

# The BCerT Solution for Bidirectional Model-Driven Transformation

Akram Idani<sup>1</sup>

<sup>1</sup>Univ. Grenoble Alpes, Grenoble INP, CNRS, VERIMAG, F-38000 Grenoble, France

## Abstract

The TTC'2026 Families to Persons case study focuses on bidirectional transformations with an emphasis on concurrent model synchronization, where simultaneous edits may occur on both source and target models. In this context, ensuring consistency while handling potentially conflicting updates is a key challenge. We present the BCerT solution providing a formal state-based approach to model transformation using the B method. Transformation rules are specified as invariants and preconditioned operations that are proved correct via theorem proving. Regarding execution, our approach follows a rule-based propagation strategy, where synchronization is achieved by leveraging the animation and constraint solving capabilities of the ProB model-checker. Our solution supports both forward and backward transformations as classical transformation cases, while extending to concurrent synchronization. Correctness is ensured by construction: since the transformation is proven, it obviously preserves both structural and semantic invariants. Testing is therefore not required for verification purposes, but is used in our approach for validation. To this end, we combine B with CSP, where CSP acts as an orchestrator of B operations, enabling systematic validation against the TTC benchmark test suite.

## Keywords

Bidirectional model transformation, B method, Model synchronization, CSP, ProB

**Artifacts and tutorials:** <https://bcert-meeduse.github.io/ttc2026.html>

## 1. Introduction

Bidirectional model transformations aim at maintaining consistency between two related models when one or both of them evolve. They are central in MDE, where different views, languages, or artefacts may represent complementary aspects of the same system. Beyond classical forward and backward transformations, modern approaches must also support concurrent synchronization, where both models may be modified independently before consistency is restored.

The TTC'2026 Families to Persons (F2P) case study provides a benchmark for evaluating such transformations. The source domain describes families and their members, while the target domain describes persons. Besides forward and backward propagation, the benchmark includes incremental, roundtrip, and concurrent synchronization scenarios specified through pre- and post-conditions.

Our solution is based on the B method [1], a state-based formal method in which systems are described using abstract machines, variables, invariants, and operations. In our approach, invariants define structural and semantic consistency properties, while operations define state transformations that are proved to preserve them. We rely on BCerT [2], a model-transformation framework built on top of B and integrating EMF with the ProB model-checker [3]. We used the F2P transformation as a pedagogical example in a paper accepted at SLE 2026 [4] in order to illustrate the approach. The TTC'2026 benchmark provides an opportunity to evaluate both the approach and the underlying tool on a broader range of realistic transformation scenarios, including forward, backward, incremental, roundtrip, and concurrent synchronization cases.

A central aspect of our contribution is the distinction between verification and validation. Transformation rules are specified in B and proved correct with respect to the invariants of the formal model, guaranteeing structural consistency by construction. However, the TTC benchmark evaluates observable behaviours expressed through pre- and post-conditions. To address this aspect, we combine

B with CSP<sup>1</sup> [5]. Transformation rules are expressed as B operations, while CSP orchestrates their execution and enables the systematic validation of all benchmark scenarios.

The rest of the paper is organized as follows. Section 2 briefly introduces Meeduse and the BCerT extension. Section 3 presents the proposed solution to the TTC’2026 F2P case, including the involved meta-models, transformation rules, propagation strategies, and CSP-based test suites. Section 4 reports execution results and implementation metrics. Finally, Section 5 concludes the paper and discusses the scalability issue.

## 2. Brief introduction of Meeduse and BCerT

### 2.1. Models, meta-models, and semantics in Meeduse

The B method [1] is a state-based formal method in which systems are described using abstract machines, variables, invariants, and operations. Correctness is established through proof obligations ensuring that operations preserve the specified invariants.

Our tool, Meeduse [6], combines EMF and B through a compilation process that translates meta-models and models into executable B specifications. The overall architecture is illustrated in Figure 1. From an ECore meta-model (or an Xtext grammar), Meeduse generates a functional B machine (referred to as *Static Semantics*) that captures the structural aspects of the language. The operational semantics can then be defined either in the same machine or in a separate machine (referred to as *Dynamic Semantics*) that includes the functional one. This semantic layer enables formal reasoning about the correctness of the DSL using theorem-proving tools such as Atelier B.

