

# Securing Systems via Design and Proof



# AMPLSE

# Software Infrastructure



# Software Infrastructure is Shaky

# Software Infrastructure is Shaky

## The New York Times

### Cars' Computer Systems Called at Risk to Hackers

By JOHN MARKOFF  
Published: May 14, 2010

Automobiles, which will be increasingly of the near future, could be vulnerable to hackers, now, two teams of computer scientists are presenting next week.

The scientists to remotely control functions, was running security in an adversarial way, "We demonstrated an automated adversary including disabling the brakes, selectively the engine, and so on," they wrote in the Modern Automobile."

In the paper, which will be presented at Oakland, Calif., computer security specialists at the University of California, San Diego, reported engineering in the design of their computers to the potential threat of hackers who might

### Medical Devices

Medical Device Safety

- 2012 Medical Device Recalls
- 2011 Medical Device Recalls
- 2010 Medical Device Recalls
- 2009 Medical Device Recalls
- 2008 Medical Device Recalls
- 2007 Medical Device Recalls
- 2006 Medical Device Recalls
- 2001-2005 Medical Device Recalls

Recent Medical Devices

- GE Healthcare, LLC, GE Healthcare's Tricore Resuscitation Systems
- St. Jude Medical, AMPLATZER TorqVue FX Delivery System
- Hamilton Medical, Inc., HAMILTON-T1 Ventilators with Software Versions 1.2 and Lower
- Vycoor Medical, Inc., Vycoor ViewSite Brain Access System (VBAS)
- Bausch and Lomb 370 Sterile Cannula Packed in Bausch and Lomb Amvisc 1.2% Sodium Hyaluronate (Model 59051, 59051L, 59051LL) and Amvisc Plus 1.4% Sodium Hyaluronate

## Bloomberg Businessweek Markets & Finance

### Software Bug Made Swedish Exchange Go Bork, Bork, Bork

By Karen Weise on November 29, 2012

A computer error stalks the markets—again. An order on a relatively obscure derivatives index in Stockholm yesterday was asking to buy futures contracts on Swedish stocks valued at **131 times the country's entire GDP**. The order made the exchange go "bananas" and caused Nasdaq OMX to stop trading in Swedish derivatives for four hours.

This was no "fat finger" incident, where a trader accidentally types an extra few digits or the wrong numbers in an order. Instead, a software glitch magnified an order, Nasdaq OMX spokesman Carl Norell told Bloomberg News. "Our system misinterpreted a certain order category and communicated a value that was way too high into the book," he said.

The interruption was in a small corner of the market, but it's just the latest in a string of technical problems that have halted trading. As more trading is driven by the algorithms of high-frequency traders, one glitch or bad order can spark major disruptions. The 2010 flash crash caused \$862 billion in stock values to vanish from the market temporarily, and technical problems have

## The New York Times

### Cars' Computer Systems Called at Risk to Hackers

By JOHN MARKOFF  
Published: May 14, 2010

Automobiles, which will be increasingly of the near future, could be vulnerable to hackers, now, two teams of computer scientists are presenting next week.

The scientists to remotely control functions, was running security in an adversarial way, "We demonstrated an automated adversary including disabling the brakes, selectively the engine, and so on," they wrote in the Modern Automobile."

In the paper, which will be presented at Oakland, Calif., computer security specialists at the University of California, San Diego, reported engineering in the design of their computers to the potential threat of hackers who might

## Bloomberg Businessweek Markets & Finance

### Software Bug Made Swedish Exchange Go Bork, Bork, Bork

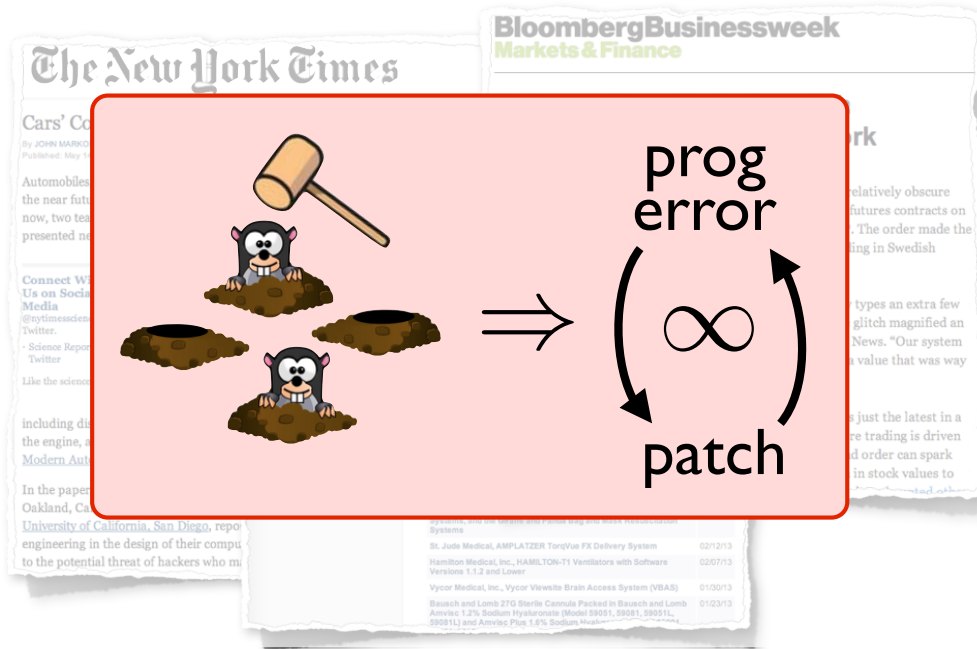
By Karen Weise on November 29, 2012

A computer error stalks the markets—again. An order on a relatively obscure derivatives index in Stockholm yesterday was asking to buy futures contracts on Swedish stocks valued at **131 times the country's entire GDP**. The order made the exchange go "bananas" and caused Nasdaq OMX to stop trading in Swedish derivatives for four hours.


This was no "fat finger" incident, where a trader accidentally types an extra few digits or the wrong numbers in an order. Instead, a software glitch magnified an order, Nasdaq OMX spokesman Carl Norell told Bloomberg News. "Our system misinterpreted a certain order category and communicated a value that was way too high into the book," he said.

The interruption was in a small corner of the market, but it's just the latest in a string of technical problems that have halted trading. As more trading is driven by the algorithms of high-frequency traders, one glitch or bad order can spark major disruptions. The 2010 flash crash caused \$862 billion in stock values to vanish from the market temporarily, and technical problems have


# Software Infrastructure is Shaky



*When exhaustive testing is impossible, our trust can only be based on proof.*



**Edsger W. Dijkstra**  
Under the Spell of Leibniz's Dream



**... not just a dream!**

## Proof Assistant Based Verification

Code in language suited for reasoning

Develop correctness proof in synch

Fully formal, *machine checkable* proof

## Proof Assistant Based Verification

Verified Compiler: **CompCert** [Leroy POPL 06]

Compiler	Bugs Found
GCC	122
LLVM	181
CompCert	?

[Yang et al. PLDI 11]

# Proof Assistant Based Verification

Verified Compiler: **CompCert** [Leroy POPL 06]

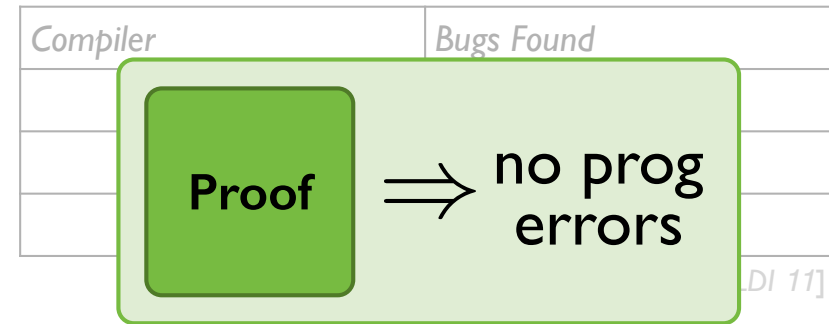
Compiler	Bugs Found
GCC	122
LLVM	181
CompCert	0

[Yang et al. PLDI 11]  
[Vu et al. PLDI 14]

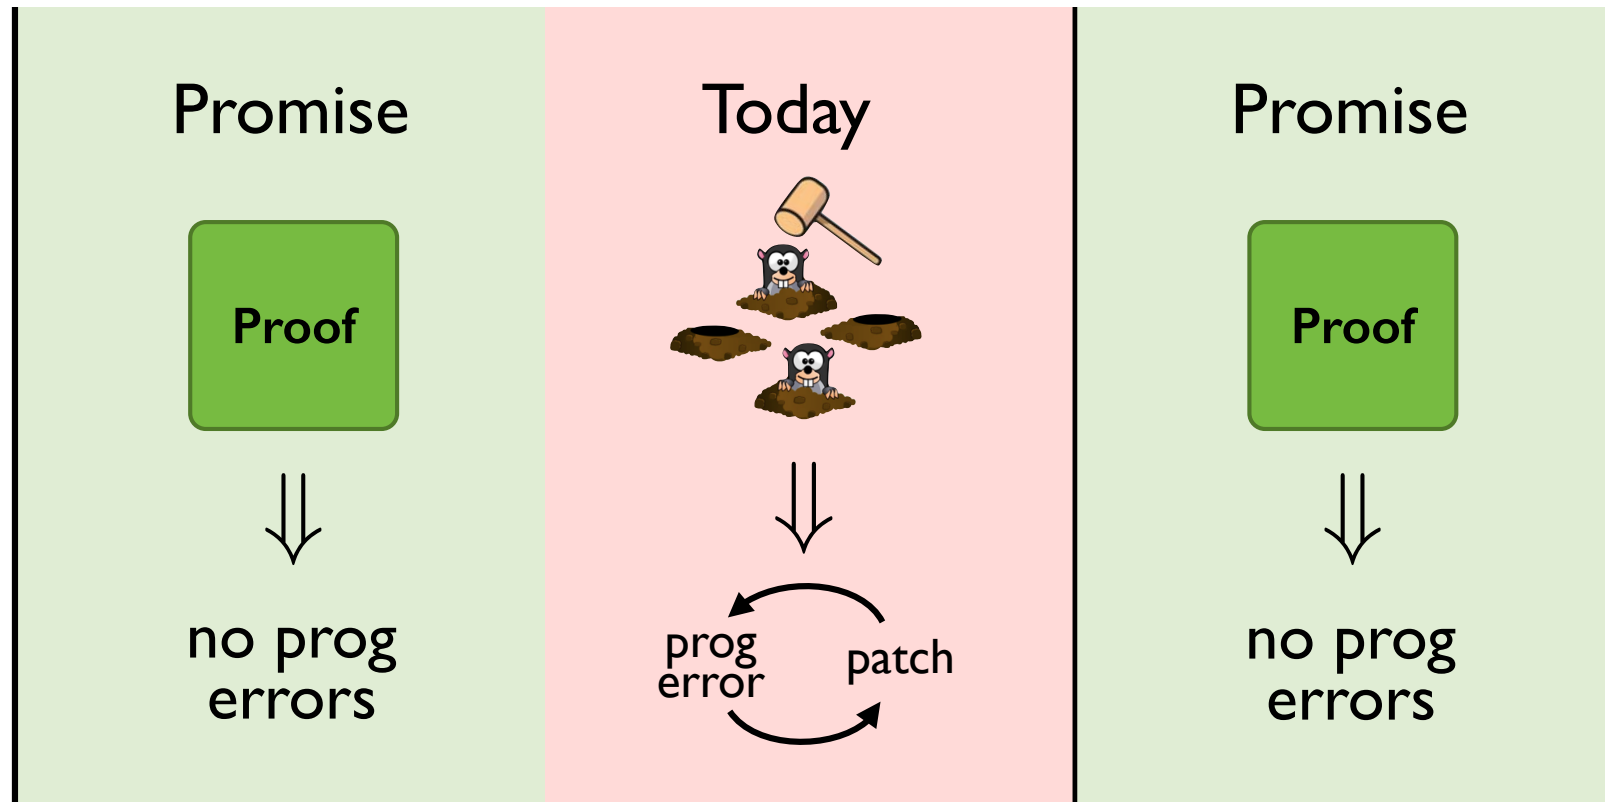
Verified OS kernel: **seL4** [Klein et al. SOSP 09]  
*realistic implementation guaranteed bug free*

