PO costs 38 euros! In itself this is not a show stopper, until one realises that there are over 11,000 proof obligations to discharge<sup>1</sup>. Our aim is, given expert proofs of more difficult lemmas together with (meta-)proof data collected with our tools, to

---

<sup>1</sup> These numbers were reported at a private meeting with colleagues from a French company dealing with a (unpublished) system involving railways.

## 2. BACKGROUND

enable less experienced proof engineers to identify remainder common situations and tackle them effectively by the use of proof patterns, hence reducing the overall proof cost and effort, both of which are usually the argument against formal methods application in industry.

In Software Engineering, *design patterns* are a highly successful technique for providing solutions to frequently occurring problems in software design [19]. Patterns give a template description of how to solve a particular problem that can be used in many different situations. Patterns have evolved to become best practices that must be implemented and are easily recognised when providing a common vocabulary for describing solutions. Similarly, mathematicians do not often talk about the details of their proofs: they have a common parlance of high-level proof patterns (e.g. by induction,  $\epsilon$ - $\delta$  proofs in analysis) that enables fellow mathematicians to understand and recreate proofs. It is our aim in this paper to introduce *proof patterns for formal methods*. Specifically, we hypothesise:

*Similar to software engineering design patterns, proof patterns exist for formal methods proof obligations. Furthermore, under the right circumstances, they are applicable over multiple methods.*

This paper presents several proof patterns useful for discharging proof obligations across formal methods. We describe these patterns similarly to software design patterns, and although we use a concrete example problem for explaining the patterns, we believe that they transfer across problems, and sometimes across provers. We provide examples of each pattern in action using a VDM model of a heap memory manager [27, Ch. 7], which we have formalised in the Isabelle and Z/EVES theorem provers [30,31]. We believe that too small an example is unlikely to clarify the issues with patterns in proof obligations, and that industrial examples do not fit in a paper (yet are amenable to our patterns). For instance, we drew on our experience from proofs of the Software Verification Grand Challenge [24] pilot projects [9,12,14,15,16,17,18,37] to identify, summarise and categorise the patterns presented here.

In the next section, we briefly introduce software design patterns, formal methods, and our running example. Then, Section 3 describes our core proof patterns with examples of each in action. We further exemplify our proof patterns with a worked example on a feasibility proof obligation in Section 4. Finally, we conclude with related and future work in Section 5.

## 2 Background

### 2.1 Patterns

Like with design patterns for software engineering [19], we see proof patterns as a combination of informal description, examples of use, and a set of attributes of discourse explaining the conditions for which the proof pattern may apply. A classic example of a design pattern is the *iterator pattern*.

We are trying to do for proof what design patterns did for software development: create a discourse of ideas and processes that might omit the specifics of “how”, unless one is happy to look at the gory details (*i.e.* not only of large proofs, but of large amounts of failed proof attempts). We try to give precise and accountable details as much as possible, yet we are still some way from having an expressive enough declarative proof language to capture proof intent.

## 2.2 Heap problem

In this section, we provide an overview of a heap memory manager, modelled in VDM [27, Ch.7], that we use throughout this paper to exemplify our proof patterns. We have formalised and proved all proof obligations associated with the first two levels of refinement for this model in the Isabelle [30,29] and Z/EVES [31] theorem provers. A full description of this formalisation, including a detailed description of the translation between a VDM model and its representation in Isabelle, can be found in [11]. The model consists of two datatypes and two operations:

**Loc:** the type of a single adjacent memory location, represented as  $\mathbb{N}$ .

**Free:** the type of the heap as a collection of all free locations. At *level 0*, it is represented as the set  $Free0 \triangleq Loc\text{-set}$ , whereas at *level 1*, it is represented as a map from start location to size that is *disjoint* and *separate*:

$$\begin{aligned} Free1 &= Loc \xrightarrow{m} \mathbb{N}_1 \\ \mathbf{inv} (f) &\triangleq \mathit{disj}(f) \wedge \mathit{sep}(f) \\ \mathit{disj}(f) &\triangleq \forall l, l' \in \mathbf{dom} f \cdot l \neq l' \longrightarrow \mathit{locs-of}(l, f(l)) \cap \mathit{locs-of}(l', f(l')) = \phi \\ \mathit{sep}(f) &\triangleq \forall l \in \mathbf{dom} f \cdot (l + f(l)) \notin \mathbf{dom} f \end{aligned}$$

The invariant conditions ensure that the range of locations identified by any two map elements (defined as  $\{l \dots l + f(l) - 1\}$  by  $\mathit{locs-of}$ ) do not intersect ( $\mathit{disj}$ ) and that contiguous memory regions are as large as possible ( $\mathit{sep}$ ). That is, for any element  $l$  in the map, the location immediately to the right of its memory region ( $l + f(l)$ ) is not the start location for another region:  $l + f(l) \notin \mathbf{dom} (f)$ . We write  $F1\text{-inv}(f)$  to refer to this invariant of  $Free1$ .

**NEW:** takes a size and heap as input and returns a starting location for a contiguous chunk of memory of the appropriate size after updating the state.

**DISPOSE:** returns a contiguous chunk of memory back to the heap. This operation takes a start location and size as parameters, updating the state.

