



HAL
open science

Interactive, Iterative, Tooled, Rule-Based Migration of Microsoft Access to Web Technologies

Santiago Bragagnolo, Nicolas Anquetil, Stéphane Ducasse, Mustapha Derras

► **To cite this version:**

Santiago Bragagnolo, Nicolas Anquetil, Stéphane Ducasse, Mustapha Derras. Interactive, Iterative, Tooled, Rule-Based Migration of Microsoft Access to Web Technologies. 2023. hal-04181591v1

HAL Id: hal-04181591

<https://inria.hal.science/hal-04181591v1>

Preprint submitted on 22 Aug 2023 (v1), last revised 6 Sep 2023 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Interactive, Iterative, Tooled, Rule-Based Migration of Microsoft Access to Web Technologies

Santiago Bragagnolo*^{1,2} | Nicolas Anquetil¹ | Stephane Ducasse¹ | Abderrahmane Serai² | Mustapha Derras²

¹ Université de Lille, CNRS, Inria, Centrale Lille, UMR 9189 – CRISTAL, Lille, France

² Berger-Levrault, France

Correspondence

*Santiago Bragagnolo, Email: santiago.bragagnolo@berger-levrault.com

Present Address

Parc scientifique de la Haute-Borne 40, avenue Halley - Bât B - Park Plaza 59650 Villeneuve d'Ascq – France

In the context of a collaboration with Berger-Levrault, an IT company producing information systems, we are working on migrating Microsoft Access monolithic applications to the web front-end and microservices back-end. Like in most software migrations, developers must learn the target technology, and they will be in charge of the evolution of the migrated system in the future. To respond to this problem, we propose the developers take over the migration project. To enable the developers to drive the migration to the target systems, we propose an Interactive, Iterative, Tooled, Rule-Based Migration approach.

The contributions of this article are (i) an iterative, interactive process to language, library, GUI and architectural migration; (ii) proposal of a set of artefacts required to support such an approach; (iii) three different validations of the approach: (a) library and paradigm usage migration to Java and Pharo, (b) tables and queries migration to Java and Typescript, (c) form migration to Java Springboot and Typescript Angular.

KEYWORDS:

Software Migration, Interactive, Iterative, Multi-target

1 | INTRODUCTION

In the context of a collaboration with Berger-Levrault, an IT company, we are migrating Microsoft Access monolithic applications to Typescript and Angular front end; Java and SpringBoot back end.

The migration constraints are those often encountered in the industry¹: The software quality is uncertain, development must continue during the migration, resources are limited, the desired target architecture is sketchy, developers are yet to master the target technology fully, and they will be in charge of the evolution of the migrated system in the future.

The last point is considered vital. Despite the changes in architecture and technology, the development team must be able to retain or migrate its knowledge of the application to continue evolving it in the future. To ensure this, we propose that the development team be the driving force behind the migration, choosing what to migrate, where and how to migrate. Yet, as already stated, they are not specialists in the target technologies (Typescript and Angular) and software migration. We propose an interactive, iterative/incremental and tooled migration process to allow them to achieve this migration.

Abstraction. A tooled approach allows software engineers to concentrate on what they want to achieve without considering how to achieve it. For example, they can transform a function into a method without worrying about transformation rules or the syntax of methods in the target language.

Low barrier. An interactive approach enables developers to point to a software artefact, e.g., a function, and decide where to migrate it, e.g., a class, without worrying *how* to do it. The fact that the approach is based on rules reduces the need for migration experts, as experts have to put their knowledge into reusable fine-tuned rules.

Gradual learning. An iterative and incremental approach allows developers to migrate screens or functionalities individually. This, plus the feedback provided by a tooling approach, allows one to gradually get acquainted with the migrated solution, its new architecture, and its new technology.

In this paper, we present an approach and a tool to realise this interactive and iterative migration. The paper is organised as follows:

- In Section 2, we give a first high-level view of the approach and its main components;
- In Section 3, we present the different elements of our migration infrastructure, which make our approach possible: models, directives, mappings and rules.
- In Section 4, we present the engine which puts all the artefacts together and enables the application of directives.
- In Section 5, we validate the approach by presenting results on three different kinds of migration.
- In Section 6, we discuss different aspects and decisions of our approach.
- In Section 7, we propose future work and conclude.

2 | OVERVIEW

We introduce the challenges and the expectations of our migration approach in Section 2.1, and we sketch the proposed process and its different parts in Section 2.2.

2.1 | Challenges & Requirements

The migration project from a monolithic Microsoft Access application to a front-end (Typescript + Angular) and back-end (Java + SpringBoot) application entails the following challenges:

Programming language migration: From Visual Basic for Applications (VBA) to multiple targets: HTML, TypeScript, and Java; To translate incompatible structures^{2,3}. VBA actively supports and encourages the usage of language structures that are undesirable or have no equivalent in our target languages, such as GO-TO.

Paradigm migration: From procedural paradigm (VBA) to object-oriented paradigm (Java and Typescript). All translated codes must be expressed in terms of classes, methods and objects^{4,3}. Note that we do not pay close attention to this challenge in the paper.

Library migration: The target systems must use target libraries^{3,2,5,6}. e.g., The “long” type in Microsoft Access must be migrated to “BigInteger” in Java or “bigint” in Typescript. “MsgBox” function usages are transformed to “alert” in Typescript or a Log4j “logger” in Java;

Architectural Migrations: The migration includes the splitting of a single stand-alone application into a front-end and a back-end application and establishing connections between the parts^{7,8,9,10}; it also includes splitting both front-end and back-end concerns into multiple micro-service applications.

GUI migration. To migrate to the target system, the source system’s GUI visual and behavioural aspects from desktop GUI to web GUI; ^{11,12,13,14}.

Migrating the expertise: As much as possible (considering the architectural changes), artefacts (classes, functions, forms) should be migrated so that developers recognise them afterwards and know where to find them¹¹. ;

Integrated in the typical lifecycle: The source and target projects are expected to have their lifecycles so that new functionalities and bug fixes can be processed with the slightest disturbances and shipped and deployed independently.

2.2 | An Iterative & Interactive Process for software migration

The process that we propose is depicted in Figure 1; it has the following steps:

Step 1. The migration tool loads the source and target applications and creates a model for each of them. These models are presented to the software developers in the tool as a tree of software artefacts and matching source code (see, for example, Figure 2).

Step 2. The user navigates the source hierarchy and selects an artefact to migrate. It could be a module, a form¹, a function,...

¹In Access, forms are first-class citizens of the programming language

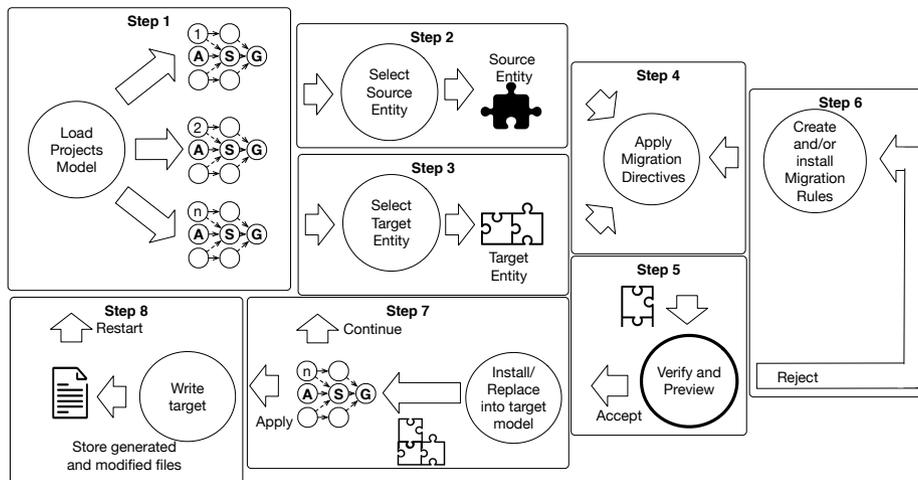


Figure 1 Iterative Interactive Migration Process Steps.

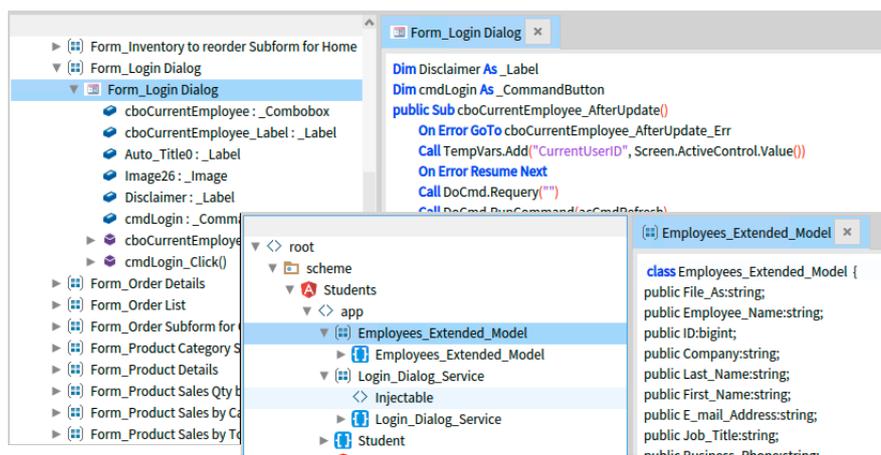


Figure 2 Microsoft Access model (top left) and Typescript Angular model (bottom right) presented to the user when watching the source code

- Step 3.** The user manually chooses where in the target hierarchy he wants to migrate the artefact. The tool provides migration actions for the user to perform. These actions are named *directives* and introduced in Section 3.2. The tool accepts any location; however, to be meaningful, it must be coherent with the artefact to migrate, *i.e.*, a function can be migrated into a class and become a method.
- Step 4.** The tool responds to the migration actions by applying migration rules (transformations) to perform the migration. This mechanism is explained in Section 4.2 and Section 4.3.
- Step 5.** The user has immediate feedback. If the tool can apply the directive properly, the user can immediately analyse and verify the code in the targeted model. If the tool fails, the user is informed, and the tool rolls back any intermediate change done by the failing directive.
- Step 6.** If the user is unsatisfied with the result after verifying the code, he can refuse it by rolling back to the previous state. Whatever the reason for a rollback (error or user choice), the migration expert (typically not the application developers) will need to add or change a rule for the process to restart.
- Step 7.** If the change is accepted, it is introduced in the target model. The user may continue applying directives, going back to step 2 or generating the code.
- Step 8.** The user chooses to apply all the modifications to the project's source code. This modifies the project files to comply with the new model. After this step, we restart the whole process from step 1.

