



# The CISE Tool: Proving Weakly-Consistent Applications Correct

Mahsa Najafzadeh, Alexey Gotsman, Hongseok Yang, Carla Ferreira, Marc Shapiro

## ► To cite this version:

Mahsa Najafzadeh, Alexey Gotsman, Hongseok Yang, Carla Ferreira, Marc Shapiro. The CISE Tool: Proving Weakly-Consistent Applications Correct. PaPoC 2016 - 2nd Workshop on the Principles and Practice of Consistency for Distributed Data , Apr 2016, Londres, United Kingdom. ACM, <10.1145/2911151.2911160>. <hal-01350636>

**HAL Id: hal-01350636**

**<https://hal.inria.fr/hal-01350636>**

Submitted on 1 Aug 2016

**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.

# The CISE Tool: Proving Weakly-Consistent Applications Correct

Mahsa Najafzadeh  
Sorbonne-Universités-UPMC  
& Inria  
France  
mahsa.najafzadeh@lip6.fr

Alexey Gotsman  
MDEA Software Institute  
Spain  
alexey.gotsman@imdea.org

Hongseok Yang  
University of Oxford  
UK  
hongseok00@gmail.com

Carla Ferreira  
NOVA LINCS, DI, FCT, U.  
NOVA de Lisboa  
Portugal  
carla.ferreira@fct.unl.pt

Marc Shapiro  
Sorbonne-Universités-UPMC  
& Inria  
France  
marc.shapiro@acm.org

## ABSTRACT

Designers of a replicated database face a vexing choice between strong consistency, which ensures certain application invariants but is slow and fragile, and asynchronous replication, which is highly available and responsive, but exposes the programmer to unfamiliar behaviours. To bypass this conundrum, recent research has studied hybrid consistency models, in which updates are asynchronous by default, but synchronisation is available upon request. To help programmers exploit hybrid consistency, we propose the first static analysis tool for proving integrity invariants of applications using databases with hybrid consistency models. This allows a programmer to find minimal consistency guarantees sufficient for application correctness.

## Keywords

Program Analysis; Consistency; Invariant

## 1. INTRODUCTION

To achieve availability and scalability, many modern distributed systems rely on *replicated databases*, which maintain multiple replicas of shared data. Clients can access the data at any of the replicas, and these replicas communicate changes to each other using message passing. For example, large-scale Internet services use data replicas in geographically distinct locations, and applications for mobile devices keep replicas locally to support offline use. Ideally, we would like replicated databases to provide *strong consistency*, i.e., to behave as if a single centralised node handles all operations. However, achieving this ideal usually requires synchronisation among replicas, which slows down the database and even makes it unavailable if network connections between replicas fail [1, 7].

For this reason, modern replicated databases often eschew synchronisation completely; such databases are commonly dubbed

*eventually consistent* [14]. In these databases, a replica performs an operation requested by a client locally without any synchronisation with other replicas and immediately returns to the client; the *effect* of the operation is propagated to the other replicas only *eventually*. Unfortunately, this way of processing operations exposes applications to undesirable concurrency behaviours, which may cause bugs such as state divergence or invariant violation [6].

For instance, consider a bank account replicated at different bank branches, which supports operations deposit and withdraw. A programmer would like to ensure an *integrity invariant* that the balance is never negative. Assume that the balance is initially €100. Eventual consistency will allow two users to concurrently withdraw €60 at different branches and thus violate the integrity invariant. To ensure the invariant in this example, we have to introduce synchronisation between replicas, and, since synchronisation is expensive, we would like to introduce it sparingly. To allow this, some research [3, 9, 12, 13] and commercial [2, 4, 10] databases now provide *hybrid* consistency models that allow the programmer to request stronger consistency for certain operations and thereby introduce synchronisation. For example, to preserve the integrity invariant in our banking application, only withdraw operations need to use strong consistency, and hence, synchronise to ensure that the account is not overdrawn; deposit operations may use eventual consistency and hence proceed without synchronisation.

Unfortunately, using hybrid consistency models effectively is far from trivial. Requesting stronger consistency in too many places may hurt performance and availability, and requesting it in too few places may violate correctness. Striking the right balance requires the programmer to reason about the application behaviour on the subtle semantics of the consistency model, taking into account which anomalies are disallowed by a particular consistency strengthening and whether disallowing these anomalies is enough to ensure correctness.