At level 0, these operations are defined as

$$\begin{array}{ll} NEW0 (s: \mathbb{N}_1) r: Loc & DISPOSE0 (l: Loc, s: \mathbb{N}_1) \\ \mathbf{ext} \ \mathbf{wr} \ f_0 : Free0 & \mathbf{ext} \ \mathbf{wr} \ f_0 : Free0 \\ \mathbf{pre} \ \exists l \in Loc \cdot \mathit{locs-of}(l, s) \subseteq f_0 & \mathbf{pre} \ \mathit{locs-of}(l, s) \cap f_0 = \{ \} \\ \mathbf{post} \ \mathit{locs-of}(r, s) \subseteq \overleftarrow{f_0} \wedge & \mathbf{post} \ f_0 = \overleftarrow{f_0} \cup \mathit{locs-of}(l, s) \\ \quad f_0 = \overleftarrow{f_0} - \mathit{locs-of}(r, s) & \end{array}$$

Set difference and set union characterise the removal and addition of elements to the heap. The precondition on  $NEW0$  ensures that there is a contiguous region

### 3. PROOF PATTERNS

of the appropriate size, whereas *DISPOSE0* ensures the range of locations being returned is not already free. At level 1, the *NEW* operation is:

```

NEW1 (s: ℕ1) r: Loc
  ext wr f1 : Free1
  pre ∃l ∈ dom f1 · f1(l) ≥ s
  post r ∈ dom  $\overleftarrow{f_1}$  ∧ ( $\overleftarrow{f_1}(r) = s$  ∧ f1 = {r} ⇐  $\overleftarrow{f_1}$  ∨
     $\overleftarrow{f_1}(r) > s$  ∧ f1 = ({r} ⇐  $\overleftarrow{f_1}$ ) ∪m {r + s ↦  $\overleftarrow{f_1}(r) - s$ )

```

*NEW1* has two behaviours depending on whether a location of exactly the required size or larger has been located. If the size matches, then that element is removed from the map; if the map element refers to a larger region, then the remaining locations in the region must be added back to the heap (hence the map union). The precondition captures both cases using  $\geq$ . We describe *DISPOSE1* in Section 4 as part of a worked example of solving a feasibility proof obligation with proof patterns. The *retrieve* between these two levels of refinement is given by function  $f_0 = locs(f_1)$ : it generalises *locs-of* over the domain of  $f_1$  using distributed union (i.e.  $locs = \bigcup \{locs\text{-of}(x, f_1(x)) \mid x \in \mathbf{dom} f_1\}$ ).

## 3 Proof patterns

In the following sections we describe our proof patterns. Despite the specificity of the Heap example, these patterns apply in most problems of interest within formal methods POs. Wherever patterns are specific enough, we provide proof snippets in Isar: a human readable formal proof language for Isabelle [35].

### 3.1 Witnessing

One of the most important steps to solve feasibility (and reification) POs is finding appropriate witnesses for the outputs and updated state. The general form of these POs is  $\forall \overleftarrow{\sigma}, \bar{i}. pre\text{-}OP(\overleftarrow{\sigma}, \bar{i}) \longrightarrow \exists \bar{\sigma}, \bar{o}. post\text{-}OP(\overleftarrow{\sigma}, \bar{i}, \bar{\sigma}, \bar{o})$ , where  $\overleftarrow{\sigma}$  and  $\bar{i}$  are the initial state and inputs and  $\bar{\sigma}$  and  $\bar{o}$  represent a sequence of updated state variables and outputs. In general, finding a witness is a difficult task, but there are two common patterns that allow some of the existentials to be discharged easily and a third pattern to help the engineer “discover” the right unknown witness.

*One point.* Often the value of an updated state variable is given explicitly in the postcondition as an equality,  $\sigma_i = t$ , where  $t$  is an expression in terms of the initial state. This can be discharged by a generalised version of the “one-point” proof rule in [36, Sect. 4.2], which Z/EVES implements and could be encoded in Isabelle, though at present we discharge it manually. The rule avoids providing any explicit instantiation, regardless of variable order. If  $t$  is a complicated expression, one may wish to introduce an informative name  $x' = t$  into the assumptions and use the one-point rule with  $x'$  instead, since  $t$  is substituted for  $\sigma_i$  everywhere it is used in the postcondition.

### 3. PROOF PATTERNS

*Existential precondition.* A more subtle situation involves an existential precondition (e.g. *pre-OP*  $\triangleq \exists x \cdot P x$ ). This often means that  $x$ , once eliminated, is ‘supposed to be’ mapped to a particular witness. Notice that inputs to preconditions can also be viewed as existentially quantified assumptions and are also suitable. This pattern occurs when a nondeterministic choice for the postcondition is required and a value satisfying the precondition can be picked. This means that we can often match an “existential” precondition variable by trying to match the  $P x$  from the precondition within the postconditions.

*Dummy.* When a witness guess is unknown, making progress in the proof serves to clarify the appropriate choice. To make progress, one should instantiate the existential quantifier with an arbitrary variable, then proceed. Once the goal is rewritten taking the witness into account, resulting subgoals can be analysed for evidence pointing to appropriate instantiations.

*Example.* In the heap case study, both one-point and existential precondition patterns occur in the feasibility PO for *NEW1* and *DISPOSE1*. After performing case analysis, the first PO is as follows:

$$\begin{aligned} & \forall f s \cdot F1\text{-inv}(f) \wedge (\exists l \in \mathbf{dom} f \cdot f(l) = s) \longrightarrow \\ & (\exists f' r' \cdot F1\text{-inv}(f') \wedge r' \in \mathbf{dom} f \wedge f(r') = s \wedge f' = \{r\} \triangleleft f) \end{aligned}$$

The witness for  $f'$  can be found using the one-point rule for  $f' = \{r\} \triangleleft f$ . This entails some existential introduction massaging, which Z/EVES does automatically and we encode in Isabelle. Notice we cannot use  $r$  in the witness as it is also being quantified. The witness for  $r$  is the  $l$  introduced by existential elimination on the precondition. The witnessing pattern reduces this feasibility proof obligation to showing that the invariant holds on the updated state as  $F1\text{-inv}(\{l\} \triangleleft f)$ , which is ready for invariant breakdown as described in the next section.