The tool allows executing the model modifications over the different application's source code at any moment, applying manual changes over any of the applications, and reloading the models.

3 | ELEMENTS OF OUR MIGRATION INFRASTRUCTURE

Our approach enables developers to execute arbitrary partial migrations between two different application models (see Section 3.1) through the interactive application of *directives* (see Section 3.2), which are resolved by the immediate or delayed application of rules (see Section 3.4).

This approach is an extension of the one presented by *Bragagnolo et al*¹⁵, enabling interaction and migrations beyond language migration.

3.1 | Modeling source and target applications

Our model-driven approach applies modifications to target models based on entities and knowledge extracted from source models. All models are instantiated from the same heterogeneous unified meta-model.

To achieve that, we load a model for each application in the migration project. In our primary study case, we target a client/server architecture. We instantiate three different models, one per application: (i) monolithic legacy application (Microsoft Access), (ii) migrated back-end application (Java Springboot) and (iii) migrated front-end application (Typescript Angular).

3.1.1 | The Heterogenous Unified Meta-model

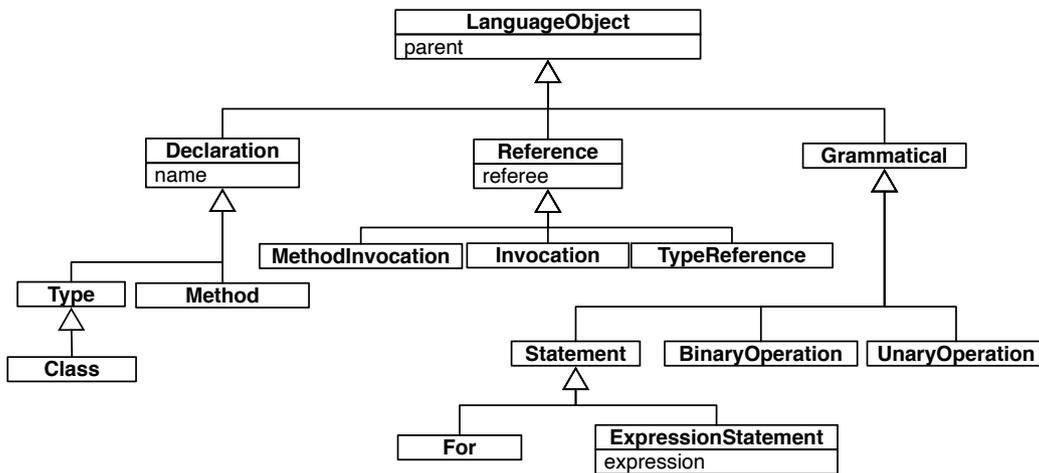


Figure 3 Exempt from our Heterogenous Unified Meta-Model, including the three core entities: Declaration, Reference and Grammatical. (See Text)

We *distinguish* different categories of types. Figure 3 shows a fragment of the meta-model. Those that can be referred to (e.g., variable declaration), that we call *declaration nodes*; and nodes for elements of the language (like statements or arithmetic expressions) that we call *grammatical*. On top of that, references are reified, giving a third kind: the *reference*. There is an implicit dependency between the nature of a declared artefact and the kind of reference able to use this artefact. To activate a function which receives one parameter, we use a function invocation with one argument.

Abstract Semantic Graph. Our approach leverages the heterogeneous unified meta-model presented by *Bragagnolo et al*¹⁵. This meta-model consists of an Abstract Semantic Graph (ASG)² that uses the same entities for different languages as long as the instances do have the same semantics: the modelling rationale behind is to model likeness and unlikeness of the different constructs. When this is not the case (e.g., Java classes and TypeScript classes), we represent them with different types to be able to differentiate them in the migration rules.

Representing inconsistent intermediate states. Our approach also requires representing *inconsistent* intermediate states as it is iterative and interactive. The migration engineer must be able to (i) write a migration rule and see the outcome, whether it is or is not correct. (ii) productive rules that produce partially correct outcomes due to insufficient information. e.g., Let us consider that we have to migrate a function invocation. Suppose we do not know how the invoked function will be migrated. In that case, we cannot decide how to translate the invocation: method invocation, static method invocation, function invocation, or attribute access.

²An ASG is an Abstract Syntax Tree extended by linking uses with definitions or declarations, making it into a graph

Representing Primitive and third party types and behaviours. Modelling of an application in a given programming language also includes creating declaration instances for the programming language's primitive types and any external library functions used. These declarations assert the existence of the library entities and have no complete definition; they are used to be referred by the parts of the application that require them.

3.2 | Directives: User explicit actions

Our engine exposes two main actions to the user: *Produce* and *Map*. These actions establish a relationship between two entities belonging to two different models. Within the duration of the action, one entity plays the role of the source, and the other plays the role of the target. To convey the direction of the action, we call them *Directives*.

Directives.

1. *Produce*: Given a *source entity* and a *target context*, this directive instructs to produce the given *target entity* inside the given *target context*, based on the *source entity*.
2. *Map*: Given a *source declaration*, a *target declaration*, and a *scope of validity* (in the target model), this directive establishes a semantic equivalence between the two declarations, meaning that —within the given scope— past and future references to the source declaration will be replaced by references to the target declaration.

Directives are implemented by migration rules and mappings, as shown in the following sections. These two directives offer the user the activation of two migrating actions.

Actions.

1. the *declaration* of a software entity in a model (e.g., a function in the legacy application) must be created as the *declaration* of another software artefact in another model (e.g., a method in the migrated application);
2. *references* to a software artefact (e.g., a primitive type or a library function in the legacy application) must be replaced by *references* to another equivalent software artefact (e.g., a matching primitive type or a replacing library function in the migrated application).

Note that the first action (re-declaration of a software artefact) also entails the second, as the previous references to the original artefact must be changed to references to the migrated one.

3.3 | Cross-model Mappings & Stubs

Mappings and Stubs represent two different kinds of equivalences between two different models.

Mappings. Mappings are *scoped* objects representing a *semantic equivalence* between two models: a source entity is equivalent to a target entity. For example, (long => bigint); (MsgBox => alert); etc.

By *scoped*, we mean that they have a scope of validity. This is to say that A is equivalent to B for a specific project, package, class, method, etc. This context is defined when recording these mappings. When the mapping is recorded automatically, this scope is the whole project. When the mapping is recorded manually, the user defines its scope.

Some rules use mappings (See Section 3.4). In our approach, a mapping association is a result of a map directive (explained in Section 3.3) or the execution of a Productive Rule, as explained in Section 3.4.

Stubs. A reference node can only point to a declaration node within the same model. When a reference in a model (e.g., a migrated application model) should point to a declaration that does not exist in this model (typically a declaration that has not been migrated yet), we create a *stub declaration* in the same model as the reference. The reference may then points to this stub, and the stub itself points to the declaration of the entity in the other model. Therefore *stubs* are bridges between models. A model cannot generate valid source code as long as it contains stubs because some references still need to be migrated. When the referred entities are created, a specific mechanism replaces these stubs with actual declaration nodes (Section 4.2).

3.4 | Rules

Rules are *scoped* conditional operations consisting of a Condition and an Operation. *Condition* consists of a predicate that satisfies the operation's assumptions and requirements. *Operation* consists of any systematic modification over the target. By *scoped*, we mean that they have a scope of validity. This scope is all migration models, as there are default rules, a specific project, package, class, or method. This context is defined when installing the rules. We have two families or rules: Productive and Adaptive.

Productive Rules These rules create a migrated version of a source entity in the context of a target entity, that is to say, as a child of this target entity. Applying a Productive Rule also entails automatically mapping the source entity and the produced target entity.

Condition consists of a predicate that tells if the rule can produce this migrated version within the specific target entity.

Operation consists of modifying the target entity by defining the source entity's migrated version. These rules are typically applied from the legacy application model –containing the source entity– to a migrated application model –containing the target entity.

For example, a rule is executed when the source entity is a function, and the target context is a class, and it operates a migration by defining a method in the target class based on the source function.

Adaptive rules We said before that *references* to a software artefact must be replaced as *references* to another equivalent software artefact. To enable this replacement mechanism, we copy the *references* as they are, knowing that this may be a wrong decision, and delay the migration criteria to when this other software artefact is available. Adaptive rules detect the apparition of semantically equivalent declarations and modify or replace reference objects based on the nature of the software artefact.

Condition consists of a predicate that tells if the rule can modify or replace a target reference to refer to a given target declaration.

Operation consists of the adaptation of the reference object.

For example, let us consider a function that just migrated as a static method. Right after, an adaptive rule will detect that this static method is recognised as equivalent to the function and therefore detect all the uses (migrated before and after the migration of the function) and adapt them from the form `function()` to the form `ClassName.function()`.