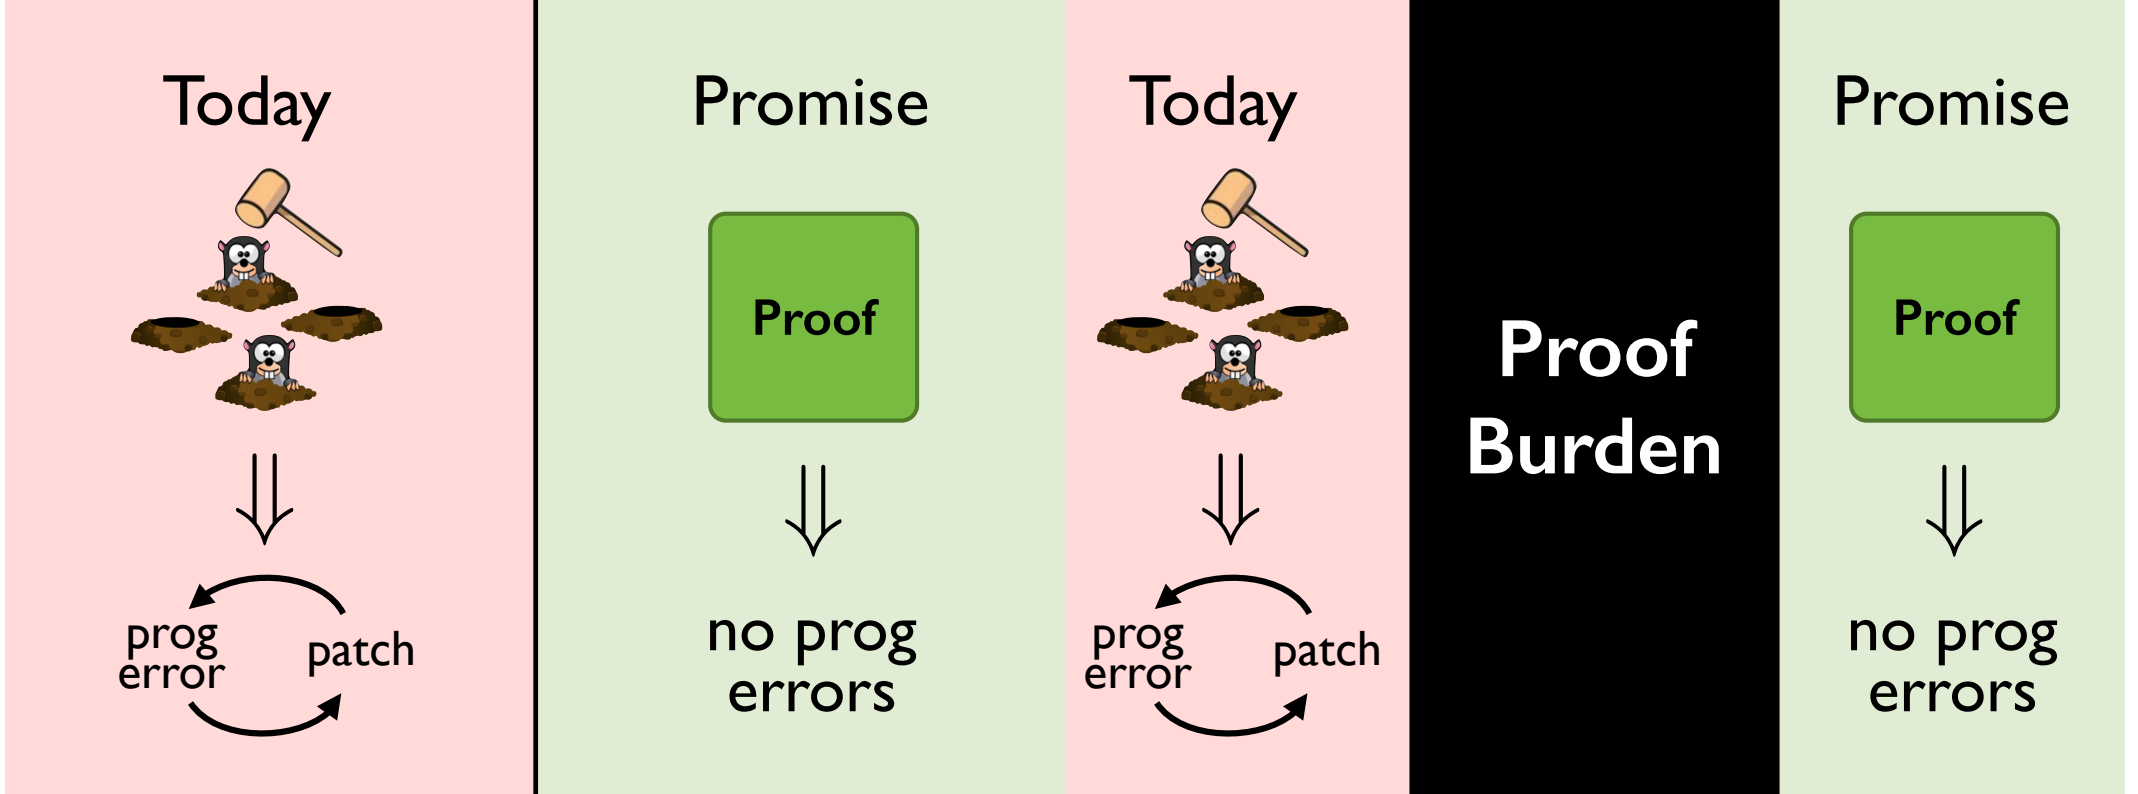
# Proof Assistant Based Verification

Verified Compiler: **CompCert** [Leroy POPL 06]



Verified OS kernel: **seL4** [Klein et al. SOSP 09]  
*realistic implementation guaranteed bug free*





## The Burden of Proof

### 1. Initial proofs require heroic effort

*CompCert: 70% proof, vast majority of effort*

*seL4: 200,000 line proof for 9,000 lines of C*

### 2. Code updates require re-proving

*CompCert: adding opts [Tristan POPL 08, PLDI 09, POPL 10]*

*seL4: changing RPC took 17% of proof effort*

## Mitigating the Burden of Proof

### 1: Scaling proofs to critical infrastructure

➡ *Formal shim verification for large apps*


*QUARK: browser with security guarantees*

### 2: Evolving formally verified systems

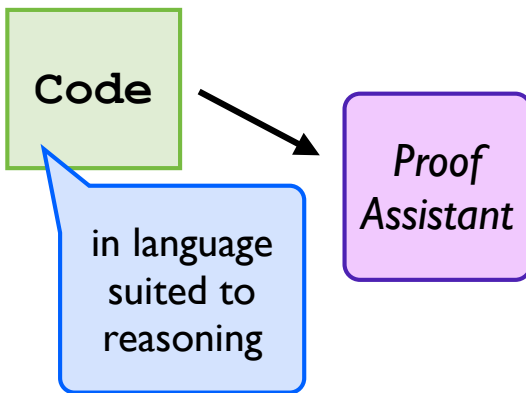
*Reflex DSL exploits domain for proof auto*

# Fully Formal Verification

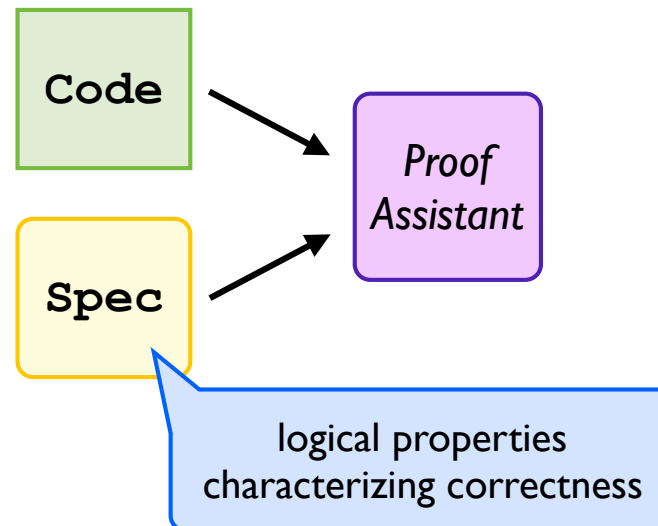
# Fully Formal Verification



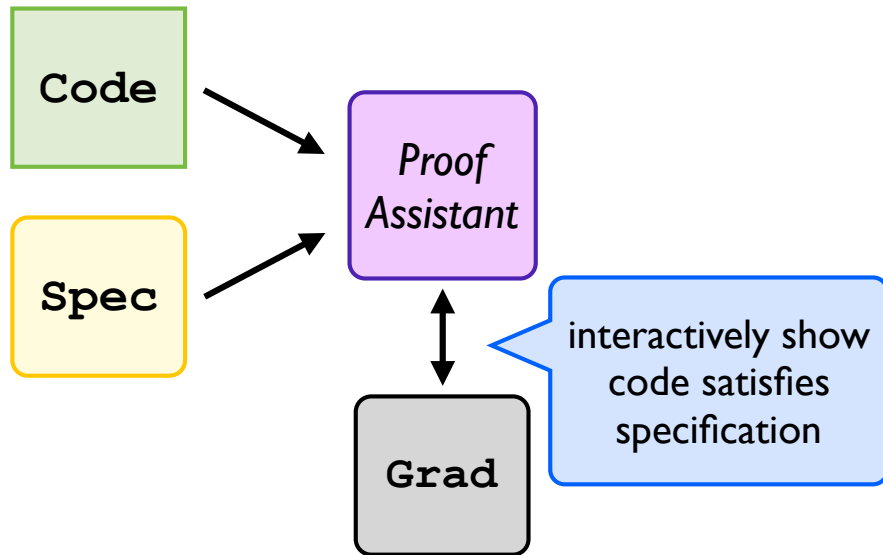
# Fully Formal Verification



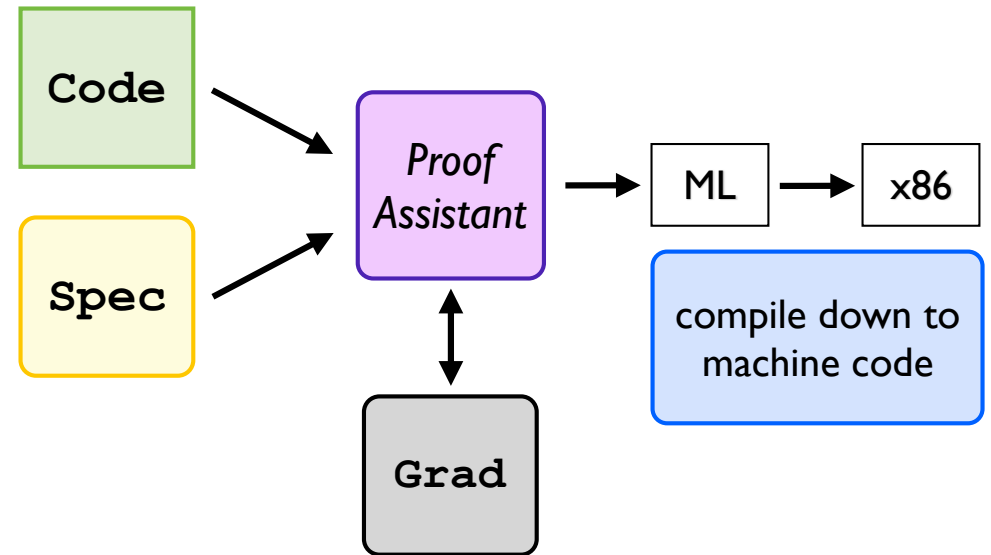
# Fully Formal Verification



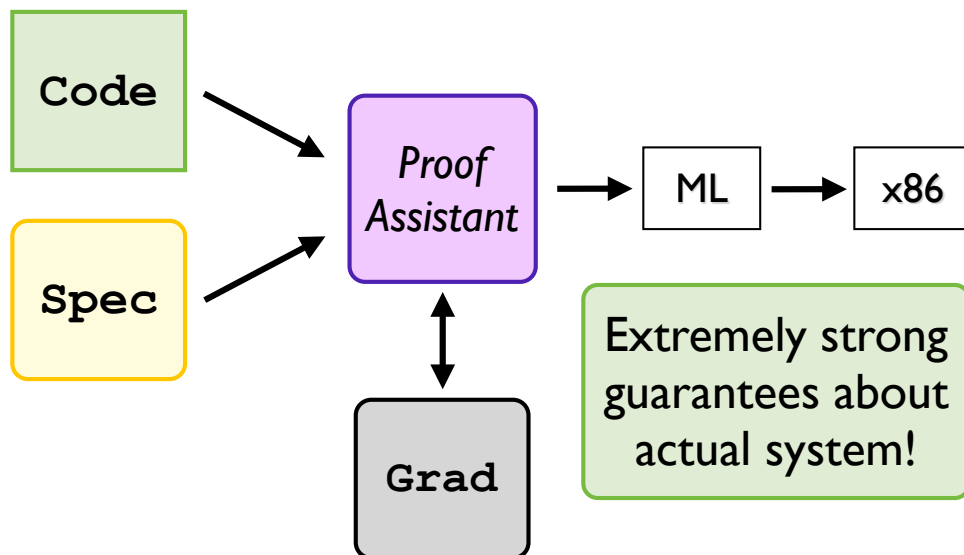
# Fully Formal Verification




# Fully Formal Verification



# Fully Formal Verification



# Fully Formal Verification



# Fully Formal Verification

```
emacs@gear-ubuntu-32
File Edit Options Buffers Tools Coq Proof-General

Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
  end.
```

program in a purely functional language

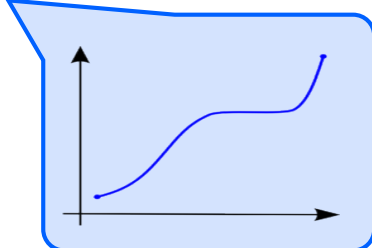
# Fully Formal Verification

```
emacs@gear-ubuntu-32
File Edit Options Buffers Tools Coq Proof-General Holes Help

Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
  end.

Definition monotonic f :=
  forall a b,
  a <= b ->
  f a <= f b.
```

specification characterizes desired behavior



# Fully Formal Verification

```
emacs@gear-ubuntu-32
File Edit Options Buffers Tools Coq Proof-General Holes Help

Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
  end.

Definition monotonic f :=
  forall a b,
  a <= b ->
  f a <= f b.

Theorem example :
  monotonic factorial.
Proof.
  ...
```

claim program satisfies spec

construct proof interactively

# Fully Formal Verification

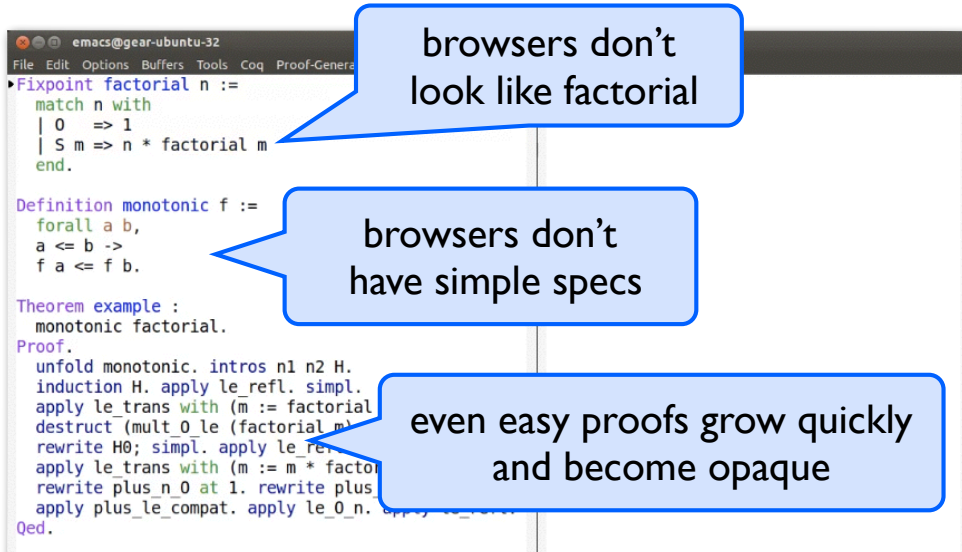
```
emacs@gear-ubuntu-32
File Edit Options Buffers Tools Coq Proof-General Holes Help

Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
  end.

Definition monotonic f :=
  forall a b,
  a <= b ->
  f a <= f b.

Theorem example :
  monotonic factorial.
Proof.
  unfold monotonic. intros n1 n2 H.
  induction H. apply le_refl. simpl.
  apply le_trans with (m := factorial m); auto.
  destruct (mult_0_le (factorial m) m).
  rewrite H0; simpl. apply le_refl.
  apply le_trans with (m := m * factorial m); auto.
  rewrite plus_n_0 at 1. rewrite plus_comm.
  apply plus_le_compat. apply le_0_n. apply le_refl.
Qed.
```

