



Automating Formal Proofs for Reactive Systems

Daniel Ricketts, Valentin Robert,
Dongseok Jang, Sorin Lerner

University of California, San Diego

Dr. Zachary Tatlock

University of Washington

Proof Assistant Based Verification

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

<i>Compiler</i>	<i>Bugs Found</i>
-----------------	-------------------

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

<i>Compiler</i>	<i>Bugs Found</i>
GCC	122

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

<i>Compiler</i>	<i>Bugs Found</i>
GCC	122
Clang/LLVM	181

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

<i>Compiler</i>	<i>Bugs Found</i>
GCC	122
Clang/LLVM	181
CompCert	?

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]

<i>Compiler</i>	<i>Bugs Found</i>
GCC	122
Clang/LLVM	181
CompCert	0

Proof Assistant Based Verification

Testing tool for C compilers [Yang et al. PLDI 11]
[Le et al. PLDI 14]

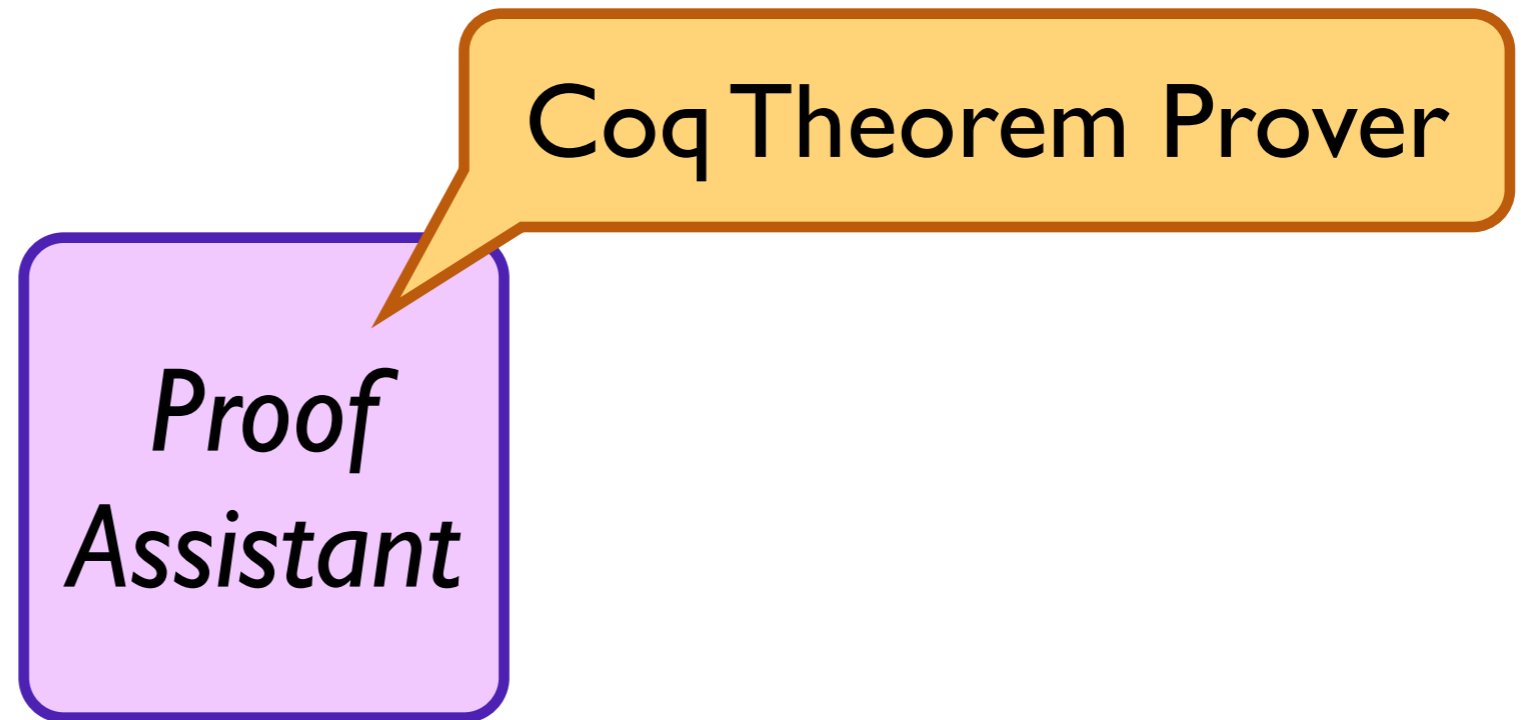
<i>Compiler</i>	<i>Bugs Found</i>
GCC	122
Clang/LLVM	181
CompCert	0