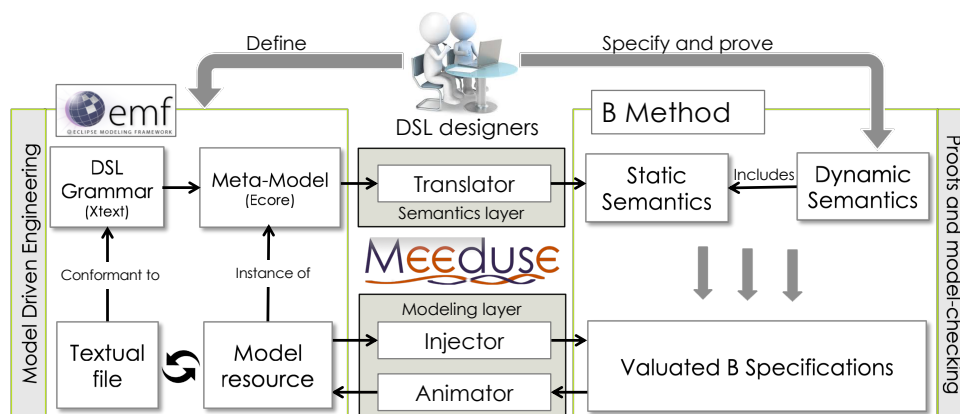


Figure 1: Meeduse architecture (taken from [6])


To support model execution, Meeduse embeds ProB [3], an animator and model checker of the B method; and uses it transparently from the user’s perspective. Given a model, Meeduse produces a data refinement of the functional B machine. This results in a valued B specification that is semantically equivalent to the input model. Figure 2 illustrates this process: (1) the structural part of machine *MeeduseTuto* is generated from a meta-model containing a class **Counter** with an integer attribute **value**; and (2) the refinement *MeeduseTuto\_r* is generated from a model containing one instance of this class, where **value** is equal to 0. The operations **inc** and **dec** define the dynamic semantics of the meta-model incrementing and decrementing attribute value, within the bounds specified by the invariant.

From the valued specification (right hand-side of Figure 2), ProB computes the enabled operations in the current state and produces an execution trace corresponding to all possible applications of these operations. At each step, Meeduse observes the state transition computed by ProB and propagates the corresponding updates back to the EMF model.

<sup>1</sup>Communicating Sequential Processes



and a CSP process. As a consequence, an operation can be executed only when it is both enabled by the B machine and permitted by the CSP controller. The process shown in Figure 4 enforces an initial choice between incrementing and decrementing, and then restricts the execution to the corresponding direction. In CSP, the operator  $\square$  denotes an external choice between alternative behaviors, while the prefix operator  $a \rightarrow P$  specifies that event  $a$  must occur before the process continues as  $P$ . The resulting behavior is illustrated on the right-hand side of the figure, where the state space of the counter is reduced to a monotonic evolution depending on the initial choice.



**Figure 4:** CSP control of the execution flow of the MeduseTuto machine

Recent versions of the ProB Java API provide support for the CSP||B combination. This feature is exploited in the BCerT extension to allow the use of CSP for orchestrating the execution of transformation rules.

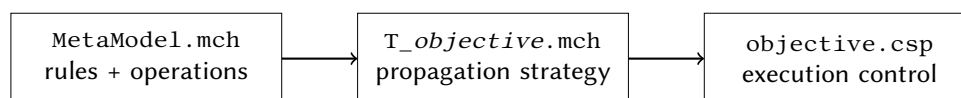
### 3. The BCerT solution to the TTC'2026 F2P case

Detailed B and CSP specifications are available in the online artefacts<sup>2</sup> and can be consulted for a complete view of the formal models (look inside folder `pivot/model/files` \* .mch and \* .csp). This paper does not aim at providing a full formalisation of the solution; instead, it focuses on giving an overview of how our approach addresses the main challenges of the benchmark.

#### 3.1. Architecture

The TTC'2026 F2P case defines several transformation objectives, including batch forward, batch backward, incremental forward, incremental backward, concurrent synchronization, and roundtrip consistency. Preliminary scalability results are reported in Section 4. In our proposal, each objective is addressed following a common architectural pattern, structured into three layers, as illustrated in Figure 5:

- `MetaModel.mch` defines the core transformation rules together with basic model manipulation operations. It captures both the structural and behavioral aspects of the transformation, including invariants and operations that are formally proved correct.
- A family of machines `T_objective.mch` is introduced, each corresponding to a specific TTC objective. These machines define the propagation strategies that govern how transformation rules are applied in a given context.
- For each test suite of the benchmark, a CSP specification orchestrates the execution of the B operations defined in `T_objective.mch` by explicitly encoding the required execution scenarios and control flow.



**Figure 5:** Generic architecture used for each TTC objective

<sup>2</sup>[https://bcert-meeduse.github.io/bcert-beta/2026\\_TTC\\_fp.zip](https://bcert-meeduse.github.io/bcert-beta/2026_TTC_fp.zip)

This structure is uniformly applied across all objectives, ensuring a clear separation between transformation rules, their operational use, and the control of their execution.