#### 3.2 Invariant breakdown

One often needs to show that the updated state preserves the invariant, as in the example above. When the updated state is defined in terms of the original state, we move the invariant predicate towards the original state in order to use the assumption held by the original invariant. This can be seen as a specialised form of the *rippling* proof plan for solving step cases of induction proofs [5]. Rather than delve into the details of rippling, we explain this pattern as an operation on invariant proofs.

The situation that triggers this pattern is as follows. We need to solve  $inv(h(\sigma))$  where we know  $inv(\sigma)$  holds and the updated state is  $h(\sigma)$ . As a simple example, imagine the state is a set of natural numbers  $X$  and the updated state is  $f(X) \cap g(X)$ , giving us a proof obligation  $inv(X) \longrightarrow inv(f(X) \cap g(X))$ . To apply the invariant breakdown pattern, we aim to move the invariant predicate *closer* to the original invariant terms: distributing it over set intersection in this case. Furthermore, the subterms that contain the original state would be generalised to an arbitrary element. This means speculating a lemma:  $P \longrightarrow inv(A) \longrightarrow inv(B) \longrightarrow inv(A \cap B)$ , where  $P$  expresses side-conditions

### 3. PROOF PATTERNS

under which the lemma must hold. The application of this lemma leads to  $inv(f(X))$  and  $inv(g(X))$  as new subgoals, we need to apply the invariant breakdown pattern again until we get  $inv(X)$  itself as assumptions to the speculated lemma. While the invariant breakdown pattern does not solve the goal, it provides a set of lemmas that, if proven, will lead to a proof of the top-level goal. We call them weakening lemmas (see Section 3.3) and they are available to function symbols either unknown or with little automation.

The process of discovering such side-conditions  $P$  is non-trivial and requires expertise. It can also be helped, however, by model-checking and counter-example checking: take  $P$  as **true** and run a counterexample checker, for example. Another source of useful information is the definitions of the operators involved, in the case above  $\{inv, f, g, \cap\}$ . Together with counter-examples found, they expose the clues to the appropriate side-conditions, which might themselves be lemmas to be proved. In the worst case, when side-conditions are difficult to guess, one might need to create a new concept specific to the domain of the problem.

*Example.* In the heap, this pattern was used frequently, since the two components of the invariant (*separateness* and *disjointedness*) needed to be proved for the updated state in both the *NEW* and *DISPOSE* operations. In *NEW1*, for example, we are required to show that the following goal holds.

$$Disjoint(\{\{r\} \Leftarrow \underline{f_1}\} \cup_m \{r + s \mapsto \underline{f_1}(r) - s\})$$

where we have a single occurrence of the original state  $f_1$  *as itself*, which gives us some indication of how to break down this formula by distributing *Disjoint* over map union ( $\cup_m$ ). Generalising, we speculate a lemma:

$$Q \longrightarrow Disjoint(f) \longrightarrow Disjoint(\{a \mapsto b\}) \longrightarrow Disjoint(f \cup_m \{a \mapsto b\})$$

Ignoring side-conditions for the moment, this gives us the following subgoals:

$$Disjoint(\{\{r\} \Leftarrow \underline{f_1}\}) \quad Disjoint(\{r + s \mapsto \underline{f_1}(r) - s\})$$

The second goal does not contain the original state, and can be solved trivially. For the first goal, we need to repeat the invariant breakdown process on domain filtering ( $\Leftarrow$ ). This allows us to solve the goal using the assumption.

#### 3.3 Weakening lemmas

When one does not have enough information about function symbols appearing in invariant breakdown, we need weakening lemmas relating these symbols.

For instance, POs over complex states often include records and data structures the prover knows little about (e.g. a map from a record to a list). Naively dealing with the presence of these (novel use of) symbols often leads to either polluted (and repetitive) proofs, or to overly specific lemmas. Instead, we need lemmas that weaken specific parts of the goal (for backward reasoning) or specific parts of the hypothesis (for forward reasoning). This breaks down the task to manageable pieces, up to the point where the prover has automation for function symbols involved, as in the example above involving  $inv(\_ \cap \_)$ .

When discharging proof obligations for the heap we came across surprising points of failure. The refinement and feasibility POs for *NEW* and *DISPOSE*

### 3. PROOF PATTERNS

motivated the creation of lemmas for both Z/EVES and Isabelle, which are fully documented in [11].

*Isabelle proofs.* From the pattern for the feasibility proof, we provide weakening lemmas to enable automation on our given witnesses. This unpicking of the various parts of the feasibility proof obligation leads to the suggestion of lemma shapes up to the point where available lemmas apply. Once these linking (weakening) lemmas are in place, Isabelle knows enough about involved operators and can automatically discharge them. Bridging this gap is where expert input is needed. For example, to represent VDM maps in Isabelle we use  $(Loc \rightarrow \mathbb{N})$ , which is a total function with an optional range type. Isabelle has useful operators like map update ( $\dagger$ ), yet not map union ( $\cup_m$ ), which we define as map update over maps with disjoint domains. Map union features in proofs for *NEW1* and *DISPOSE1*, and we needed lemmas linking the new operator to a representation known to Isabelle, namely map update. Thus, while performing *invariant breakdown*, we identify the need for weakening lemmas about distributing VDM operators over invariant subparts. Similar lemmas for other function symbols were also added, where different side-conditions determine where such weakening lemmas can be used. In total we have 51 of them for Isabelle.

*Z/EVES proofs.* Arguably, given Z already contains a mathematical toolkit akin to VDM's, it is easier to represent a VDM model in Z/EVES. This means many of the weakening lemmas we need are already available. The lesson from this though is that weakening lemmas over known (or reused) function symbols do indeed transfer across problems, and that is our experience. Some of the Isabelle weakening lemma were informed by previous experience with Z/EVES.

