

Using Reachability Logic to Verify Distributed Systems

Stephen Skeirik, Andrei Stefanescu,
and José Meseguer

ACC Weekly Research Seminar
Nov. 2, 2016

Outline

1 Introduction

2 Reachability Logic

3 Current Progress

4 Future Directions

Introduction

ACC System Design

- Building complex cloud-based systems is challenging
 - many more *failure modes* than traditional servers
 - often *combinatorial explosion* in size of state space
- Testing *alone* is *inadequate* for verifying cloud-based systems
 - can only cover *finite* part of the state space
 - testing requires *essentially complete* artifacts, yet early design mistakes most costly
 - requirements *evolve* with system

Q: What methods can assist designing/verifying cloud systems?

Introduction

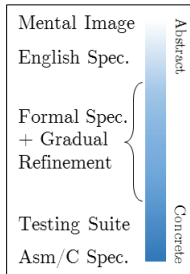
Formal Methods

Q: What methods can assist designing/verifying cloud systems?

A: We propose *formal specification/verification methods*

Goal: a *precise, readable, testable, refinable* version of our mental model

Essentially, a kind of logic/high-level *declarative* programming language



Q: What *frameworks* to use for formal specification/verification?
How to use formal methods to *specify/verify* systems?

Introduction

Classifying & Analyzing Systems

Q: What frameworks to use for formal specification/verification?

A: One possible option is *Rewriting Logic* (RWL), because it is:

- *Expressive*—models complex systems with few K lines
- *Modular*—supports parameterization by theories
- *Executable*—can be tested; several interpreters available

Q: How to use formal methods to specify/verify systems?

A: Specify via *declarative programming*; verify via either
Model checking: automatable, finiteness requirements or
Deductive methods: less automatable, admits most systems

Introduction

Previous Work

Previous ACC projects include analysis/verification of...

- Apache Zookeeper-based Group Key Management
- Google Megastore and extension Megastore-CGC
- Apache Cassandra
- RAMP (Read-Atomic-Multi-Partition) Systems

Q: What do the above four projects have in *common*?

A: (1) They are all *distributed databases*

(2) Analysis framework was *Rewriting Logic*

(3) They were all analyzed via *model checking*

Introduction

Previous Work

Q: These systems are clearly infinite state.
How can they be model-checked?

A: A few simplifications were necessary:

- Bound all non-determinism
- Abstract away non-essential behaviors
- Consider a finite set of initial configurations

Q: Model checking already provides high assurance;
However, it cannot cover all possible configurations.
What methods might provide even greater assurance?

A: We believe *Reachability Logic* (RL) is one such method.

Outline

- 1 Introduction
- 2 Reachability Logic**
- 3 Current Progress
- 4 Future Directions

Reachability Logic


Rewriting Logic

Q: What is Rewriting Logic?

A: A logic of *distributed states* and *concurrent interactions*

Ex: A mutual exclusion alg. QLOCK specified by 3 rewrite rules:

$$\begin{aligned} n2w &: \langle n \ i \ | \ w \ | \ c \ | \ q \ \rangle \rightarrow \langle n \ | \ w \ i \ | \ c \ | \ q \ ; \ i \rangle \\ w2c &: \langle n \ | \ w \ i \ | \ c \ | \ i \ ; \ q \rangle \rightarrow \langle n \ | \ w \ | \ c \ i \ | \ i \ ; \ q \rangle \\ c2n &: \langle n \ | \ w \ | \ c \ i \ | \ i \ ; \ q \rangle \rightarrow \langle n \ i \ | \ w \ | \ c \ | \ q \ \rangle \end{aligned}$$



Reachability Logic

Rewriting Logic Notation

QLOCK can be formalized by set of *rewrite rules* $R = \{l_i \rightarrow r_i\}_{i \in I}$

Concrete instances of rewrite rules are called *states/terms*

Rules induce *transition system* over concrete instances

Example transition (*rewrite step*):

$$\begin{array}{ccc} n2w : \langle n \ i \ | \ w \ | \ c \ | \ q \ \rangle & \rightarrow & \langle n \ | \ w \ i \ | \ c \ | \ q \ ; \ i \ \rangle \\ & & \downarrow \qquad \qquad \downarrow \\ & \langle 1 \ 3 \ 2 \ | \ mt \ | \ mt \ | \ nil \ \rangle & \rightarrow \langle 1 \ 2 \ | \ 3 \ | \ mt \ | \ 3 \ \rangle \end{array}$$

A *rewrite path* is a sequence $s_0 \rightarrow s_1 \rightarrow \dots \rightarrow s_{i-1} \rightarrow s_i$

Reachability Logic

Symbolic States

Q: How to define symbolic states?