### 3.2. Involved meta-models

The meta-models *Families* and *Persons* are provided by the contest. To bridge them as required by our infrastructure, we introduce a pivot meta-model (Figure 6). The central element is the class `Pivot`, which aggregates references to both `FamilyRegister` and `PersonRegister` through `familyModel` and `personModel` attributes. It also contains configuration parameters controlling the execution of the transformation, such as the direction (attribute `strategie`) and synchronization flags (`FAMILY_TO_NEW` and `PARENT_TO_CHILD`).

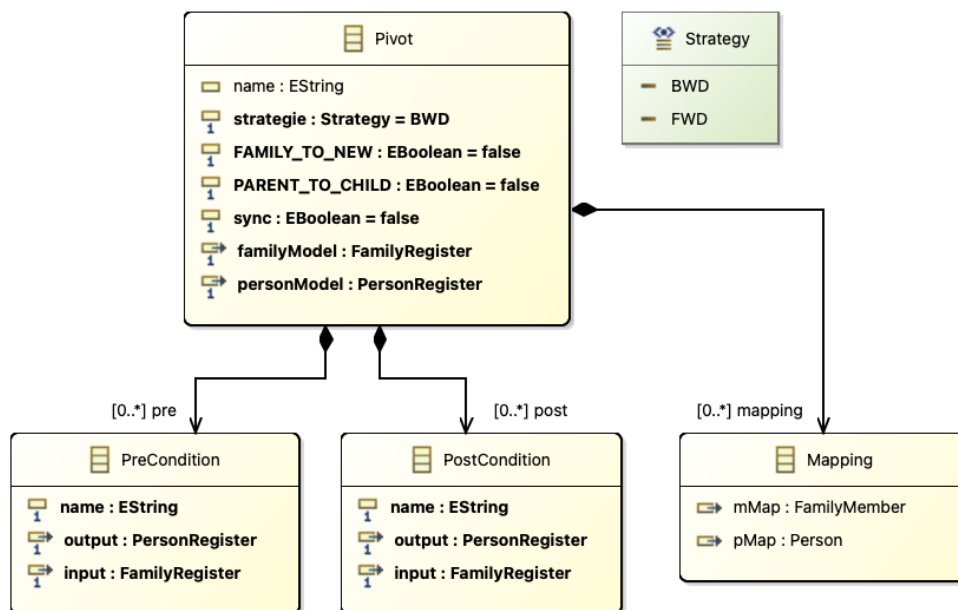


Figure 6: The pivot meta-model integrating Families and Persons

The pivot meta-model further introduces three key classes. First, `Mapping` elements explicitly represent correspondences between source and target elements, linking `FamilyMember` instances to `Person` instances. Second, `PreCondition` and `PostCondition` elements are used to describe the expected input and output configurations of test scenarios. These elements reference both the family and person registers and provide a structured way to express the conditions that must hold before and after the execution of transformation rules. Multiple `PreCondition` and `PostCondition` elements can be associated with a given `Pivot` instance. Each condition is identified by its name attribute.

### 3.3. Transformation rules in B

The B data structures (sets and relations) generated for the three meta-models are defined in the machine `MetaModel.mch`. Transformation rules are then specified as B operations over these structures, expressing how elements are created, updated, and related, while ensuring the preservation of structural and semantic invariants. These invariants capture both typing constraints (e.g., well-formed relations between families, members, and persons), structural properties (e.g., every `FamilyMember` belongs to exactly one `Family`, and every `Person` belongs to a `PersonRegister`), and consistency properties between the two domains (e.g., mapped elements must respect gender and role constraints). As all transformation rules are proved to preserve these invariants, every reachable state of the system is guaranteed to be correct by construction. We distinguish three categories of rules: forward transformation rules, backward transformation rules, and rules dedicated to concurrent synchronization and conflict resolution.

