



# CStar: Unifying Programming and Verification in C

Di Wang

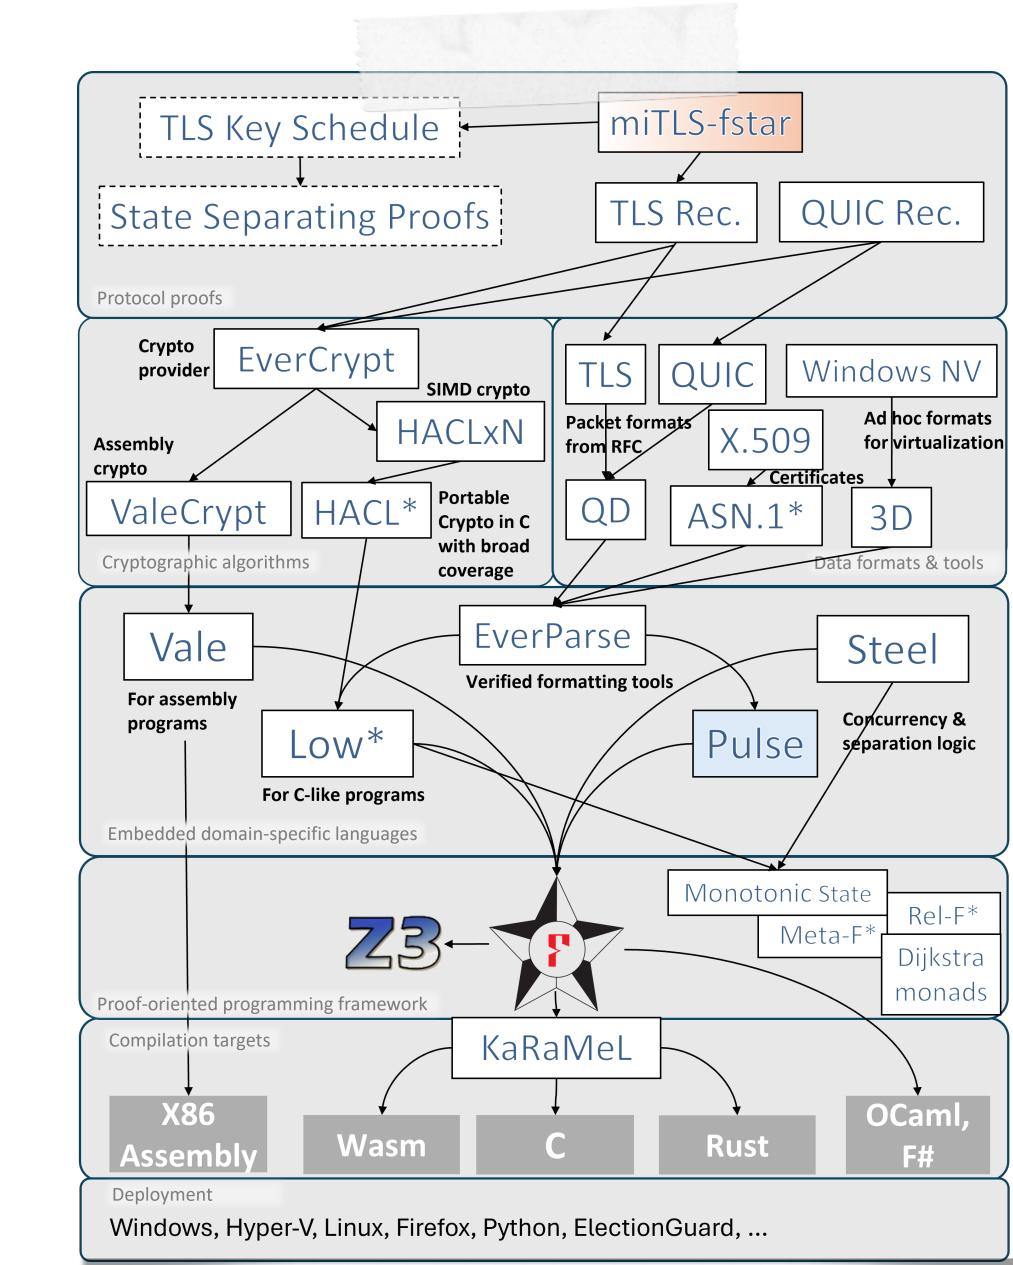
Programming Languages Lab, Peking University

# Towards Safer Systems Software

formal verification to the rescue?



COMPCERT



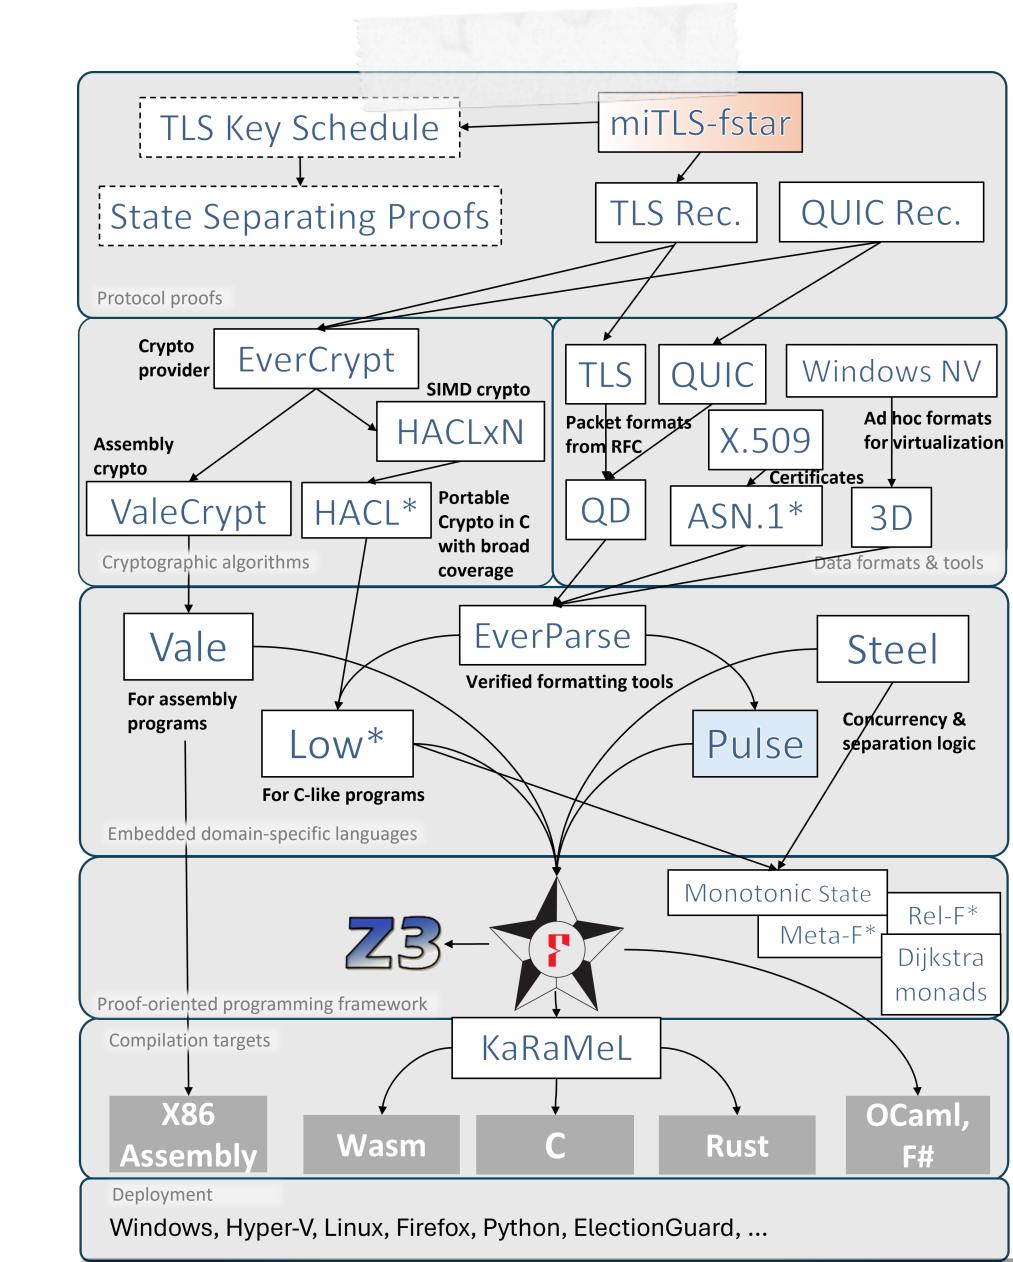
# Towards Safer Systems Software

formal verification to the rescue?



COMPCERT

verified kernels



# Towards Safer Systems Software

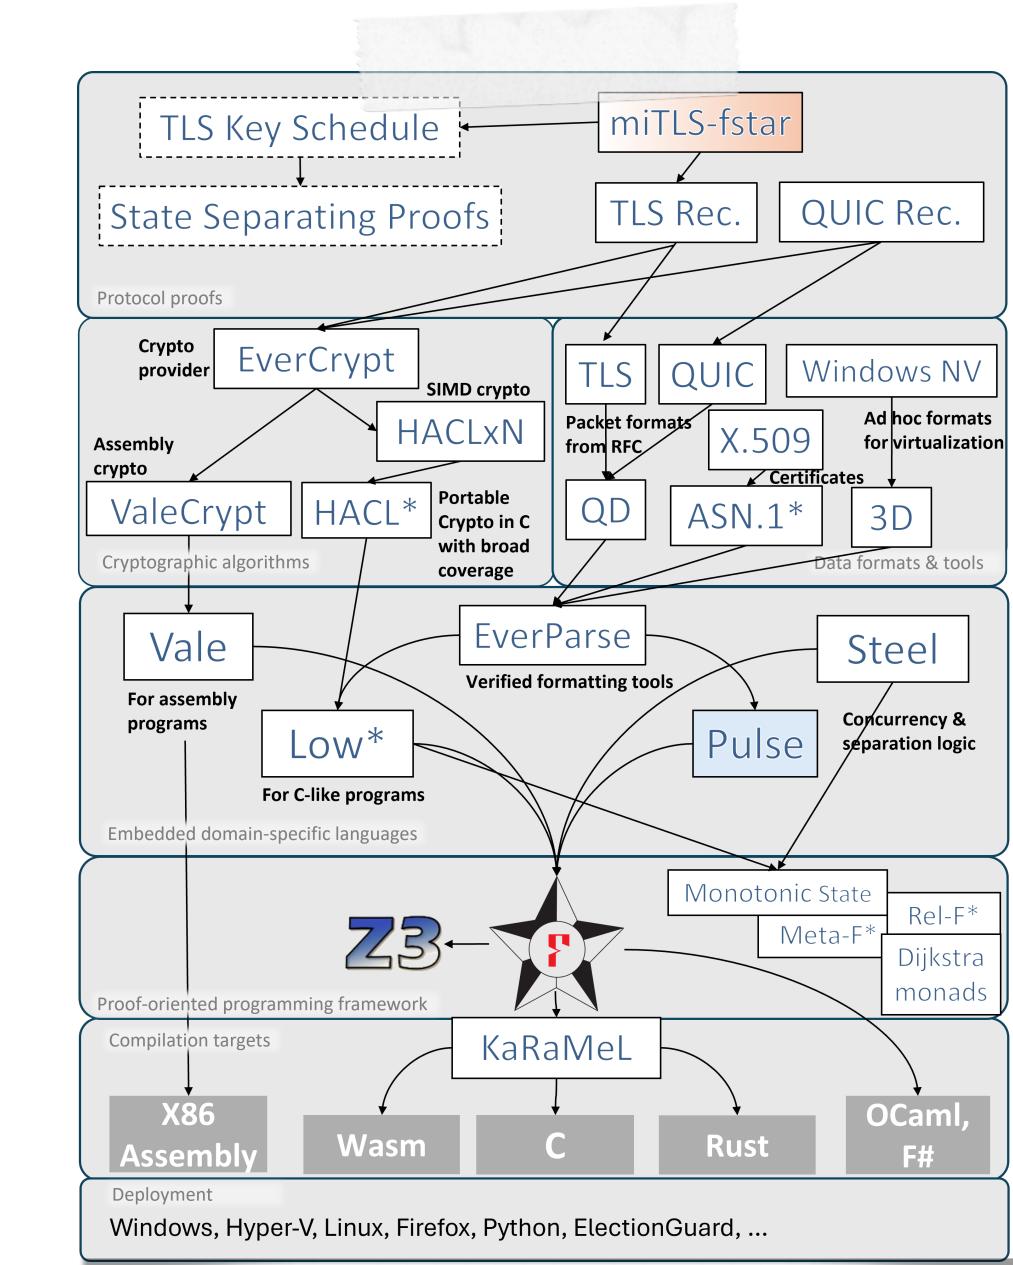
formal verification to the rescue?



verified kernels

COMPCERT

verified compilers



# Towards Safer Systems Software

formal verification to the rescue?

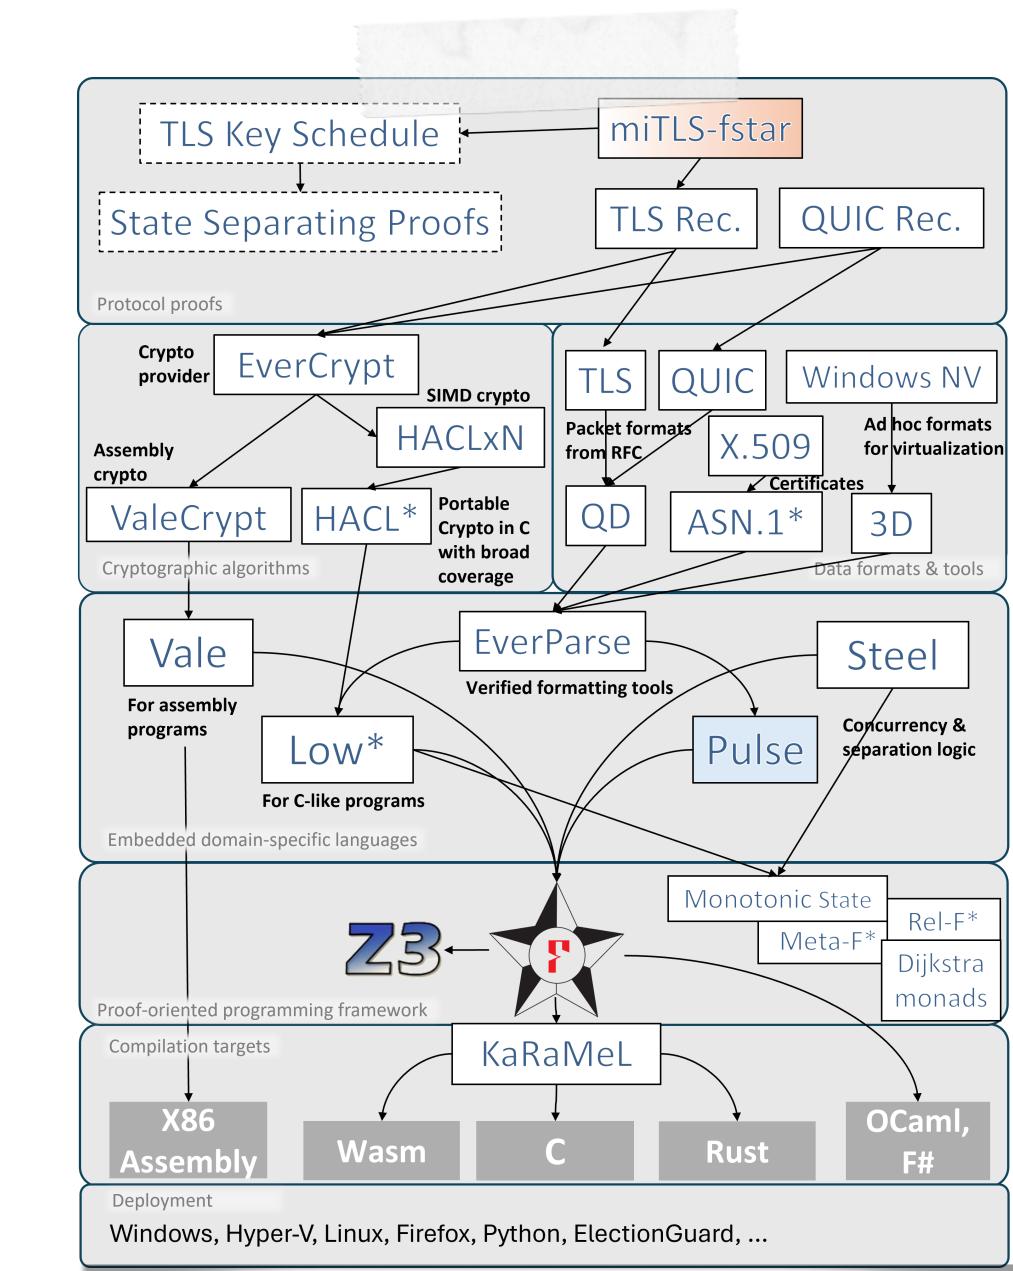


verified kernels



COMPCERT

verified compilers

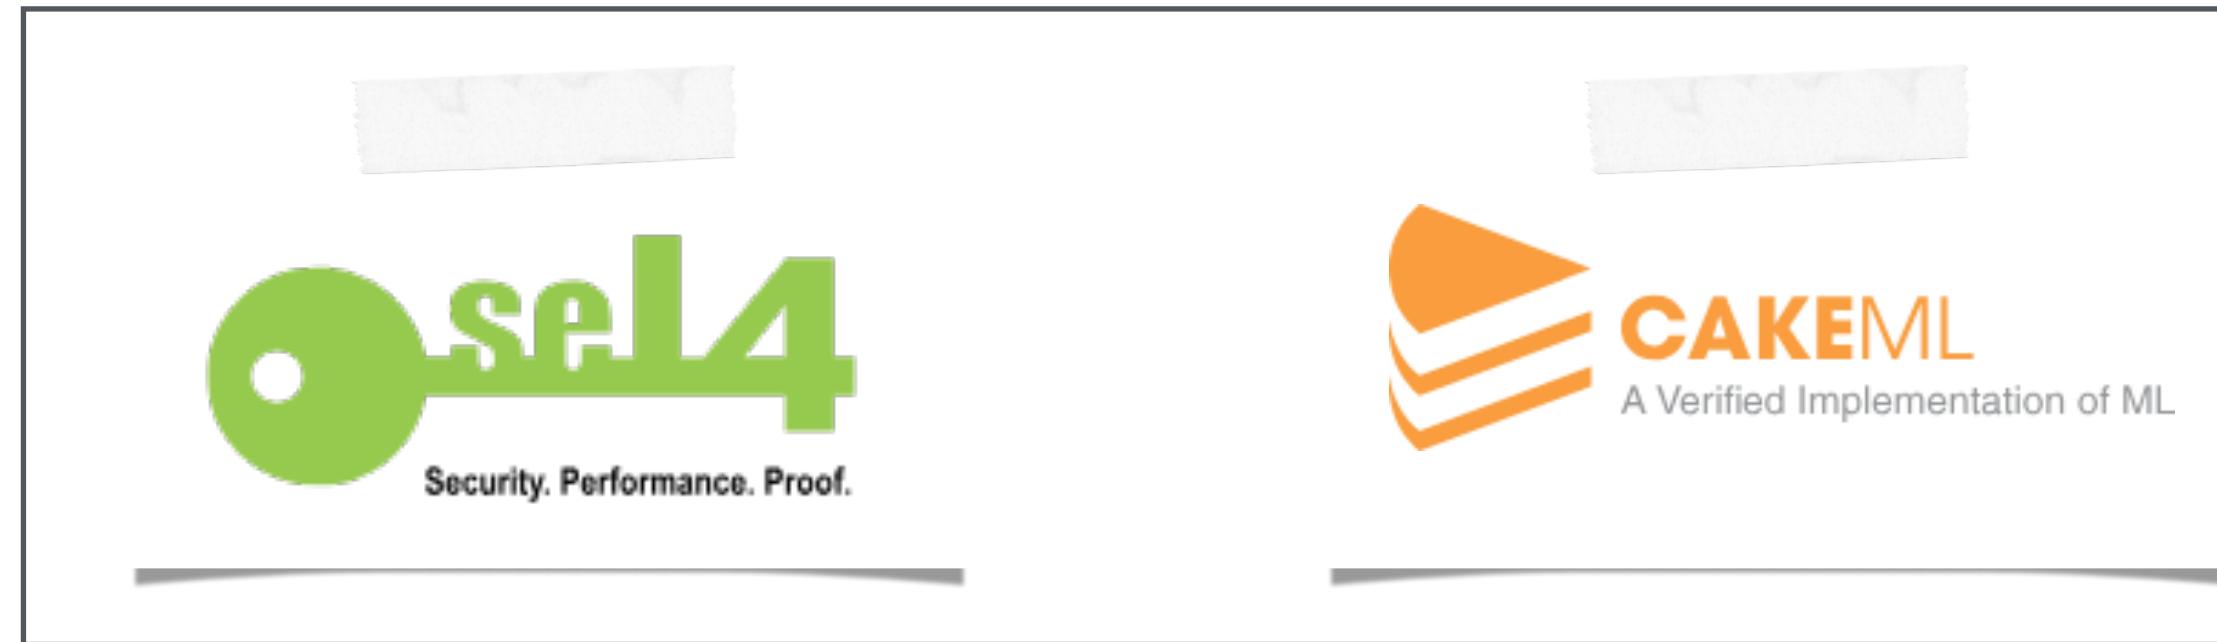


verified protocols

# Towards Safer Systems Software

formal verification to the rescue?

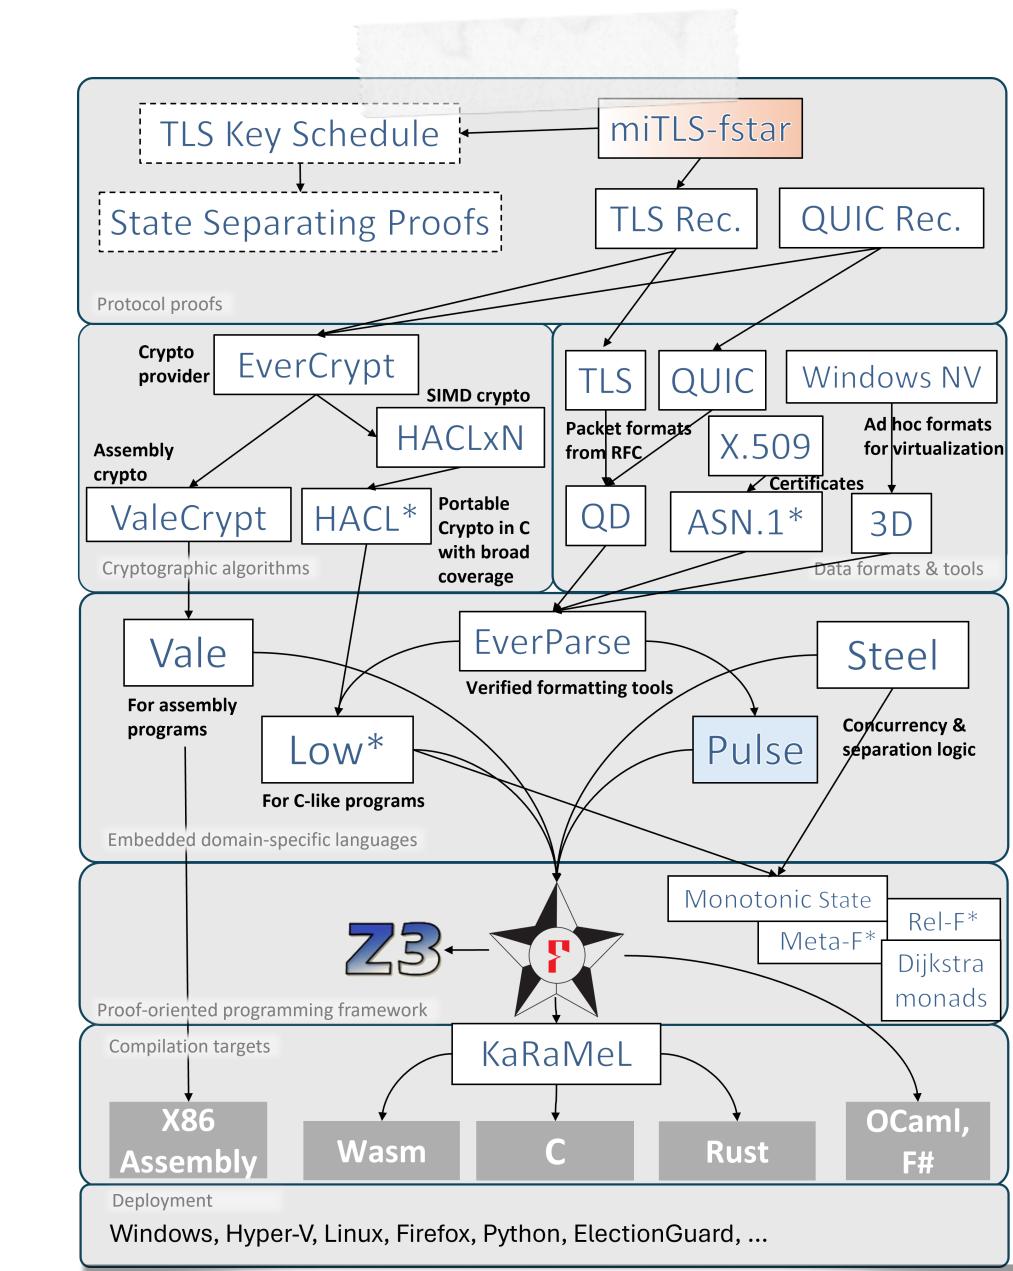
Isabelle/HOL  
HOL4



verified kernels

COMPCERT

verified compilers

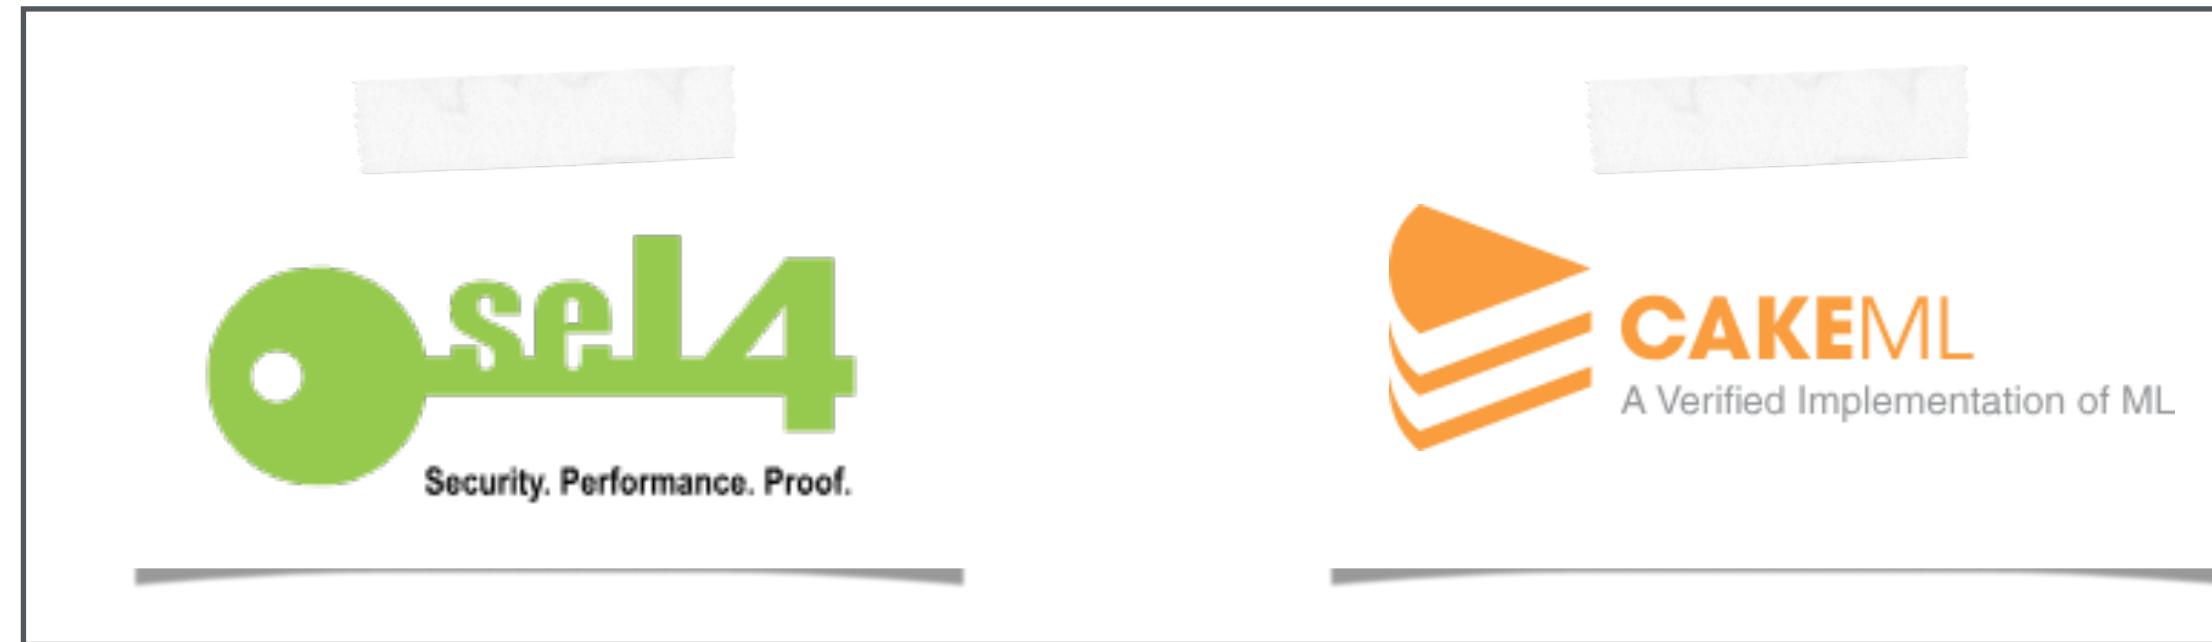


verified protocols

# Towards Safer Systems Software

formal verification to the rescue?