3.5 | Migration engine: Executing directives with rules look-up and default behaviour

At the heart of our migrating approach is the *engine* controlling the application of the directives, the selection of the rules and their application. In our approach, rules and mappings are contextualised, meaning they only affect a part. We can install a rule or a mapping; thus, it is always applicable or only applicable in the context of specific construction (project, package, class, etc.).

3.5.1 | Executing the produce directive

When the user initiates a produce directive on a source entity and a target context, the engine searches in the target context for a *Productive Rule* that can handle the migrating case. This is to say that the condition predicate accepts the source entity and target context. As we explained in Section 3.4, rules are contextualised, so we look up rules by context. The search starts in the target context and goes up the ASG as needed. We present in detail its application in Section 4.2.

The produce directive is direct order of transformation from the user. Therefore, it must be accomplished immediately, which means we must have a default rule.

Default Productive Rule: AnyCopy.

As we explained in Section 3.1.1 we base our approach on a heterogenous unified meta-model, meaning that we use the same meta-model for all projects. This choice is no accident. We want only to provide rules for the cases that require a transformation. For the rest, we provide the AnyCopy Productive Rule.

AnyCopy.

Context: Root context (available regardless of which project is the target)

Condition:

1. Always true

Operation:

1. Make an instance of the same class as the source entity within the target context.
2. Migrate each of the children of the source entity using the freshly created instance as the target context.

3.5.2 | Executing the Map directive

When a map directive is initiated by the user, on a source and target declarations, within a context, the engine installs a mapping in the given target context (for example, function F is equivalent to method M only in the package P) right after; the engine searches all the references within the given context to be fixed. Each of them searches for any *adaptive rules* that can be applied after the mapping has been registered, starting with the reference-to-fix context and going up the ASG as needed. Adaptive rules are executed only when there is a mapping available. Therefore there is no default rule. We present in detail its application in Section 4.3.

4 | THE APPROACH IN ACTION

We illustrate the process of the migration of the VBA sub-procedure “showName” (Listing 1), to the context of a Java class named “MyDestination” in the back end application (Listing 2). The source sub-procedure is a simple piece of Visual Basic that pops up a dialogue showing the content of a variable `name`, concatenated with the string literal “Ms”. Because this is migrated to the back-end application, the migrated class cannot access a GUI. Any information that was previously displayed should now be logged.

Listing 1: VBA sub-procedure that pops up a dialog

```
Dim name as String
public Sub showName()
    Call MsgBox ("Ms " & name)
End Sub
```

Listing 2: MyService Java class

```
package MyPackage;
class MyDestination {
    public static void log (String) {
        ...
    }
}
```

4.1 | Engine setup

We consider the engine configured with specific rules and no previously existing mapping for this example. We now present the installed productive rules AnyCopy, CopyAsStaticMethod and CopyReplaceOperator and the adaptive rule RenameAdaptToStaticReceiver.

4.1.1 | Productive Rules

AnyCopy.

was presented in Section 3.5.1.

CopyAsStaticMethod.

Context: Java migrated application

Condition:

1. The source entity is a sub-procedure **AND**
2. The target context is a class

Operation:

1. Define in the target entity a static method with the same selector as the source entity and using void as returning type reference.
2. Migrate all the children of the source entity using the method as the target context.

CopyReplaceOperator.

This rule requires two parameters when it is instantiated: The operator to detect (OtD) and the operator to replace it (OtR).

Context: Java migrated application

Condition:

1. The source entity is a binary operation
2. The operator matches OtD.

Operation:

1. Define a binary operation in the target entity, using OtR as an operator.
2. Migrate all the children of the source entity using the binary operation as the target context.

4.1.2 | Adaptive Rule

For this example, we use only one adaptive rule, defined before by *Bragagnolo et al*¹⁵.

RenameAdaptToStaticReceiver – Replace by method invocation with static receiver .

Context: Java migrated application

Condition:

1. The target reference is a function invocation **AND**
2. The given map target declaration is a method **AND**
3. The given map target declaration is *static*.

Operation:

1. Define a method invocation expression (a reference) using the method's parent as a receiver.
2. Set the arguments used by the function invocation in the new method invocation.
3. Set the new target reference to the mapped target declaration.

4.1.3 | Rule installation

To complete the engine setup, we must install the rules we want for our example. As we said in Section 3.4, rules are *contextualized*. As we said in Section 3.5 the engine performs a lookup through the target to find rules and mappings. Section 4 depicts our example installation. AnyCopy is always available regardless of the targeted project. In the context of the target project, we find the rules defined above: CopyAsStaticMethod, CopyReplaceOperator and RenameAdaptToStaticReceiver.

4.2 | Example of the application of a produce directive

The produce directive starts when a user asks to migrate the specific source declaration, the sub-procedure “showName”, into a specific target declaration, “MyDestination” class.

The result of the application of the produce directive is displayed in Figure 5 and can be used to follow the migration. We break down the example into eight steps. Figure 5 shows the source model and the target model after the execution of this directive. Each number attached to a target model entity is the step where they are created.

Step 1 The engine searches for a *Productive Rule* in the context of the “MyDestination” class (of the target model) and does not find it; next, it looks in its parent context and up in the context hierarchy until the top level node of the Java migrated application model where the CopyAsStaticMethod rule is found that matches: following its definition (see Section 4.1.1), it accepts a subprocedure being migrated to a class. The rule creates a method declaration node as a child of the class declaration node (the target context) and gives the name of the sub-procedure to the method. In this case, the return type is set to void as VBA sub-procedures do not return any value. The method is also defined as *static* to allow invoking it without an object. Then the rule delegates the migration of all children to the engine.

Step 2 Here, “showName” has no parameter, so the only child is the single statement in the sub-procedure’s body. The engine now looks for a *Productive Rule* accepting the statement as source entity and the newly created method declaration node as target context. The rule matching these source and target entities is found at the top level of the model. It is the generic AnyCopy rule (Defined in Section 3.5.1) that copies the source sub-tree into the destination context. So it will copy the “expression-statement” node (a grammatical node, see Section 3.1) of the sub-procedure. And again, delegate handling of its child (a “function call” reference node).

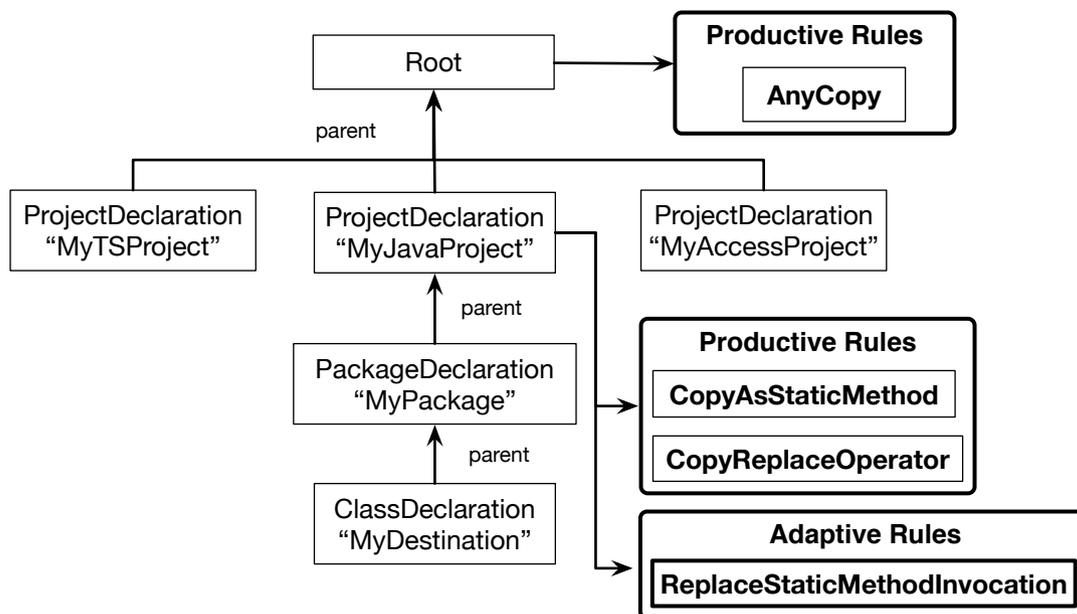


Figure 4 Rules by target context. Between the contexts, we find all the projects. We only analyse the rules related to the Java project. At root, we find the rules that apply to all targets. At the Java level, we find Productive and Adaptive rules.

Step 3 This also falls back to the `AnyCopy` rule. But this time, it must be noted that the copied node is not valid in a Java model (no function calls in Java, only method calls). This problem will be corrected later (by a map directive). For now, the reference remains a function call and points to no declaration node.

Step 4 The next step is to handle the children of the function call. In this case, it is the argument of the function: the binary operation. This time `CopyReplaceOperator` (defined in Section 4.1.1) rule is found that matches. This rule has been configured to detect and replace the “&” binary operator (string concatenation), which is lookup in the top context of the Java model that generates a “+” operator.

Step 5 and 6 Two new iterations of this process will apply the `AnyCopy` rule on the operator’s arguments. As the first argument, this gives a literal string node (valid in Java) and, as the second argument, a reference to a variable named “name”. Note that at this stage, the reference node points to nothing (no declaration node).

Step 7 The engine is now done with the creation of nodes in the target model, the entire AST of the “showName” sub-procedure declaration has been recreated (in a modified form) in the Java model, and there are two references (variable access and function call) pointing to nothing.

The engine then looks at the references that were created. For the variable reference, it searches whether an adaptive rule can make this variable reference fit the target model (it should be pointing to an existing variable). Supposing there is none (the variable still needs to be migrated), the engine creates a stub declaration to which the variable reference can point, which points to the variable reference in the source model. Similarly, for the function call, the engine looks for an existing adaptive rule able to adapt the call to `MsgBox` to fit the target model (it should be pointing to an existing method) and does not find one. A stub function declaration is created that points to the function invocation in the source model. The state of the models at this stage is pictured in Figure 5.