1. Forward transformation rules propagate changes from the *Families* model to the *Persons* model:
  - `Member2Person` creates a new `Person` from an unmapped `FamilyMember`, initializes its attributes, and establishes the corresponding mapping.
  - `ForwardCreateMissingPerson` completes an existing mapping by creating the missing `Person` when a `FamilyMember` is already mapped but no corresponding person exists.
  - `ForwardRename` updates the name of a `Person` to reflect changes in the corresponding `FamilyMember`.
  - `ForwardDelete` deletes a `Person` when the corresponding `FamilyMember` no longer exists, and removes the associated mapping.
  - `ForwardRepairGender` repairs inconsistencies between the gender of a `Person` and the corresponding `FamilyMember`.
2. Backward transformation rules propagate changes from the *Persons* model to the *Families* model.
  - `Person2MemberNewFamily` creates a new `Family` and a corresponding `FamilyMember` from an unmapped `Person`, when no suitable family exists.
  - `Person2MemberExistingFamily` creates a new `FamilyMember` inside an existing family that matches the `Person`'s family name.
  - `BackwardRenameMemberName` updates the name of a `FamilyMember` to match the corresponding `Person`.
  - `BackwardDelete` deletes a `FamilyMember` when the corresponding `Person` no longer exists, and removes the associated mapping.
  - `BackwardMoveMaleToExistingFamily` and `BackwardMoveFemaleToExistingFamily` move a `FamilyMember` to an existing family when the associated `Person` refers to a different family.
  - `BackwardMoveMaleToNewFamily` and `BackwardMoveFemaleToNewFamily` create a new family and move the `FamilyMember` into it when the target family does not exist.
3. Concurrent synchronization and consistency rules handle situations where changes occur independently on both models and must be reconciled.
  - `ConcurrentMatchMemberPerson` creates a mapping between an existing `FamilyMember` and an existing `Person` when they match but are not yet linked.
  - `ConcurrentDeleteMatched` removes obsolete mappings when both the source and target elements have already been deleted.
  - `ConcurrentTargetRenameWins` resolves rename conflicts by prioritizing the target side and propagating the change back to the family model.
  - `ConcurrentDeleteWins` resolves conflicts between deletion and renaming by prioritizing deletion and removing the corresponding elements and mappings.

For illustration we show the example of rule `Member2Person` in Figure 7. This forward rule is triggered when a `FamilyMember` is not yet mapped to any `Person` (i.e.,  $aMember \notin \text{ran}(mMap)$ ), in forward mode ( $\text{strategie}(pvtRoot) = \text{FWD}$ ) and outside synchronization ( $\text{sync}(pvtRoot) = \text{FALSE}$ ). It creates a fresh `Person` ( $aPerson \notin \text{Person}$ ), assigns its gender according to the source element ( $\text{isFemale}(aMember)$ ), and initializes its attributes, including a default birthdate ( $\text{birthday} := \text{birthday} \Leftarrow \{aPerson \mapsto \text{"DEFAULT"}\}$ ) and a name constructed from the family and member names. The new person is also inserted into the persons register ( $\text{persons} := \text{persons} \Leftarrow \{aPerson \mapsto \text{prsRoot}\}$ ). Finally, a new mapping is created ( $aMapping \notin \text{Mapping}$ ) and recorded consistently ( $mMap := mMap \Leftarrow \{aMapping \mapsto aMember\}$ ,  $pMap := pMap \Leftarrow \{aMapping \mapsto aPerson\}$ ), ensuring that the source element is now linked to its target counterpart.

```

Member2Person =
  ANY aMember, aPerson, aMapping WHERE
    aMember ∈ theMembers
    ∧ aMember ∉ ran(mMap)
    ∧ aPerson ∈ PERSON
    ∧ aPerson ∉ Person
    ∧ aMapping ∈ MAPPING
    ∧ aMapping ∉ Mapping
    ∧ strategie(pvtRoot) = FWD
    ∧ sync(pvtRoot) = FALSE
  THEN
    IF isFemale(aMember) THEN
      Female := Female ∪ {aPerson}
    ELSE
      Male := Male ∪ {aPerson}
    END
    Person := Person ∪ {aPerson}
    || birthday := birthday ⇐ {aPerson ↦ "DEFAULT"}
    || Persons_Person_name(aPerson) :=
      STRING_CONC([Families_Family_name(familyOf(aMember)),
        ", ", Families_FamilyMember_name(aMember)])
    || persons := persons ⇐ {aPerson ↦ prsRoot}
    || Mapping := Mapping ∪ {aMapping}
    || mapping := mapping ⇐ {aMapping ↦ pvtRoot}
    || mMap := mMap ⇐ {aMapping ↦ aMember}
    || pMap := pMap ⇐ {aMapping ↦ aPerson}
  END ;

```

Figure 7: Example of a forward transformation rule specified in B

### 3.4. Rule propagation

Depending on the considered objective, different propagation strategies are defined. For example, the basic propagation mechanism is defined by the machine `T0_MainTransformation.mch`. Unlike the machines (Figure 5) dedicated to benchmark test suites, this machine is not tied to a predefined scenario. It is intended for interactive use: the user edits an EMF model manually (a person model or a family model), launches the transformation in BCeRT, and the tool computes via ProB the operations that are enabled in the current state. All enabled rules (those issued from clause `PROMOTES`) are then proposed to the user, who chooses which one to apply.