# Fully Formal Verification



browsers don't look like factorial

browsers don't have simple specs

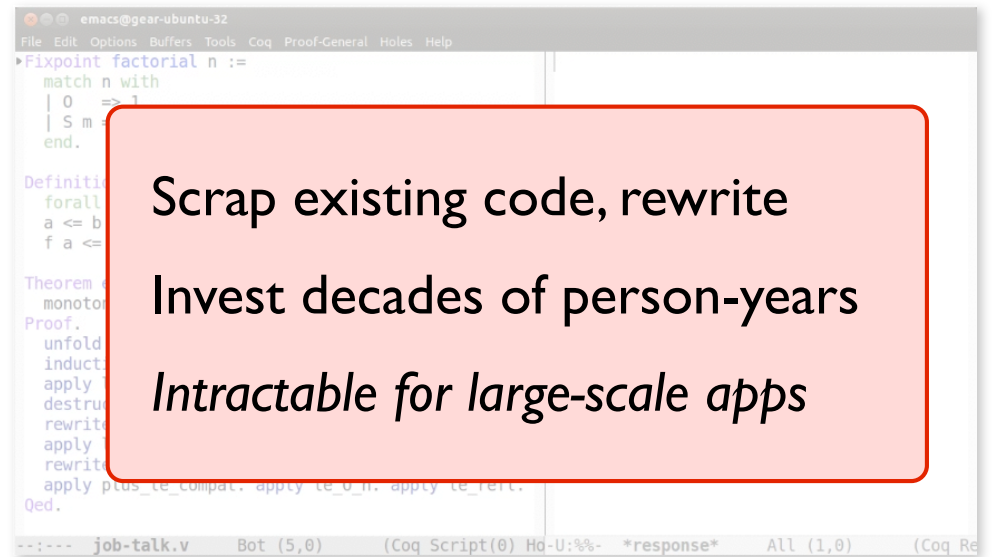
even easy proofs grow quickly and become opaque

```
Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
end.

Definition monotonic f :=
  forall a b,
  a <= b ->
  f a <= f b.

Theorem example :
  monotonic factorial.
Proof.
  unfold monotonic. intros n1 n2 H.
  induction H. apply le_refl. simpl.
  apply le_trans with (m := factorial
  destruct (mult_0 le (factorial m)
  rewrite H0; simpl. apply le_refl.
  apply le_trans with (m := m * factor
  rewrite plus_n_0 at 1. rewrite plus
  apply plus_le_compat. apply le_0_n.
Qed.
```

# Fully Formal Verification



Scrap existing code, rewrite

Invest decades of person-years

*Intractable for large-scale apps*

```
Fixpoint factorial n :=
  match n with
  | 0 => 1
  | S m => n * factorial m
end.

Definition monotonic f :=
  forall a b,
  a <= b ->
  f a <= f b.

Theorem example :
  monotonic factorial.
Proof.
  unfold monotonic. intros n1 n2 H.
  induction H. apply le_refl. simpl.
  apply le_trans with (m := factorial
  destruct (mult_0 le (factorial m)
  rewrite H0; simpl. apply le_refl.
  apply le_trans with (m := m * factor
  rewrite plus_n_0 at 1. rewrite plus
  apply plus_le_compat. apply le_0_n.
Qed.
```

Formally Verify a Browser?!

Formally Verify a Browser?!

Millions of LOC


Web  
Browser



# Formally Verify a Browser?!

Millions of LOC

High performance




# Formally Verify a Browser?!

Millions of LOC

High performance

Loose access policy




# Formally Verify a Browser?!

Millions of LOC

High performance

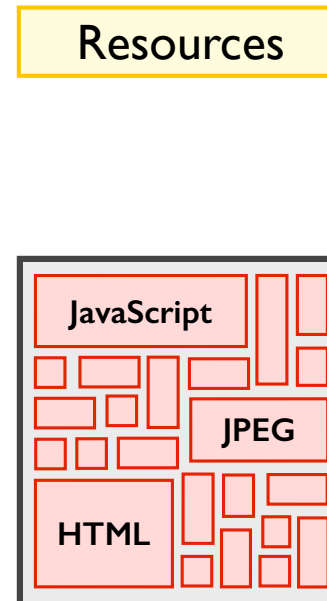
Loose access policy

Constant evolution




# Formally Verify a Browser?!

Isolate  
*sandbox untrusted code*




# Formally Verify a Browser?!



Isolate  
*sandbox untrusted code*

Implement shim  
*guards resource access*

# Formally Verify a Browser?!




Isolate  
*sandbox untrusted code*

Implement shim  
*guards resource access*

Verify shim  
*prove security policy*

## Formal Shim Verification

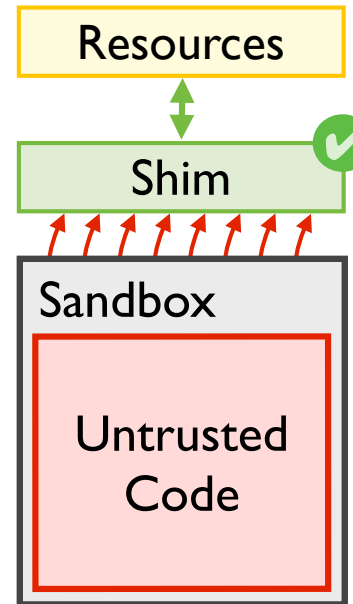


Isolate  
*sandbox untrusted code*

Implement shim  
*guards resource access*

Verify shim  
*prove security policy*

## Formal Shim Verification



Isolate  
Implement shim  
Verify shim

Applies when:

1. *sys fits architecture*
  2. *policy over resources*
- browser, httpd, sshd, ...*

# Formal Shim Verification

## Key Insight: *Focus Effort*

Guarantee sec props for entire system

Only implement and prove small shim

Radically ease verification burden

Prove *actual code* correct

# Mitigating the Burden of Proof

1: Scaling proofs to critical infrastructure

*Formal shim verification for large apps*

➔ *QUARK: browser with security guarantees*

2: Evolving formally verified systems

*Reflex DSL exploits domain for proof auto*

# Mitigating the Burden of Proof

1: Scaling proofs to critical infrastructure

➔ *Formal shim verification for large apps*

*QUARK: browser with security guarantees*

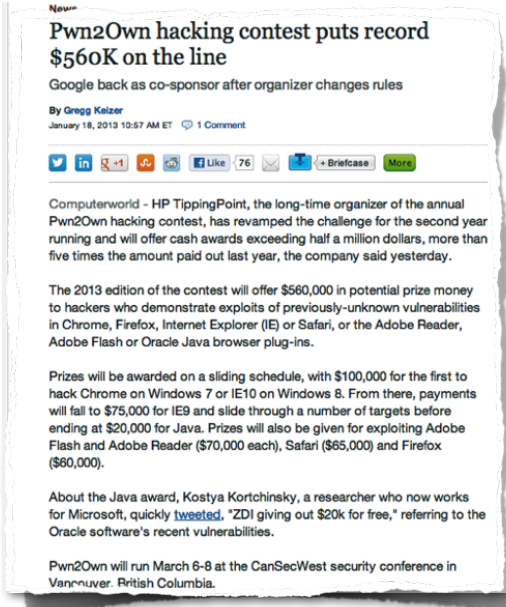
2: Evolving formally verified systems

*Reflex DSL exploits domain for proof auto*

## Browsers: Critical Infrastructure



# Browsers: Vulnerable




## Defenses / Policies:


- [Jang et al. W2SP]
- [Stamm et al. WWW]
- [Jackson et al. W2SP]
- [Barth et al. CCS]
- [Singh et al. OAKLAND]
- ...