Proof Assistant Based Verification

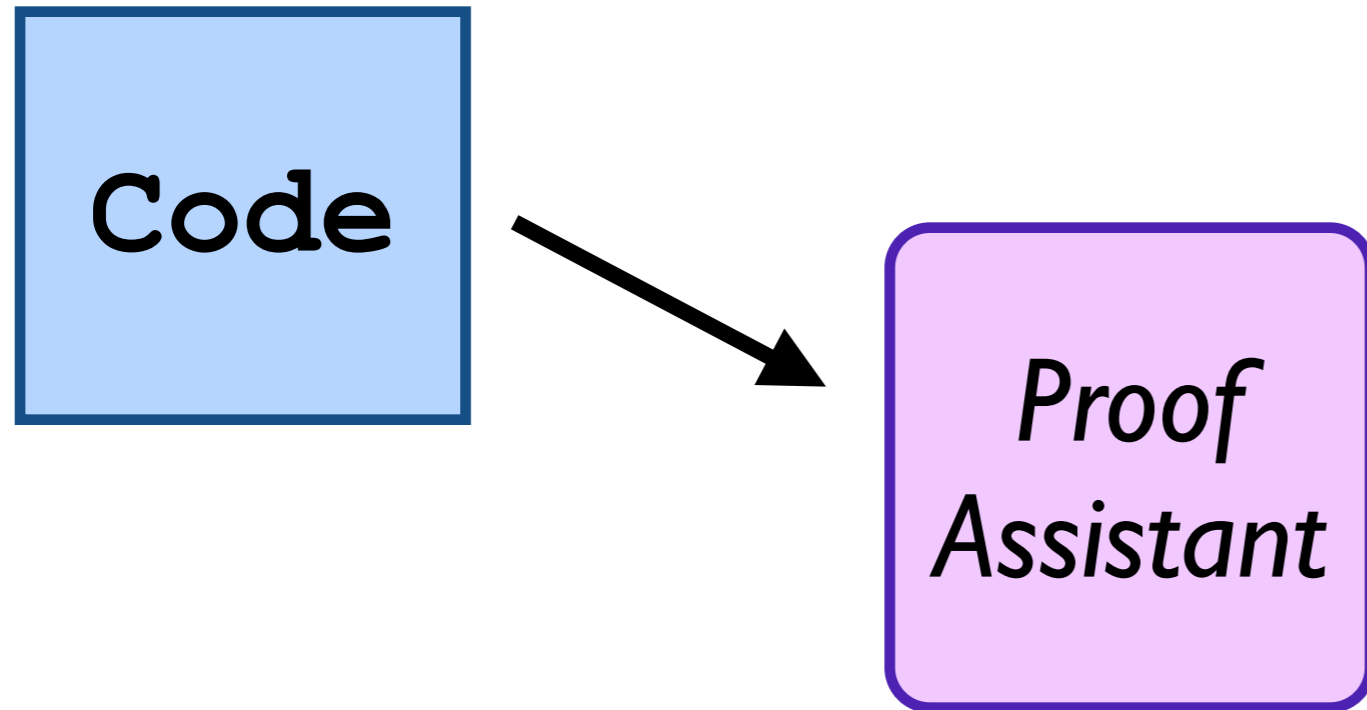
Proof Assistant Based Verification

*Proof
Assistant*

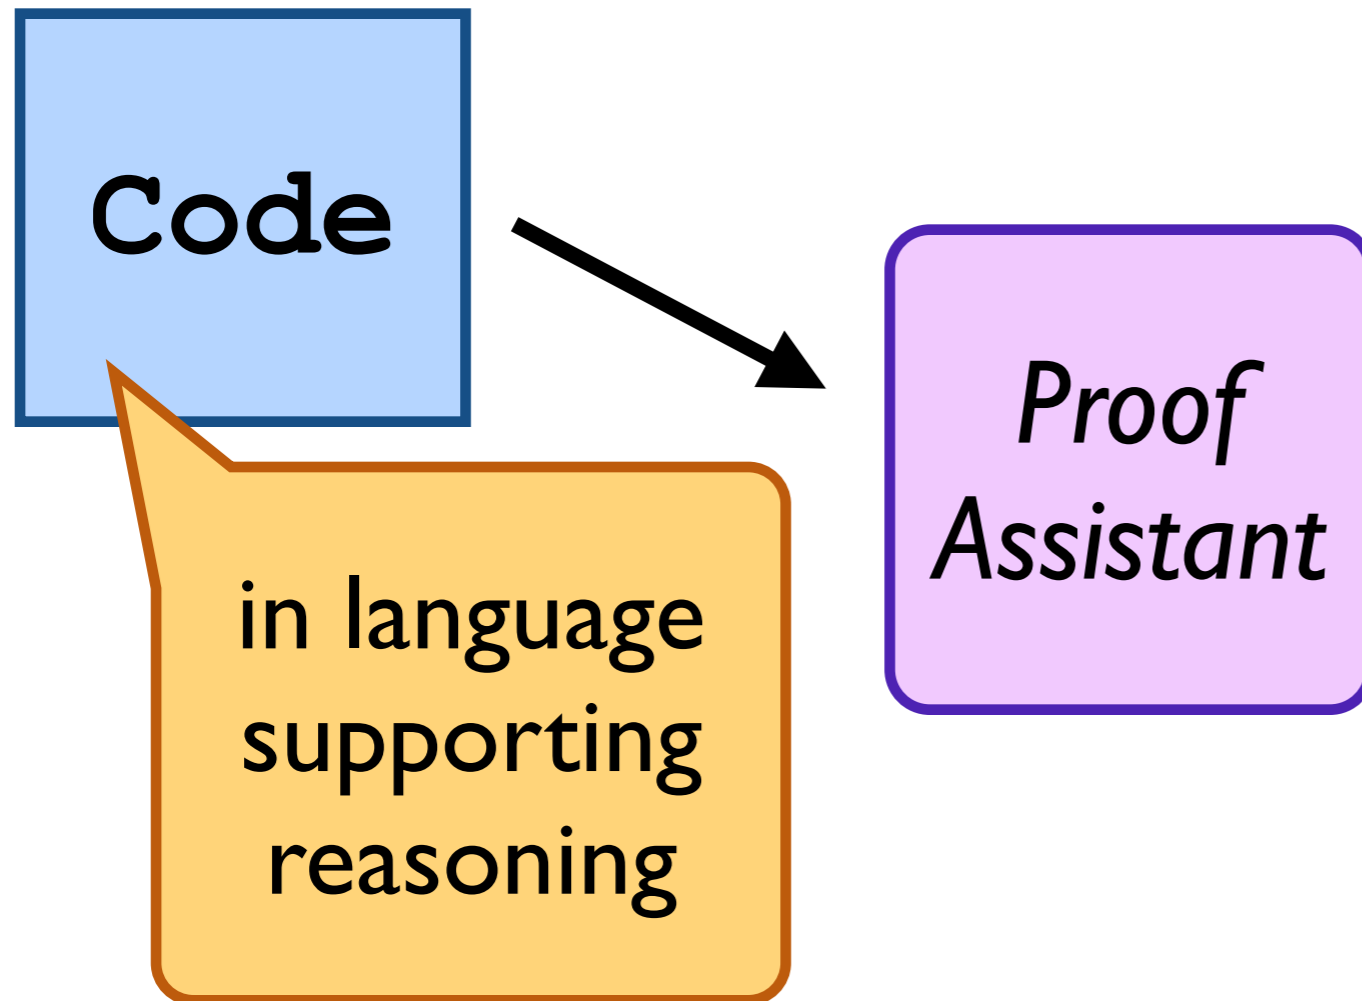
Proof Assistant Based Verification



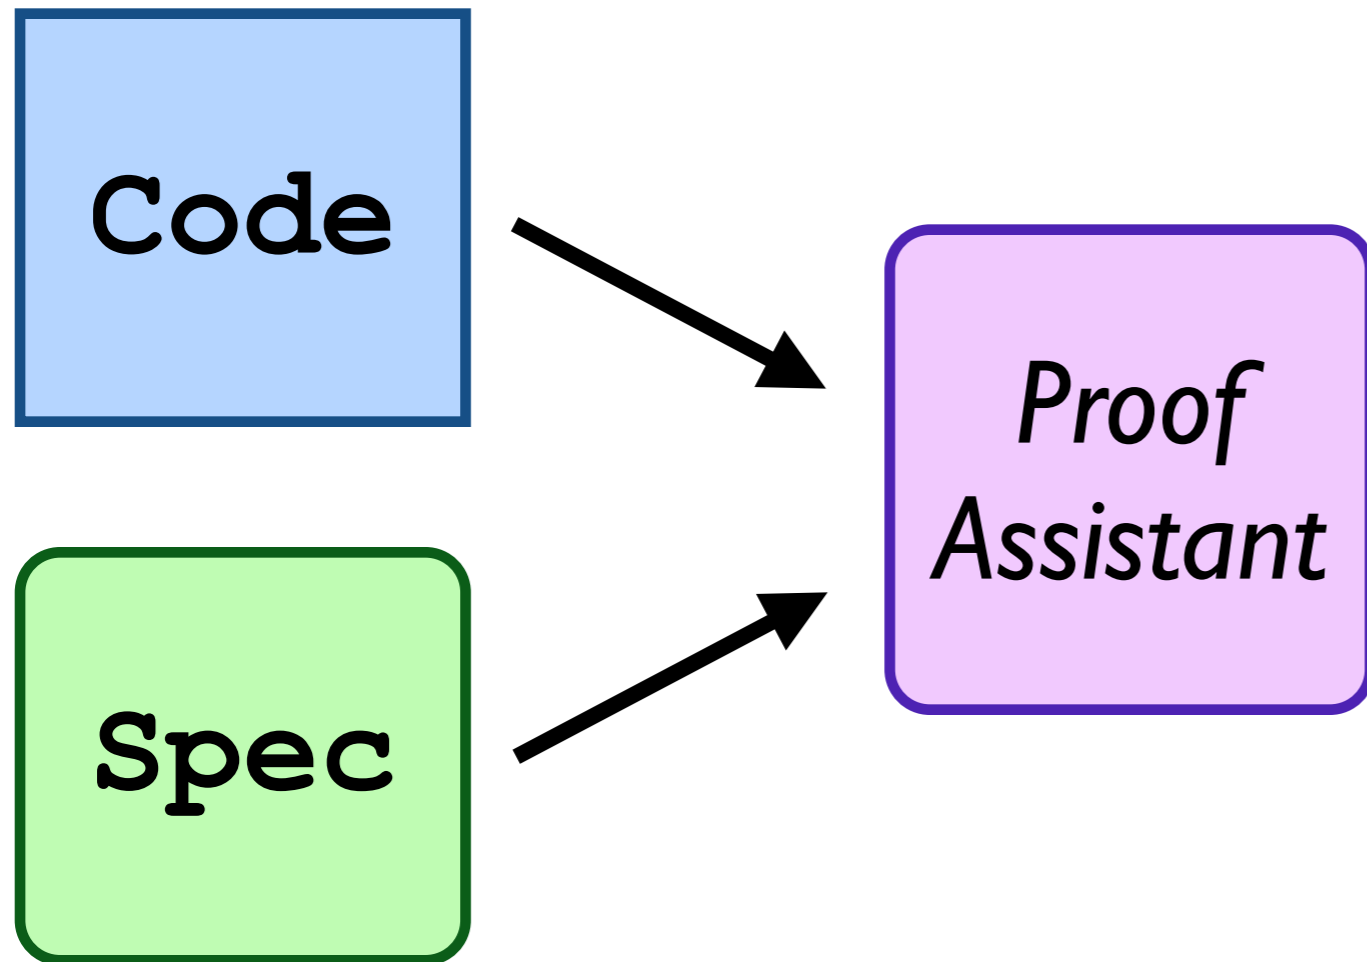
Proof Assistant Based Verification



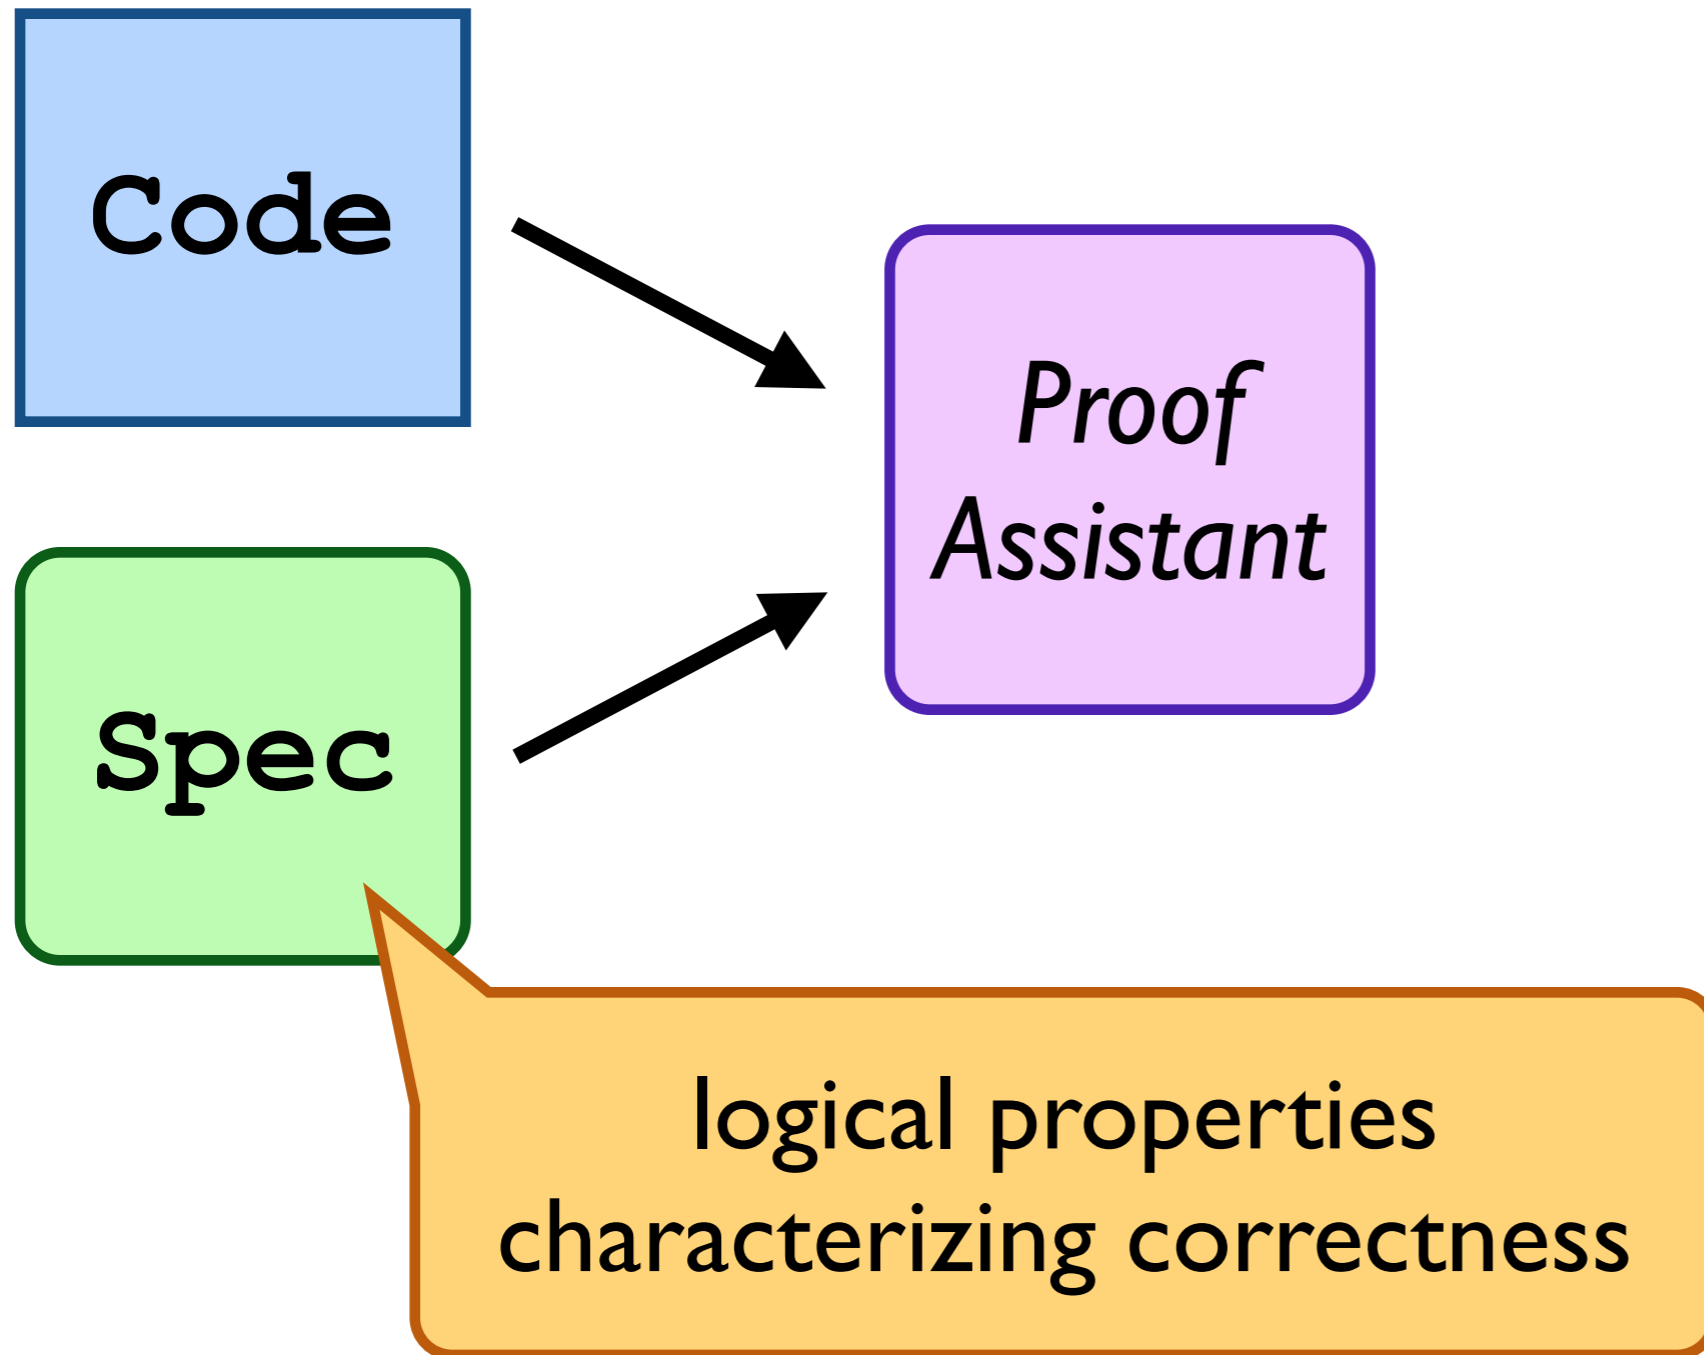
Proof Assistant Based Verification



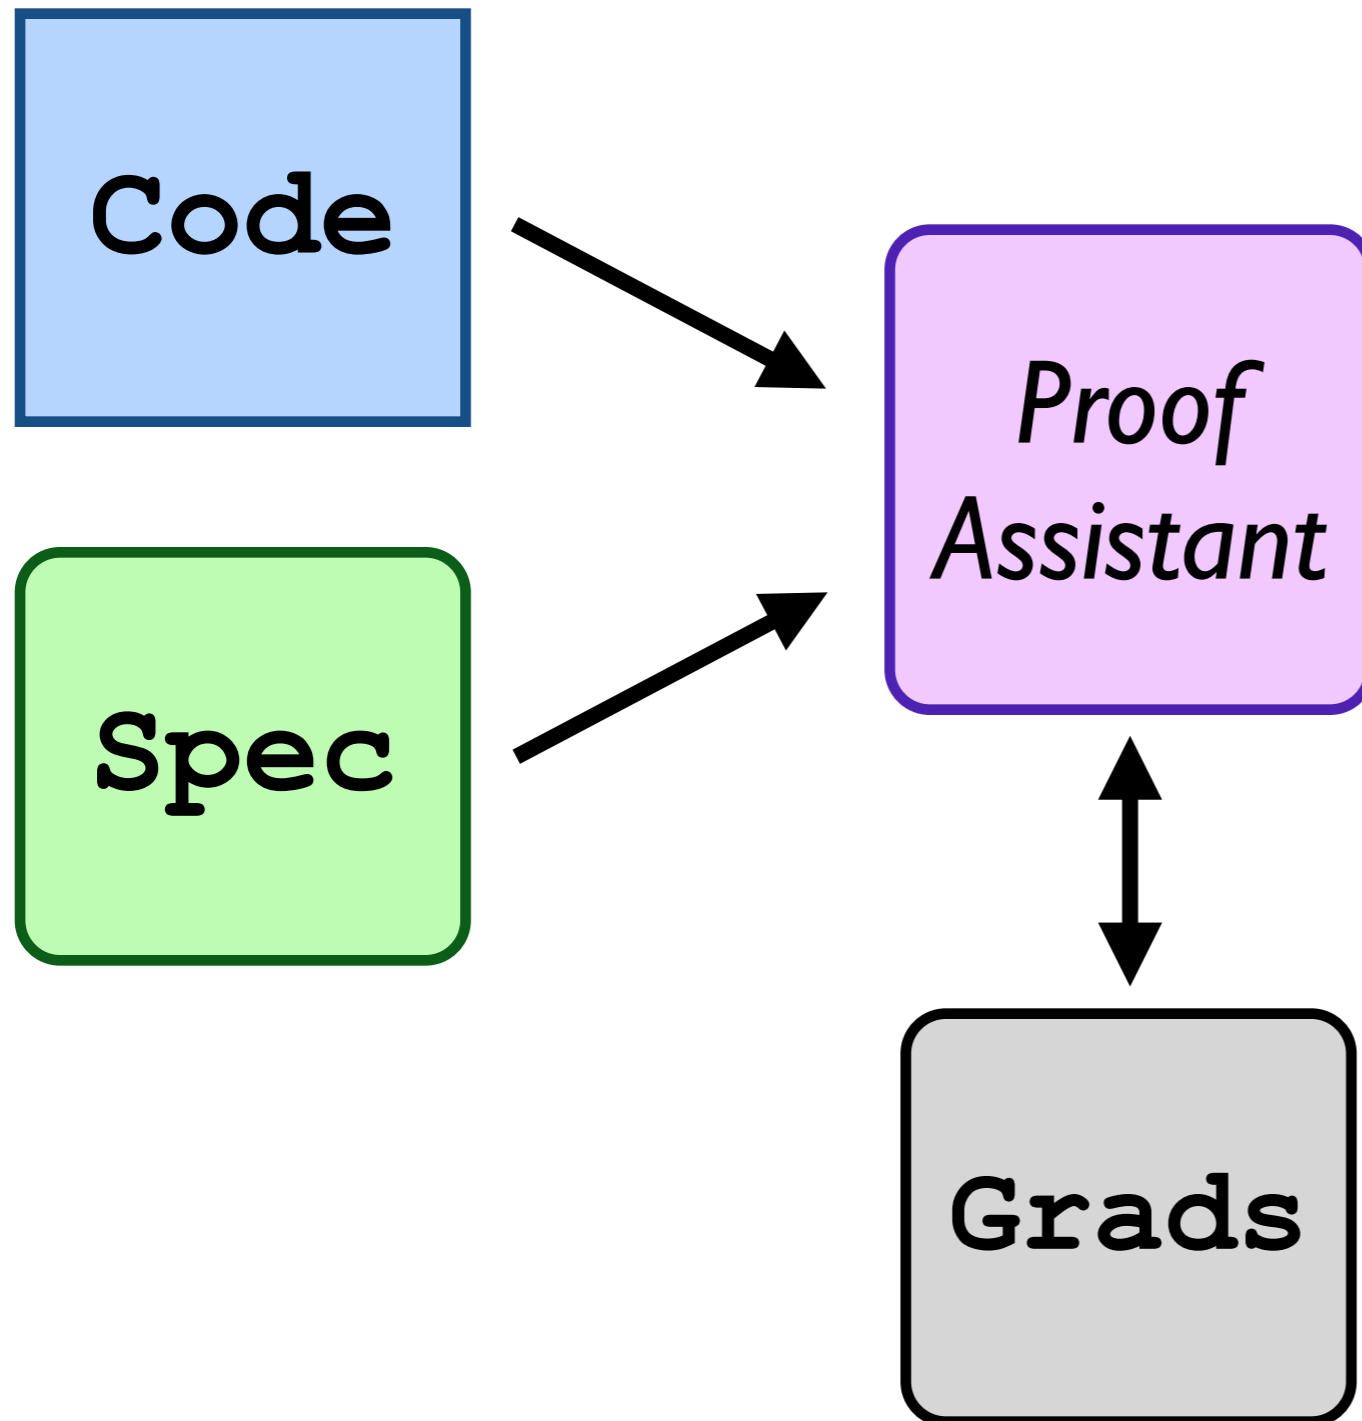
Proof Assistant Based Verification



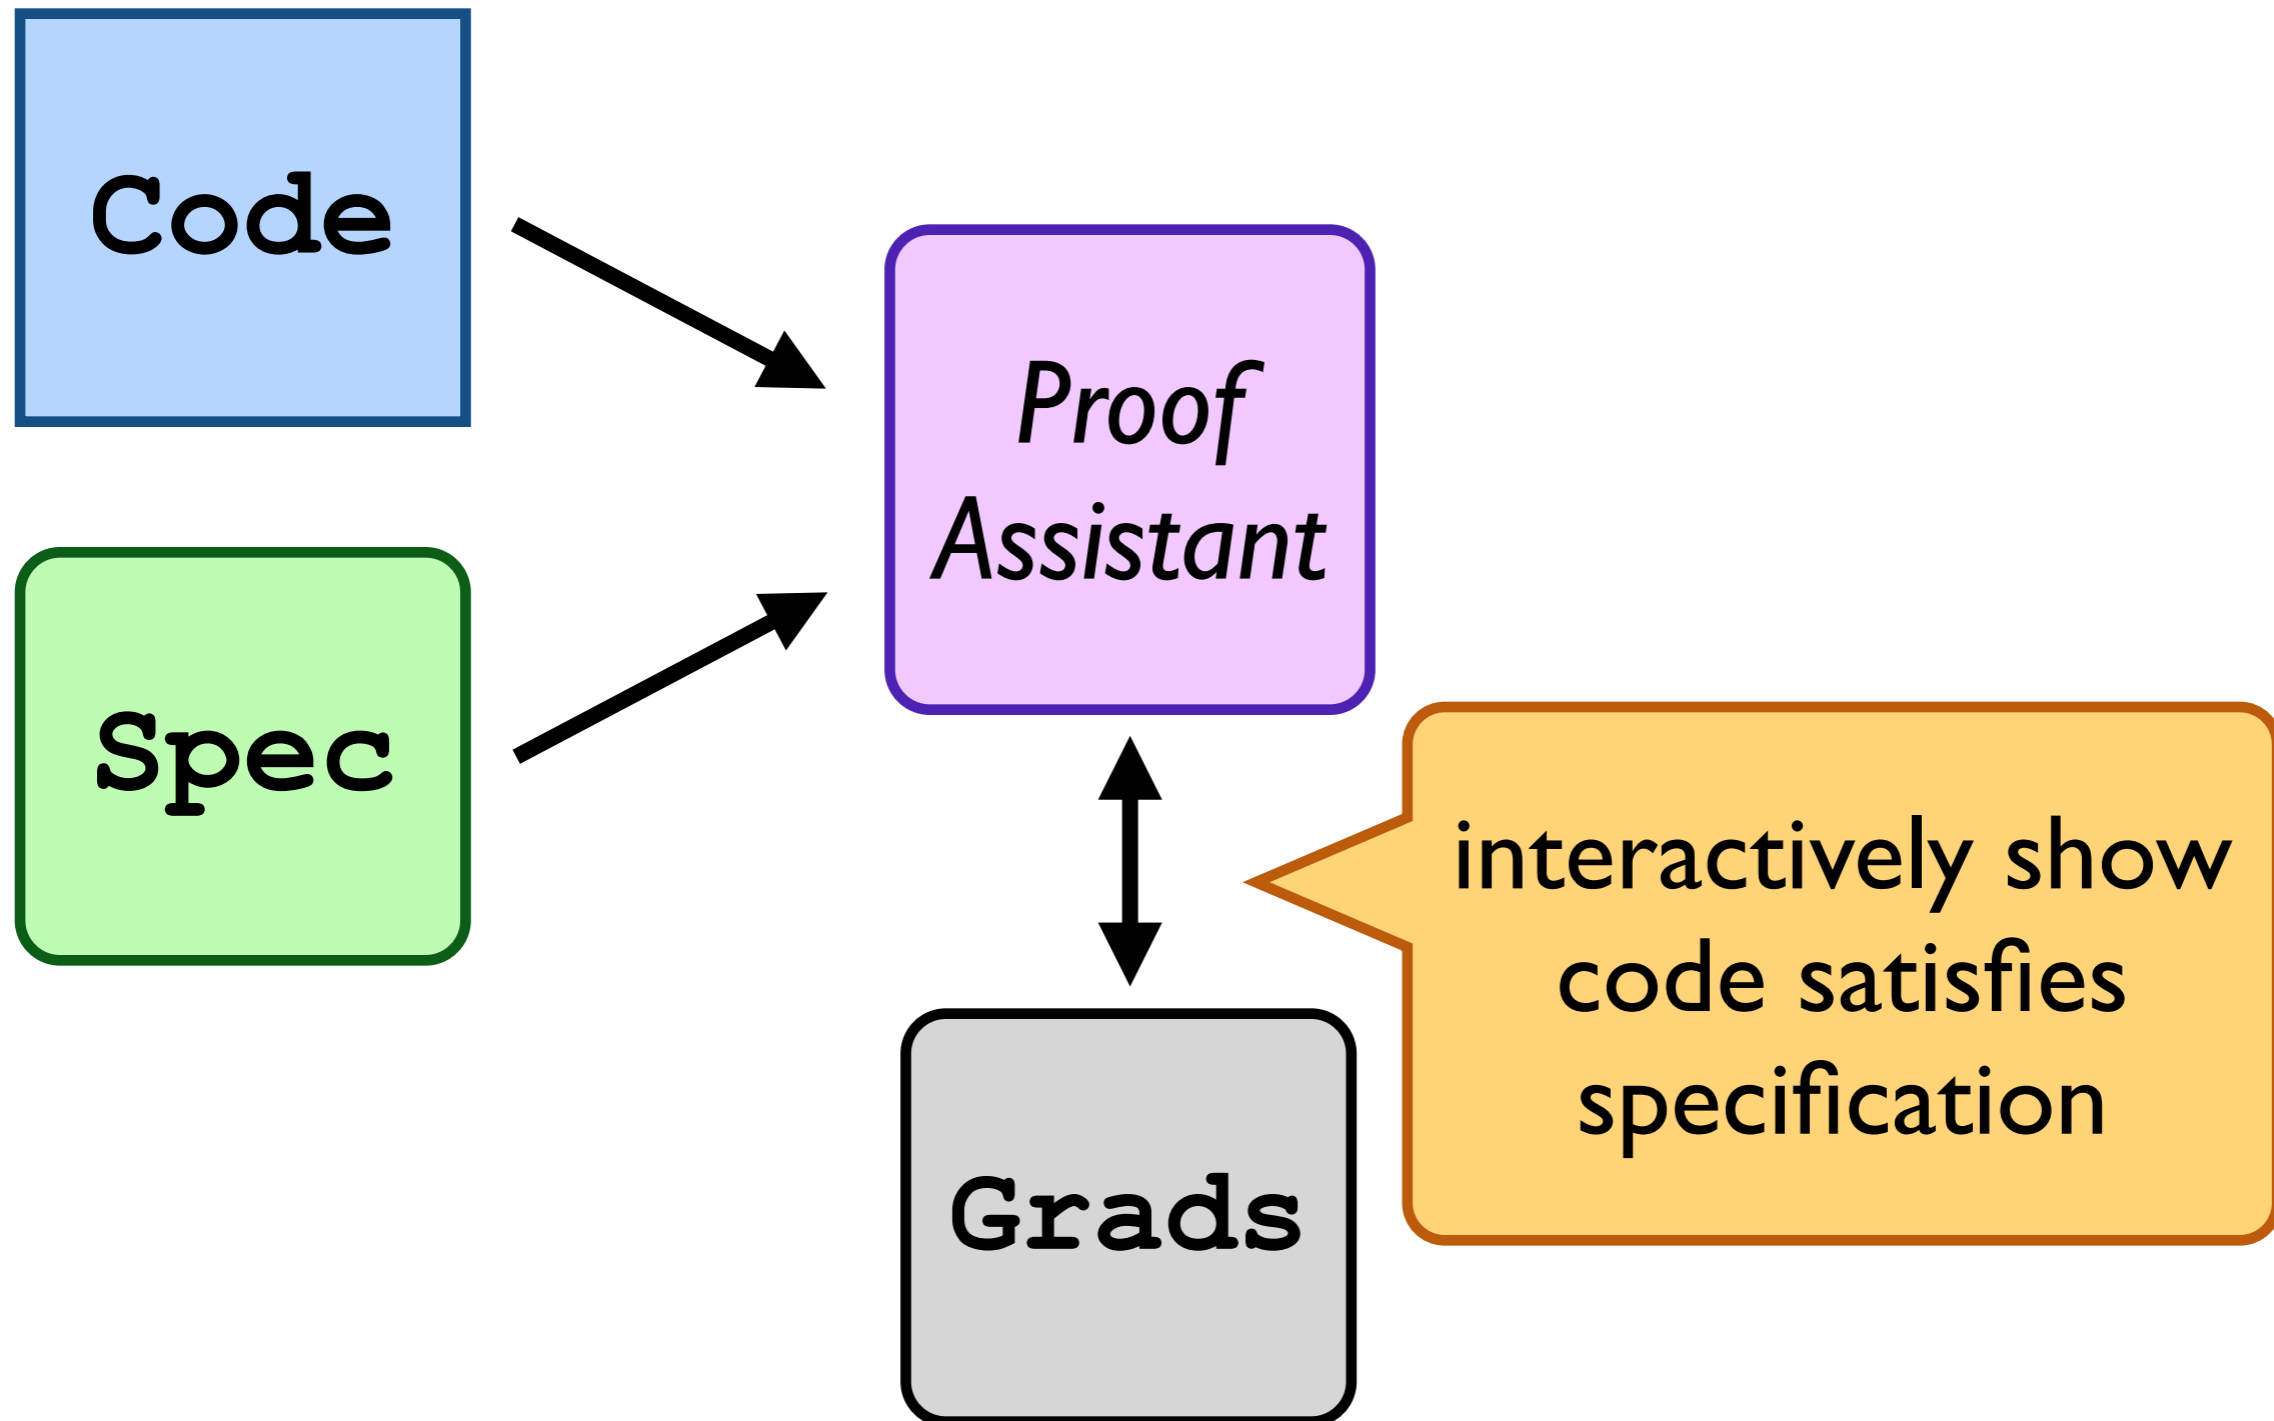
Proof Assistant Based Verification



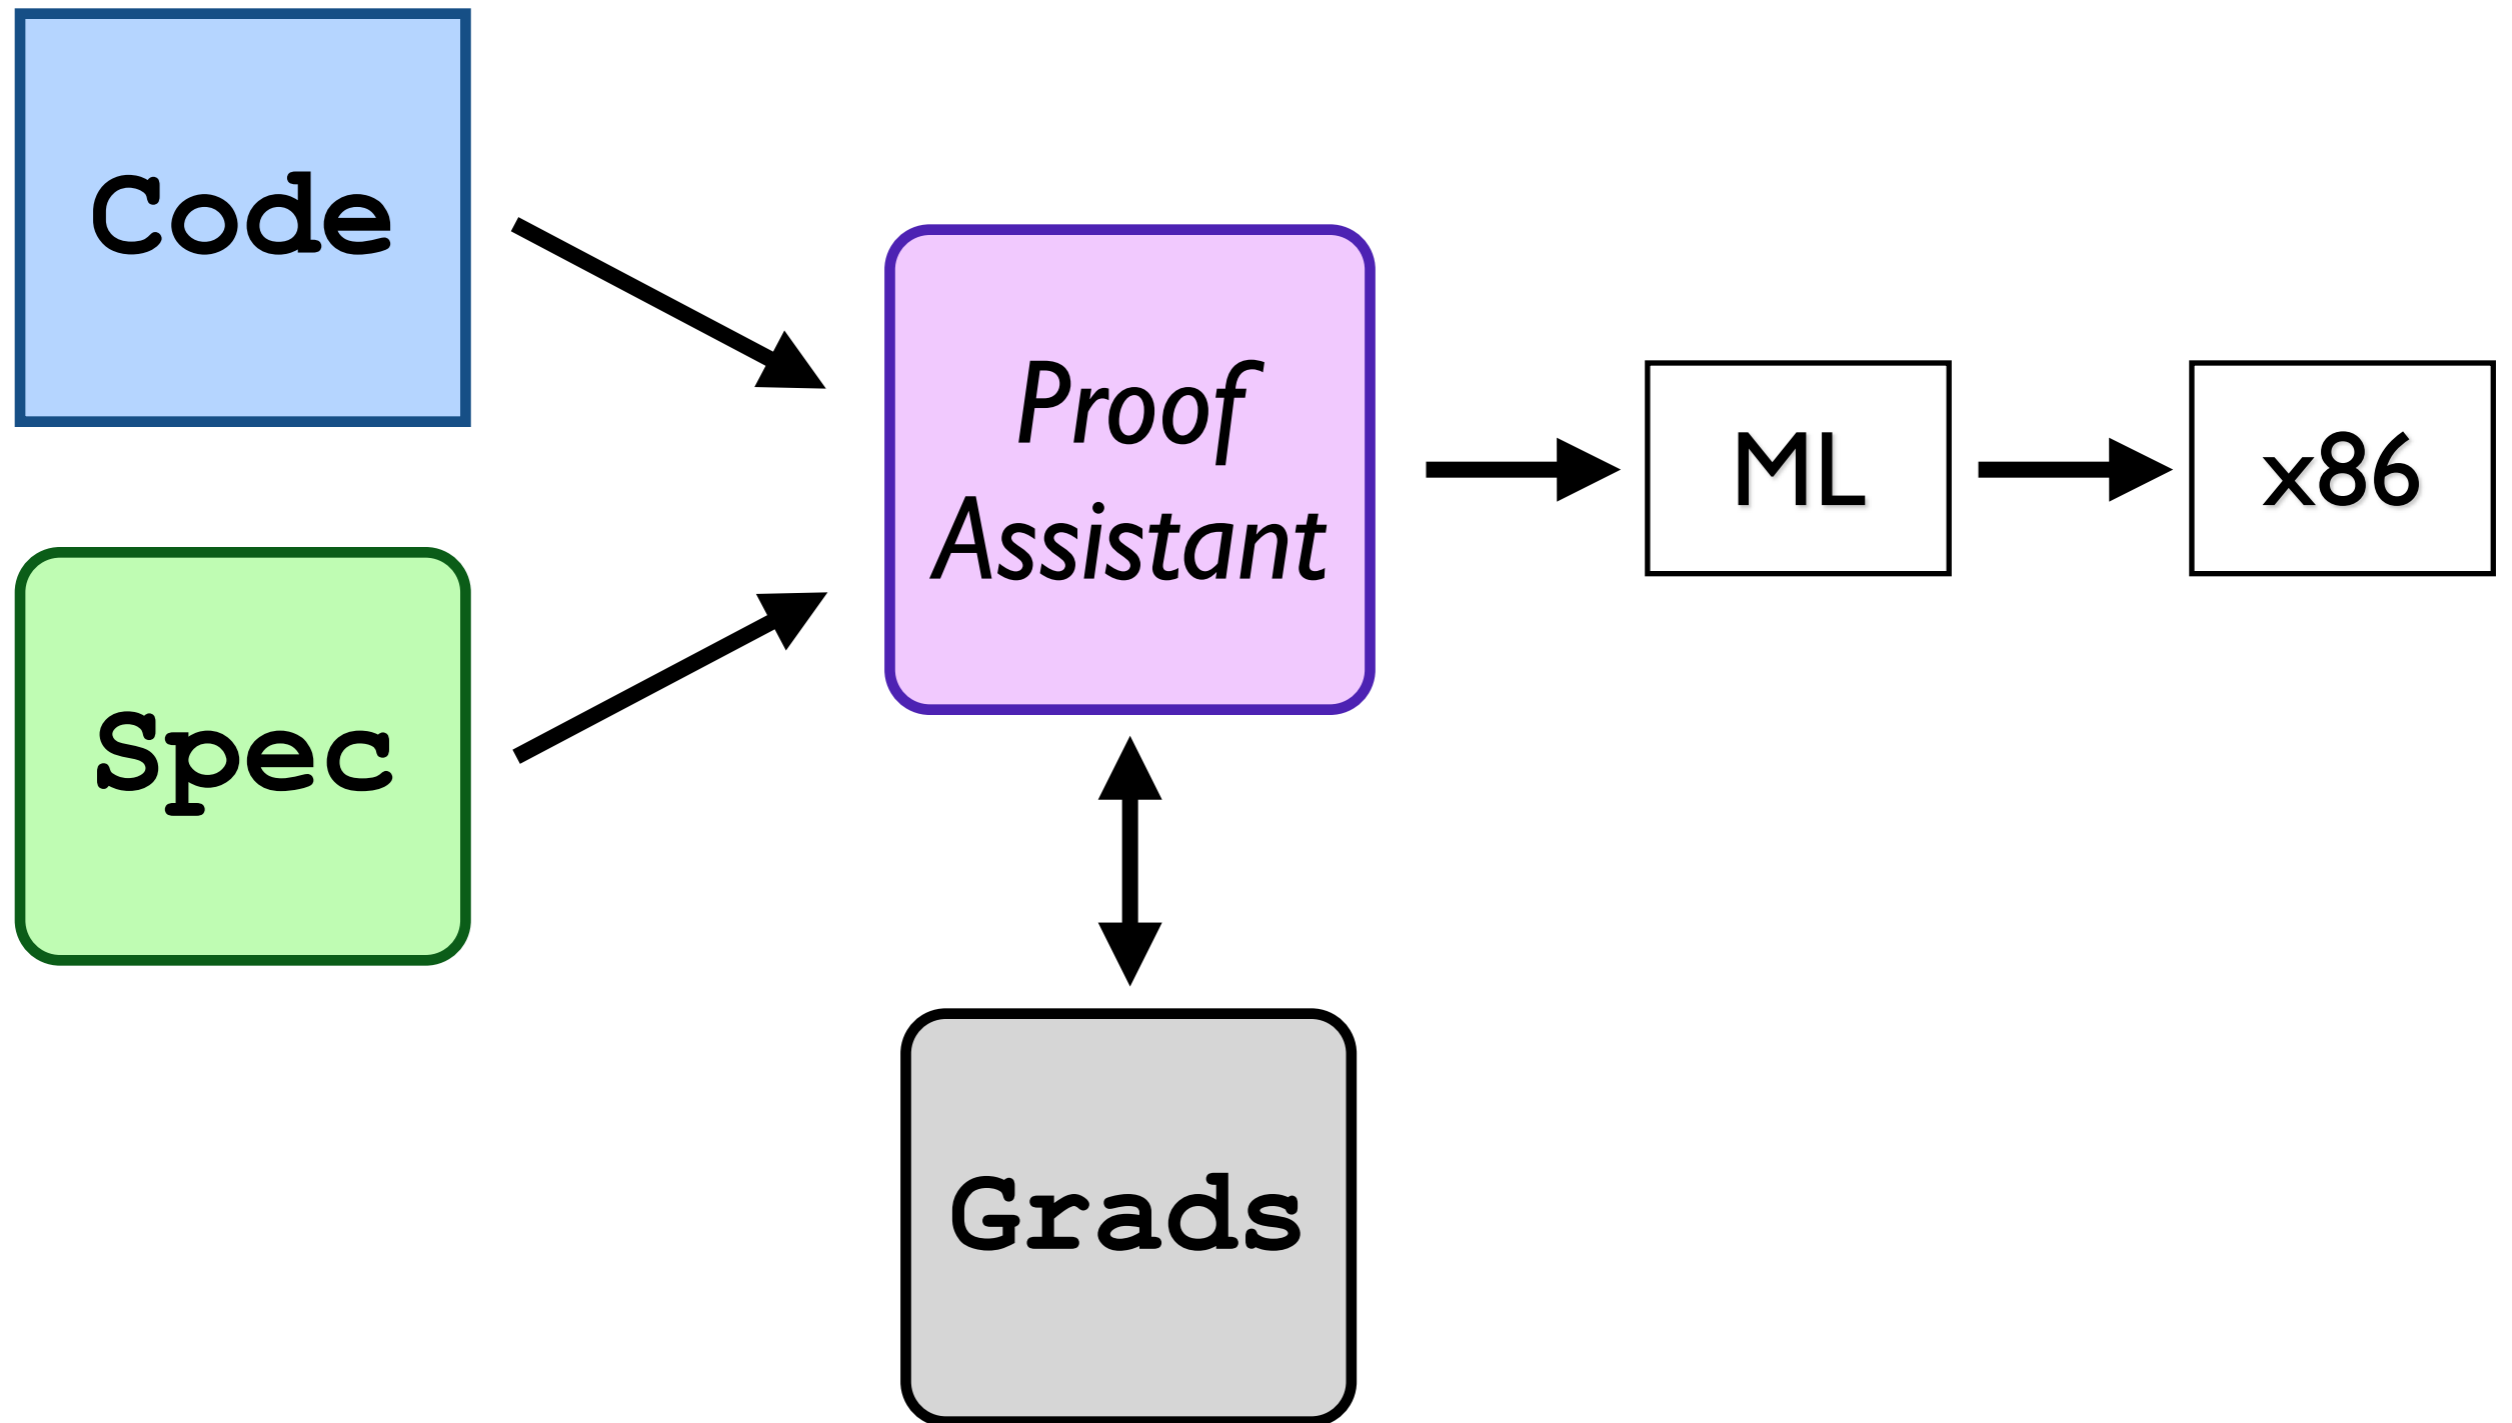
Proof Assistant Based Verification



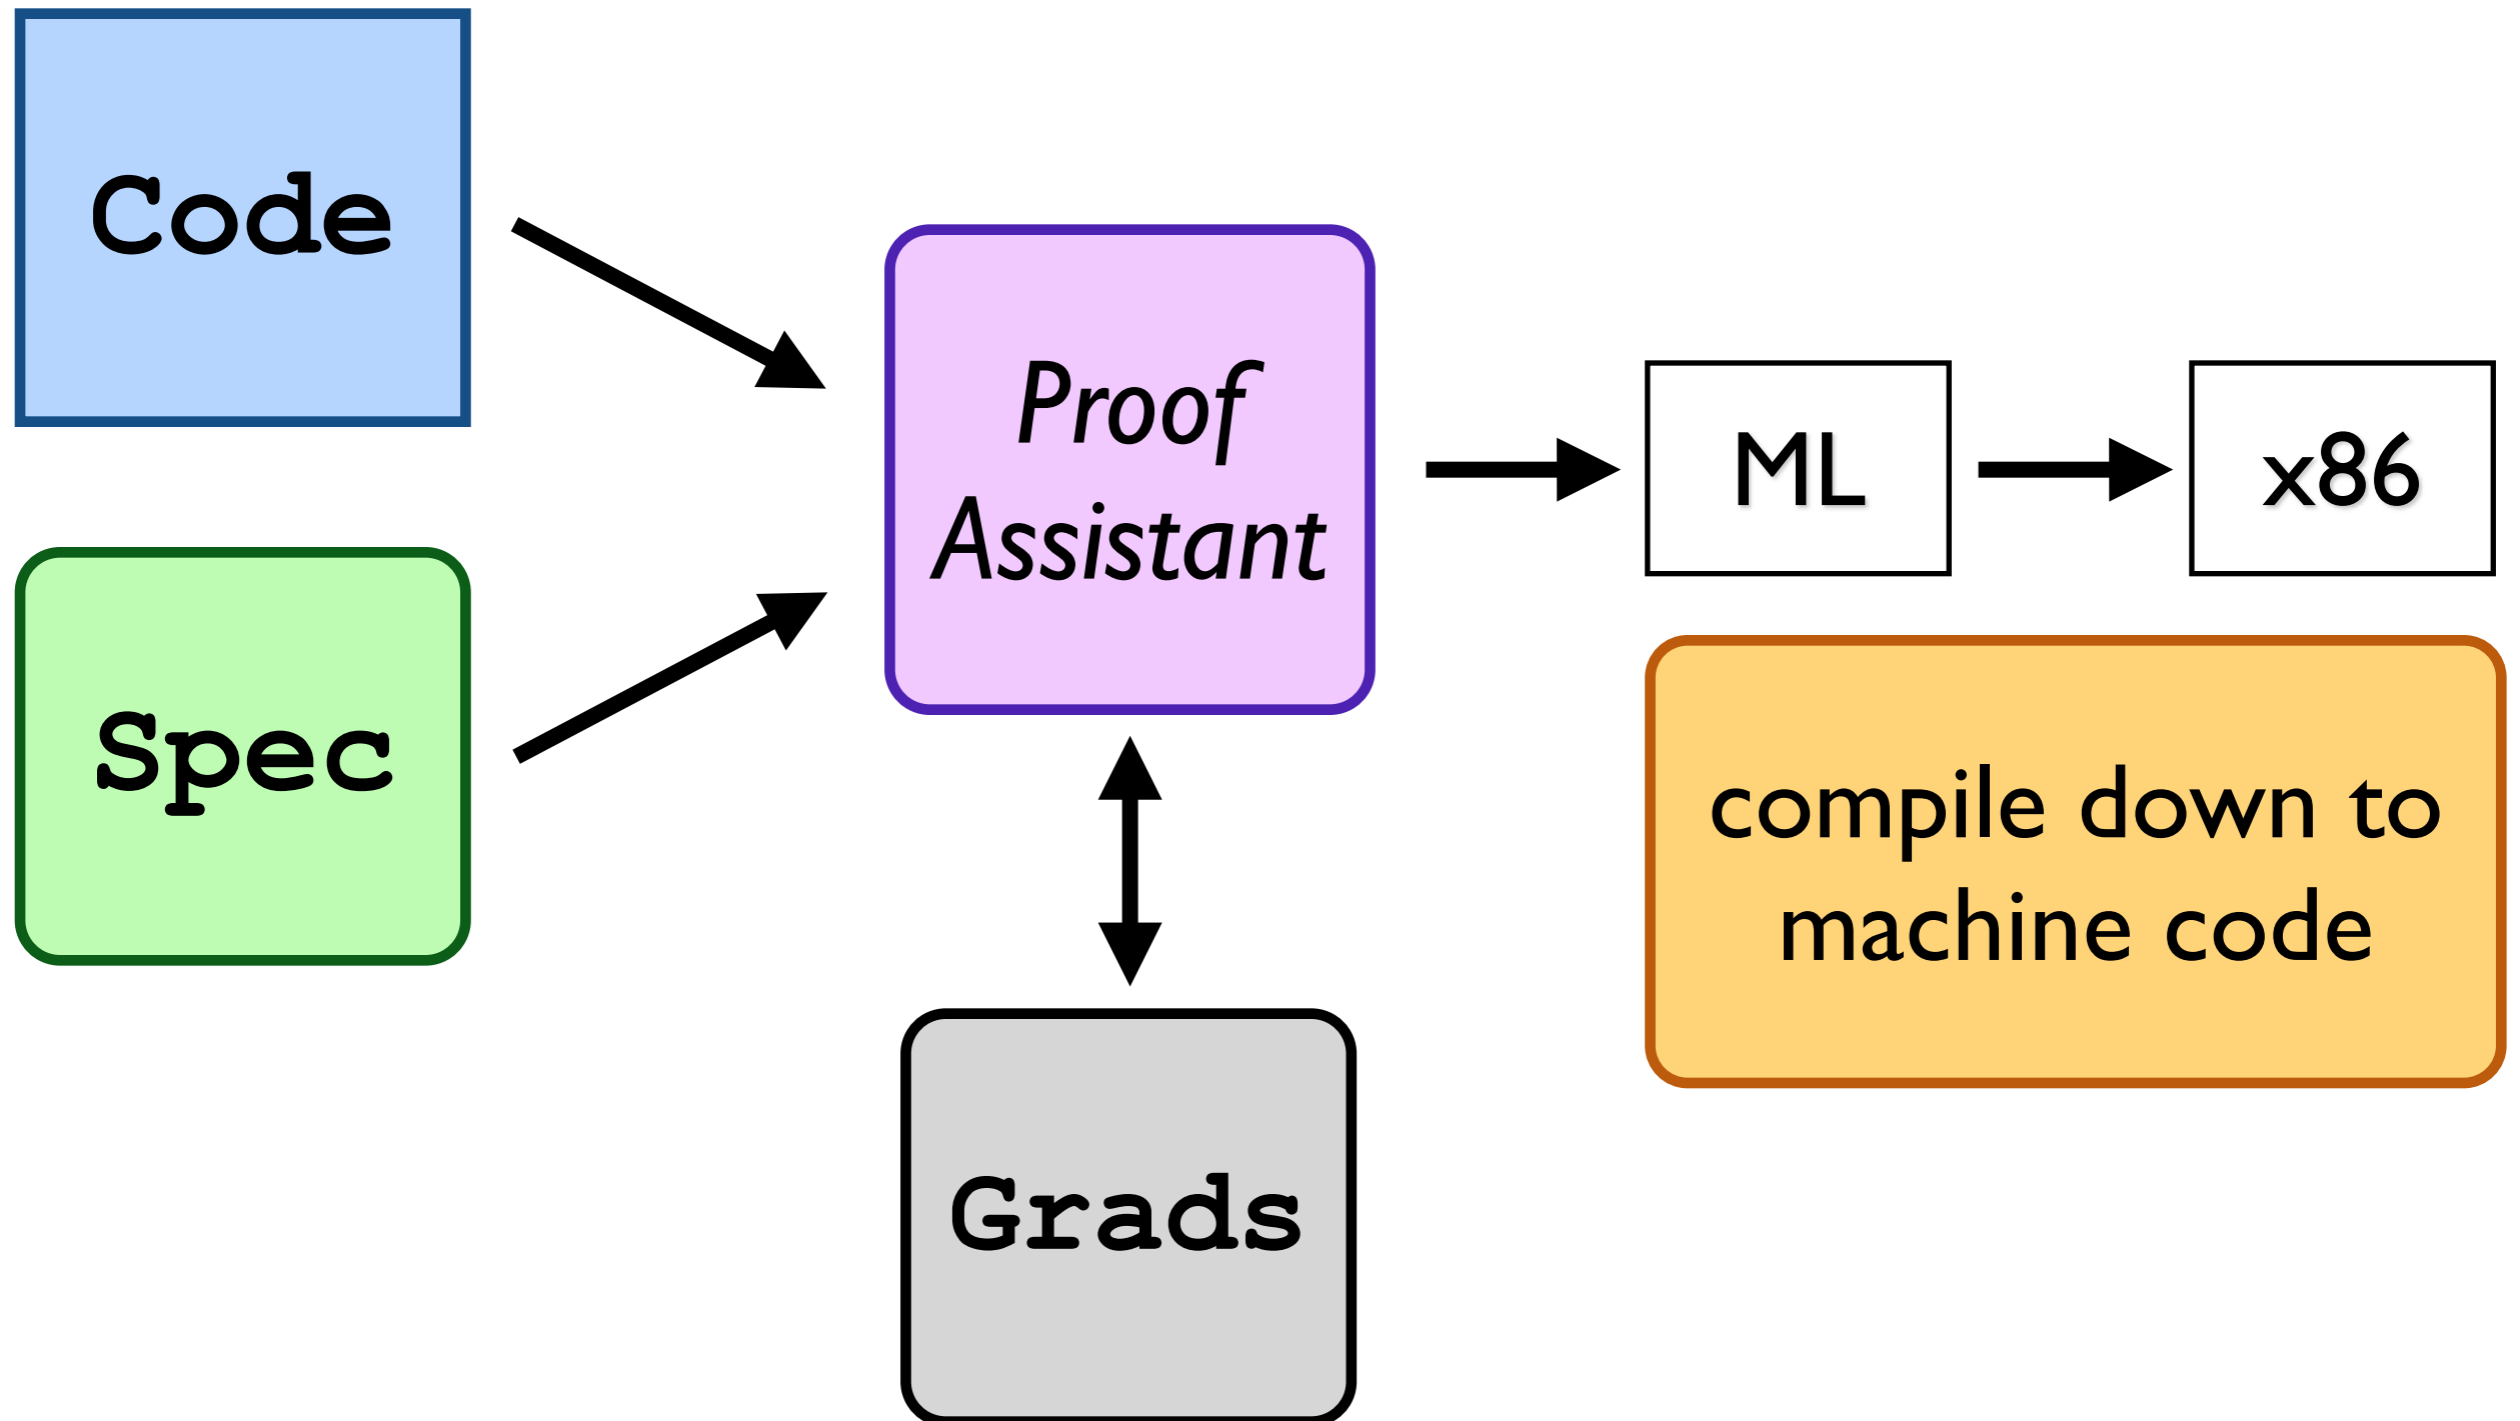