Obviously, some choices are non-deterministic. For instance, if a mapped pair exists but their names differ, both `ForwardRename` (propagating changes from the family model) and `BackwardRenameMemberName` (propagating changes from the person model) may be applicable. In the absence of additional control, the choice between these rules is left non-deterministic.

Regarding the test objectives of the TTC benchmark, the propagation of the underlying transformation rules is controlled by a dedicated operation `propagate` defined in each `T_objective.mch` machine.

- For the **batch forward** objective (`T1_BatchForward.mch`), the strategy consists in repeatedly applying the rule `Member2Person`. This reflects the fact that the transformation is performed in a single direction without considering deletions or updates.
- For the **batch backward** objective (`T2_BatchBackward.mch`), the strategy selects between two rules depending on the configuration. If the flag `FAMILY_TO_NEW` is enabled and a suitable family already exists (*i.e.*, a family with a matching name is found), the rule `Person2MemberExistingFamily` is applied. Otherwise, the rule `Person2MemberNewFamily` is used to create a new family. This introduces a conditional and data-dependent propagation strategy.

- For the **incremental objectives** (`T3_IncrementalForward.mch` and `T4_IncrementalBackward.mch`), the propagation is non-deterministic and defined using a CHOICE construct. At each step, one applicable rule is selected depending on the current state. In the forward case, this includes rules such as `Member2Person`, `ForwardDelete`, `ForwardRename`, and `ForwardRepairGender`, while in the backward case it includes creation, deletion, renaming, and movement rules (e.g., `Person2Member*`, `BackwardDelete`, `BackwardRenameMemberName`). No explicit priority is enforced between rules. Since all transformation rules are proved to preserve the invariants of the formal model, this non-determinism affects the order in which enabled rules are applied, but not the correctness of reachable states.
- The **concurrent** objective (`T5_Concurrent.mch`) introduces a priority-based propagation strategy. Rules are not applied arbitrarily but according to a fixed ordering. First, consistency-repair rules are considered: `ConcurrentDeleteMatched`, `ConcurrentMatchMemberPerson`, and `ConcurrentTargetRenameWins`. These rules resolve inconsistencies between the source and target models. Only when none of these rules is applicable does the system fall back to standard forward or backward transformations. In this second phase, an explicit direction is chosen (`SetFWD` or `SetBWD`), and rules are applied with internal priorities (e.g., deletion before renaming, and renaming before creation). This ensures that inconsistencies are resolved before structural updates are performed. The conflict-resolution policies themselves are encoded explicitly as B operations. Consequently, alternative policies could be introduced by replacing these operations and adapting the propagation priorities, without modifying the pivot meta-model. Such changes would however require re-validation and proof obligations to ensure that the resulting transformation still preserves the specified invariants.
- Finally, the **roundtrip** objective combines both forward and backward rules in a fully non-deterministic manner. The `propagate` operation allows any applicable rule from both directions, as well as selected concurrent rules, to be applied. This reflects the absence of a fixed direction and enables exploring arbitrary sequences of transformations. However, in the test suite, the direction is explicitly controlled by invoking `SetFWD` or `SetBWD` before calling `propagate`.

Overall, propagation strategies range from deterministic (batch) to non-deterministic (incremental and roundtrip), and to priority-driven (concurrent), providing different levels of control over rule application depending on the test objective considered by the test suite.

### 3.5. Test suites

#### 3.5.1. Model mutations

In the test suites provided by the TTC benchmark, the general approach consists in starting from an initial configuration where both the *Families* and the *Persons* models are empty. Depending on the test objective, these models are then progressively evolved through a sequence of mutations (e.g., creation, deletion, and renaming) before applying propagation or synchronization steps, and finally checking pre- and post-conditions. In the reference implementation, these mutations are encoded in Java using helper operations such as `srcEdit(helperFamily::createSimpsonFamily)`.