Step 8 Finally, in the context associated to “MyDestination” target class, a mapping is registered from the legacy “showName” function to the migrated “showName” method. At this stage, the produce directive is finished; it gives a target model that is not coherent as it contains nodes that are not valid in Java. Invalid nodes are allowed temporarily if one does not try to export the model to the source code.

Please, note that two features of our meta-model are leveraged in this explanation. (i) Four out of six rule applications use the `AnyCopy` rule. This is due to our application meta-model rationale based on similarities and differences. (ii) The application of the directive finished in an expected yet incorrect state. This is possible thanks to the model’s ability to represent *inconsistent* intermediate states.

Such is the interest in using a heterogenous unified meta-model, able to copy as default behaviour and represent inconsistent intermediate states.

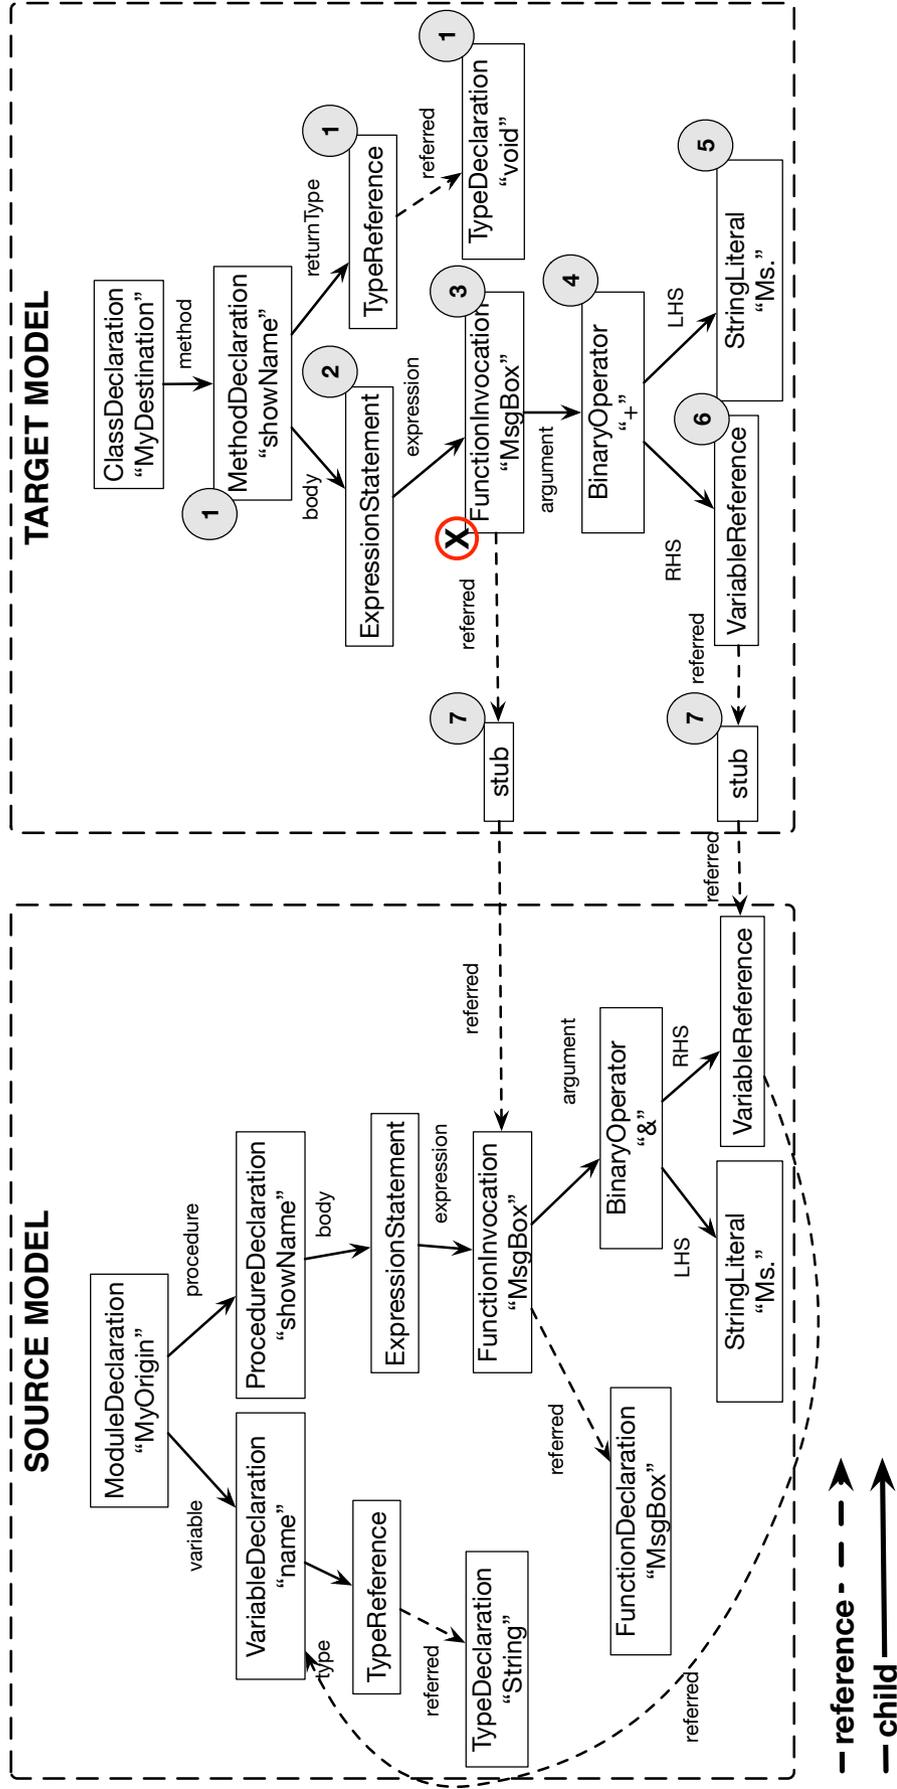


Figure 5 Produced "showName" method from "showName" sub-procedure. BinaryOperator and StringLiteral are grammatical nodes; the FunctionInvocation node is not valid in the target model. The numbers in the visualisation refer to the step when the different elements have been created. We note the difference between children and references.

4.3 | Example of the application of a map directive

Having converted the VBA “showName” sub-procedure to a Java method, the user is notified that this method is incomplete: it requires the resolution of two artefacts in the target model: (i) the invoked method “MsgBox” (ii) the referred variable “name”.

For helping to resolve these artefacts, our approach includes an action to manually express that these missing artefacts have an equivalent in the target model: the map directive.

To do this, the user applies a map directive (see Section 3.2) that is going to register a mapping between the source declaration of “MsgBox” (a *declaration* that contains no definition as it is a VBA library routine), and the target declaration of Java method “log”, within the context of “MyDestination” class. Each time a mapping is registered, either by the map directive or as a result of migration a declaration (step 8 of produce directive Section 4.2), the engine starts what we call an *adaptive phase* (systematic application of adaptive rules).

We break down the map directive into four steps, including registering the mapping and the *adaptive phase*; the first step and the two sub-steps 2.1 and 2.2 are illustrated in Figure 6.

The application of adaptive rules is made with a kind of “double lookup”.

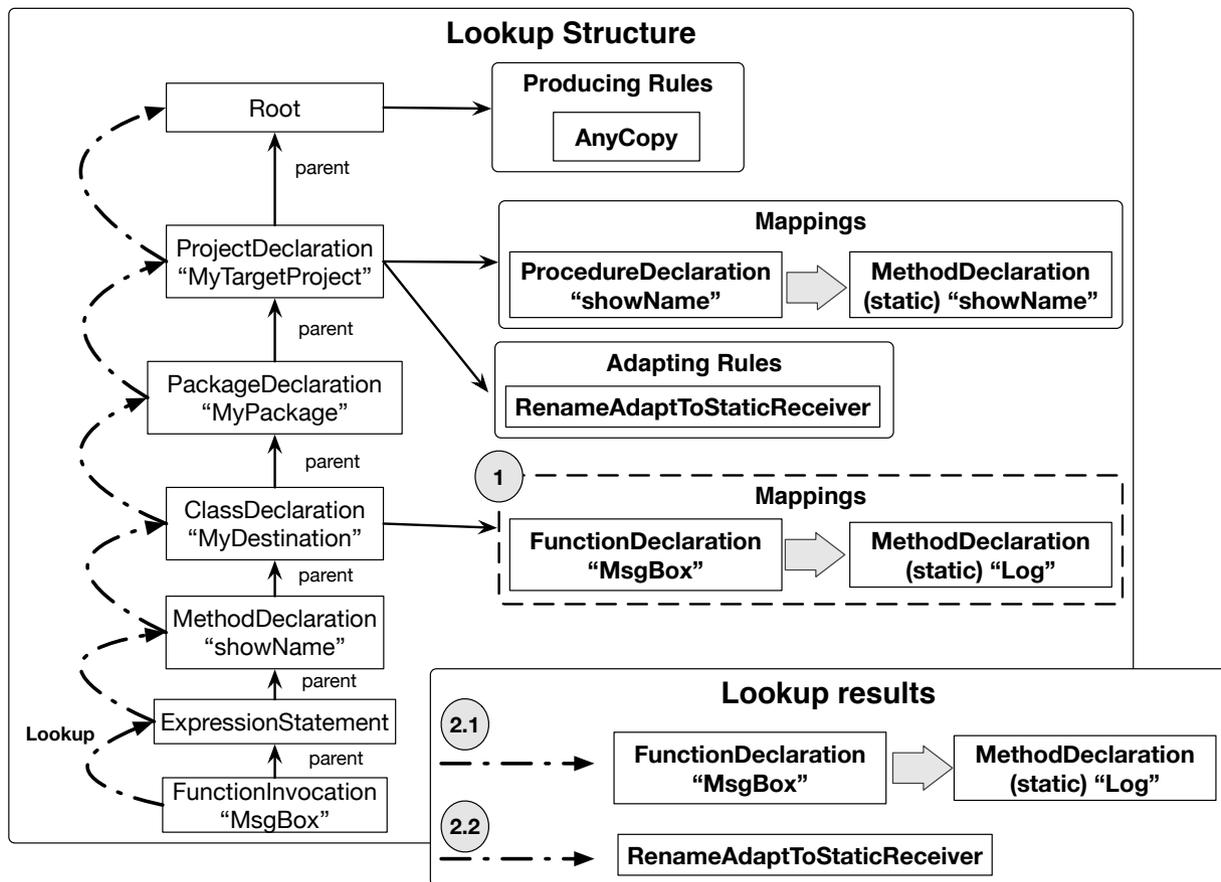


Figure 6 Map directive application: Mapping registering, mappings lookup, adaptive rule lookup. The mapping (ProcedureDeclaration “showName” => MethodDeclaration “showName”) was automatically registered during the Produce directive. The numbers in the visualisation refer to the step related. Step 1 creates a mapping. Step 2 does a double lookup selecting specific items.