Proof Assistant Based Verification



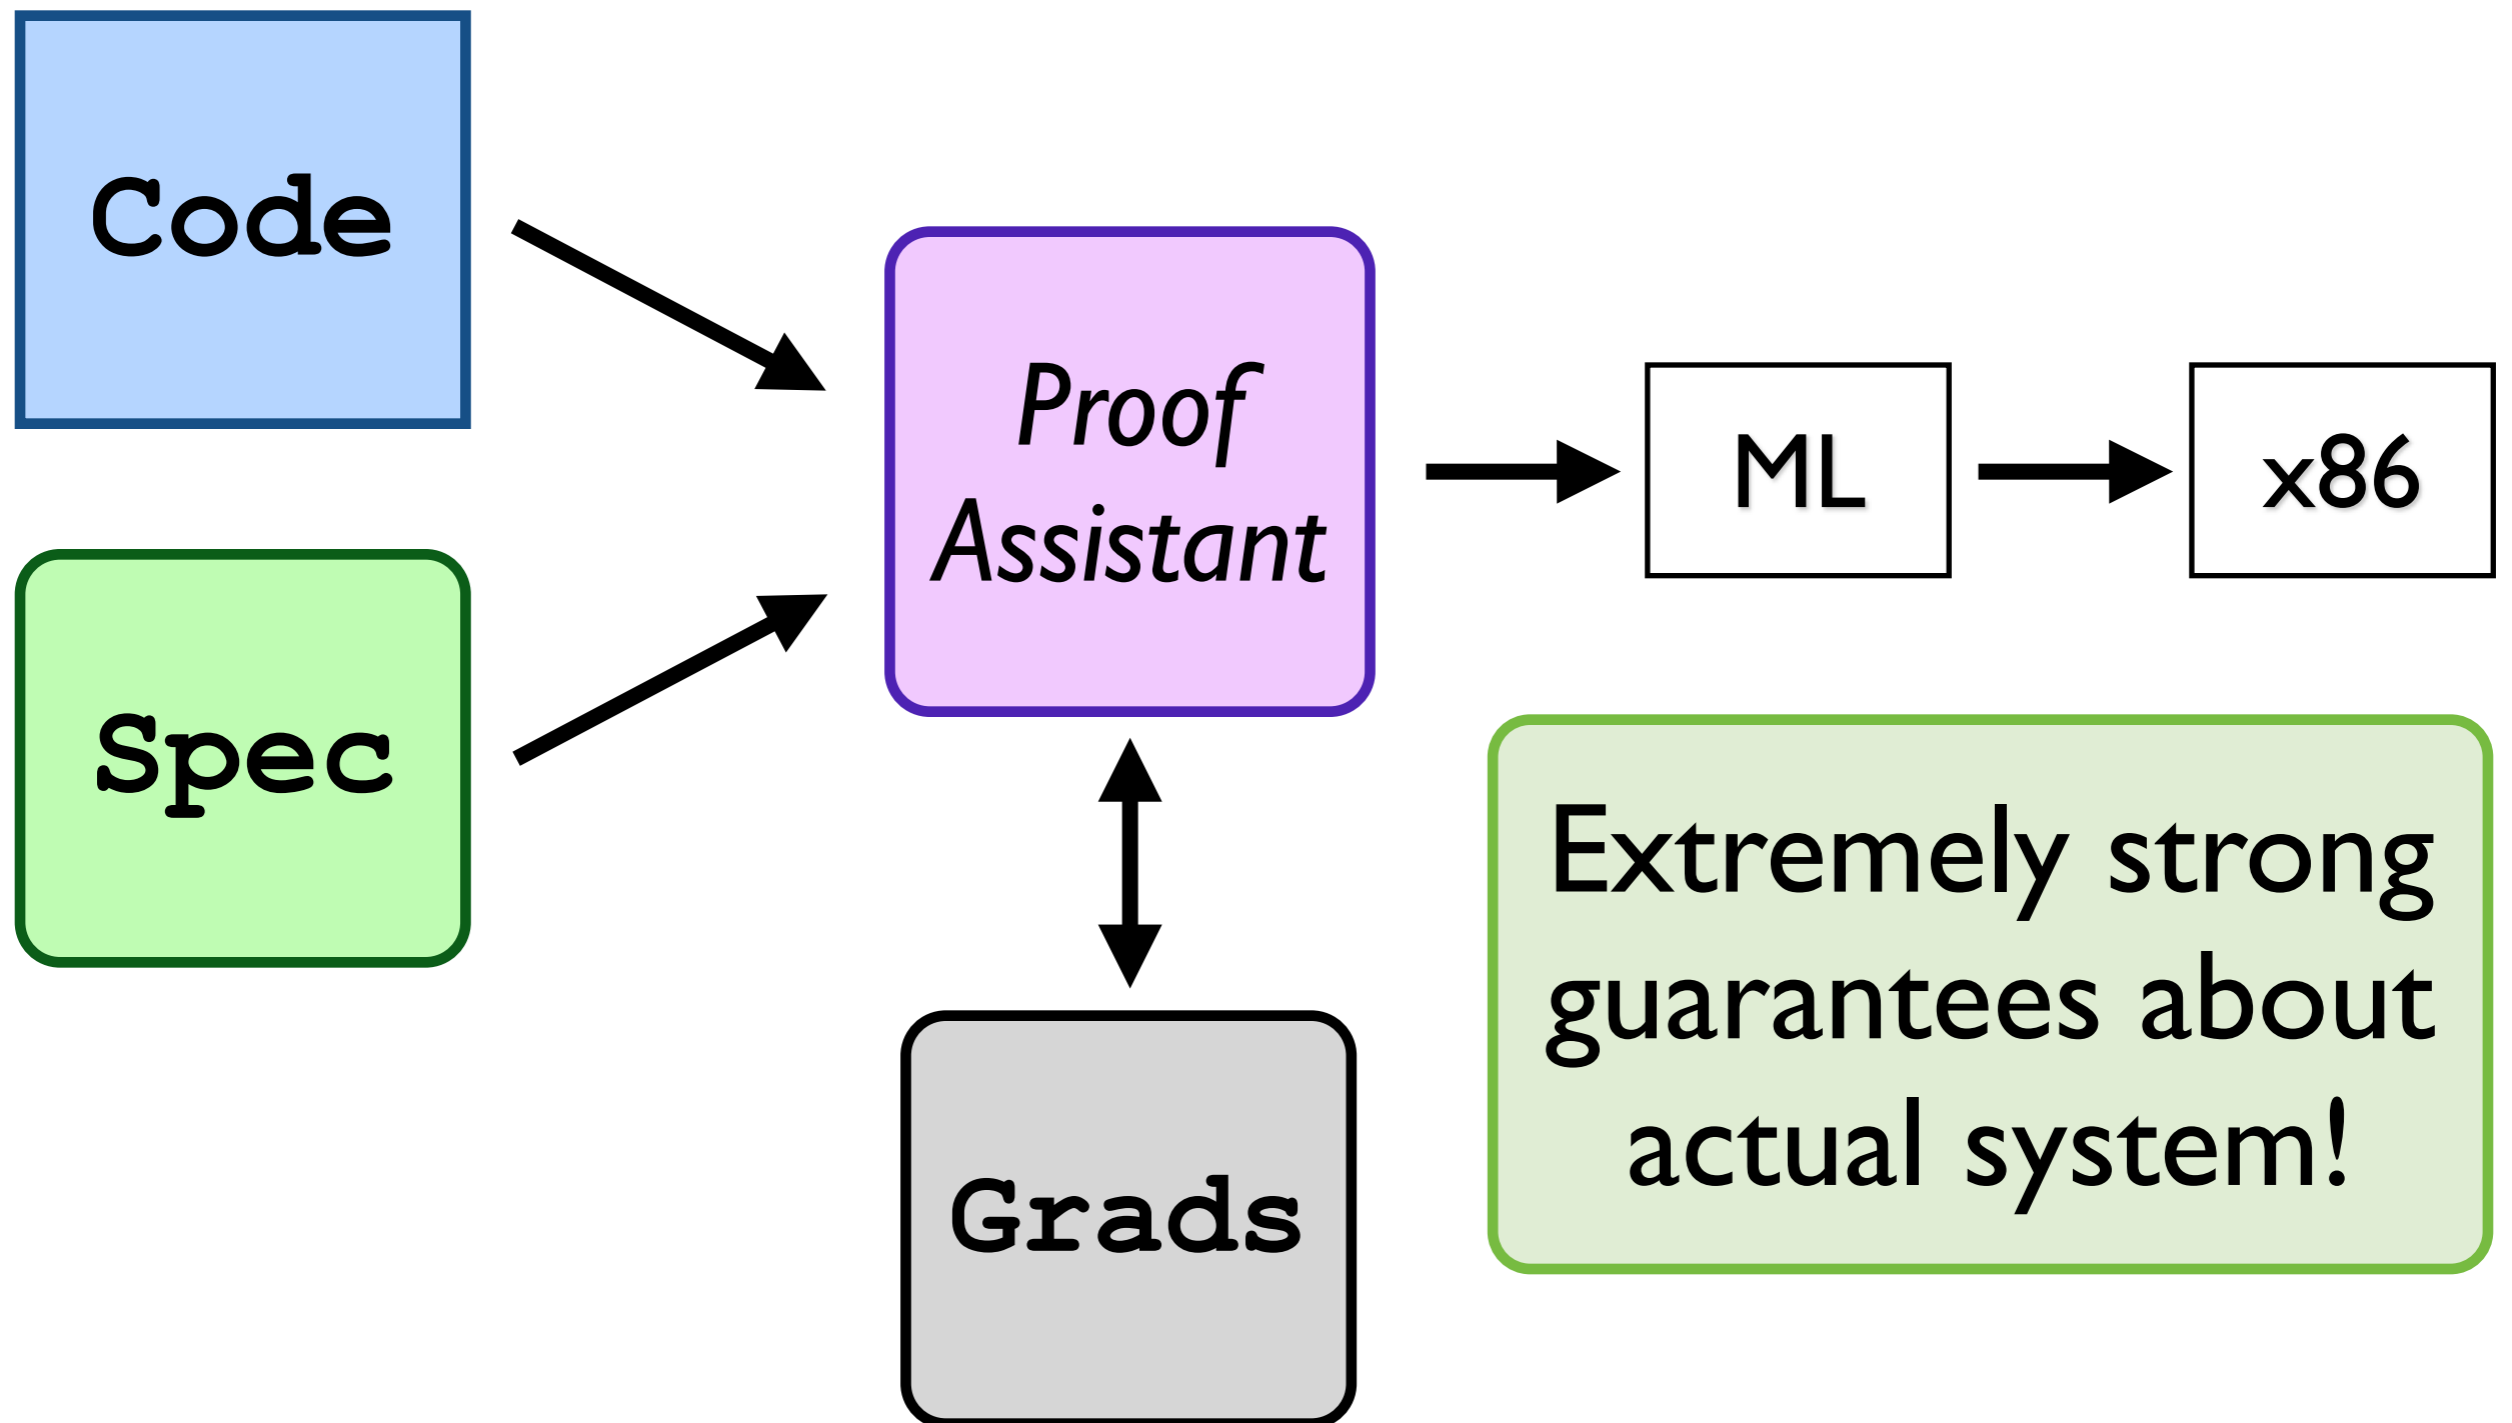
Proof Assistant Based Verification



Proof Assistant Based Verification



Proof Assistant Based Verification



Proof Assistant Based Verification

Verified Compiler: **CompCert**

[Leroy POPL 06]

Proof Assistant Based Verification

Verified Compiler: **CompCert**

[Leroy POPL 06]

Verified OS micro-kernel: **seL4**

[Klein et al. SOSP 09]

Proof Assistant Based Verification

Verified Compiler: **CompCert**

[Leroy POPL 06]


Verified OS micro-kernel: **seL4**

[Klein et al. SOSP 09]

Verified Web browser: **Quark**

[Jang et al. Security 12]