A: Recall set-builder notation from mathematics. Ex:

$$Even = \{x \in \mathbb{N} \mid x \text{ mod } 2 = 0\}$$

We can also *constrain* terms. Let $dupl(s \ s \ s_0) = tt$. Ex:

$$\underbrace{\langle n \mid w \mid c \mid q \rangle}_{\text{term}} \quad \Bigg| \quad \underbrace{dupl(n \ w \ c) \neq tt}_{\text{constraint}}$$

This constrained term refers to states *without* duplicate ids

Reachability Logic

Introduction

Reachability Logic (RL) is:

- parameterized over an underlying *rewrite theory* \mathcal{R}
- considers formulas $A \longrightarrow^* B$ between *constrained terms* A, B
- a generalization of Hoare Logic *partial correctness*,
i.e. can view sequents $A \longrightarrow^* B$ as $\{A\}\mathcal{R}\{B\}$
- directly captures *circular* behavior in *any* theory \mathcal{R} ,
unlike Hoare Logic, special rules for loops, etc, *unnecessary*

Reachability Logic

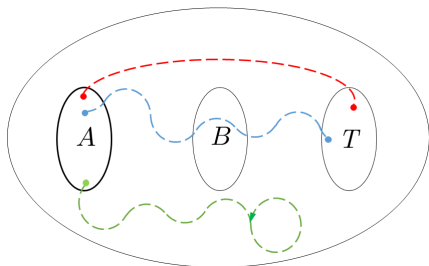
Sequents

Q: What does the relation $A \longrightarrow^* B$ mean?

A: Suppose we have:

- (1) a rewrite theory \mathcal{R}
- (2) constrained terms A, B
- (3) and terminating states T

Then $A \longrightarrow^* B$ means:
for each concrete t in A ,
for each t rewrite path p ,
either: (1) p crosses B
(2) p is infinite



- - - indicates counterex.
- - - satisfies $A \longrightarrow^* B$
- - - *vacuously* satisfies

Reachability Logic

Proof Rules

Q: Then given RWL theory \mathcal{R} , how do we prove $A \longrightarrow^* B$?

A: Perhaps surprisingly, two proof rules are enough

- A rule that traces *rewrite steps* of *symbolic* states in \mathcal{R}
- A rule that captures *circular behavior* of \mathcal{R}

We call these two rules *Step+Subsumption* and *Axiom* resp.

Reachability Logic

Proof Rules

$$\frac{\bigwedge_{(j,\alpha) \in \text{UNIFY}(u|\varphi', R)} [\mathcal{A} \cup \mathcal{C}, \emptyset] \vdash_T (r_j | \varphi' \wedge \phi_j)\alpha \longrightarrow^* \bigvee_i (v_i | \psi_i)\alpha}{[\mathcal{A}, \mathcal{C}] \vdash_T u | \varphi \longrightarrow^* \bigvee_i v_i | \psi_i}$$

$$\frac{\bigwedge_j [\{u' | \varphi' \longrightarrow^* \bigvee_j v'_j | \psi'_j\} \cup \mathcal{A}, \emptyset] \vdash_T v'_j\alpha | \varphi \wedge \psi'_j\alpha \longrightarrow^* \bigvee_i v_i | \psi_i}{[\{u' | \varphi' \longrightarrow^* \bigvee_j v'_j | \psi'_j\} \cup \mathcal{A}, \emptyset] \vdash_T u | \varphi \longrightarrow^* \bigvee_i v_i | \psi_i}$$

The *Step+Subsumption* and *Axiom* Rules

Outline

- 1 Introduction
- 2 Reachability Logic
- 3 Current Progress**
- 4 Future Directions

Current Progress

Q: So what work has been done already?

A: A substantial RL framework is already in place with:

- full semantics for RL developed in terms of RWL
- soundness proof for proof system and semantics
- working Maude prototype of proof system
- a small but growing collection of case studies

Current Progress

Case Studies

Example Name	# Goals/Lemmas
Thermostat	4
Bounded Retransmission Protocol	1
Dijkstra's Mutual Ex. Alg.	4
Fault-Tolerant Comm. Protocol	6
QLOCK Mutual Ex. Alg.	3
Readers/Writers	3
Fixed-Size Token Ring	3
Unbounded Lamport's Bakery	11
IMP Swap Function	2
IMP Min Function	2
IMP Max Function	2

Outline

- 1 Introduction
- 2 Reachability Logic
- 3 Current Progress
- 4 Future Directions**

Future Directions

Q: So, how to scale up to more complex cloud-based systems?