Complex +  
Implementation Bugs

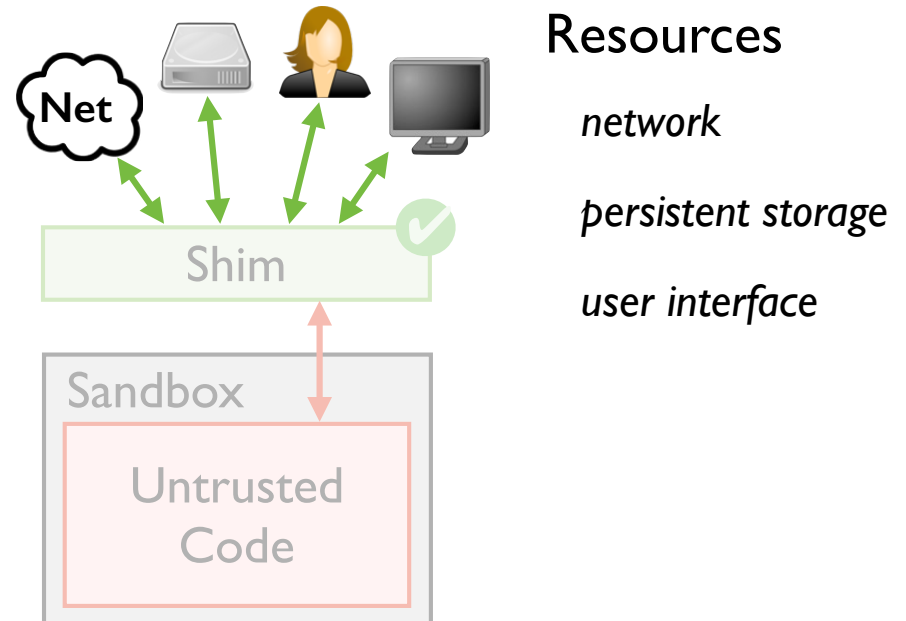
# Quark: Verified Browser



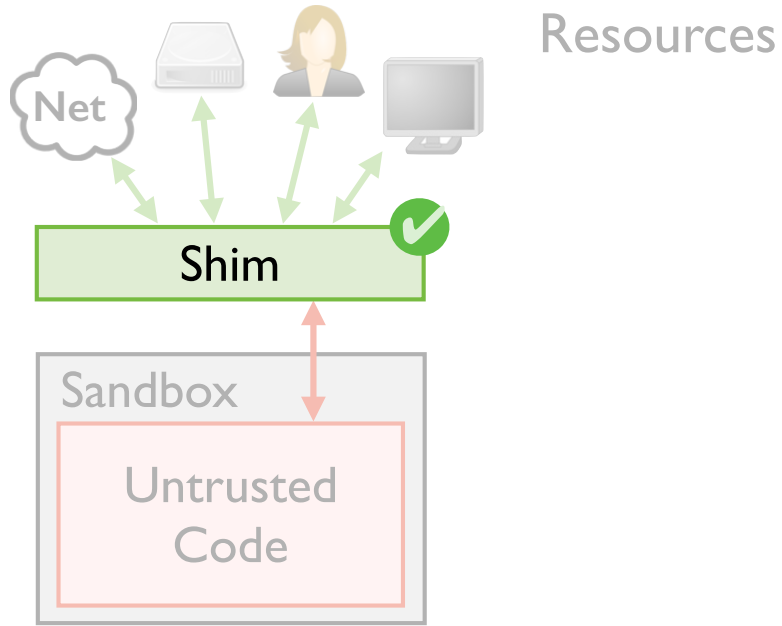
# Quark: Verified Browser



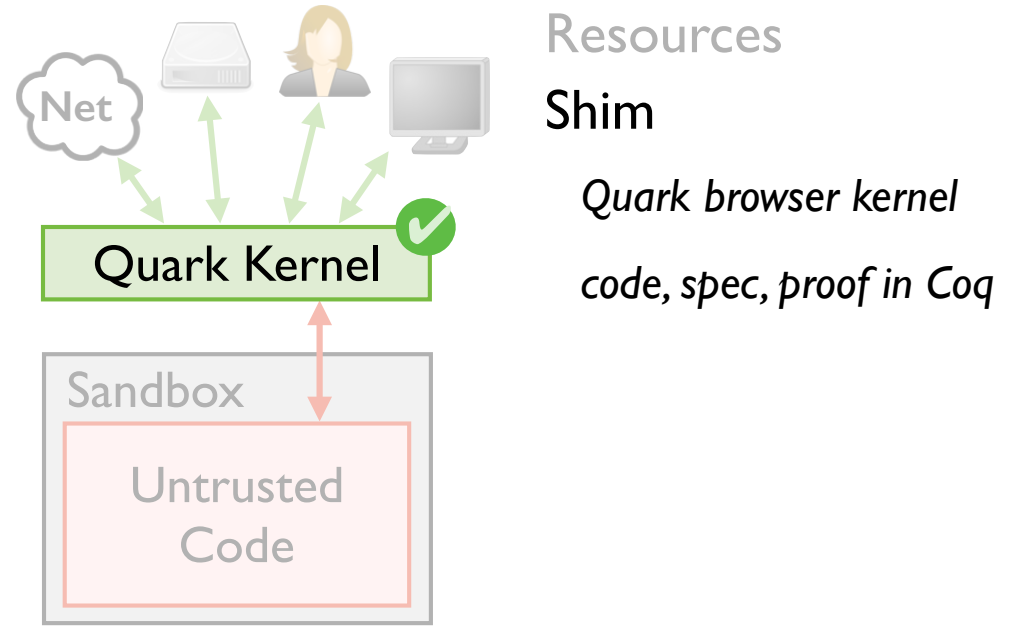
# Quark: Verified Browser



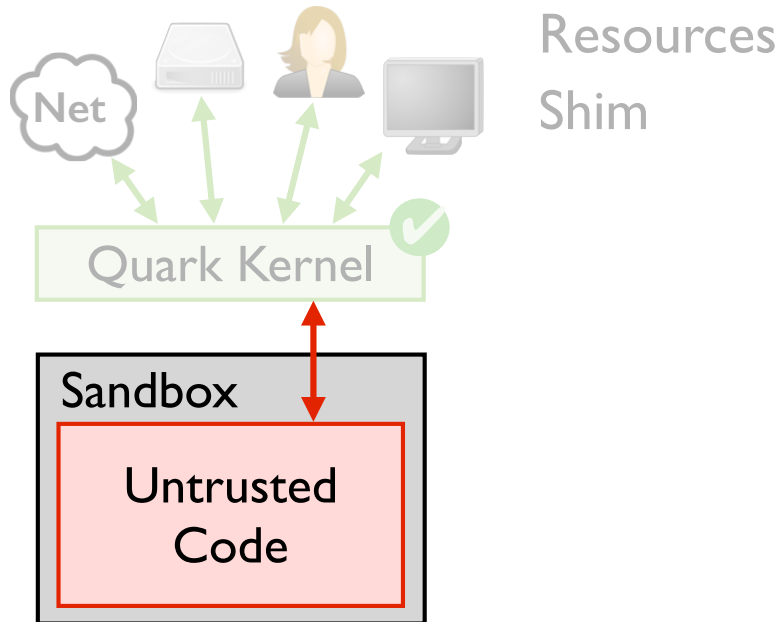
# Quark: Verified Browser



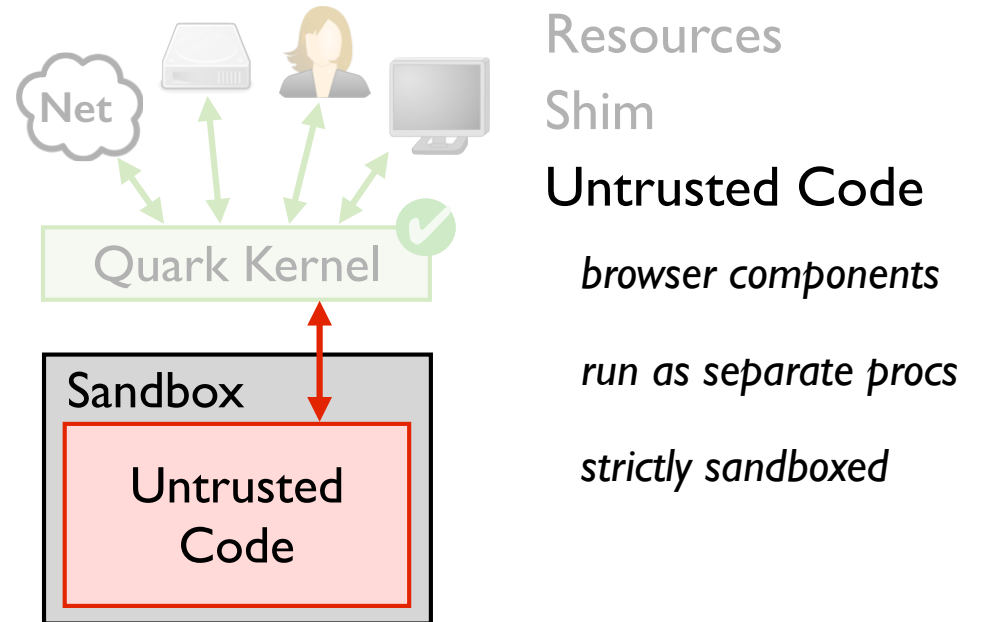
# Quark: Verified Browser




# Quark: Verified Browser



# Quark: Verified Browser

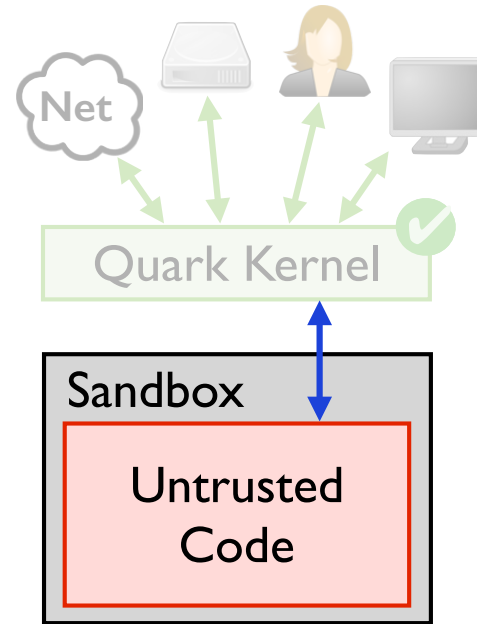


# Quark:Verified Browser



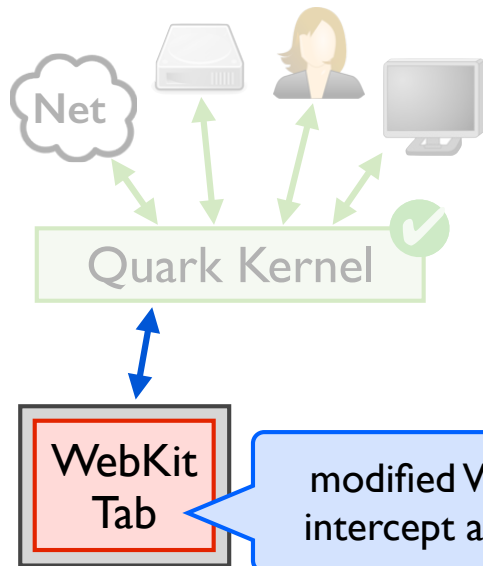
Resources  
Shim  
Untrusted Code  
*browser components*  
*run as separate procs*  
*strictly sandboxed*  
*talk to kernel over pipe*

# Quark:Verified Browser



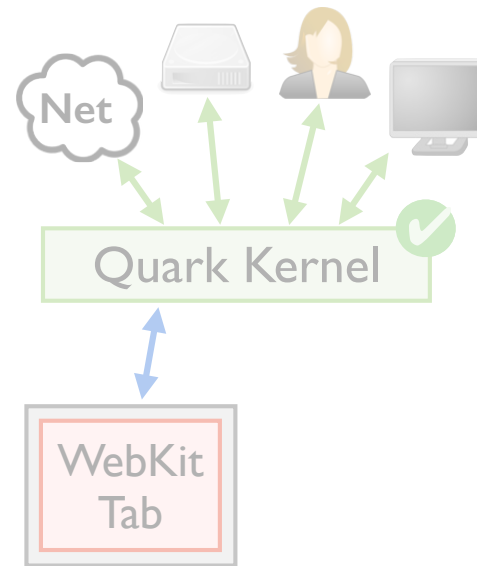
Resources  
Shim  
Untrusted Code  
*two component types*

# Quark:Verified Browser



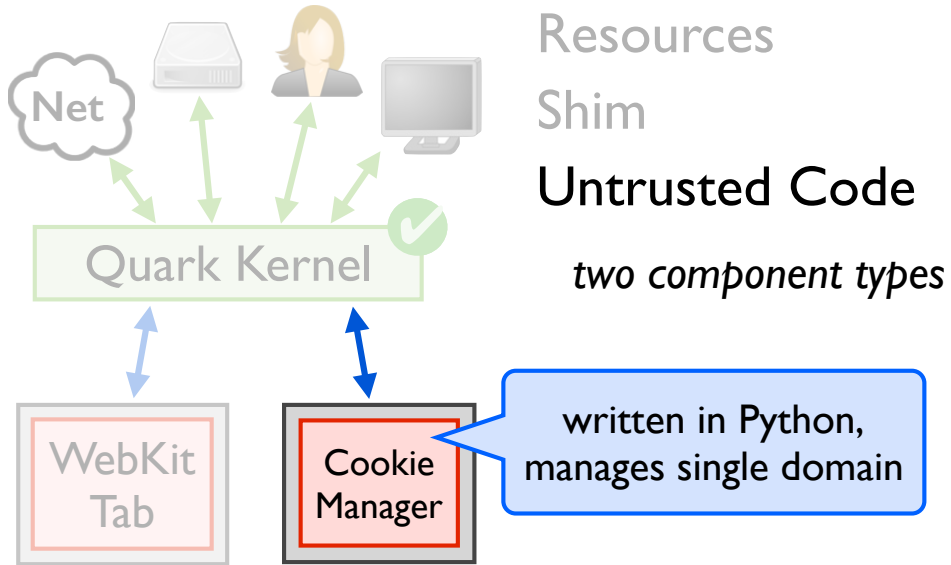
Resources  
Shim  
Untrusted Code  
*two component types*