Manual Proof Burden



Manual Proof Burden

Verified OS micro-kernel: **seL4**

Manual Proof Burden

Verified OS micro-kernel: **seL4**

9,000 lines of C code


Manual Proof Burden

Verified OS micro-kernel: **seL4**

9,000 lines of C code

20 person-years for verification

Manual Proof Burden



Manual Proof Burden

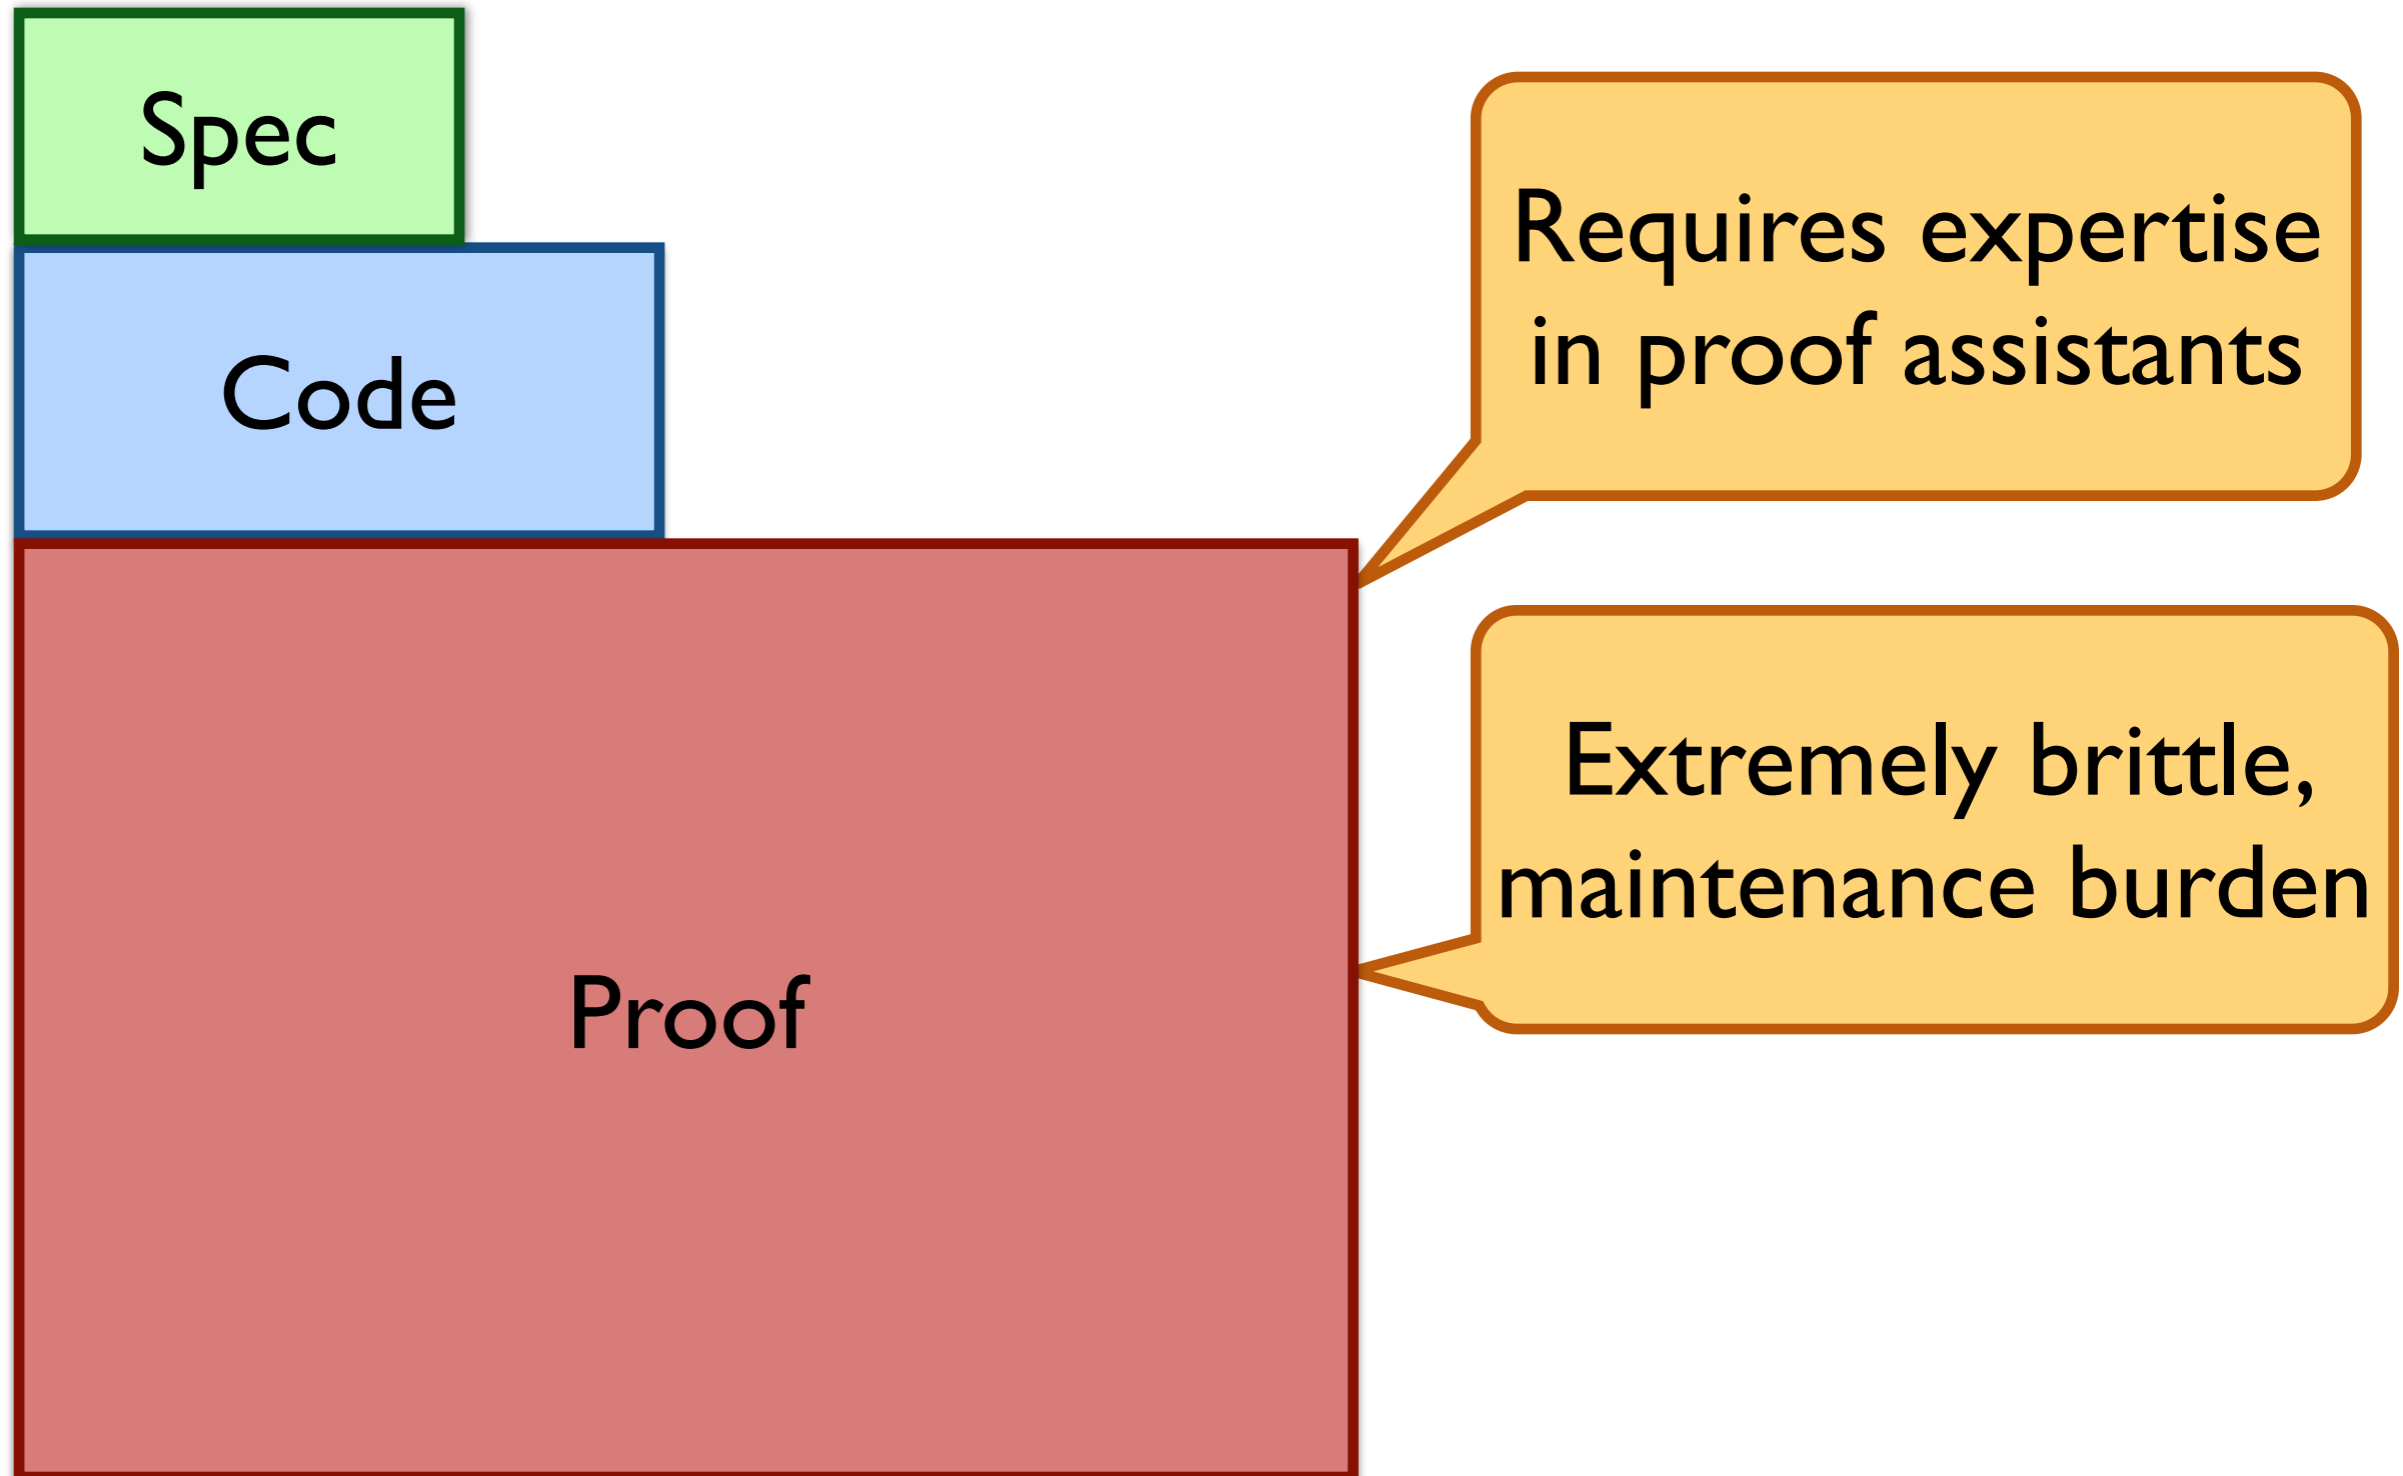
Spec

Code


Proof

Requires expertise
in proof assistants


Manual Proof Burden



Manual Proof Burden



Manual Proof Burden



Single application domain

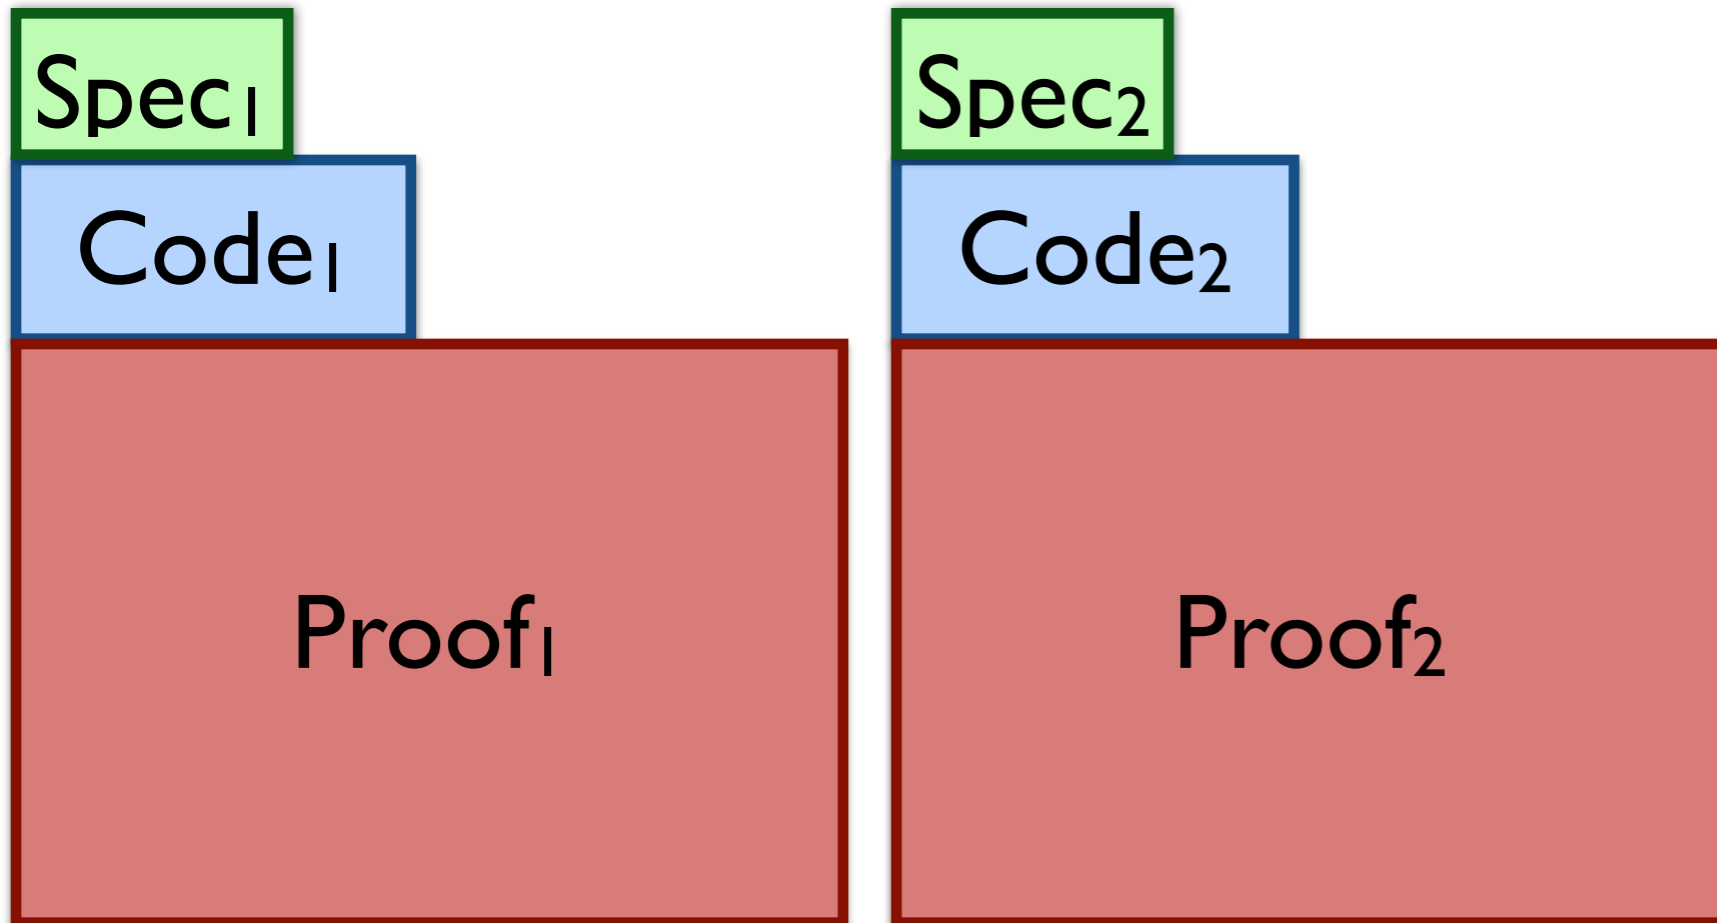
Single application domain

Spec₁


Code₁

Proof₁

Single application domain

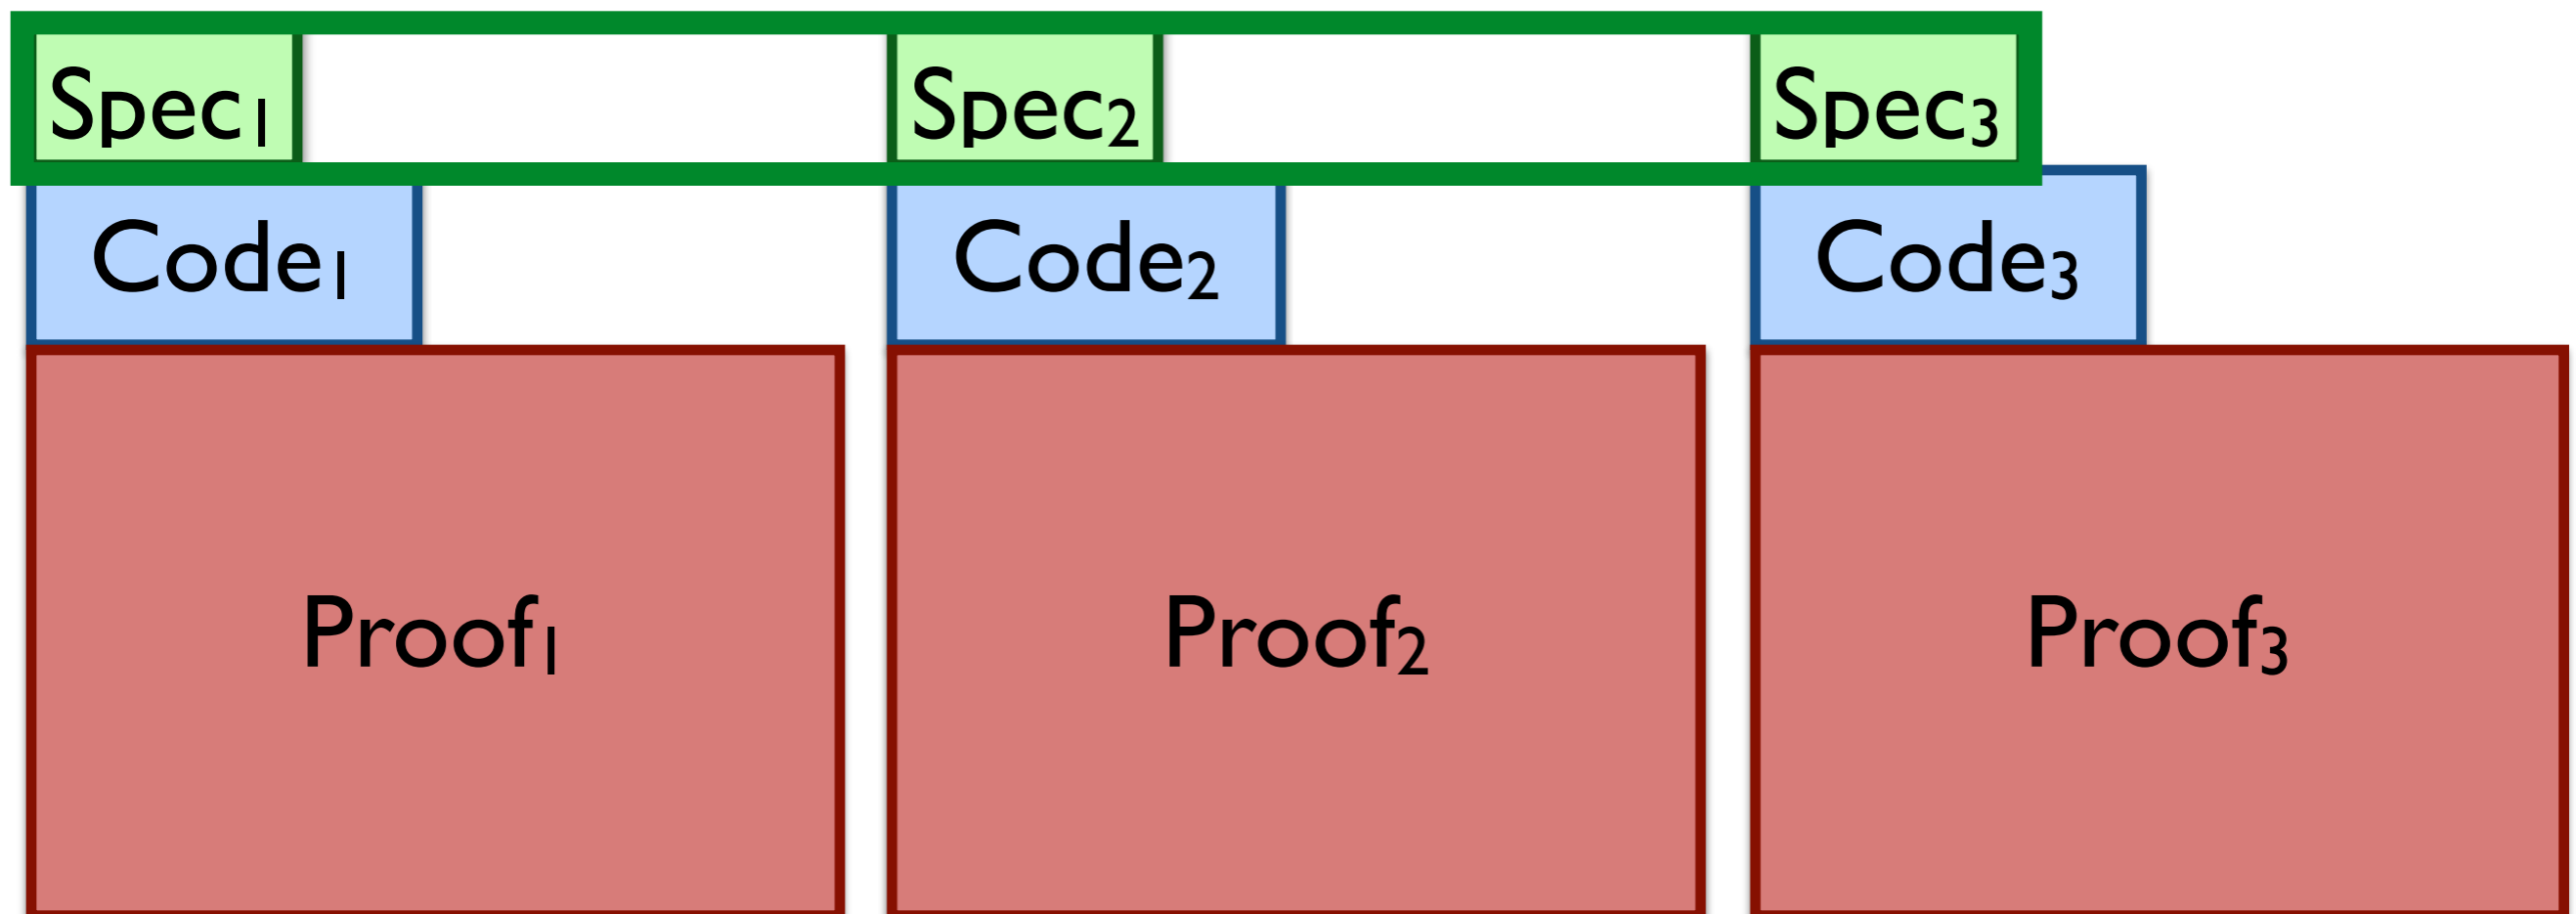


Single application domain



Single application domain

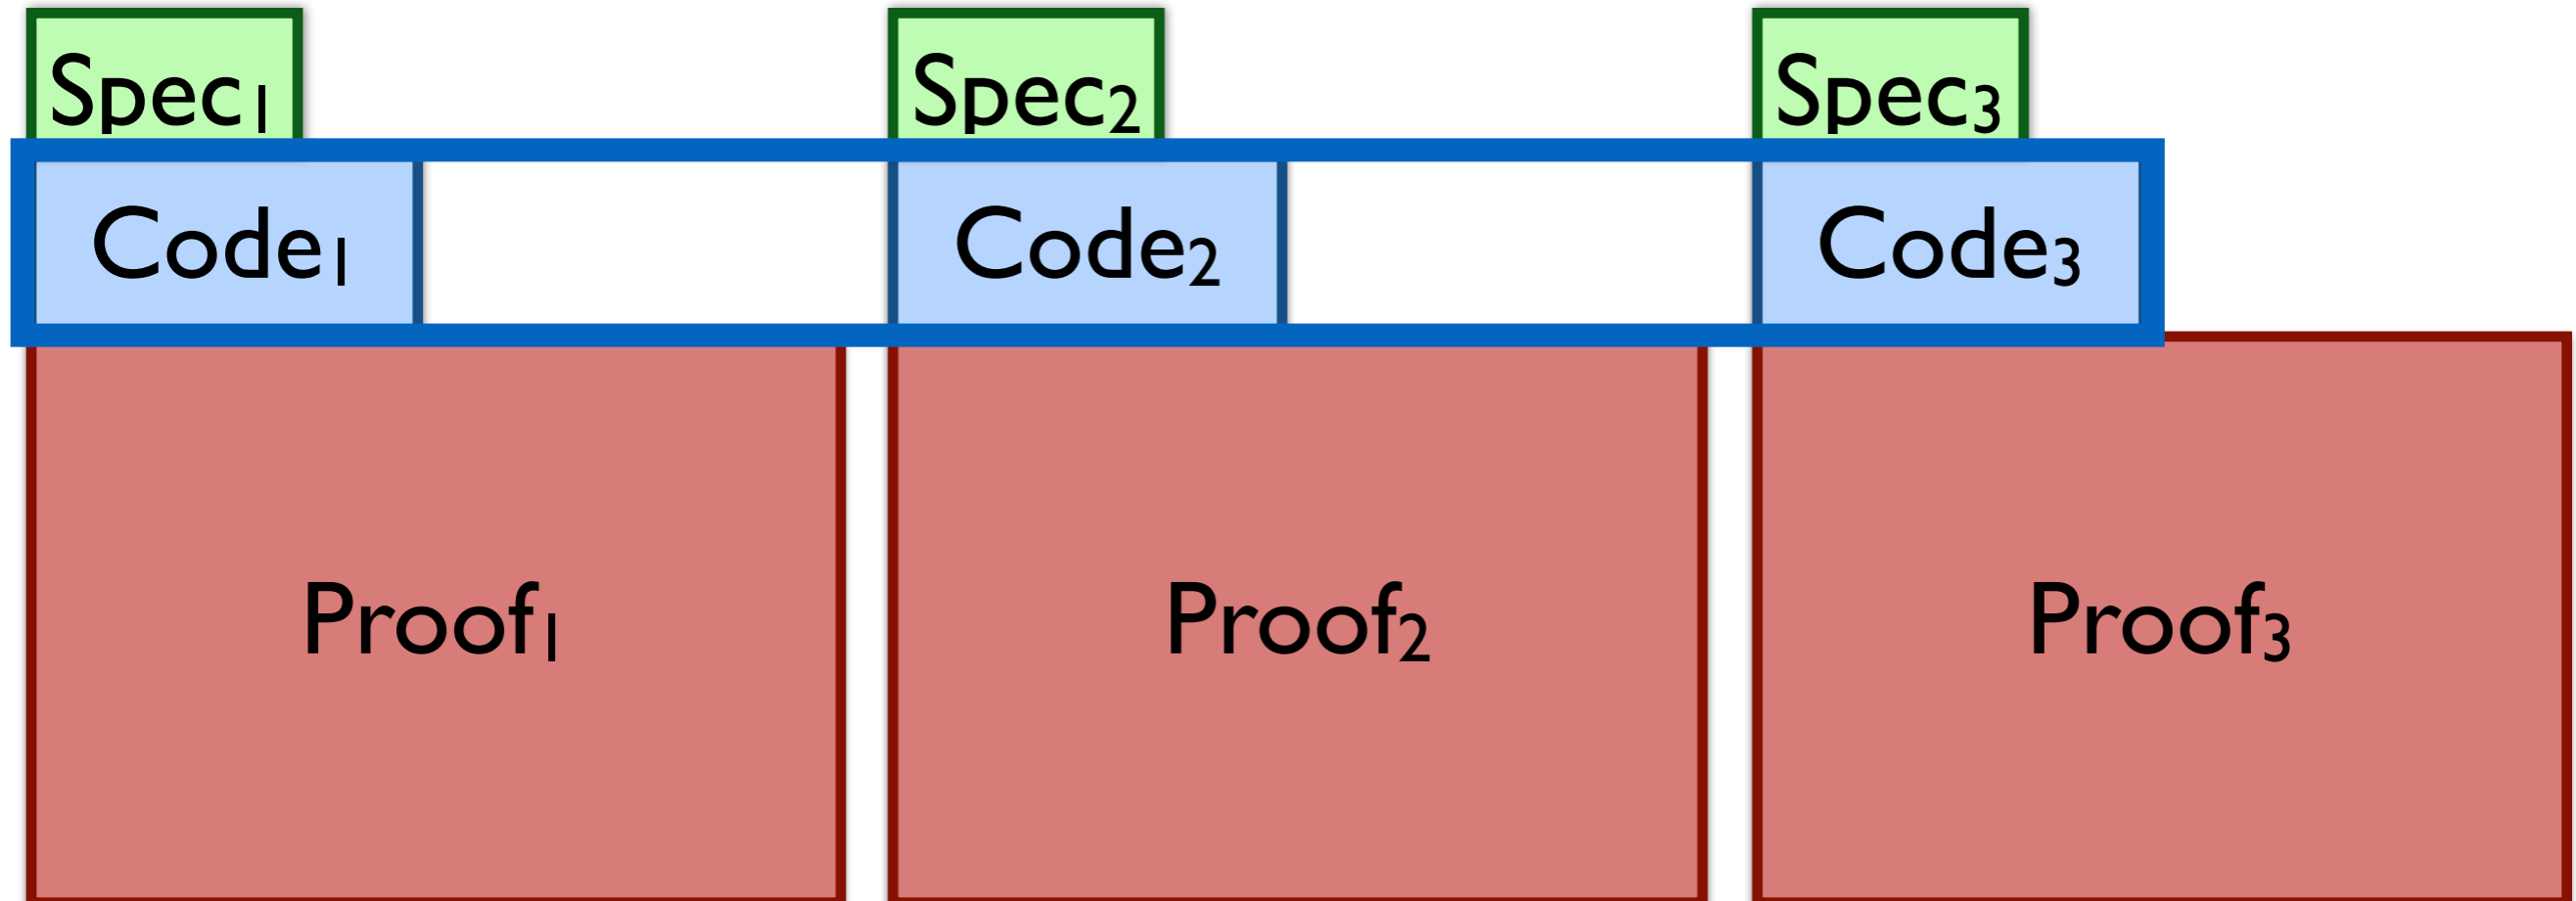
Similar properties



Single application domain

Similar properties

Similar architecture



Single application domain

Similar properties

Similar architecture

Similar reasoning

Spec₁

Code₁

Spec₂

Code₂

Spec₃

Code₃

Proof₁

Proof₂

Proof₃

Single application domain

DSL for Specs

Similar architecture

Similar reasoning

Spec₁

Code₁

Proof₁

Spec₂

Code₂

Proof₂

Spec₃

Code₃

Proof₃

Single application domain

DSL for Specs

DSL for Code

Similar reasoning

Spec₁

Code₁

Proof₁

Spec₂

Code₂

Proof₂

Spec₃

Code₃

Proof₃

Single application domain

DSL for Specs