To reproduce these mutations within our formal framework, we introduce a set of basic model manipulation operations directly in B. For example, the operation `SonNEW` (Figure 8) creates a new `FamilyMember` and attaches it as a son to an existing family identified by its name. Based on these primitives, model mutations used in the reference implementation are reproduced by invoking our B operations with concrete parameter values in the objective-specific machines. For instance, the following operation from `T1_BatchForward.mch` creates a new family named “Flanders” and adds a son “Rod” (Figure 9).

```

SonNEW(aFamilyName, aSonName) =
PRE
  aSonName ∈ STRING ∧
  aFamilyName ∈ STRING
THEN
  ANY aFamily, aMember WHERE
    aFamily ∈ Family ∧ Families_Family_name(aFamily) = aFamilyName
    ∧ families(aFamily) = fmlRoot
    ∧ aMember ∈ FAMILYMEMBER
    ∧ aMember ∉ FamilyMember
  THEN
    FamilyMember := FamilyMember ∪ {aMember} ||
    theSons := theSons ∪ {aMember ↦ aFamily} ||
    Families_FamilyMember_name := Families_FamilyMember_name ⇐ {aMember ↦ aSonName}
  END
END;

```

**Figure 8:** Example of a basic model mutation operation defined in B

```

fam ← createFlandersFamilySonRod =
BEGIN
  fam ← FamilyNEW("Flanders");
  SonNEW("Flanders", "Rod")
END;

```

**Figure 9:** Example of a call to SonNEW.

### 3.5.2. Test suites in CSP

Each *T\_objective.mch* machine is associated with a CSP specification through the definition *CSP\_GUIDE\_FILE* in the header of the machine. This means that the execution of B operations is guided by a CSP controller. The role of CSP is to orchestrate three aspects: (i) the application of model mutations, (ii) the triggering of transformation rules through the *propagate* operation, and (iii) the verification of pre- and post-conditions. Instead of leaving rule application fully non-deterministic, CSP constrains the execution by defining admissible sequences of events.

A CSP test suite is structured as a set of scenarios, each corresponding to a TTC test case. A scenario consists of a sequence of mutation events followed by propagation phases and condition checks. The following excerpt illustrates this structure.

```

MAIN =
  testIncrementalInserts.true
  -> createInitialFamilies
  -> PROPAGATE_UNTIL_SYNC ;
  setBirthdayOfRod
  -> setBirthdayOfFatherBart
  -> createNewFamilySimpsonWithMembers -> unsync
  -> PROPAGATE_UNTIL_SYNC ;
  changeAllBirthdays
  -> createSonBart -> unsync
  -> PROPAGATE_UNTIL_STOP

```

**Figure 10:** Excerpt of a CSP test scenario (file *IncrementalForward.csp*)

In this example, the execution starts by selecting a specific test case (here *testIncrementalInserts*). A mutation is then applied to the source model (e.g., *createInitialFamilies*). A propagation phase is then entered, during which transforma-

tion rules are repeatedly applied via the process `PROPAGATE_UNTIL_SYNC`, defined in Figure 11. This process triggers rule applications through events of the form `propagate!r`, where `r` denotes a transformation rule. The execution continues until a synchronization event (`sync`) is reached. After this phase, additional mutations are applied, followed by another propagation phase. Finally, the execution ends with a last propagation phase, `PROPAGATE_UNTIL_STOP`, which is similar to `PROPAGATE_UNTIL_SYNC` but does not allow synchronization. As a result, rule applications continue until no further rule is enabled.

```
PROPAGATE_UNTIL_SYNC =
  [] r : RULE @ propagate!r -> PROPAGATE_UNTIL_SYNC
  [] sync -> SKIP
```

