

# AutoBash: Improving Configuration Management with Operating System Causality Analysis

Ya-Yunn Su, Mona Attariyan, and Jason Flinn  
University of Michigan

Presented by: Fahad Arshad




Slide 1/20



## Motivation

- Configuration management is frustrating!
- Users may have to
  - Change environment variables
  - Edit configuration files
  - Manage inter-application dependencies
- Current approach:
  - Ask friends, search on-line, read manual, ...
  - Try potential solutions
  - Carefully undo wrong solutions



Slide 2/20



## Problems with current approach AutoBash solves these problems

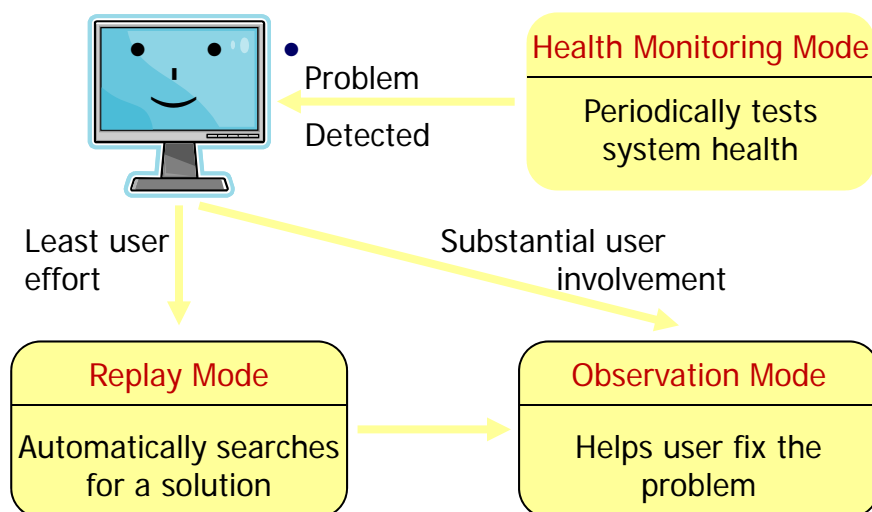
- Applying solutions is time-consuming  
Automatically tries many solutions
- Undoing a wrong solution can be hard  
Provides undo capability
- Hard to know how a problem was solved  
Explains solution to user
- A “solution” may cause new problems  
Automatically runs regression tests



Slide 3/20



## AutoBash overview



Slide 4/20



## Outline

- Motivation
- AutoBash design and implementation
  - Observation mode
  - Replay mode
  - Health monitoring mode
- Evaluation
- Conclusion




Slide 5/20



## Observation mode

- A modified bash shell
  - User types in commands to solve the problem



```
% command 1
% test if app works
% undo testing
% undo command 1
% command 2
```



Slide 6/20



## Verifying a solution is tedious

- AutoBash automatically tests using *predicates*
- Predicate:
  - Tests if an application functions correctly
  - Returns true/false if the test passes/fails

```
% command 1
% test if app works
% undo testing
% rollback command 1
% command 2
```

Predicate to server



## Undoing testing is tedious


- Predicate testing has no side effects
  - Executed speculatively and rolled back
- Speculator [SOSP '05]
  - Process-level speculative execution

```
% command 1
% test if app works
% undo testing
% rollback command 1
% command 2
```



## Undo can be hard

- AutoBash speculatively executes each action
  - Light-weight checkpoint and rollback



- Speculative execution makes undo easy




Slide 9/20



## Regression testing is hard

- AutoBash automatically runs regression tests
  - Executes predicates in the predicate database
  - Ensures all predicates pass



Slide 10/20



## Regression tests can be slow

- Problem: running all predicates can be slow



- Only need to run predicates affected by an action
  - Uses causality tracking to find affected predicates



Slide 11/20



## Tracking causality

- Output set
  - kernel objects an action causally affects

Action: touch foo

- Input set Output set = {file foo}
  - kernel objects a predicate causally depends on

