

Analysis is necessary – but far from sufficient

Jon Pincus
Reliability Group (PPRC)
Microsoft Research

Why are so few successful real-world development and testing tools influenced by academic research?

My definition of “real world”:

- Commercial or quasi-commercial
- Software and net services
(Excluding IT or Mil-Aero – I have no background there)

Jon Pincus (Microsoft Research)

2

Outline

- What makes a tool successful?
- Characteristics of successful tools
- Analysis – in context
- Implications for analysis
- Summary

Jon Pincus (Microsoft Research)

3

Success!

- Cute diagram

Jon Pincus (Microsoft Research)

4

Success: a real-world view

- A tool is successful if people use it
 - Not if people buy it but don't use it (“Shelfware”)
 - Not if people try it but don't use it

Jon Pincus (Microsoft Research)

5

Some examples of success

(drawn from *defect detection* space, because
that's my background)

- Purify
- BoundsChecker
- PREFIX (2.X and later)
 - Especially interesting because 1.0 was unsuccessful

Jon Pincus (Microsoft Research)

6

Why do people use a tool? If

- it helps them get their work done ...
- ... more efficiently than they would otherwise
- ... without making them look bad.

- Think in terms of
 - Goals
 - Value proposition

Jon Pincus (Microsoft Research)

7

Goals

- *Organizational goals*
 - “compensate for not being able to find enough good developers/QA engineers”
 - “get higher-quality software to market”
 - “get high-quality software to market more quickly”
 - “avoid the memory leaks that plagued our last release”
 - “stop breaking the build”
- *Personal goals*
 - “stop having to waste my time on others’ mistakes”
 - “find that killer bug”
 - “stop getting blamed for breaking the build”

Jon Pincus (Microsoft Research)

8

Three definite non-goals

- Looking stupid
- Not being able to use techniques I already know
- Creating additional work

For more on goals: Alan Cooper’s *About Face*
http://www.cooper.com/books/01_goal_directed_design.html

Jon Pincus (Microsoft Research)

9

An example

- People want to fix the key defects as easily as possible
- PREFIX 1.0: “detects defects”
 - Too far from goal to be generally useful
 - Not successful [although it found a lot of real defects]
- PREFIX 2.X and later: focus on *prioritizing and understanding* defects
 - Successful (although fixing them would be even better)
 - Techniques: focus on user interaction, data storage, repository, understandability, prioritization, ...

Jon Pincus (Microsoft Research)

10

Look at the *value proposition*

- (Value – Cost) must be
 - Positive
 - More positive than any alternatives
- Initially, cost will exceed value; how long until payback?
- Value: the benefit of the tool
 - E.g., higher quality
- Cost: more complex

Jon Pincus (Microsoft Research)

11

Cost

- Time investment
 - Initial use
 - Steady-state use
 - Training
 - Existing employees
 - New employees
- Process changes
- Licensing cost
 - Zero for “free” software
 - Typically much smaller than the others

Jon Pincus (Microsoft Research)

12

An example

- Purify 1.0:
 - Virtually zero initial cost on most code bases
 - Immediate value
 - Companies invested (substantially!) to leverage
 - E.g., changing memory allocators to better match Purify's

Jon Pincus (Microsoft Research)

13

An example

- PREFIX 1.0:
 - Major initial cost
 - Usability issues increase ongoing costs
 - There's value; but nobody got to sustained value
- PREFIX 2.X:
 - Lowered initial cost
 - Improved usability decreased steady-state and training cost
 - Value – to people who care a lot about reliability
- Ongoing work: decrease costs further

Jon Pincus (Microsoft Research)

14

Outline

- What makes a tool successful?
- Characteristics of successful tools
- Analysis – in context
- Implications for analysis
- Summary

Jon Pincus (Microsoft Research)

15

Characteristics of successful tools

- Successful tools
 - address significant problems,
 - on real code bases,
 - give something for (almost) nothing,
 - and are easy to use.

