



Improving the "Correct Eventual Consistency" Tool

Sreeja Nair, Marc Shapiro

► **To cite this version:**

Sreeja Nair, Marc Shapiro. Improving the "Correct Eventual Consistency" Tool. [Research Report] RR-9191, Sorbonne Université. 2018. hal-01832888

HAL Id: hal-01832888

<https://hal.inria.fr/hal-01832888>

Submitted on 9 Jul 2018

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.



Improving the “Correct Eventual Consistency” Tool

Sreeja Nair, Marc Shapiro

**RESEARCH
REPORT**

N° 9191

July 2018

Project-Teams DELYS



Improving the “Correct Eventual Consistency” Tool

Sreeja Nair, Marc Shapiro

Project-Teams DELYS

Research Report n° 9191 — July 2018 — 8 pages

Abstract: Preserving invariants while designing distributed applications under weak consistency models is difficult. The CEC (*Correct Eventual Consistency Tool*) is meant to aid the application designer in this task. It provides information about the errors during concurrent operations and suggestions on how and where to synchronize operations. This report presents two features of the tool: providing a counterexample for debugging and concurrency control suggestions.

Key-words: Consistency, Verification, Distributed applications

**RESEARCH CENTRE
PARIS**

2 rue Simone Iff - CS 42112
75589 Paris Cedex 12

Résumé : Dans les applications réparties s'exécutant dans un modèle de cohérence faible, le maintien des invariants est un problème difficile. L'outil CEC (*Correct Eventual Consistency*) est destiné à aider le développeur d'application dans cette tâche. Il fournit des informations sur les erreurs lors des opérations simultanées et des suggestions sur comment et où synchroniser les opérations. Dans ce rapport, nous présentons deux fonctionnalités de l'outil : la première fournit un contre-exemple lors du débogage et la seconde des suggestions permettant un meilleur contrôle de concurrence.

Mots-clés : cohérence, vérification, applications distribuées

1 Introduction

Reasoning about the correctness of distributed applications while ensuring high availability is a non-trivial task. CISE, introduced by Gotsman et al. [1], is a logic for reasoning about the correctness of a distributed application operating on top of a causally-consistent database. In cases where the developer needs to control concurrent operations, he/she can use concurrency control named tokens. According to the CISE logic, a distributed application is guaranteed to uphold its invariant if:

1. Each operation is sequentially correct.
2. The precondition of each operation is stable under any concurrent operation.
3. Concurrent operations commute.

Based on the CISE logic, a first tool was developed by Najafzadeh et al. [2]. This initial CISE tool is difficult because the user needs to use the low-level Z3 APIs directly [3]. Subsequently, Marcelino, Balesgas, and Ferreira [4] designed a second generation CISE tool, called Correct Eventual Consistency (CEC), which works on the same principle and provides a high-level developer friendly verification language, an extension of Boogie [5] with specific annotations.

An application written as a CEC specification consists of the following parts:

- Data structures and properties
- Variables
- Invariants
- Operations with pre and post conditions

The tool first checks the specification for errors in the sequential specification. The first step checks whether:

- Each individual component of the specification is syntactically correct (*syntax check*).
- Each operation satisfies the invariant individually (*safety check*).
- The specification contains any contradictory clauses (*anomaly check*).
- Every variable that is modified by an operation has a properly defined value when the operation terminates (*completeness check*).

If the first phase passes, the analysis proceeds onto the second stage to check for concurrency. This stage verifies the CISE consistency conditions. It checks every pair of operations to see whether:

- The two operations of the pair have opposing preconditions including the tokens they acquire (*opposition check*). If so, they will never run concurrently. Otherwise (the two operations can run concurrently), the next two checks are applied.
- The precondition of one operation is preserved even under concurrent execution of the other operation with tokens (*stability check*). According to the CISE logic, this check ensures that the application’s invariant is preserved.
- The two operations commute (*commutativity check*). This ensures convergence.

A first report on the evaluation of the tool is available [6]. The current paper documents the improvements since the previous report in the following areas:

- Counter example for failed verification.
- Optimized token generation.

2 Counterexample for failed verification

The previous report[6] highlighted the lack of a counterexample when a verification fails. The tool informed the user about the type of check that failed, it was difficult for the user to understand why. The improved version of the tool now provides a counterexample with the following information:

- The statement that failed.
- The values of the parameters of the operation.
- The values of all the variables in the failed expression.

Let us illustrate with an example of a bank account. It has two operations - deposit and withdraw. Suppose that the developer has specified a precondition for withdraw operation, but not for deposit.

```

@init
type Client;

@variables
var balances : [Client]int;

@equals [Client]int @as forall c: Client :: @this[c] == @other[c];

@invariant
forall c: Client :: balances[c] >= 0;

@operations

procedure deposit(accountId: Client, amount: int)
  modifies balances;
  ensures
    forall c: Client ::
      ((c != accountId) ==> (balances[c] == old(balances)[c])) &&
      ((c == accountId) ==> (balances[accountId] == old(balances)[accountId]
        + amount));

procedure withdraw(accountId: Client, amount: int)
  modifies balances;
  requires balances[accountId] - amount >= 0;
  ensures
    forall c: Client ::
      ((c != accountId) ==> (balances[c] == old(balances)[c])) &&
      ((c == accountId) ==> (balances[accountId] == old(balances)[accountId]
        - amount));

```

The `@init` includes the initialization section of the specification. This sections contains datatype declarations, axioms and function declarations. In our case, there is a single datatype declaration here. The next section, `@variable` contains all the variables used in the specification. The variable we use here is a mapping from each `Client` to their balance (which is an integer).