Step 1 The engine registers the map between the source and target declarations in the list of mappings of the application context. (MsgBox => MyDestination.log)

Step 2 First step of the adaptive phase. The engine finds all unresolved uses of the source entity in the target model: all the reference objects pointing to stubs pointing to this source entity.

Step 2.1 The first lookup aggregates all the possible mappings affecting each of the references. As shown by Figure 6, this lookup starts from the context of each reference. This aggregated information is an ordered list. The most concrete mapping is tested first. In our example, the only mapping responding to our case is the one just installed.

Step 2.2 The second lookup tests all possible adaptive rules with each reference and each mapping in order. As shown by Figure 6, this lookup starts from the context of each reference. This process yields the first rule that tested positive for its application. If no rule is found, the process finishes. In our example, the first (and only) rule testing positive is `RenameAdaptToStaticReceiver`.

Step 3 The process finishes here if no rule is found. If a rule is found, the engine applies it. In our example, the engine applies the `RenameAdaptToStaticReceiver` parametrised with the function call and the mapping that represents equivalence between “`MsgBox`” and “`MyDestination.log`”. As an outcome, the rule replaces the function call with a method invocation, using the same arguments as the original call, and setting `MyDestination.log` method as referee.

Step 4 Last step of the adaptive phase. The engine removes all the useless stubs.

In this way, the mechanism can be applied retroactively (by searching for past Stub nodes) and to future migrations (because the produce directive registers mappings for all migrated declarations).

5 | VALIDATION

We implemented the engine and interactive UI to validate our approach and conducted several experiments. In this section, we overview the implementation and present different results.

5.1 | Implementation

To validate the approach, we implemented the heterogeneous unified meta-model ASG, the migration engine and the UI tool allowing the user to express the migrating directives.

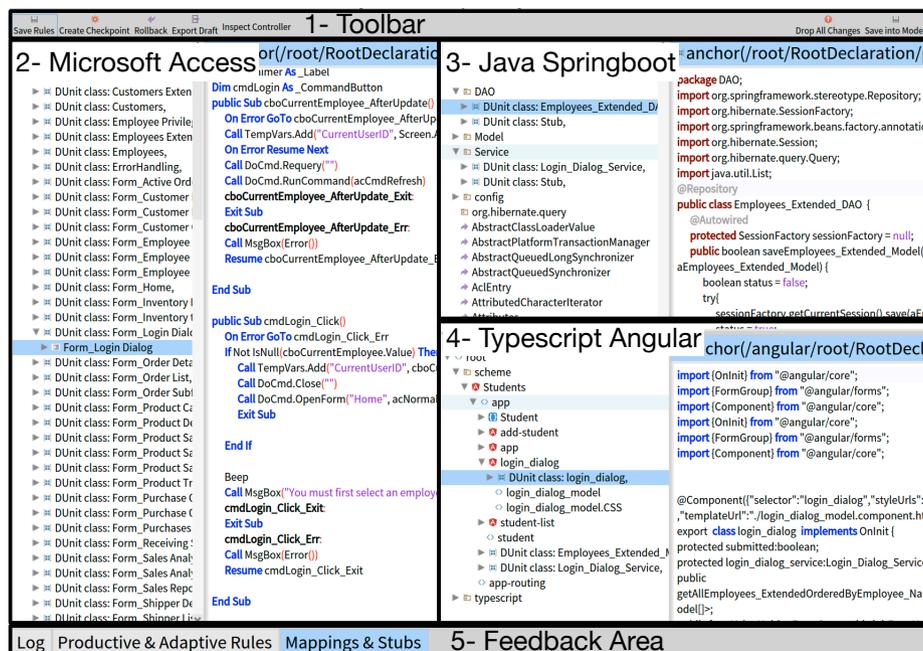


Figure 7 Tool Screenshot. The tool consists of a toolbar, three areas of model visualisation: one for MS Access, one for Java and one for Typescript, and an area for feedback

On the one hand, Figure 7 depicts the central area of the migrating tool. On the other hand, Figure 8, Figure 9, and Figure 10 depict the three feedback areas of our implementation.

Log		Productive & Adaptive Rules	Mappings & Stubs
* Stamp	Operation	Detail	
2023-02-07 10:32:41	Produce	#Result an Array(class: Login_Dialog_Service AngularComponent: 'login_dialog')	
2023-02-07 10:32:31	Produce	#Result class: Login_Dialog_Service	
2023-02-07 10:32:24	Produce	#Result class: Login_Dialog_Controller	
2023-02-07 10:32:15	Produce	#Result an Array(class: Employees_Extended_DAO)	
2023-02-07 10:32:03	CheckPoint	'Successfull checkpoint. Version: 2'	
2023-02-07 10:31:53	Produce	#Result class: Employees_Extended_Model	
2023-02-07 10:31:46	Produce	#Result class: Employees_Extended_DAO	
2023-02-07 10:31:40	Produce	#Result class: Employees_Extended_Model	

Figure 8 Tool feedback: The log panel. Here we see the outcome of the different user interactions.

Log		Productive & Adaptive Rules	Mappings & Stubs
* Pr	Origin	Rule	* Pr Origin Rule
1	8	RootDeclaration if Binary {L} Is {R} ==> {L} =={R}	1 0 RootDeclaration Ren
2	8	RootDeclaration if Binary {L} And {R} ==> {L} &&{R}	2 0 RootDeclaration Ren
3	8	RootDeclaration if Binary {L} Or {R} ==> {L} {R}	3 0 RootDeclaration Ren
4	8	RootDeclaration if Binary {L} = {R} ==> {L} =={R}	4 0 RootDeclaration Ren
5	8	RootDeclaration if Binary {L} & {R} ==> {L} +{R}	
6	8	RootDeclaration if Binary {L} < {R} ==> {L} !={R}	
7	7	RootDeclaration if Reference named vbCrLf ==> Replace with String literal '\r\n'	
8	7	RootDeclaration if Type reference named Boolean ==> Rename type reference as bool	
9	7	RootDeclaration if ModelTypeIsAnyOf:(MOModule) ==> Transform into Class	
10	7	RootDeclaration if Type reference named #_TypeUndefined ==> Rename type reference as Object	
11	7	RootDeclaration if ModelTypeIsAnyOf:(MOCClassFunction MOCClassProcedure MOFunction MOProcedure) ==>	
12	7	RootDeclaration if ModelTypeIsAnyOf:(MOGlobalVariable MOGlobalConstant) ==> Transform into Attribute	
13	7	RootDeclaration if ModelTypeIsAnyOf:(MOAlceTable MOAlceQuery) ==> a FylgjaDBEntityToAngularModel	
14	7	RootDeclaration if ModelTypeIsAnyOf:(MOAlceForm) ==> Apply Ops : Produce Angular Component based on F	
15	0	RootDeclaration if Is Unary operator prefix Not ==> {L} !{R}	
16	0	RootDeclaration a FylgjaInvocationToIsNullToNullComparison	
17	0	RootDeclaration if On Error Go to used as Try/Catch ==> Replace On Error with Try catch	
18	0	RootDeclaration if a FylgjaConditionAnd ==> a FylgjaFromJavaControllerToAngularService	
19	6	angular if ModelTypeIsAnyOf:(MOCClass) ==> Transform into TS Class	
20	0	if AlwaysTrue ==> Copy	

Figure 9 Tool feedback: The rules panel. On the left panel, we observe the productive rules. On the right one, the adaptive rules

- 1- Toolbar: tools for the project.
- 2- The left panel shows the structure and code of the MS Access project.
- 3- The right-top panel shows the structure and code of the Java-SpringBoot project.
- 4- The right-top panel shows the structure and code of the TypeScript-Angular project.
- 5- Feedback Area Figure 8 shows logs of the actions applied, informing errors and successes. Figure 9 shows the rules valid in a selected context (a class by example). Figure 10 shows the mappings and stubs tracked so far and valid in a given context.

In the feedback area, we can find three tabs. In the log tab, we find all the successful and failed interactions.

The production tab is bound to the last selected declaration within any model. It shows all the applicable productive rules available from this declaration context.

The mappings & references tab is also related to the last selected declaration within any model. The right side shows the applicable adaptive rules available from this declaration context. The left side shows all the references that are not resolved within this declaration.