- No one right answer; by way of example...
let's intuitively sketch a fragment of Apache Cassandra

Q: So what is Apache Cassandra?

Future Directions

Apache Cassandra

Q: So what is Apache Cassandra?

A: A distributed key value store where:

- peers are *equal*; no master/slave
- data is copied among user-configurable # of *replicas*
- writes are *cheap*; inconsistencies handled during reads
- data tagged with timestamp; merges by *last-write-wins*

Q: Does Cassandra satisfy *eventual consistency*?

How to RL to *verify* it satisfies eventual consistency?

Future Directions

Modelling Apache Cassandra

Two rewrite rules: request generation/server request consumption

$$\langle wr [ts, k, v] \mid m \mid \{i \mid ks\} s \mid c \rangle \rightarrow$$
$$\langle wr \mid i \leftarrow [c + d \mid ts, k, v] m \mid \{i \mid ks\} s \mid c \rangle$$
$$\langle nil \mid i \leftarrow [c \mid ts, k, v] m \mid \{i \mid ks\} s \mid c \rangle \rightarrow$$
$$\langle nil \mid fwd(s, ts, k, v) \quad m \mid \{i \mid ins(ts, k, v, ks)\} s \mid c \rangle$$

Future Directions

Modelling Apache Cassandra

- Assume an initial state $A = \langle wr \mid mt \mid s \mid c \rangle$

- A consistency violation can be specified by

$$cv(\{i \mid k = v : ks\} \{j \mid k = v' : ks\} s) = tt$$

- To show eventual consistency, it is enough to prove

$$A \longrightarrow^{\otimes} \langle nil \mid mt \mid s' \mid c' \rangle \mid cv(s') \neq tt$$

Future Directions

Conclusions

There are several concrete *action steps* in the pipeline

- Support constrained terms over *undecidable theories*
- Increase automation by integrating *inductive theorem prover*
- Implement *optimizations* for built-in satisfiability methods
- Further *exploit* above techniques in analyzing cloud systems

The End

Questions?

Example

Verifying QLOCK

$n2w : \langle n \ i \mid w \mid c \mid q \rangle \rightarrow \langle n \mid w \ i \mid c \mid q ; i \rangle$

$w2c : \langle n \mid w \ i \mid c \mid i ; q \rangle \rightarrow \langle n \mid w \mid c \ i \mid i ; q \rangle$

$c2n : \langle n \mid w \mid c \ i \mid i ; q \rangle \rightarrow \langle n \ i \mid w \mid c \mid q \rangle$

Mutual exclusion property $P = \langle n \mid w \mid c \mid q \rangle \mid size(c) \leq 1$

Would like to show $\langle n \mid mt \mid mt \mid nil \rangle \xrightarrow{*} P$

But since QLOCK *never* terminates, it is impossible to verify...

Q: What can be done?

Example

Verifying QLOCK

Q: What can be done?

A: A simple theory transformation shown below

$$n2w : \langle n \ i \ | \ w \ | \ c \ | \ q \ \rangle \rightarrow \langle n \ | \ w \ i \ | \ c \ | \ q \ ; \ i \ \rangle$$
$$w2c : \langle n \ | \ w \ i \ | \ c \ | \ i \ ; \ q \ \rangle \rightarrow \langle n \ | \ w \ | \ c \ i \ | \ i \ ; \ q \ \rangle$$
$$c2n : \langle n \ | \ w \ | \ c \ i \ | \ i \ ; \ q \ \rangle \rightarrow \langle n \ i \ | \ w \ | \ c \ | \ q \ \rangle$$
$$term : \langle n \ | \ w \ | \ c \ | \ q \ \rangle \rightarrow [n \ | \ w \ | \ c \ | \ q \]$$

New mutual ex. property $[P] = [n \ | \ w \ | \ c \ | \ q] \mid size(c) \leq 1$

Then prove $\langle n \ | \ mt \ | \ mt \ | \ nil \ \rangle \mid dupl(n) \neq tt \xrightarrow{*} [P]$

Needed lemma $\langle n \ | \ w \ | \ c \ | \ q \ \rangle \mid size(c) \leq 1 \xrightarrow{*} [P]$