# Quark:Verified Browser

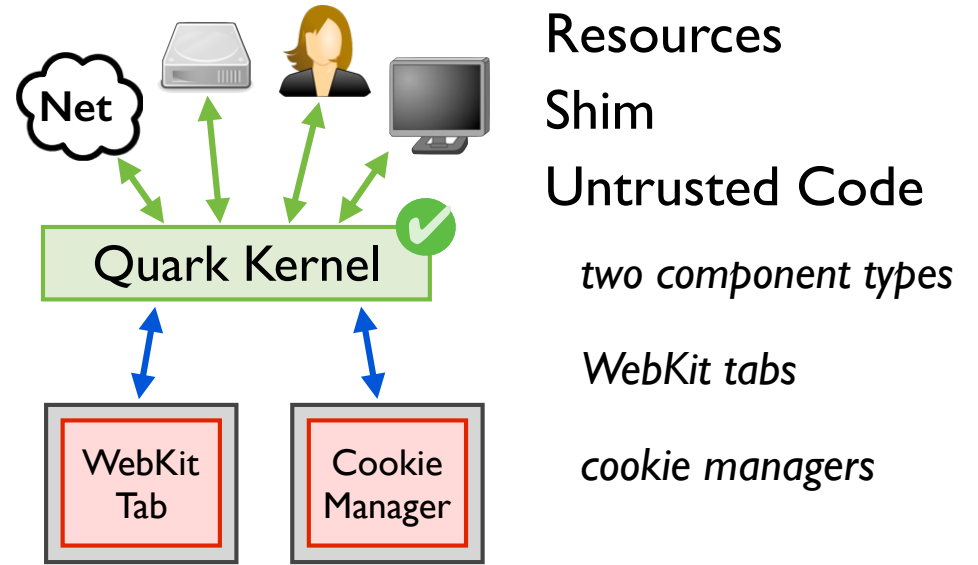


Resources  
Shim  
Untrusted Code  
*two component types*

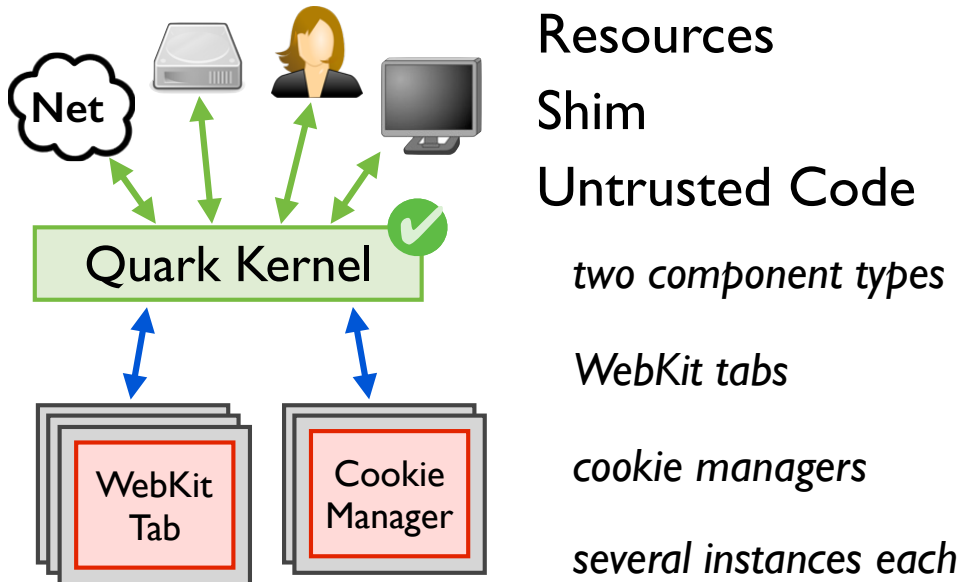
# Quark:Verified Browser



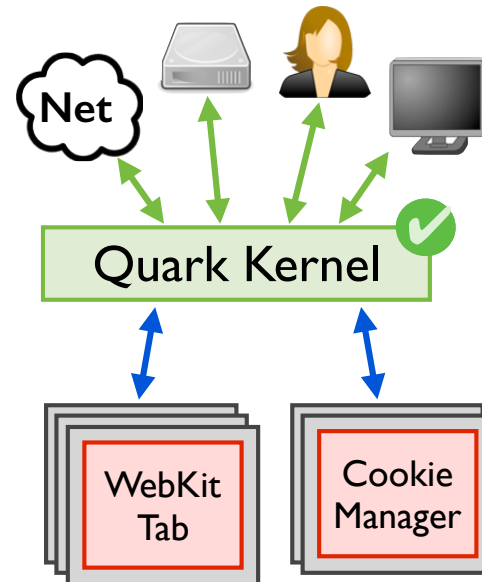
# Quark:Verified Browser



# Quark:Verified Browser




# Quark:Verified Browser



# Quark: Verified Browser


# Quark Kernel


Quark Kernel 

Quark Kernel 

Quark Kernel: Code, Spec, Proof

Quark Kernel: *Code*, Spec, Proof

Quark Kernel 

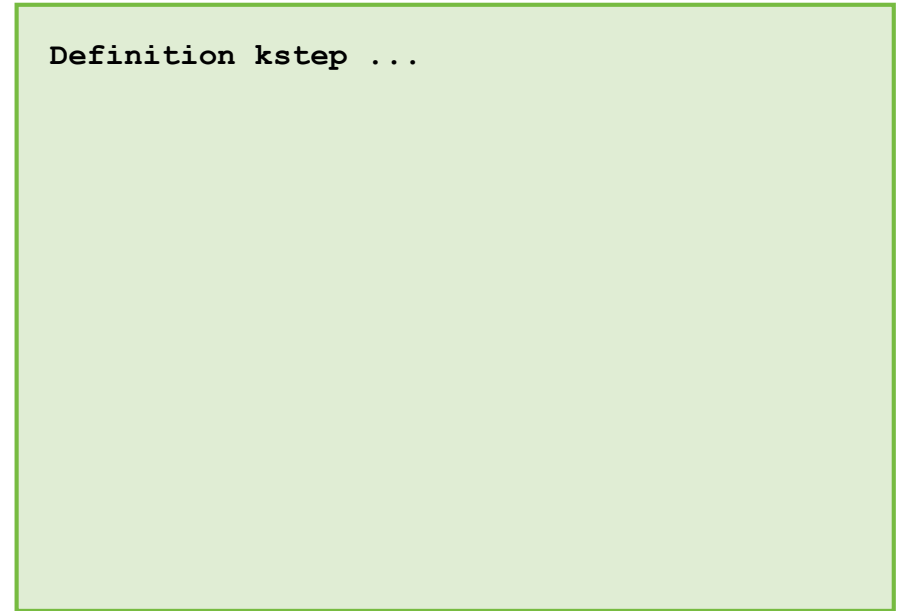
Quark Kernel 



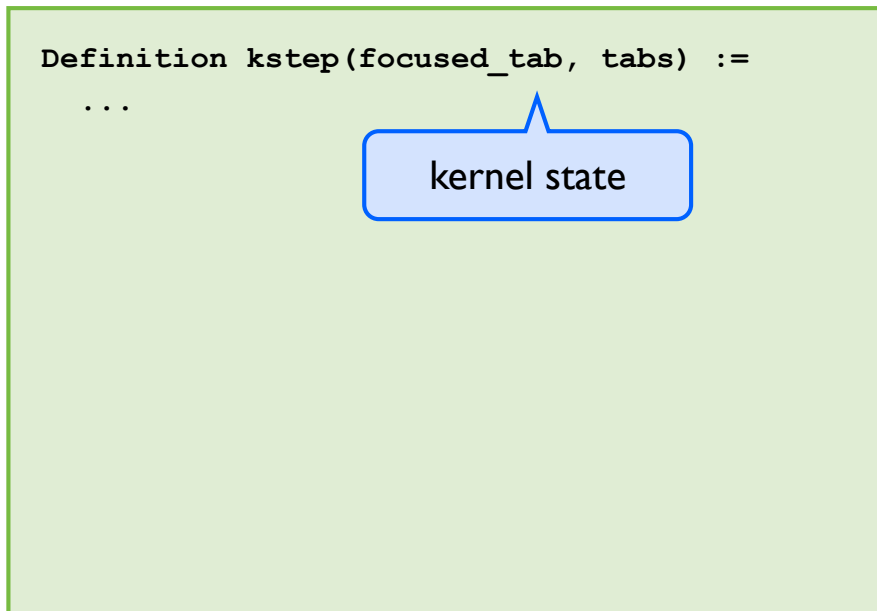
## Quark Kernel: *Code*, Spec, Proof



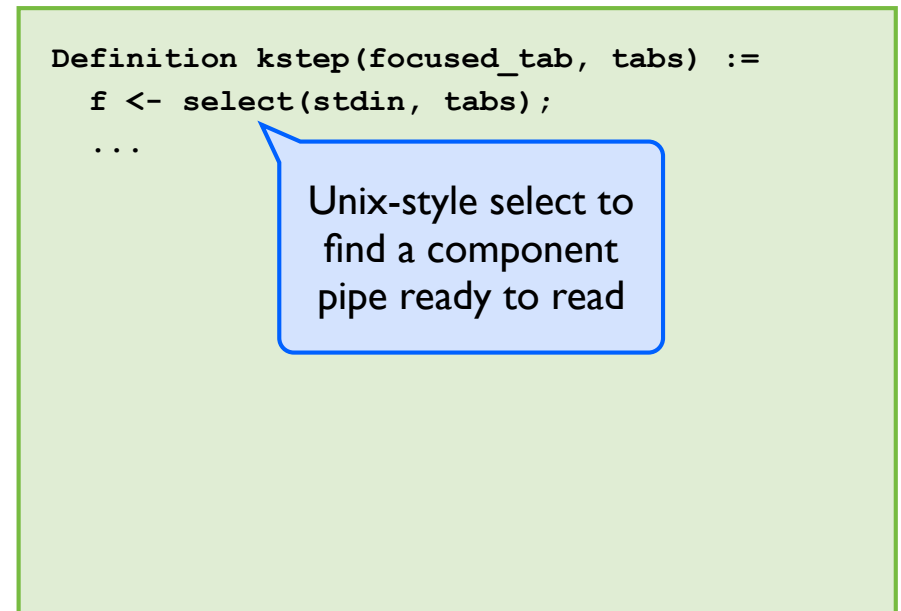
## Quark Kernel: *Code*, Spec, Proof



## Quark Kernel: *Code*, Spec, Proof



## Quark Kernel: *Code*, Spec, Proof



## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    ...  
  | Tab t =>  
    ...
```

case:  $f$  is user input

case:  $f$  is tab pipe

## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    ...  
  | Tab t =>  
    ...
```

read command from  
user over `stdin`

## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      ...  
    | ...  
  | Tab t =>  
    ...
```

user wants to create  
and focus a new tab

## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      ...  
    | ...  
  | Tab t =>  
    ...
```

create a new tab

# Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      ...  
    | ...  
  | Tab t =>  
    ...
```

tell new tab to  
render itself

# Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      return (t, t::tabs)  
    | ...  
  | Tab t =>  
    ...
```

return updated state

# Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      return (t, t::tabs)  
    | ...  
  | Tab t =>  
    ...
```

# Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      return (t, t::tabs)  
    | ...  
  | Tab t =>  
    ...
```

handle other  
user commands

## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      return (t, tabs);  
    | ...  
  | Tab t =>  
    ...
```

handle requests  
from tabs

## Quark Kernel: *Code*, Spec, Proof

```
Definition kstep(focused_tab, tabs) :=  
  f <- select(stdin, tabs);  
  match f with  
  | Stdin =>  
    cmd <- read_cmd(stdin);  
    match cmd with  
    | AddTab =>  
      t <- mk_tab();  
      write_msg(t, Render);  
      return (t, t::tabs);  
    | ...  
  | Tab t =>  
    ...
```

## Quark Kernel: *Code*, Spec, Proof

## Quark Kernel: Code, *Spec*, Proof


# Quark Kernel: Code, *Spec*, Proof

Safety properties to mitigate attacks

*restrict kernel behavior to only safe executions*

Example: mitigate phishing attacks

*prevent tricks that get users to divulge secrets*



# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs

`read(), write(), open(), write(), ...`


# Quark Kernel: Code, *Spec*, Proof

Safety properties to mitigate attacks

*restrict kernel behavior to only safe executions*

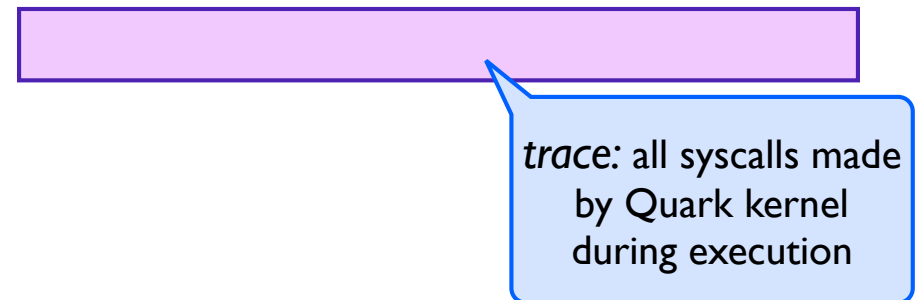
Example: mitigate phishing attacks

*prevent tricks that get users to divulge secrets*




# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs




# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs




# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



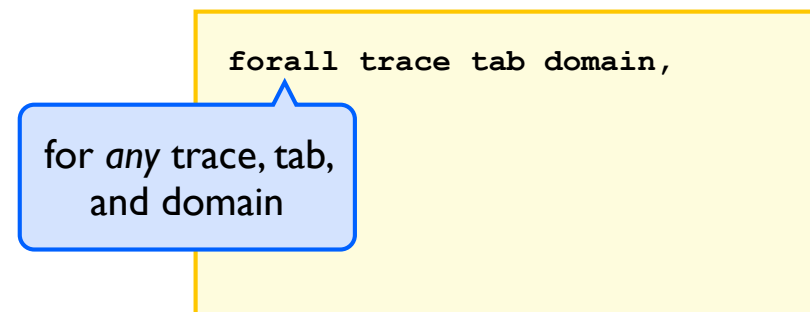
Example: address bar correctness

# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



Example: address bar correctness



# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



*structure of produceable traces supports spec & proof*

Example: address bar correctness

```
forall trace tab domain,  
  quark_produced(trace)  ^  
  ...
```

*if Quark could have produced this trace*

# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



*structure of produceable traces supports spec & proof*

Example: address bar correctness

```
forall trace tab domain,  
  quark_produced(trace)  ^  
  tab = cur_tab(trace)  ^  
  ...
```

*and tab is the selected tab in this trace*

# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



*structure of produceable traces supports spec & proof*

Example: address bar correctness

```
forall trace tab domain,  
  quark_produced(trace)  ^  
  tab = cur_tab(trace)  ^  
  domain = addr_bar(trace) ->  
  ...
```

*and domain displayed in address bar for this trace*

# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



*structure of produceable traces supports spec & proof*

Example: address bar correctness

```
forall trace tab domain,  
  quark_produced(trace)  ^  
  tab = cur_tab(trace)  ^  
  domain = addr_bar(trace) ->  
  domain = tab_domain(tab)
```

*then domain is the domain of the focused tab*

# Quark Kernel: Code, *Spec*, Proof

Specify correct behavior wrt syscall seqs



*structure of produceable traces supports spec & proof*

**Example: address bar correctness**