DSL for Code

Proof Automation

Spec₁

Code₁

Proof₁

Spec₂


Code₂

Proof₂

Spec₃

Code₃

Proof₃



Single application domain

DSL for Specs

DSL for Code

Spec₁

Code₁

Spec₂

Code₂

Spec₃

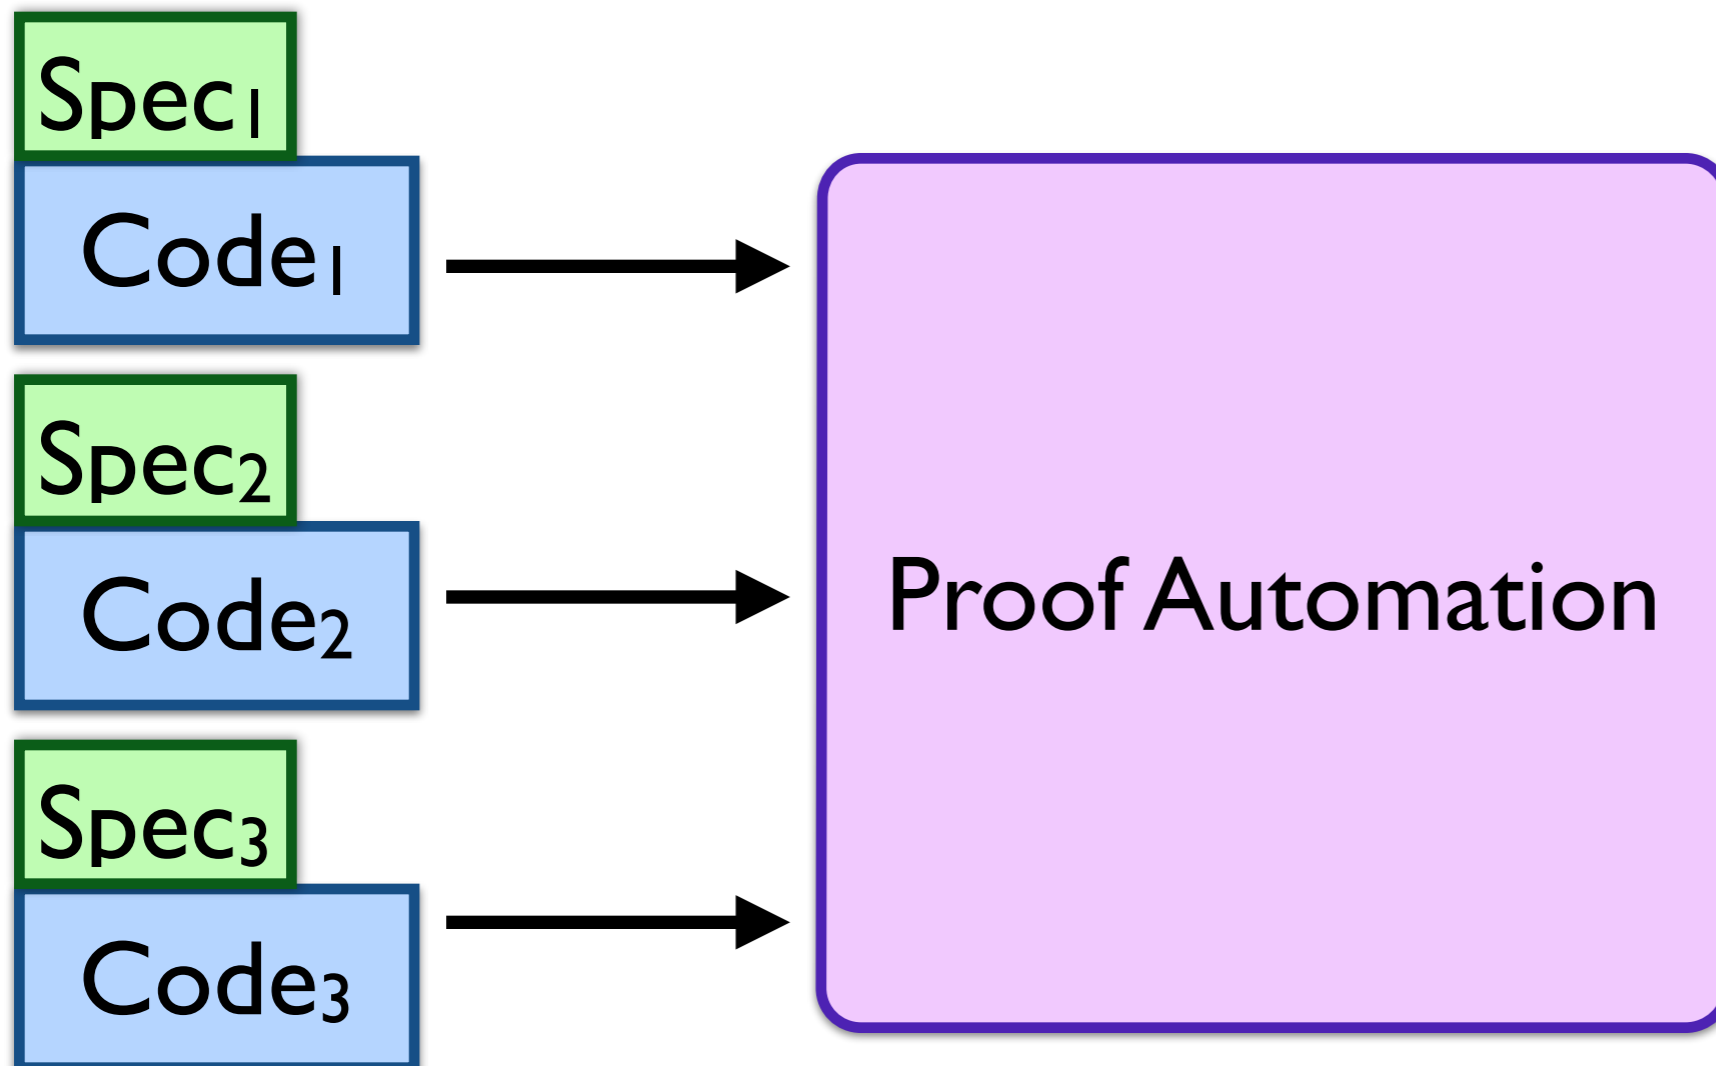
Code₃

Proof Automation

Single application domain

DSL for Specs

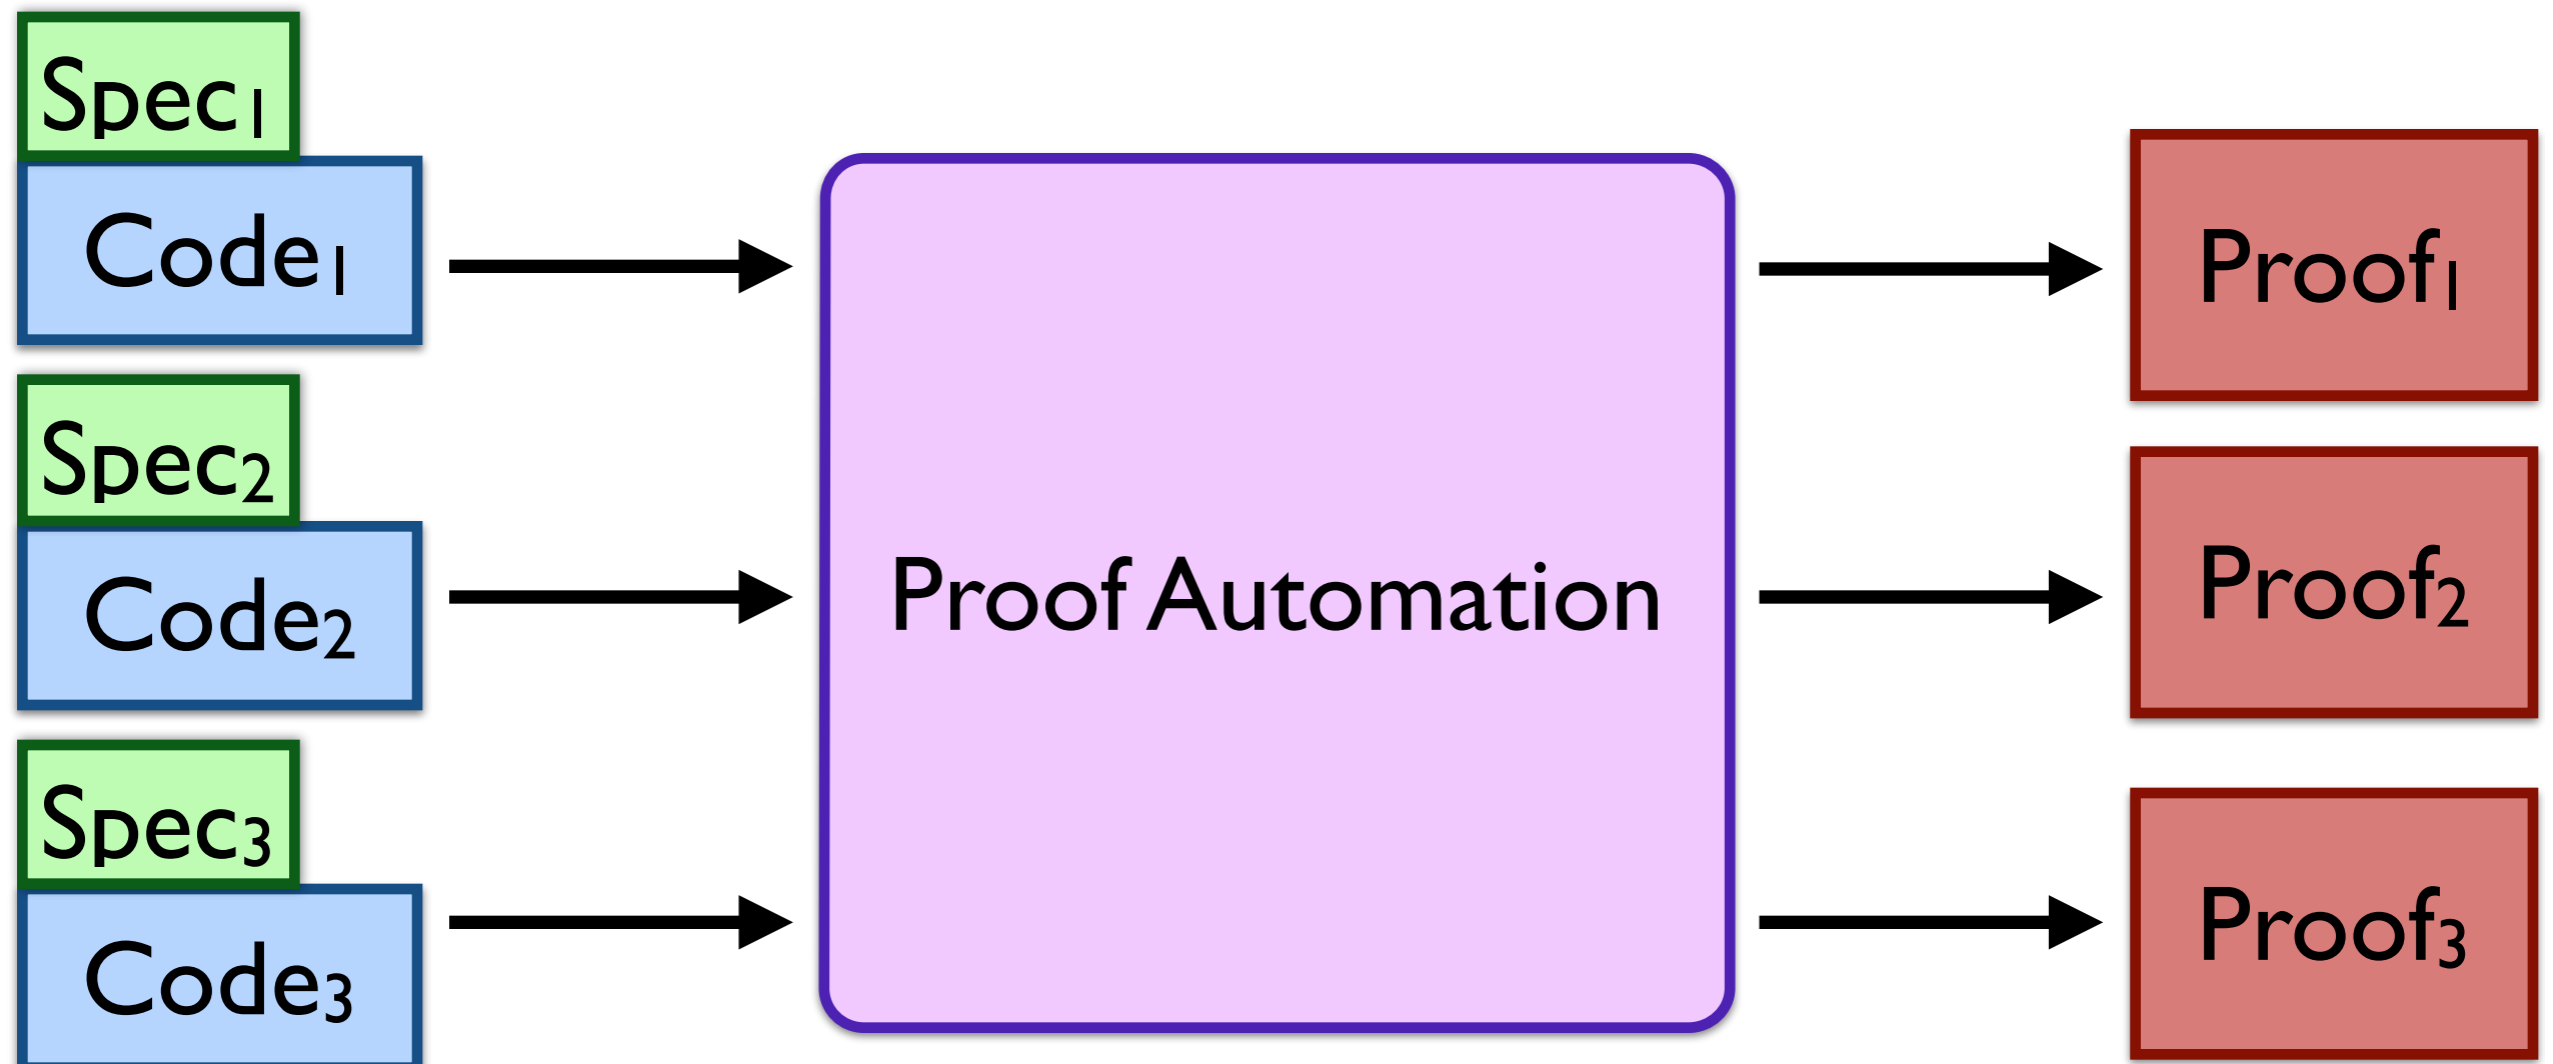
DSL for Code



Single application domain

DSL for Specs

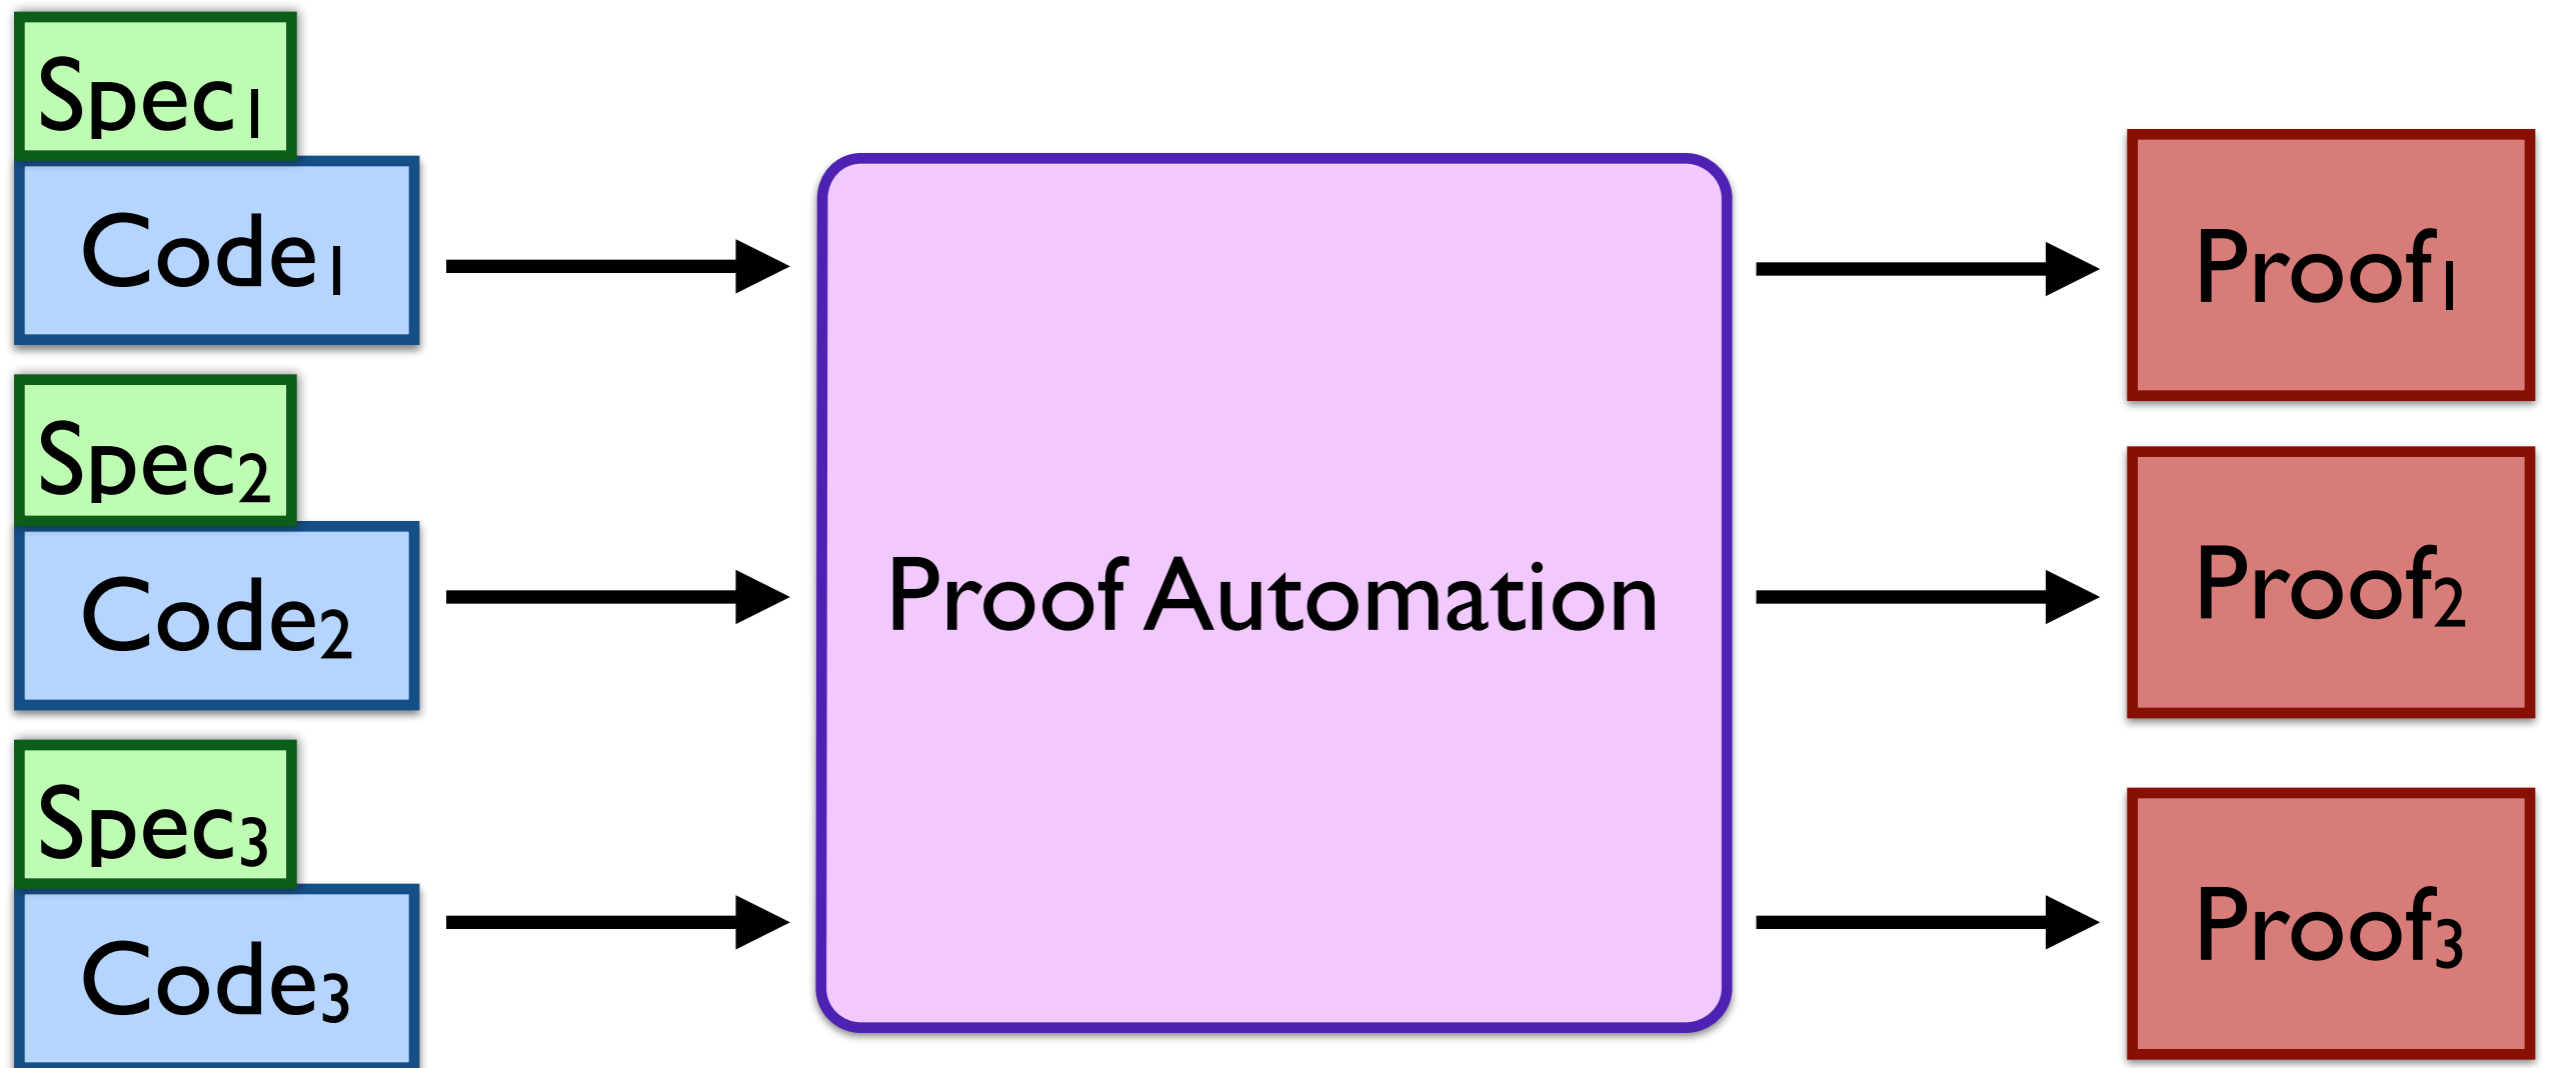
DSL for Code



Reactive systems

DSL for Specs

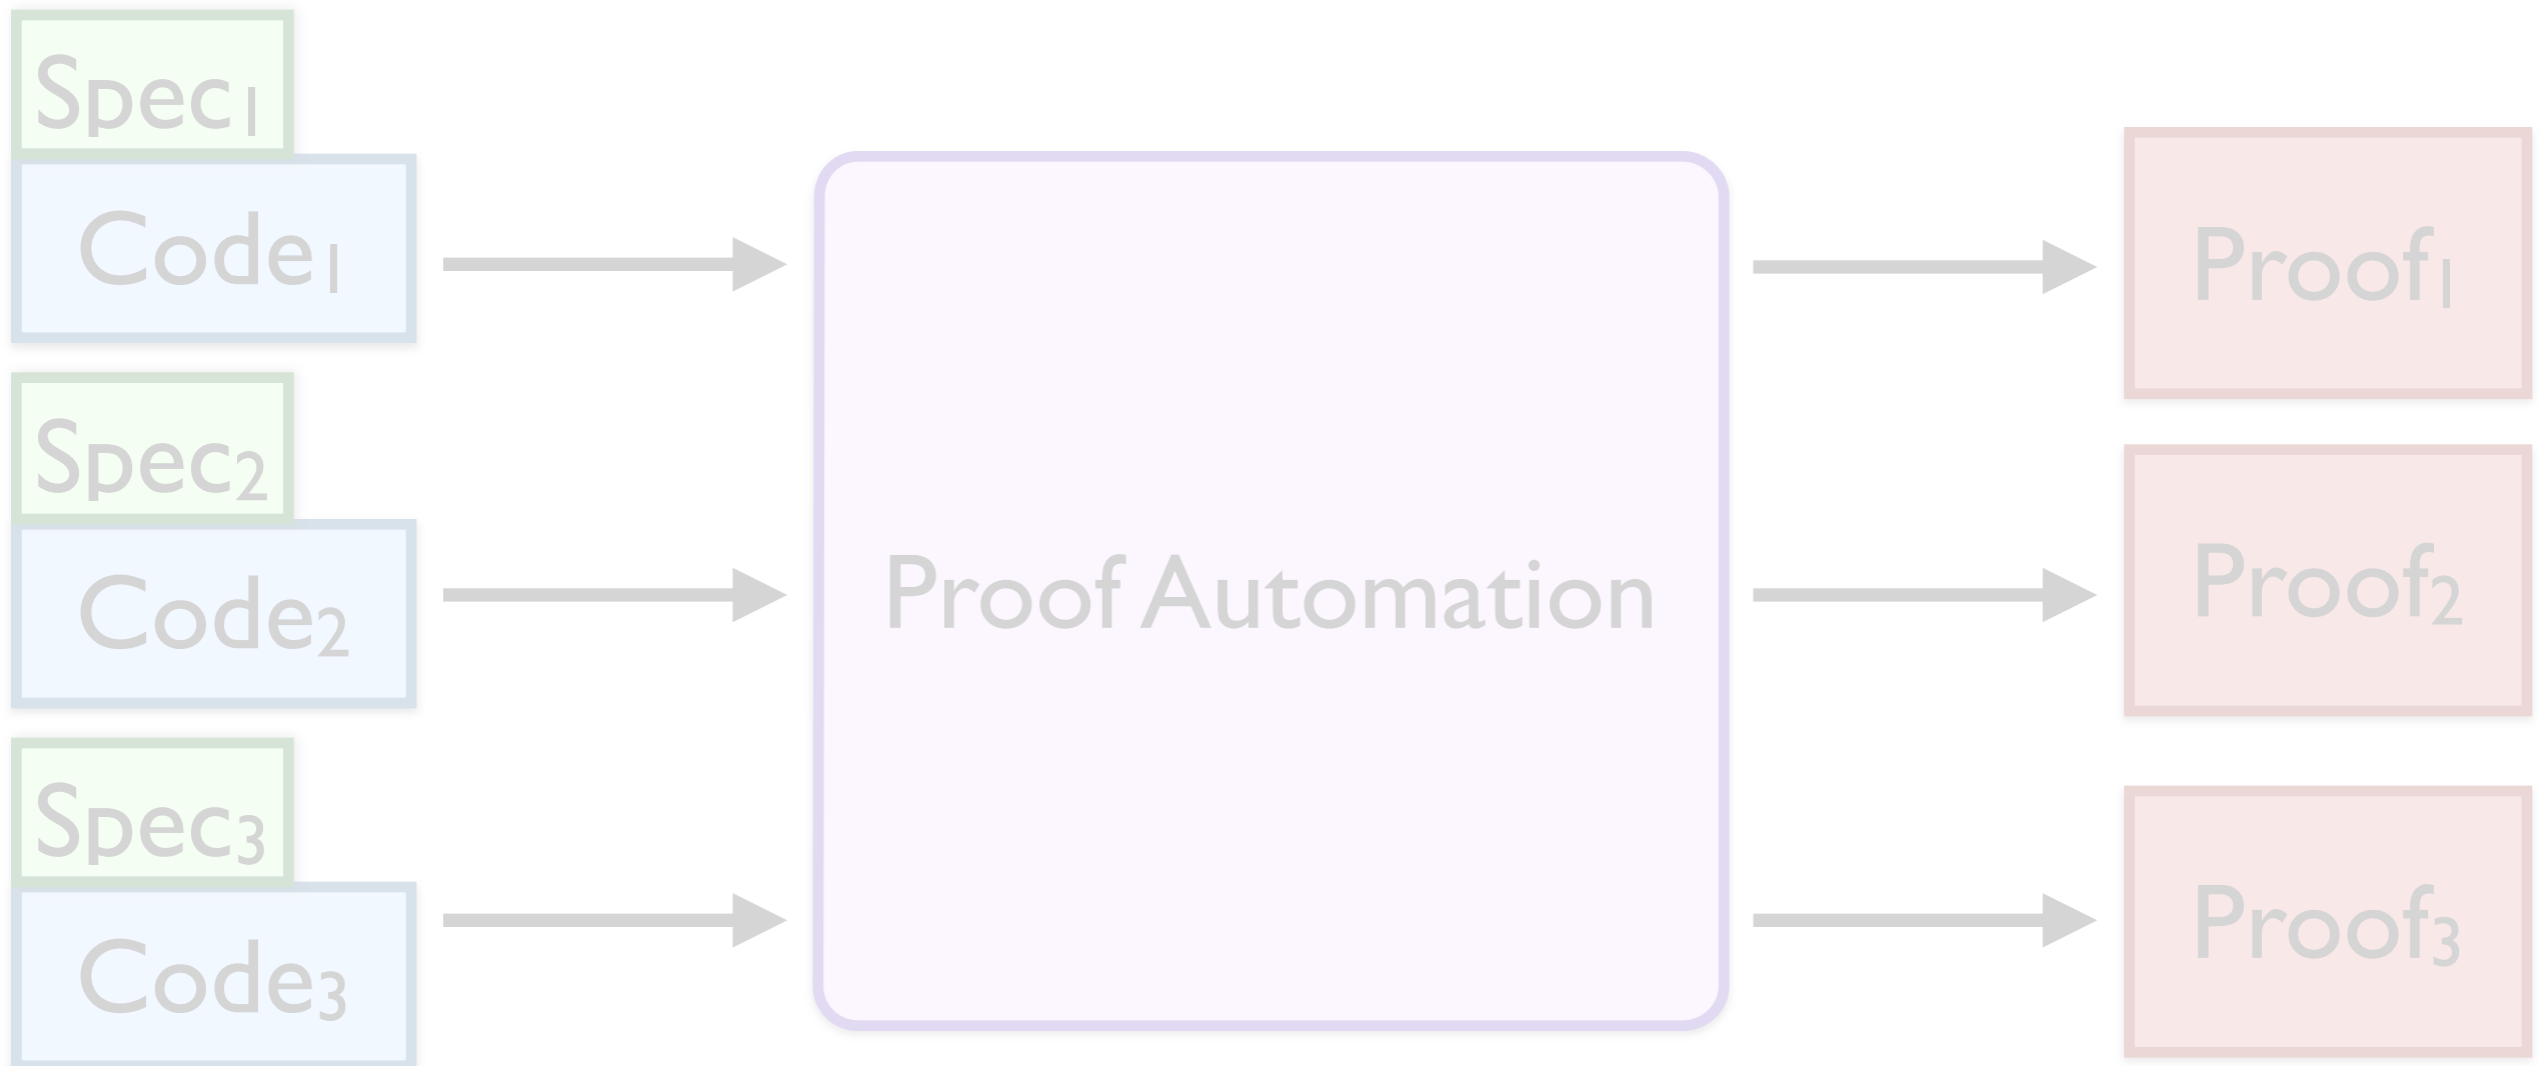
DSL for Code



Reactive systems

DSL for Specs

DSL for Code



Reactive systems

REFLEX

Sp

C

Sp

C

Sp

Code₃

of₁

of₂

Proof₃

Reactive systems

REFLEX

No manual proofs,

Sp

C

Sp

C

Sp

Code₃

of₁

of₂

Proof₃

Reactive systems

REFLEX

No manual proofs,
yet proof assistant guarantee.

Sp

C

Sp

C

Sp

Code₃

of₁

of₂

Proof₃

Reactive systems

REFLEX

No manual proofs,
yet proof assistant guarantee.

Automation incomplete,

Sp

C

Sp

C

Sp

Code₃

of₁

of₂

Proof₃

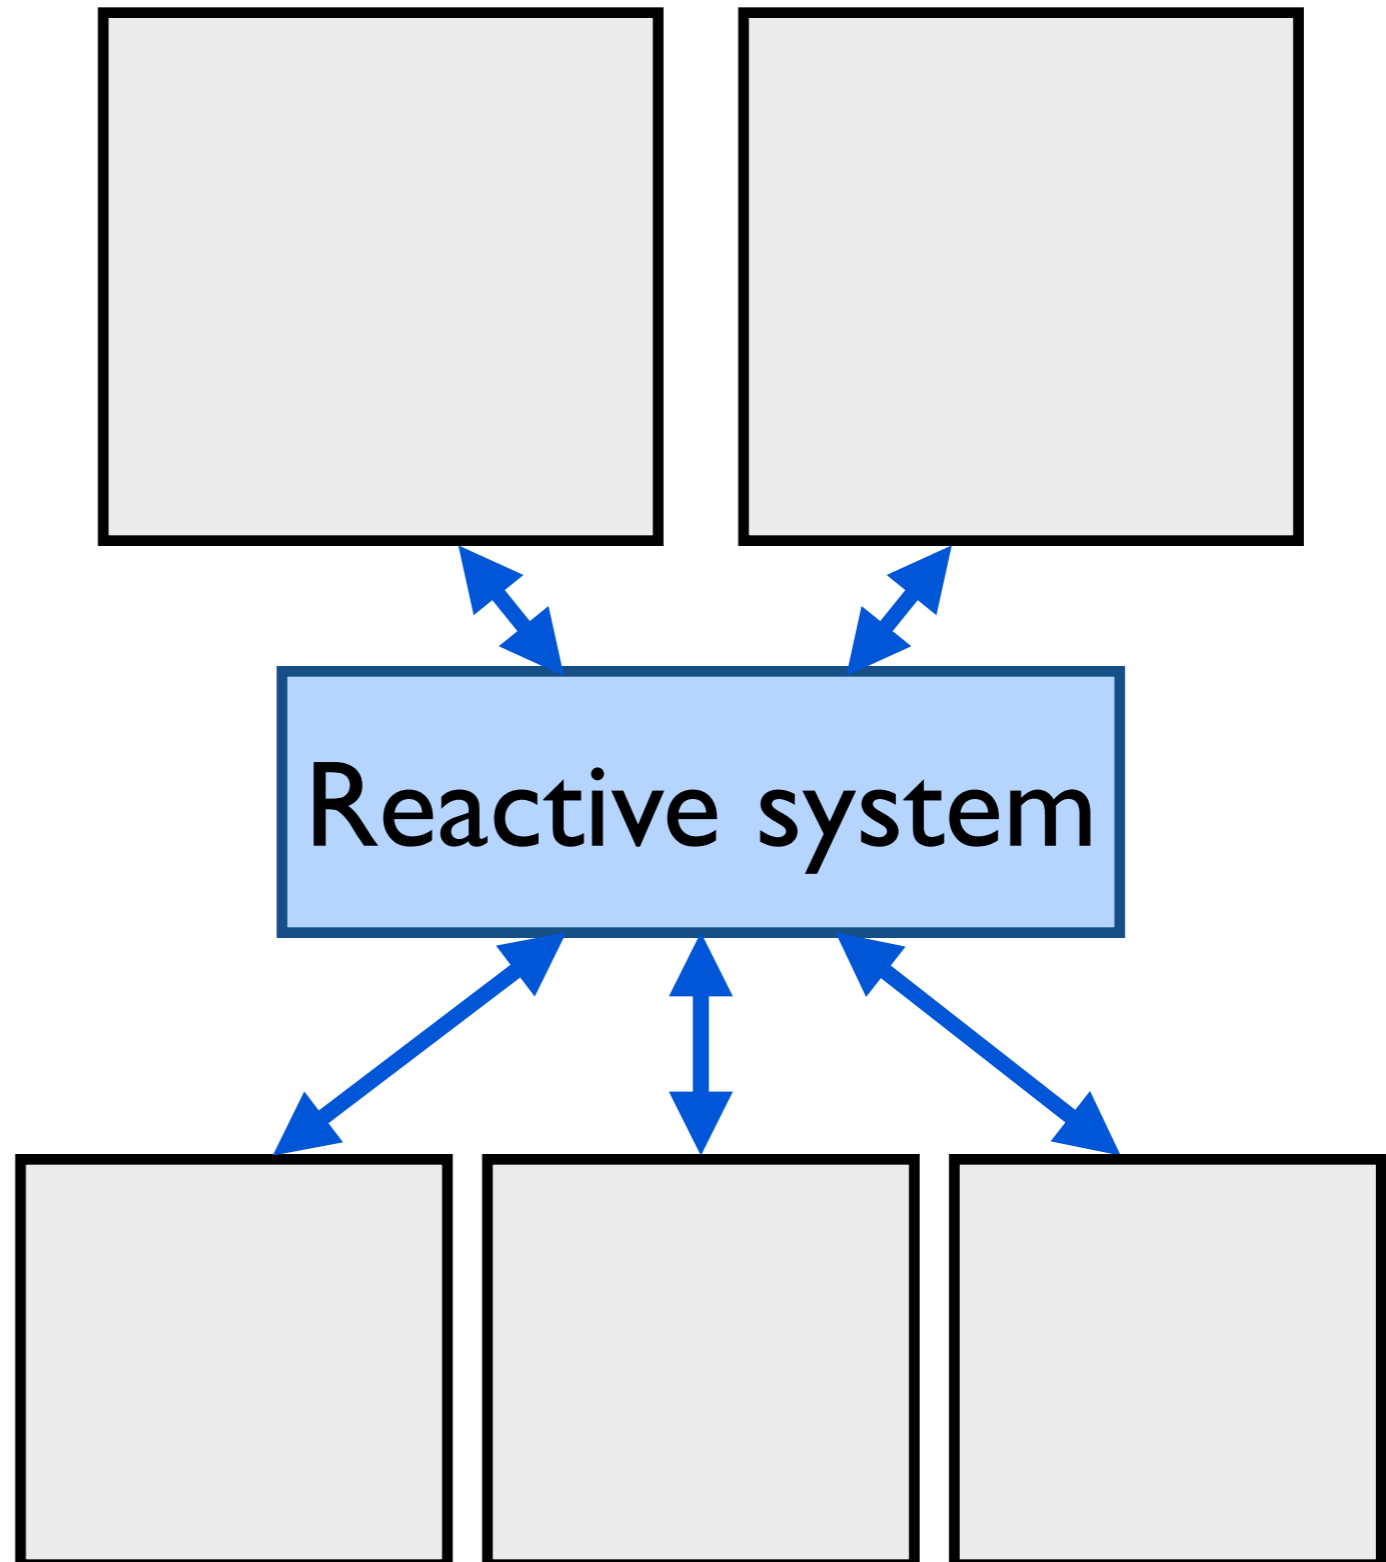
Reactive systems

REFLEX

No manual proofs,
yet proof assistant guarantee.


Automation incomplete,
but verified browser, ssh, web server.

Reactive systems




Reactive systems

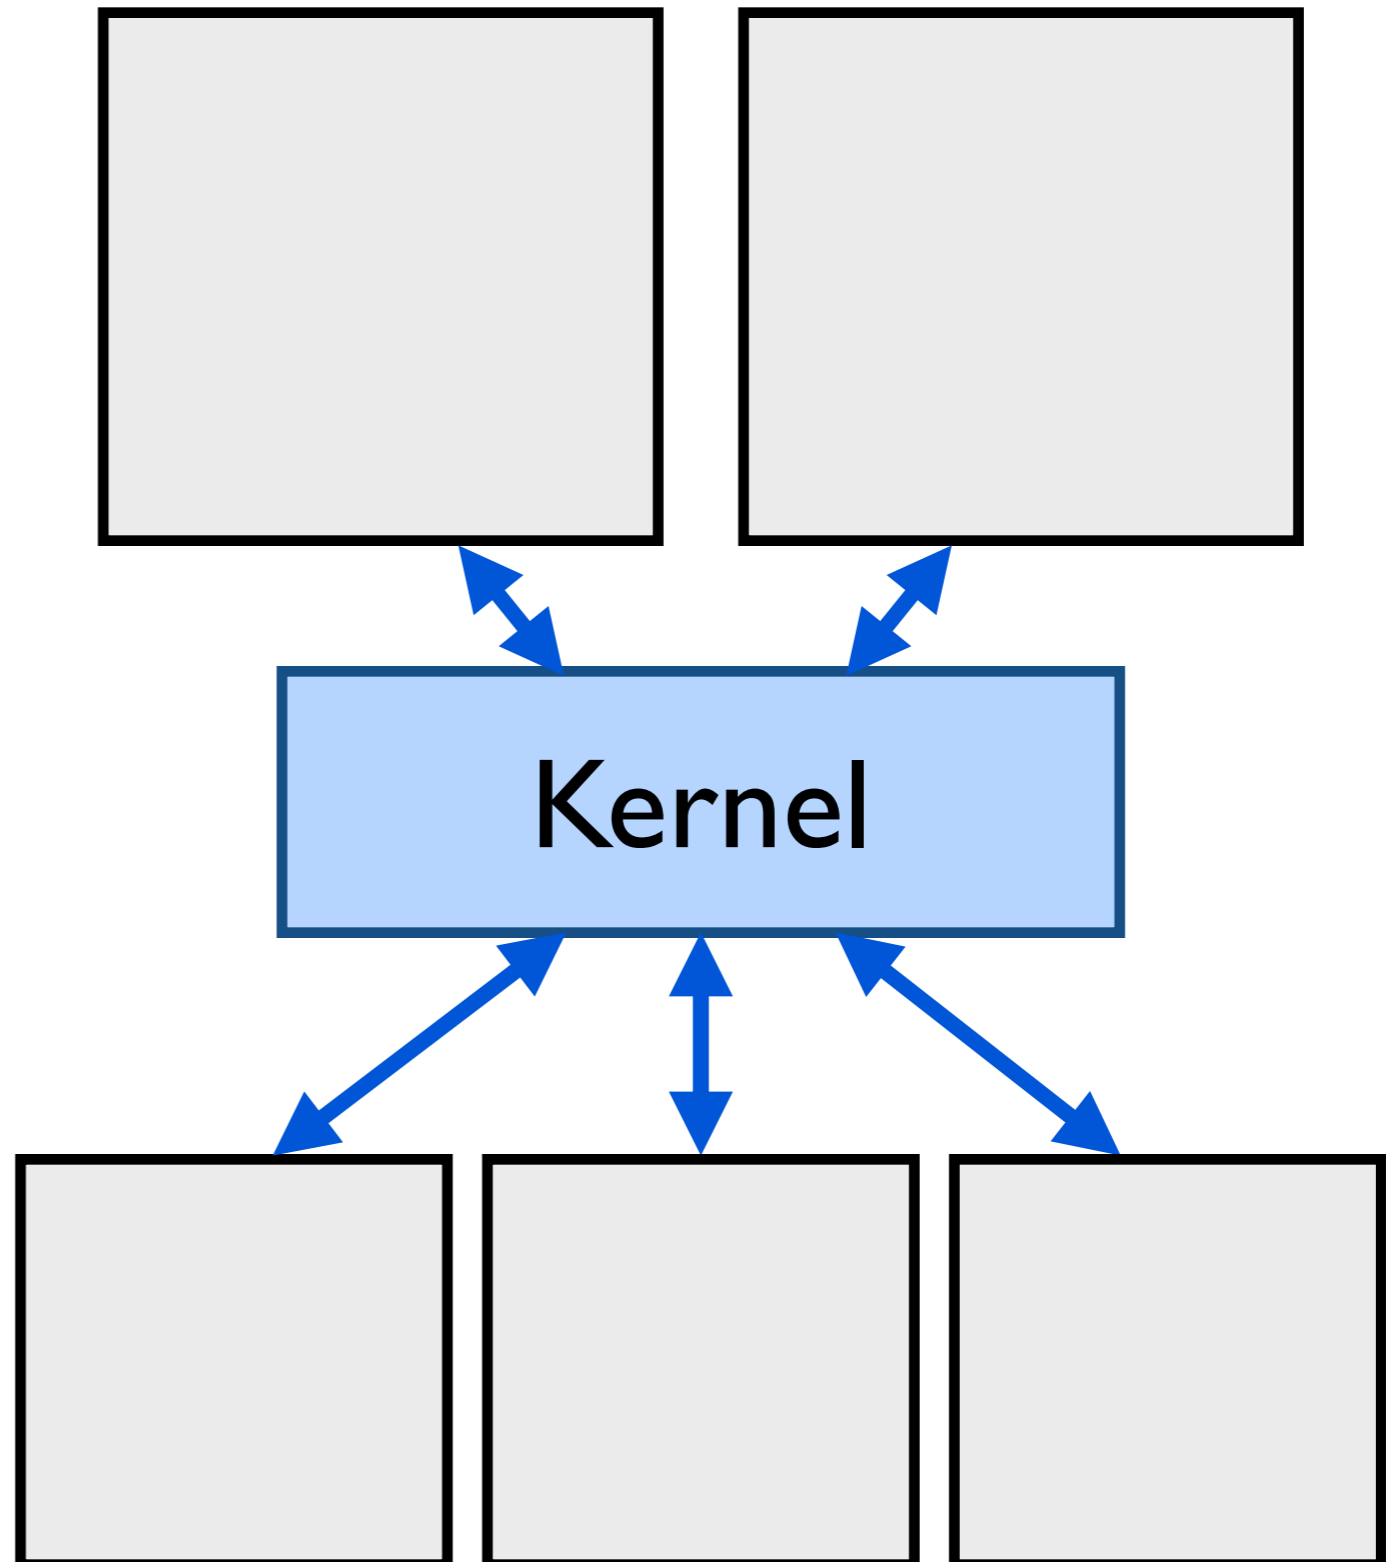
Continuously read messages from components and send messages to components