To help programmers exploit hybrid consistency models, we propose the first static analysis tool (called CISE: 'Cause I'm Strong Enough) for proving integrity invariants of applications using replicated databases with a range of hybrid models. Our tool is based on a novel proof rule, which we have proved sound [8]. The tool automates the proof rule by discharging its obligations using an SMT solver. If an obligation fails, the tool provides a counter-example, which the developer can use to understand the source of the problem. Using the tool, we have verified several example applications that require strengthening consistency in nontrivial ways [8]. These

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PaPOC'16, April 18-21, 2016, London, United Kingdom

© 2016 ACM. ISBN 978-1-4503-4296-4/16/04...\$15.00

DOI: <http://dx.doi.org/10.1145/2911151.2911160>

include an extension of the above banking application, an online auction service and a course registration system. A demo of the tool is available online [11].

In the rest of this paper, we explain our static analysis by the example of the above banking application.

## 2. SYSTEM MODEL

An application consists of a set of operations  $\text{Op}$  over some set of objects, and invariants over the objects. The database system consists of a set of replicas, each maintaining a full copy of the database state  $\text{State}$ .

The replication model uses a Read-One-Write-All (ROWA) approach [5]. A client operation is initially executed at a single replica, which we refer to as its *origin* replica. This updates the replica state deterministically, and immediately returns a value to the client. After this, the replica sends a message to all other replicas containing the *effector* of the operation, which describes the updates done by the operation to the database state. Upon receipt, the replicas apply the effector to their state. Effectors of causally-dependent operations are executed in the same order at every replica; effectors of independent (concurrent) operations are executed in any order.

More precisely, the semantics of operations is defined by a partial function

$$\mathcal{F} \in \text{Op} \rightarrow (\text{State} \rightarrow (\text{Val} \times (\text{State} \rightarrow \text{State}) \times \mathcal{P}(\text{Token}))).$$

Given a state  $\sigma \in \text{State}$  in which an operation  $o \in \text{Op}$  executes at its origin replica,  $\mathcal{F}(o)(\sigma)$  determines:

- The return value of the operation, from a set  $\text{Val}$ . We use a special value  $\perp$  for operations that return no value.
- A function defining the *effector* of the operation. This will be applied by every replica to its state: immediately at the origin replica, and after receiving the corresponding message at all other replicas.
- A set of *tokens*, used to introduce synchronisation. We explain them later.

For instance, consider the naïve banking application in Figure 1. A client can read the balance from the local replica, make deposits to and withdrawals from the account, and compute interest, all without communicating with the other replicas. Each operation is associated with a *precondition*—a predicate over the state of its origin replica and parameters that determines when the operation can be safely executed (and the  $\mathcal{F}$  function defined). A minimal precondition of the  $\text{deposit}(\text{amount})$  and  $\text{withdraw}(\text{amount})$  operations is  $\text{amount} \geq 0$ . Their effectors add amount to (respectively, subtract it from) the balance. The interest operation’s precondition is  $\text{true}$  and its effector multiplies the balance by the interest rate. (As we will see later, the analysis shows that the precondition of  $\text{withdraw}$  needs to be strengthened, and that this effector of interest is unsafe.)

## 3. CISE ANALYSIS

### 3.1 Safety Analysis

The first CISE proof obligation, called the effector safety analysis, verifies that the effector of every operation maintains the invariant when applied to *any* state where the operation’s precondition is true (not necessarily the one in which the operation was generated).

Let’s try out the effector safety analysis on the simple banking application of Figure 1. According to the analysis, the effectors of  $\text{deposit}$  and  $\text{interest}$  always maintain the invariant. However,

$$\begin{aligned} \sigma_{\text{init}} &= 0 \\ I &= (\text{balance} \geq 0) \\ \text{Token} &= \emptyset \\ \mathcal{F}_{\text{deposit}(\text{amount})}(\text{balance}) &= (\perp, (\lambda \text{balance}'. \text{balance} + \text{amount}), \emptyset) \\ \mathcal{F}_{\text{interest}()}(\text{balance}) &= (\perp, (\lambda \text{balance}'. (1.05 * \text{balance}')), \emptyset) \\ \mathcal{F}_{\text{withdraw}(\text{amount})}(\text{balance}) &= (\perp, (\lambda \text{balance}'. \text{balance} - \text{amount}), \emptyset) \end{aligned}$$