Isabelle/HOL  
HOL4

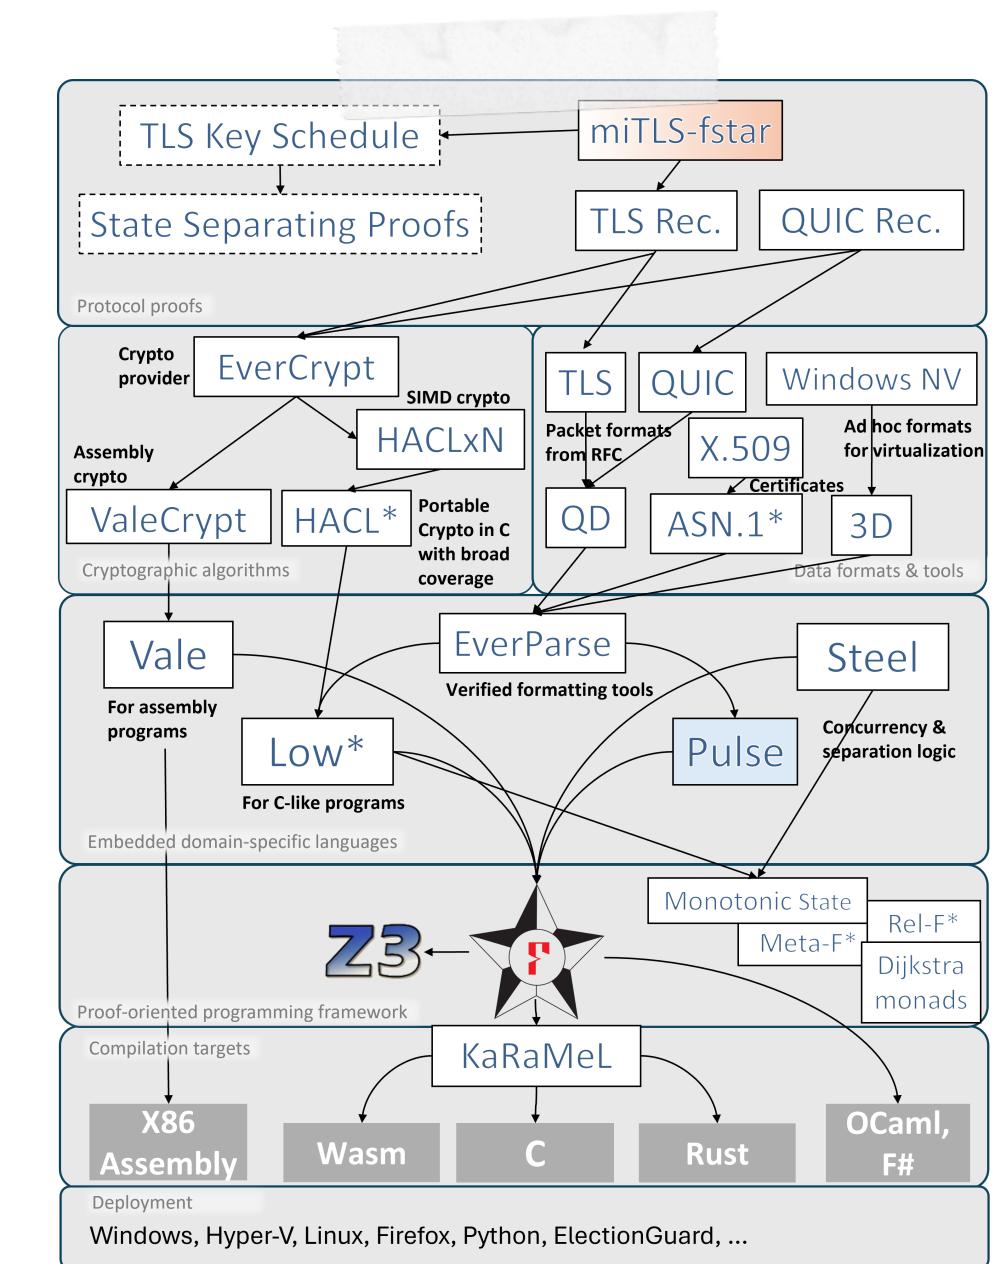


Rocq  
(formerly Coq)



verified kernels

verified compilers

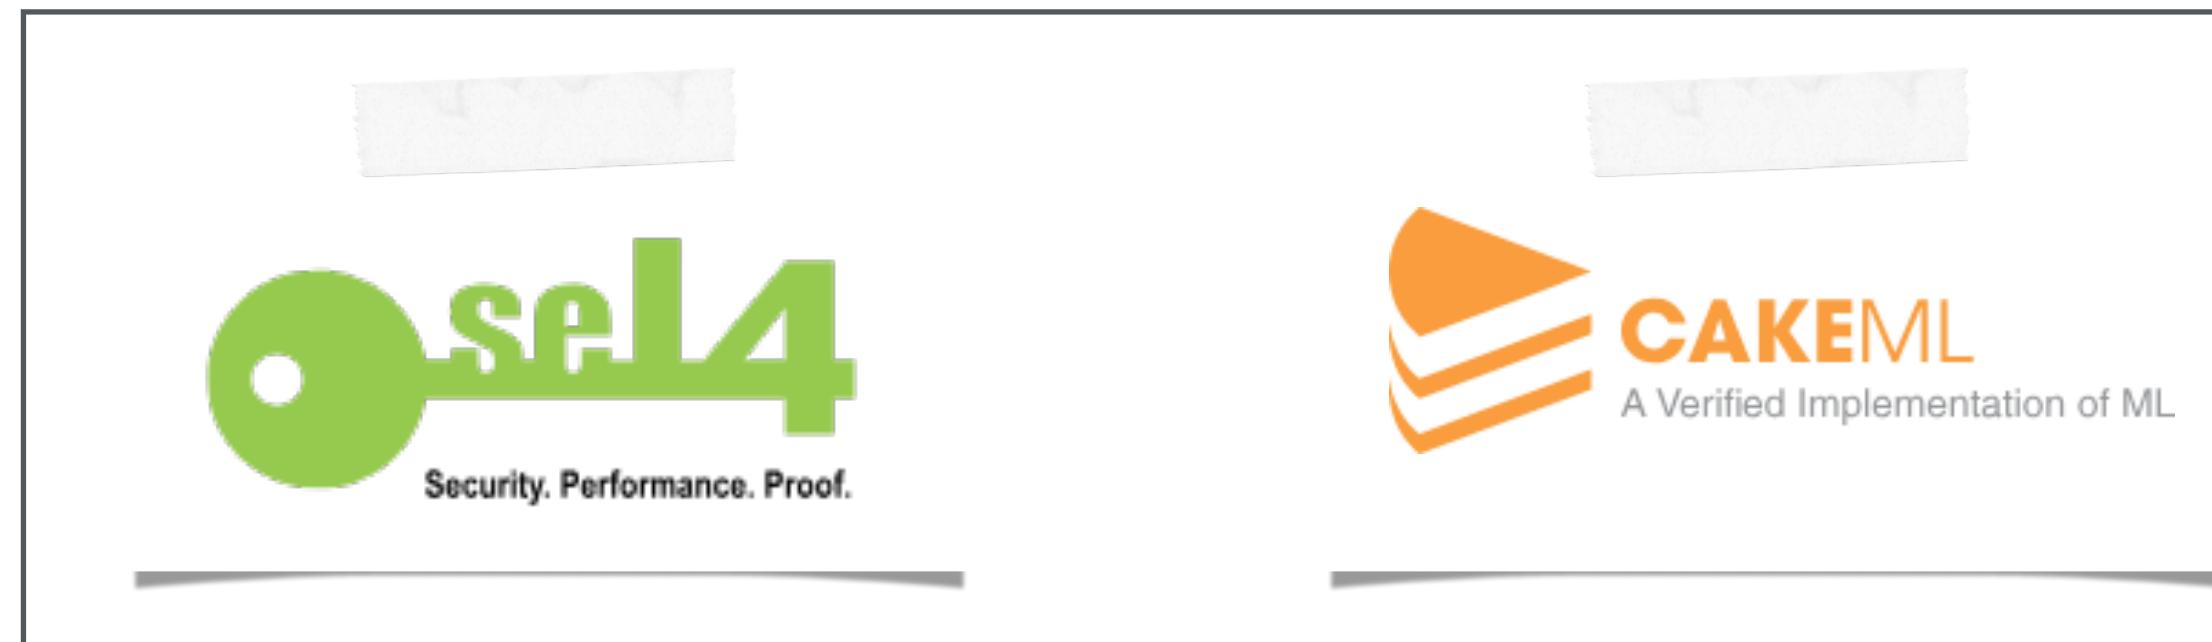


verified protocols

# Towards Safer Systems Software

formal verification to the rescue?

Isabelle/HOL  
HOL4

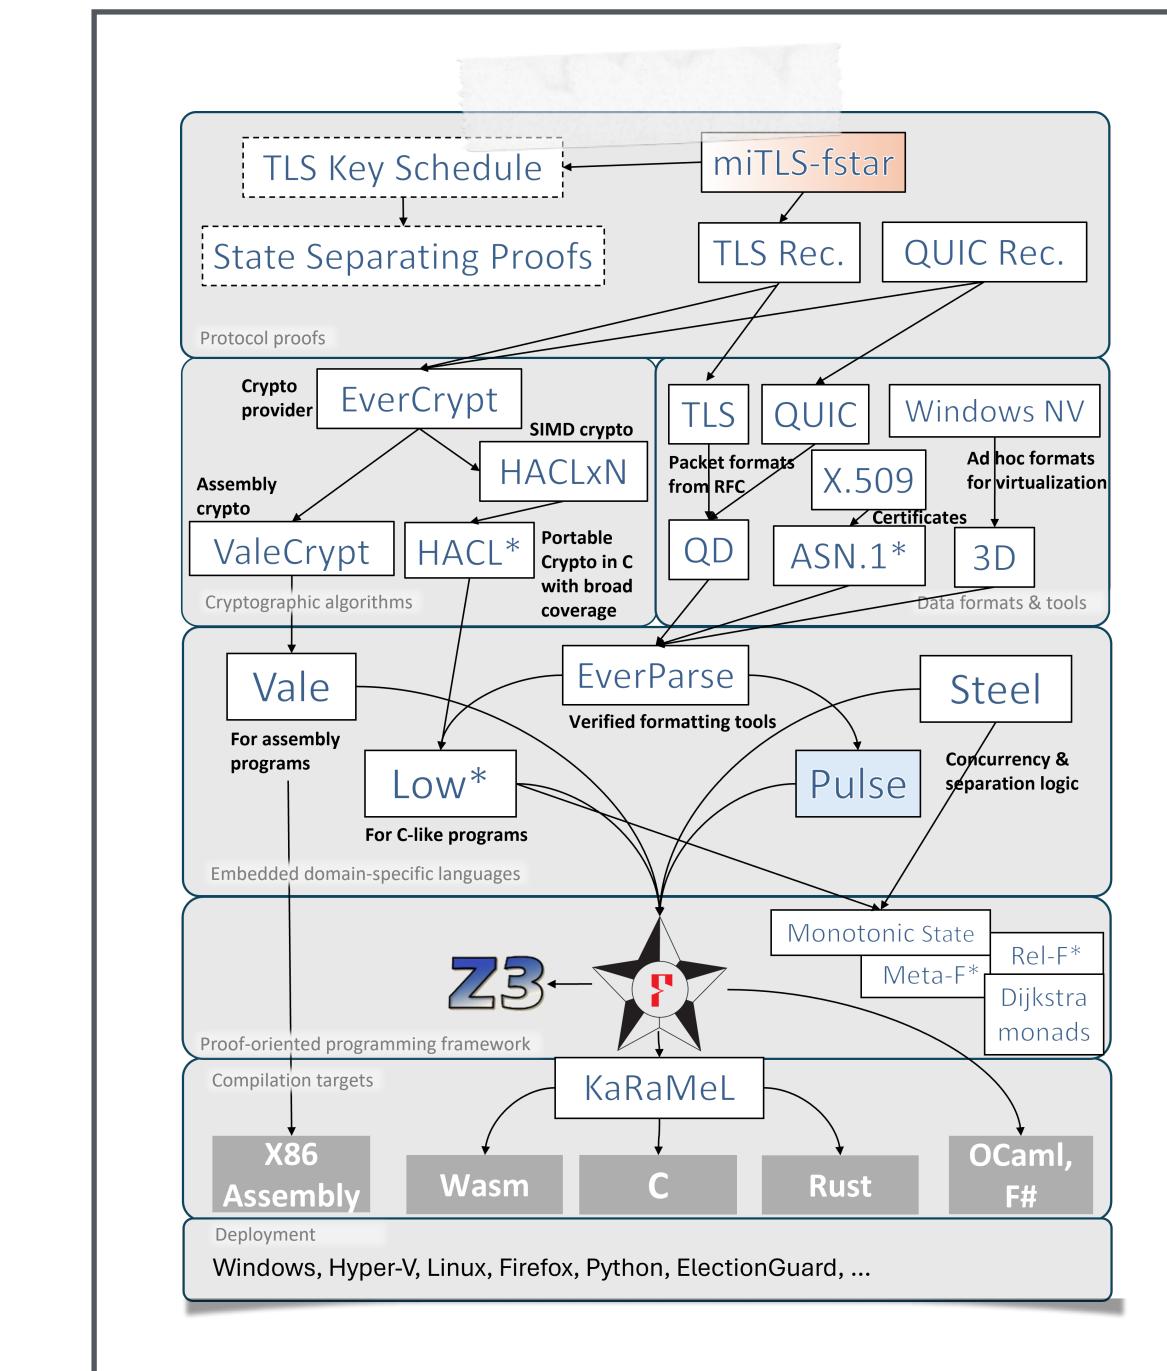


Rocq  
(formerly Coq)



verified kernels

verified compilers



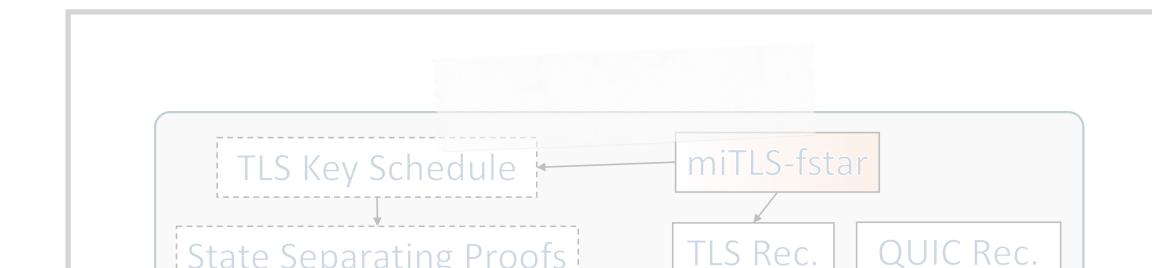
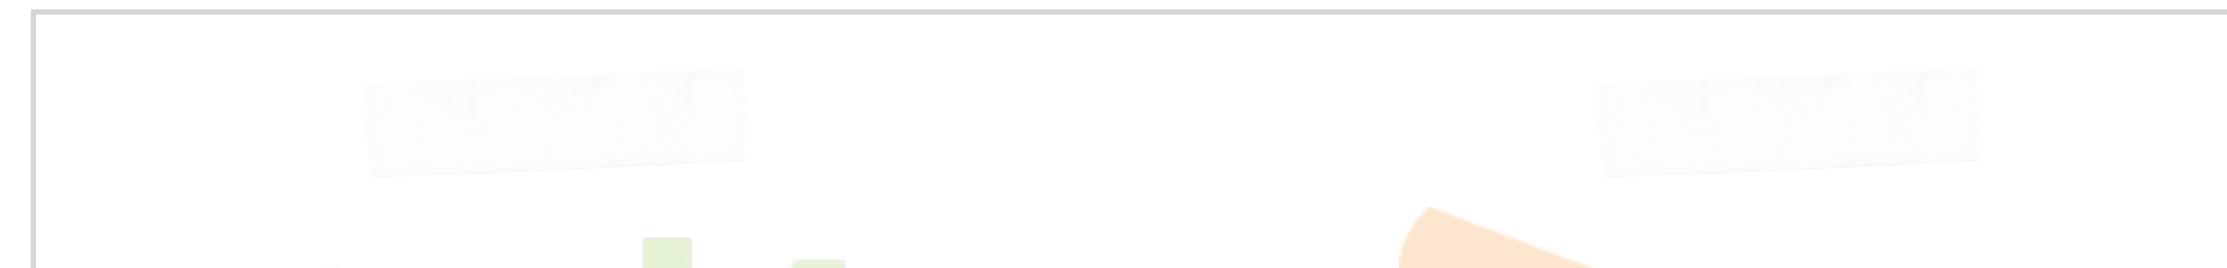
verified protocols

Dafny  
F\*

# Towards Safer Systems Software

formal verification to the rescue?

Isabelle/HOL  
HOL4



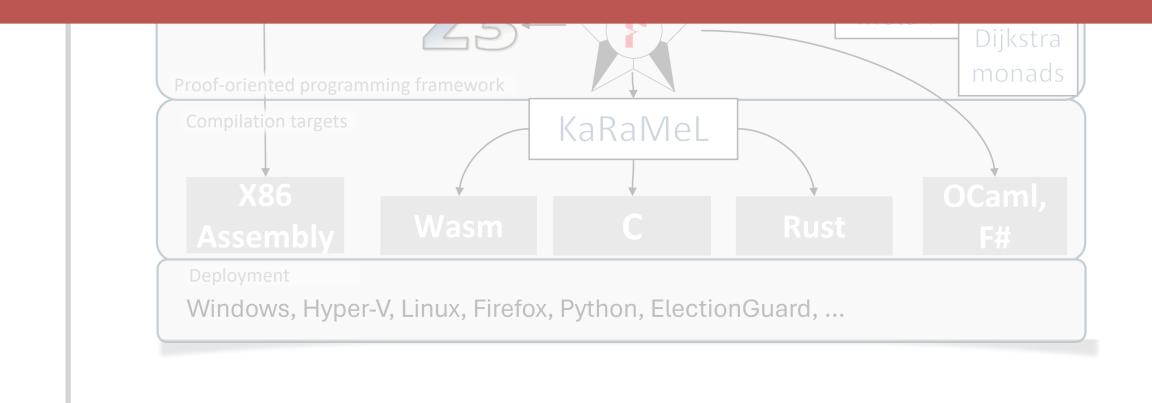
afny  
T\*

"theorem prover" sounds scary

Ro  
(formerly CoqJ)



COMPCERT



verified kernels

verified compilers

verified protocols

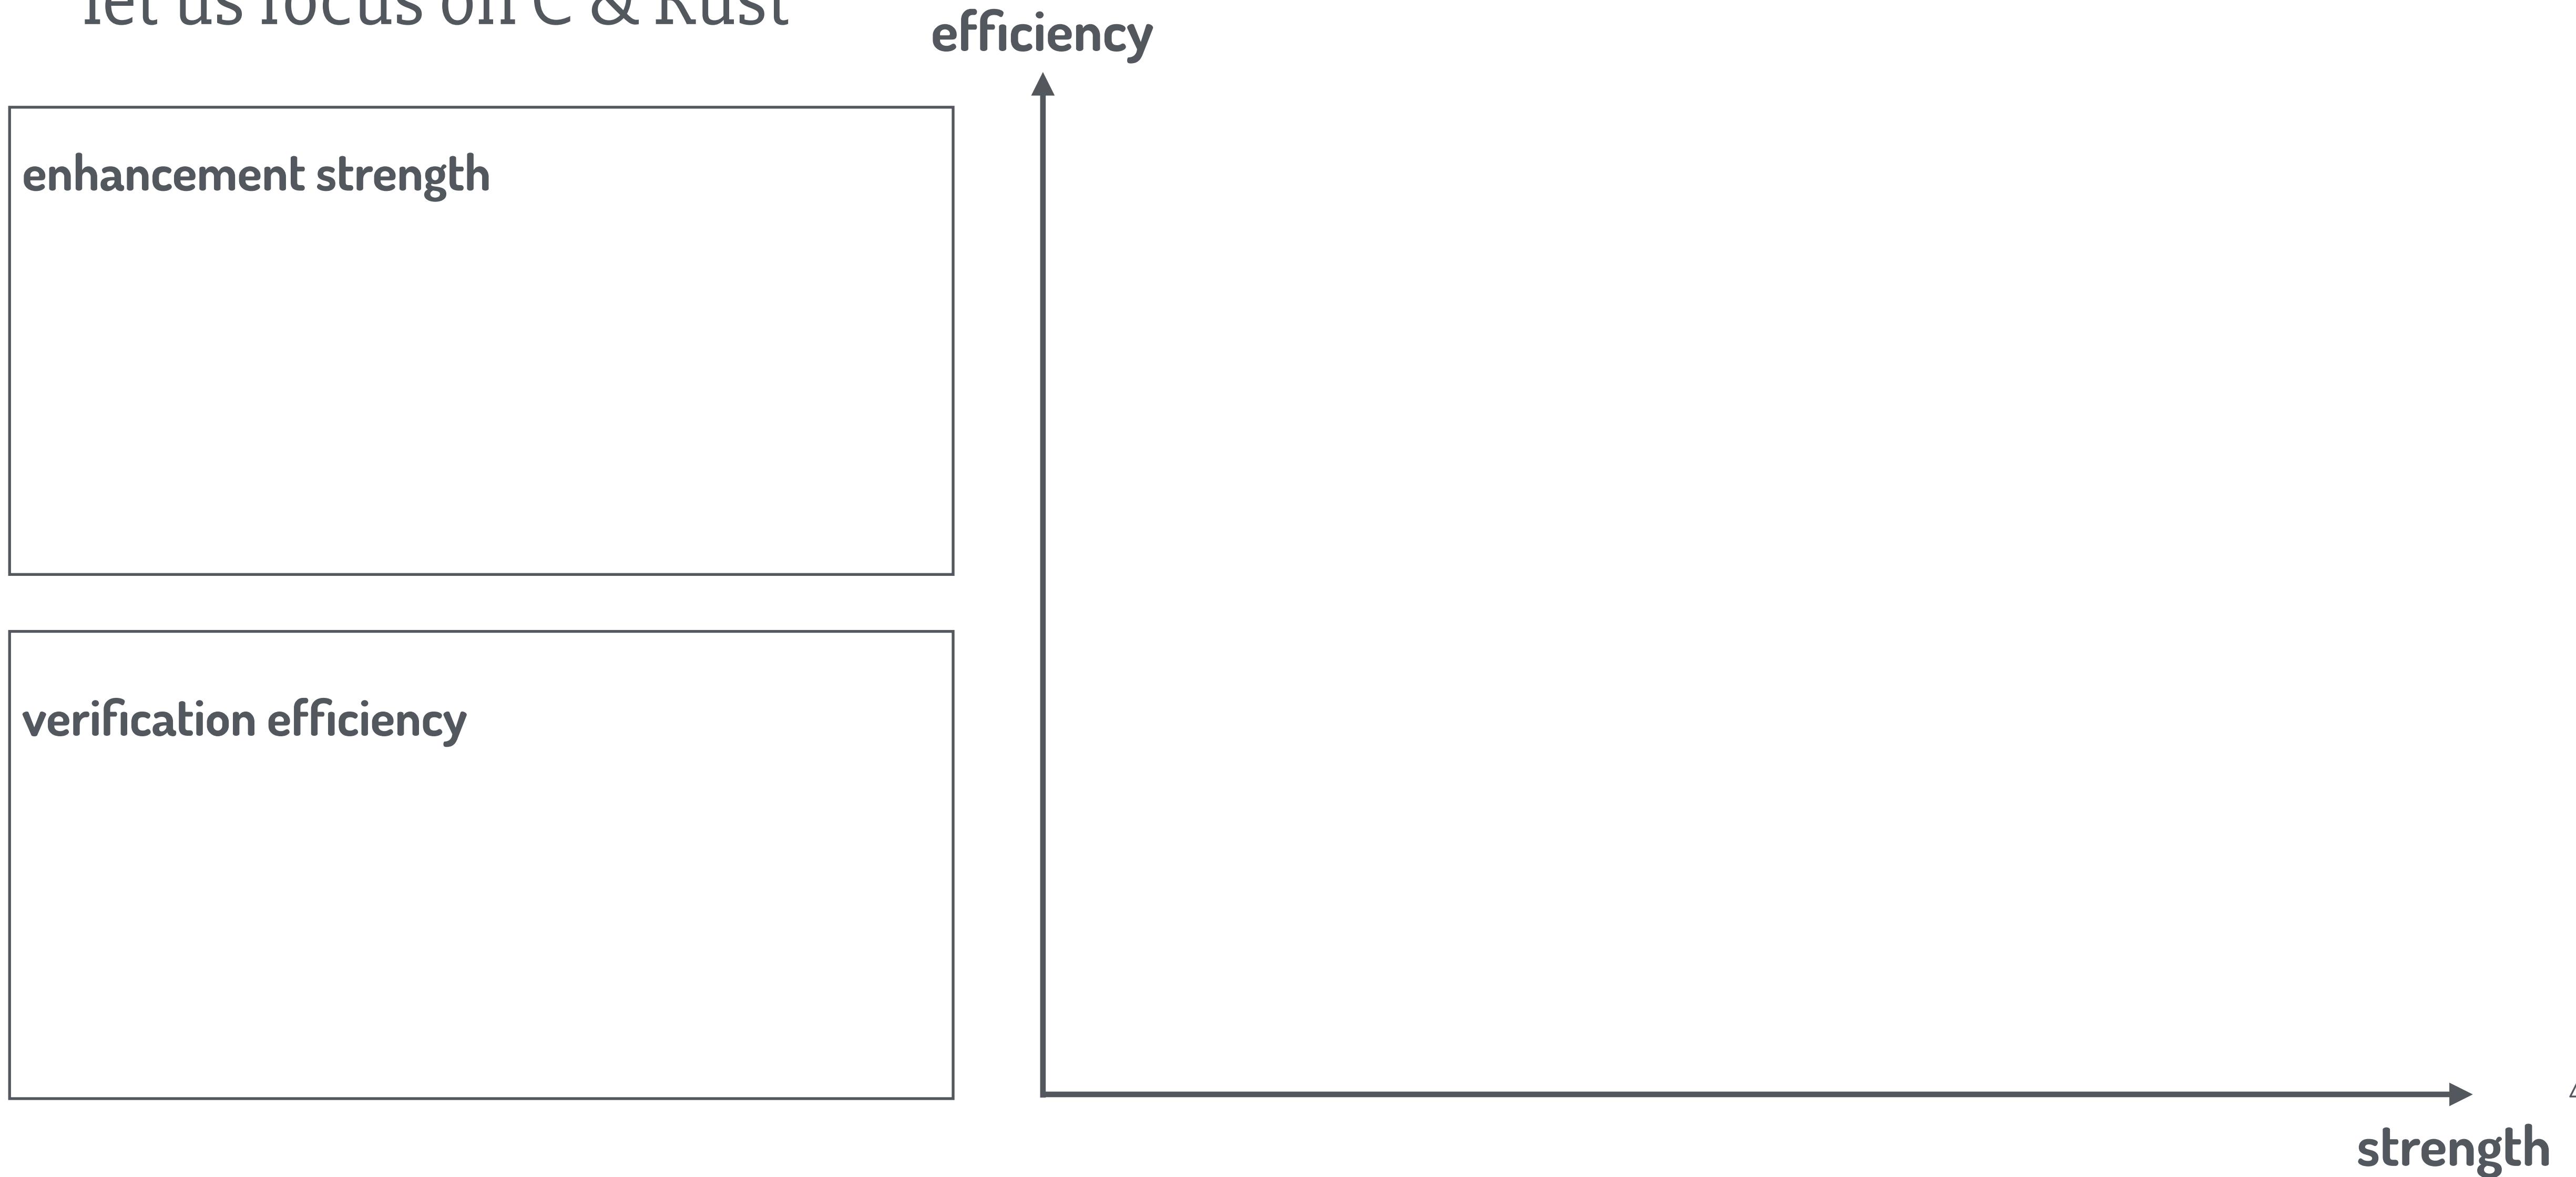
# Language Enhancement?

