Report From the CoalFace

Lessons learnt building OPUS, a General-Purpose, Always-on Provenance System

Nikilesh Balakrishnan, Thomas Bytheway
Ripduman Sohan, Andy Hopper

University of Cambridge
Computer Laboratory

What is OPUS?

What is OPUS?

User-space

What is OPUS?

User-space

Motivation

What is OPUS?

User-space

Motivation

Non-intrusive

What is OPUS?

User-space

Motivation

Non-intrusive

Easy to deploy

What is this paper about?

What is this paper about?

Implementation challenges

What is this paper about?

Implementation challenges

Library level interposition

What is this paper about?

Implementation challenges

Library level interposition

General principles

Challenges

Challenges

1. Interposition at scale

Challenges

1. Interposition at scale
2. Semantic Equivalence

Challenges

1. Interposition at scale
2. Semantic Equivalence
   (i) vfork

Challenges

1. Interposition at scale
2. Semantic Equivalence
   (i) vfork
   (ii) signals

Interposition at scale

write

Interposition at scale

Interposition at scale

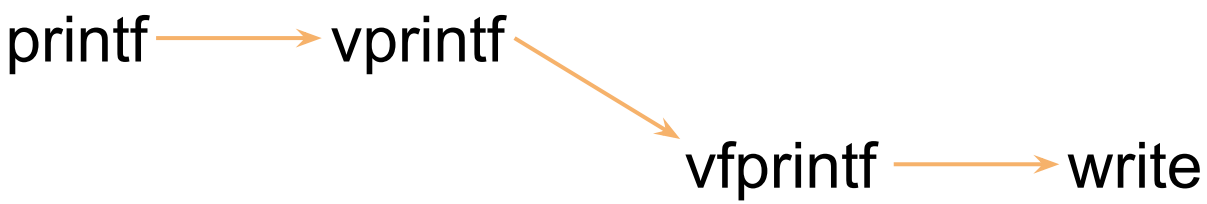

Interposition at scale

