


EXECUTING CODE IN THE PAST: EFFICIENT IN-MEMORY OBJECT GRAPH VERSIONING


PLUQUET FRÉDÉRIC
STEFAN LANGERMAN
ROEL WUYTS

IN-MEMORY OBJECT VERSIONING


IN-MEMORY OBJECT VERSIONING




IN-MEMORY OBJECT VERSIONING




IN-MEMORY OBJECT VERSIONING



IN-MEMORY OBJECT VERSIONING




IN-MEMORY OBJECT VERSIONING



«If I want revisit these 3 versions, how can I do it ?»

IN-MEMORY OBJECT VERSIONING



«If I want revisit these 3 versions, how can I do it ?»

⇒ An ad hoc solution must be built

BUT...

- It is a recurring problem !
 - Eiffel-like Checked Post-Conditions
 - Stateful Tracer, Debugger, Google Wave
 - More than 20 applications in the theoretical algorithmic domain (computational geometry,...)
 - ...

OUR SOLUTION

Introduce In-Memory Object Versioning in any Object-Oriented Language

- What are the challenges ?
 - Be generally applicable !
 - Be expressive !
 - Be efficient in time and in space !

INTUITIVE APPROACHES

INTUITIVE APPROACHES

- Deep copy of the system after each change

INTUITIVE APPROACHES

- Deep copy of the system after each change
 - ➔ The memory fills up very quickly

INTUITIVE APPROACHES

- Deep copy of the system after each change
 - ➔ The memory fills up very quickly
- Deep copy of the system at each given interval of time

INTUITIVE APPROACHES

- Deep copy of the system after each change
 - ➔ The memory fills up very quickly
- Deep copy of the system at each given interval of time
 - ➔ Some useful values can be not saved

INTUITIVE APPROACHES

- Deep copy of the system after each change
 - ➔ The memory fills up very quickly
- Deep copy of the system at each given interval of time
 - ➔ Some useful values can be not saved
- Use database

INTUITIVE APPROACHES

- Deep copy of the system after each change
 - ➔ The memory fills up very quickly
- Deep copy of the system at each given interval of time
 - ➔ Some useful values can be not saved
- Use database
 - ➔ Not the same goal

WHAT WE REALLY WANT

WHAT WE REALLY WANT

- Only keep

WHAT WE REALLY WANT

- Only keep
- interesting old values

WHAT WE REALLY WANT


- Only keep
 - interesting old values
 - of interesting fields

WHAT WE REALLY WANT

- Only keep
 - interesting old values
 - of interesting fields
 - of interesting objects

WHAT WE REALLY WANT

- Only keep
 - interesting old values
 - of interesting fields
 - of interesting objects



WHAT WE REALLY WANT

- Only keep

- **interesting** old values

⇒ *Snapshots*


- of **interesting** fields

⇒ *Field selection*


- of **interesting** objects

EXAMPLE


EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE




EXAMPLE



EXAMPLE




EXAMPLE




HOW TO EFFICIENTLY STORE OLD VALUES ?

- *Fat Node Method* [Driscoll et al,1986].
 - Store the old values of a field in the field itself



HOW TO EFFICIENTLY STORE OLD VALUES ?



Insert a new value	$O(1)$
Access last value	$O(1)$
Look up for a value	$O(\log(\#values))$

COMPLEXITIES

COMPLEXITIES

	<i>Worst case</i>
Take a snapshot	$O(1)$

COMPLEXITIES

	<i>Worst case</i>
Take a snapshot	$O(1)$
Perform an ephemeral update	$O(1)$

COMPLEXITIES

	<i>Worst case</i>
Take a snapshot	$O(1)$
Perform an ephemeral update	$O(1)$
Perform a versioned update	$O(1)$

COMPLEXITIES

	<i>Worst case</i>
Take a snapshot	$O(1)$
Perform an ephemeral update	$O(1)$
Perform a versioned update	$O(1)$
Read last value of a field	$O(1)$


COMPLEXITIES

	<i>Worst case</i>
Take a snapshot	$O(1)$
Perform an ephemeral update	$O(1)$
Perform a versioned update	$O(1)$
Read last value of a field	$O(1)$
Read old value of a field	$O(\log \# \text{values})$


AND OF COURSE

- Don't lose time to save useless data
- Developer selects !