We declare weakening lemmas in Z/EVES as a rewrite rule, which the prover uses automatically during simplification and is akin to Isabelle's `simp` attribute on lemmas. These lemmas are not quite solving, but distilling the problem through the proof engineering process described by our proof patterns. These kinds of lemmas are not usually transferable across problems, yet the general principles/patterns behind them are, as our experience with the GC pilot project experiments shows [9,12,14,15,16,17,18,37].

#### 3.4 Type bridging and zooming

When discharging weakening lemmas and/or discovering side-conditions of invariant breakdown, one often needs to add explicit (novel, if obvious) information about type relationships and their layers of abstraction/representation. These lemmas establish algebraic properties between new user-defined operators, and known (set theory, say) operators.

Provers have preferred directions for reasoning, be that left-right simplification, or a “waterfall” [28] involving generalisation and simplification. We call “type bridging” all those auxiliary lemmas necessary for weakening the goal (in backward proof) towards *true* by substituting “simpler” goals up to the point they meet the available hypothesis. For example, when a conditional rule fails to match, it is often because its side-condition could not be discharged. Often

### 3. PROOF PATTERNS

these conditions are simply type checking (e.g. parameters are within the involved function symbols’ declared types). This gives rise to a set of specialised type-inference lemmas for the expressions involved.

The mathematical toolkits of Z, Event-B, and VDM are defined in terms of maps, sequences, sets, and relations. User-defined functions are often given in terms of operators with higher automation, in which case the appropriate expansion lemma to the function symbol with most automation is needed. In methods such as Z and VDM, the notion of records is ubiquitous yet goals involving records do not require the prover to know about all the record structures, hence specific lemmas exposing a record’s properties are often needed to streamline proofs (*i.e.* in [15], such record slicing reduced a 45-page long cluttered goal into 16-chunks of related intent about a page each).

This type bridging, where the lemma is there to help bridge notions between operators of interest (*i.e.* user-defined and set theory in this case) enhances proof automation. Similarly, when moving between layers of abstraction/representation, the right level of definition expansion needs to be taken into account. That is, we need to instruct the prover what “zoom” to use, and we do not want to expand all definitions to sets or predicates, but rather keep definitions at different “zoom” levels, adding lemmas between levels as needed.

Such type judgements can also work in forward proof by strengthening hypotheses of interest. For instance, given a goal involving an injective map inverse ( $f \sim (f(x))$ ) and an assumption  $x \in \mathbf{dom} f$ , we could extend the hypothesis to say that  $f \sim (f(x)) = x$ . This usually makes the prover substitute the goal with the simpler right hand side (RHS) involving  $x$ . The amount and shape of these auxiliary lemmas are determined by the direction a prover takes, as well as by the amount of previously available information for the given operators. Finally, certain invariants are mathematically sensible (*i.e.* a sequence of size up to 10,  $s \in T^* \wedge \mathit{card}(s) \leq 10$ ), yet hard to use in proof because of tricky operators like cardinality ( $\mathit{card}(s)$ ) that involve bijective functions. Equivalent versions of the same invariant using just set theory (*i.e.* sequence indices range from 1 up to 10,  $\mathbf{inds} s \subseteq \{1, \dots, 10\}$ ) can be proved as type bridging lemmas, hence providing better automation.

*Example.* In the Isabelle proof for the heap example we added a (congruence) lemma that required no extra side-conditions, hence directly simplifying the goal by removing one of the operators. Within Isabelle libraries, the term “*congruence lemma*” is used to refer to lemmas that weaken the goal structure with respect to some easier (sub-)term that is preferred. For example, in Isabelle lemma `imp_cong` can be used to rewrite a goal involving an implication such as  $?P \longrightarrow ?Q$  to a corresponding goal such as  $?P' \longrightarrow ?Q'$ , so long as one can find that  $?P$  is related (in the example, equivalent) to  $?P'$  and the same for  $?Q'$ . Weaker relationships between involved terms are also used.

In our case, such a lemma states that subtracting from the map’s domain preserves separation (*i.e.*  $\mathit{sep}(x \triangleleft f) \longrightarrow \mathit{sep}(f)$ ). The presence of the operator to absorb on the left hand side (LHS) tells the prover our preference for the RHS expression as a result. Moreover, the free variables in weakening lemmas for *sep*



### 3. PROOF PATTERNS

needed for the feasibility of *NEW1* were reused in *DISPOSE1*, hence making the feasibility proof script itself much like the one for *NEW1* (*i.e.* their common strategy being reused modulo the key lemmas discovered).

#### 3.5 Retrieve state update

When refining a model, a *retrieve* relation maps concrete to abstract state representations. A key PO is that if the postcondition holds at the concrete level, then it holds in the abstract. This proof obligation is called *narrow postcondition* in VDM (correctness in Z). A common feature in refinement is a *type jump* between the abstract and the concrete levels. For example, the heap is represented as a set of locations at level 0 and refined to a partial map in level 1.

In these situations where the updated state at the concrete level is described using an equality  $\sigma = f(\overline{\sigma})$  for some  $f$ , that is, when we have a functional retrieve, we can apply a pattern called *retrieve state update* to conjecture lemmas that map across the type jump. To prove the postcondition we can use one-point witnessing under the retrieve to solve  $\text{retr}(\sigma) = g(\text{retr}(\overline{\sigma}))$  for some function  $g$  on the original state, where  $\text{retr}(\text{con}) = \text{abs}$  is a function from the concrete to the abstract state. After the state update with the equality above, this is really:  $\text{retr}(f(\overline{\sigma})) = g(\text{retr}(\overline{\sigma}))$ . In this situation, we have three pieces of information: a) the structure of the retrieve function mapping concrete and abstract; b) a set of operators at the concrete level (used in  $f$ ); and, c) a set of operators at the abstract level (used in  $\text{retr}$  and  $g$ ). To solve this goal, we must first translate the operators in  $f$  to those of  $g$  by distributing the retrieve function. For example, the narrow postcondition PO for the NEW operation on the heap is:

$$\text{retr}(\{r\} \triangleleft \overline{f}) = \text{locs}(\overline{f}) - \text{locs-of}(r, s)$$