let us focus on C & Rust



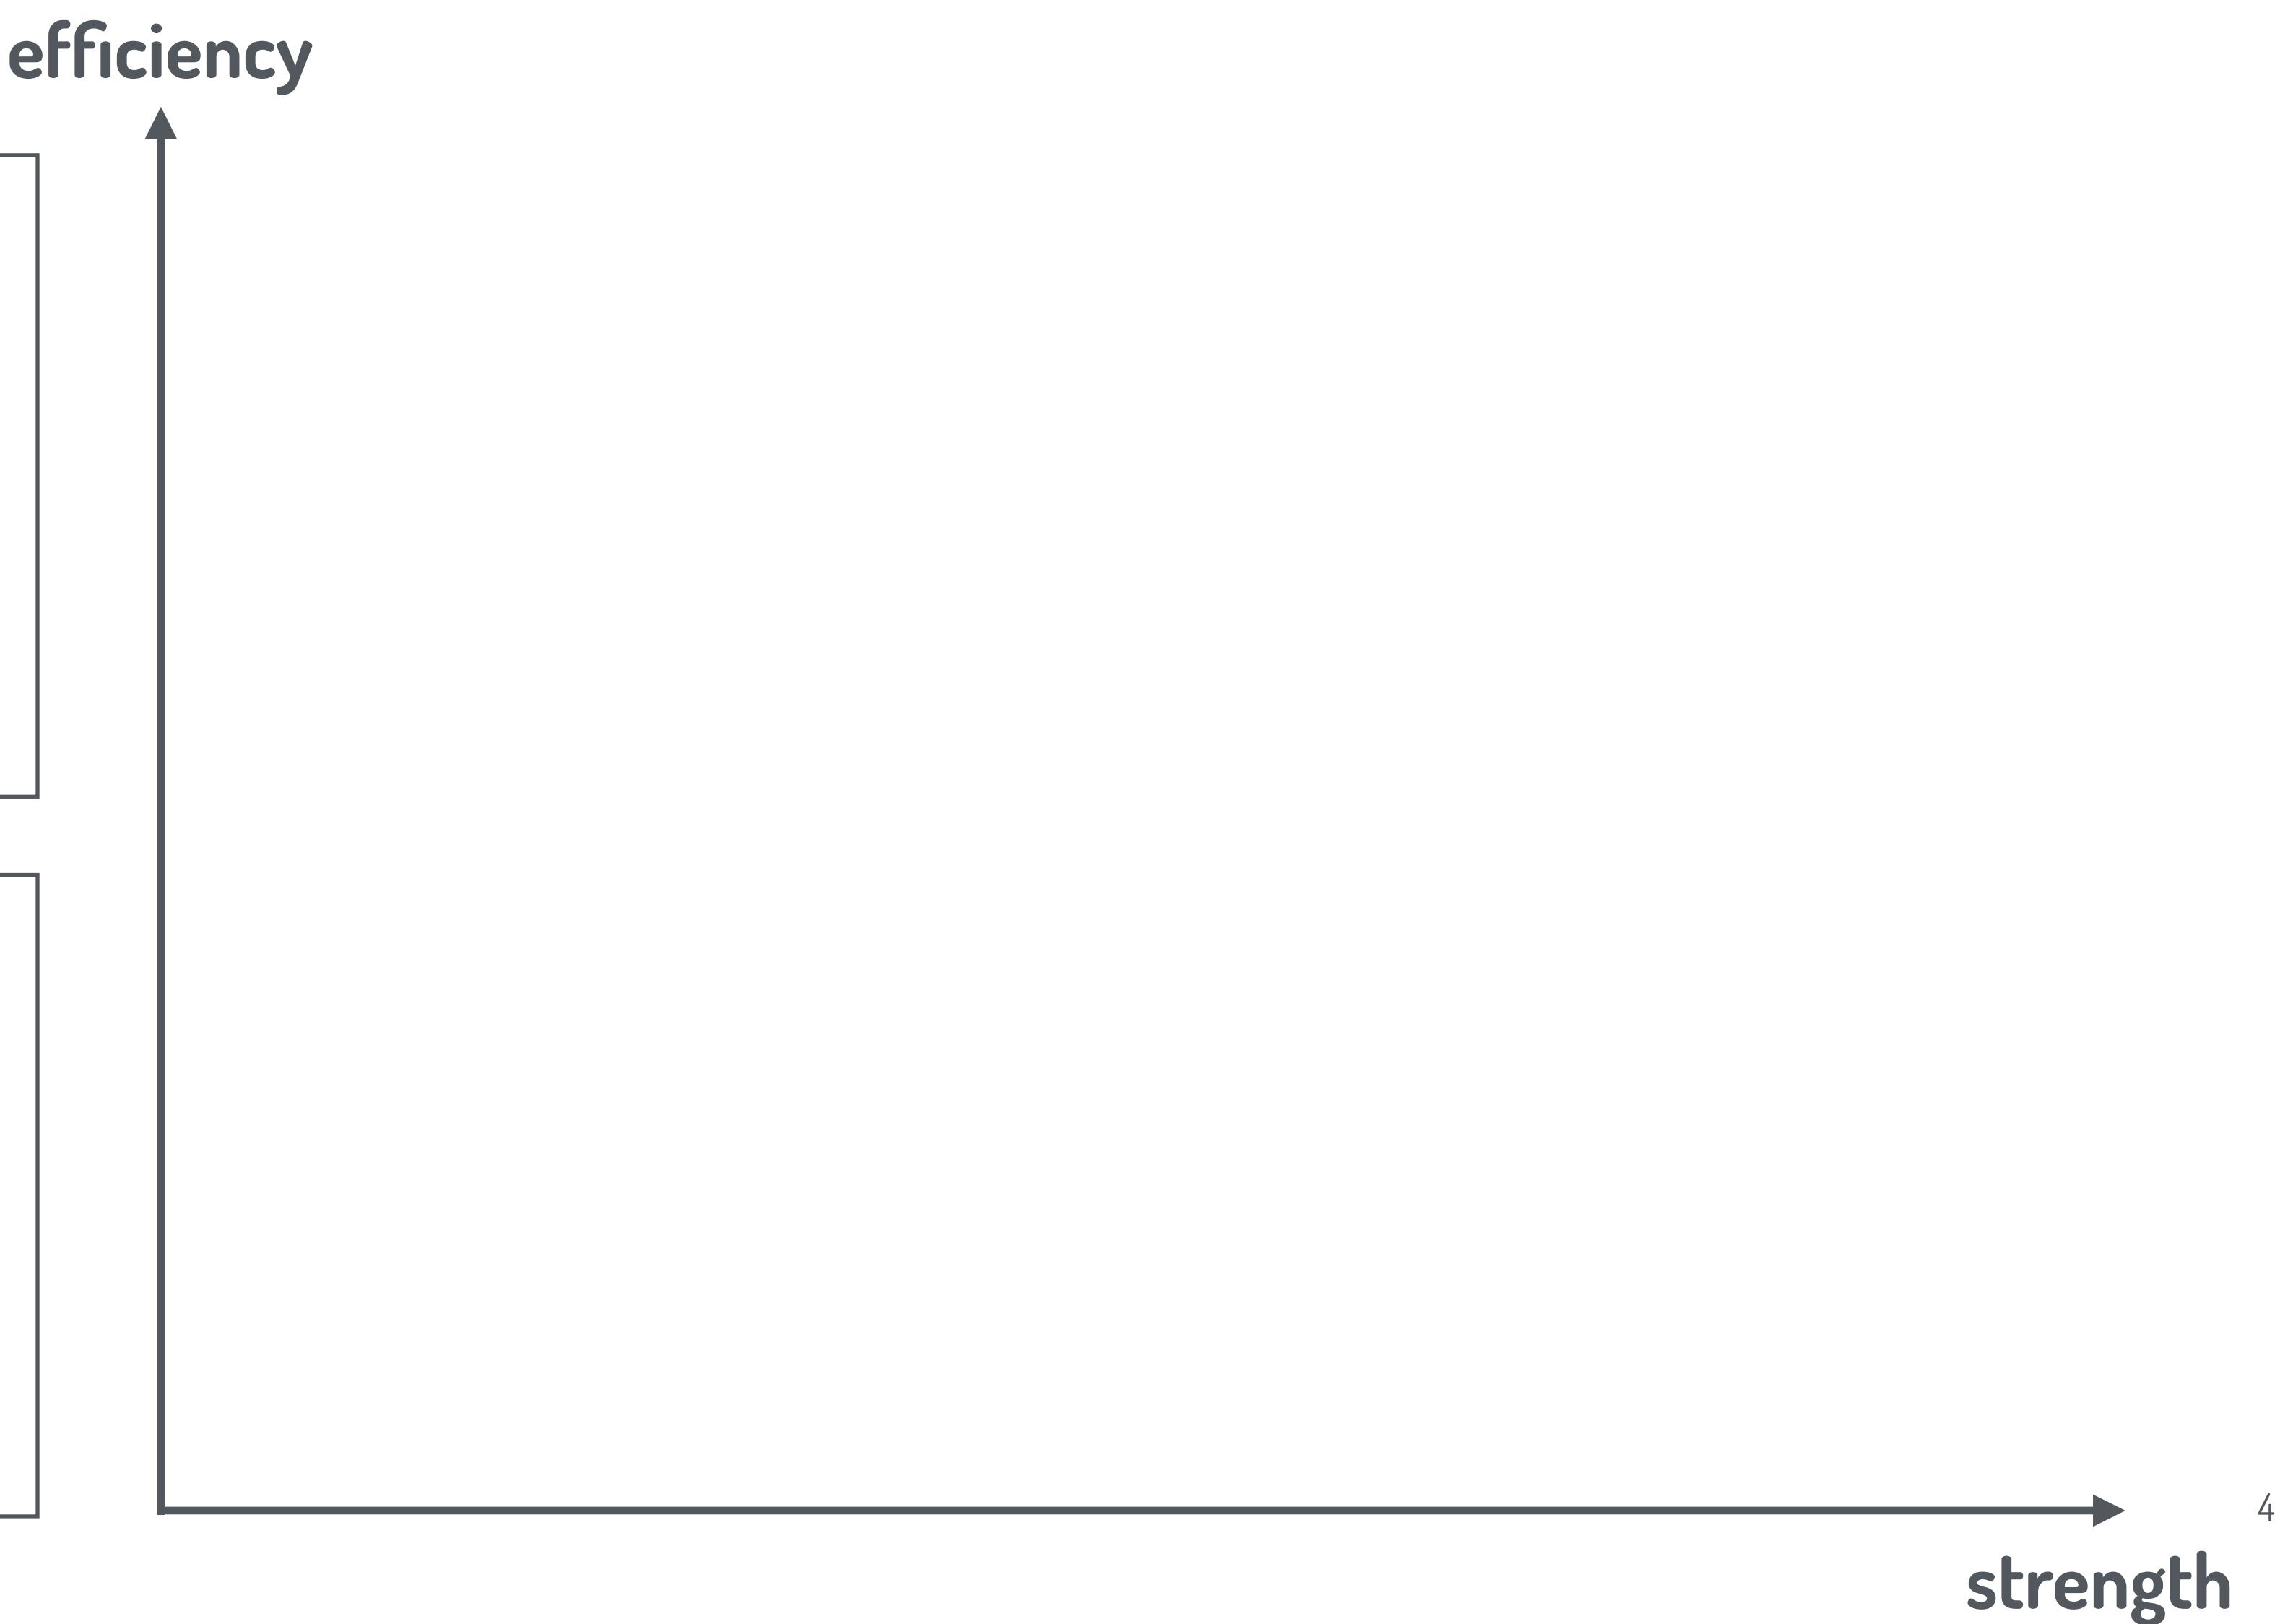
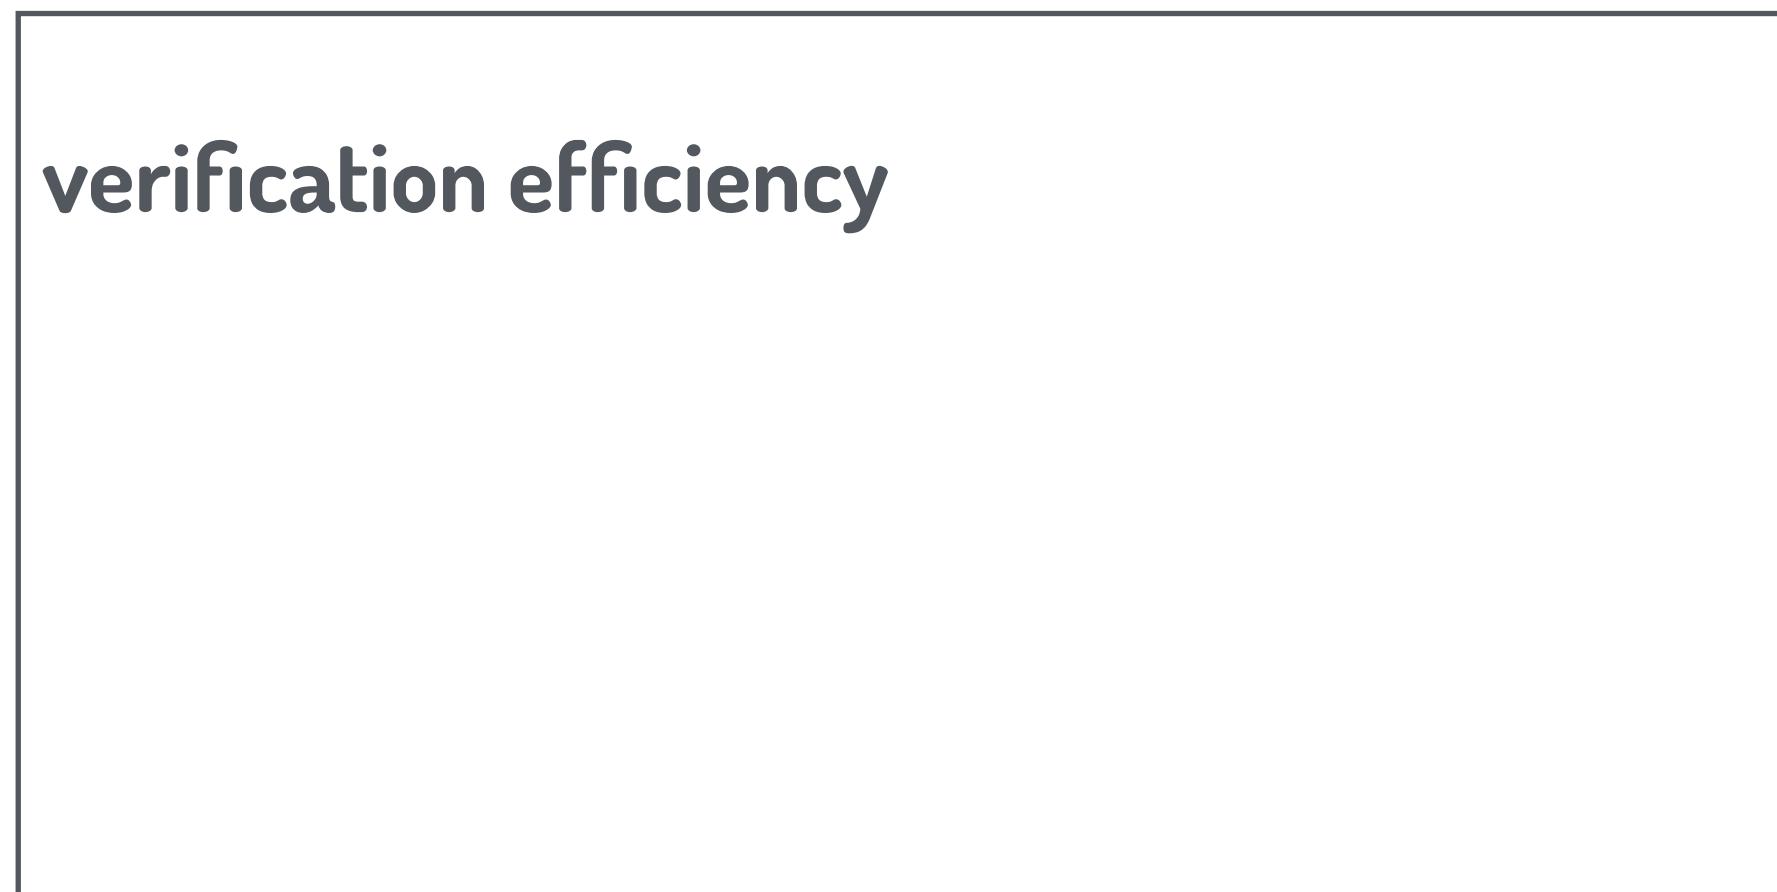
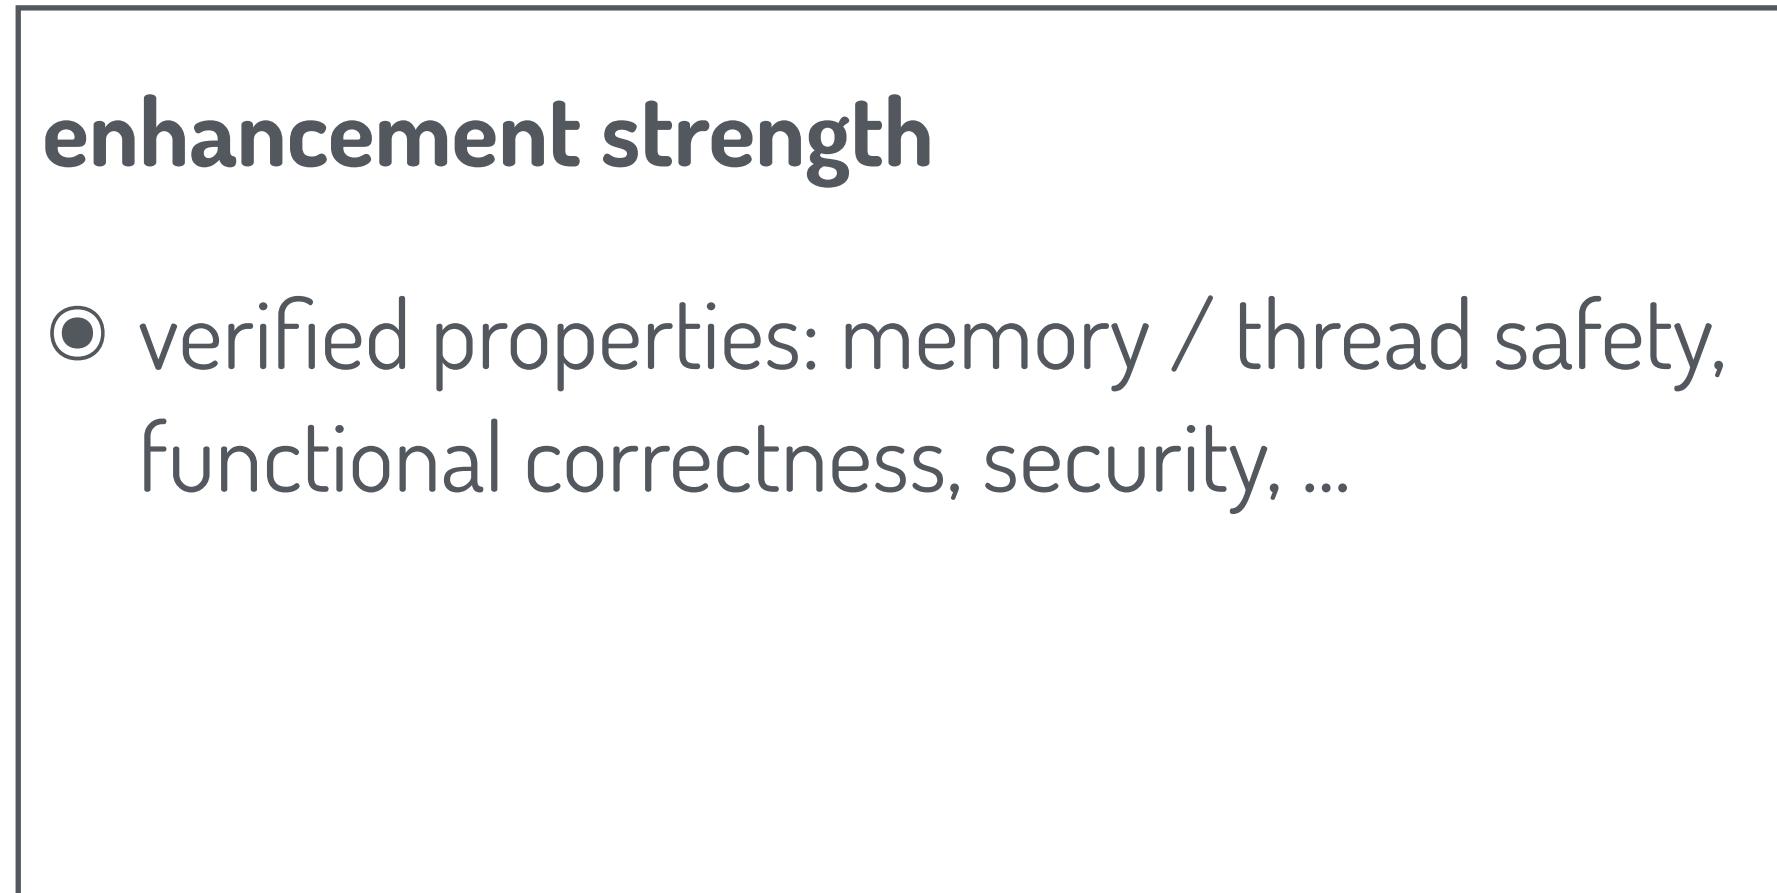
# Language Enhancement?

let us focus on C & Rust



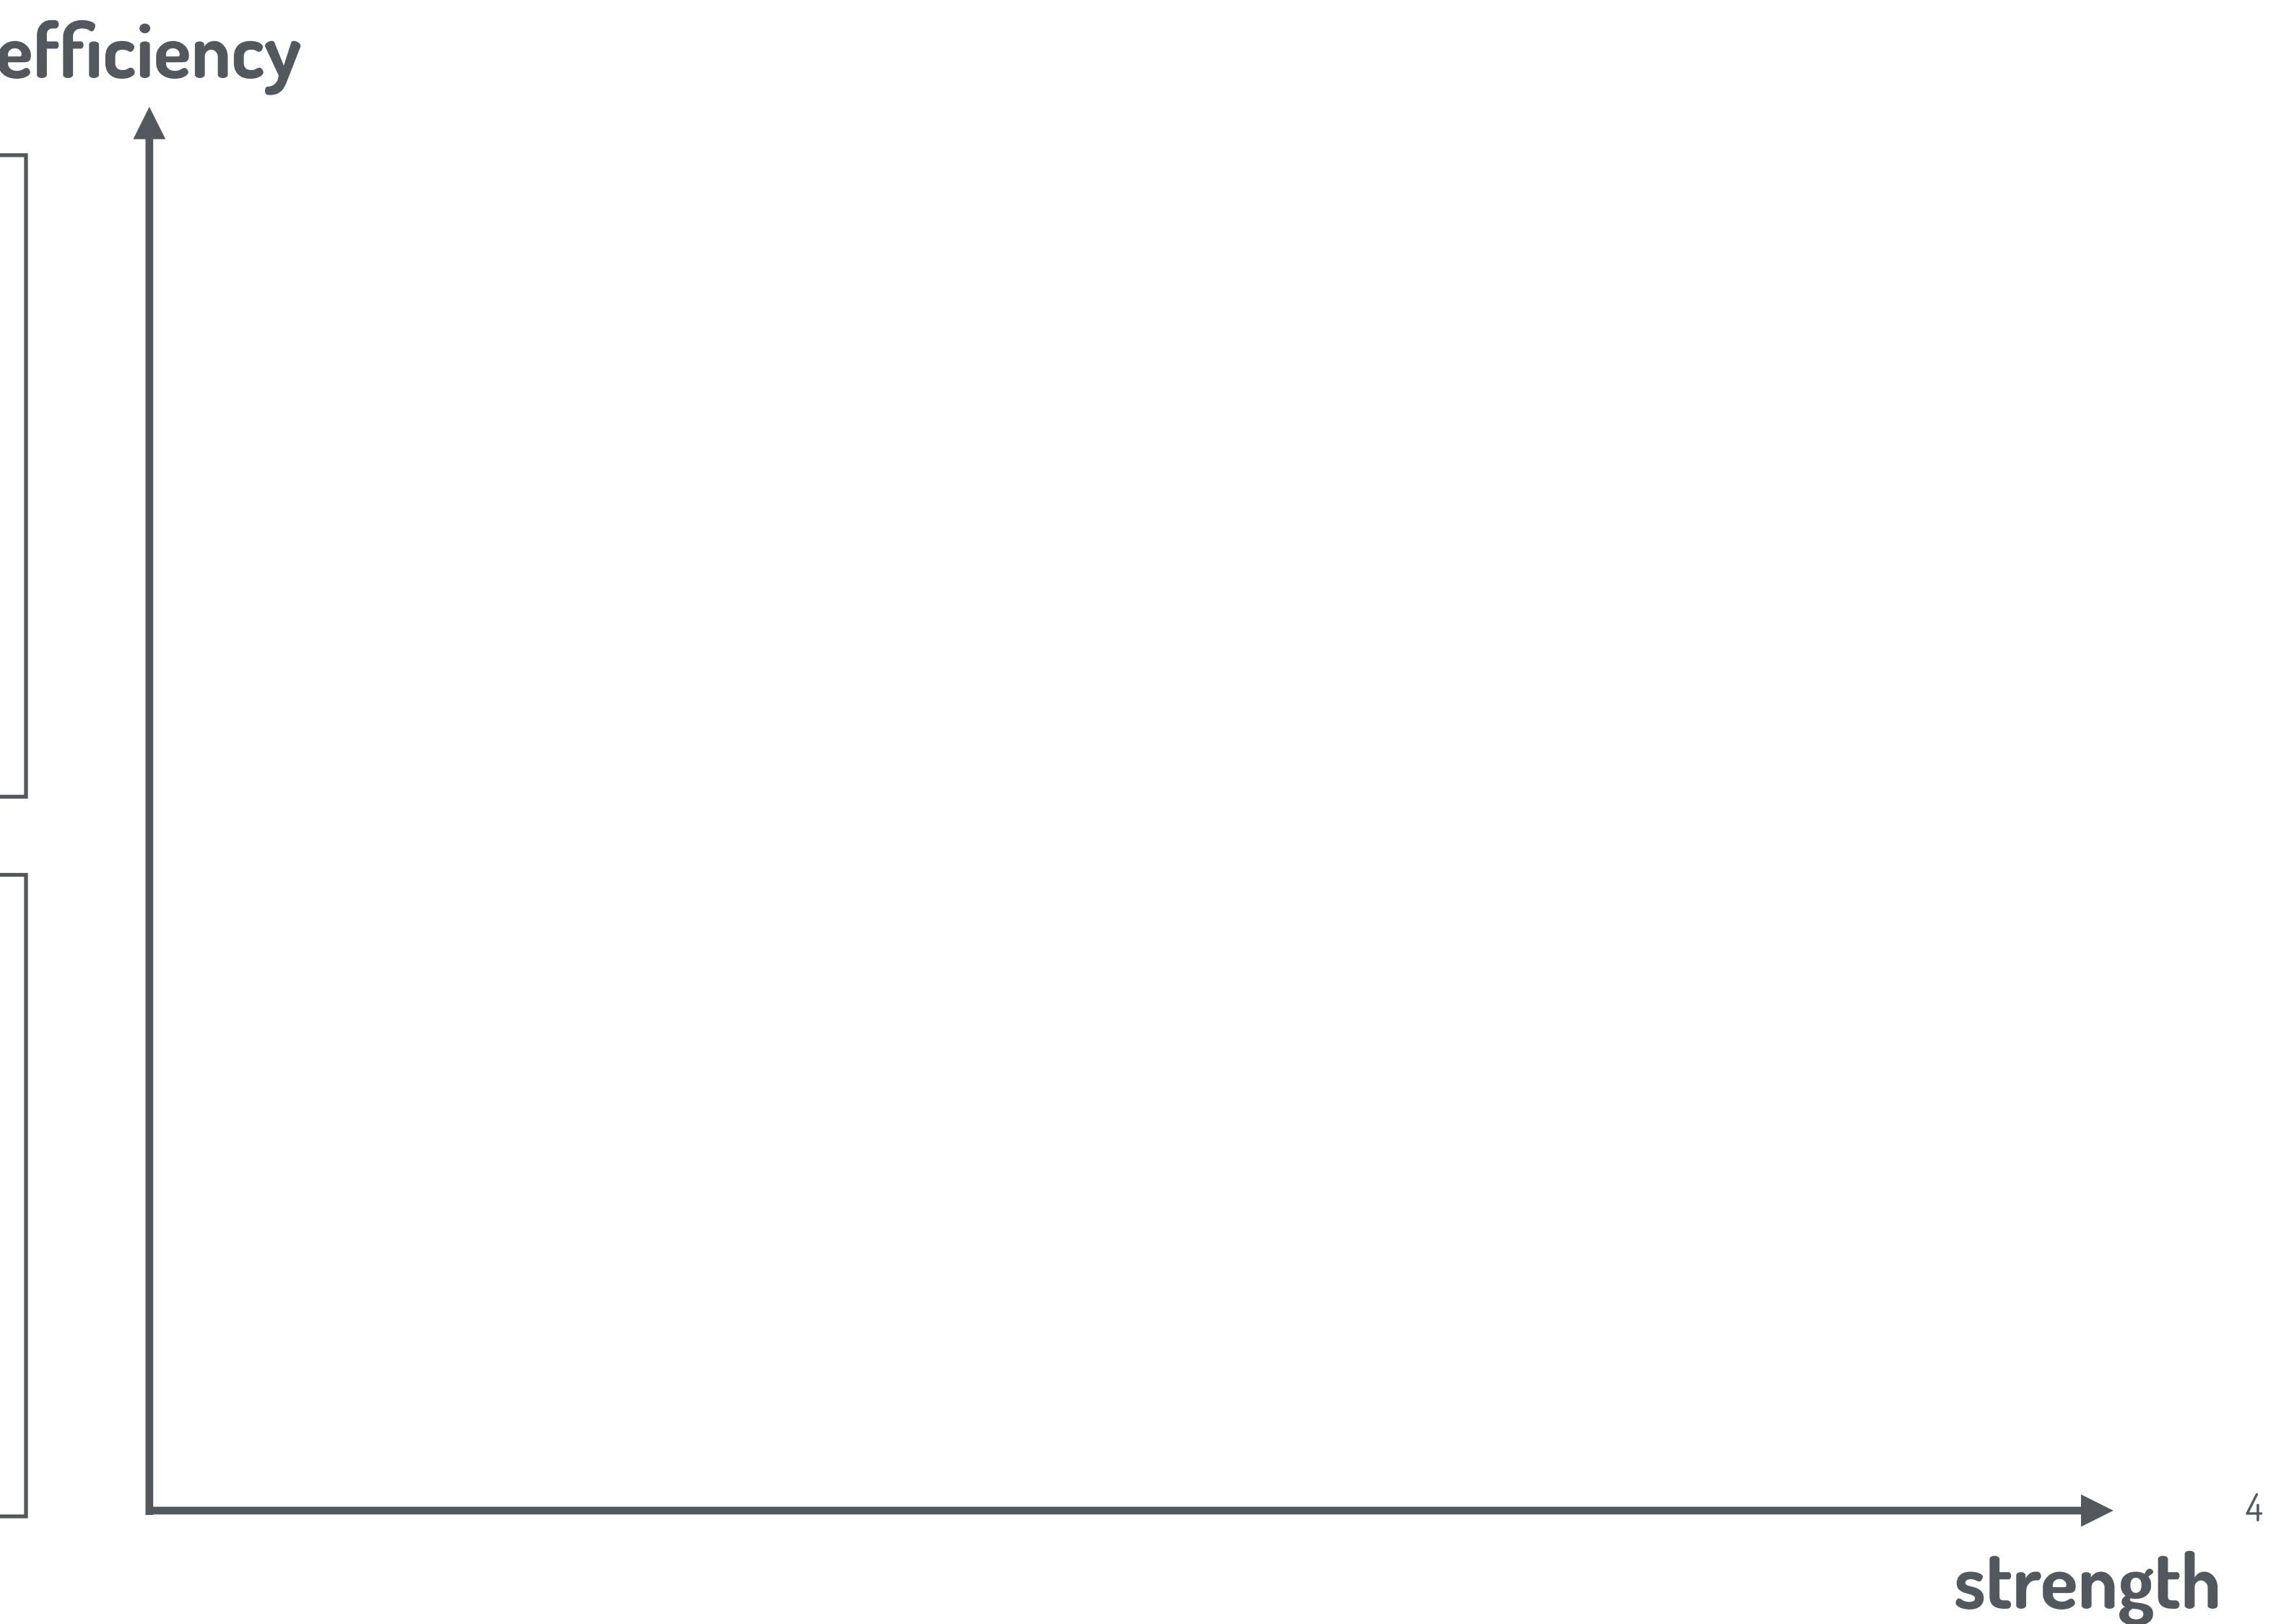
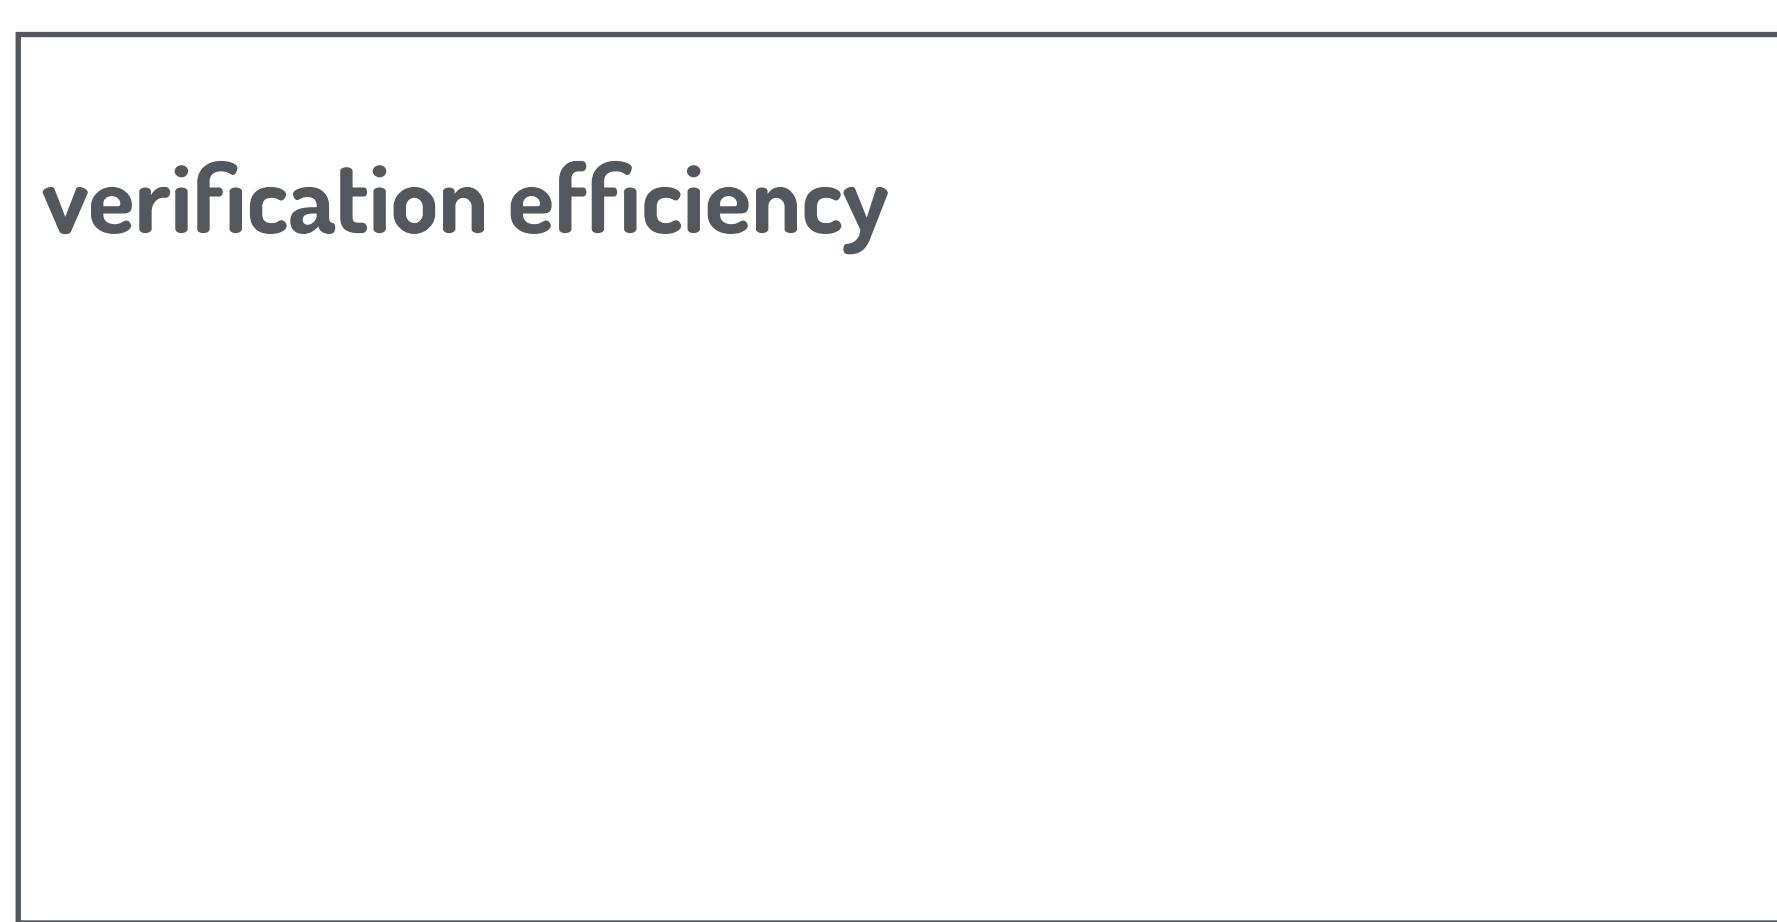
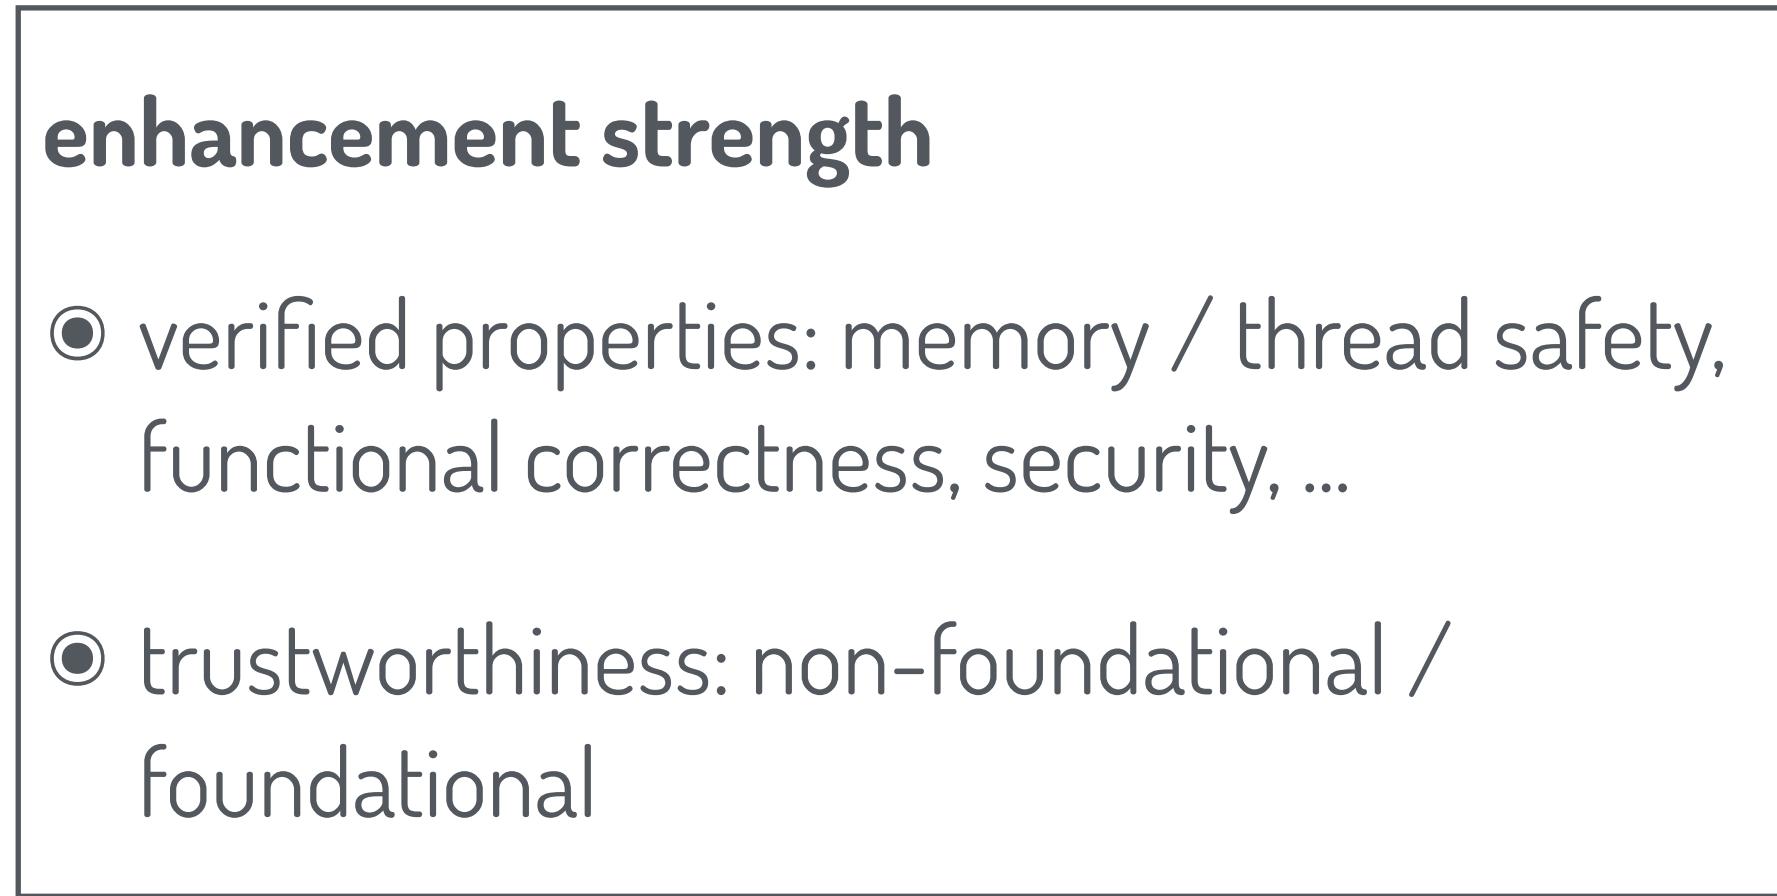
# Language Enhancement?