Jon Pincus (Microsoft Research)

16

Significant problems

- Nobody fixes all the bugs. What are the key ones?
 - Often based on most recent scars
 - Often based on development or business goals

Jon Pincus (Microsoft Research)

17

Examples: significant problems

- Purify: memory leaks
- BoundsChecker: bounds violations
- PREFIX: defects not found by existing tests
- Lint (back in K&R days): portability issues

Jon Pincus (Microsoft Research)

18

Example: insufficiently significant problems

- Vanilla lint (today): ???
(Note: PC-Lint/FlexeLint extend Lint to attack today's problems)
- Pointer analyses for its own sake
(although it may be useful for solving another problem)
- C/C++ metrics tools

Jon Pincus (Microsoft Research)

19

Real code bases

- Usually large (1M+LOC) or very large (10M+ LOC).
 - In some areas, “reality” is smaller – e.g., 100s of lines of DHTML/JScript
- In nasty languages (e.g., C/C++),
 - using nasty features (e.g., casts between pointers and ints, unions, bit fields, gotos, ...)
 - with nasty extensions (GCC, MS)
 - and non-ANSI-compliant code (GCC, Sun, MS)

Jon Pincus (Microsoft Research)

20

Examples: insufficiently real code bases

- “For a subset of C, excluding pointers, structs, and unions ...”
- “Assuming the whole program text is available (i.e., there are no calls to system libraries) ...”
- “We have tested on our approach on programs up to several thousand lines of Scheme ...”

Jon Pincus (Microsoft Research)

21

A “fully general solution” includes

(Please assume appropriate trademark/copyright symbols)

- Perl
- C [K&R, ANSI], C++ [Cfront, GCC extensions, MSVC extensions], Java, C#
- VB, TCL
- ECMAScript [JScript, JavaScript], VBScript, Python
- HTML [HTML3.2, Netscape/MS variants], DHTML
- SQL
- XML, XSL, XSL-T, XML Schemas
- FORTRAN, COBOL for legacy code
- Make, sh, InstallShield, IDL, Excel macros, ...

Jon Pincus (Microsoft Research)

22

Something for (almost) nothing

- Engineering time is the single most critical resources at most (successful) companies
- Engineers need to be convinced before investing non-trivial amounts of time
- So don't even *think* about *requiring* significant up-front investment
 - code modifications
 - process changes

Jon Pincus (Microsoft Research)

23

Examples: something for (almost) nothing

- Purify for UNIX: just relink
- BoundsChecker: you don't even need to relink!
- PREfix 3.5: just type “prefix”
 - (Uh, most of the time, anyhow)
- A non-technology solution: “we'll do it for you”
 - Commercial variant: an initial benchmark for \$X, money back if it doesn't work
 - In many cases, money may be cheaper than engineering time ...

Jon Pincus (Microsoft Research)

24

Examples: too far from nothing

- “Once the programmer modifies the code to include calls to the appropriate functions ...”
- “The programmer simply inserts the annotations to be checked as conventional comments ...”
- PREFIX 1.0: “the following changes to your build process may prove necessary ...”

Jon Pincus (Microsoft Research)

25

Ease of use

- Admittedly, the bar is low here ...
[not sure exactly what to say; it seems so self-evident ...]

Jon Pincus (Microsoft Research)

26

Examples: ease of use

- [use an example from the CAD space. Place and route is highly algorithmic; but these days, it’s ease of use which gets the best results – because engineers can make one more iteration in the time allotted for a benchmark]

Jon Pincus (Microsoft Research)

27

Outline

- What makes a tool successful?
- Characteristics of successful tools
- Analysis – in context
- Implications for analysis
- Summary

Jon Pincus (Microsoft Research)

28

Counterintuitively ...

Actual analysis is only a small part of any program analysis tool.

In PREFIX, < 10% of the “code mass”

Jon Pincus (Microsoft Research)

29

PREFIX architecture

- Pretty picture here

Jon Pincus (Microsoft Research)

30

PREfix' key operations

- Parsing
- Calculating function dependencies
- Walking paths through functions
- Tracking memory during simulation
- Generating and storing models
- Generating and storing defect information
- Viewing/sorting/filtering sets of defects
- Viewing paths through source code
- Build integration

Jon Pincus (Microsoft Research)

31

3 key non-analysis issues

- User interaction
 - Information presentation
 - Navigation
 - Control
- Integration
 - Build process
 - Defect tracking system
 - SCM system
- Parsing

Jon Pincus (Microsoft Research)

32

User interaction

- Engineers must be able to use the results of the analysis
 - Understanding individual defects
 - Prioritizing, sorting, and filtering sets of defects
 - Interacting with other engineers
 - Controlling the analysis (because analyses aren't perfect)
- Today, the bar is ridiculously low
 - A good place to make progress!