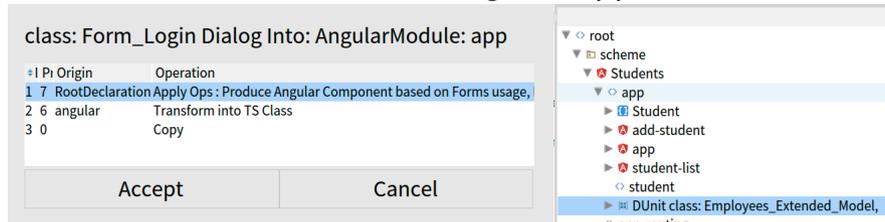
Using the tool.

The user can drag and drop a Declaration from one model to another. Once the drop is done, a popup asks the user to tell which directive she wants to apply: map or produce. In implementing the approach, we decided to allow the user to have different levels of control over the lookup of rules when applying the produce directive. (i) automatic lookup: the lookup uses the first found result. (ii) multiple choice: the lookup prompts the user to choose the rule to be applied. (iii) debugging: the lookup prompts the user to choose each rule.

Log		Productive & Adaptive Rules	Mappings & Stubs
		Auto Bind	
References (26)		Mapping Rules (36)	
#	Reference	#	Origin Mapping
1	AcCommand->EnumEntry:acCmdRefresh(->AcCommand)	1	RootDeclaration a Set(!;Form_Login Dialog.c
2	AcFormView->EnumEntry:acNormal(->AcFormView)	2	RootDeclaration a Set(!;cmdLogin_Click.MO
3	Conversion->Function: Error/1->BSTRString	3	RootDeclaration a Set(T;DUnit class: Employ
4	DoCmd->Class Function: Close/3pt: Void	4	RootDeclaration a Set(V;cboCurrentEmploye
5	DoCmd->Class Function: Requery/1pt: Void	5	RootDeclaration a Set(V;if.MOBlock => V;if.Ui
6	DoCmd->Class Function: RunCommand/1pt: Void	6	RootDeclaration a Set(V;if.MOEmptyBlock =>
7	Form_Login Dialog->Form_Login Dialog->ATTR:cboCurrentEmployee(->_Combobox)	7	RootDeclaration a Set(!;Form_Login Dialog.c
8	Interaction->Function: Beep/0pt: Void	8	angular a Set(T;NWD.int16 => T;angu
9	Interaction->Function: MsgBox/5->USERDEFINED	9	angular a Set(T;NWD.BSTRString =>
10	VbMsgBoxStyle->EnumEntry:vbOKOnly(->VbMsgBoxStyle)	10	angular a Set(T;NWD.Collection => T
11	_Application->_Application->ATTR:DoCmd(->DoCmd)	11	angular a Set(T;NWD.Byte => T;angu
12	_Application->_Application->ATTR:Screen(->USERDEFINED*)	12	angular a Set(T;NWD.LongLong => T
13	_Application->_Application->ATTR:TempVars(->TempVars)	13	angular a Set(T;NWD.ErrObject => T;
14	_Combobox->_Combobox->ATTR:Value(->Win32Variant)	14	angular a Set(T;NWD.int => T;angula
15	_TempVars->Class Function: Add/2pt: Void	15	angular a Set(T;NWD.Object => T;an
		16	angular a Set(T;NWD.uint32 => T;an
		17	angular a Set(T;NWD.Err => T;angul

Figure 10 Tool feedback: The stubs and mappings panel. On the left part, we see the Reference table, which shows the elements referred to by the stubs. On the right panel, are the established mappings.

(a) produce directive:
from Microsoft Access Form Login to app module



(b) produce directive:
from Microsoft Access Form Login to Java DAO package

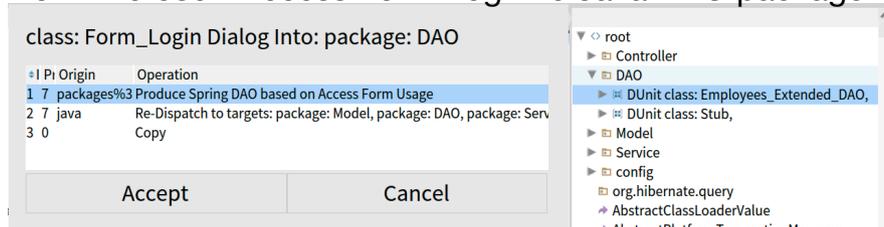


Figure 11 Choosing a rule. The figure shows the GUI prompting the user to choose the rule to apply. The rules displayed in the table are all available in the scope with positive conditions. We can see that regardless being the source the same entity, the rules in (a) are different to those in (b) as the target is different.

In both (ii) and (iii) cases, the user chooses from all the rules available in the scope with positive conditions.

Figure 11 shows the GUI prompting the user to choose the rule to be applied in two different cases: Figure 11 (a) prompts for the rule when migrating the Form login to the “app” Angular module while Figure 11 (b) prompts for the rule when migrating it to the “DAO” Java package.

5.2 | Library and paradigm partial migration

Presented in *Bragagnolo et al*¹⁵. We produced 53 test unit tests in Microsoft Access using the 53 most used library artefacts in an industrial project, based on analysing and sorting the artefacts by the number of uses. All the function calls are automatically transformed into method invocation. We defined five productive rules and five adaptive rules. We performed two migrations of these tests: one to Java and one to Pharo.

Rules.

We defined four *Productive* rules: AnyCopy, ModuleToClass, FunctionToMethod, GlobalToAttribute and CopyReplaceBinaryOperator. We defined five *Adaptive* rules: SimpleRename, RenameAdaptToStaticReceiver, RenameAdaptToThisReceiver, RenameAdaptToArgumentReceiver and Autowrap. The autowrap rule generates the skeleton of whatever used element defined in a library that has no mapped equivalent in the target. We installed **the same rules** for both targets, Pharo and Java. The only difference between the configurations was the source and target ASG mappings.

Mappings.

The mappings differ by target since the available libraries are different. Table 1 shows few of them. We found mappings for functions, types and constants. Please address *Bragagnolo et al*¹⁵ for the mapping details.

Table 1 Type mappings Mapping examples table

VBA	Java	Pharo
VBA.Single	java.lang.Float	Kernel-Numbers.Float
VBA.String	java.lang.String	Collections-Strings.String
VBA.ubyte	java.lang.Byte	Kernel-Numbers.SmallInteger

Results.

Table 2 Translation Results. Mapped: number of library elements mapped. Parsed: number of results that are verified by a parser. Compiled: number of results that are successfully compiled in the target technology. Test Result: Number of compiled tests's results: Success, Failure or Error.

Test case	Total	Mapped	Parsed	Compiled	Test Result		
					Success	Failure	Error
Constants	14	1	14	1	1	0	0
Functions	21	11	21	16	8	3	5
Type	18	13	18	9	4	4	1
Java Total	53	25	53	26	13	7	6
Constants	14	4	14	14	3	0	11
Functions	21	14	21	21	10	3	8
Type	18	13	18	18	7	10	1
Pharo Total	53	31	53	53	20	13	20

Table 2 gives the results per target. By column: *Mapped*: the number of library elements mapped. *Parsed*: the number of results verified by a parser. *Compiled*: the number of results successfully compiled in the target technology. *Test Result*: the number of compiled tests' results: Success, Failure or Error.

There is a direct relation between the number of mapped entities and the results of the number of migrated tests that are successful, failure or error. For Java, we mapped 47% of entities. 58% of them in Pharo. For both targets, the produced source code is verified by the SmaCC Java³ parser and the Pharo parser. With ten rules that are the same for two different languages, we produce 100% of parseable code. The Java compiler compiles 47% of the translated testing methods. The Pharo compiler compiles 100%. The compiling rate seems tightly related to the mapping in the case of Java since the compiling errors come from wrong typing. We argue that most of them stem from the MS Access type named Variant, which we mapped to Object in Java. MS Access Variant accepts any value but keeps the type of the element (object or not). In Java, the type Object accepts only objects (no primitive values allowed), but if a variable is typed as Object, it loses all singularity.

³<https://github.com/j-brant/SmaCC>

5.3 | Migrating Tables and Queries to a Java back end and a Typescript frontend

We migrate all the tables and queries of the Microsoft Northwind project ⁴. Northwind consists of 20 tables and 27 queries.

Rules.

In this case, the *Productive* rules differ for the Java back end and Typescript front end.

Back end

TableOrQueryToJavaPersistentModel Defines a class with persistent mappings based on the given Table or Query. It adds public accessors.

TableOrQueryToDAO Defines a DAO which uses the class produced by the previous rule. It defines default methods for CRUD actions.

Front end

TableOrQueryToClass Defines an exported public class based on the given Table or Query. All the attributes are public for accessing the content directly.

The *Adaptive* rule is the same for both the front and back end: SimpleRename. This rule is in charge of adapting all the types based on mappings.

Mappings.

The mappings for this experiment are presented in Table 3.

Table 3 Mapping examples table

VBA	Java	Typescript
Access.dbText	java.lang.String	string
Access.dbMemo	java.lang.String	string
Access.dbDate	java.time.LocalDate	Date
Access.dbAttachment	java.sql.Blob	any
Access.dbInt	int	number
Access.dbDouble	double	number

Results.

We present the experiment's results in Table 4. On the *Java* part of the experiment, we want to note that the queries generated by DTO were not compiling because we did not recognise the primary key of these tables, making it impossible to add the expected annotation `javax.persistence.@Id` to any attribute. We solved this problem by manually adding the annotations.

We tested each DAO manually, having successful results, which means that the code parses, compiles and executes as expected. On the *Typescript* part of the experiment, we successfully produced all the required classes without any parsing or compiling errors.

Table 4 DAO and DTO creation results.