let us focus on C & Rust



# Language Enhancement?

let us focus on C & Rust



# Language Enhancement?

let us focus on C & Rust

## enhancement strength

- verified properties: memory / thread safety, functional correctness, security, ...
- trustworthiness: non-foundational / foundational

## verification efficiency

- learning curve

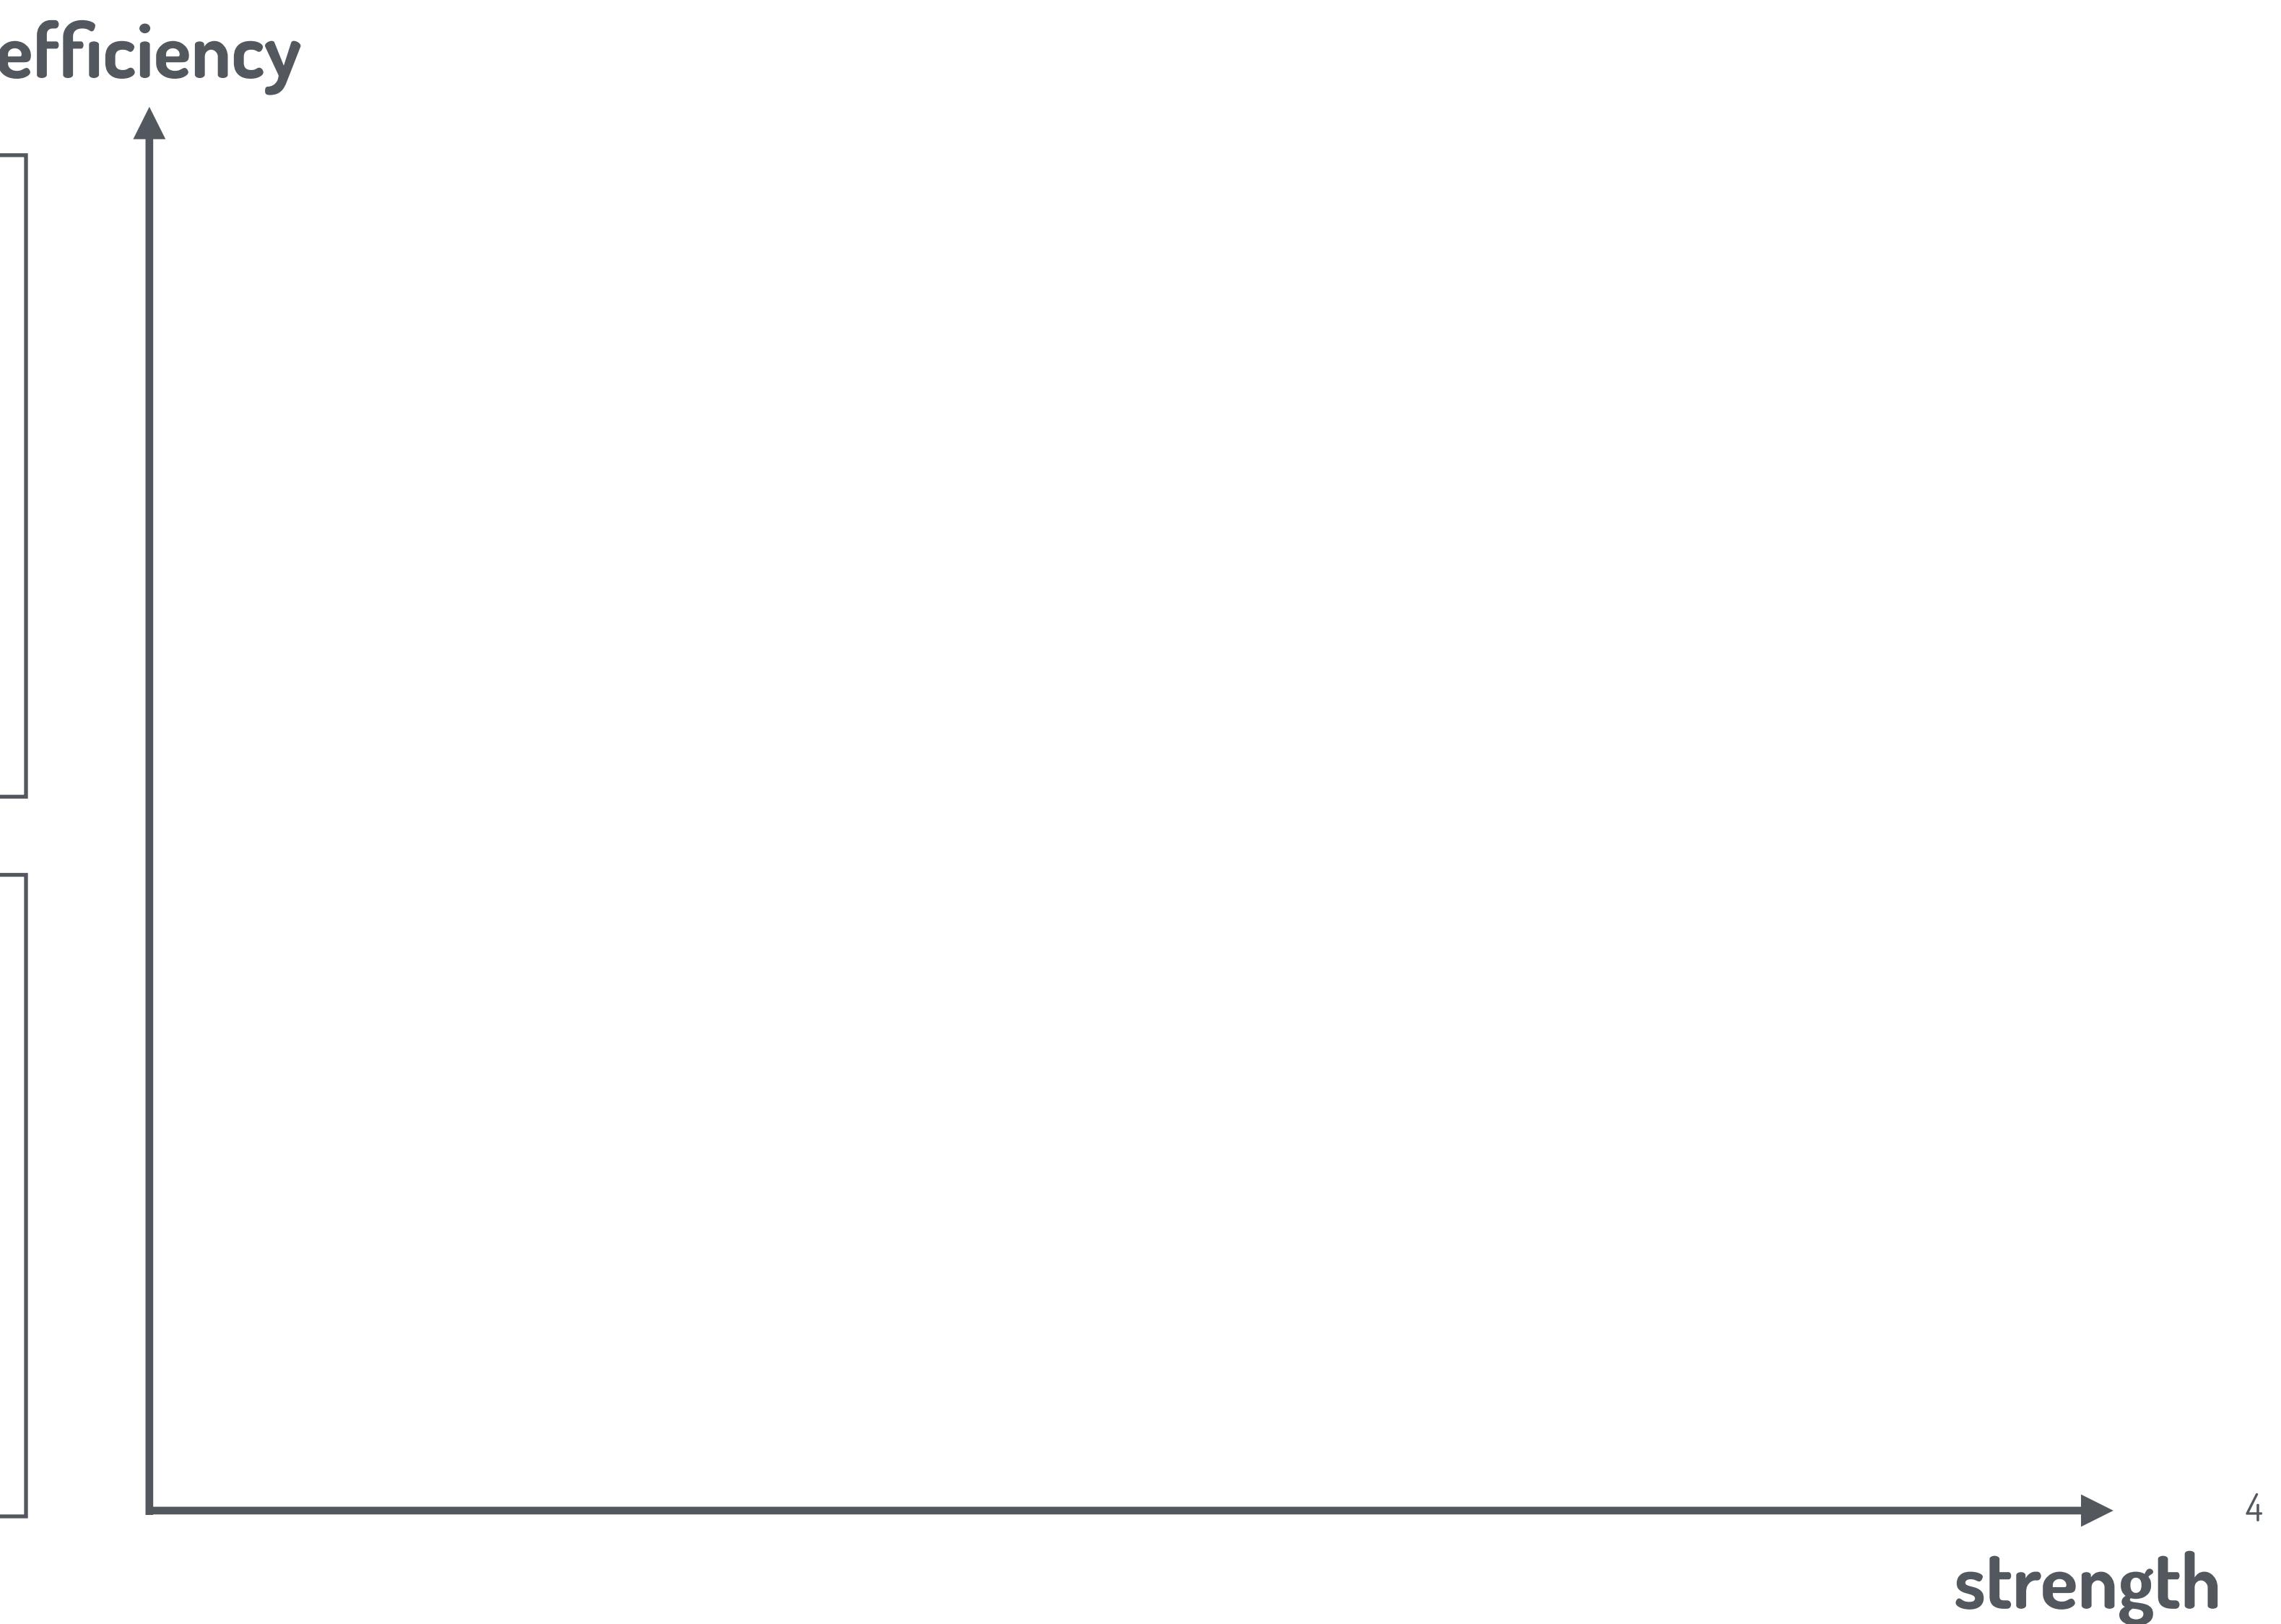
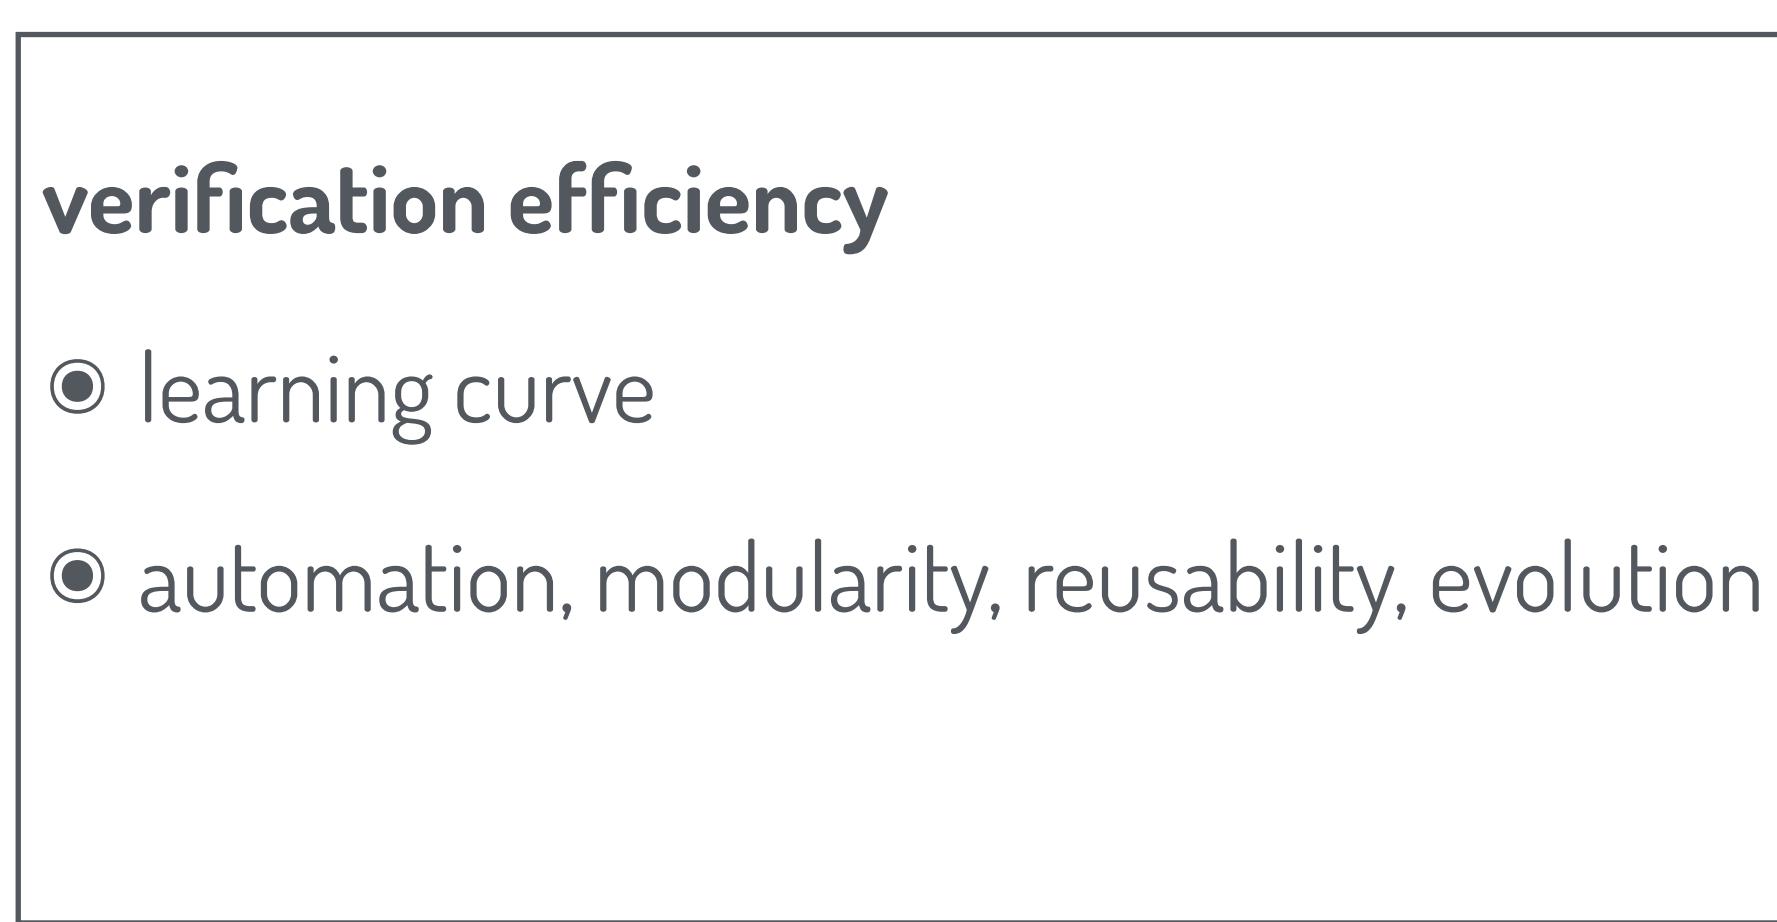
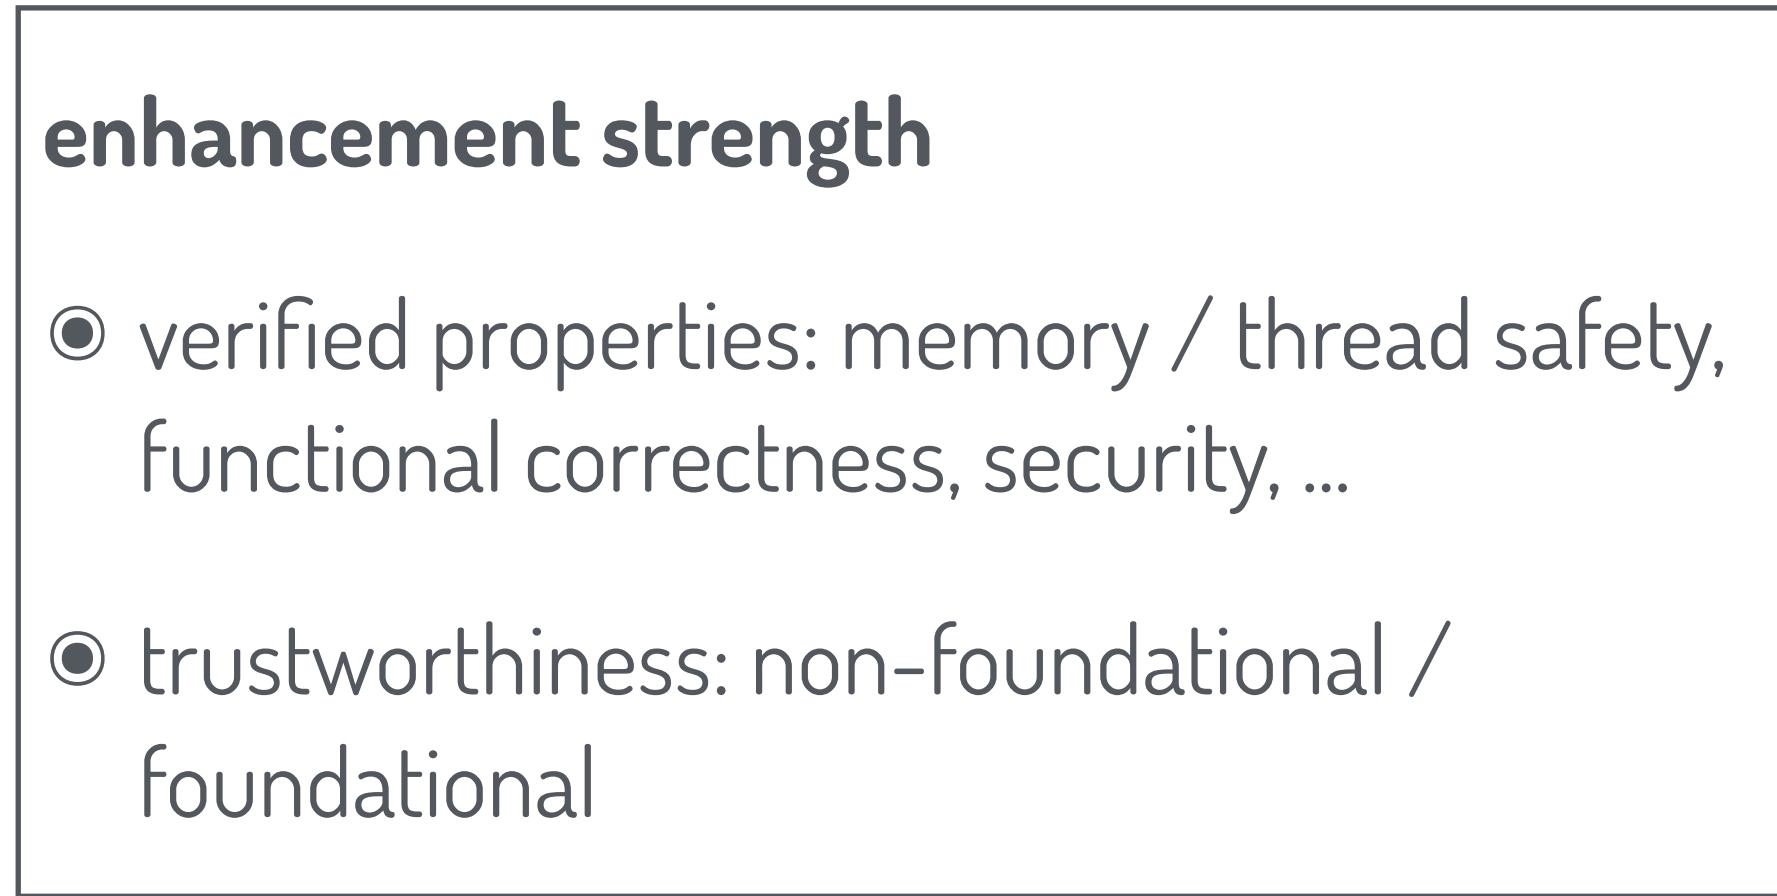
efficiency