```
forall trace tab domain,  
  quark_produced(trace)  ^  
  tab = cur_tab(trace)   ^  
  domain = addr_bar(trace) ->  
  domain = tab_domain(tab)
```

# Quark Kernel: Code, *Spec*, Proof

# Quark Kernel: Code, *Spec*, Proof

## Formal Security Properties

### Tab Non-Interference

*no tab affects kernel interaction with another tab*

### Cookie Confidentiality and Integrity

*cookies only accessed by tabs of same domain*

### Address Bar Integrity and Correctness

*address bar accurate, only modified by user action*

# Quark Kernel: Code, Spec, *Proof*



# Quark Kernel: Code, Spec, *Proof*

Prove kernel code satisfies sec props  
by induction on traces Quark can produce

# Quark Kernel: Code, Spec, *Proof*

Prove kernel code satisfies sec props  
by induction on traces Quark can produce




induction hypothesis:  
trace valid up to this point

# Quark Kernel: Code, Spec, *Proof*

Prove kernel code satisfies sec props  
by induction on traces Quark can produce

# Quark Kernel: Code, Spec, *Proof*



induction hypothesis:  
trace valid up to this point

proof obligation:  
still valid after step?



induction hypothesis:  
trace valid up to this point

proof obligation:  
still valid after step?

Proceed by case analysis on `kstep()`

*what syscalls can be appended to trace?*

*will they still satisfy all security properties?*

*prove each case interactively in proof assistant*

# Quark Kernel: Code, Spec, *Proof*

Proving required diverse range of tools

*monads* encoding I/O in functional language

*Hoare logic* reasoning about imperative programs


*op. semantics* defining correctness of Quark kernel

*linear logic* proving resources created / destroyed

**YNot**

[Naneveski et al. ICFP 08]

## Formally Verified Browser!



# Quark Kernel: Code, Spec, Proof

## Key Insight: *FSV Effective*

Guarantee sec props for browser

Use state-of-the-art components

Only prove simple browser kernel

## Extending Quark

Filesystem access, sound, history

*could be implemented w/out major redesign*

Finer grained resource accesses

*support mashups and plugins*

Liveness properties

*no blocking, kernel eventually services all requests*

# Trusted Computing Base

Infrastructure we assume correct

*bugs here can invalidate our formal guarantees*

Fundamental

Statement of security properties  
Coq (soundness, proof checker)

Eventually  
Verified  
[active research]

OCaml [VeriML]  
Tab Sandbox [RockSalt]  
Operating System [seL4]  
...

# Quark Development Effort

150 lines of security props

900 lines of kernel code

4,500 lines of proofs

1,000,000 lines of WebKit

week

months

# Quark Development Effort

150 lines of security props

900 lines of kernel code

4,500 lines of proofs

1,000,000 lines of WebKit

## Mitigating the Burden of Proof

1: Scaling proofs to critical infrastructure

*Formal shim verification for large apps*

→ QUARK: browser with security guarantees

2: Evolving formally verified systems

*Reflex DSL exploits domain for proof auto*

# Mitigating the Burden of Proof

1: Scaling proofs to critical infrastructure

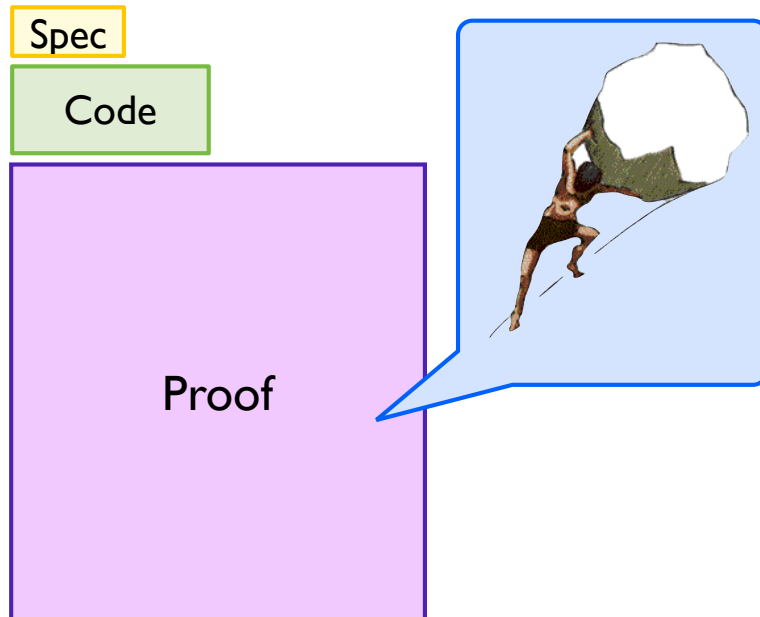
*Formal shim verification for large apps*

*QUARK: browser with security guarantees*

2: Evolving formally verified systems

➡ *Reflex DSL exploits domain for proof auto*

## Division of Labor *(to scale)*



# Struggle Against Formality Inertia

Adding cookies to Quark quite difficult

*all the pieces already there, still took over a month*


Proof updates repetitive and shallow

*sensitive proof scripts, changes not mechanical*

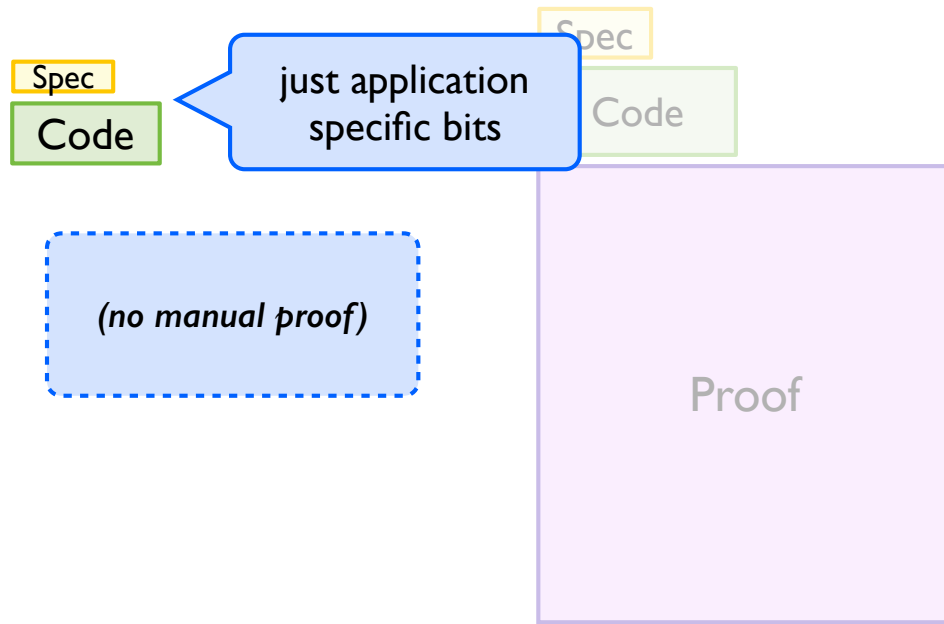
```
match svec_ith PAYREST i as _vi return
forall (EQ: (svec_ith (projT2 (existT vdesc' ENVD_SIZE PAYREST)) i) = _vi),
match _vi as __d return (base_term (existT vdesc' ENVD_SIZE PAYREST) __d -> Prop)
with
| Desc d => fun _ => True
| Comp c => fun b => FdSet.In
  (comp_fd (projT1 (eval_base_term (envd:=existT _ ENVD_SIZE PAYREST) crest b))) fds end
match EQ in _ = _vi return base_term _ _vi with Logic.eq_refl =>
  Var (existT vdesc' ENVD_SIZE PAYREST) i end
->
match _vi as __d return (base_term (existT vdesc' (S ENVD_SIZE) (PAYO, PAYREST)) __d -> Prop) with
| Desc d => fun _ => True
| Comp c => fun b =>
  FdSet.In (comp_fd (projT1 (eval_base_term (envd:=existT _ (S ENVD_SIZE) (PAYO, PAYREST)) (e0, crest) b))) fds end
match EQ in _ = _vi return base_term _ _vi with Logic.eq_refl =>
  Var (existT vdesc' (S ENVD_SIZE) (PAYO, PAYREST)) (Some i) end
with
| Desc d => _ | Comp c => _ end (Logic.eq_refl _)
```

## Division of Labor

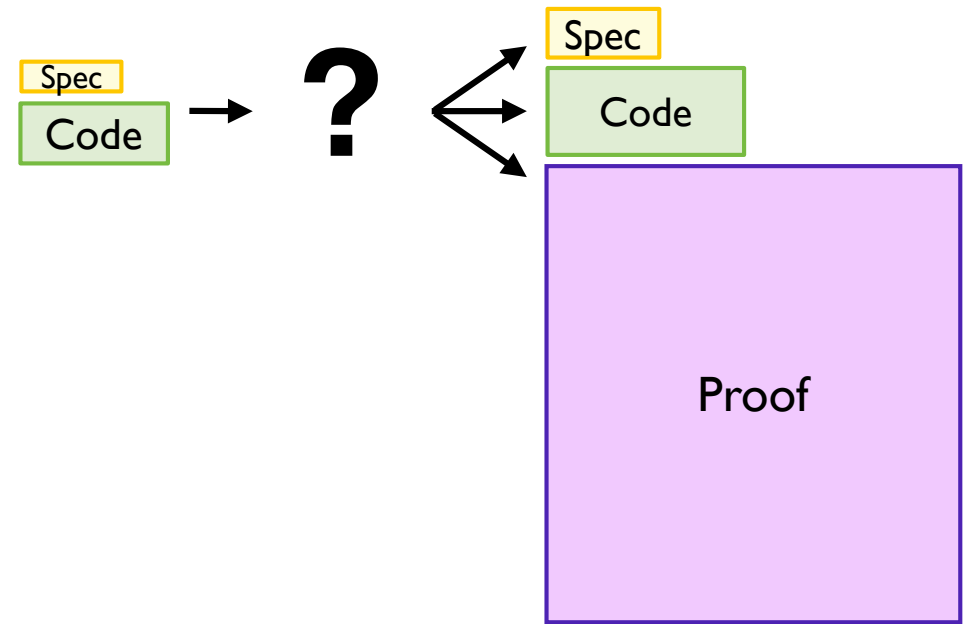
Ideal?



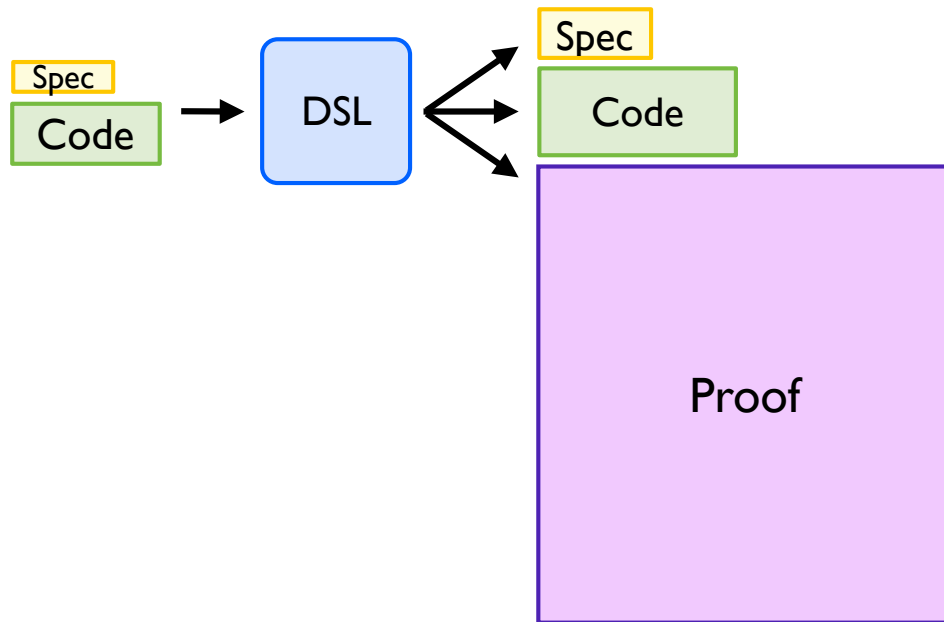
## Division of Labor



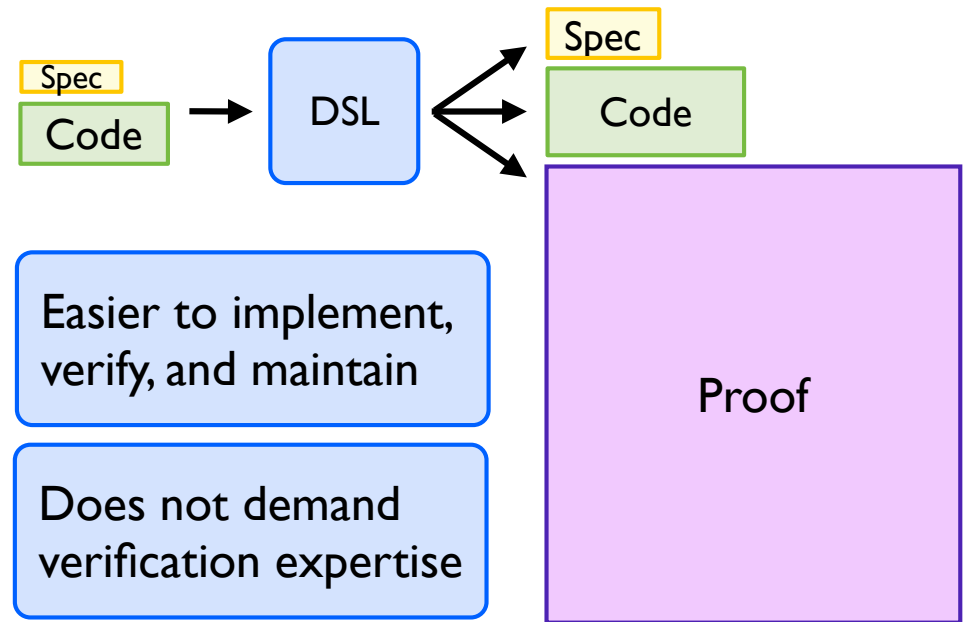
## Division of Labor



## Division of Labor



## Division of Labor

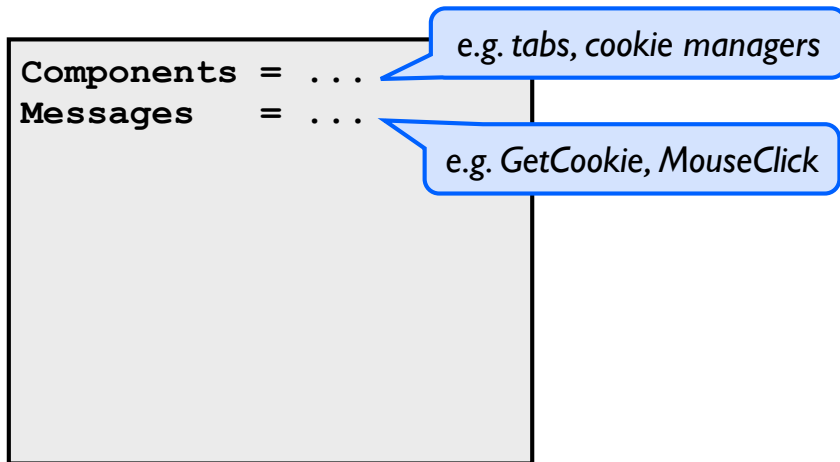


# Reflex: a DSL for Reactive Systems

[PLDI 14]

Exploit structure of app domain

*kernel based archs, well suited to FSV design*

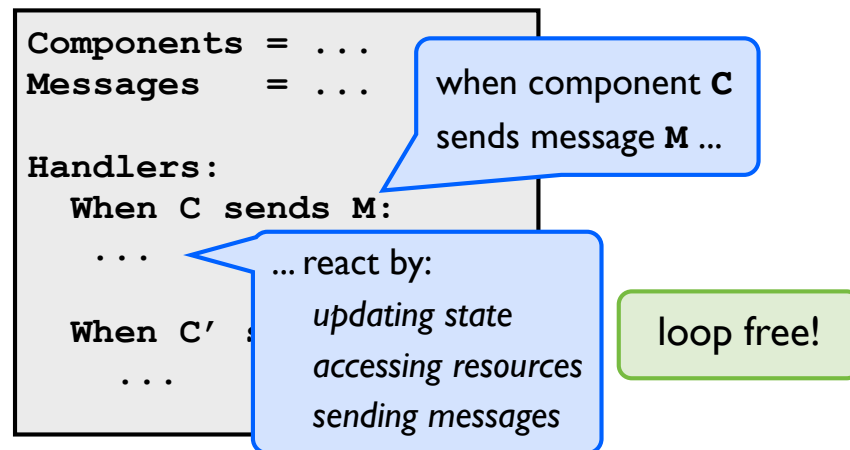


# Reflex: a DSL for Reactive Systems

[PLDI 14]

Exploit structure of app domain

*kernel based archs, well suited to FSV design*



# Reflex: a DSL for Reactive Systems


[PLDI 14]

Exploit structure of app domain

*kernel based archs, well suited to FSV design*

Provide expressive spec language

*subset of LTL and non-interference properties*



# Reflex: a DSL for Reactive Systems

[PLDI 14]

Exploit structure of app domain

*kernel based archs, well suited to FSV design*

Provide expressive spec language

*subset of LTL and non-interference properties*

Auto prove user-provided specs

*exploit domain, ensure all traces match spec*

Counterexample-driven search discovers invariants.

# Reflex: a DSL for Reactive Systems [PLDI 14]

**Reflex Effective:**

- Prototype sshd, browser, httpd
- Specify basic access controls
- Auto prove user-provided specs

# Reflex: Evaluation

<b>Web browser</b>	Domains do not interfere, Cookie integrity, ... <i>auto prove non-interference</i>
<b>SSH server</b>	No PTY access before authentication, At most 3 authentication attempts, ... <i>auto prove non-local props</i>
<b>Web server</b>	Clients only spawned after successful login, File requests guarded by access control, ...

Auto verified 33 properties (80% in < 2 minutes)

# Reflex: Development Effort

**Reflex :** *Many reactive systems*  
7500 lines of Coq

<b>Web browser</b>	<b>SSH server</b>	<b>Web server</b>
--------------------	-------------------	-------------------

**Quark Web browser :**  
5500 lines of Coq *Single reactive system*

# Mitigating the Burden of Proof


1: Scaling proofs to critical infrastructure

*Formal shim verification for large apps*

*QUARK: browser with security guarantees*

2: Evolving formally verified systems

➔ *Reflex DSL exploits domain for proof auto*




# Double Trouble