Jon Pincus (Microsoft Research)

33

Example

- [single-line Dereferencing NULL Pointer message]

Jon Pincus (Microsoft Research)

34

A better example

- [Complex code path failing to check new]

Jon Pincus (Microsoft Research)

35

A still better example

- Message describing the problem

Jon Pincus (Microsoft Research)

36

Noise

- [definition of noise]
- [deal with it at analysis level, or at rest of system level?]

Some interesting questions ...

- How to summarize information usefully?
- How to visualize (sets of) (partial) paths through code?
- Can analysis refine presentation?

Integration

- A tool is useless if people can't use it
 - Implied: "use it *in their existing environment*"
- "Environment" includes
 - Configuration management (SCM)
 - A build process (makefiles, scripts, ...)
 - Policies
 - A defect tracking system
- People have invested *a lot* in their environment
 - They probably won't change it just for one tool

Approaches to integration

- Special-case methods seem the norm
 - E.g., "intercepting" build commands; special purpose scripts
 - PREFIX' latest attempt:
 - treat build information as first-class data – tracked in database, used in analyses, ...
 - move to more general "repository"
- Can this be generalized?
- Is there a better way?

Parsing

- You can't parse better than anybody else ...
... but you *can* parse worse
- Complexities:
 - Incompatibilities
 - Extensions
 - Full language complexity
 - Language evolution

Approaches to Parsing

- Don't
 - Alternatives: EDG, GCC, Jikes, ...

Outline

- What makes a tool successful?
- Characteristics of successful tools
- Analysis – in context
- Implications for analysis
- Summary

Jon Pincus (Microsoft Research)

43

Characteristics of useful analyses

- Scalable to large system
 - Typically implies *incomplete, unsound, incremental, and/or very simple*
- Produce information usable by typical engineer
 - If there's a violation, where? How?
 - Post-processing output can be useful
 - Remember: half the engineers are below average
- “Accurate enough” for the particular task
- Handle full language complexity
 - Or can compensate for unhandled constructs

Jon Pincus (Microsoft Research)

44

Different tradeoffs from compilers

- Focus on information, not just results
 - Compilers don't have to explain what they did and why
- Incompleteness and unsoundness may be okay
- Intra-procedural analysis often not enough

Jon Pincus (Microsoft Research)

45

A spectrum of analyses

[make this a chart?]

- Flow- and context-insensitive:
 - Typically scales well.
 - Not particularly accurate (but clearly accurate enough for some tasks)
 - Often hard to understand or prioritize the output [no path information; no callstack]
- Flow- and context-sensitive
 - Scaling problems.
 - More accurate; still issues with path-sensitivity
 - Info may be more understandable
- Path-sensitive
 - Non-PREfix Examples?

Jon Pincus (Microsoft Research)

46

Examples of analysis tradeoffs

- Purify/BoundsChecker: very simple
- PREfix: incomplete, somewhat unsound, incremental
- Lint (without post-processing): not “accurate enough”, often not usable by typical engineer
- PREfast: consciously tradeoff completeness for performance

Jon Pincus (Microsoft Research)

47

Some interesting questions ...

- Which analyses are right for which problems?
 - No such thing as the right pointer analysis – it depends what you want to do with the results
- Are there opportunities to combine analyses?
 - Can we use a cheap flow-insensitive algorithm to focus a more expensive algorithm on juicy places?
 - Can we use expensive local path-sensitive algorithms to improve global flow-insensitive algorithms?

Jon Pincus (Microsoft Research)

48

Outline

- What makes a tool successful?
- Characteristics of successful tools
- Analysis – in context
- Implications for analysis
- Summary

Jon Pincus (Microsoft Research)

49

Summary

- People use tools to accomplish their tasks
- Successful tools must
 - address real problems,
 - on real code bases,
 - give something for (almost) nothing,
 - and be easy to use
- Analysis is only one piece of a tool
- Information is useless if it's not presented well

Jon Pincus (Microsoft Research)

50

Why are there so few successful real-world programming and testing tools based on academic research?

These are not where research has focused.
Can – and should – that change?

Jon Pincus (Microsoft Research)

51

Questions?

Jon Pincus (Microsoft Research)

52

Analysis is necessary –
but far from sufficient

Jon Pincus
Reliability Group (PPRC)
Microsoft Research