strength

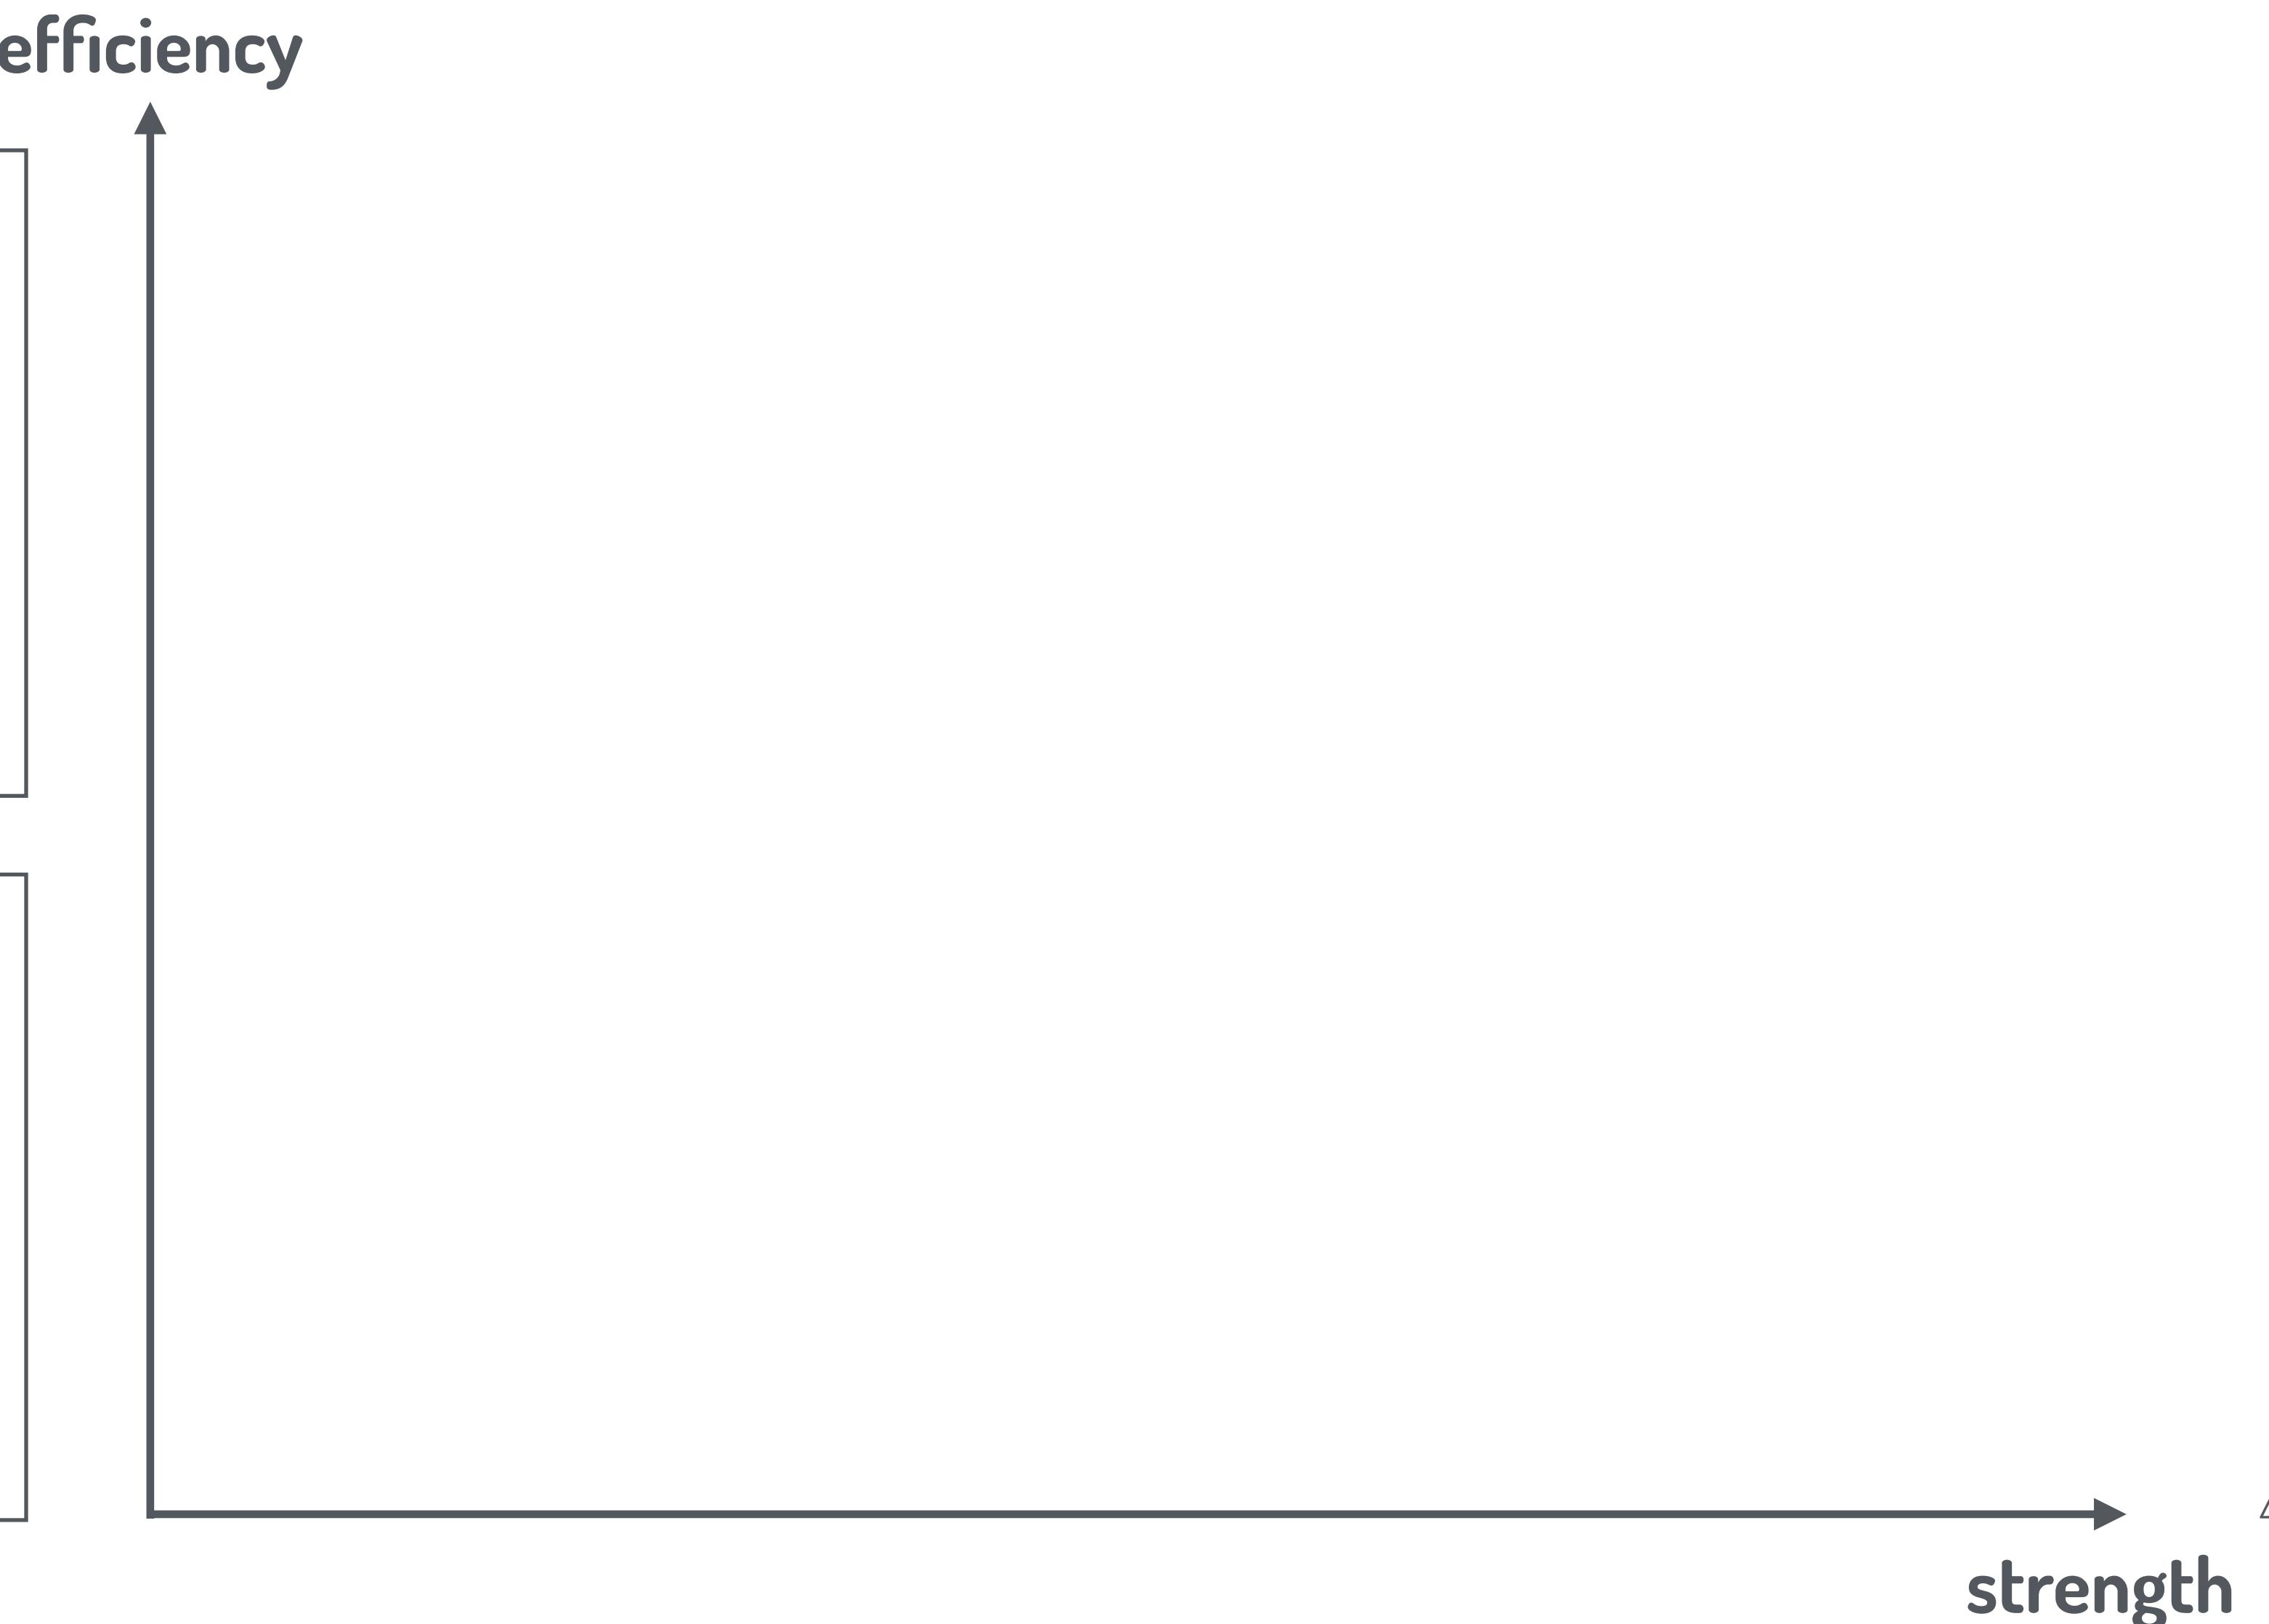
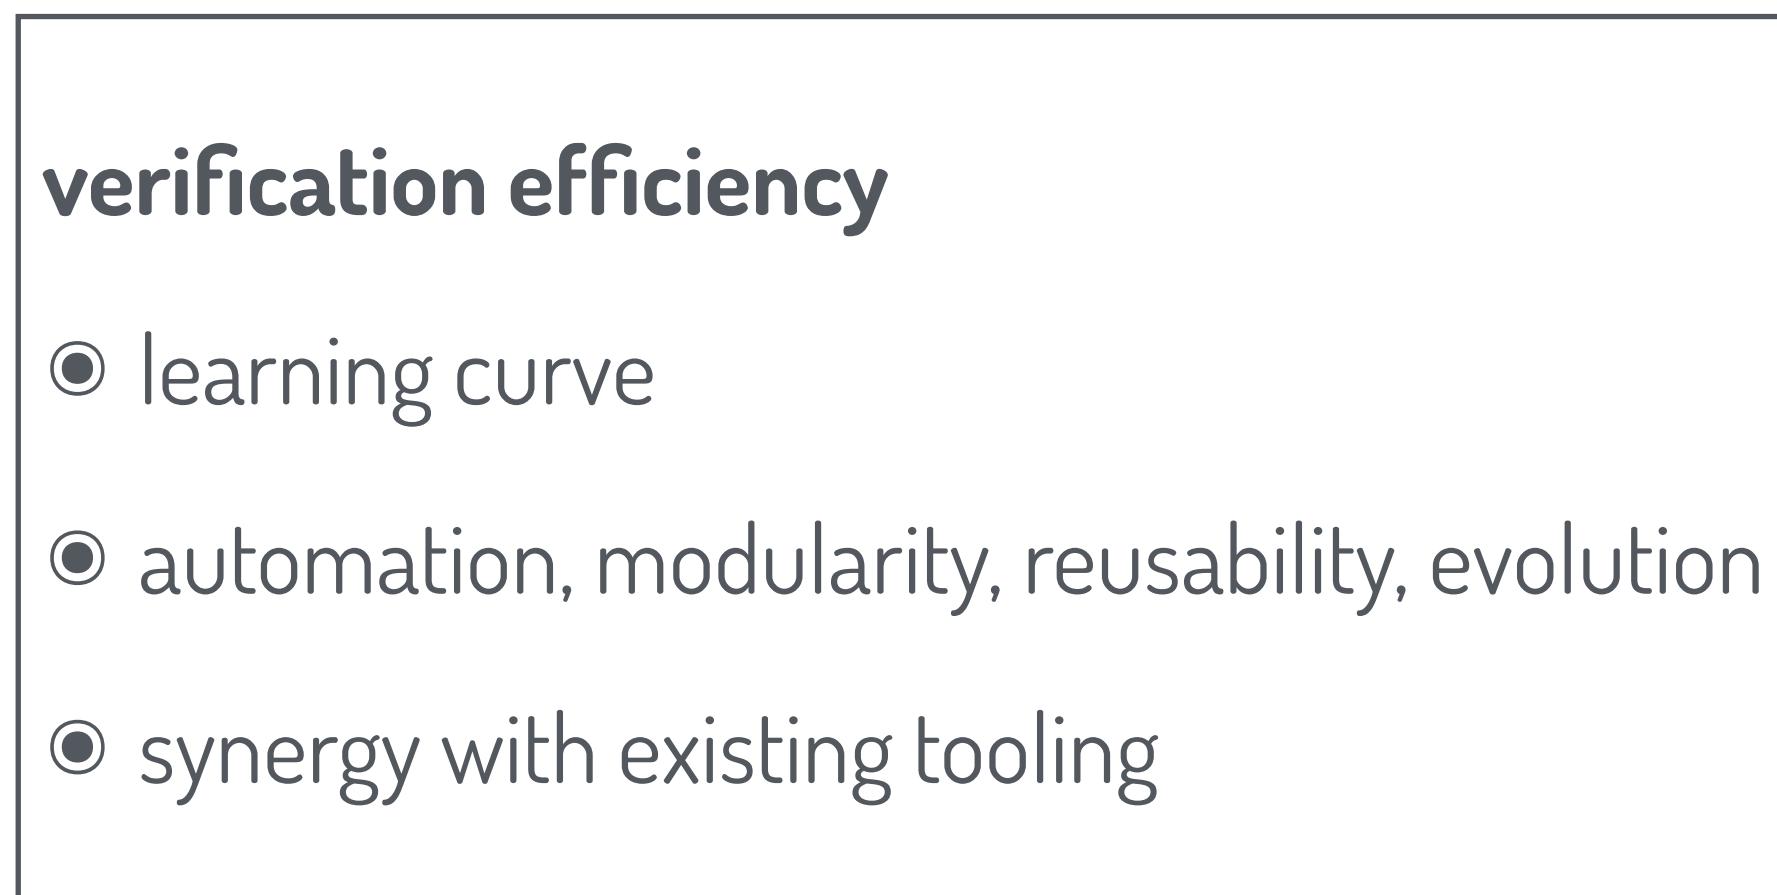
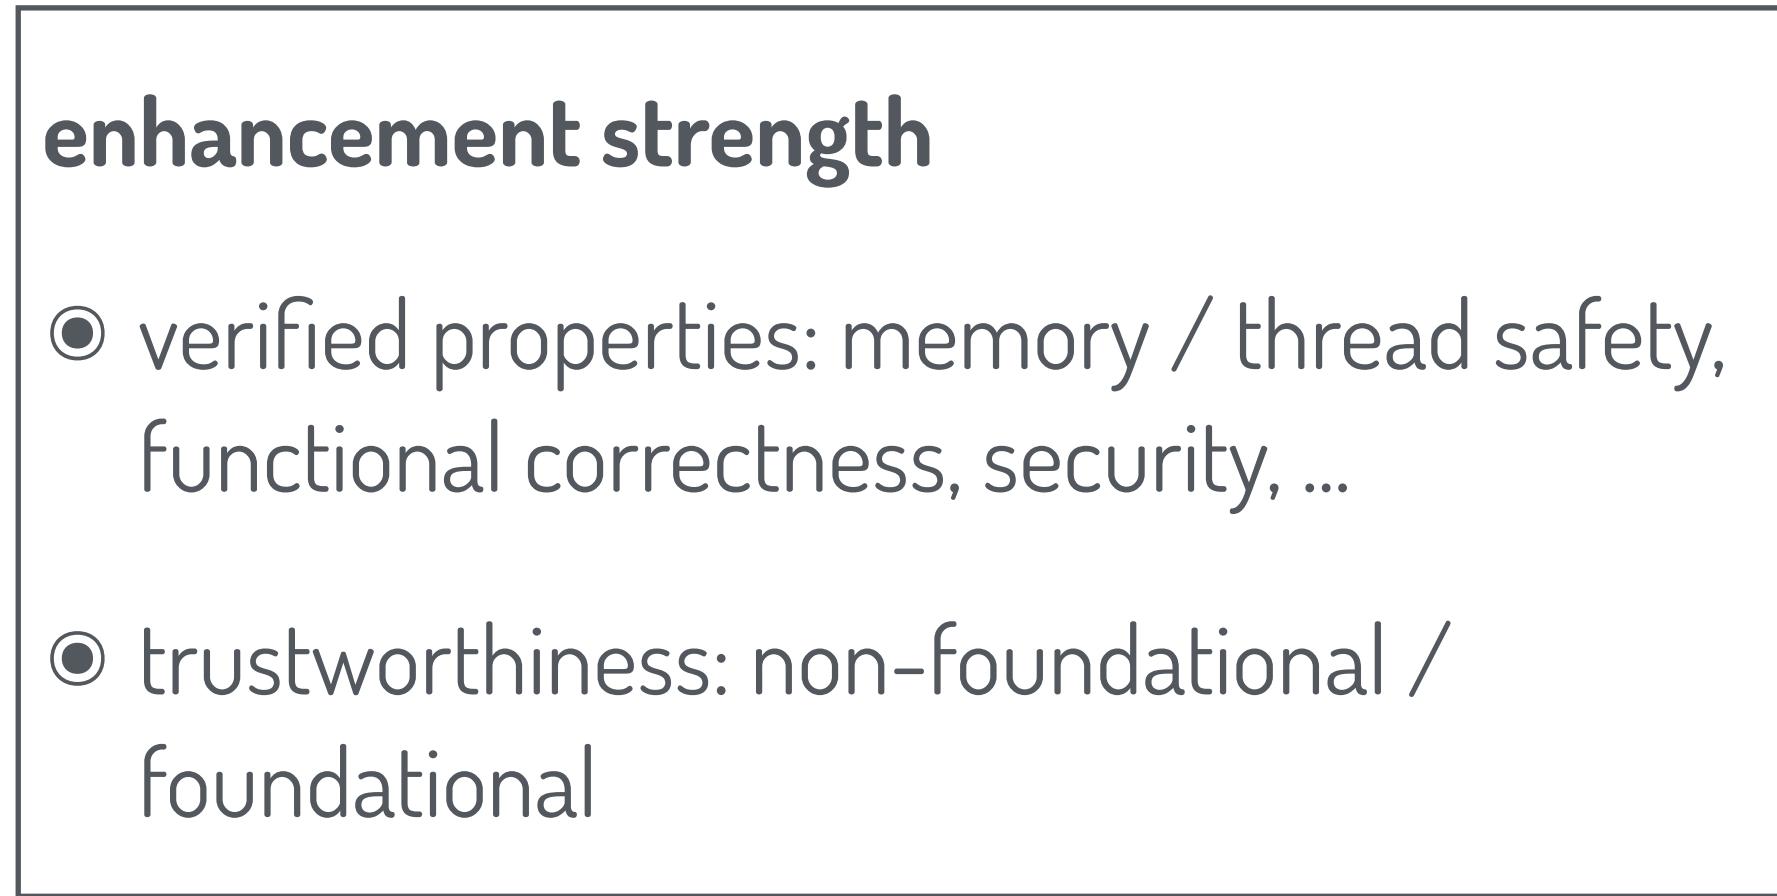
# Language Enhancement?

let us focus on C & Rust



# Language Enhancement?

let us focus on C & Rust



# Language Enhancement?

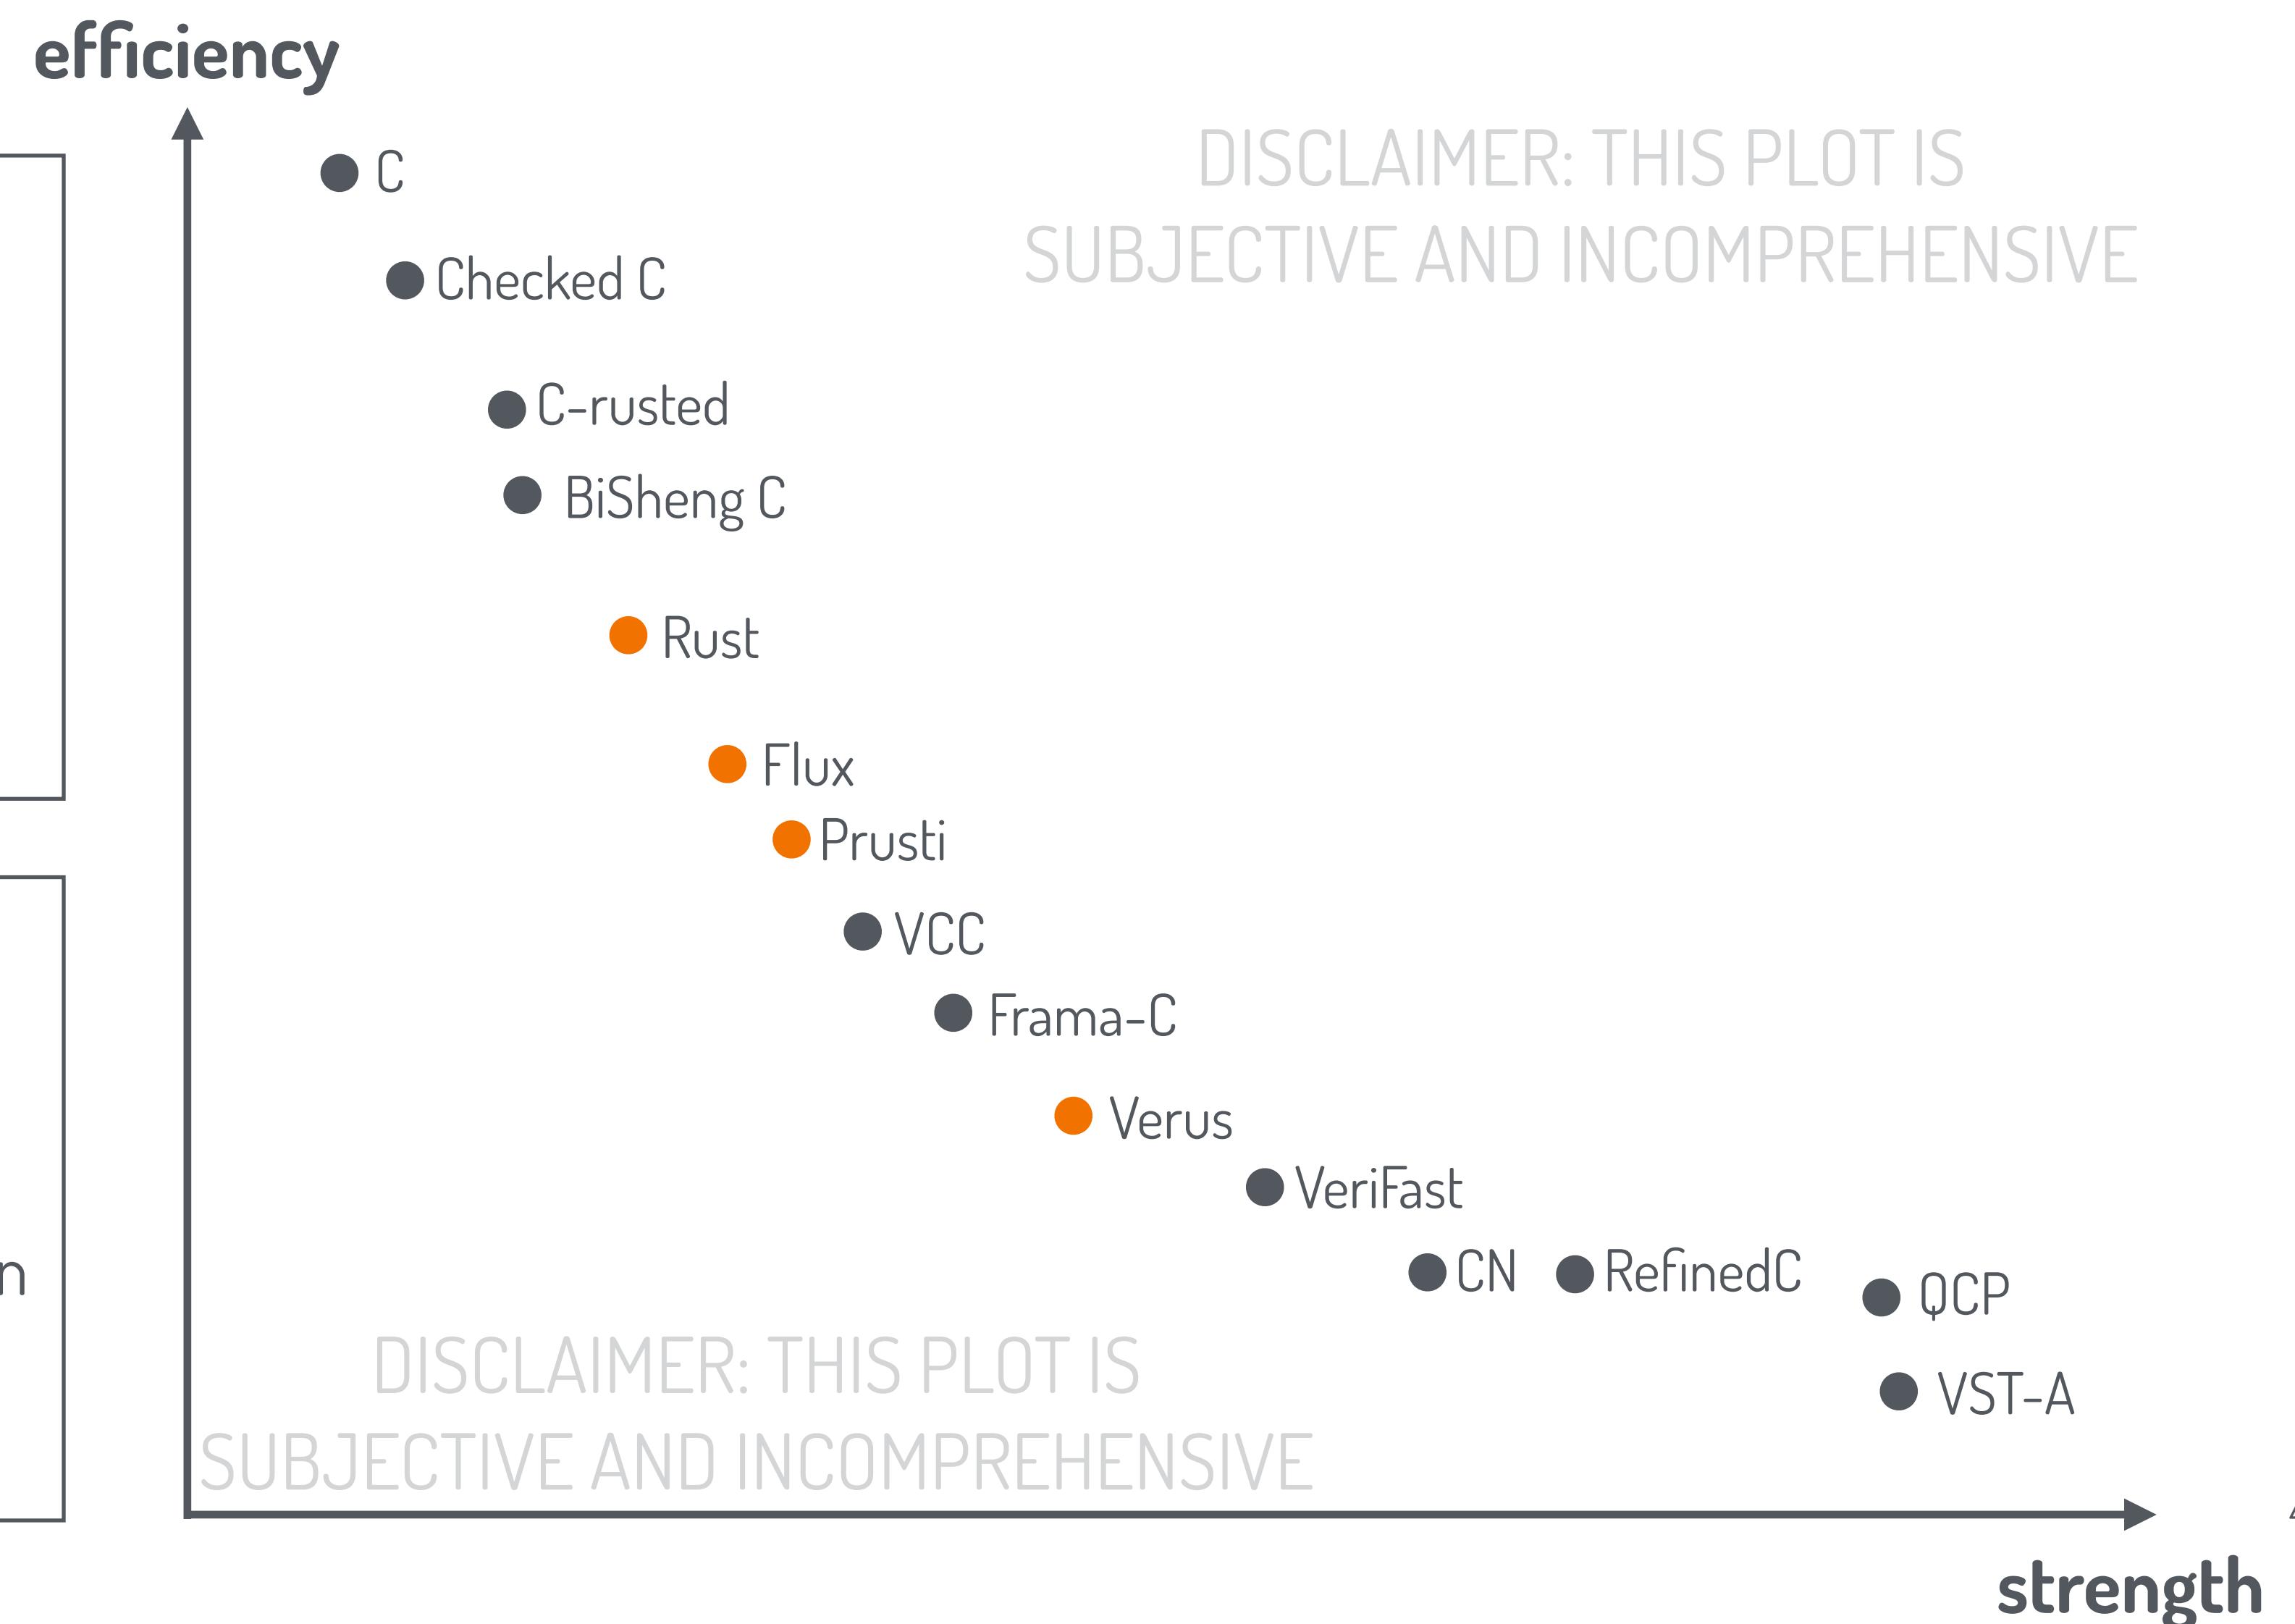
let us focus on C & Rust

## enhancement strength

- verified properties: memory / thread safety, functional correctness, security, ...
- trustworthiness: non-foundational / foundational