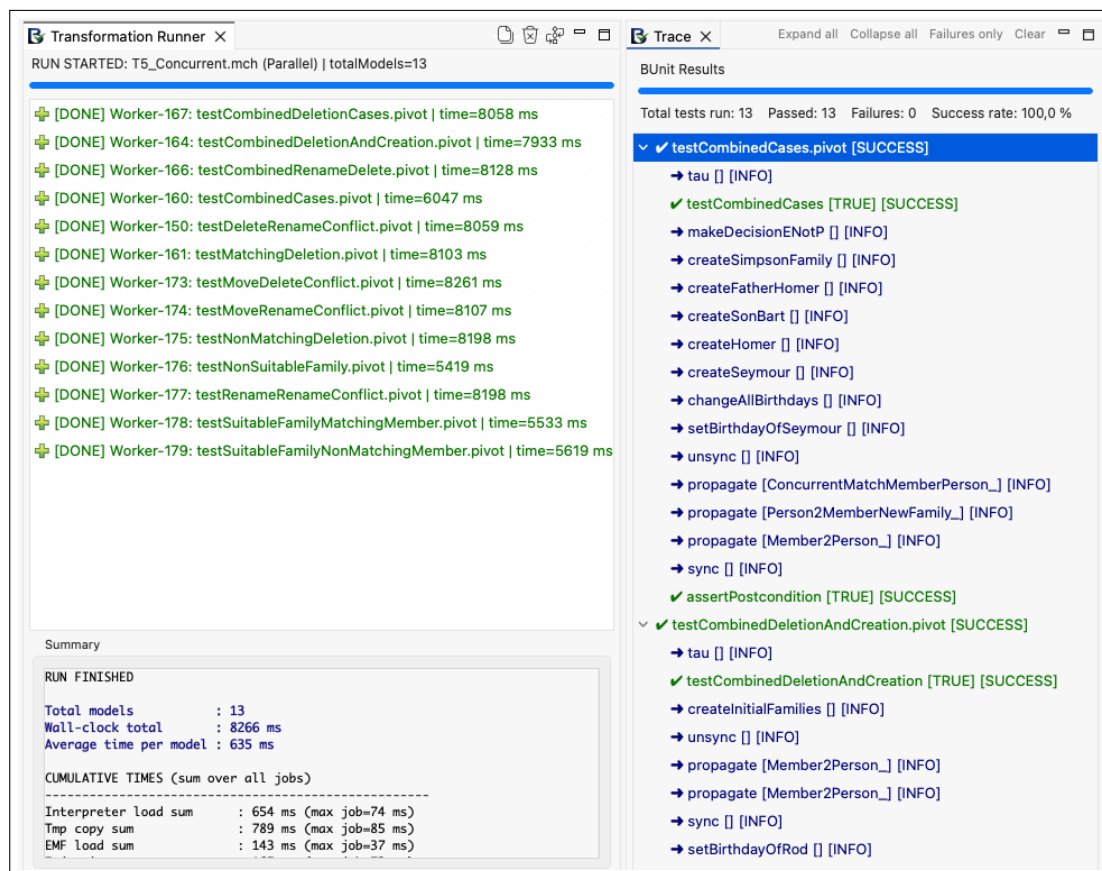
**Figure 11:** CSP control of rule application

Pre- and post-conditions are encoded as explicit checks using events such as `assertPrecondition` and `assertPostcondition`. Depending on their evaluation, the execution either continues or stops, ensuring that each scenario satisfies the expected behavior defined by the benchmark.

## 4. Results

### 4.1. Executability and functional validation

BCerT provides full support for the execution of CSP||B transformation scenarios. It ensures that model mutations, rule propagation, and condition checking are executed and reproducible. Figure 12 illustrates the execution of CSP||B model within BCerT of the *Concurrent* test suite.



**Figure 12:** Execution of a CSP||B test suite in BCerT

The left-hand side shows the execution summary, where each test scenario is run as an independent job. The Concurrent test suite contains 13 scenarios, all of which are successfully executed in parallel. The total execution time is 8266 ms, corresponding to an average execution time of 635 ms per scenario. These measurements are reported to characterize the execution of the TTC validation scenarios. Since the benchmark scenarios involve models of moderate size and are primarily intended for functional validation, these results should not be interpreted as scalability results. The right-hand side provides a detailed execution trace for each scenario. It shows the sequence of events triggered during execution, including model mutations, synchronization steps (`sync/unsync`), rule applications via `propagate`, and the evaluation of pre- and post-conditions. For instance, one can observe successive rule applications such as `ConcurrentMatchMemberPerson`, `Person2MemberNewFamily`, and `Member2Person`, followed by a synchronization point and a successful post-condition check. This execution trace demonstrates how CSP orchestrates the application of B-based transformation rules while ensuring that each scenario satisfies its expected properties. It also highlights the ability of BCerT to execute multiple scenarios concurrently and to provide detailed feedback on the correctness of each execution.

Table 1 summarizes the size of the main B and CSP specifications involved in our solution. The overall specification consists of 3275 lines of B code and 1006 lines of CSP code. The `MetaModel.mch` machine represents the core of the solution, as it defines the data structures, invariants, and transformation rules. Objective-specific machines (`T_objective.mch`) mainly define propagation strategies and remain relatively compact in comparison.

<b>B Machines</b>	<b>LOC</b>	<b>CSP Specifications</b>	<b>LOC</b>	<b>Test scenarios</b>
BatchBackward.mch	205	BatchForward.csp	81	7
MetaModel.mch	1278	BatchBackward.csp	124	11
RESET.mch	6	IncrementalForward.csp	170	8
T0_MainTransformation.mch	24	IncrementalBackward.csp	250	8
T1_BatchForward.mch	109	Concurrent.csp	265	13
T2_BatchBackward.mch	209	RoundTrip.csp	116	3
T3_IncrementalForward.mch	346	<b>Total</b>	<b>1006</b>	<b>50</b>
T4_IncrementalBackward.mch	367			
T5_Concurrent.mch	502			
T6_RoundTrip.mch	229			
<b>Total</b>	<b>3275</b>			

**Table 1**

Size of the B machines and CSP specifications, with number of test scenarios

The size reported in Table 1 should be interpreted with care. The core transformation logic is concentrated in the transformation rules of `MetaModel.mch`, namely the forward, backward, and concurrent rules described in Section 3. This corresponds to 17 B operations (about 460 LOC). The remaining parts of the B and CSP specifications mainly encode utility operations and model mutations together with TTC-specific test scenarios. In particular, the CSP files are not part of the transformation logic itself; they play a role comparable to the Java test code of the reference benchmark by orchestrating and validating the TTC scenarios.