Predicate: grep "test" bar

Input set = {file bar}

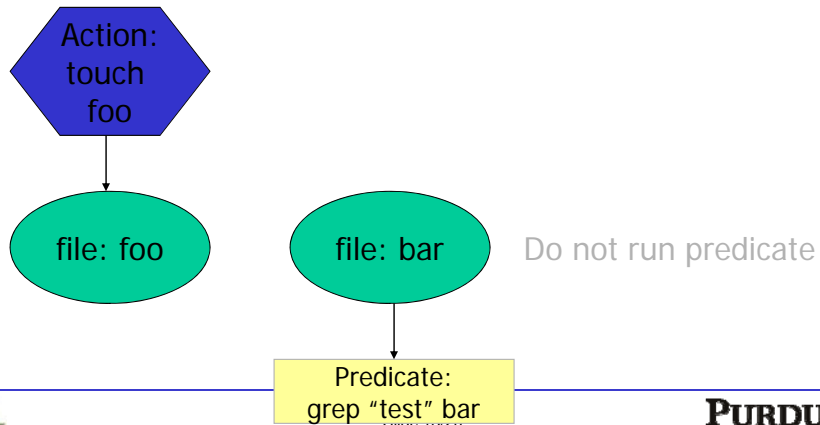


Slide 12/20



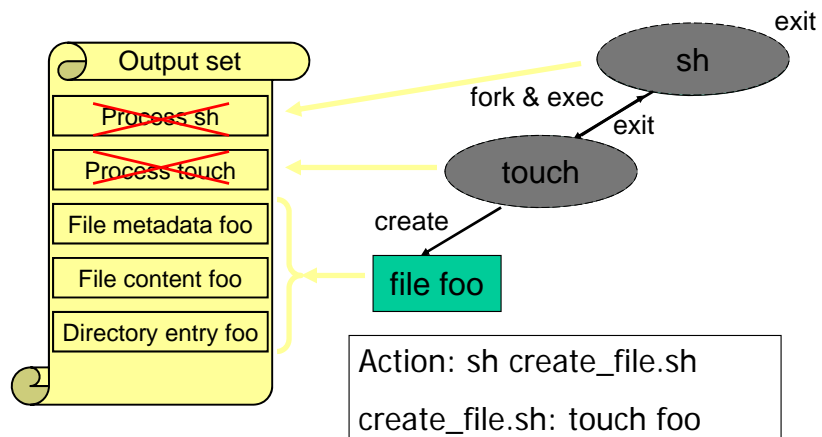
## Analyzing causality

- AutoBash calculates the intersection
  - Determines which predicates to run



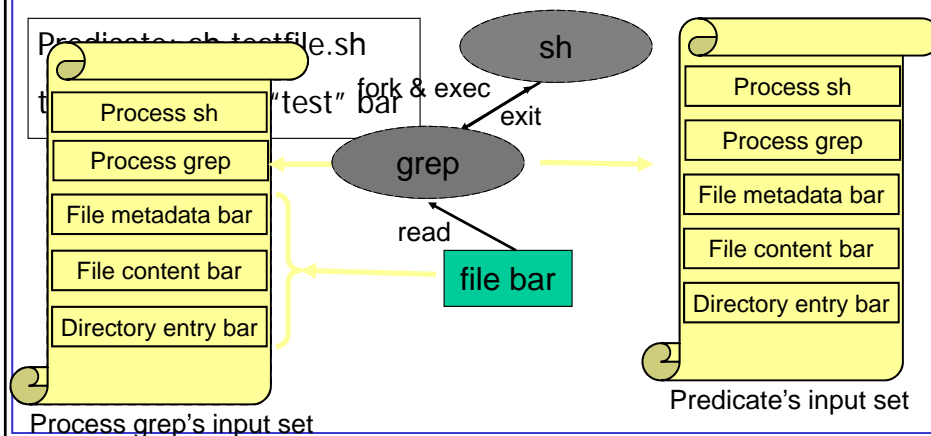
## Tracking output sets

- An output set is tracked for each action



## Tracking input sets

- An input set is tracked for each predicate



Slide 15/20



## Understanding solutions can be hard

- AutoBash generates causal explanation
  - Analyzes input and output sets




Slide 16/20






## Causal explanation

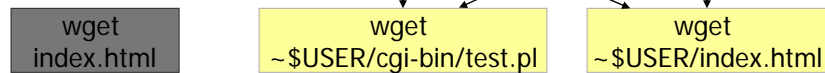
Actions



Kernel objects



Predicates



Slide 17/20



## Outline

- *Motivation*
- AutoBash design and implementation
  - *Observation mode*
  - *Replay mode*
  - *Health monitoring mode*
- **Evaluation**
- **Conclusion**



Slide 18/20



## Replay mode


- Problem: finding a solution is time-consuming
- Automatically searches for a solution
  - No user input needed
- Speculative execution provides isolation
  - User continues foreground task
  - AutoBash runs replay mode in background



Slide 19/20



## How replay mode works



(1) Initial predicate testing:


- Tracks input set for each predicate
- Determines passed/failed predicates



Slide 20/20



## How replay mode works



(2) Solution execution:


- Speculatively executes a solution
- Tracks solution output set



Slide 21/20




## How replay mode works



(3) Verifying solution:

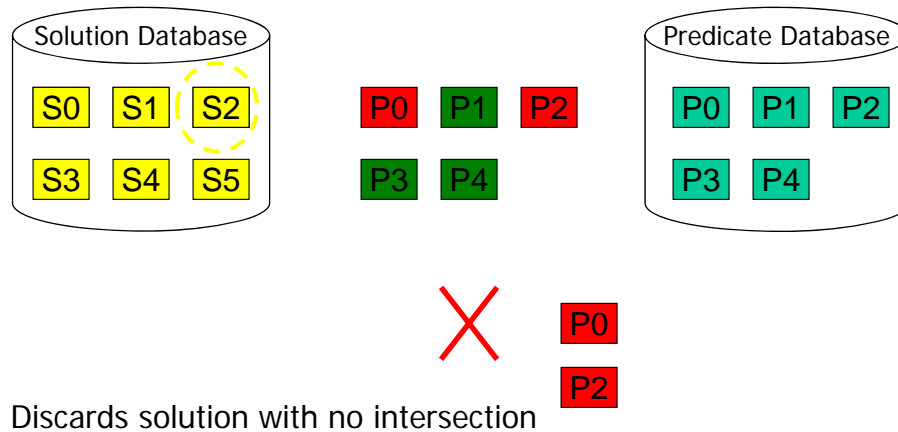
- Calculates intersection
- Runs predicates with intersection



Slide 22/20



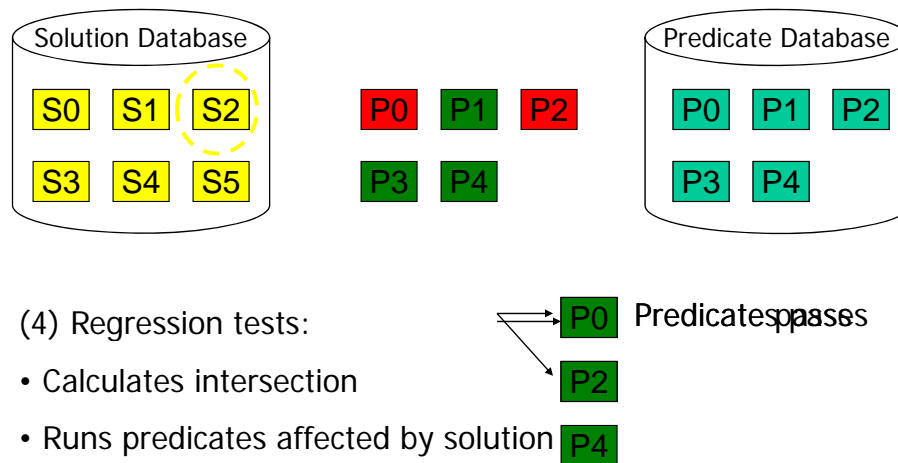
## How replay mode works



Slide 23/20



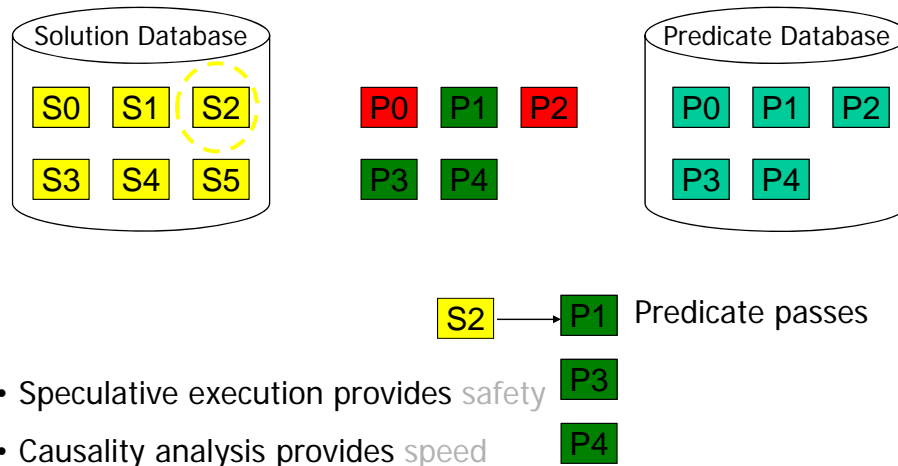
## How replay mode works



Slide 24/20



## How replay mode works



Slide 25/20



## Health monitoring mode

- Periodically executes all predicates
- If any predicate fails, AutoBash
  - Runs replay mode to search for a solution
  - Reports to the user to run observation mode



Slide 26/20



## Outline

- *Motivation*
- *AutoBash Design and Implementation*
  - *Observation mode*
  - *Replay mode*
  - *Health monitoring mode*
- *Evaluation*
- **Conclusion**



Slide 27/20



## Evaluation


- **Questions:**
  - What is the overhead of speculative execution?
  - How effective is causality analysis?
- **Methodology:**
  - Evaluated CVS, gcc cross compiler, web server
  - Manually created 10 bugs and 10 solutions
  - Manually created 5-8 predicates



Slide 28/20



## Total replay time (GCC)




Speculative execution overhead is negligible

Slide 29/20

PURDUE  
UNIVERSITY

## Total replay time (GCC)



Causal analysis improves predicate re-testing time by 67-99%

Slide 30/20

PURDUE  
UNIVERSITY

## Conclusion

- Configuration management is frustrating
- AutoBash automates most tedious parts
- Speculative execution makes AutoBash safe
- Causality analysis makes AutoBash fast



Slide 31/20



**Microsoft®**



Slide 32/20