The ‘zoom’ level for the abstract state is that of sets and the level of the concrete is maps. The application of this pattern suggests distributing *retr* over domain filtering, which suggests a possible lemma of the form

$$P \longrightarrow \text{retr}(\{r\} \triangleleft \overline{f}) = Q(\text{retr}(\overline{f}), \{r\})$$

where  $P$  expresses side-conditions and  $Q$  is some undetermined operation on sets. At this point the retrieve function is operating solely on the state, so we stop applying the pattern. Comparison with the RHS of our goal makes clear what  $Q$  should be, but this is not always so straightforward. Thus, just like invariant breakdown, the proof pattern can help suggest an attack, but it requires input from the proof engineer.

#### 3.6 Hidden case analysis

Hidden case analysis is the insertion of a lemma of the form  $P \vee \neg P$  to hypotheses and the subsequent breakdown of the proof obligation into the case where  $P$  holds and the case where  $\neg P$  holds. This pattern is common to mathematics as well. It is often used in situations where it is not clear what information the precondition provides to prove the postcondition holds, despite it being available, if hidden. A few examples of where this occurs are:

#### 4. PATTERNS IN THE HEAP

- Certain predicate calculus patterns in either pre/postconditions lend themselves to hidden cases analysis. Conditional postconditions ( $P \longrightarrow Q$ ) indicate a case split on the condition. Disjunctive postconditions are often associated with hidden disjuncts of the precondition (e.g. as in *NEW1*).
- Non-linear arithmetic operators occur in the precondition, such as  $\geq$ , and also when negation and non-linear arithmetic are combined.
- Sets, sequences, and maps in preconditions might require case analysis for emptiness; and so on.

For example, the disjunction in the postcondition for the *NEW1* operations

$$\overline{f_1}(r) = s \wedge f_1 = \{r\} \triangleleft \overline{f_1} \vee (\overline{f_1}(r) > s \wedge f_1 = (\{r\} \triangleleft \overline{f_1}) \cup \{r+s \mapsto \overline{f_1}(r)-s\})$$

suggests that a case analysis must be performed on the preconditions. In this case, introducing  $f_1(l) = s \vee f_1(l) \neq s$  allows us to derive from the precondition ( $\exists l \in \mathbf{dom} f_1 \cdot f_1(l) \geq s$ ) that  $f_1(l) > s \vee f_1(l) = s$ . Each case in the precondition is explicit as a goal in each side of the disjunction, although the disjunct with the precondition was hidden by  $\geq$ .

### 3.7 Shaping

In formal methods, proof obligations tend to be large with lots of information present in the goal and assumptions. This can often obscure the overall structure of a goal. In the shaping proof pattern, we utilise some of this information to simplify the goal before applying an important lemma or applying another proof pattern, such as invariant breakdown. A shaping pattern consists of a set of shaping lemmas, which are equalities between a sub term of the goal and a simpler representation, and a set of targeted rewrites that simplify the goal. The benefit of using this proof pattern before applying an important weakening lemma is that it will considerably simplify the resulting subgoals. In mathematics, this often occurs when trying to get rid of some difficult operator like square root, so one squares both sides of an equation, say.

In the heap example, shaping was used frequently to simplify some of the details of the *DISPOSE1* postcondition in the feasibility and refinement proofs, where prior case analysis had provided further information about the structure. We give an example of shaping (and the case analysis that triggers it) next.

## 4 Patterns in the heap

The top-level POs for a model, such as operation feasibility, tend to have similar structure that can be exploited to increase proof effectiveness. We call such similarities *methodological patterns*. Rather than describe these patterns, we give a worked example of how the patterns described in the previous section can be composed to help solve the feasibility proof obligation for the *DISPOSE* operation at level 1. In a companion technical report, we provide a detailed presentation of methodological patterns [13]. The VDM operation is as follows:

```

DISPOSE1 (d: Loc, s: N1)
ext wr f : Free1
pre is-disj(locs-of(d, s), locs(f))
post ∃below, above, ext ∈ Loc  $\xrightarrow{m}$  N1 .
    below = {l | l ∈ dom f ∧ l +  $\overleftarrow{f}$ (l) = d} ◁ f ∧
    above = {l | l ∈ dom f ∧ l = d + s} ◁ f ∧ ext = above ∪m below ∪m {d ↦ s} ∧
    f = (dom below ∪ dom above ◁  $\overleftarrow{f}$ ) ∪m {min-loc(ext) ↦ sum-size(ext)}

```

The inputs  $d$  and  $s$  are the start location and size of region to add back to the heap. The precondition is similar to level 0: this time using  $locs$  to construct the set of all free locations from the heap map.

The complexity in *DISPOSE1* arises from the fact that the memory region being added back to the heap may adjoin zero, one, or two other regions already in the heap. Thus, to preserve the *sep* part of the invariant (e.g. memory regions should be as large as possible), we must join them together. The map *above* will adjoin the end of the region being added; *below* defines the map of elements adjoining from the start. The *extended* map then consists of *above*, *below*, and the disposed region in the middle. Updating the state is then a case of removing above and below (using domain filtering) and adding a region that corresponds to the minimum starting location and the sum of the sizes of the elements of *ext*. To illustrate, for a heap  $f = \{0 \mapsto 4, 8 \mapsto 3\}$ , *DISPOSE1*(4, 4) would result in the updated  $f = \{0 \mapsto 11\}$ . This is because  $above = \{8 \mapsto 3\}$ ,  $below = \{0 \mapsto 4\}$ , and  $ext = \{0 \mapsto 4, 4 \mapsto 4, 8 \mapsto 3\}$ .