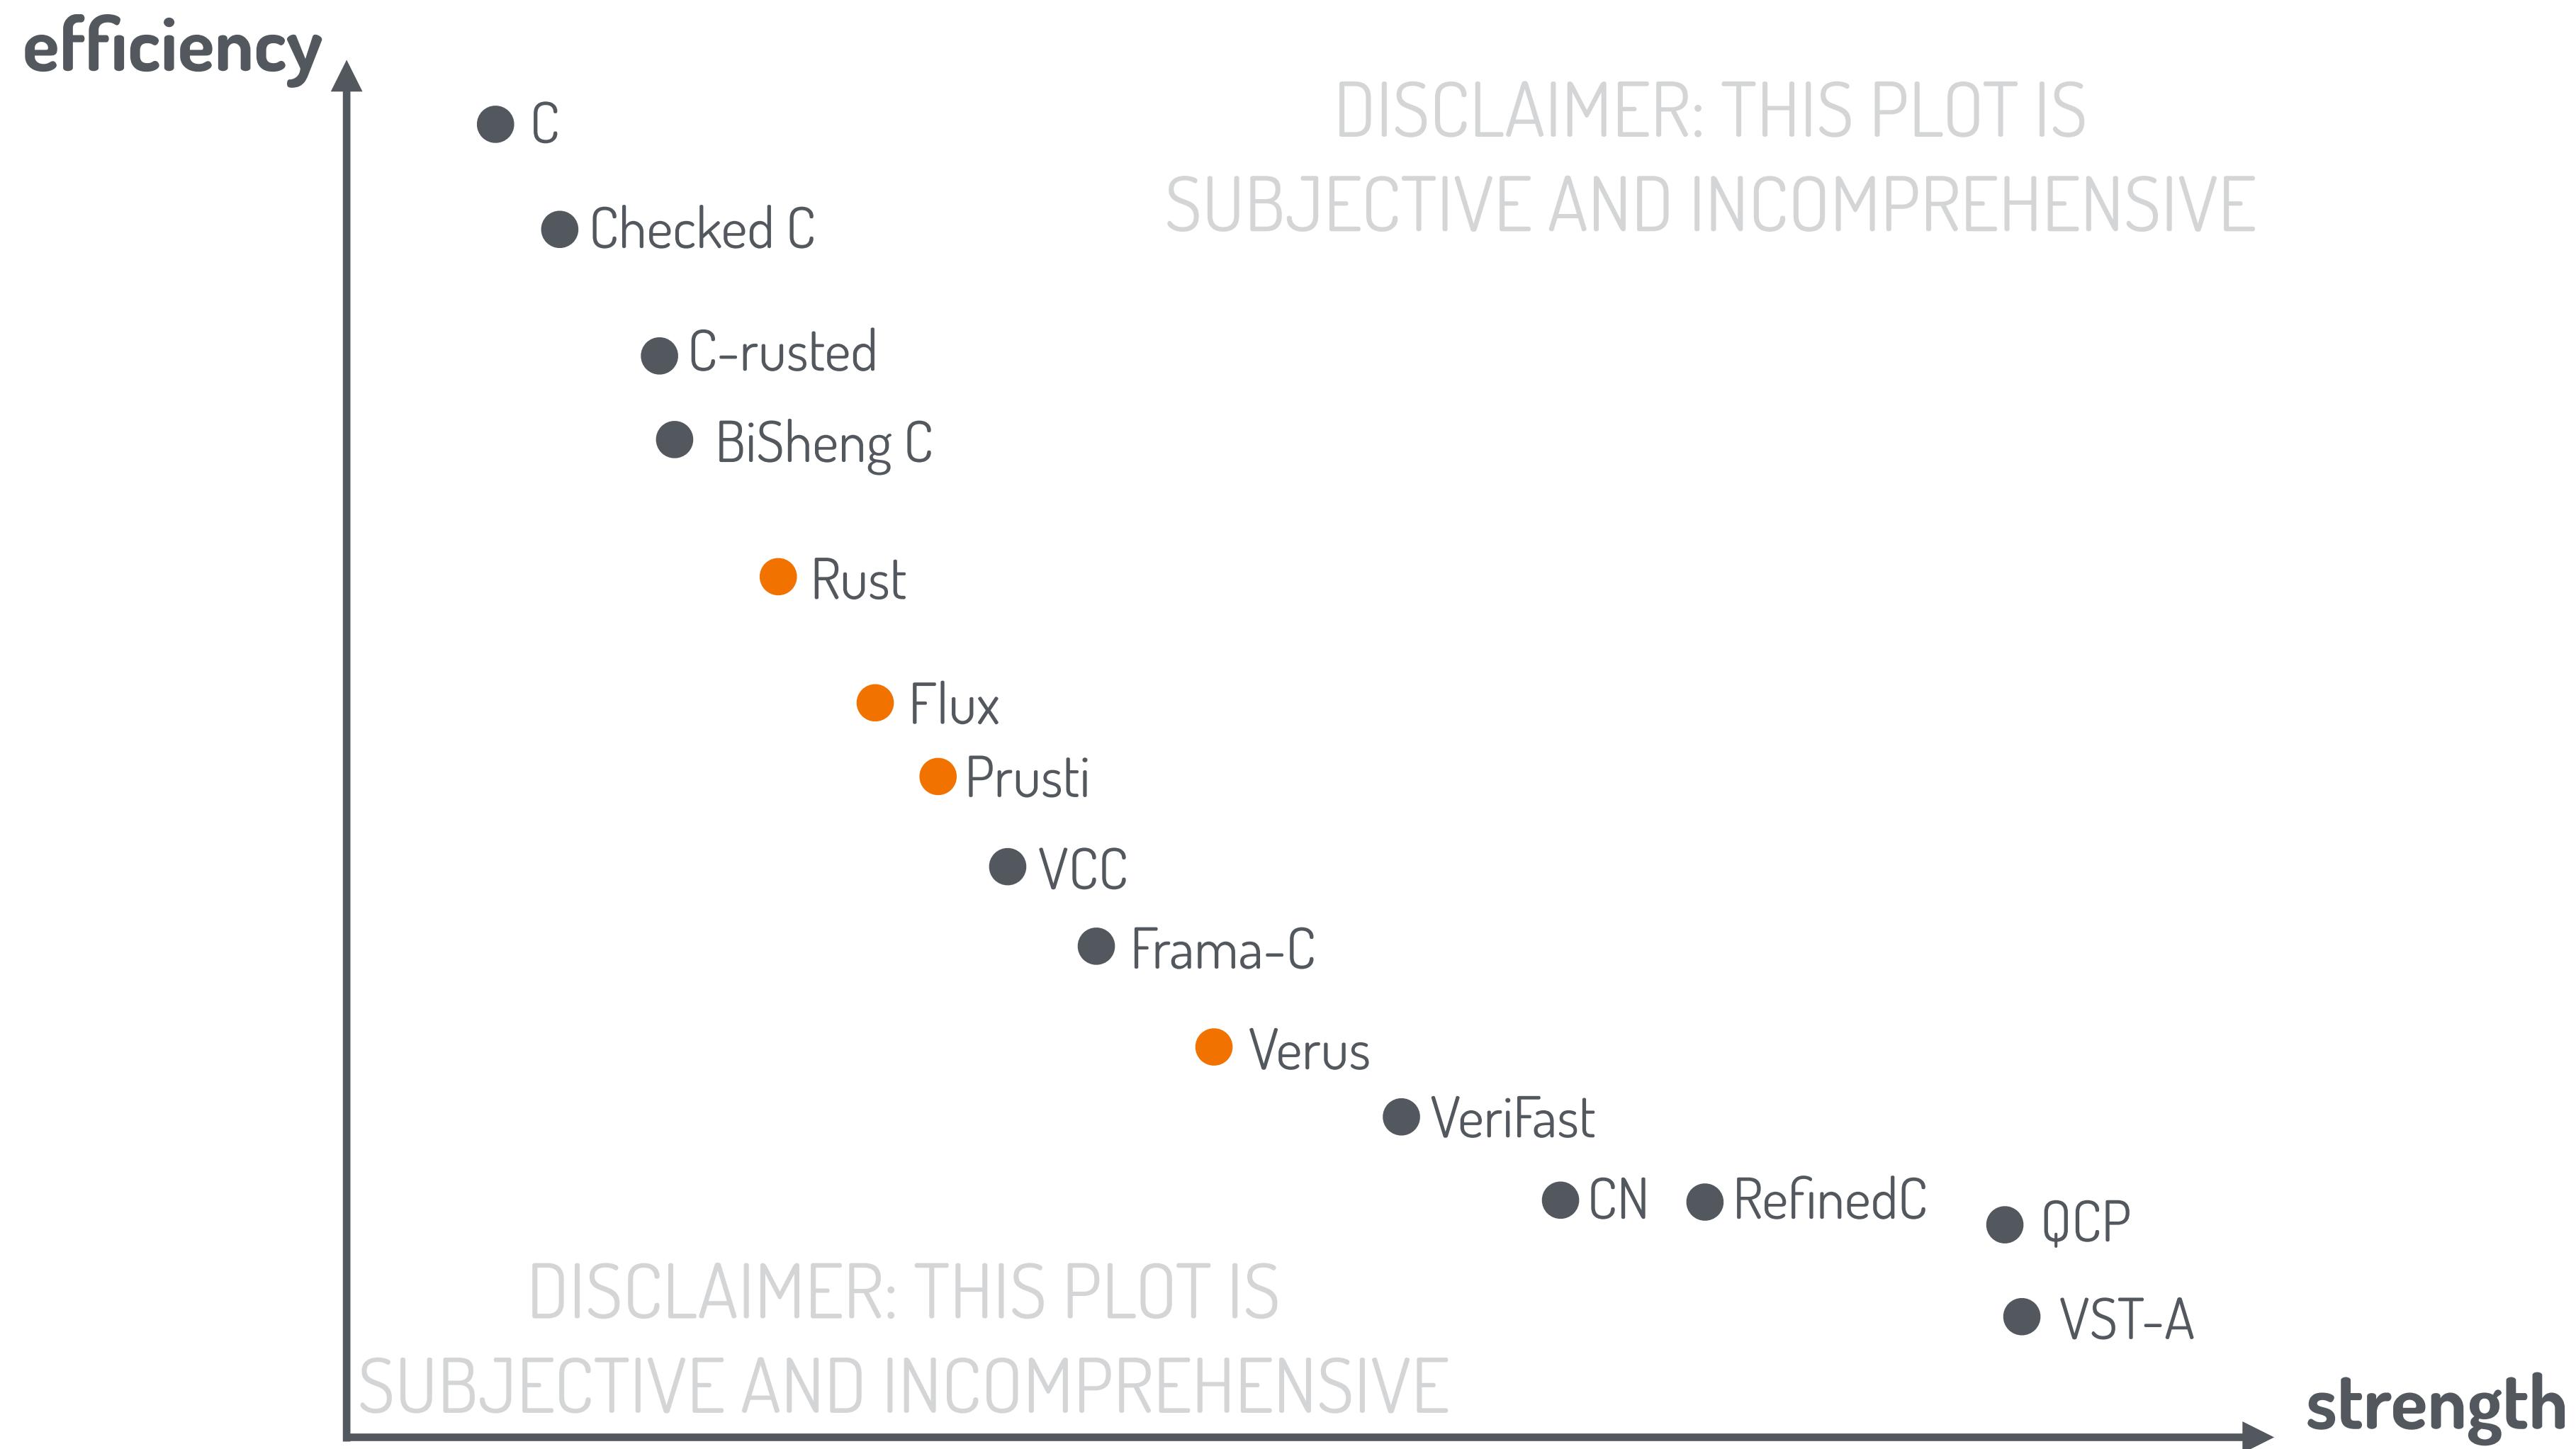
## verification efficiency

- learning curve
- automation, modularity, reusability, evolution
- synergy with existing tooling