The section `@equals` gives the guidance to the tool about how to compare two variables for equality. In this specification, two variables of type `[Client]int` as considered equal if each `Client` has the same balance. The application invariant here is that the balances should always be non-negative. `@invariant` states this application invariant and `@operations` contains a list of operations of the application. Each operation will have a precondition which should be true at the origin replica (denoted by `requires` clause) and the effector function which is the result of executing the operation (denoted by the `ensures` clause).

The tool performs an analysis on this specification. In this case the safety test fails during the first stage (verifying the sequential specification). The tool returns information on the assertion that failed, and the values of the variables involved in the assertion. It also provides the values of the parameters of the failing operation. Below is the corresponding output of the failure.

```
##### SPECIFICATION CORRECTION TESTS #####
-----BASE VERIFICATION-----
IMPORTS TEST: (PASSED)
INITIALIZER TEST: (PASSED)
VARIABLES TEST: (PASSED)
INVARIANTS TEST: (PASSED)
deposit OPERATION TEST: (PASSED)
withdraw OPERATION TEST: (PASSED)
FULL TEST: (PASSED)

-----SEQUENTIAL VERIFICATION-----
deposit SAFETY TEST: (FAILED)
DEBUG INFO :
The statements which failed are :::
ensures (forall c: Client :: balances[c] >= 0);
assume(forall c: Client :: balances[c] >= 0);
The value of accountId is T@Client!val!0.
The value of balances is |T@[ Client ]Int!val!1|.
The value of amount is (- 1200).
The value of c is .
The value of Client is .

withdraw SAFETY TEST: (PASSED)
deposit ABSURD TEST: (PASSED)
withdraw ABSURD TEST: (PASSED)
deposit COMPLETENESS TEST: (PASSED)
withdraw COMPLETENESS TEST: (PASSED)
FAILED LOGIC SPECIFICATION CORRECTION TEST
EXECUTION STOPPED
```

The first section of the result, the base verification, performs a syntax check for each part of the specification. The second part, the sequential verification checks whether the specification is safe sequentially.

The result tells that the tool was unable to verify that deposit operation ensures a positive balance, the application invariant. The parameters of the failing deposit operation are `amount` and `accountId`. The value of `amount` is -1200, a negative value. The developer takes this as a hint that negative values for `amount` are problematic. To test this hypothesis, he adds a precondition to indicate that `amount` should be positive. Then the specification looks like the following:

```
@init
type Client;

@variables
var balances : [Client]int;
```



```

@equals [Client]int @as forall c: Client :: @this[c] == @other[c];

@invariant
forall c: Client :: balances[c] >= 0;

@operations

procedure deposit(accountId: Client, amount: int)
  modifies balances;
  requires amount >= 0;
  ensures
    forall c: Client ::
      ((c != accountId) ==> (balances[c] == old(balances)[c])) &&
      ((c == accountId) ==> (balances[accountId] == old(balances)[accountId]
        + amount));

procedure withdraw(accountId: Client, amount: int)
  modifies balances;
  requires balances[accountId] - amount >= 0;
  ensures
    forall c: Client ::
      ((c != accountId) ==> (balances[c] == old(balances)[c])) &&
      ((c == accountId) ==> (balances[accountId] == old(balances)[accountId]
        - amount));

```

Now we rerun the tool with the corrected specification. We can see that the sequential safety test passes and we get the following result.

```

##### SPECIFICATION CORRECTION TESTS #####

-----BASE VERIFICATION-----
IMPORTS TEST: (PASSED)
INITIALIZER TEST: (PASSED)
VARIABLES TEST: (PASSED)
INVARIANTS TEST: (PASSED)
deposit OPERATION TEST: (PASSED)
withdraw OPERATION TEST: (PASSED)
FULL TEST: (PASSED)

-----SEQUENTIAL VERIFICATION-----
deposit SAFETY TEST: (PASSED)
withdraw SAFETY TEST: (PASSED)
deposit ABSURD TEST: (PASSED)
withdraw ABSURD TEST: (PASSED)
deposit COMPLETENESS TEST: (PASSED)
withdraw COMPLETENESS TEST: (PASSED)
##### CONSISTENCY TESTS #####

-----PAIR OPPOSITION VERIFICATION-----
deposit deposit OPPOSITION TEST: (PASSED)
deposit withdraw OPPOSITION TEST: (PASSED)
withdraw withdraw OPPOSITION TEST: (PASSED)

-----PAIR STABILITY VERIFICATION-----
deposit deposit STABILITY TEST: (PASSED)
deposit withdraw STABILITY TEST: (PASSED)
withdraw withdraw STABILITY TEST: (FAILED)
DEBUG INFO :
The statements which failed are ::
requires (balances[accountId] - amount >= 0);
assume(forall c: Client :: balances[c] >= 0);
The value of accountId is T@Client!val!0,T@Client!val!0.