Reactive systems



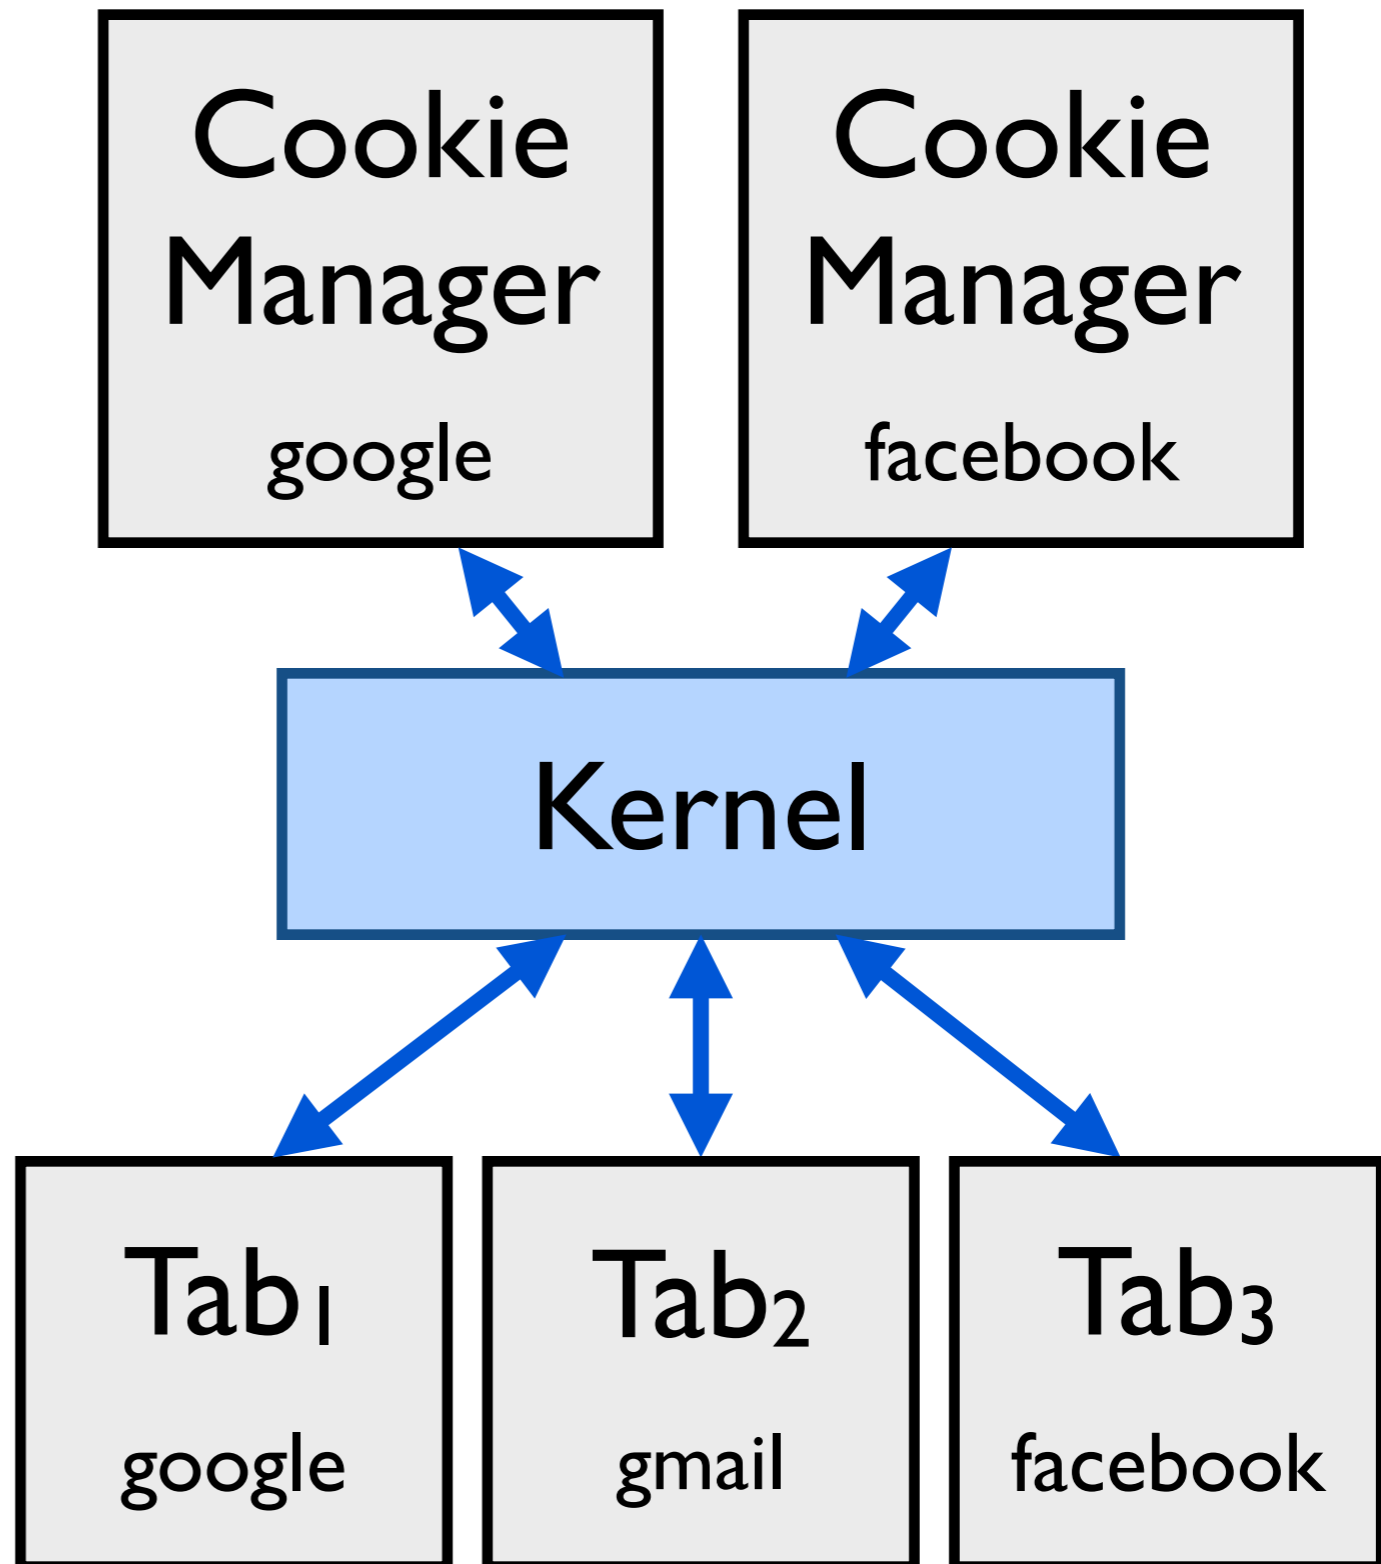
Example: Web browser kernel



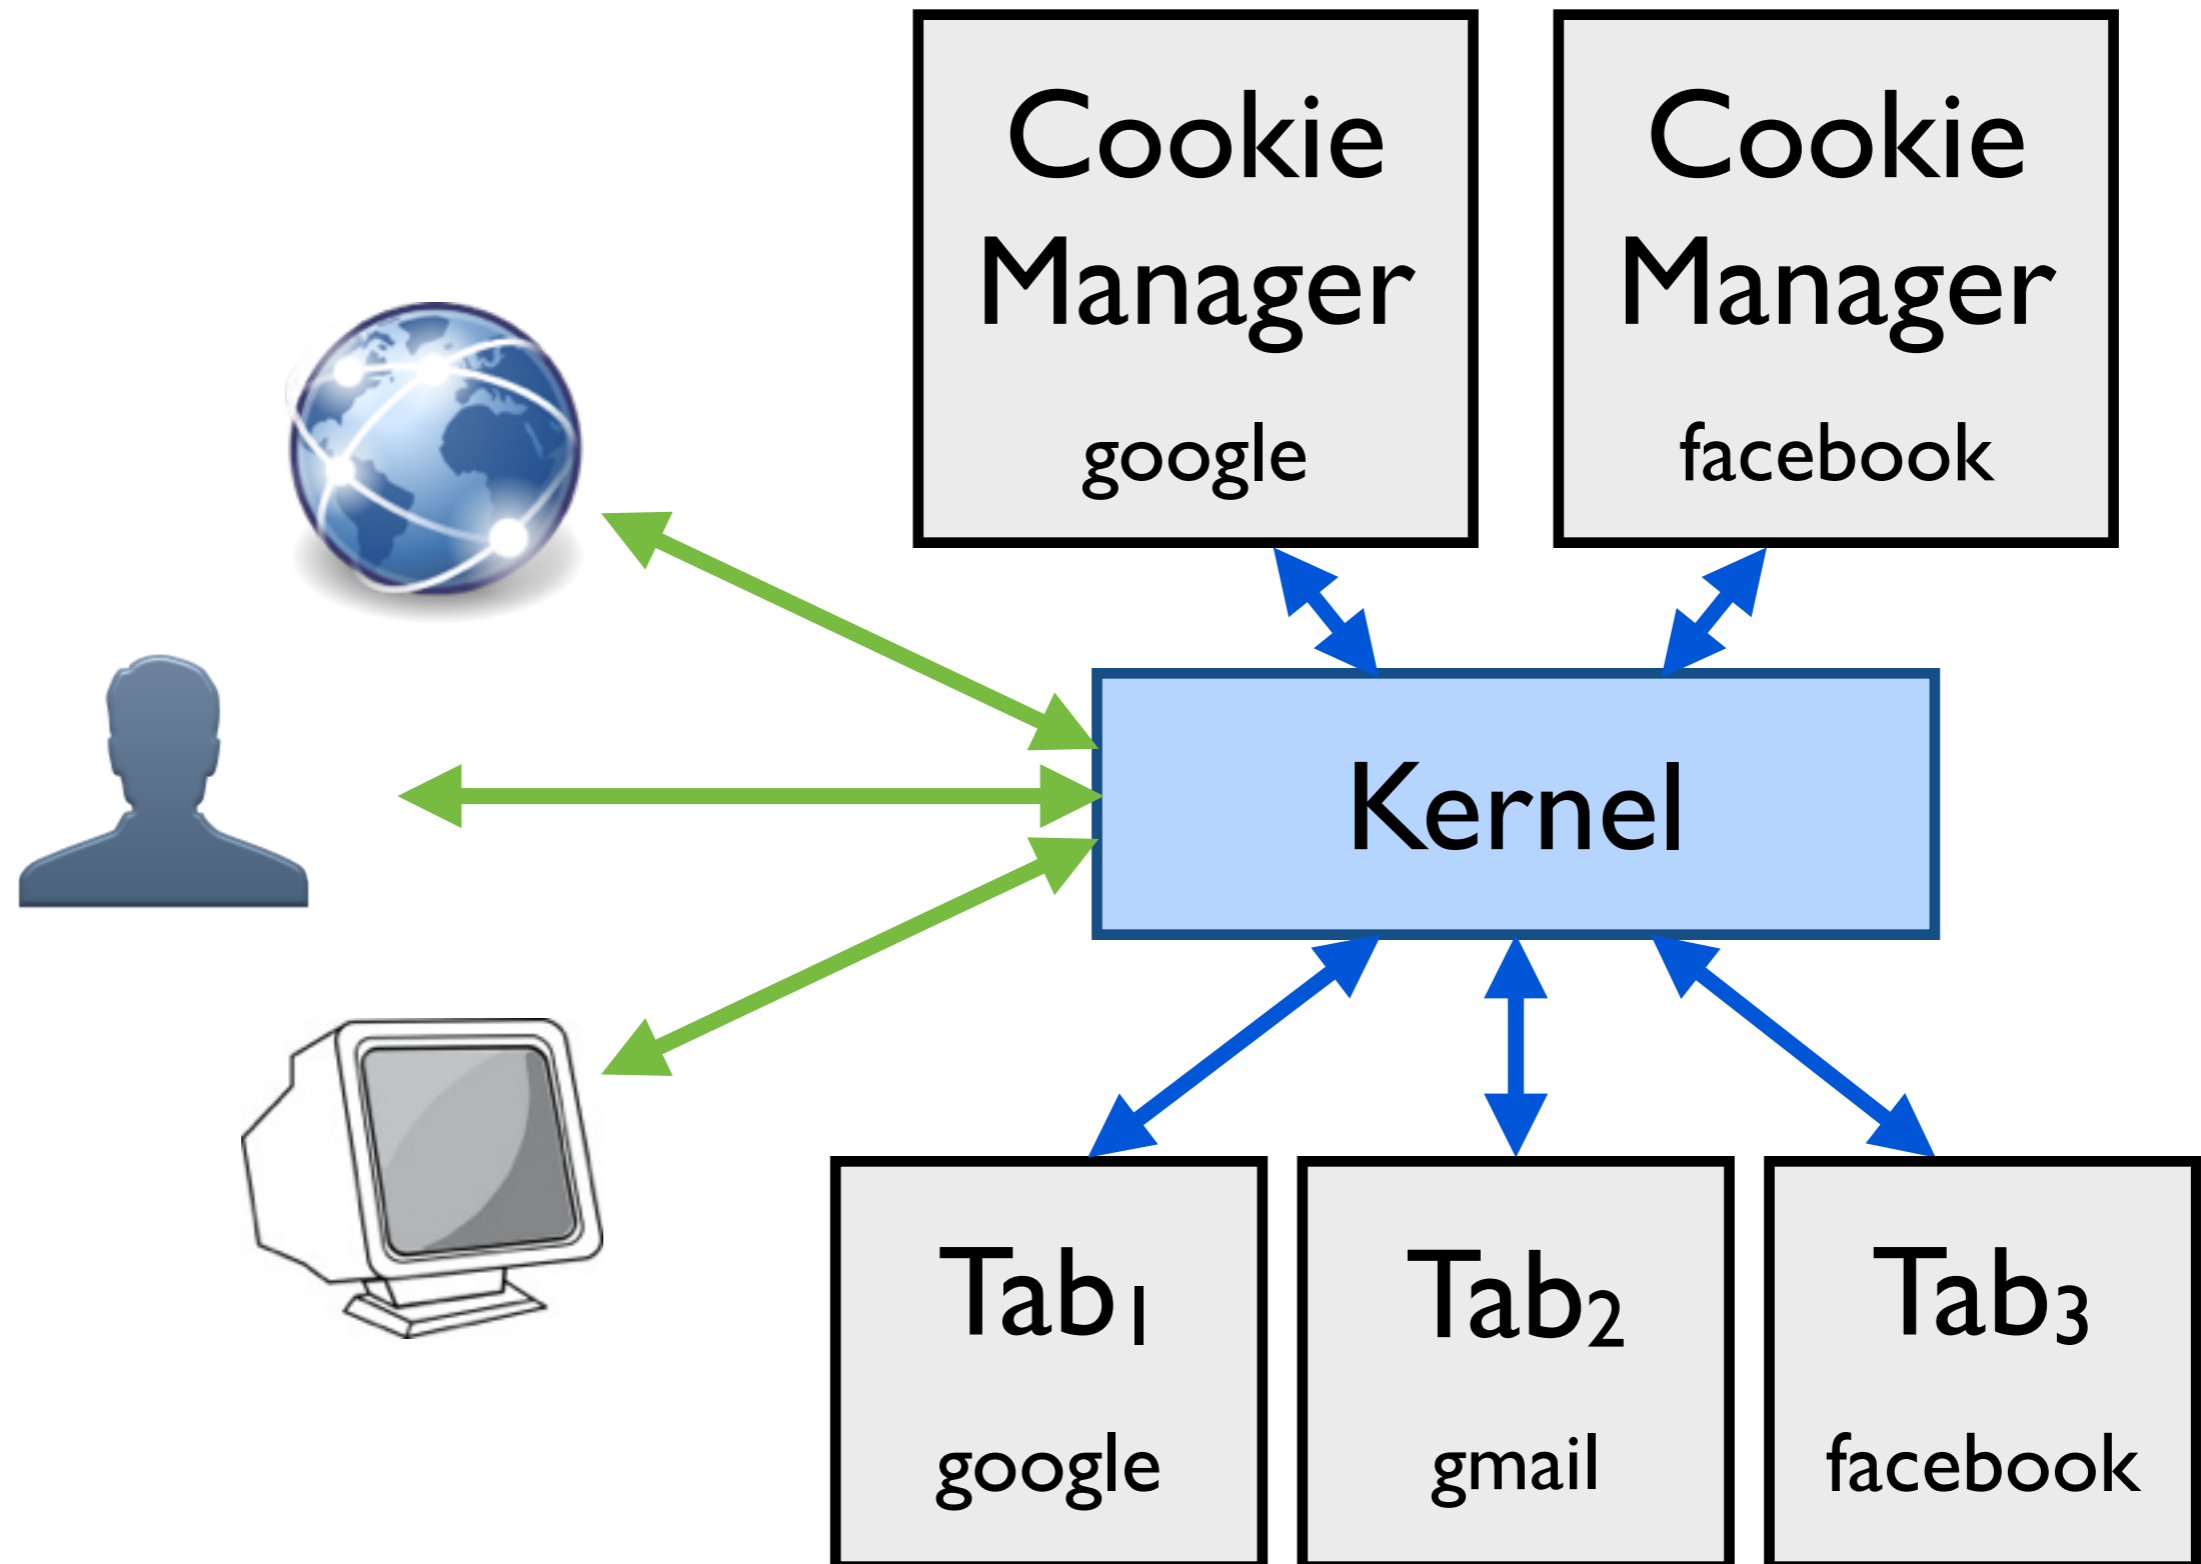
Example: Web browser kernel



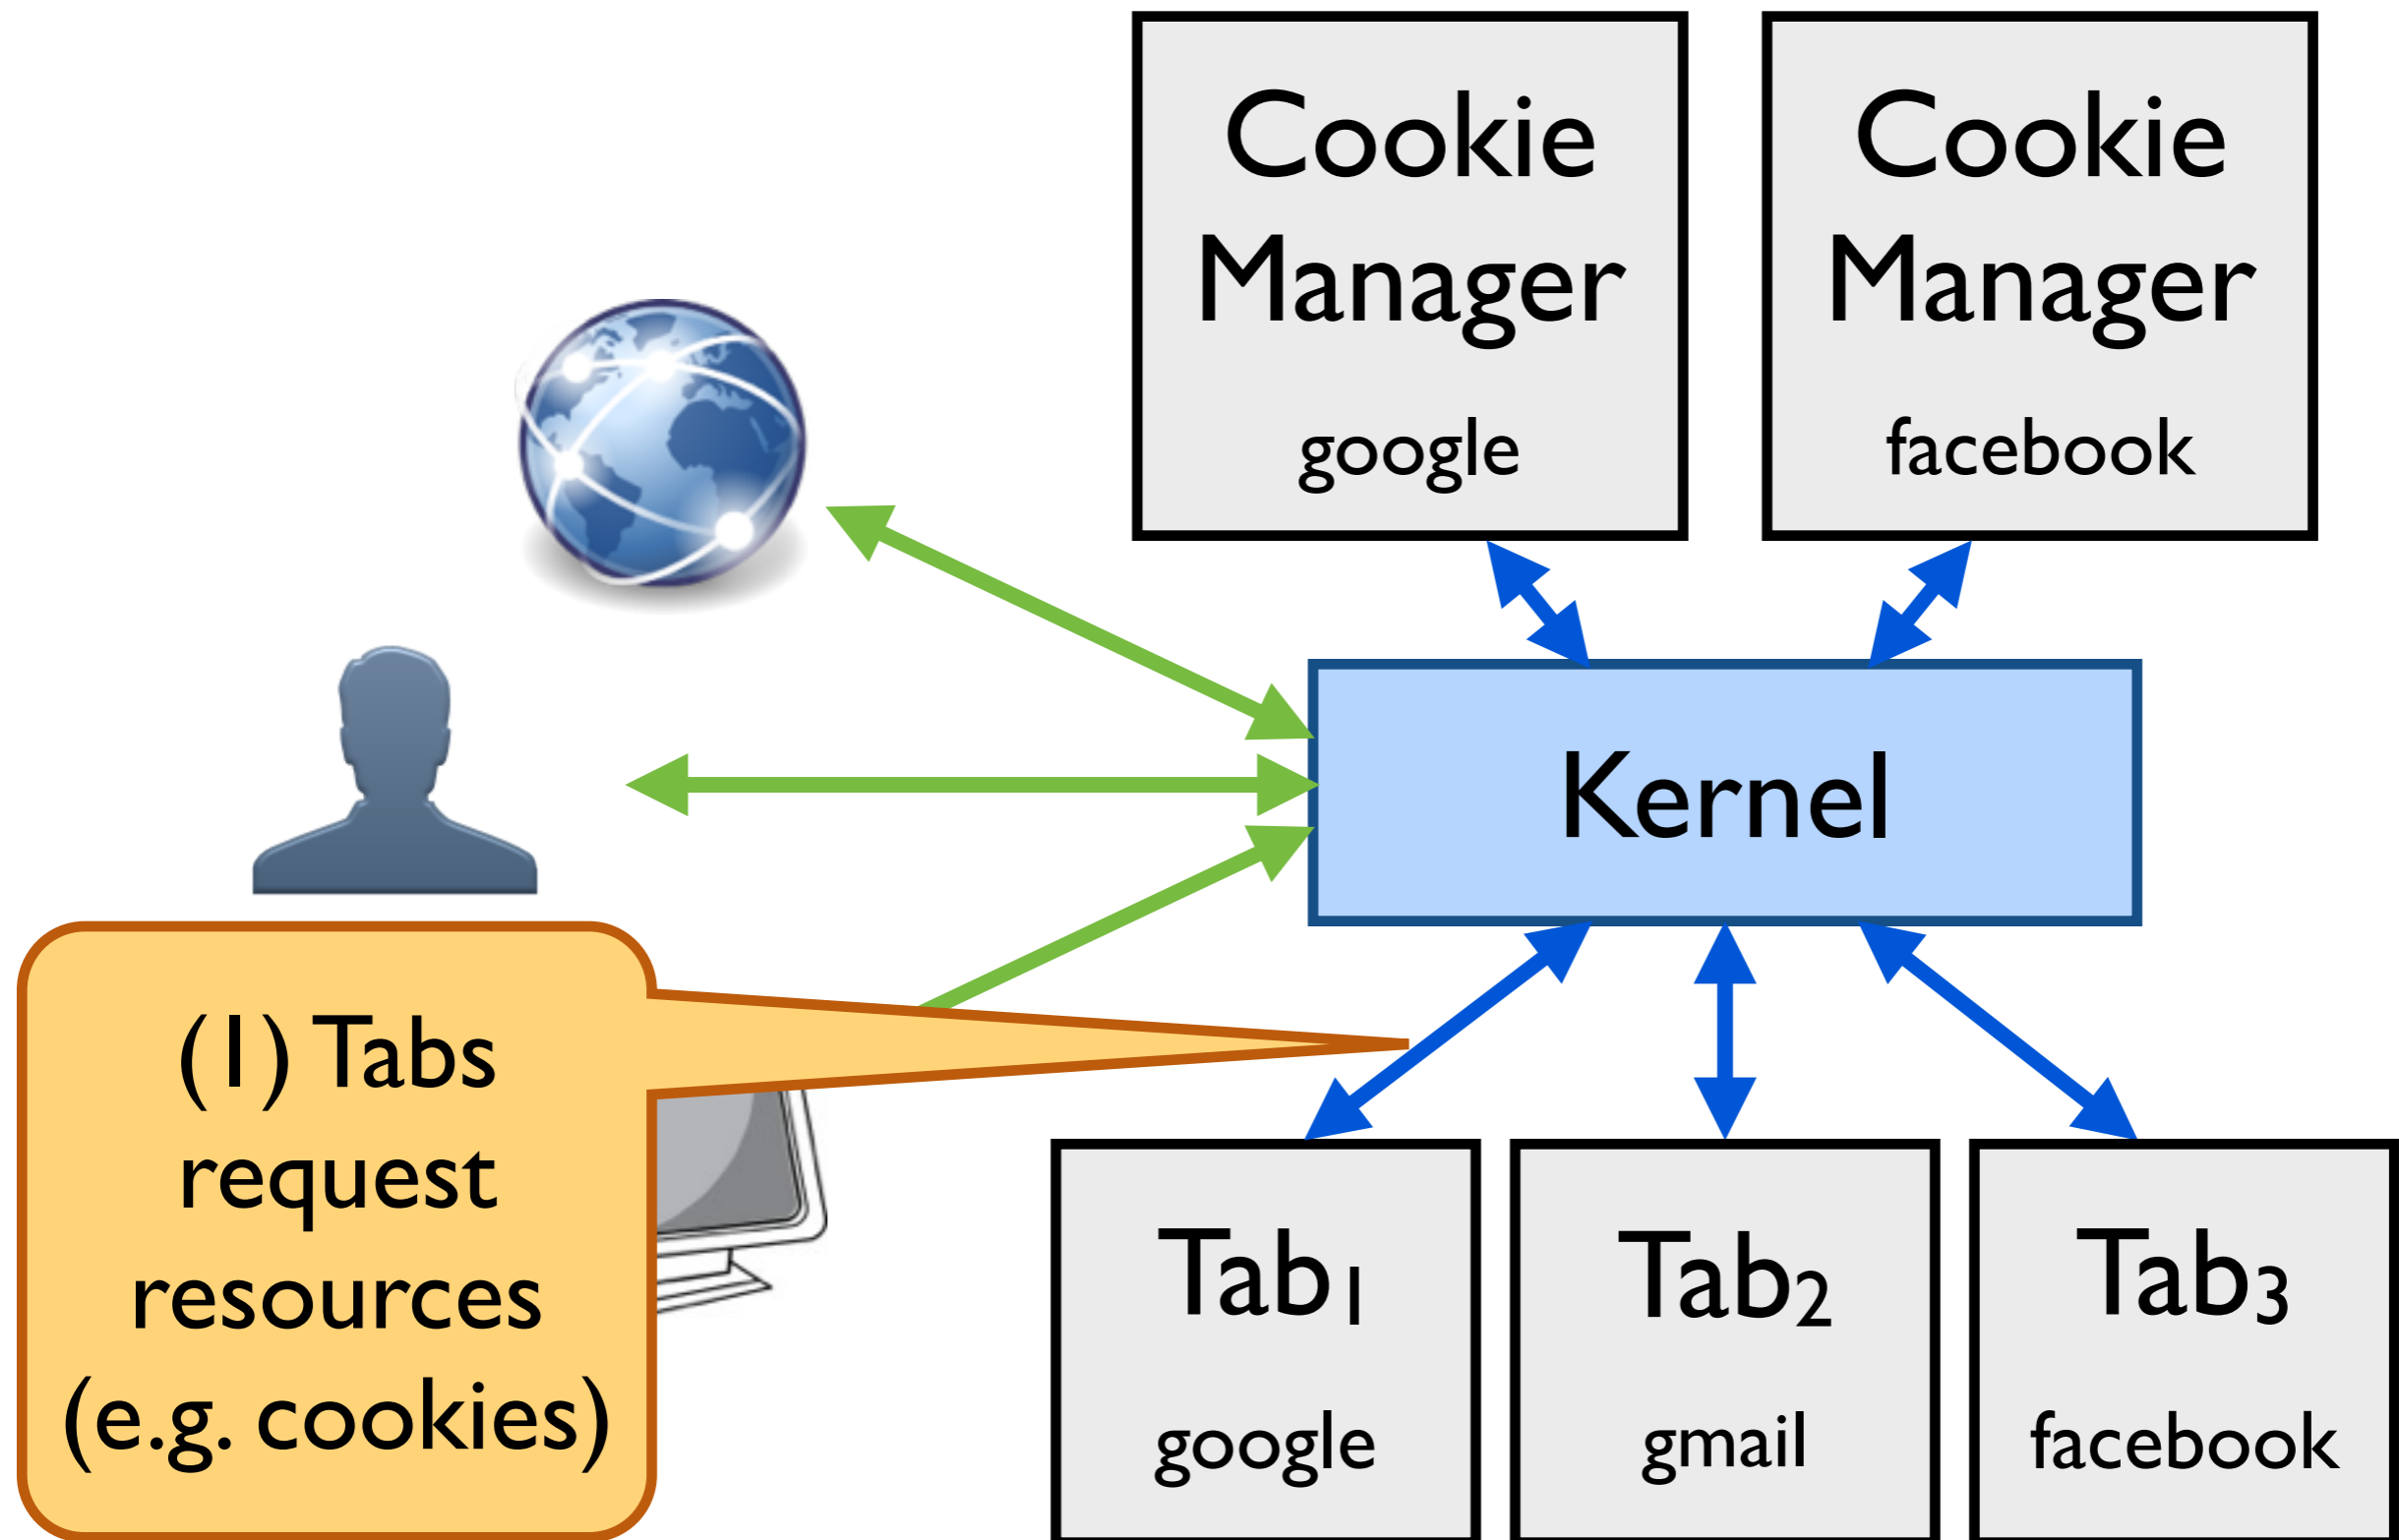
Example: Web browser kernel



Example: Web browser kernel



Example: Web browser kernel



Example: Web browser kernel

(2) Kernel grants access subject to access controls (e.g. domain checks)

Cookie Manager
google

Cookie Manager
facebook

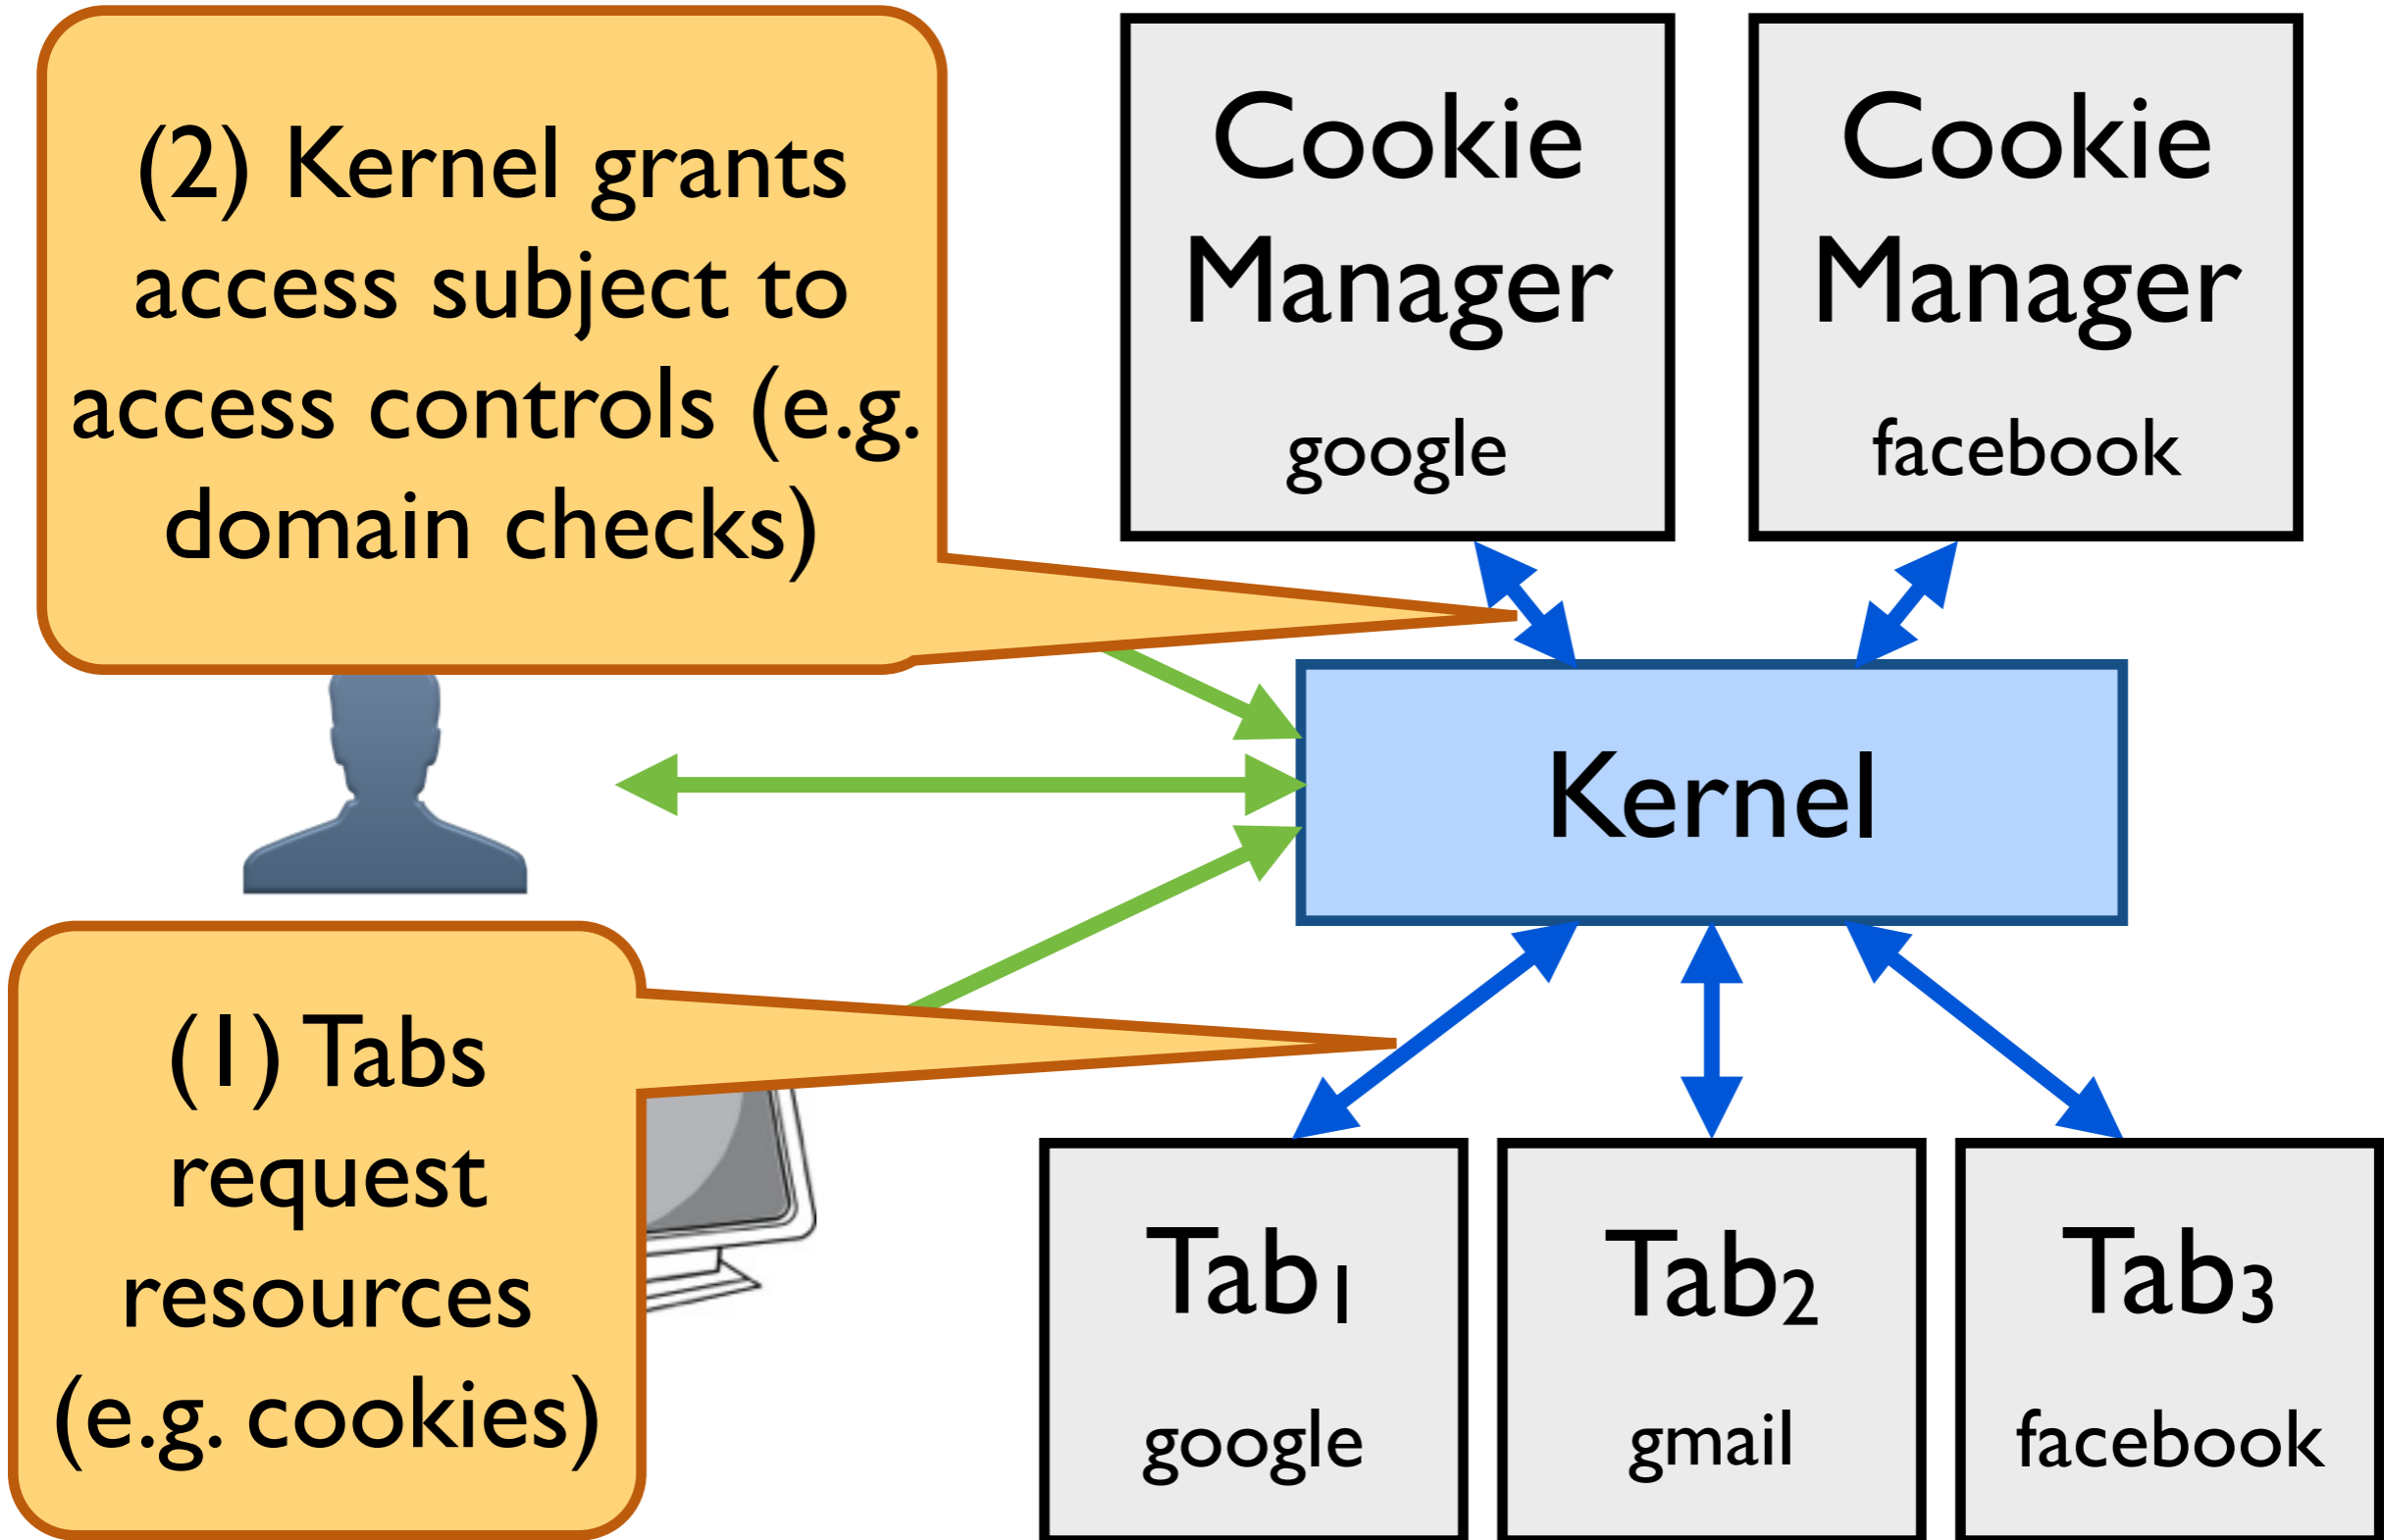
Kernel

(1) Tabs request resources (e.g. cookies)

Tab₁
google

Tab₂
gmail

Tab₃
facebook



Example: Web browser kernel



Example: Web browser kernel

```
Components = Tab | CookieMgr | ...
```



Types of
components

Example: Web browser kernel

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```



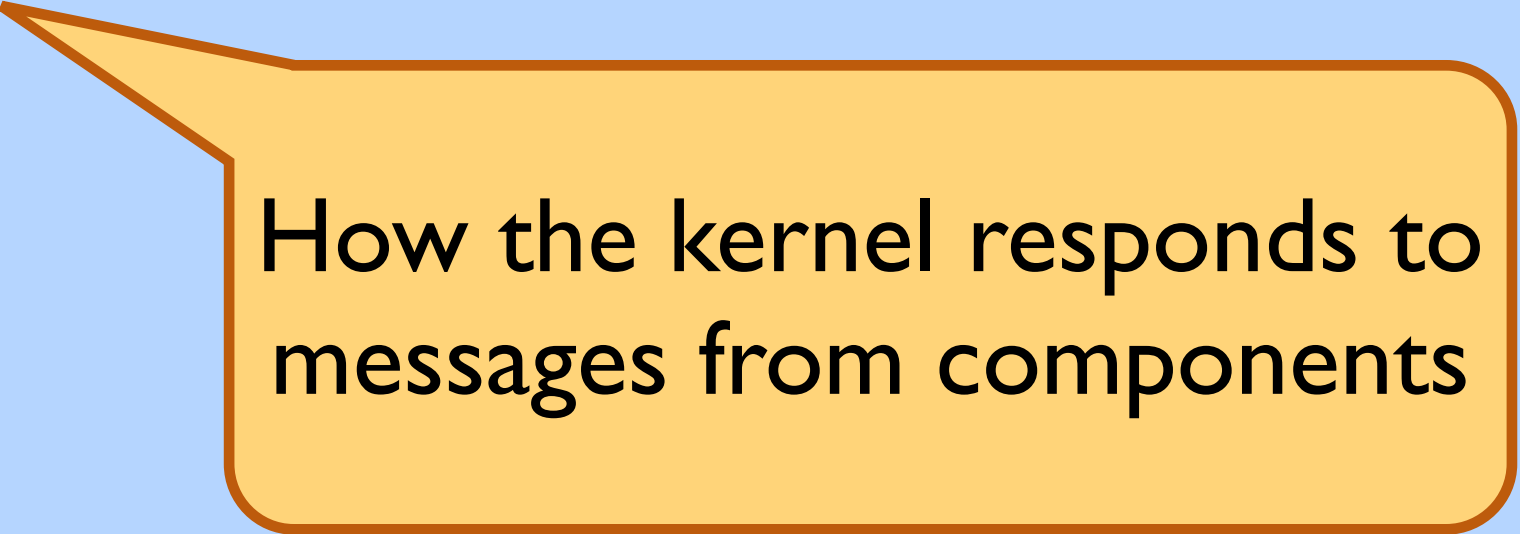
Types of messages

Example: Web browser kernel

`Components = Tab | CookieMgr | ...`

`Messages = CookieSet | CookieGet | ...`

`Handlers:`



How the kernel responds to messages from components

Example: Web browser kernel

`Components = Tab | CookieMgr | ...`

`Messages = CookieSet | CookieGet | ...`

Handlers:

`When [Tab t] sends CookieSet(c):`

When tab `t` sends the kernel a `CookieSet` message with payload `c`

Example: Web browser kernel

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):  
    cp <- find CookieMgr(t.domain)
```

Get existing cookie manager
with domain of `t` or spawn a
new one

Example: Web browser kernel

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Send the cookie `c` to the found cookie manager

Example: Web browser kernel

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
  ...
```

More handlers

Properties

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):  
    cp <- find CookieMgr(t.domain)  
    send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):  
    ...
```

Properties

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
  ...
```

Specify allowed behaviors

Properties

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
  ...
```

Specify allowed behaviors wrt sequence of system calls

Properties

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
  ...
```

Specify allowed behaviors wrt sequence of system calls

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls



Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls



Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far



Time

Properties

```
When [Tab t] sends CookieSet(c):  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far



Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far



Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far

`Recv(Tab, CookieSet(c))`

...

Time



Properties

```
When [Tab t] sends CookieSet(c):  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far

```
Recv(Tab, CookieSet(c))
```

...

Time



Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

The system
calls so far

Recv(Tab, CookieSet(c))

...

Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

Spawn CookieMgr(t.domain)

Recv(Tab, CookieSet(c))

...

The system
calls so far

Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

Spawn CookieMgr(t.domain)

Recv(Tab, CookieSet(c))

...