### Proof of *DISPOSE1* feasibility by patterns.

*Step 1: representation transformation.* This proof pattern can be used optionally at the start of any proof in order to make a representational change to simplify the proof. All that is required is a lemma which equates both POs. In the case of *DISPOSE1*, we gave explicit definitions for *above*, *below*, and *ext*, which makes it easier to deal with existential quantifiers, and we prove this alternative definition equal to the original. It is worth mentioning one ought not change the model for the sake of proof alone, as this often impairs model clarity.

*Step 2: safe decomposition.* Not a formal methods proof pattern per se, but *safe decomposition* is used as a standard technique in automated reasoning to break down fixed variables, hypotheses and conclusions of a goal. We extend this technique in our Isabelle development to unfold the definitions of *pre-DIS1* and *post-DIS1-ALT* by a single zoom level. The result, given as a declarative Isabelle/Isar proof script is as follows:

```

theorem DIS1-feas: (∀ · f d s . pre-DIS1 f d s → (∃ · f' . post-DIS1 f d s f'))
proof (subst dispose-feas-transform) — The transformation step
show (∀ · f d s . pre-DIS1 f d s → (∃ · f' . post-DIS1-ALT f d s f'))
unfolding pre-DIS1-def post-DIS1-ALT-def
proof (intro allI impI, elim conjE) — Safe decomposition

```

#### 4. PATTERNS IN THE HEAP

```

fix  $f d s$  assume  $inv: F1-inv f$  and  $pre: (locs-of d s) \cap (locs f) = \{\}$ 
show  $\exists f'. f' = (dom (below f d) \cup dom (above f d s)) \rightarrow f \cup m$ 
       $[min-loc (ext f d s) \mapsto sum-size (ext f d s)] \wedge F1-inv f'$ 
gap — The gap represents the area of the proof still to solve
qed
qed

```

*Step 3: witnessing.* In this case, witnessing is straightforward as we have a one-point existential. The resulting proof script replaces the previous **gap**.

```

show  $\exists f'. f' = (dom (below f d) \cup dom (above f d s)) \rightarrow f \cup m$ 
       $[min-loc (ext f d s) \mapsto sum-size (ext f d s)] \wedge F1-inv f'$ 
proof ( $rule\ exI, rule\ conjI, rule\ refl$ ) — Single-point witnessing with  $exI$ 
show  $F1-inv ((dom (below f d) \cup dom (above f d s)) \rightarrow f \cup m \dots)$ 
gap
qed

```

*Step 4: hidden case analysis.* At this point, we could apply the invariant breakdown pattern, but an expert proof engineer notes that the definitions of *above* and *below* mean it is either empty or a singleton map. Thus, the next proof step is to apply case analysis on both *above* and *below*, resulting in four separate cases to be solved. We focus on the first, where  $above = below = \phi$ .

*Step 5: shaping.* We again postpone invariant breakdown and perform the shaping pattern since we have added the hidden case analysis information above to the hypotheses. We now know, for example, that filtering the (empty) domains of *above* and *below* from  $f$  will have no effect. The result of case analysis and shaping leaves another proof gap as:

```

proof ( $cases\ below\ f\ d = empty, cases\ above\ f\ d\ s = empty$ ) — case analysis
assume  $below-empty: below\ f\ d = empty$  and  $above-empty: above\ f\ d\ s = empty$ 
have  $ab-shape: (dom (below f d) \cup dom (above f d s)) \rightarrow f = f$  gap
have  $min-loc-shape: min-loc (ext f d s) = d$  gap
have  $sum-size-shape: sum-size (ext f d s) = s$  gap
show  $?thesis$ 
proof ( $subst\ ab-shape, subst\ min-loc-shape, subst\ sum-size-shape$ )
show  $F1-inv (f \cup m [d \mapsto s])$ 
gap
qed
... — the other three cases are not shown
qed

```

*Step 6: zooming and decomposition.* We wish to attack each part of the invariant (*sep* and *disj*) independently, so we unfold definitions and decompose the conjunction accordingly.

*Step 7: invariant breakdown.* We are now in a position to apply the invariant breakdown proof pattern. We describe the application of the pattern for *disj*, but the *sep* part is similar. Recall, from above, that we have  $F1-inv(f)$  as an

## 5. RELATED WORK AND CONCLUSIONS

assumption (and thus also  $disj\ f$ ). The aim of invariant breakdown is to speculate lemmas that expose the hypothesis in our conclusion. In this case, we have the original state on the LHS of the map union operator, but not the right. This helps us to speculate a lemma:  $P \rightarrow disj(f) \rightarrow disj(f \cup_m [d \mapsto s])$  where  $P$  represents unknown side-conditions. In the case of  $disj$ , the important condition is  $(locs\text{-of}\ d\ s) \cap (locs\ f) = \{\}$ , which is exactly the precondition for  $DISPOSE1$ . The (partial) proof script is:

```

show  $disj\ (f \cup_m [d \mapsto s])$ 
proof (rule unionm-singleton-disj) — The speculated lemma
show  $sep\ f$  using  $inv$  by assumption — Use of hypothesis
show  $(locs\text{-of}\ d\ s) \cap (locs\ f) = \{\}$  by (rule pre) — Precondition as side-condition
... — Additional side-conditions
qed

```

The use of invariant breakdown does not solve the invariant subgoals, but it directs the proof engineer to the appropriate lemma structures to speculate.

*Step 8: weakening and type bridging.* Because of the repeated nature of some POs, weakening lemmas speculated by invariant breakdown are often reusable. The lemma *unionm-singleton-disj* is used in each of the other three cases introduced by hidden case analysis.