```

```

The value of balances is |T@[ Client | Int ! val ! 0 |.
The value of amount is 797,797.
The value of c is .
The value of Client is .

```

```

—————PAIR COMMUTATIVITY VERIFICATION—————
deposit deposit COMMUTATIVITY TEST: (PASSED)
deposit withdraw COMMUTATIVITY TEST: (PASSED)

```

The stability test failed for concurrent `withdraw` operations. Indeed, the precondition of the second `withdraw` is not stable under the execution of the first. We can see that the values of `accountId` are the same for both operations. This is a hint for the developer to insert concurrency control for concurrent `withdraw` operations for the same `accountId`.

3 Synthesis of concurrency control

The tool now leverages the counterexample obtained in Section 2 to suggest a concurrency control token.

Let us continue with our bank account example. From the counter example, the tool infers two possible restrictions:

- The `amounts` should be different
and/or
- The `accountIds` should be different

The tool automatically reruns the analysis with these added restrictions. Even when the two `amounts` are different, the analysis still violates the stability check. Therefore this is not a fruitful exercise. When the `accountIds` are constrained to be different the specification passes the stability test. This shows that concurrent withdraws to different accounts are ok, and suggests that concurrency should be disallowed for the same `accountId`. The CISE abstract for concurrency control is called a token.

The output of the token generator is :

```

—————COUNTER EXAMPLE GUIDED SOLVER—————
withdraw(accountId: Client, amount: int) withdraw(accountId: Client, amount:
  int) COUNTER EXAMPLE GUIDED SOLVER: (PASSED)
Solutions:
( @1 != @3 )

—————TOKEN MODEL—————
Tokens per operations:
withdraw: withdraw_accountId

Token conflicts:
withdraw_accountId: withdraw_accountId

```

The first part of the output tells that in the list of parameters, the first parameter should not be equal to the third one. The token model presents the same information grouped in the form of tokens needed for each operation and conflicts for each token. This complies with the token specifications in CISE. The token model tells that `withdraw` operation needs to acquire a token for each `account_id`. For a developer, this means the `withdraw` operations operating on one `account_id` need to synchronize.

4 Conclusion

This report summarises some recent improvements made on the Correct Eventual Consistency (CEC) tool. The tool now provides more comprehensible counter examples. The tool also suggests tokens by utilising the information from the counter examples in an optimized fashion.

Currently the tool is providing suggestions only based on imposing inequality restrictions on the parameters. This can be improved considering more relations between parameters. The next step would be to develop a complementary tool for analysing the applications which use state-based update propagation mechanism.

Acknowledgement

The authors would like to thank Carla Ferreira, Gustavo Petri, Gonçalo Marcelino and Yann Thierry-Mieg for their inputs. This research is supported in part by the RainbowFS project (*Agence Nationale de la Recherche*, France, number ANR-16-CE25-0013-01) and by European H2020 project 732 505 LightKone (2017–2020).

References

- [1] Alexey Gotsman, Hongseok Yang, Carla Ferreira, Mahsa Najafzadeh, and Marc Shapiro. “Cause I’m Strong Enough: Reasoning About Consistency Choices in Distributed Systems”. In: *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. POPL ’16. St. Petersburg, FL, USA: ACM, 2016, pp. 371–384. ISBN: 978-1-4503-3549-2. DOI: 10.1145/2837614.2837625. URL: <http://doi.acm.org/10.1145/2837614.2837625>.
- [2] Mahsa Najafzadeh, Alexey Gotsman, Hongseok Yang, Carla Ferreira, and Marc Shapiro. “The CISE Tool: Proving Weakly-consistent Applications Correct”. In: *Proceedings of the 2Nd Workshop on the Principles and Practice of Consistency for Distributed Data*. PaPoC ’16. London, United Kingdom: ACM, 2016, 2:1–2:3.
- [3] Mahsa Najafzadeh. “The Analysis and Co-design of Weakly-Consistent Applications”. Theses. Université Pierre et Marie Curie, Apr. 2016.
- [4] Gonçalo Marcelino, Valter Balegas, and Carla Ferreira. “Bringing Hybrid Consistency Closer to Programmers”. In: *Proceedings of the 3rd International Workshop on Principles and Practice of Consistency for Distributed Data*. PaPoC’17. Belgrade, Serbia: ACM, 2017, 6:1–6:4.
- [5] *Boogie*. 2017. URL: <https://github.com/boogie-org/boogie>.
- [6] Sreeja S Nair. *Evaluation of the CEC (Correct Eventual Consistency) Tool*. Research Report RR-9111. Inria Paris ; LIP6 UMR 7606, UPMC Sorbonne Universités, France, Nov. 2017, pp. 1–27. URL: <https://hal.inria.fr/hal-01628719>.



**RESEARCH CENTRE
PARIS**

2 rue Simone Iff - CS 42112
75589 Paris Cedex 12

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399