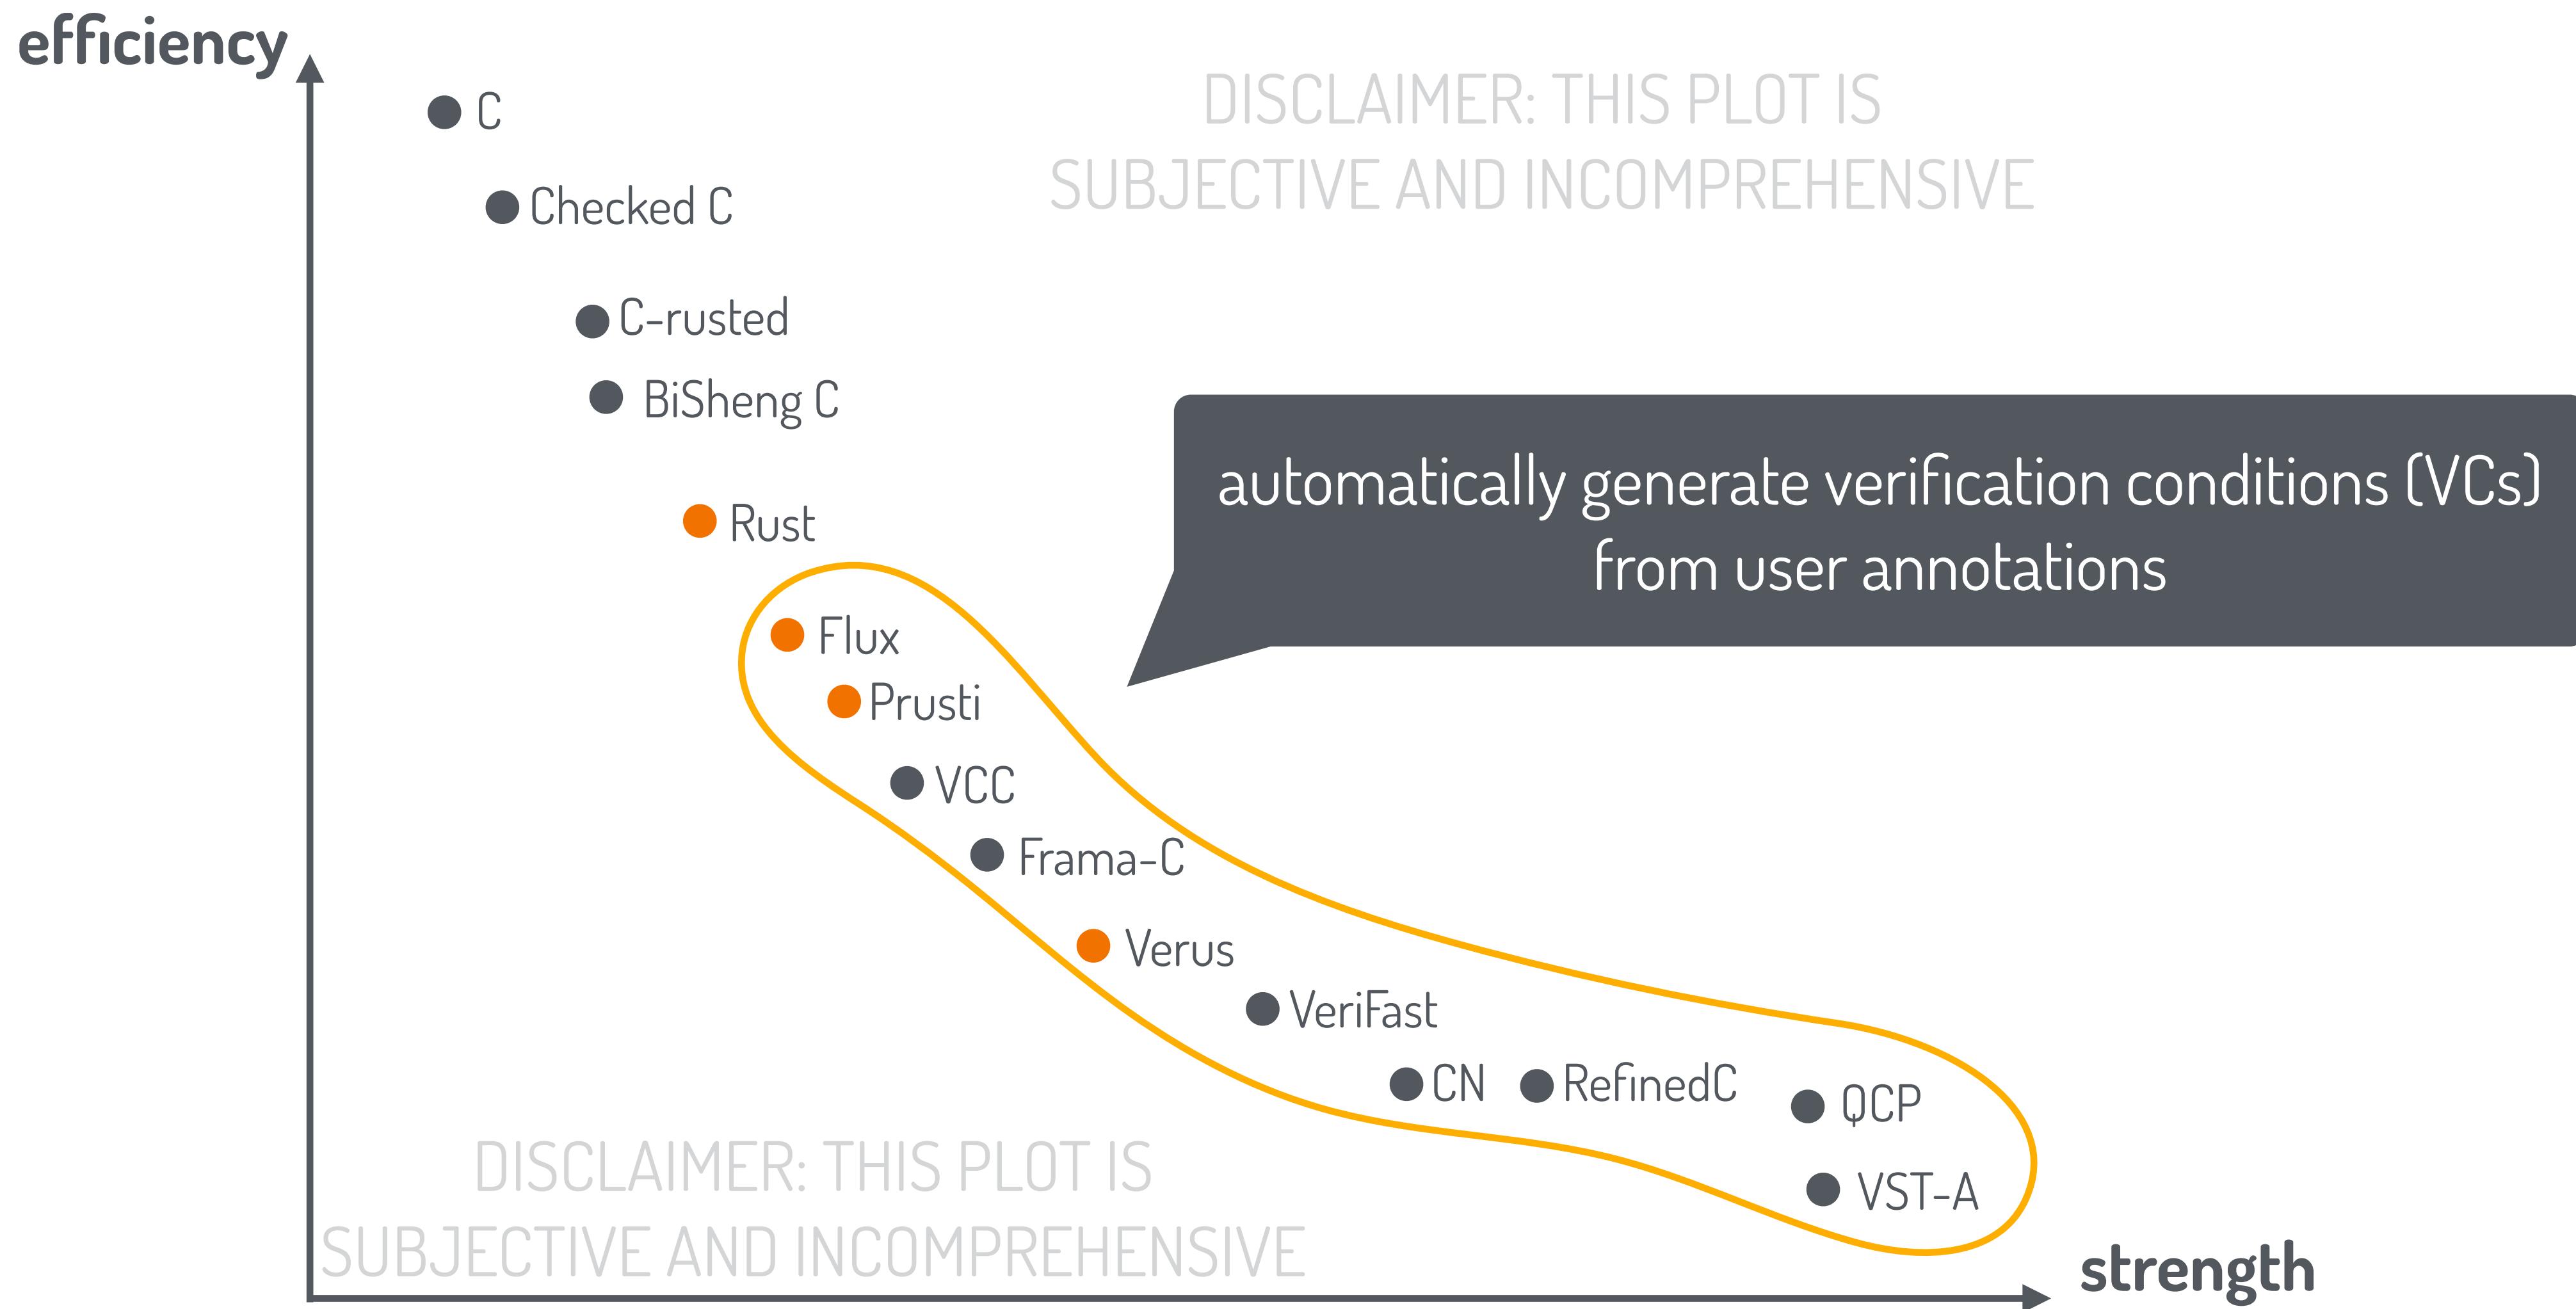
# Obs I: VC Generation is Required

if aiming for high efficiency



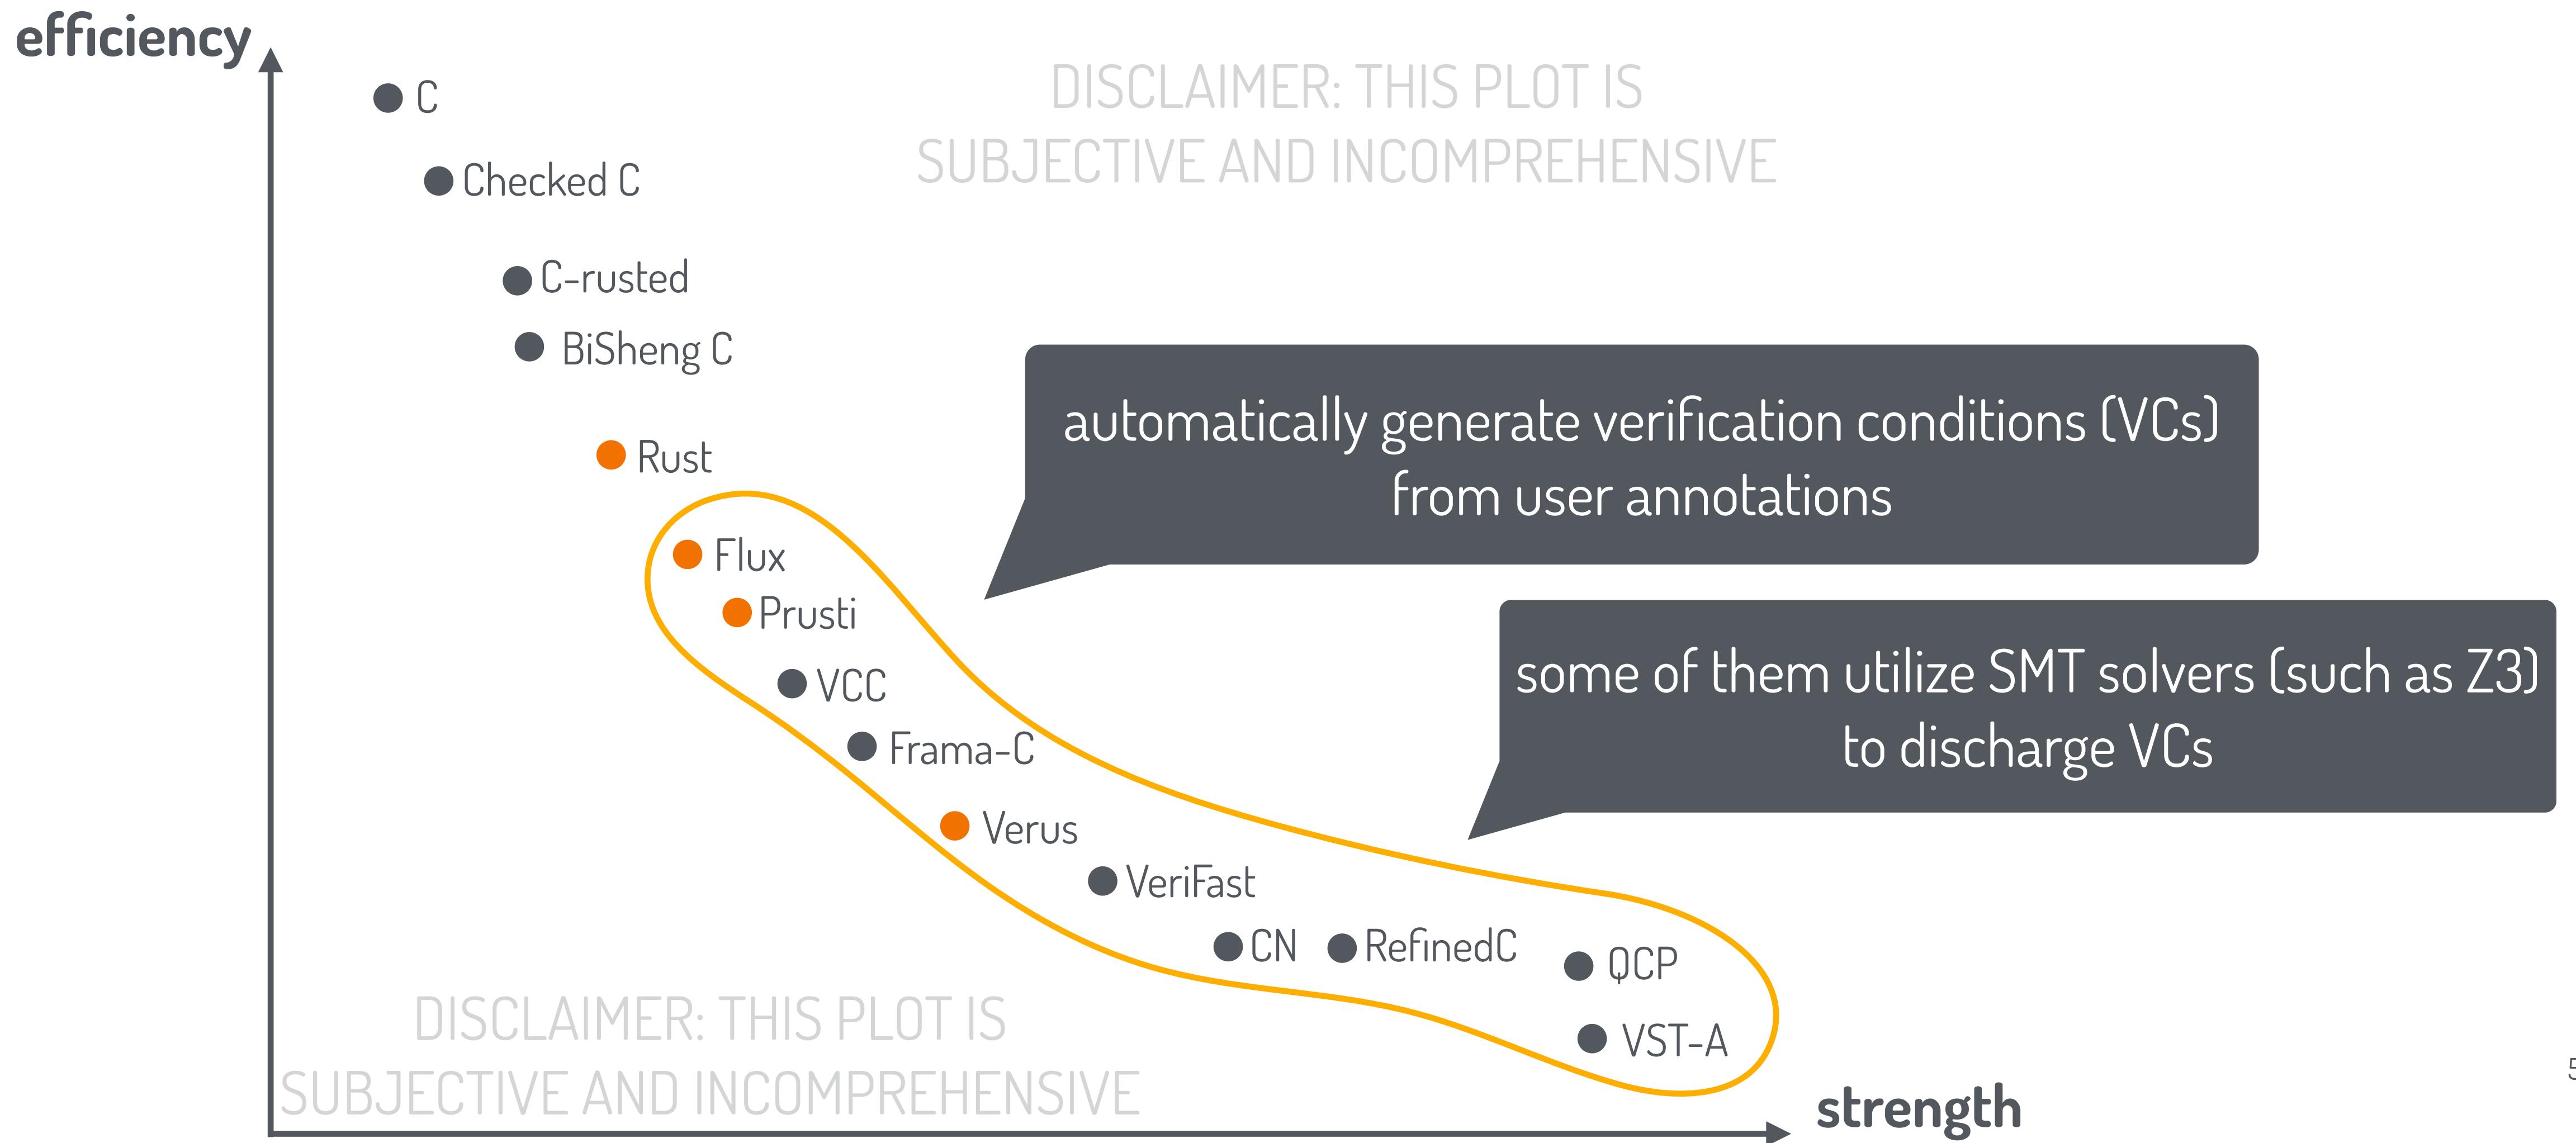
# Obs I: VC Generation is Required

if aiming for high efficiency



# Obs I: VC Generation is Required

if aiming for high efficiency



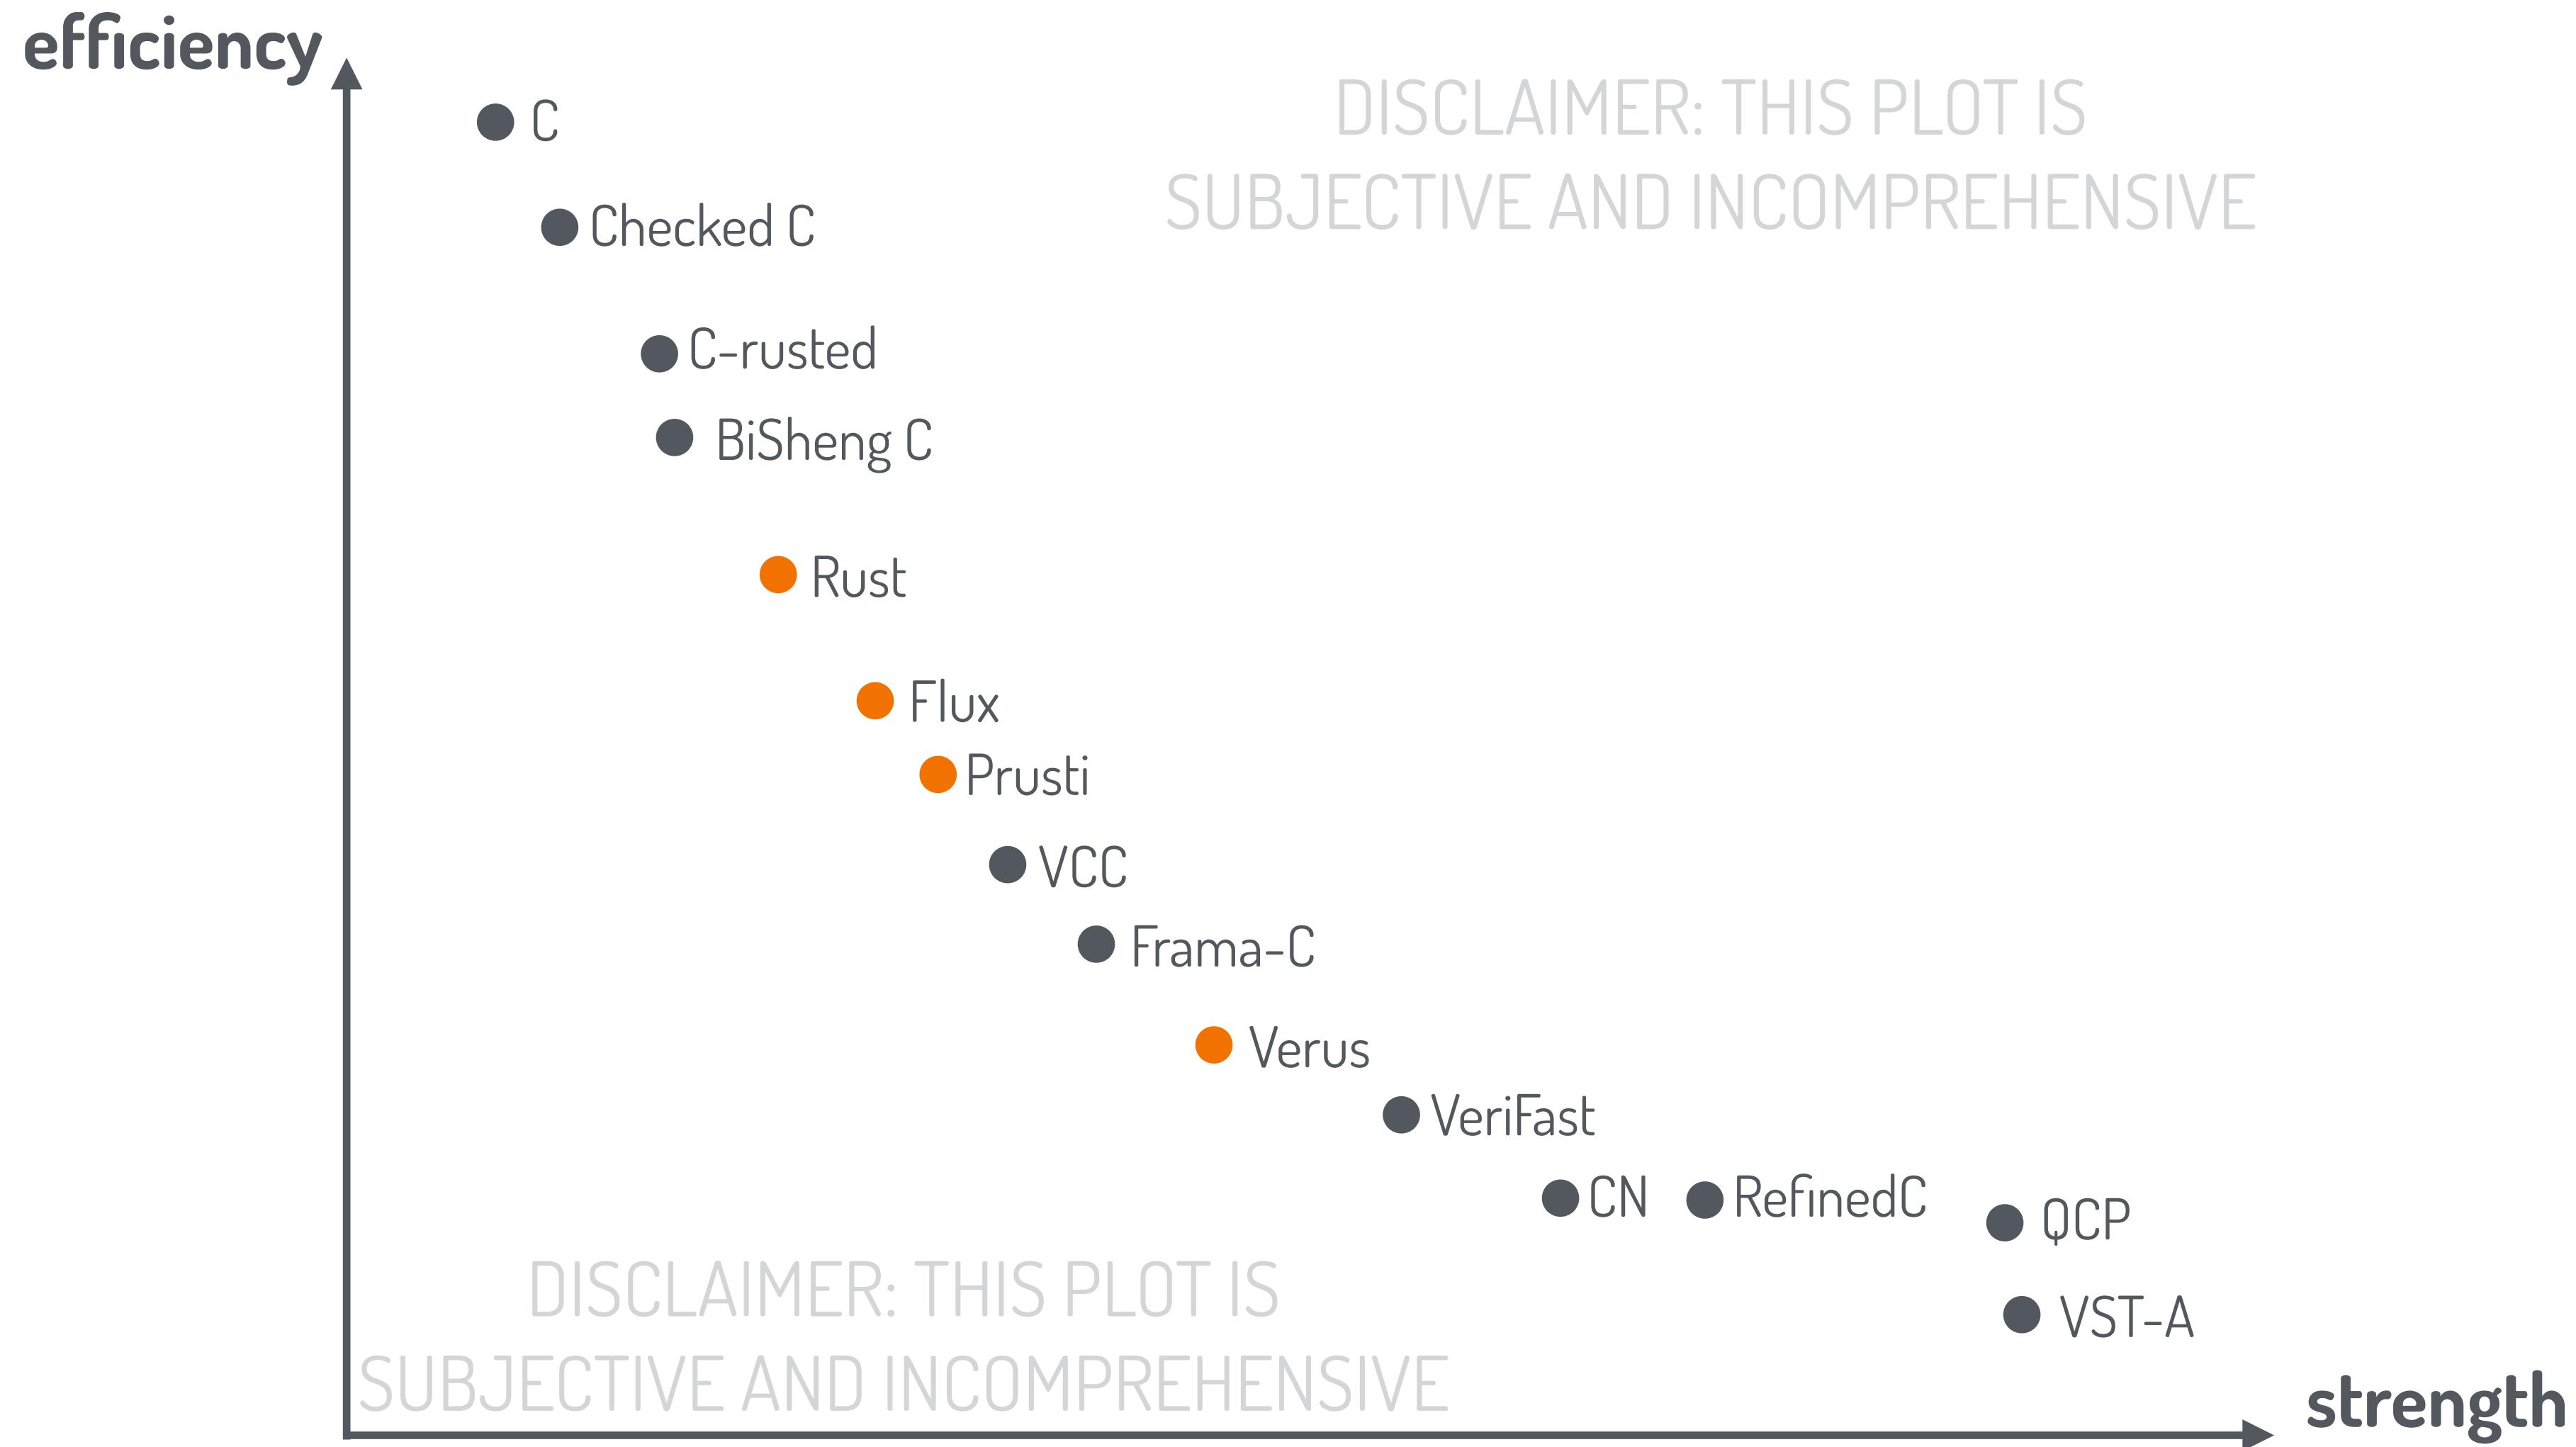


# CStar Demo: VC Generation

nothing new yet

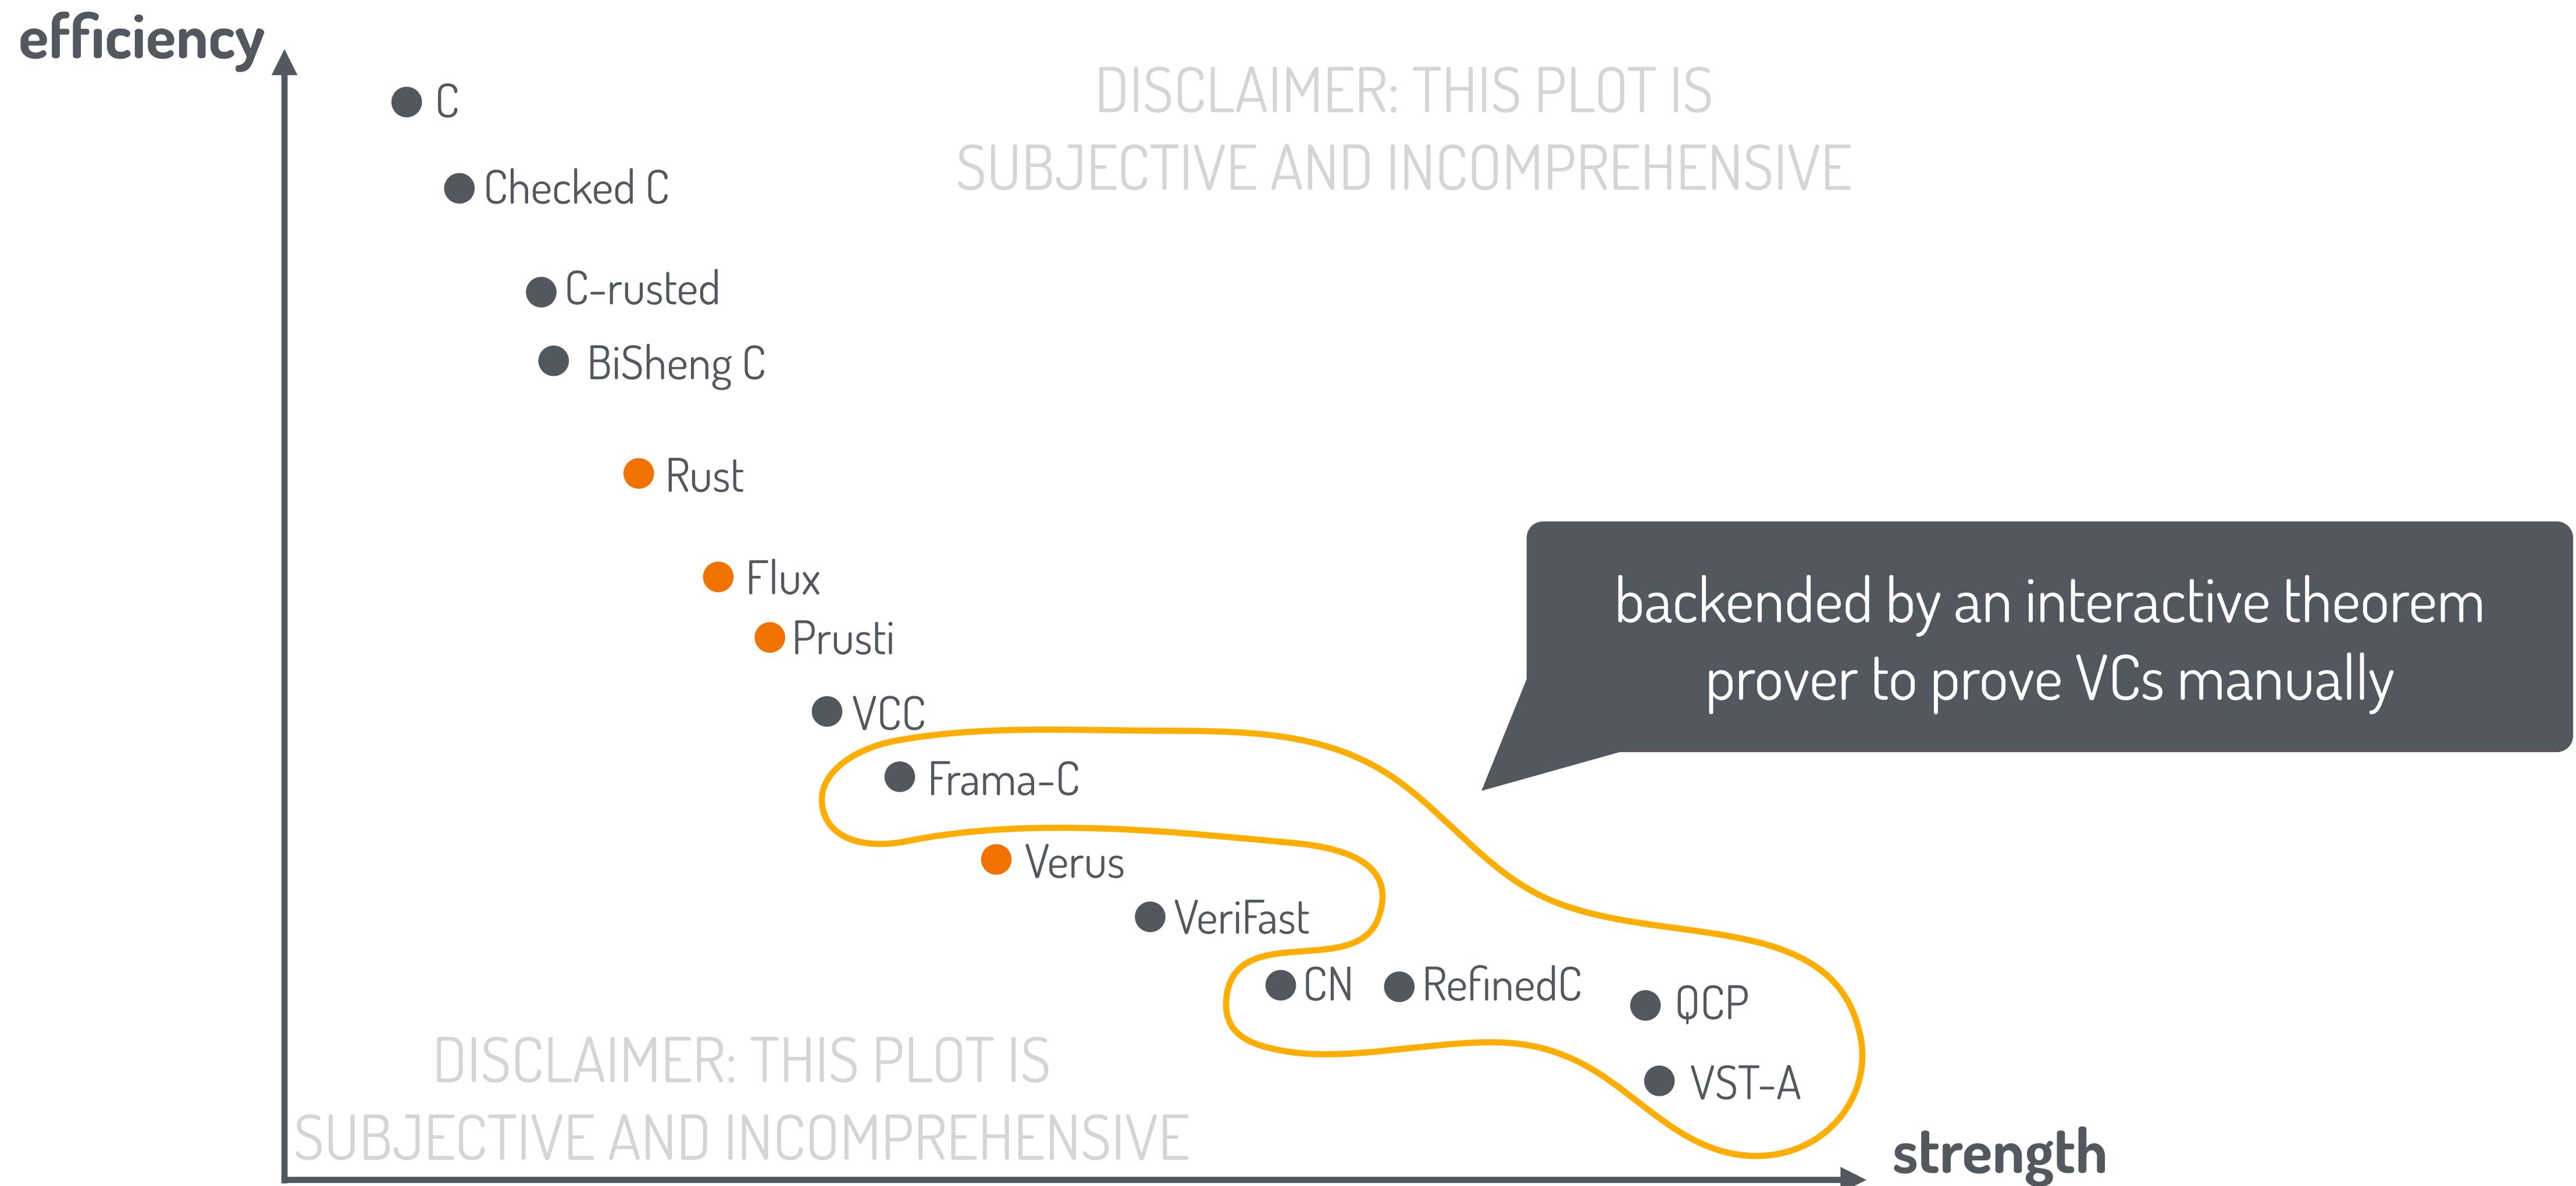
# Obs: Manual Proof is Unavoidable

if aiming at strong strength



# Obs: Manual Proof is Unavoidable

if aiming at strong strength





# Prob I: Proof Encapsulation

theorem provers are nice, but ...

current status in CN (and many others):  
proofs are delegated to Rocq

```
lemma append_nil (datatype seq l1)
  requires true;
  ensures append(l1, Nil {}) == l1;
```

```
lemma append_cons (datatype seq l1, i32 x, datatype seq l2)
  requires true;
  ensures append(l1, Cons {head: x, tail: l2})
    == append(snoc(l1, x), l2);
```

```
struct node *tmp = cur->tail;
cur->tail = last;
last = cur;
cur = tmp;
/*@ unfold rev(L2); */
/*@ apply append_cons (rev (tl(L2)), hd(L2), L1); */
```



# Prob I: Proof Encapsulation

theorem provers are nice, but ...

- obs: libraries constitute the hidden substrate of modern (systems) software

current status in CN (and many others):  
proofs are delegated to Rocq

```
lemma append_nil (datatype seq l1)
  requires true;
  ensures append(l1, Nil {}) == l1;
```

```
lemma append_cons (datatype seq l1, i32 x, datatype seq l2)
  requires true;
  ensures append(l1, Cons {head: x, tail: l2})
    == append(snoc(l1, x), l2);
```

```
struct node *tmp = cur->tail;
cur->tail = last;
last = cur;
cur = tmp;
/*@ unfold rev(L2); */
/*@ apply append_cons (rev (tl(L2)), hd(L2), L1); */
```



# Prob I: Proof Encapsulation

theorem provers are nice, but ...

- obs: libraries constitute the hidden substrate of modern (systems) software
- obs: multilingual libraries, on the other hand, might cause cross-language-boundary issues

current status in CN (and many others):  
proofs are delegated to Rocq

```
lemma append_nil (datatype seq l1)
  requires true;
  ensures append(l1, Nil {}) == l1;

lemma append_cons (datatype seq l1, i32 x, datatype seq l2)
  requires true;
  ensures append(l1, Cons {head: x, tail: l2})
    == append(snoc(l1, x), l2);
```

```
struct node *tmp = cur->tail;
cur->tail = last;
last = cur;
cur = tmp;
/*@ unfold rev(L2); */
/*@ apply append_cons (rev (tl(L2)), hd(L2), L1); */
```



# Prob I: Proof Encapsulation

theorem provers are nice, but ...

- obs: libraries constitute the hidden substrate of modern (systems) software
- obs: multilingual libraries, on the other hand, might cause cross-language-boundary issues

principle: integrate **proof-specification-implementation** together into one language

current status in CN (and many others):  
proofs are delegated to Rocq

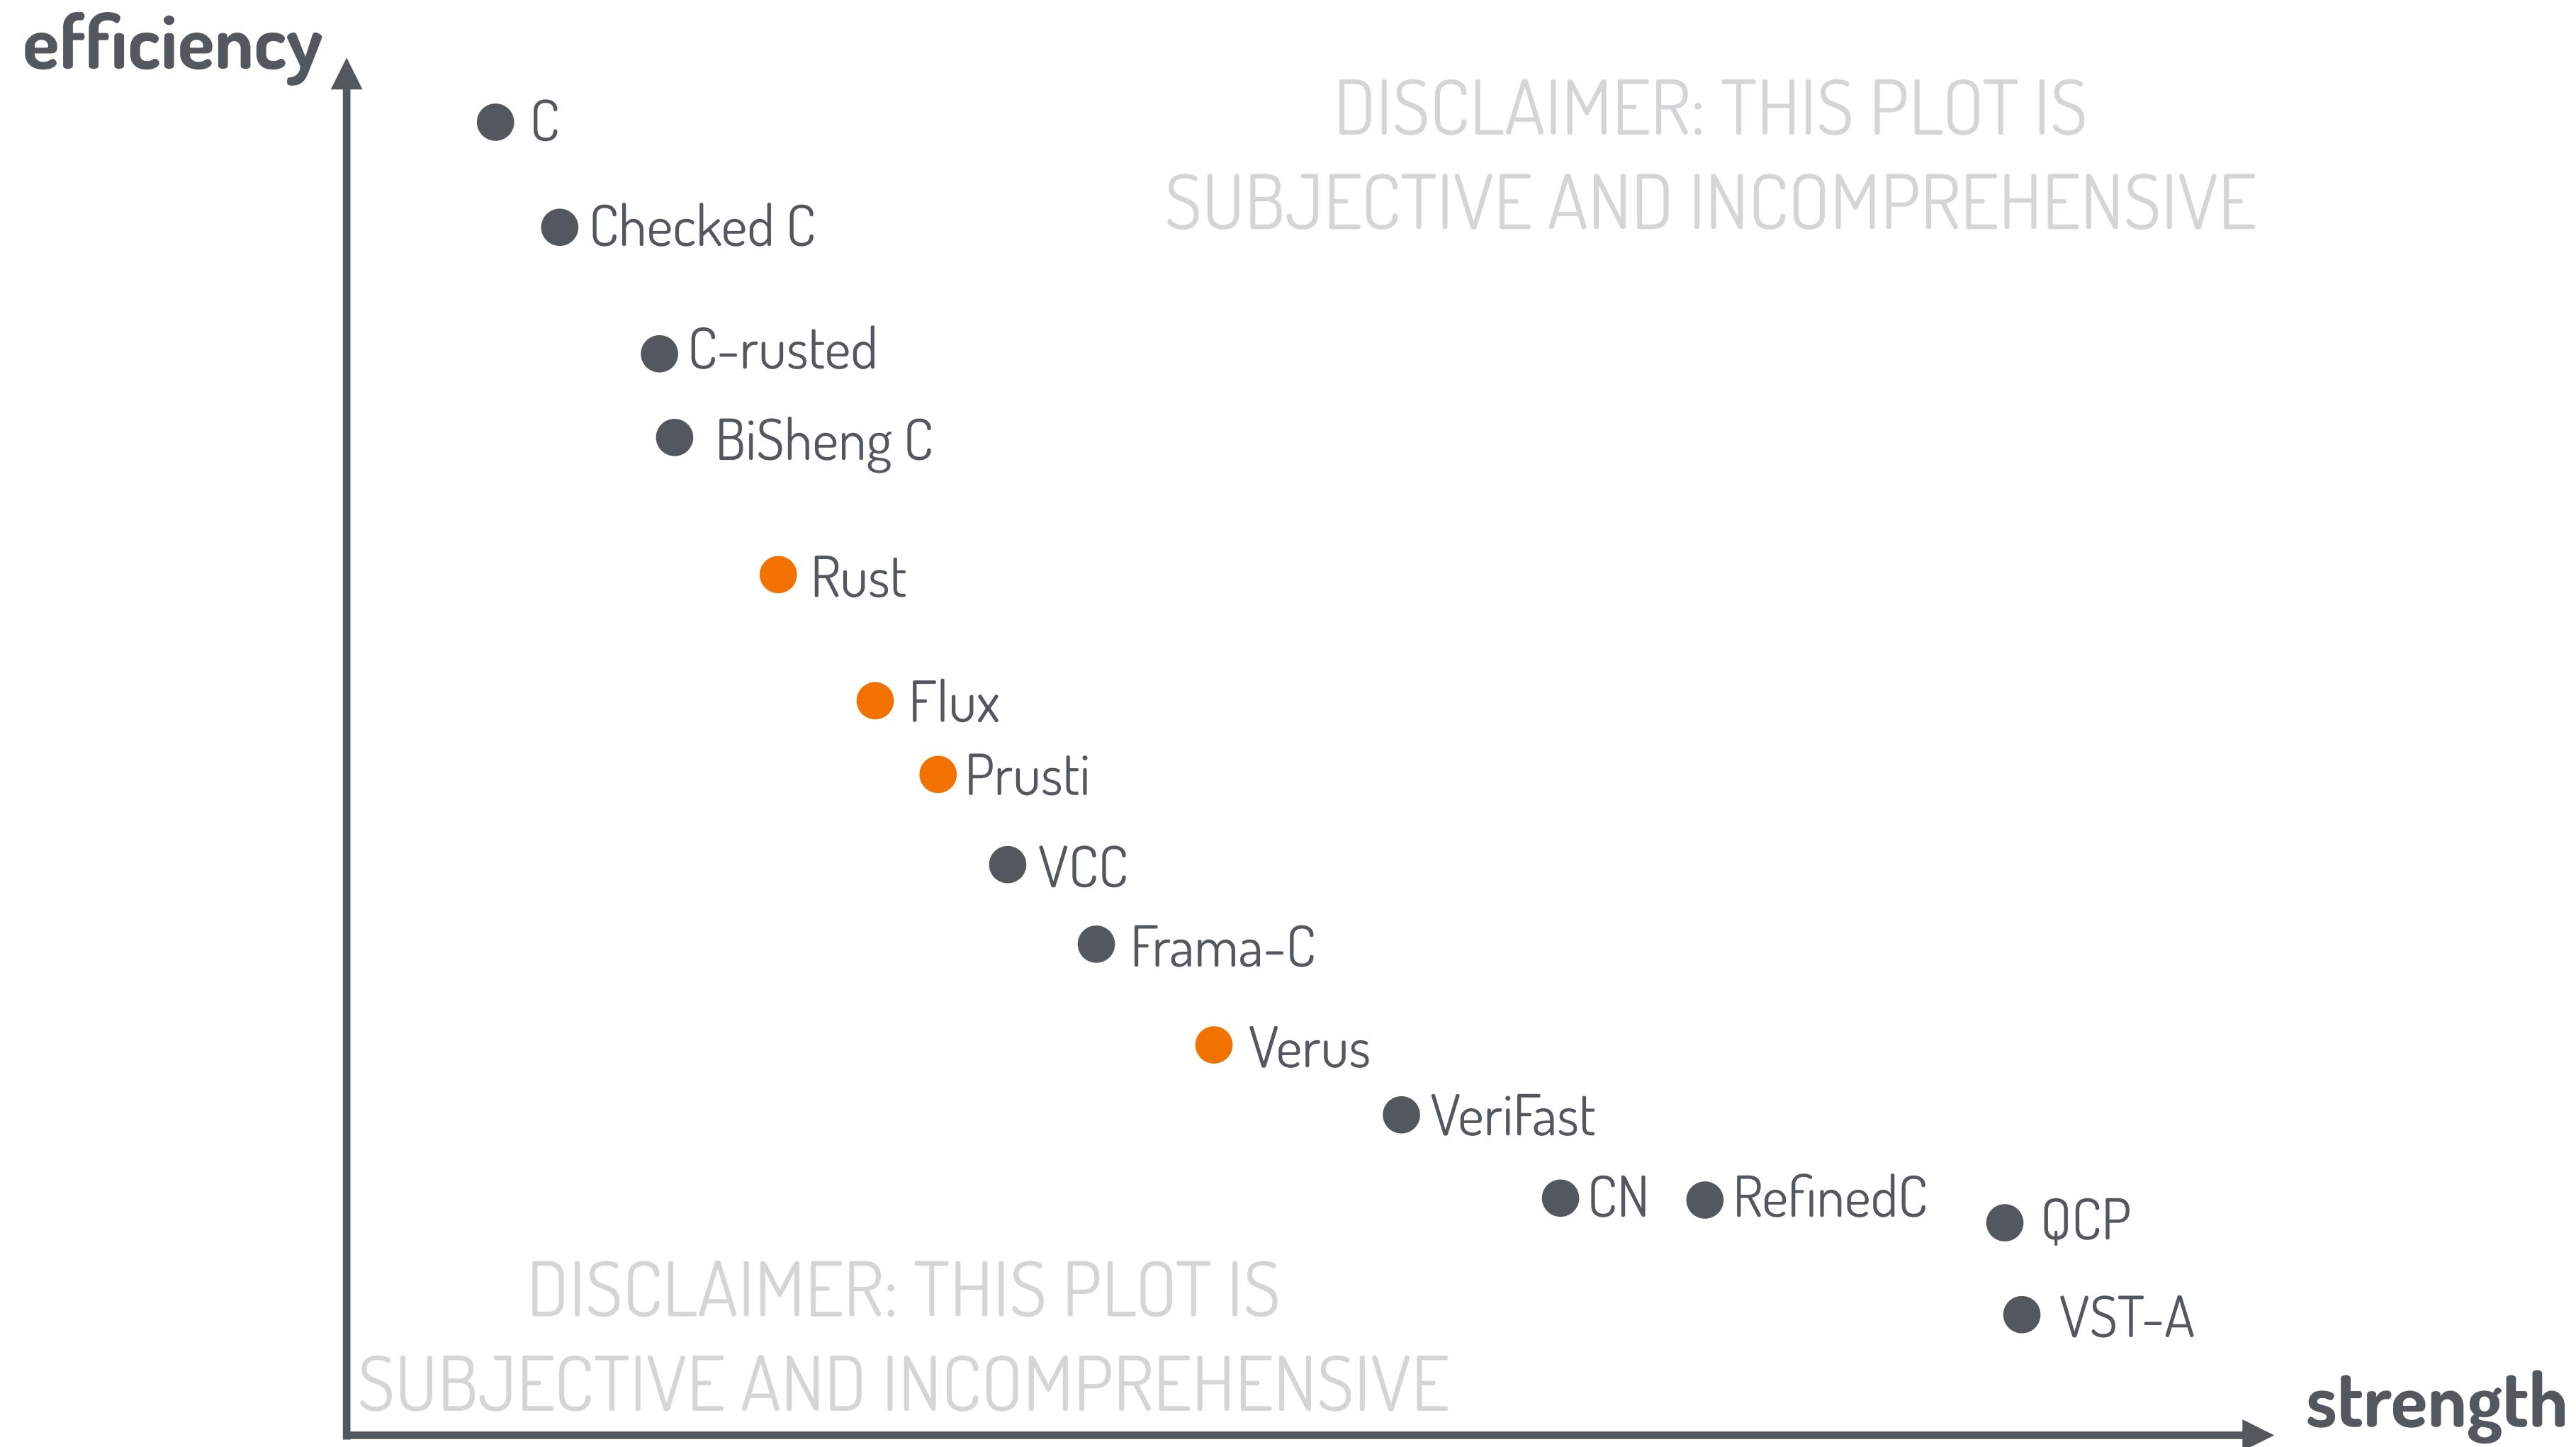
```
lemma append_nil (datatype seq l1)
  requires true;
  ensures append(l1, Nil {}) == l1;

lemma append_cons (datatype seq l1, i32 x, datatype seq l2)
  requires true;
  ensures append(l1, Cons {head: x, tail: l2})
    == append(snoc(l1, x), l2);

struct node *tmp = cur->tail;
cur->tail = last;
last = cur;
cur = tmp;
/*@ unfold rev(L2); */
/*@ apply append_cons (rev (tl(L2)), hd(L2), L1); */
```