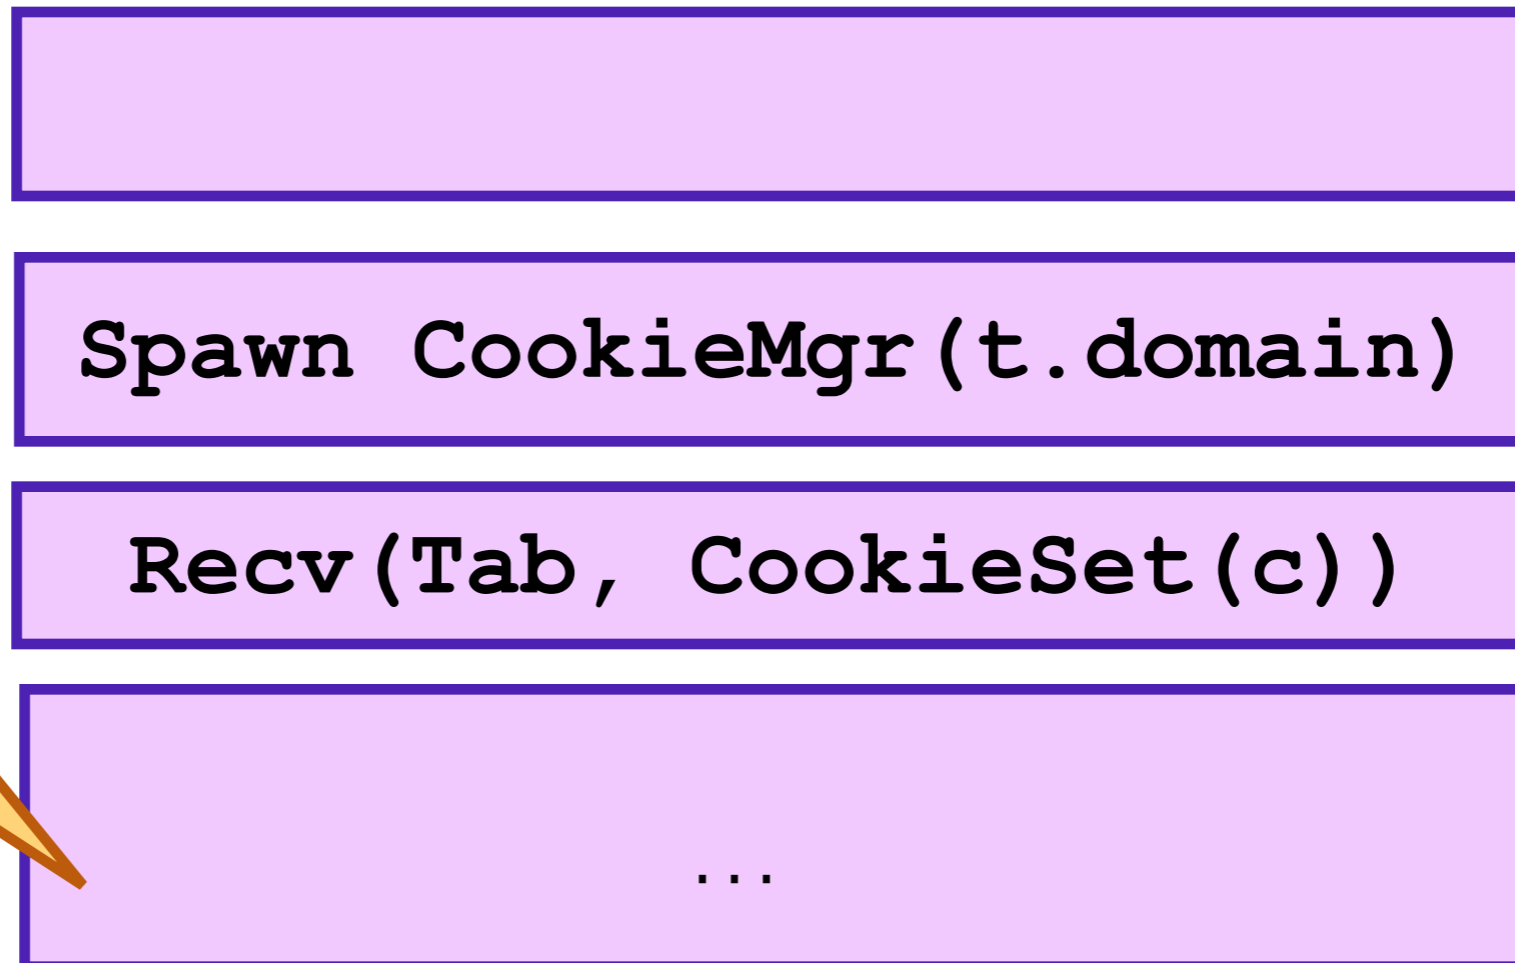
The system
calls so far

Time

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls



Time

The system
calls so far

Properties

```
When [Tab t] sends CookieSet(c) :  
  cp ← find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

Specify allowed behaviors wrt sequence of system calls

Send(cp, CookieSet(c))

Spawn CookieMgr(t.domain)

Recv(Tab, CookieSet(c))

...

The system
calls so far

Time

Example: Web browser kernel

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
    cp <- find CookieMgr(t.domain)  
    send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
    ...
```


Example: Web browser kernel

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):
```

```
    cp <- find CookieMgr(t.domain)
```

```
    send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):
```

```
    ...
```

Example: Web browser kernel

Specify cookie
integrity

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
  ...
```

Example: Web browser kernel



```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c) :  
    cp <- find CookieMgr(t.domain)  
    send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c) :  
    ...
```

Example: Web browser kernel

```
forall d c,
```

For any domain
d and cookie c

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):
```

```
  cp <- find CookieMgr(t.domain)
```

```
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):
```

```
  ...
```

Example: Web browser kernel

```
forall d c,
```

The kernel sends the cookie manager for domain d a cookie c

```
[Send(CookieMgr(d), CookieSet(c))]
```

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):
```

```
  cp <- find CookieMgr(t.domain)
```

```
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):
```

```
  ...
```

Example: Web browser kernel

```
forall d c,
```

Only if

```
Enables
```

```
[Send(CookieMgr(d), CookieSet(c))]
```

```
Components = Tab | CookieMgr | ...
```

```
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):
```

```
  cp <- find CookieMgr(t.domain)
```

```
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):
```

```
  ...
```

Example: Web browser kernel

```
forall d c,  
  [Recv (Tab (d), CookieSet (c))]  
  Enables  
  [Send (CookieMgr (d), CookieSet (c))]
```

Component
Message

The kernel already received a
cookie *c* from a tab of domain *d* . . .

. . .

Handlers:

```
When [Tab t] sends CookieSet (c) :  
  cp <- find CookieMgr (t.domain)  
  send (cp, CookieSet (c))
```

```
When [Tab t] sends CookieGet (c) :
```

. . .

Example: Web browser kernel

```
forall d c,  
  [Recv(Tab(d), CookieSet(c))]  
  Enables  
  [Send(CookieMgr(d), CookieSet(c))]
```

C
M
H

A Enables B
iff

...

every sys call B is preceded by sys call A

```
when [Tab t] sends CookieGet(c):  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

When [Tab t] sends CookieGet(c):

...

Example: Web browser kernel

```
forall d c,  
  [Recv(Tab(d), CookieSet(c))]  
  Enables  
  [Send(CookieMgr(d), CookieSet(c))]
```

```
Components = Tab | CookieMgr | ...  
Messages = CookieSet | CookieGet | ...
```

Handlers:

```
When [Tab t] sends CookieSet(c):  
  cp <- find CookieMgr(t.domain)  
  send(cp, CookieSet(c))
```