## 4.2. Scalability experiment

To provide preliminary scalability insights, we evaluated the incremental forward transformation on generated Families models. Each generated model contains  $N$  families, where each family has five members: a father, a mother, one son, and two daughters. The corresponding Persons model is initially empty, and the pivot model is configured in forward mode.

These results show that BCerT can execute transformations on models significantly larger than those used in the TTC functional test suites. However, they also confirm that scalability remains a limitation of the current implementation. The execution time increases substantially with the number

Families	Members	Execution time
3	15	1.59 s
5	25	1.85 s
10	50	2.60 s
50	250	22.55 s
100	500	1.28 min
300	1500	10.65 min
500	2500	26.56 min
1000	5000	48.49 min
5000	25000	1.90 h
10000	50000	3.71 h

**Table 2**

Preliminary scalability results for incremental forward transformation

of model elements. This is expected, since rule application relies on ProB’s animation and constraint-solving mechanisms. These measurements should therefore be interpreted as preliminary scalability observations rather than as optimized performance results.

## 5. Conclusion

This paper presented our BCerT solution to the TTC’2026 F2P case study. The solution is specified using B and CSP, and is fully executable within the BCerT environment. The core transformation logic is expressed as B operations, while invariants capture the structural and semantic consistency properties that must be preserved. Since the transformation rules are proved correct with respect to these invariants, the approach provides correctness by construction solution at the level of the formal model.

Beyond verification, the TTC benchmark requires validation against observable scenarios. To address this aspect, we used CSP as an orchestration layer for B operations. CSP controls model mutations, rule propagation, synchronization phases, and pre-/post-condition checking. This makes it possible to execute the TTC test suites in a reproducible way, while keeping the rule-based execution mechanism of BCerT and ProB. The solution covers batch forward, batch backward, incremental forward, incremental backward, roundtrip, and concurrent synchronization scenarios.

The current version of the solution also provides preliminary scalability results for generated Families models. The largest completed experiment involves 10,000 families, corresponding to 50,000 family members, and requires 3.71 hours for the incremental forward transformation. These results show that the approach can handle models larger than the functional scenarios, but they also confirm that scalability remains a limitation of the current implementation. Since BCerT relies on ProB’s constraint-solving and animation mechanisms, further work is needed to optimize rule selection and reduce the cost of enabled-operation computation.

## Declaration on Generative AI

The author used generative AI in order to support latex formatting and grammar correction, during the preparation of this paper. The author reviewed and edited the content as needed and takes full responsibility for the publication’s content.

## References

- [1] J.-R. Abrial, *The B-book: Assigning Programs to Meanings*, Cambridge University Press, New York, NY, USA, 1996.

- [2] A. Idani, Transpilation Meets the B Method, in: E. Börger, Y. A. Ameur (Eds.), Festschrift in Honor of Jean-Raymond Abrial, Springer, 2026. Accepted book chapter.
- [3] M. Leuschel, M. Butler, ProB: an automated analysis toolset for the B method, *International Journal on Software Tools for Technology Transfer* 10 (2008) 185–203. URL: <https://doi.org/10.1007/s10009-007-0063-9>. doi:10.1007/s10009-007-0063-9.
- [4] A. Idani, G. Vega, Engineering verified model transformations through a proof-based language workbench, in: 19th ACM SIGPLAN International Conference on Software Language Engineering (SLE'26), ACM, 2026. doi:10.1145/3806383.3815523, accepted long paper.
- [5] C. A. R. Hoare, *Communicating Sequential Processes*, Prentice Hall, Englewood Cliffs, NJ, 1985.
- [6] A. Idani, Meeduse: A tool to build and run proved DSLs, in: 16th International Conference on Integrated Formal Methods (IFM), volume 12546 of *LNCS*, Springer, 2020, pp. 349–367. doi:10.1007/978-3-030-63461-2\_19.
- [7] M. Butler, M. Leuschel, Combining CSP and B for specification and property verification, in: *International Conference on Formal Methods*, volume 3582, Springer, 2005, p. 221–236. doi:10.1007/11526841\_16.