EFFICIENCY IN SIZE




EFFICIENCY IN SIZE




- States are shared between snapshots

EFFICIENCY IN SIZE



- States are shared between snapshots

EFFICIENCY IN SIZE



- States are shared between snapshots
- The states are in the object itself
 - If an object is no longer reachable, it is garbage collected
 - Including all of its fields history

IMPLEMENTATION

HISTOORY

- Early version in Java
- Stable version in Smalltalk

SMALLTALK INTEGRATION

- A standard library
 - Can be loaded in any application
 - Existing code does not need modification
- Three primitives
 - `anObject selectFields`
 - `s := HSnapshot atNow`
 - `s execute: aBlock`
- Easy to learn and use

SMALLTALK INTEGRATION

borrowing | aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.

returning | aBook state: 'dirty'.
aBook borrower: nil.

HISTORY

SMALLTALK INTEGRATION

borrowing

```
aClient := Client named: 'John'.  
aBook := Book titled: '1985'.  
aBook state: 'clean'.  
aBook borrower: aClient.  
aBook selectFields: {#state. #borrower}.
```

returning

```
aBook state: 'dirty'.  
aBook borrower: nil.
```

HISTORY

SMALLTALK INTEGRATION

borrowing | aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.
aBook selectFields: {#state. #borrower}.

returning | s1 := HSnapshot atNow.
aBook state: 'dirty'.
aBook borrower: nil.

HISTORY

SMALLTALK INTEGRATION

borrowing aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.
aBook selectFields: {#state. #borrower}.

returning **s1 := HSnapshot atNow.**
aBook state: 'dirty'.
aBook borrower: nil.

aBook state.

SMALLTALK INTEGRATION

borrowing | aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.
aBook selectFields: {#state. #borrower}.

returning | s1 := HSnapshot atNow.
aBook state: 'dirty'.
aBook borrower: nil.

aBook state. —————→ 'dirty'

SMALLTALK INTEGRATION

borrowing | aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.
aBook selectFields: {#state. #borrower}.

returning | s1 := HSnapshot atNow.
aBook state: 'dirty'.
aBook borrower: nil.

aBook state. —————→ 'dirty'
s1 execute: [aBook state].

SMALLTALK INTEGRATION

borrowing aClient := Client named: 'John'.
aBook := Book titled: '1985'.
aBook state: 'clean'.
aBook borrower: aClient.
aBook selectFields: {#state. #borrower}.

returning s1 := HSnapshot atNow.
aBook state: 'dirty'.
aBook borrower: nil.

aBook state. —————→ 'dirty'
s1 execute: [aBook state]. —→ 'clean'

PERFORMANCE

PERFORMANCE

	<i>Slowdown factor</i>
Synthetic worst case update	7


PERFORMANCE

	<i>Slowdown factor</i>
Synthetic worst case update	7
Real cases update	1.3 to 2

APPLICATIONS


- We implemented
 - Eiffel-like Checked Post-Conditions for Smalltalk
 - Execution tracer with states
 - Planar Point Location

BENCHMARKS




Stores messages sent in a collection

BENCHMARKS




Stores messages sent in a collection
Also stores states of receiver and arguments

BENCHMARKS



Stores messages sent in a collection
Also stores states of receiver and arguments

BENCHMARKS



Only 30%
slower !

Stores messages sent in a collection
Also stores states of receiver and arguments

CONCLUSION

- A general and efficient model for in-memory object versioning for object-oriented languages
- We implemented it in Smalltalk
 - 3 primitives
 - Non-intrusive library
 - Efficient

THANK YOU !

- fpluquet@ulb.ac.be
- HIST[⌚][⌚]ORY on Google
- *Executing code in the past: Efficient in-memory object graph versioning*, F. Pluquet, S. Langerman, and R. Wuyts, In Proceedings of the 2009 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'09).
- *Implementing Partial Persistence In Object-Oriented Languages*, F. Pluquet, S. Langerman, A. Marot and R. Wuyts, In Proceedings of ALENEX'08, 2008.