```
When [Tab t] sends CookieGet(c):  
  ...
```

Example: Web browser kernel



Example: Web browser kernel

REFLEX Benefits:

for

D

E

D

Com

Mes

Han

W

W

...

Example: Web browser kernel

REFLEX Benefits:

Proofs fully automated

Example: Web browser kernel

REFLEX Benefits:

Proofs fully automated

No lemmas

...

Example: Web browser kernel

REFLEX Benefits:

Proofs fully automated

No lemmas

No invariants

...

Example: Web browser kernel

REFLEX Benefits:

Proofs fully automated

No lemmas

No invariants

No manual proofs

...

Evaluation

Evaluation

**Web
browser**

**SSH
server**

**Web
server**

Evaluation

**Web
browser**

**SSH
server**

**Web
server**

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

Evaluation

Web browser	Domains do not interfere, Cookie integrity, ...
SSH server	No PTY access before authentication, At most 3 authentication attempts, ...
Web server	Clients only spawned after successful login, File requests guarded by access control, ...

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

Evaluation

Web browser	Domains do not interfere, Cookie integrity, ...
SSH	No root access before authentication, at 3 authentication attempts, ...
server	Only spawned after successful login, File requests guarded by access control, ...

Able to automate proofs of non-interference

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

Evaluation

Web browser	Domains do not interfere, Cookie integrity, ...
SSH server	No PTY access before authentication, At most 3 authentication attempts, ...
Web server	Clients only spawned after successful login, File requests guarded by access control, ...

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

Evaluation

Web browser	Domains do Cookie in	Able to automate proofs of non-local properties
SSH server	No PTY access before authentication, At most 3 authentication attempts, ...	
Web server	Clients only spawned after successful login, File requests guarded by access control, ...	

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

Development Effort: Systems

Development Effort: Systems

<i>Benchmark</i>	<i>Language</i>	<i>Lines of Code/Spec</i>	
Browser	REFLEX	81	37
	C++, Python	970,240	-
SSH server	REFLEX	64	22
	C, Python	89,567	-
Web server	REFLEX	56	29
	Python	386	-

Development Effort: Systems

<i>Benchmark</i>	<i>Language</i>	<i>Lines of Code/Spec</i>	
Browser	REFLEX	81	37
	C++, Python	970,240	Webkit
SSH server	REFLEX	64	22
	C, Python	89,567	-
Web server	REFLEX	56	29
	Python	386	-

Development Effort: Systems

<i>Benchmark</i>	<i>Language</i>	<i>Lines of Code/Spec</i>	
Browser	REFLEX	81	37
	C++, Python	970,240	Webkit
SSH server	REFLEX	64	22
	C, Python	89,567	OpenSSH
Web server	REFLEX	56	29
	Python	386	-

Expressiveness Limitations

Expressiveness Limitations

Strict subset of temporal logic + non-interference

Expressiveness Limitations

Strict subset of temporal logic + non-interference

Loop free handlers


Expressiveness Limitations

Strict subset of temporal logic + non-interference


Loop free handlers

No user-defined unbounded data structures


SSH Server Kernel in REFLEX



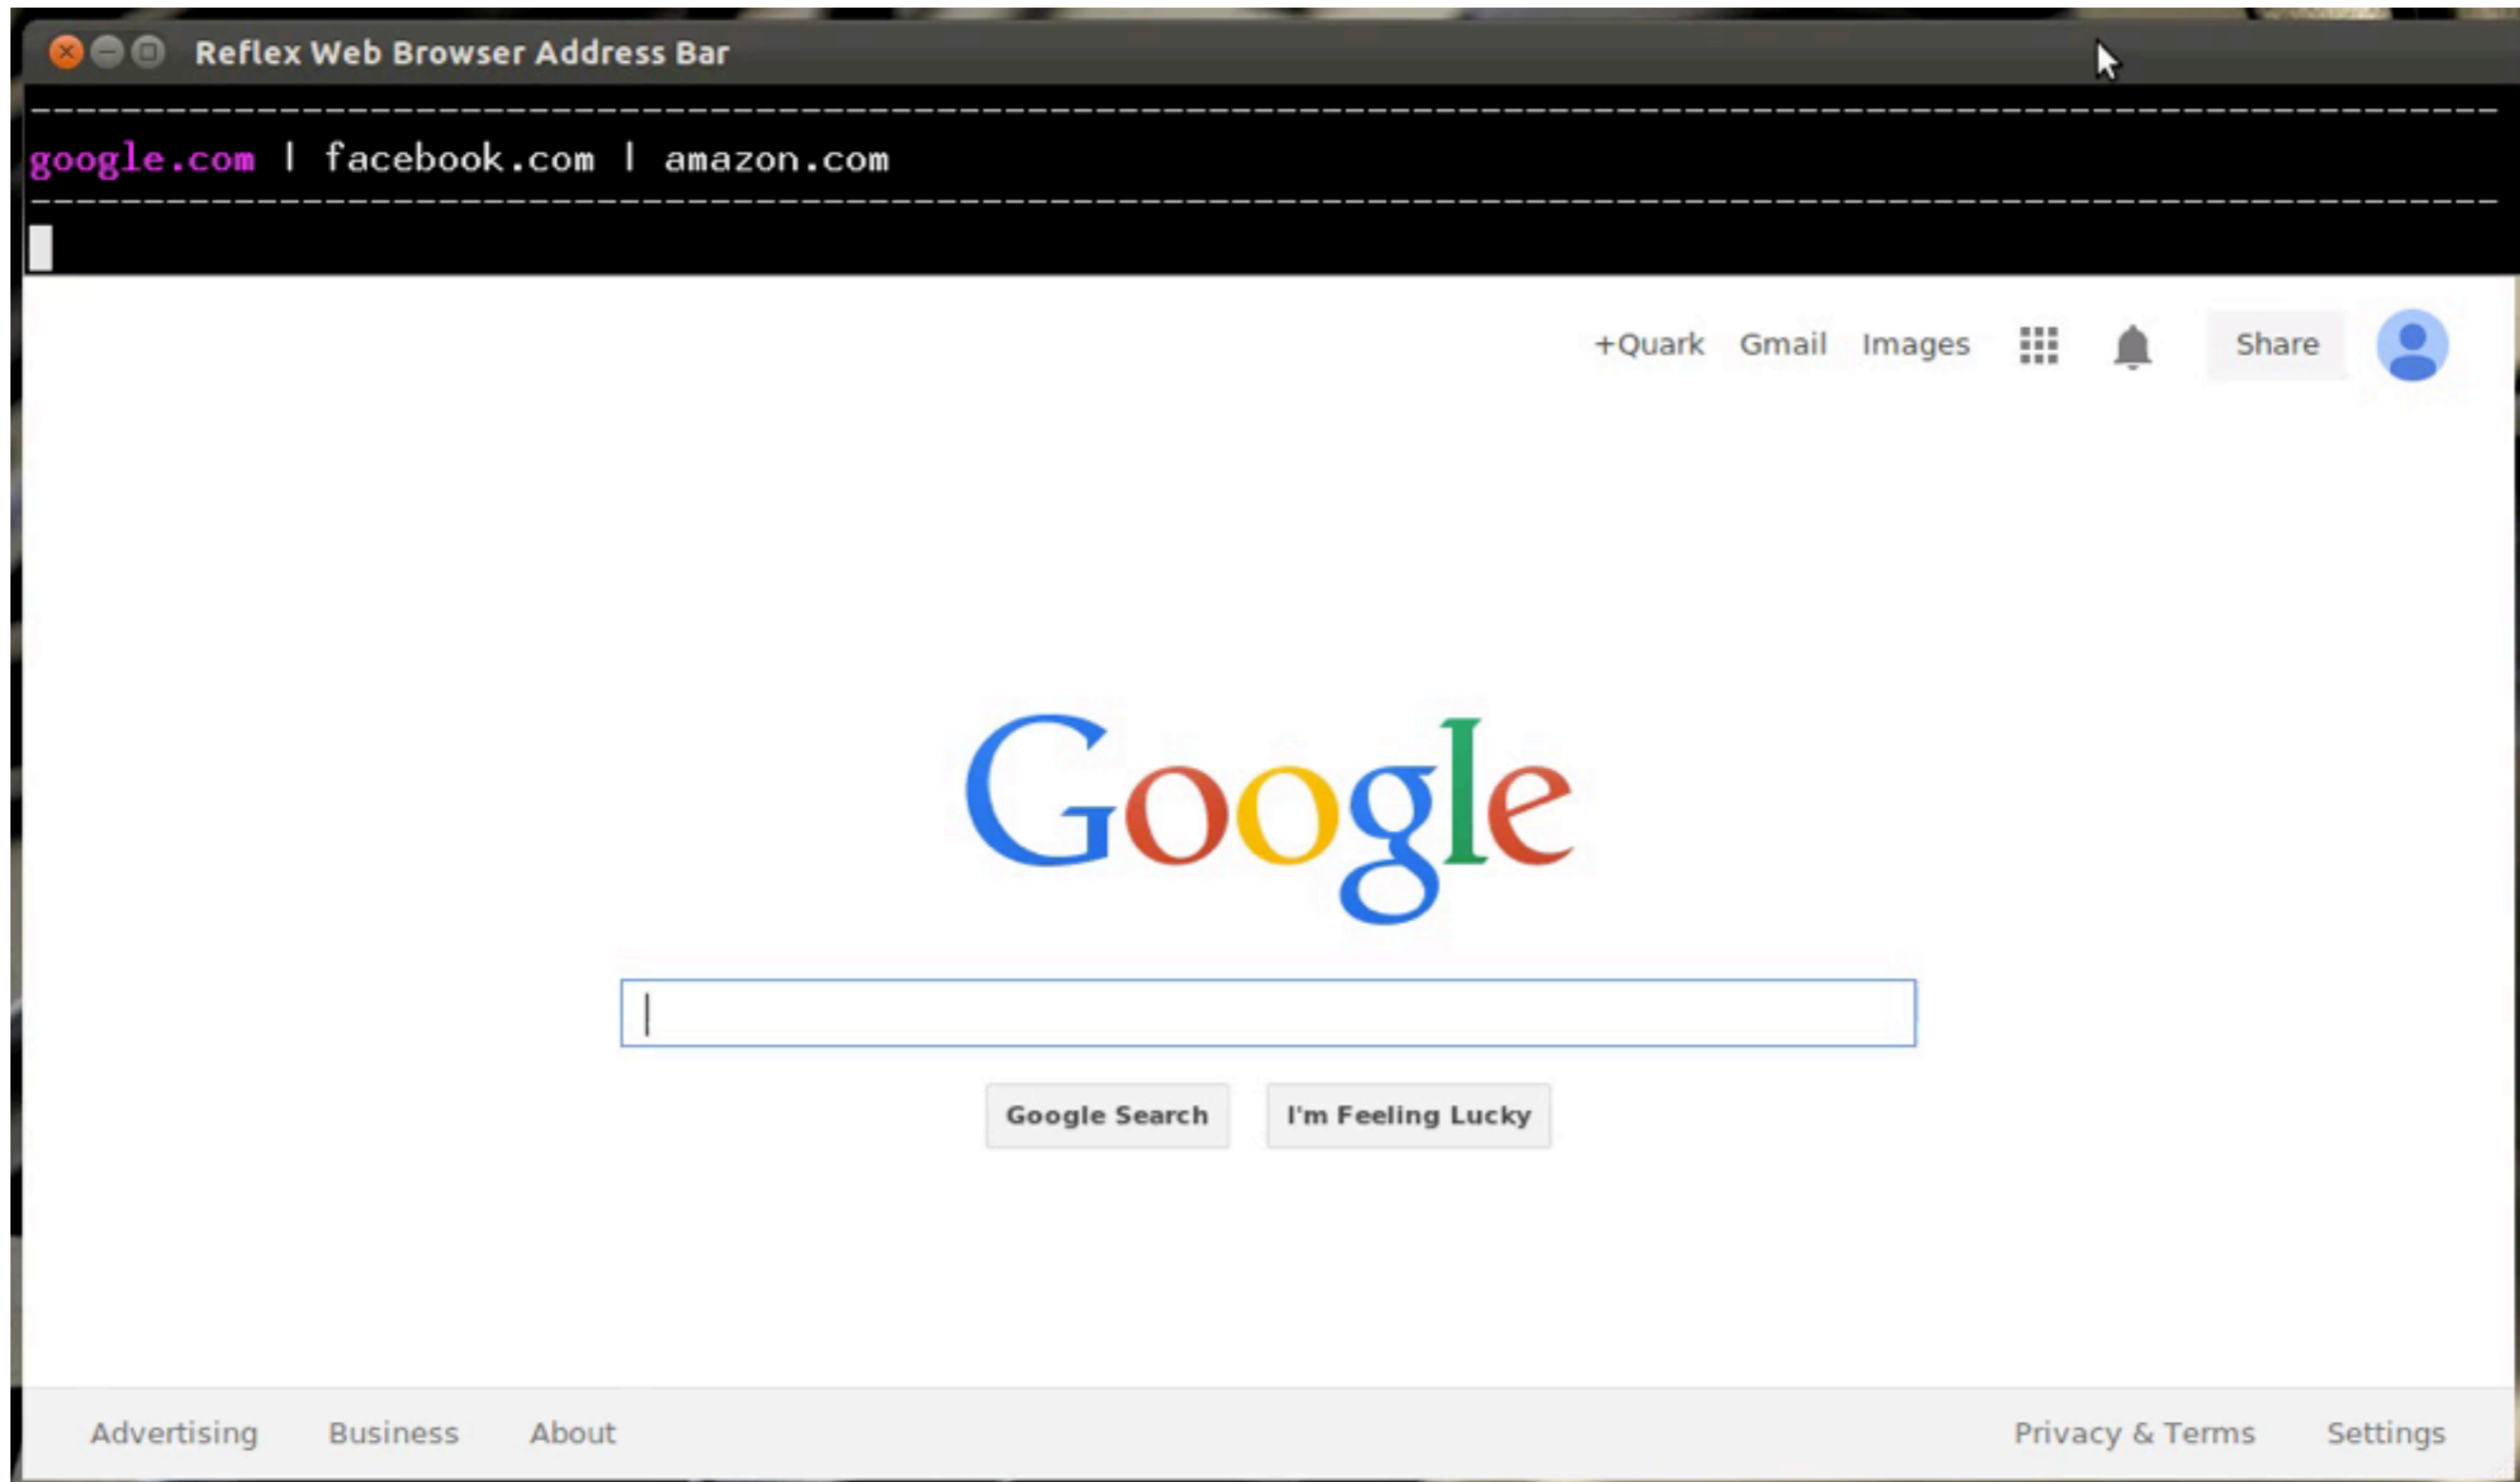
SSH Server Kernel in REFLEX



Web Browser Kernel in REFLEX



Web Browser Kernel in REFLEX



Ongoing: UAV controllers


Ongoing: UAV controllers




Ongoing: UAV controllers



Ongoing: UAV controllers



Ongoing: UAV controllers



Ongoing: UAV controllers



Ongoing: UAV controllers



Used by military, companies, and individuals

Ongoing: UAV controllers



Used by military, companies, and individuals

Safety critical


Ongoing: UAV controllers

Ongoing: UAV controllers




Controller

Ongoing: UAV controllers




Ongoing: UAV controllers



Height is always < 400 ft

Ongoing: UAV controllers



Height is always < 400 ft

UAV eventually reaches target location

Ongoing: UAV controllers

DSL for dynamics and properties

DSL for controllers

Spec


Code

Proof Automation

Ongoing: UAV controllers

DSL for dynamics and properties

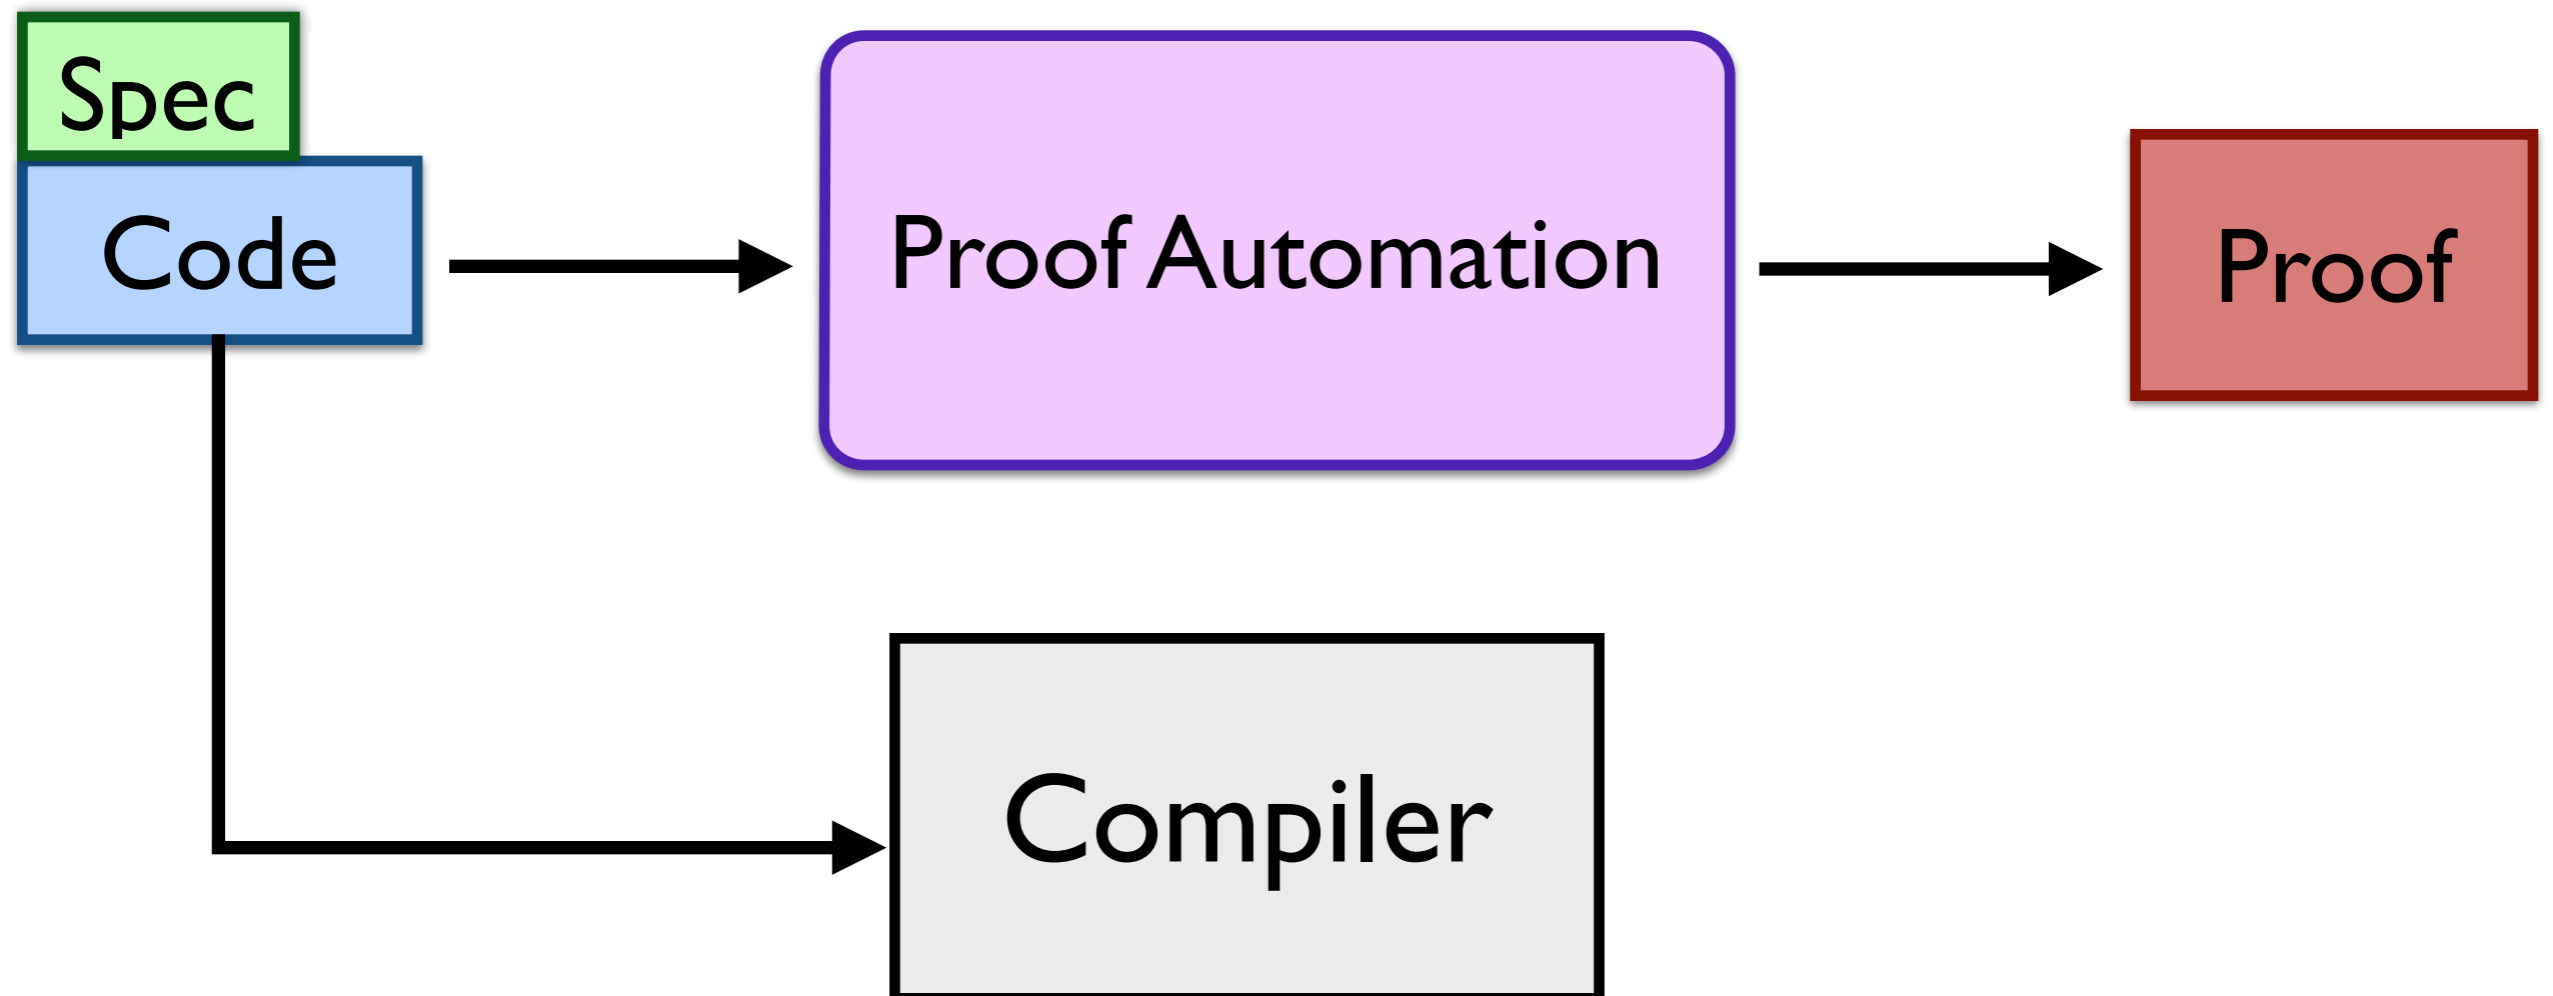
DSL for controllers



Ongoing: UAV controllers

DSL for dynamics and properties

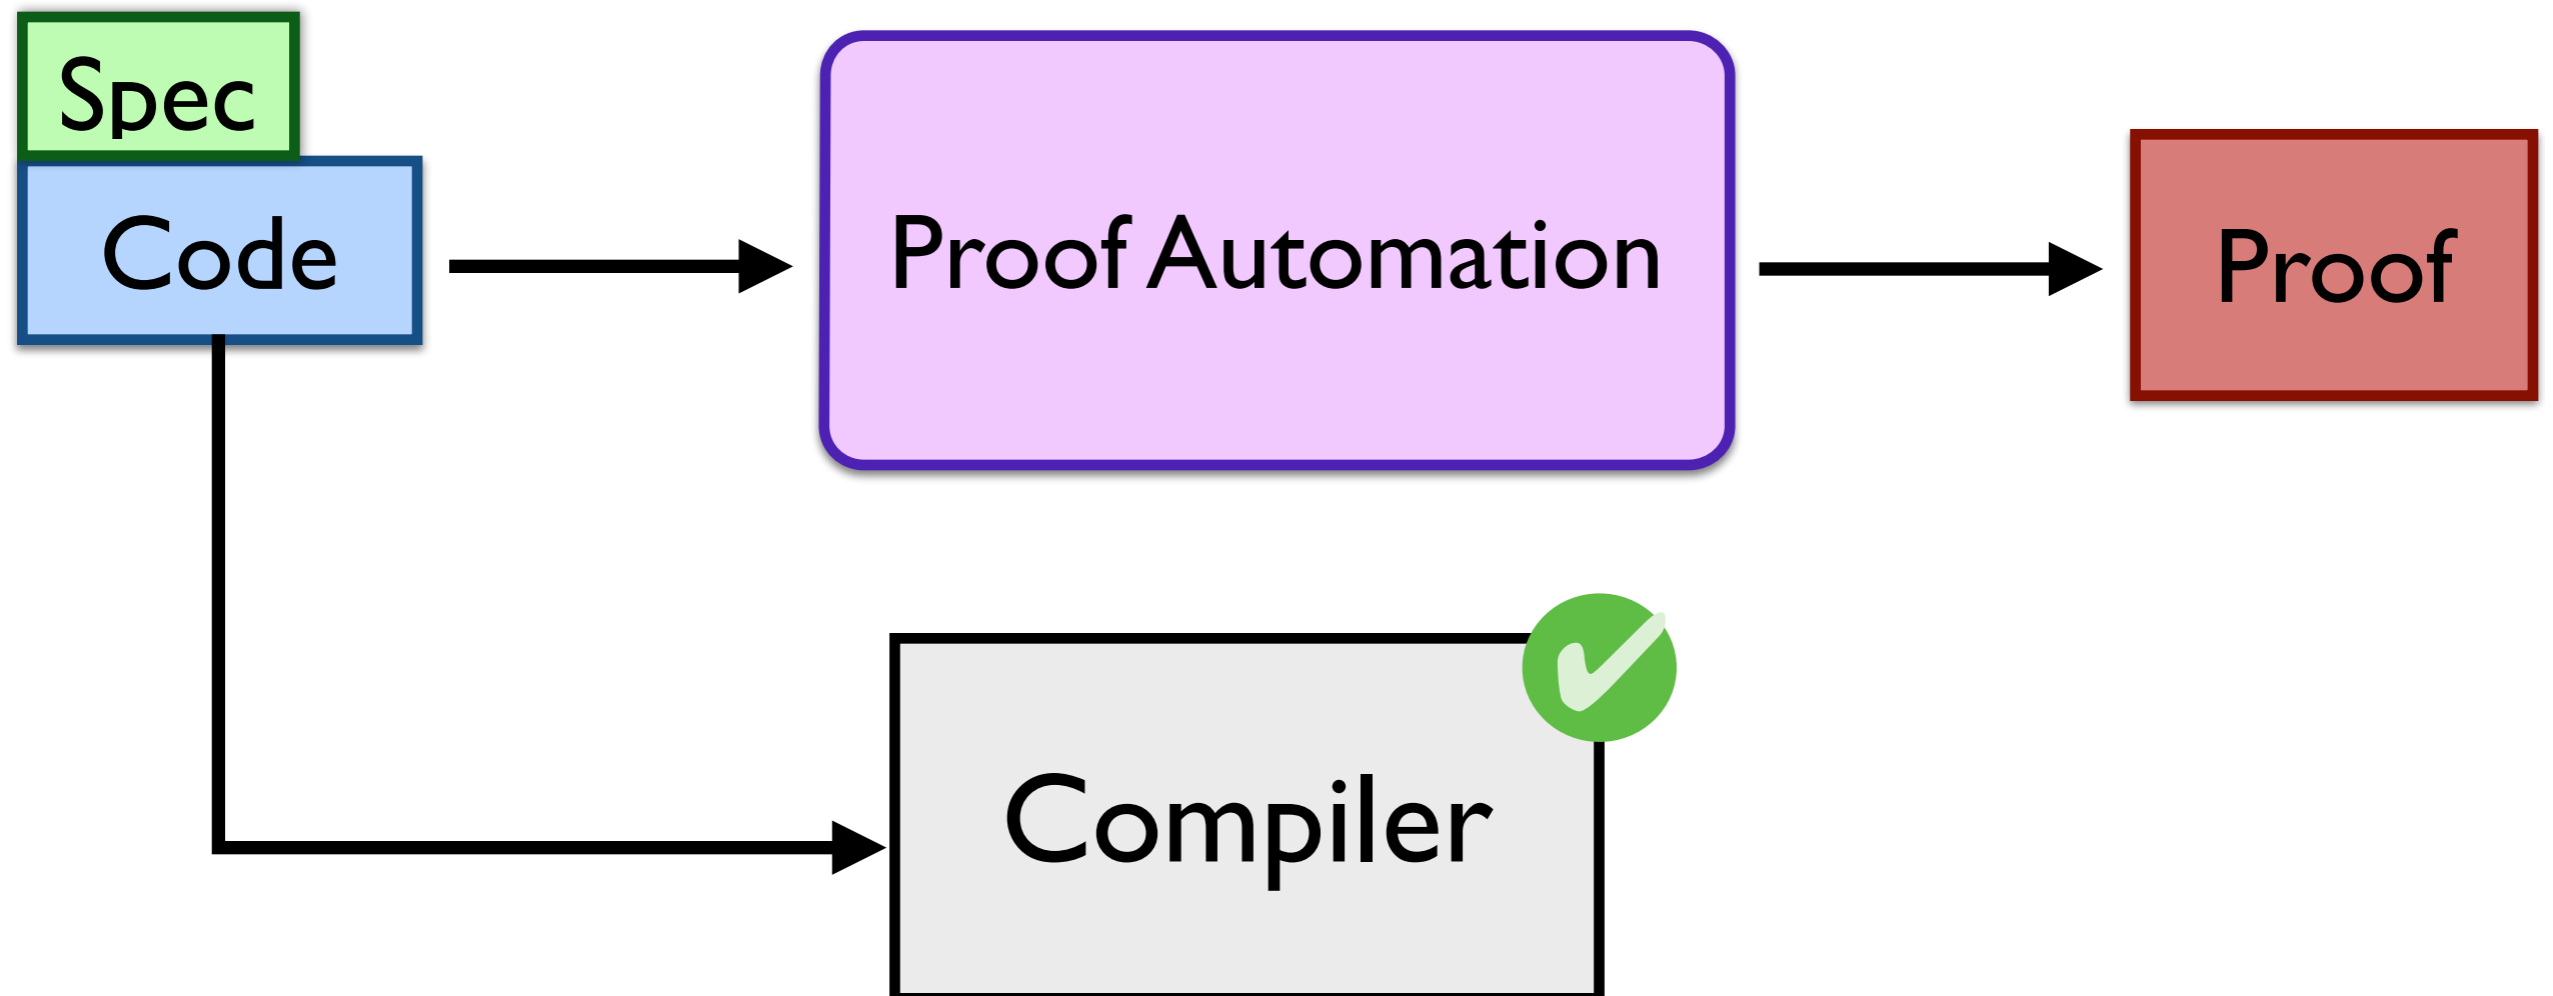
DSL for controllers



Ongoing: UAV controllers

DSL for dynamics and properties

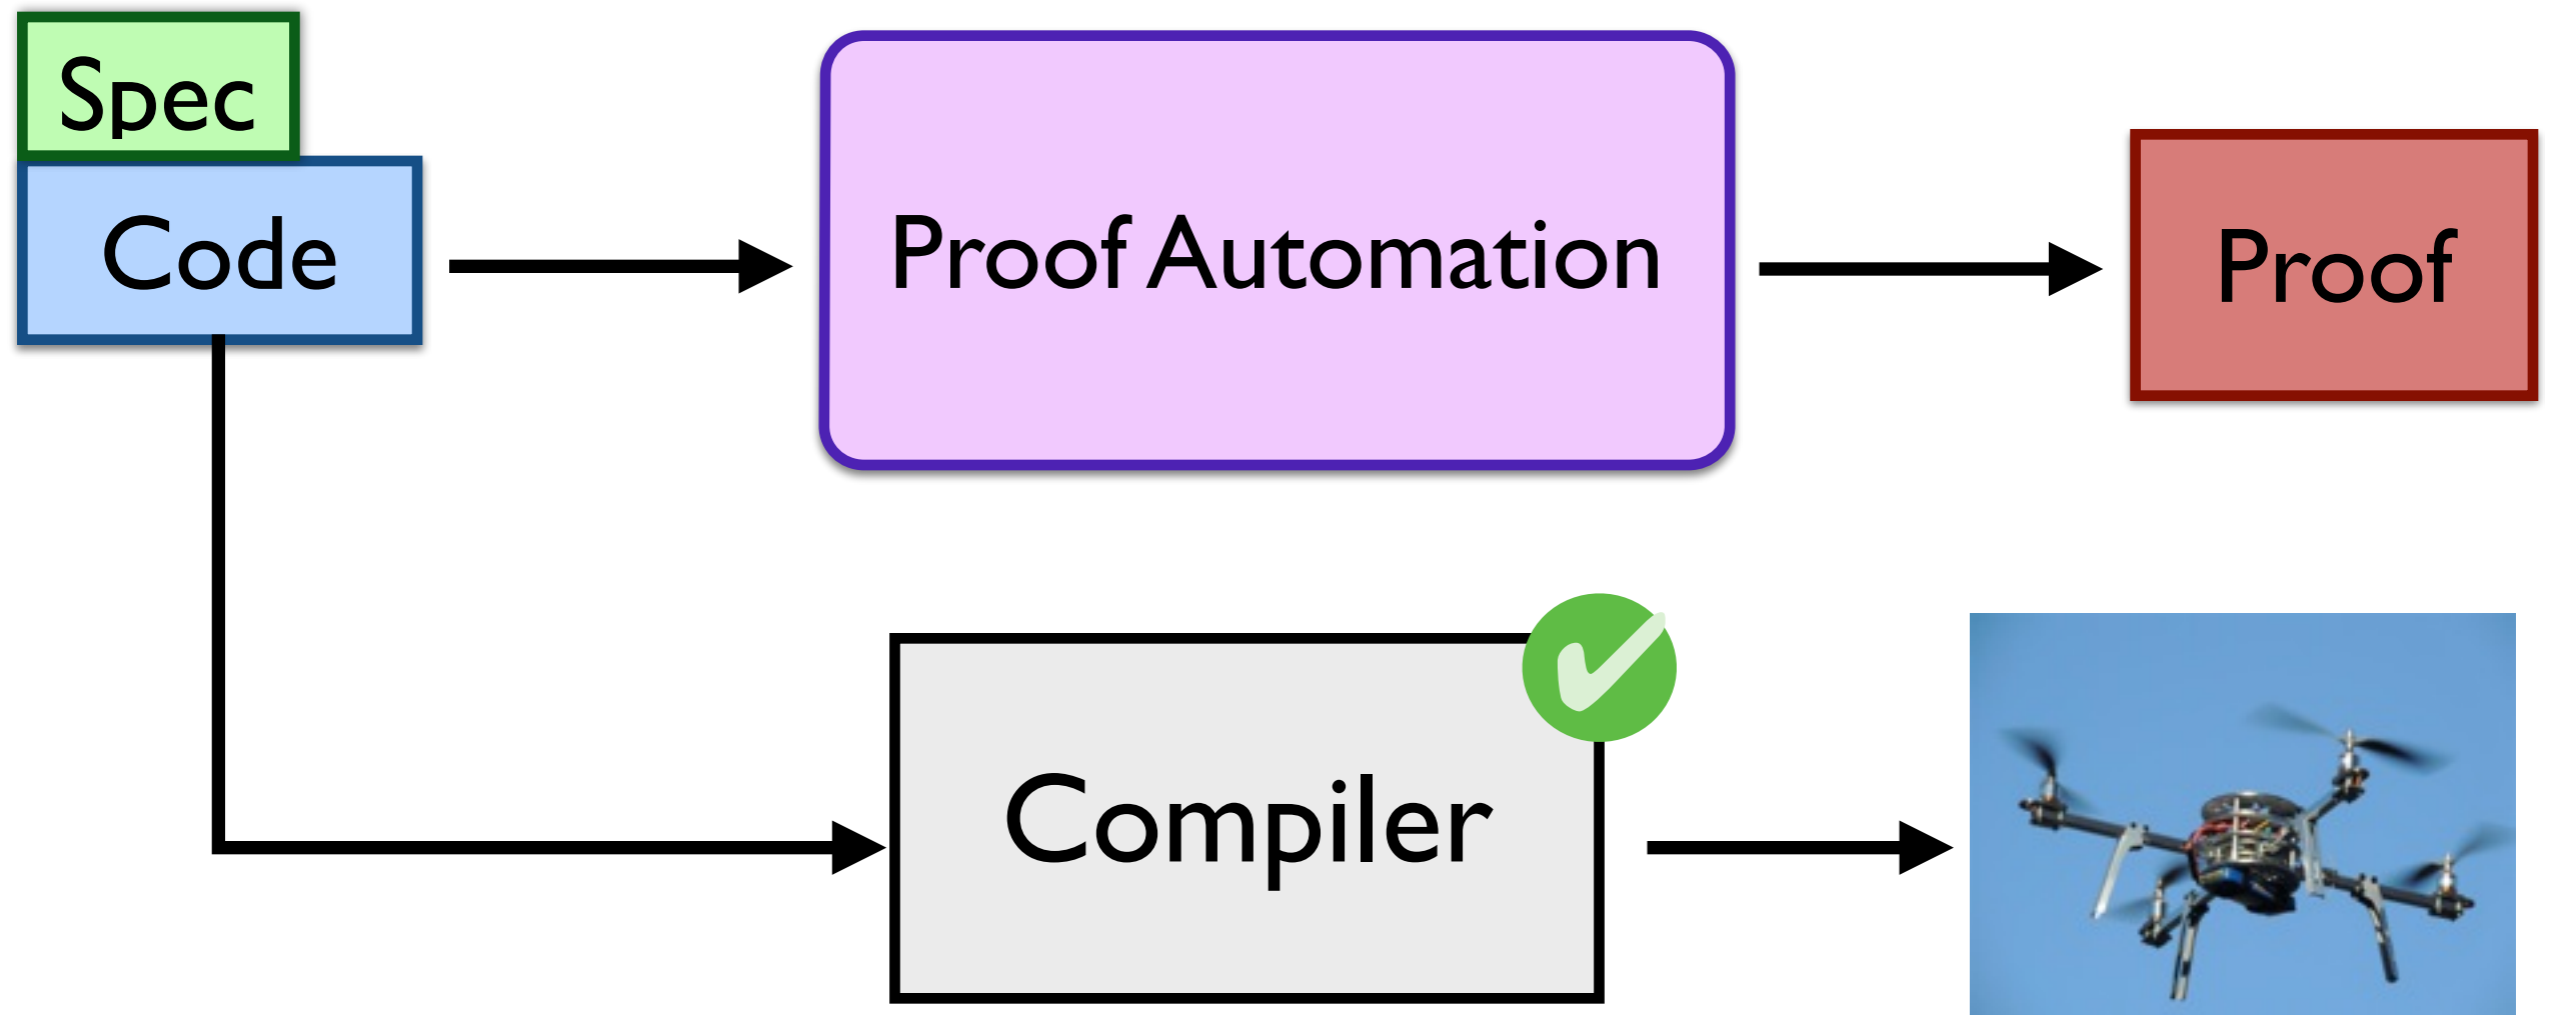
DSL for controllers



Ongoing: UAV controllers

DSL for dynamics and properties

DSL for controllers



Goal

**End-to-end
guarantee in
Coq**



Conclusion

Proof assistant based verification




Conclusion

Proof assistant based verification



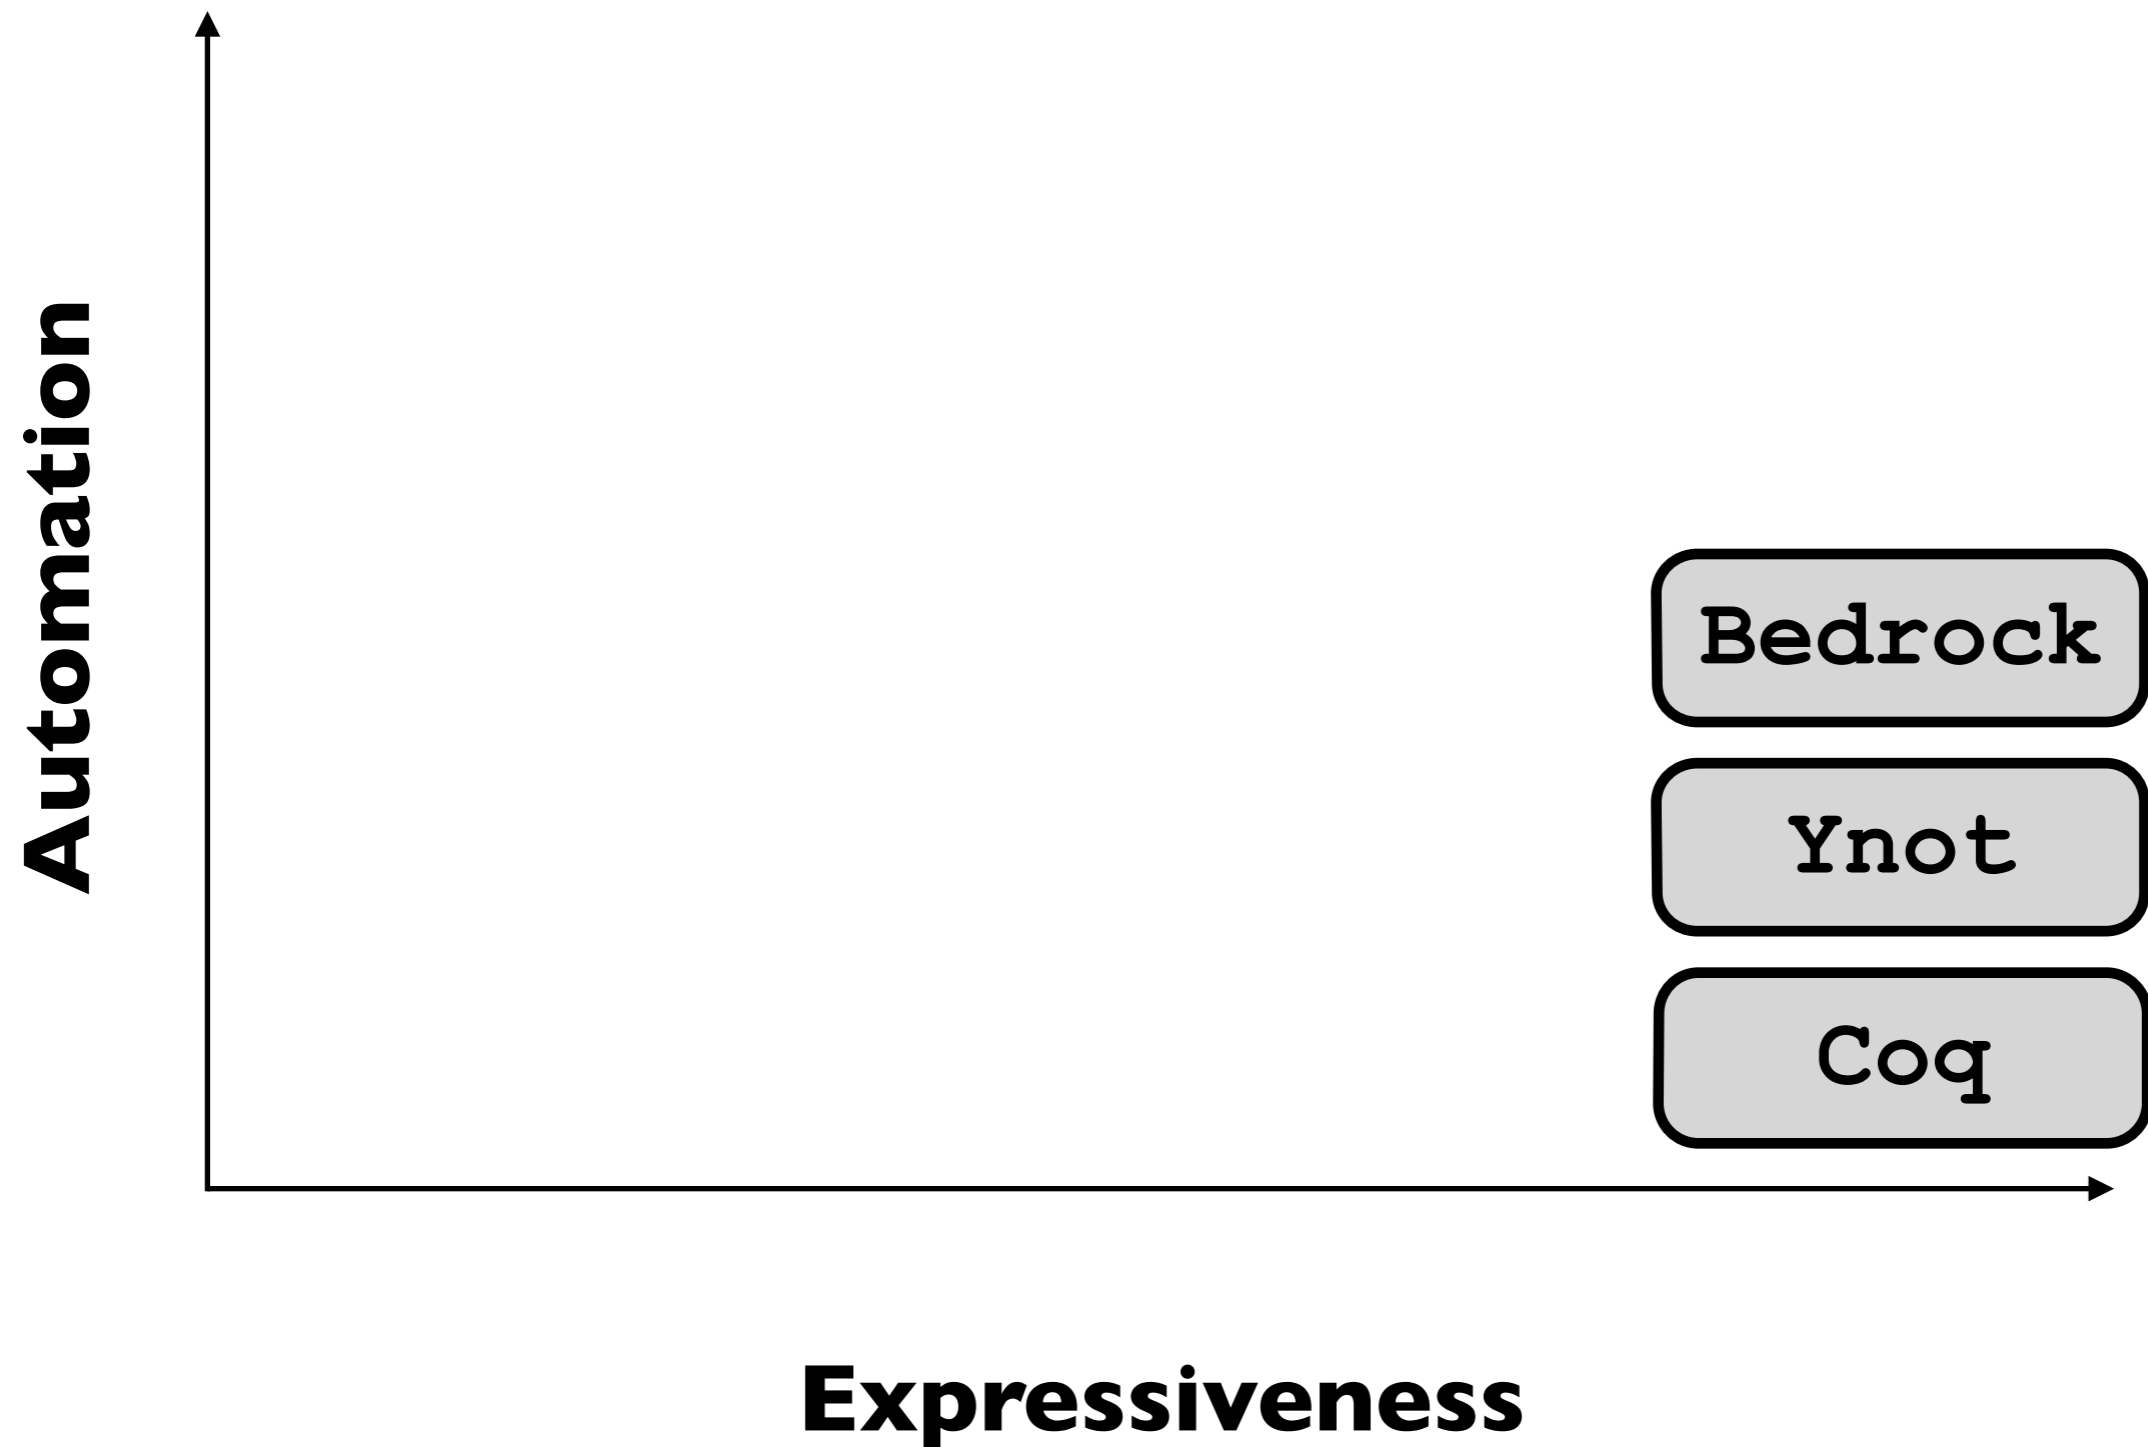
Conclusion

Proof assistant based verification



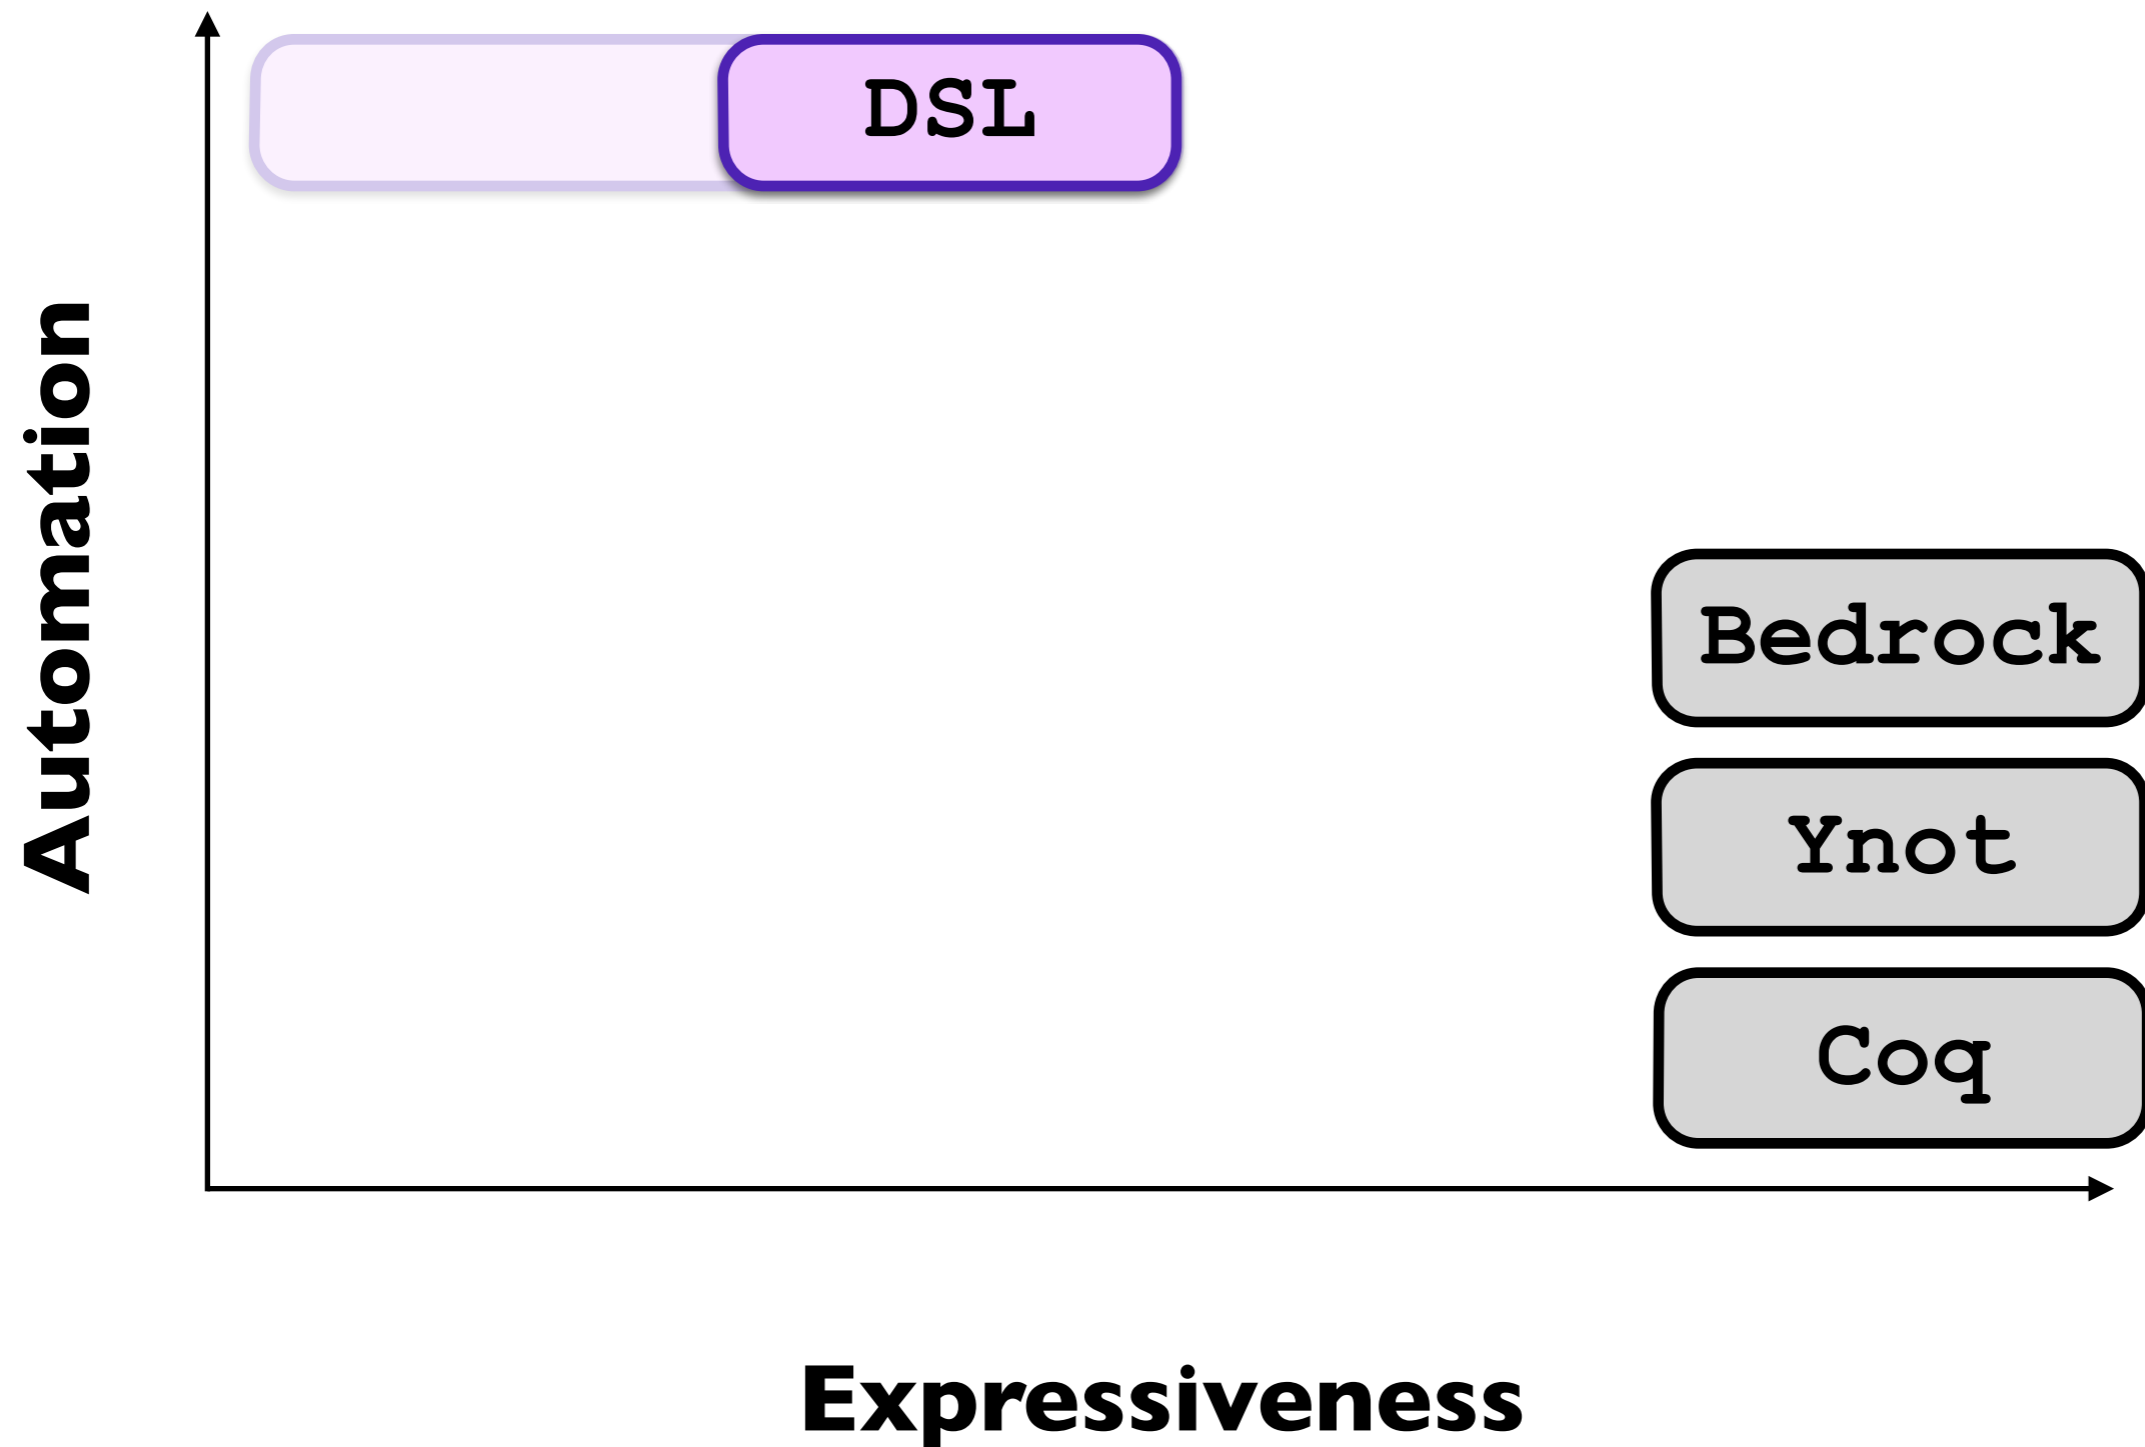
Conclusion

Proof assistant based verification



Conclusion

Proof assistant based verification




Conclusion

DSL

Conclusion




DSL



DSL expressive enough for entire domain

Conclusion

DSL






DSL expressive enough for entire domain

Automation eliminates manual proof burden

Conclusion

DSL




DSL expressive enough for entire domain

Automation eliminates manual proof burden

<http://goto.ucsd.edu/reflex/>

Thank You!

DSL



DSL expressive enough for entire domain

Automation eliminates manual proof burden

<http://goto.ucsd.edu/reflex/>