Kind	Number	Java DTO	Java DAO	Typescript DTO
Tables	20	0 error	0 error	0 error
Queries	27	27 error	0 error	0 error

Discussion.

In our migration, we do not foresee mapping java persistence for using a full object model since migrating code is primarily used as records. This is why in the back end, we create one model entity per table or query (to be used as a data transfer object - DTO) and a data access object (DAO) to list all the elements, save, update and delete. For the front end, we only need to define the DTO objects to be able to consume information from the back end. This is why we only produce classes.

⁴Microsoft Access Northwind traders, learning example

5.4 | Migrating a full Form: Form Login

We propose a validation that aims to give a holistic study case. We fully migrated the login form MS Access Northwind Traders (depicted by Figure 12) to Java+Springboot and Typescript+Angular.

This migration is simple, but it includes all the migrating aspects of our industrial case of migration: Splitting the application into the front end and the back end. Produce a back end that provides the data required by the front end. Produce a front end to consume data the back end provides and show an equivalent UI.



Figure 12 Login northwind

Rules.

This experiment requires more *Productive* rules than previous experiments and of different natures.

Language translation: 6 rules in the front end translating the source code of the event handling procedures. *Data access migration*: 2 rules in the back end and 1 in the front end creating DAOs and DTOs like in the previous experiment. *Data requirement*: 3 rules in the back end and 1 in the front end creating SpringBoot and Angular services and controllers to interact based on the data required by the Form. *UI* 1 rule in the front end leveraging casino¹⁶ that migrates the UI. Casino works from a GUI model we generate by analysing the migrating form when the rule is applied.

We provide a single *Adaptive* rule: SimpleRename in both the back and front end for migrating the types used.

Mappings.

This experiment uses the same mappings already presented by Table 3.

Results.

The visual and behavioural result of the migration is the angular component shown in a browser in Figure 13. The shown widget has a dropdown list filled up with backend content and a login button, which is also functional.

We provide the execution of the backend endpoint in Figure 14 to illustrate the migration result. We note that by choosing the target of each part of the source application, we split the application in two successfully. No GUI concern was migrated to the backend. No database concern was migrated to the front end.

Making it work.

The outcome of this example is not immediately compiling and working. We could achieve this by adding more rules. However, the missing rules would be particular and hard to reuse. Therefore, we decided to go through the last mile by hand.

The modifications done by hand are the following: *Back end* Adding imports (some imports were not calculated), fixing the column names (in access names with special characters are allowed. Our test runs in a Postgres database, which does not allow this). *Front end* Adding imports, configuring the angular modules and routing to add the component to the website. Changing a call to a method that does not exist by one that exists (to do it automatically would require a new adaptive rule); adding two functions required to fill up the HTML select control with the data coming from the back end: a function that extracts the ID from the DTO, and a function that gives a string to display from the same DTO. To make it work took approximately 30 minutes for one of the authors.

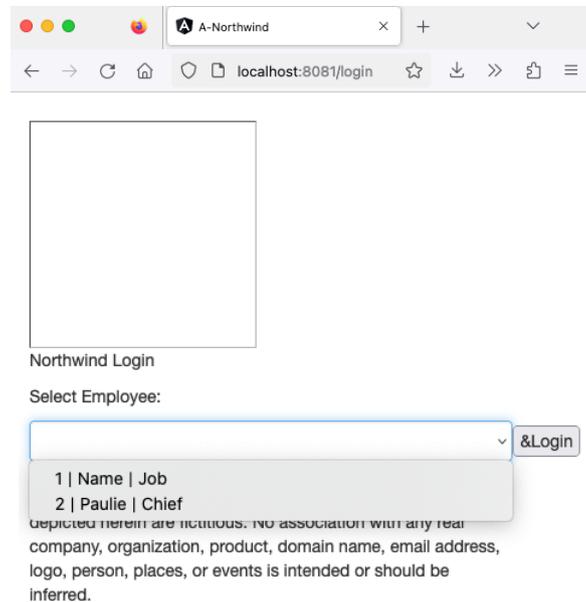


Figure 13 Generated Front End. The figure shows the visual aspect of the migrated version of the widget. It also shows that the dropdown list has been filled with backend content.

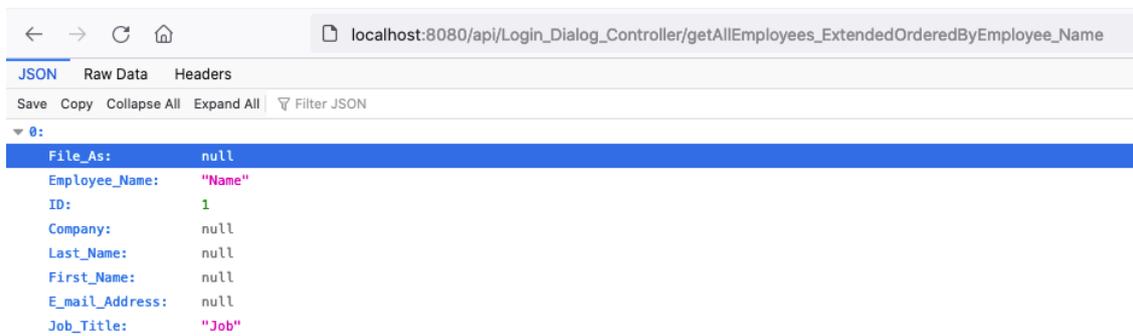


Figure 14 Generated Back-End API request. The figure shows the manual invocation of the backend. The endpoint called is the one required by the Form login to load the combo box. The response is a JSON collection.

Discussion.

We can see that the component follows a similar layout to the original and has the same amount of components. It does not follow the same visual configuration since, for this experiment, we did not make any effort to migrate the visual design to CSS. However, we can access such information from the source model, and the Casino visual model is designed to handle it.

The backend includes a controller handling the login end-point. The implementation we provide now is just an exception announcing that a method must be implemented. In the source application, the click-handling procedure delegates the information to a library function. Our backend implementation handling is a method that responds to the same signature as the source library function but has no other behaviour than a runtime exception prompting the developer to implement its behaviour.

6 | DISCUSSION

6.1 | Programming paradigm migration

In the project, we must consider migrating from a procedural paradigm (VBA) to an Object-Oriented paradigm (Java, TypeScript or Pharo). This is a complex issue that requires defining correct abstractions (classes) from variables and functions scattered in a procedural application (e.g.,^{17,18,19,4}). There is no well-accepted solution for this, and we left this issue to the migration tool user. This means the user is responsible for abstracting classes from the legacy application and migrating the relevant variables and functions to their appropriate classes. As we saw in the preceding section, the tools assist in migrating functions to methods. One tricky point remains for the invocation of methods which differ from function invocations in that they need a receiver.

We introduced three adaptive rules to handle some simple cases:

RenameAdaptToStaticReceiver: This rule converts a function invocation to a method invocation using the class of the (static) method as the receiver. Presented in Section 4.1.2;

RenameAdaptToSameClassReceiver: This rule converts a function invocation to a method invocation using “this” or “super” as receiver, according to where the method is defined;

RenameAdaptToArgumentReceiver: This rule converts a function invocation to a method invocation using one argument of the function invocation as the receiver of the method invocation. The execution of the rule includes asking the user which argument to use as the receiver.

One could think of a more complex rule (for example, one that would create an object on the spot), but it falls outside the scope of the current project.

6.2 | Rule morphology

Traditional model transformation approaches aim to be fully automatic (“push button” approach), and reproducible²⁰. For this, the engine applies all the rules to every entity in the source model that it can handle. As such, transformation rules are typically composed of at least two parts (e.g.,²¹): (i) Left-Hand Side (LHS), a selection part specifying which source model element must be matched by the rule; and (ii) Right Hand Side (RHS), a transformation part specifying what target model elements must be generated and how they should be initialised. The LHS part can be tricky to define in these approaches when one needs to precisely pinpoint a subset of entities, or a unique entity, for specific rules.

Our rules share this characteristic: They have an LHS part (what we call condition) that filters the kind of entities they can apply to. But because our rules are started by a user directive on a source entity and have an application context, the filtering part is usually straightforward, for example, only checking for a specific entity type.

6.3 | Rule application

Our approach relies on two features of the application of the rules: their ability for reentrancy and the scoping of validity.

Reentrancy. A rule can always delegate its application over an entity’s children. Like this, the rule that can create a Java method from an Access function will rely on other rules to create the parameters of the method or its body.

Scoping. Migration rules work from a source model to a target model. In our approach, the *source model* of a rule is often, but not always, the legacy application model. Section 4.3 gives examples of the source model of a rule being one of the migrated application models. Symmetrically, the *target model* will usually be one of the migrated application models (actually, it is always the case in this paper). Still, we envision cases where the legacy application model could be the target of a transformation rule, for example, in a partial migration where parts of the legacy and migrated applications run together (see, for example,²²). In such a case, the legacy application may need to be modified to reference some declaration node in the migrated application.

Rules are defined and applied within a *context*. The context can be the entire application, a specific package, a specific class, ... Concretely, the context can be any declaration node in the model of an application⁵. When the context of the rule is the top-level node of an application model, this allows one to handle the specific programming language of this application. For example, one could have an “RPC”⁶ package with special migration rules to add Spring Boot annotations automatically.

⁵In the current implementation, we store the contexts in a “twin tree” of the target AST (i.e. with the same branches). For simplicity, this paper considers that contexts are stored in the AST.

⁶Remote Procedure Call: When a method in a program calls a method defined in a different program, usually on another computer

6.4 | Modelling

Declarations and grammatical entities. The boundary between *declaration* and *grammatical* is not clear cut. It depends on whether we want to be able to handle references on an entity. For example, in Java, primitive types are part of the language's grammar. Still, they are modelled as declaration nodes because entities like variables have reference to them in their declarations and also because we might need to migrate them to other types (e.g., Java boolean to C int). Conversely, the "+" operator is currently modelled as a grammatical node. Still, it could be modelled as a function (a declaration node), for example, when this operator represents the string concatenation in Java. In this case, the "+" in the Java expression "the answer is: "+42" would be modelled as a reference to a particular java function *stringConcatenation*, with two arguments, so that one would be able to add migration rules for it.