Lemma inference is quite hard in general, yet within the proof process presented, weakening and type bridging lemmas that will make proof progress are easier to identify. For example, in the proof for the other cases in the  $DISPOSE1$  feasibility PO, because of the presence of  $inv(f \cup_m g)$ , we need a side-condition that the map domains are disjoint ( $\mathbf{dom}\ f \cap \mathbf{dom}\ g = \{\}$ ). In context this is given by the rather specific lemma:

**lemma** *l-dispose1-munion-disjoint*:  
 $dom\ ((\mathbf{dom}\ (dispose1\text{-below}\ f1\ d1) \cup \mathbf{dom}\ (dispose1\text{-above}\ f1\ d1\ s1)) \rightarrow f1) \cap$   
 $dom\ [min\text{-loc}\ (dispose1\text{-ext}\ f1\ d1\ s1) \mapsto sum\text{-size}\ (dispose1\text{-ext}\ f1\ d1\ s1)] = \{\}$

This lemma is an instantiation of a side-condition that features in all subgoals when breaking down the invariant for  $DISPOSE1$ , hence is the key enabler of progress for the invariant breakdown of that feasibility proof.

## 5 Related work and conclusions

This paper introduced the concept of *proof patterns in formal methods* and described several patterns of proof that commonly occur. We tested this hypothesis by an experiment: the proof of a VDM heap memory manager [27]. Just as in software design patterns [19], proof patterns are informal and are described generically, without reference to an individual problem. We also describe the composition of a set of patterns in a worked example of a feasibility proof obligation. We believe that a small collection of proof patterns is all that is needed to increase the automation of most formal methods proof obligations. This **will**

## 5. RELATED WORK AND CONCLUSIONS

**not** eventually remove the burden of proof, yet we believe proof patterns offer good support in de-skilling the process, as well as increasing proof effort reuse.

Before getting to the nub of the problem within industrial-scale proof obligations, which almost always involve large formulae (i.e. tens of pages long) and multiple (i.e. over 100) variables, we claim it is fundamental to have in place a considerable amount of machinery to enable automation to an acceptable level. Proof engineering is essential for scalability: it takes a good amount of unrelated proof effort to enable the actual proof obligations of interest to be tackled. Lemmas are useful whenever one needs to: decompose a complex problem; fine-tune the theorem prover’s rewriting abilities to given goals; generalise a solution of some related (usually more abstract) problem; or to provide alternative solutions/encodings of the same data structure/algorithm being modelled; *etc.*

*Related work.* The term design pattern originated in architecture [2], but is most widely known in software design [19]. Some languages, such as Java, even have built-in support from patterns, such as the *iterator* pattern. In software engineering, there is a wealth of research in design patterns, from architectural patterns to specific patterns for user interfaces. In [7], Buschmann introduces architectural patterns that are capable of describing large-scale software systems, such as model-view controller, which separates the representation of data (the model) and the user’s view of it. Architectural patterns are similar to our methodological patterns, which we describe as a composition of proof patterns to help solve a top-level proof obligation. We describe some closely related work in formal proof, but Buschmann et al. related work on design patterns [8].

It has been noted in the mathematical community that pencil-and-paper proofs often follow specific patterns, such as proof by contradiction, by induction, etc. These ideas have found their way into the domain of automated theorem proving. Bundy’s proof plans were an early attempt at capture patterns of inductive proof [4]. Proof plans have been implemented many times, most recently in Isaplanner [10]. We would like to investigate whether the planning language in Isaplanner is expressive enough to formalise some of our proof patterns. In [26], we start this process by describing a language to capture meta-proof information. Also within the AI<sub>4</sub>FM project, colleagues have developed a graphical rewriting language [20] that will hopefully be amenable to proof pattern recognition. As well as the static proof patterns described here, it may be possible to learn patterns from a corpus of proofs [22,23]. Some mathematical proof patterns are so prevalent that they have been implemented as theorem provers in their own right. The Calculator theorem prover, for example, can be seen as specifically implementing the pattern of proof using Galois-connections and indirect equalities [32].

*Future work.* In the heap example, we also used AI<sub>4</sub>FM tools (under development) to collect (> 250 GB) proof process data [33,34], which we plan to analyse to identify clusters using AI techniques akin to [21]. We hope with this data to find hard evidence for proof patterns. The AI<sub>4</sub>FM hypothesis is that “*enough information-extraction can be automated from a mechanical proof that future proofs of examples of the same class can have increased automation*”. The

## 5. RELATED WORK AND CONCLUSIONS

challenge is discussed in several earlier publications including [6,26,11]. As such, an important area of future work on proof patterns is the transference of a proof, described using patterns, to help automate another similar proof obligation. We are also interested in extending our catalogue of proof patterns, as well as collecting further examples of the patterns for different formal methods. While we have given a natural language presentation of patterns in this paper, we would like to formalise a pattern language in order to present and automatically generate proofs from the patterns, similar to work in software design patterns [3]. We will analyse the proof process data captured by our tools using AI techniques.

*Acknowledgements* We are grateful to Cliff Jones for suggesting the Heap problem and for, together with Andrius Velykis, the many fruitful discussions. Other AI<sub>4</sub>FM members helped us understand important problems in automated reasoning. This work is supported by EPSRC grant EP/H024204/1.