Precondition	Operation
$\text{amount} \geq 0$	$\text{deposit}(\text{amount})$
$\text{true}$	$\text{interest}()$
$\text{amount} \geq 0$	$\text{withdraw}(\text{amount})$

Figure 1: A simple banking application (incorrect).

for  $\text{withdraw}$ , the obligation fails and our tool produces a counter-example: if the balance is zero, a non-zero  $\text{withdraw}$  operation makes the balance negative. Therefore, we must fix the issue by strengthening its precondition, so that the amount debited is less or equal than the current balance.

With this correction, the effector safety analysis succeeds. The corrected preconditions are shown at the bottom of Figure 2.

### 3.2 Commutativity Analysis

Effectors of concurrent operations may execute in different orders at different replicas. The second CISE obligation, called the commutativity analysis, checks if all pairs of effectors of such operations *commute*: executing them in any order yields the same result, whatever the starting state.

Let us check this obligation for the specification in Figure 1. Predictably, applying the commutativity analysis proves that  $\text{deposit}$  and  $\text{withdraw}$  effectors commute. However, the effector of  $\text{interest}$  does not commute with that of the other operations, and the tool returns a counter-example. Consider two replicas 1 and 2. The balance is initially €100. Replica 1 is the origin for an interest operation. Replica 2 is the origin for a  $\text{deposit}(20)$  operation. Replica 1 first applies the effector of interest and then that of  $\text{deposit}$ , whereas replica 2 applies them in the opposite order. Depending on the order of execution, the result is different, and the replicas diverge.

We fix this by changing the interest operation to compute the absolute interest at the origin replica and letting its effector add this amount to the local balance of every replica (Figure 2). With this corrected specification, our tool proves that the effector of interest does commute with those of the other operations.

### 3.3 Stability Analysis

The effector safety analysis verified that that the effector of each operation  $o$  maintains the invariant when executed in a state satisfying the precondition of the operation. The precondition holds at  $o$ ’s origin replica, but how do we know that it will hold when  $o$ ’s effector is applied at a different replica, which concurrently executes effectors of other operations? The third obligation of CISE analysis, called stability analysis, checks if executing the effector of any other operation  $o'$  maintains the precondition of  $o$ .

Let us illustrate the stability analysis of the  $\text{withdraw}$  operation in Figure 1. The precondition of the  $\text{withdraw}$  operation is stable under the effectors of  $\text{deposit}$  and  $\text{interest}$ , but it is not stable under the effector of  $\text{withdraw}$ . The tool returns the following counter-example. Let the balance be €2. The precondition to  $\text{withdraw}(1)$  is verified. However, a concurrent  $\text{withdraw}(2)$  (whose precondition

$$\begin{aligned}
\sigma_{\text{init}} &= 0 \\
I &= \text{balance} \geq 0 \\
\text{Token} &= \{\tau\} \\
\bowtie &= \{(\tau, \tau)\} \\
\mathcal{F}_{\text{deposit}(\text{amount})}(\text{balance}) &= (\perp, (\lambda \text{balance}'. \text{balance} + \\
&\quad \text{amount}), \emptyset) \\
\mathcal{F}_{\text{interest}()}(\text{balance}) &= (\perp, (\lambda \text{balance}'. (\text{balance}' + 0.05 \\
&\quad * \text{balance})), \emptyset) \\
\mathcal{F}_{\text{withdraw}(\text{amount})}(\text{balance}) &= (\perp, (\lambda \text{balance}'. \text{balance} - \\
&\quad \text{amount}), \{\tau\})
\end{aligned}$$

Precondition	Operation
$\text{amount} \geq 0$	<code>deposit(amount)</code>
true	<code>interest()</code>
$\text{balance} \geq \text{amount} \geq 0$	<code>withdraw(amount)</code>

**Figure 2: A corrected banking application.**

is also OK) at a different replica makes the balance zero, now violating the precondition of `withdraw(1)`. If we were to continue, and to apply the effector of the first withdrawal operation at the second replica, the balance would become negative.