We avoid modelling each programming language independently, trying to have, as much as possible, all language ASGs represented using the same types. In this, we diverge from what a meta-model like Famix does (one specific meta-model for each programming language²³) and are closer to what KDM of the OMG proposes (one unique meta-model²⁴). Our rationale for having only one meta-model is that when migrating from one language to another, we want to keep as much of the source ASG as possible. Having the same nodes for both languages helps us do that as it allows us to "migrate" by just copying an entire branch of the ASG (e.g. a complex numerical expression) from the source model to the target one. This allowed us to use the same rules for this experiment when migrating to Pharo and Java.

Yet, it is not always possible to do it because, as noted in²³ "various programming languages have minor semantic differences regarding how they implement [some programming concepts]". The same goes for ASG nodes; for example, some languages have an "elseif"⁷ statement (or part of the statement), and others represent it as an "else" containing a new "if". (Note that KDM or GASTM do not offer an "elseif" node.)

To accurately migrate a language, these minor semantic differences must be modelled. Therefore, we use the same ASG nodes for different languages as long as the nodes do have the same semantics. When this is not the case (e.g., Java classes and TypeScript classes), we create different ASG nodes to be able to differentiate them in the migration rules.

7 | FUTURE WORK AND CONCLUSION

7.1 | Future work

Our most probable direction is the development of a third directive: delegate. This directive is expected to propose the migration and further replacement of the source model to become an invocation of the freshly migrated piece of code. This would encourage and ease the idea of migration as a deflating process that takes the complexity out of the source.

During this work, we recognise many critical challenges involving the user experience, the thing that is chiefly important in the approach's productivity, such as: (i) What visualisations would allow the developer to acquire **enough** information to take architectural decisions? (ii) What visualisations would allow the developer to understand the immediate impact of applying a directive? (iii) What kind of validation algorithm would inform the developer of the current problems of the migration? (iv) Can migration rules be inferred from the UI interaction? (v) How do we graphically debug rules?

We also got many process-related questions for the future, such as: (i) How can we measure the process of migration? When does a migration finish? (ii) Can we generate tests automatically knowing that we have a source implementation known to be correct?

7.2 | Conclusion

Complete system migration is an overwhelming, complex and intimidating mission to achieve.

Our particular case of migration is a case with multiple layers of software migration, as we described in Section 2.1. Translate language implies using different grammatical rules for saying the same and using entirely different constructions to think of problems when the paradigm behind them is distant.

Different languages grow surrounded by different communities addressing problems in different and often incompatible ways, favouring the apparition of different libraries and frameworks, from the definition of their concerns to the definition of their API. Different deploying environments and architectural designs highlight different features and requirements impacting the inbound design in radically different ways, such as the differences between a monolithic application and a microservices one. All these gaps impact at many levels; not only must the code change ultimately its nature: the grammar, the APIs, the concepts and the concerns but the developers as well. We claim that for large mission-critical projects with such a gap, fully automatic migration is almost as undesirable and unreliable as software rewriting since an automatically migrated mission-critical

⁷For example, "elseif" in Visual Basic or PHP, "elif" in Python

project would have no immediate maintainers nor developers at most of the aspects of the target projects: source code, library, architecture, building, shipping, and deploying. Even from the point of view of the database, which we do not migrate, the code interaction would change, meaning that the knowledge of the database administrators may fail with the new versions of the code.

To tackle down, we propose an iterative, interactive migration process (Section 2.2) based on three central ideas: developer control, immediate feedback, and independent life cycle of source and target systems. Our proposal has as its major drawback the requirement of good migration planning. Architectural and design decisions have to be taken as in forward engineering. To promote experimentation (which is highly required for taking decisions), we reinforce the importance of not punishing bad decisions by being able to undo all modifications from both model and source-code points of view.

Our approach proposes to scope the kind of migrating approach within the design of each rule. In our migrating approach, language migration is compulsory as the intention is to abandon Microsoft Access completely. Therefore all our rules are based on white-box approaches.

However, this lack of black-box or grey-box approach rules does not mean that the approach is incompatible with it. A simple and naive example of producing a wrapper is the rule applied when applying the produce directive with a controller method from the back end into a service class in the front end.

In this article, we presented a process of migration fully supported by an interactive UI, which resolves the user migrating interactions with a transformation engine based on directives. For that, we introduced our infrastructure elements: directives, mappings/stubs, rules and multiple models (Section 3). We explained step-by-step how these elements interact to achieve the user's desired activity (Section 4)

We validate the approach (Section 5) by implementing the UI and transformation engine (Section 5.1) and applying it to three different cases of migration: We migrated library usage and paradigm invocations from Microsoft Access to Java and Pharo languages (Section 5.2). We migrated the table and query objects from Microsoft Access to Java + Java persistence and Typescript (Section 5.3). We migrated a form and a query from Microsoft Access to Java+Springboot and Typescript+Angular, using the tool to create (i) the back end and the front end code and (ii) the required code to interact between the front end and back end (Section 5.4).

References

1. Terekhov AA, Verhoef C. The realities of language conversions. *IEEE Software* 2000; 17(6): 111–124. doi: 10.1109/52.895180
2. Brant J, Roberts D, Plendl B, Prince J. Extreme Maintenance: Transforming Delphi into C#. In: ; 2010.
3. Kontogiannis K, Martin J, Wong K, Gregory R, Müller H, Mylopoulos J. Code Migration through Transformations: An Experience Report. In: CASCON '98. IBM Press; 1998: 13.
4. Zou Y, Kontogiannis K. A framework for migrating procedural code to object-oriented platforms. In: IEEE. ; 2001: 390 to 399.
5. Martin J, Muller HA. C to java migration experiences. In: IEEE. ; 2002: 143–153.
6. Trudel M, Furia CA, Nordio M, Meyer B, Oriol M. C to OO translation: Beyond the easy stuff. In: IEEE. ; 2012: 19–28.
7. De Lucia A, Francese R, Scanniello G, Tortora G. Developing legacy system migration methods and tools for technology transfer. *Software: Practice and Experience* 2008; 38(13): 1333-1364. doi: <https://doi.org/10.1002/spe.870>
8. Zaragoza P. *Model-Driven Software Migration of Legacy Applications to Micro-Services Architectures*. PhD thesis. Informatique Montpellier, 2022.
9. Zhang Z, Yang H. Incubating services in legacy systems for architectural migration. In: IEEE. ; 2004: 196 to 203.
10. Razavian M, Lago P. A Lean and Mean Strategy for Migration to Services. In: WICSA/ECSA '12. Association for Computing Machinery; 2012; New York, NY, USA: 61 to 68
11. Verhaeghe B. *Incremental Approach for Application GUI Migration using Metamodel*. PhD thesis. Université de Lille, 2021.
12. Garcés K, Casallas R, Álvarez C, et al. White-box modernization of legacy applications: The oracle forms case study. *Computer Standards & Interfaces* 2017: 110–122. doi: <https://doi.org/10.1016/j.csi.2017.10.004>
13. Moore , Rugaber , Seaver . Knowledge-based user interface migration. In: IEEE Comput. Soc. Press; 1994: 72–79
14. Moore MM. Rule-based detection for reverse engineering user interfaces. In: IEEE. ; 1996: 42–48.

15. Bragagnolo S, Anquetil N, Ducasse S, Mustapha D. Reporting Context Aware Partial Translation engine based on immediate and delayed Rule application. tech. rep., Inria; 2022.
16. Verhaeghe B, Shatnawi A, Seriai A, et al. From GWT to Angular: An Experiment Report on Migrating a Legacy Web Application. *IEEE Software* 2021. doi: 10.1109/MS.2021.3101249
17. Duddy K, Gerber A, Lawley M, Raymond K, Steel J. Declarative transformation for object-oriented models. *Transformation of Knowledge, Information, and Data: Theory and Applications* 2004.
18. Fox G, Li X, Qiang Z, Zhigang W. A prototype of Fortran-to-Java converter. *Concurrency: Practice and Experience* 1997; 9(11): 1047–1061. doi: 10.1002/(SICI)1096-9128(199711)9:11<1047::AID-CPE348>3.0.CO;2-V
19. Grosse-Kunstleve RW, Terwilliger TC, Sauter NK, Adams PD. Automatic Fortran to C++ conversion with FABLE. *Source Code for Biology and Medicine* 2012; 7(1): 5. doi: 10.1186/1751-0473-7-5
20. Siikarla MP, Systa TJ. Decision Reuse in an Interactive Model Transformation. In: ; 2008: 123-132
21. Czarnecki K, Helsen S. Feature-based survey of model transformation approaches. *IBM systems journal* 2006; 45(3): 621–645.
22. Verhaeghe B, Shatnawi A, Seriai A, et al. A Hybrid Architecture for the Incremental Migration of a Web Front-end. In: INSTICC. SciTePress; 2022: 101-110
23. Anquetil N, Etien A, Houekpetodji MH, et al. Modular Moose: A new generation of software reengineering platform. In: No. 12541 in LNCS. ; 2020
24. OMG . Knowledge Discovery Meta-Model (KDM). 2011.

How to cite this article: S. Bragagnolo, N. Anquetil, S. Ducasse, A. Seriai, and M. Derras (2023), Interactive, Iterative, Tooled, Rule-Based Migration of Microsoft Access to Web Technologies, *Software: Evolution and Process.*, 2023;00:1–6.