```
x = 0.1 + 0.2;
if (x != 0.3)
    printf("wat.\n");
```

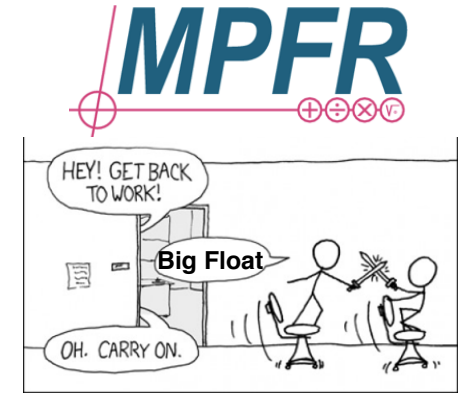
$$\frac{(-b) \pm \sqrt{b^2 - 4 \cdot (a \cdot c)}}{2 \cdot a}$$



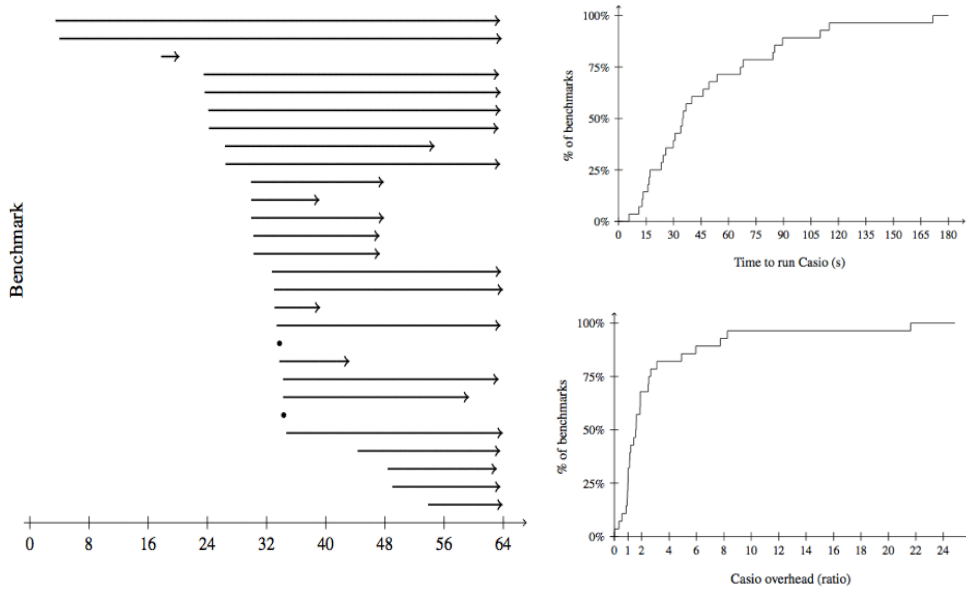
Futz



Analyze



# Less Double Trouble


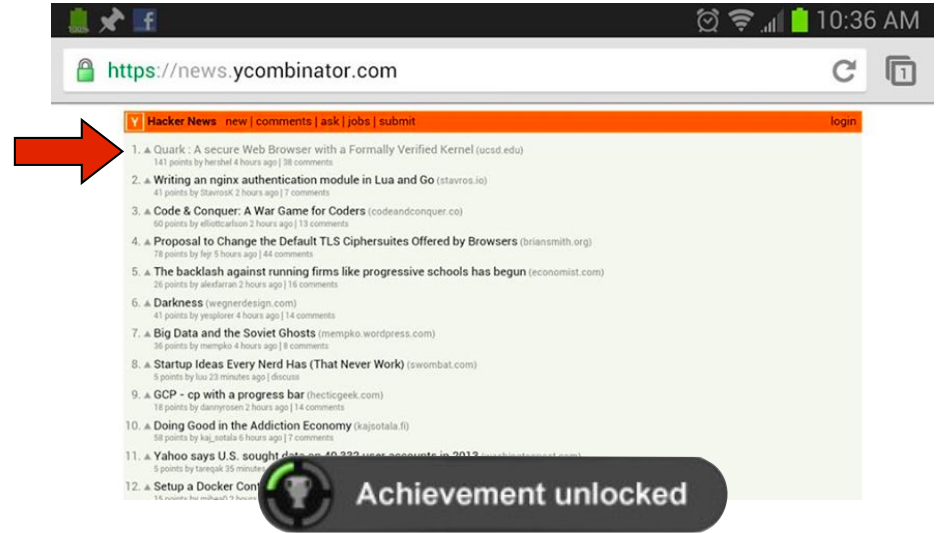



# Neutron Beams

UW Medicine  
SCHOOL OF MEDICINE





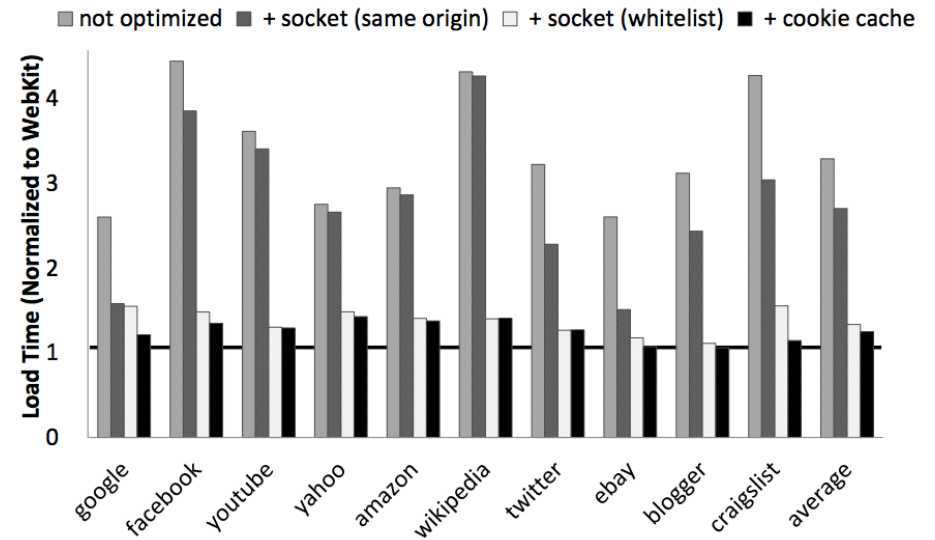


# Thank You!

Goal: mitigate formality inertia  
*address scaling and evolving formally verified systems*

1. Extend verification frontier  
*develop techniques to verify critical "pinch points"*

2. Make verification accessible  
*equip domain experts with effective tools*



## Verifying Optimizations

Rich compiler correctness history:

*McCarthy 67, Samet 75, Cousot 77, ...*

Already solved?

Compiler	Bugs Found
GCC	122
LLVM	181
CompCert	0


[Yang et al. PLDI 11]

many optimization bugs

lacks many optimizations

## Verifying Optimizations


# Verifying Optimizations



# Verifying Optimizations



Proof original and opt code equivalent.




# Verifying Optimizations



Proof original and opt code equivalent.

Construct *bisimulation relation*:





# Verifying Optimizations


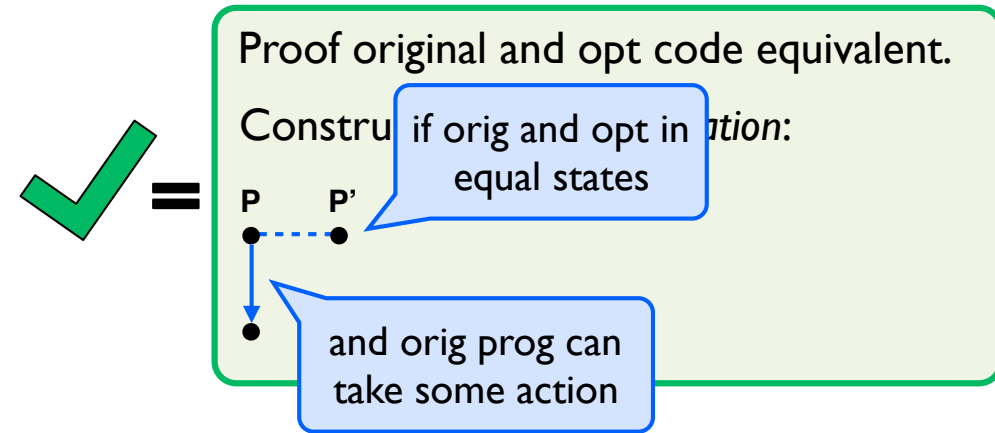


Proof original and opt code equivalent.


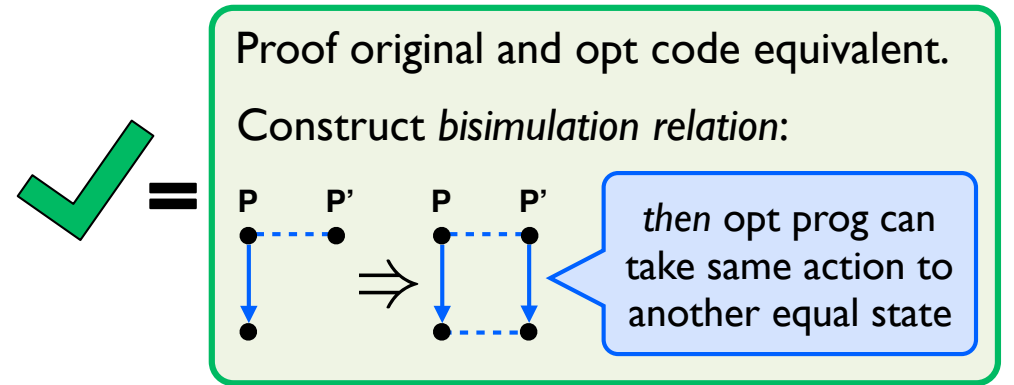