## References

1. J.-R. Abrial. *The Event-B book*. Cambridge University Press, UK, 2010.
2. C. Alexander, S. Ishikawa, and M. Silverstein. *A Pattern Language: Towns, Buildings, Construction (Center for Environmental Structure Series)*. Oxford University Press, later printing edition, Aug. 1977.
3. F. J. Budinsky, M. A. Finnie, J. Vlissides, and P. Yu. Automatic code generation from design patterns. *IBM Systems Journal*, 35(2):151–171, 1996.
4. A. Bundy. The use of explicit plans to guide inductive proofs. In E. Lusk and R. Overbeek, editors, *9th International Conference on Automated Deduction*, volume 310 of *LNCS*, pages 111–120. Springer Berlin, 1988.
5. A. Bundy, D. Basin, D. Hutter, and A. Ireland. *Rippling: Meta-level Guidance for Mathematical Reasoning*, volume 56 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2005.
6. A. Bundy et al. Learning from experts to aid the automation of proof search. In L. O’Reilly et al., editors, *PreProceedings of the 9th AVoCS’09*, CSR-2-2009, pages 229–232. Swansea University, UK, 2009.
7. F. Buschmann et al. *Pattern-oriented software architecture: a system of patterns*. John Wiley & Sons, Inc., New York, NY, USA, 1996.
8. F. Buschmann, K. Henney, and D. Schmidt. Past, present, and future trends in software patterns. *Software, IEEE*, 24(4):31–37, 2007.
9. A. Butterfield, L. Freitas, and J. Woodcock. Mechanising a formal model of flash memory. *Science of Computer Programming*, 74(4):219–237, 2009.
10. L. Dixon and J. Fleuriot. IsaPlanner: A Prototype Proof Planner in Isabelle. In *Proceedings of CADE’03*, volume 2741 of *LNCS*, pages 279–283, 2003.
11. L. Freitas, C. B. Jones, A. Velykis, and I. Whiteside. How to say why. Technical Report CS-TR-1398, Newcastle University, [www.ai4fm.org/tr](http://www.ai4fm.org/tr), November 2013.
12. L. Freitas and J. McDermott. Formal methods for security in the Xenon hypervisor. *International Journal on Software Tools for Technology Transfer*, 13(5):463–489, 2011.
13. L. Freitas and I. Whiteside. Proof patterns for formal methods. Technical Report CS-TR-1399, Newcastle University, Nov 2013.
14. L. Freitas and J. Woodcock. Proving theorems about JML classes. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 4700 LNCS:255–279, 2007.

## 5. RELATED WORK AND CONCLUSIONS

15. L. Freitas and J. Woodcock. Mechanising Mondex with Z/Eves. *Formal Aspects of Computing*, 20(1):117–139, 2008.
16. L. Freitas and J. Woodcock. A chain datatype in Z. *International Journal of Software and Informatics*, 3(2-3):357–374, 2009.
17. L. Freitas, J. Woodcock, and Y. Zhang. Verifying the CICS file control API with Z/Eves: an experiment in the verified software repository. *Science of Computer Programming*, 74(4):197–218, 2009.
18. L. Freitas, J. Woodcock, and F. Zheng. POSIX file store in Z/Eves: an experiment in the verified software repository. *Science of Computer Programming*, 74(4):238–257, 2009.
19. E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1998.
20. G. Grov, A. Kissinger, and Y. Lin. A graphical language for proof strategies. To appear in LPAR’13. Also available at arXiv:1302.6890, 2013.
21. J. Heras and E. Komendantskaya. ML4PG in computer algebra verification. In *Conferences on Intelligent Computer Mathematics*, 2013.
22. J. Heras and E. Komendantskaya. Statistical proof-patterns in Coq/SSReflect. *CoRR*, abs/1301.6039, 2013.
23. M. Jamnik et al. Automatic learning of proof methods in proof planning. *Logic Journal of the IGPL*, 11(6):647–673, 2003.
24. C. Jones, P. O’Hearn, and J. Woodcock. Verified software: a grand challenge. *IEEE Computer*, 39(4):93–95, 2006.
25. C. B. Jones. *Systematic Software Development using VDM*. Prentice Hall, 1990.
26. C. B. Jones, L. Freitas, and A. Velykis. Ours is to reason why. In Z. Liu, J. Woodcock, and H. Zhu, editors, *Theories of Programming and Formal Methods*, volume 8051 of *LNCS*, pages 227–243, 2013.
27. C. B. Jones and R. C. F. Shaw, editors. *Case Studies in Systematic Software Development*. Prentice Hall International, 1990.
28. M. Kaufmann, P. Manolios, and J. S. Moore. *ACL2 Computer-Aided Reasoning: An Approach*. University of Austin Texas, 2009.
29. T. Nipkow, L. C. Paulson, and M. Wenzel. *Isabelle/HOL — A proof assistant for higher-order logic*, volume 2283 of *LNCS*. Springer, 2002.
30. L. C. Paulson. Isabelle: A Generic Theorem Prover. *LNCS*, 828, 1994.
31. M. Saaltink. The Z/EVES system. In J. B. et al, editor, *ZUM ’97: The Z Formal Specification Notation*, volume 1212 of *Lecture Notes in Computer Science*, pages 72–85. Springer Berlin / Heidelberg, 1997.
32. P. F. Silva and J. N. Oliveira. Galculator: functional prototype of a Galois-connection based proof assistant. In S. Antoy and E. Albert, editors, *PPDP*, pages 44–55. ACM, 2008.
33. A. Velykis. Inferring the proof process. In C. Choppy, D. Delayahe, and K. Klai, editors, *FM2012 Doctoral Symposium*, Paris, France, August 2012.
34. A. Velykis. *Capturing & Inferring the Proof Process (under submission)*. PhD thesis, School of Computing Science, Newcastle University, 2014.
35. M. M. Wenzel. *Isabelle/Isar - a versatile environment for human-readable formal proof documents*. PhD thesis, Technische Universität München, 2002.
36. J. Woodcock and J. Davies. *Using Z*. Prentice Hall International, 1996.
37. J. Woodcock and L. Freitas. Linking VDM and Z. In *International Conference on Engineering of Complex Computer Systems*, pages 143–152, Belfast, 2008.