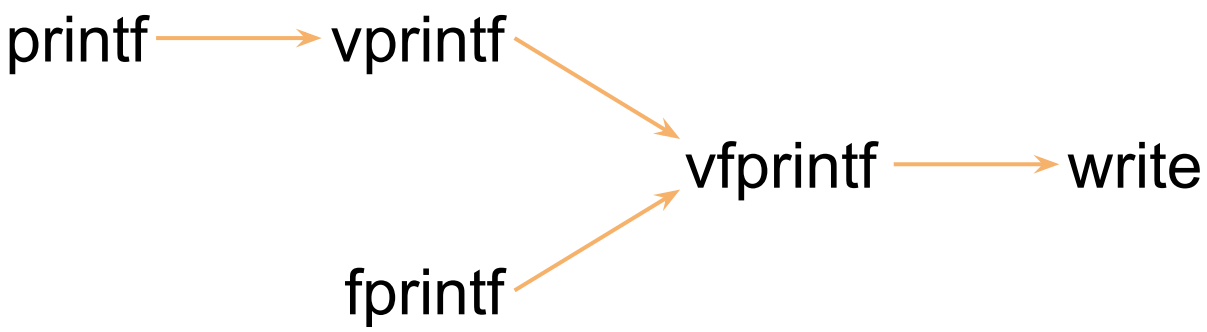

Interposition at scale

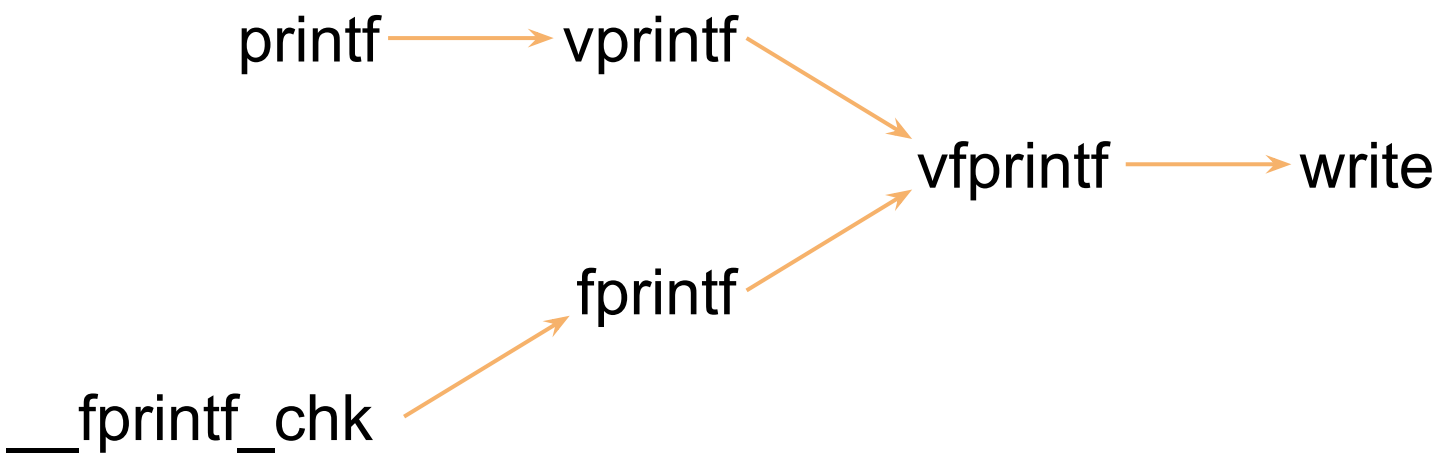

Interposition at scale

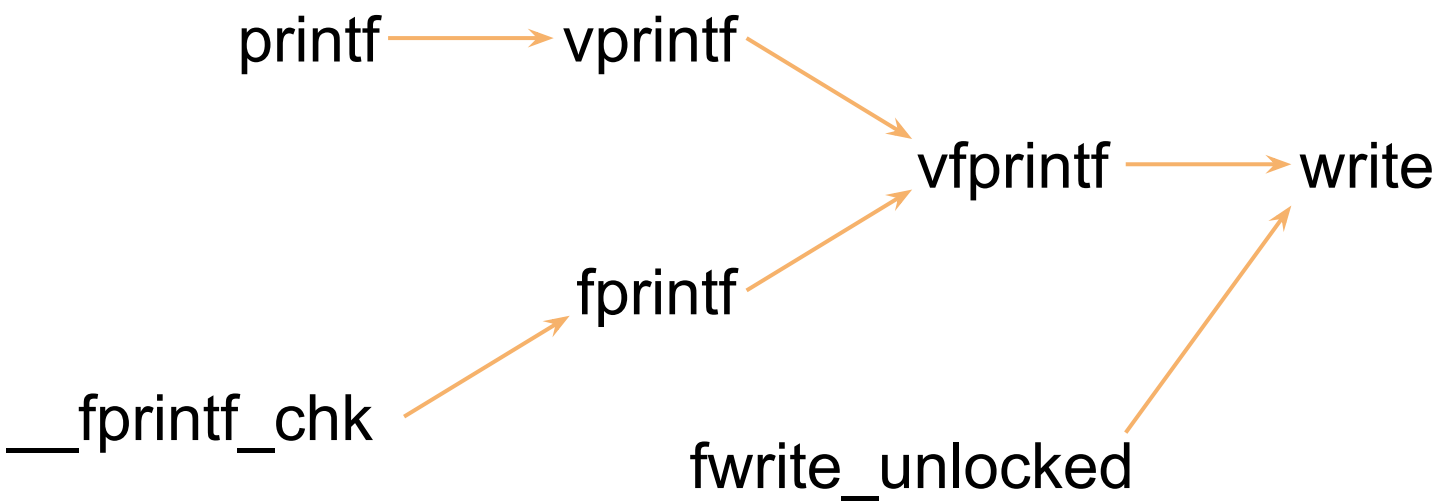

Interposition at scale

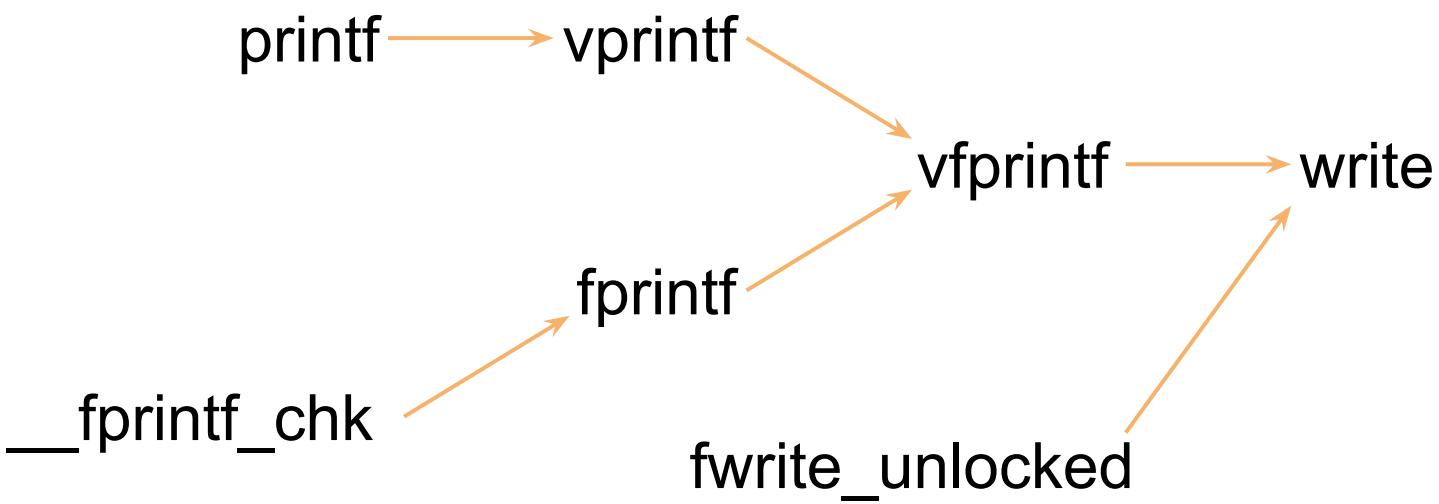

~160 functions

Interposition solution

Interposition solution

Code generation by templating

Interposition solution

Code generation by templating

- Easy to add new functions.

Interposition solution

Code generation by templating

- Easy to add new functions.
- Easy to extend to other libs.

Manual effort still required

Manual effort still required

Process related (fork, exec etc.)

Manual effort still required

Process related (fork, exec etc.)

Signal Handling

Manual effort still required

Process related (fork, exec etc.)

Signal Handling

Thread library

Semantic Equivalence - vfork

Semantic Equivalence - vfork

Legacy

Semantic Equivalence - vfork

Legacy

Deprecated

Semantic Equivalence - vfork

Legacy

Deprecated

Still being used

Semantic Equivalence - vfork

Legacy

Deprecated

Still being used

Interposition is tricky

Path to solution

Path to solution

Mimic vfork using fork

Path to solution

Mimic vfork using fork

Store in a register (glibc)

Path to solution

Mimic vfork using fork

Store in a register (glibc)

Store in memory (our solution)

Semantic Equivalence - signals

Semantic Equivalence - signals

Capture signal events

Semantic Equivalence - signals

Capture signal events

Wrap application’s handlers

Semantic Equivalence - signals

Capture signal events

Wrap application’s handlers

Arcane interface

Semantic Equivalence - signals

Capture signal events

Wrap application’s handlers

Arcane interface

Maintain extra state

Takeaways

Takeaways

Library level function capture

Takeaways

Library level function capture

Completeness

Takeaways

Library level function capture

Completeness

Semantic equivalence

Takeaways

Library level function capture

Completeness

Semantic equivalence

Engineering effort

Future of OPUS

Visualization

Future of OPUS

Visualization

Base platform

Future of OPUS

Visualization

Base platform

Open source

Thank You

For more info visit:

http://www.cl.cam.ac.uk/research/dtg/fresco/