Construct if orig and opt in *tion*:




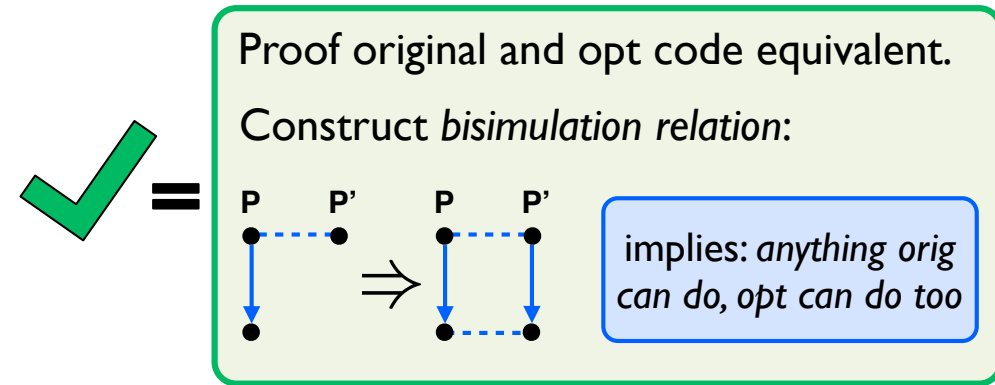
# Verifying Optimizations




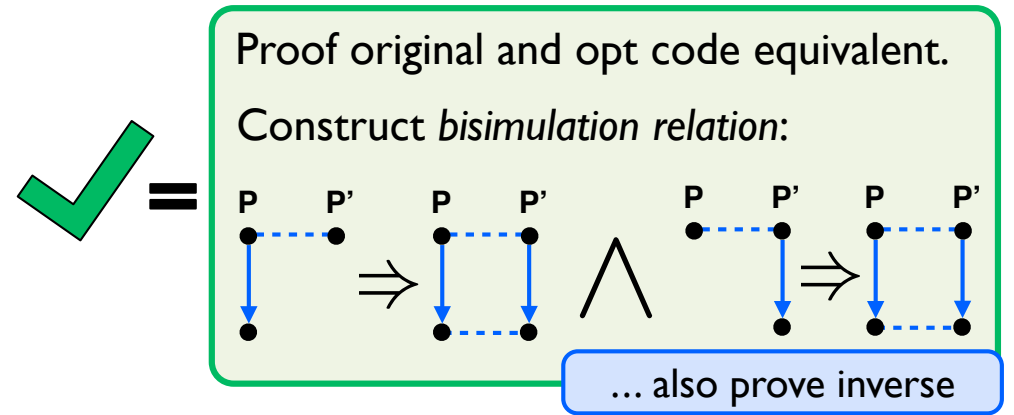
# Verifying Optimizations




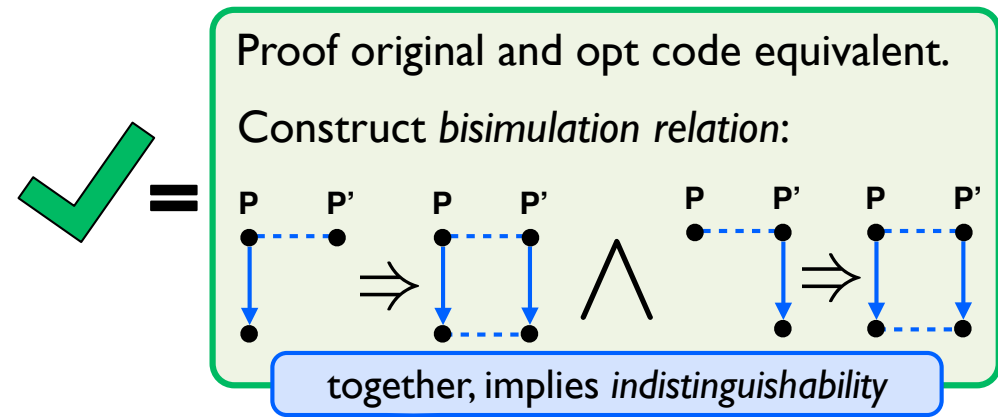
# Verifying Optimizations





# Verifying Optimizations



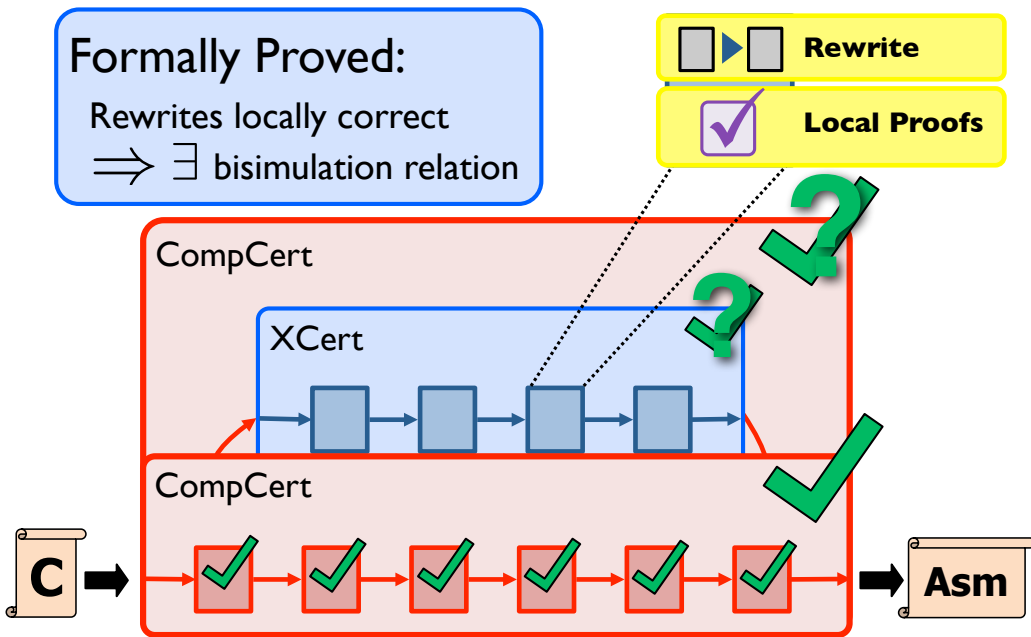
# Verifying Optimizations



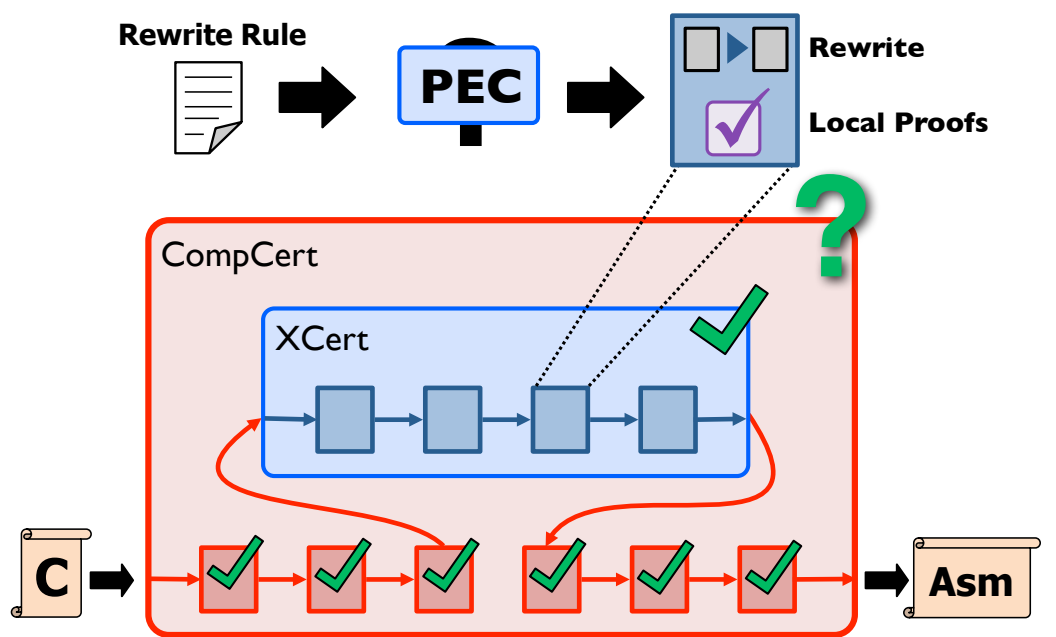
# Verifying Optimizations



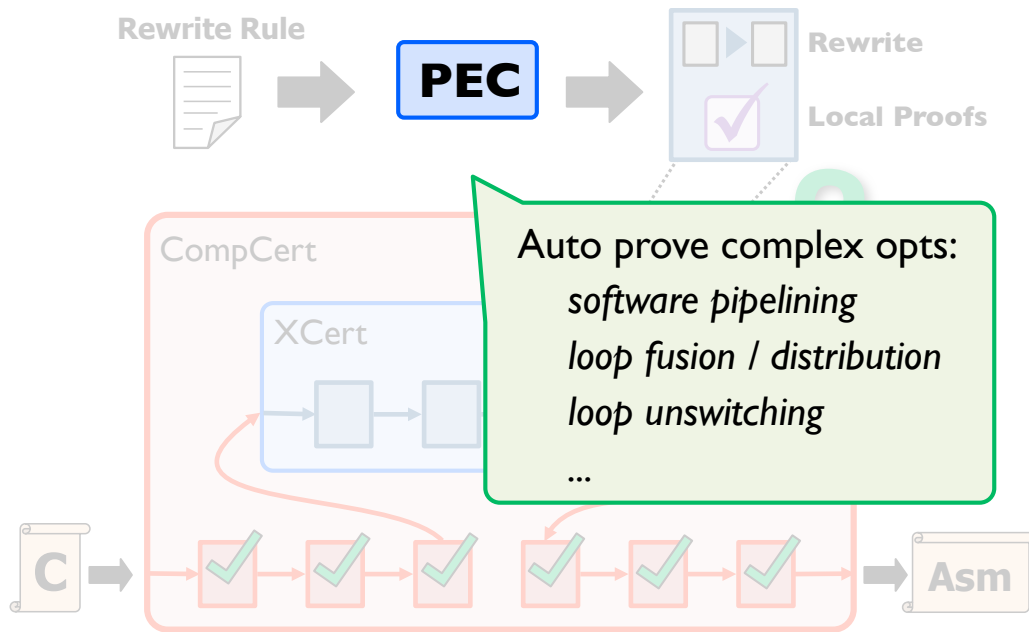
# Verifying Optimizations



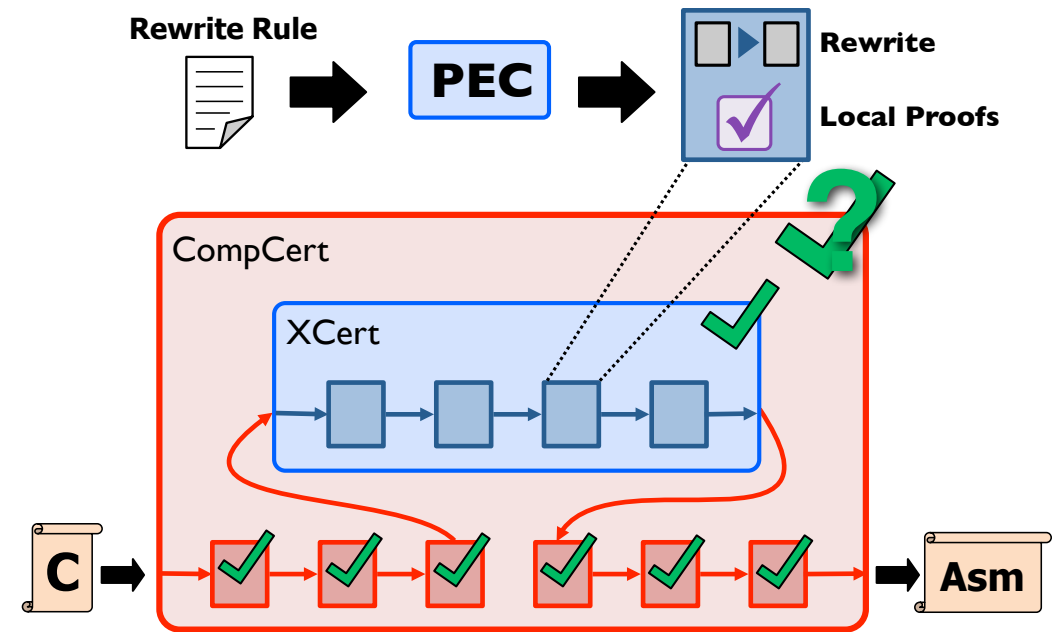
# Verifying Optimizations



# Verifying Optimizations



# Verifying Optimizations



## Future Work

Generating and evaluating specs


*techniques to ensure spec matches intuition*

Even perfect program verification can only establish that a program meets its specification... Much of the essence of building a program is in fact the debugging of the specification.

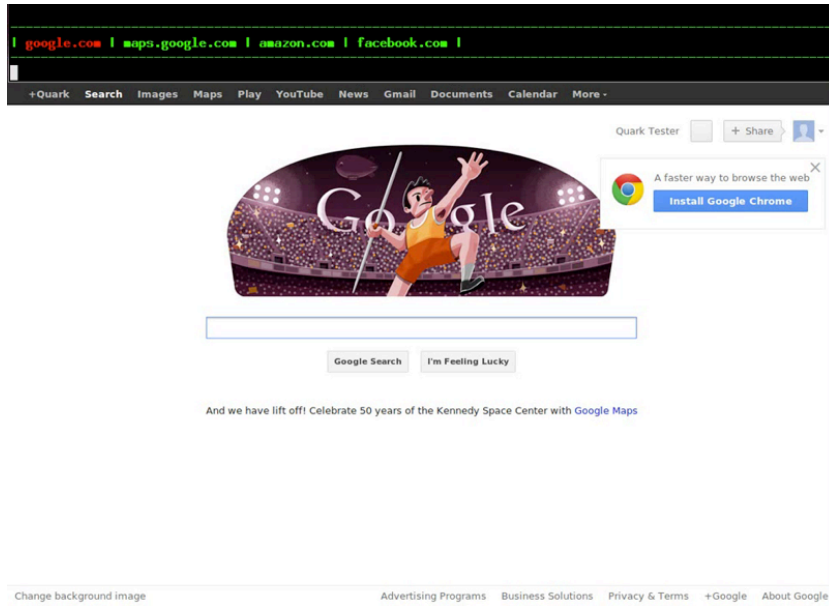
Frederick P. Brooks, Jr.  
No Silver Bullet



## Software Infrastructure



# Quark Usability



# Browsers: Critical Infrastructure

## Browsers: Critical Infrastructure


## Browsers: Critical Infrastructure



# Browsers: Critical Infrastructure



# Browsers: Critical Infrastructure



# Browsers: Critical Infrastructure