To fix the problem, the developer of the banking application may disallow the execution of withdrawals without synchronisation. To model such concurrency control, we use *tokens*  $\text{Token} = \{\tau, \dots\}$  and a symmetric *conflict relation*  $\bowtie \subseteq \text{Token} \times \text{Token}$  between pairs of them. In the banking application, we associate a token  $\tau$  to withdraw such that  $\tau \bowtie \tau$  (similarly to a mutual exclusion lock). This ensures that any two withdrawals synchronise.

Figure 2 presents a corrected banking application, incorporating all the changes outlined above. Our static analysis confirms that this application indeed maintains the integrity invariant.

## 4. CONCLUSIONS AND FUTURE WORK

We implemented the CISE Tool to automatically discharge the CISE analysis. The tool enables an application developer to check whether her application maintains crucial application invariants, over a replicated database with a given synchronisation protocol.

In the future, we plan to study proof rules for reasoning about integrity invariants on consistency models weaker than causal consistency. We also intend to automate the analysis of counter-examples in order to generate corrections semi-automatically.

## References

- [1] D. Abadi. Consistency tradeoffs in modern distributed database system design: CAP is only part of the story. *IEEE Computer*, 45(2), 2012.
- [2] Amazon. Supported operations in DynamoDB. <http://docs.aws.amazon.com/amazondynamodb/latest/developerguide/APISummary.html>, 2015.
- [3] V. Balesgas, N. Preguiça, R. Rodrigues, S. Duarte, C. Ferreira, M. Najafzadeh, and M. Shapiro. Putting consistency back into eventual consistency. In *Euro. Conf. on Comp. Sys. (EuroSys)*, pages 6:1–6:16, Bordeaux, France, Apr. 2015.
- [4] Basho Inc. Using strong consistency in Riak. <http://docs.basho.com/riak/latest/dev/advanced/strong-consistency/>, 2015.
- [5] P. Bernstein, V. Radzilacos, and V. Hadzilacos. *Concurrency Control and Recovery in Database Systems*. Addison Wesley Publishing Company, 1987.
- [6] J. C. Corbett, J. Dean, M. Epstein, A. Fikes, C. Frost, J. Furman, S. Ghemawat, A. Gubarev, C. Heiser, P. Hochschild, W. Hsieh, S. Kanthak, E. Kogan, H. Li, A. Lloyd, S. Melnik, D. Mwaure, D. Nagle, S. Quinlan, R. Rao, L. Rolig, Y. Saito, M. Szymaniak, C. Taylor, R. Wang, and D. Woodford. Spanner: Google’s globally-distributed database. In *Symp. on Op. Sys. Design and Implementation (OSDI)*, pages 251–264, Hollywood, CA, USA, Oct. 2012. Usenix.
- [7] S. Gilbert and N. Lynch. Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, 2002.
- [8] A. Gotsman, H. Yang, C. Ferreira, M. Najafzadeh, and M. Shapiro. ‘Cause I’m strong enough: Reasoning about consistency choices in distributed systems. In *Symp. on Principles of Prog. Lang. (POPL)*, pages 371–384, St. Petersburg, FL, USA, 2016.
- [9] C. Li, D. Porto, A. Clement, J. Gehrke, N. Preguiça, and R. Rodrigues. Making geo-replicated systems fast as possible, consistent when necessary. In *Symp. on Op. Sys. Design and Implementation (OSDI)*, pages 265–278, Hollywood, CA, USA, Oct. 2012.
- [10] Microsoft. Consistency levels in DocumentDB. <http://azure.microsoft.com/en-us/documentation/articles/documentdb-consistency-levels/>, 2015.
- [11] M. Najafzadeh and M. Shapiro. Demo of the CISE tool, Nov. 2015.
- [12] Y. Sovran, R. Power, M. K. Aguilera, and J. Li. Transactional storage for geo-replicated systems. In *Symp. on Op. Sys. Principles (SOSP)*, pages 385–400, Cascais, Portugal, Oct. 2011. Assoc. for Computing Machinery.
- [13] D. B. Terry, V. Prabhakaran, R. Kotla, M. Balakrishnan, M. K. Aguilera, and H. Abu-Libdeh. Consistency-based service level agreements for cloud storage. In *SOSP*, 2013.
- [14] W. Vogels. Eventually consistent. *CACM*, 52(1), 2009.