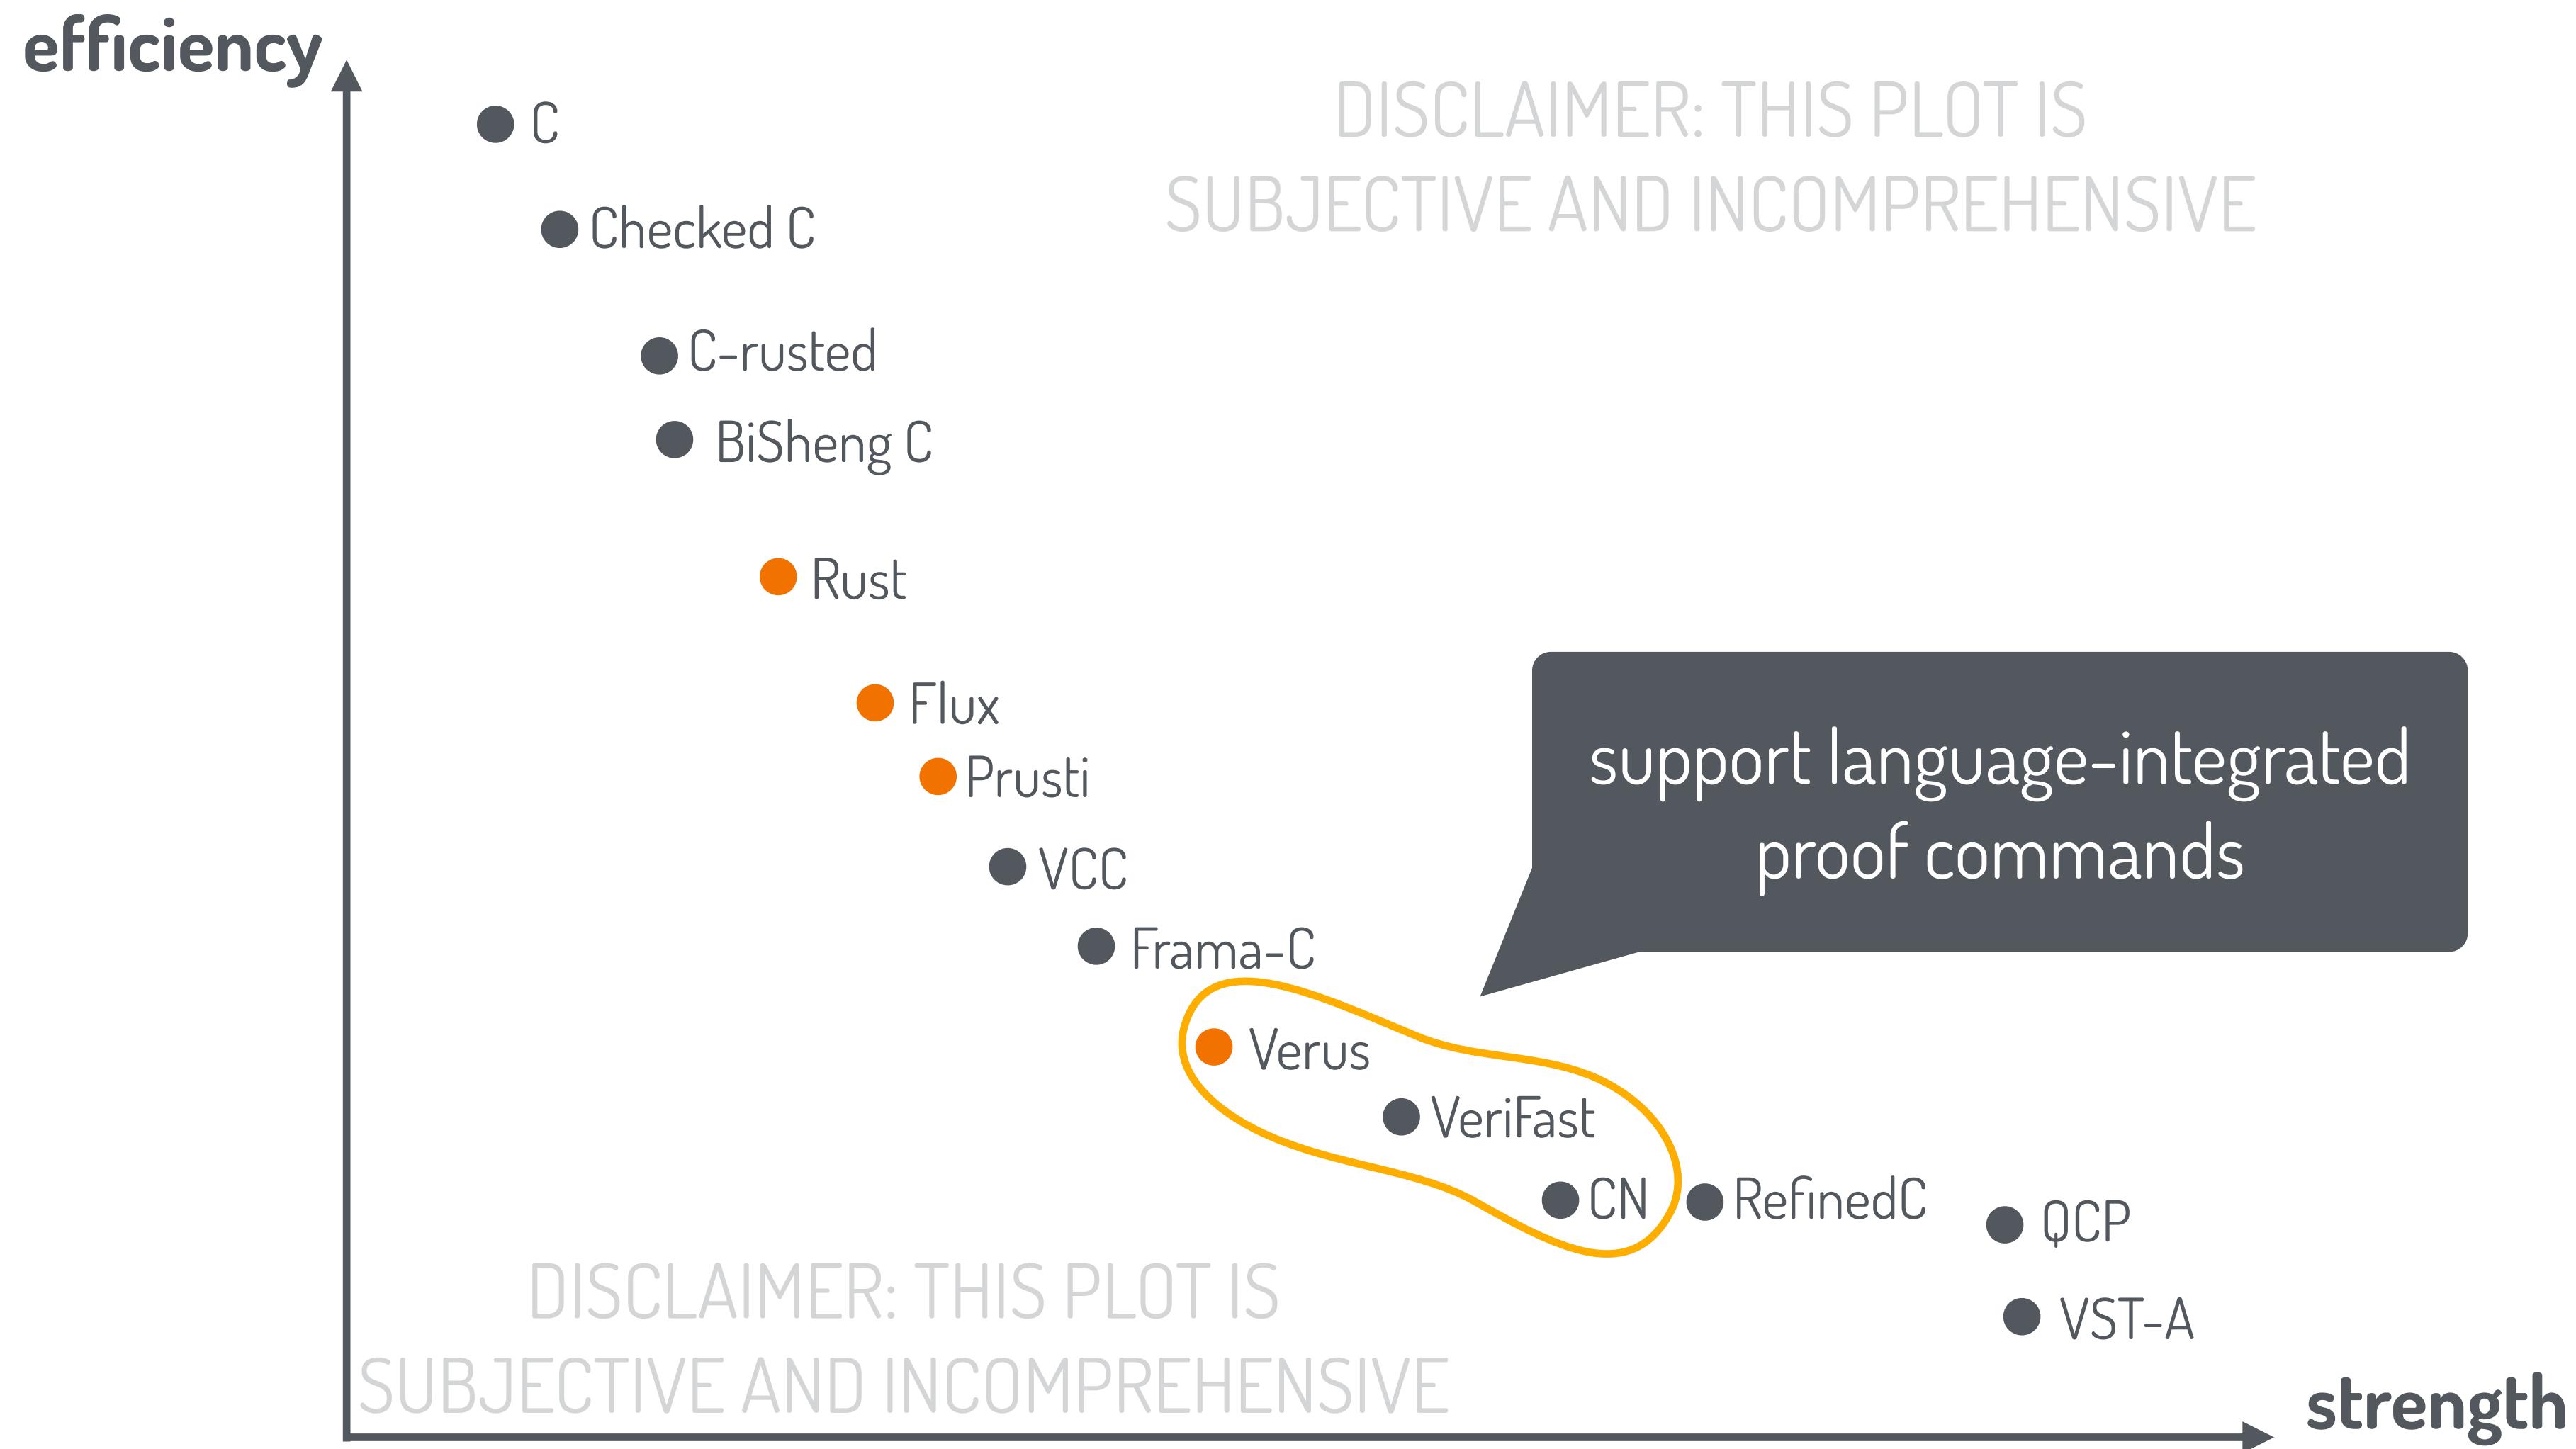
# Obs: Manual Proof is Unavoidable

if aiming at strong strength



# Obs: Manual Proof is Unavoidable

if aiming at strong strength





```
struct node *reverse(struct node *xs)
/*@ requires take L = IntList(xs);
ensures  take L_ = IntList(return);
L_ == rev(L);

@*/
{

    struct node *last = 0;
    struct node *cur = xs;
/*@ apply append_nil(rev(L)); @*/
while(1)
/*@ inv take L1 = IntList(last);
take L2 = IntList(cur);
append(rev(L2), L1) == rev(L);

@*/
{
    if (cur == 0) {
/*@ unfold rev(Nil {}); @*/
/*@ unfold append(Nil {}, L1); @*/

        return last;
    }
    struct node *tmp = cur->tail;
    cur->tail = last;
    last = cur;
    cur = tmp;
/*@ unfold rev(L2); @*/
/*@ apply append_cons (rev (tl(L2)), hd(L2), L1); @*/
}
}
```

example program with proof commands in CN

```

struct node *reverse(struct node *xs)
/*@ requires take L = IntList(xs);
ensures take L_ = IntList(return);
L_ == rev(L);
*/
{
  struct node *last = 0;
  struct node *cur = xs;
/*@ apply append_nil(rev(L)); */
while(1)
/*@ inv take L1 = IntList(last);
take L2 = IntList(cur);
append(rev(L2), L1) == rev(L);
*/
{
  if (cur == 0) {
    /*@ unfold rev(Nil {}); */
    /*@ unfold append(Nil {}, L1); */

    return last;
  }
  struct node *tmp = cur->tail;
  cur->tail = last;
  last = cur;
  cur = tmp;
  /*@ unfold rev(L2); */
  /*@ apply append_cons (rev (tl(L2)), hd(L2), L1); */
}
}
  
```

```

public boolean remove(Object o)
/*: requires " init "
  modifies content, csize
  ensures
  " ( result → ( ∃ i. (i,o) ∈ old content ∧
    (¬ ∃ j. j < i ∧ (j,o) ∈ old content) ∧
    ( ∀ j e. 0 ≤ j ∧ j < i → (j,e) ∈ content = (j,e) ∈ old content) ∧
    ( i ≤ j ∧ j < csize → (j,e) ∈ content = (j+1,e) ∈ old content) ) ∧
  (¬ result → ( content = old content ∧ ¬ ∃ i. (i,o) ∈ old content ))" */
{
  int index = 0;
  while /*: inv " ( ∀ j. 0 ≤ j ∧ j < index → o ≠ elements[j] ) ∧
  0 ≤ index ∧ size = old size" */
    ( index < size ) {
    if (elements[index] == o) {
      shift(index);
      /*: note ObjectRemoved:
      " ∀ j e. ( 0 ≤ j ∧ j < index → (j,e) ∈ content = (j,e) ∈ old content ) ∧
      ( index ≤ j ∧ j < csize → (j,e) ∈ content = (j+1,e) ∈ old content )"
      from shift_Postcondition , LoopInv , LoopCondition
      content_def , csize_def ;
      witness index for
      " ∃ i. (i,o) ∈ old content ∧ (¬ ∃ j. j < i ∧ (j,o) ∈ old content) ∧
      ( ∀ j e. ( 0 ≤ j ∧ j < i → (j,e) ∈ content = (j,e) ∈ old content ) ∧
      ( i ≤ j ∧ j < csize → (j,e) ∈ content = (j+1,e) ∈ old content ) )" */
      return true;
    }
    index = index + 1;
  }
  return false;
}
  
```

example program with proof commands in CN

example program with proof commands in Jahob

# Prob II: Programmable Proof

proof commands are nice, but ...

$p ::=$	$p_1 ; p_2$
	assert $l : F$ from $\vec{h}$
	note $l : F$ from $\vec{h}$
	localize in $(p ; \text{note } l : F)$
	mp $l : (F \rightarrow G)$
	assuming $l_F : F$ in $(p ; \text{note } l_G : G)$
	cases $\vec{F}$ for $l : G$
	showedCase $i$ of $l : F_1 \vee \dots \vee F_n$
	byContradiction $l : F$ in $p$
	contradiction $l : F$
	instantiate $l : \forall \vec{x}. F$ with $\vec{t}$
	witness $\vec{t}$ for $l : \exists \vec{x}. F$
	pickWitness $\vec{x}$ for $l_F : F$ in $(p ; \text{note } l_G : G)$
	pickAny $\vec{x}$ in $(p ; \text{note } l : F)$
	induct $l : F$ over $n$ in $p$

# Prob II: Programmable Proof

proof commands are nice, but ...

- obs: manual proofs are inevitable for complex reasoning tasks

$p ::=$	$p_1 ; p_2$ <b>assert</b> $l : F$ <b>from</b> $\vec{h}$ <b>note</b> $l : F$ <b>from</b> $\vec{h}$ <b>localize in</b> $(p ; \text{note } l : F)$ <b>mp</b> $l : (F \rightarrow G)$ <b>assuming</b> $l_F : F$ <b>in</b> $(p ; \text{note } l_G : G)$ <b>cases</b> $\vec{F}$ <b>for</b> $l : G$ <b>showedCase</b> $i$ <b>of</b> $l : F_1 \vee \dots \vee F_n$ <b>byContradiction</b> $l : F$ <b>in</b> $p$ <b>contradiction</b> $l : F$ <b>instantiate</b> $l : \forall \vec{x}. F$ <b>with</b> $\vec{t}$ <b>witness</b> $\vec{t}$ <b>for</b> $l : \exists \vec{x}. F$ <b>pickWitness</b> $\vec{x}$ <b>for</b> $l_F : F$ <b>in</b> $(p ; \text{note } l_G : G)$ <b>pickAny</b> $\vec{x}$ <b>in</b> $(p ; \text{note } l : F)$ <b>induct</b> $l : F$ <b>over</b> $n$ <b>in</b> $p$
---------	--

# Prob II: Programmable Proof

proof commands are nice, but ...

- obs: manual proofs are inevitable for complex reasoning tasks
- obs: using a fixed set of proof commands feels neither comprehensive nor extensible

$p ::=$	$p_1 ; p_2$ <b>assert</b> $l : F$ <b>from</b> $\vec{h}$ <b>note</b> $l : F$ <b>from</b> $\vec{h}$ <b>localize in</b> $(p ; \text{note } l : F)$ <b>mp</b> $l : (F \rightarrow G)$ <b>assuming</b> $l_F : F$ <b>in</b> $(p ; \text{note } l_G : G)$ <b>cases</b> $\vec{F}$ <b>for</b> $l : G$ <b>showedCase</b> $i$ <b>of</b> $l : F_1 \vee \dots \vee F_n$ <b>byContradiction</b> $l : F$ <b>in</b> $p$ <b>contradiction</b> $l : F$ <b>instantiate</b> $l : \forall \vec{x}. F$ <b>with</b> $\vec{t}$ <b>witness</b> $\vec{t}$ <b>for</b> $l : \exists \vec{x}. F$ <b>pickWitness</b> $\vec{x}$ <b>for</b> $l_F : F$ <b>in</b> $(p ; \text{note } l_G : G)$ <b>pickAny</b> $\vec{x}$ <b>in</b> $(p ; \text{note } l : F)$ <b>induct</b> $l : F$ <b>over</b> $n$ <b>in</b> $p$
---------	--

# Prob II: Programmable Proof

proof commands are nice, but ...

- obs: manual proofs are inevitable for complex reasoning tasks
- obs: using a fixed set of proof commands feels neither comprehensive nor extensible

principle: support using the full language to **programmatically** construct proof

$p$	$::=$	$p_1 ; p_2$
		<b>assert</b> $l : F$ from $\vec{h}$
		<b>note</b> $l : F$ from $\vec{h}$
		<b>localize</b> in $(p ; \text{note } l : F)$
		<b>mp</b> $l : (F \rightarrow G)$
		<b>assuming</b> $l_F : F$ in $(p ; \text{note } l_G : G)$
		<b>cases</b> $\vec{F}$ for $l : G$
		<b>showedCase</b> $i$ of $l : F_1 \vee \dots \vee F_n$
		<b>byContradiction</b> $l : F$ in $p$
		<b>contradiction</b> $l : F$
		<b>instantiate</b> $l : \forall \vec{x}. F$ with $\vec{t}$
		<b>witness</b> $\vec{t}$ for $l : \exists \vec{x}. F$
		<b>pickWitness</b> $\vec{x}$ for $l_F : F$ in $(p ; \text{note } l_G : G)$
		<b>pickAny</b> $\vec{x}$ in $(p ; \text{note } l : F)$
		<b>induct</b> $l : F$ over $n$ in $p$

# The LCF Solution

logic for computable functions





# The LCF Solution

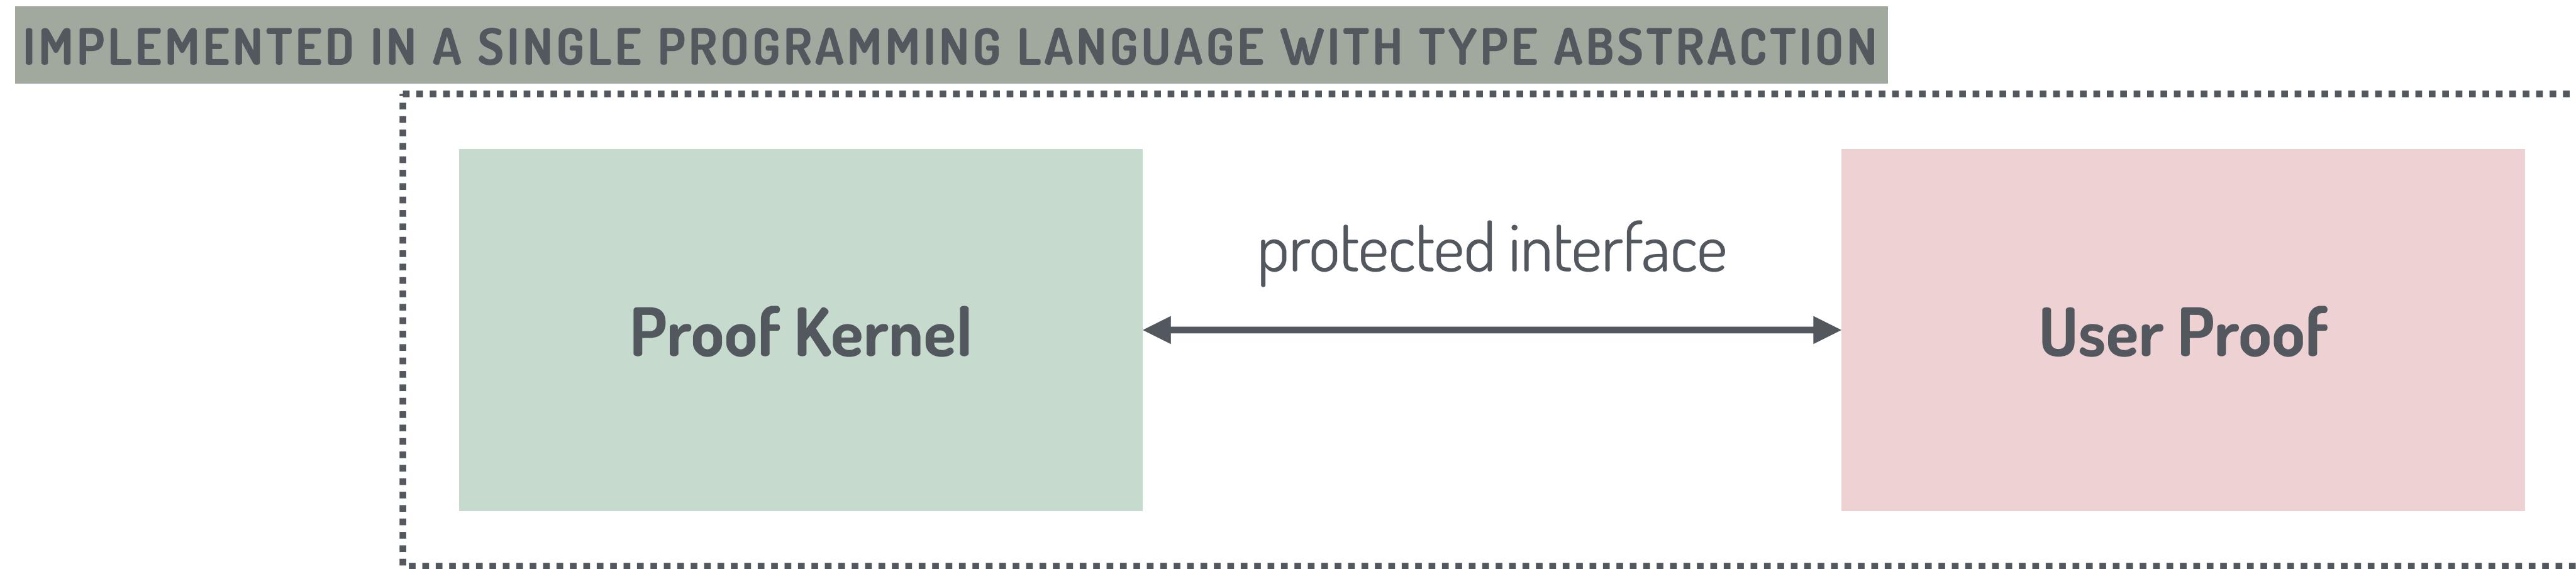
logic for computable functions

- the **LCF architecture** for theorem proving by Robin Milner in the 1970s

# The LCF Solution

logic for computable functions

- the **LCF architecture** for theorem proving by Robin Milner in the 1970s
- programming formal proof is to **compute theorems** using a **protected interface**



# The LCF Solution

## logic for computable functions

- the **LCF architecture** for theorem proving by Robin Milner in the 1970s
- programming formal proof is to **compute theorems** using a **protected interface**
- **correct-by-construction** even when extended with arbitrary user proof code

IMPLEMENTED IN A SINGLE PROGRAMMING LANGUAGE WITH TYPE ABSTRACTION





# CStar Demo: Proof Programming

a practice of LCF in C





it is now already possible to prove VCs in C



it is now already possible to prove VCs in C  
but we can do better by introducing  
program-proof states



# CStar Demo: Program-Proof State

essentially, a **symbolic** representation of the program state





a VC is a prop "pre |-- post" at a program point



both are program-proof states

a VC is a prop "pre |-- post" at a program point



both are program-proof states

a VC is a prop "pre |-- post" at a program point

proving a VC is **coupled** with programming in CStar



both are program-proof states

a VC is a prop "pre |-- post" at a program point

proving a VC is **coupled** with programming in CStar  
you write a proof to transform from **pre** to **post**



# CStar Demo: Program Proof

unify program and proof

# The Architecture of CStar

how to ensure correct-by-construction





# The Architecture of CStar

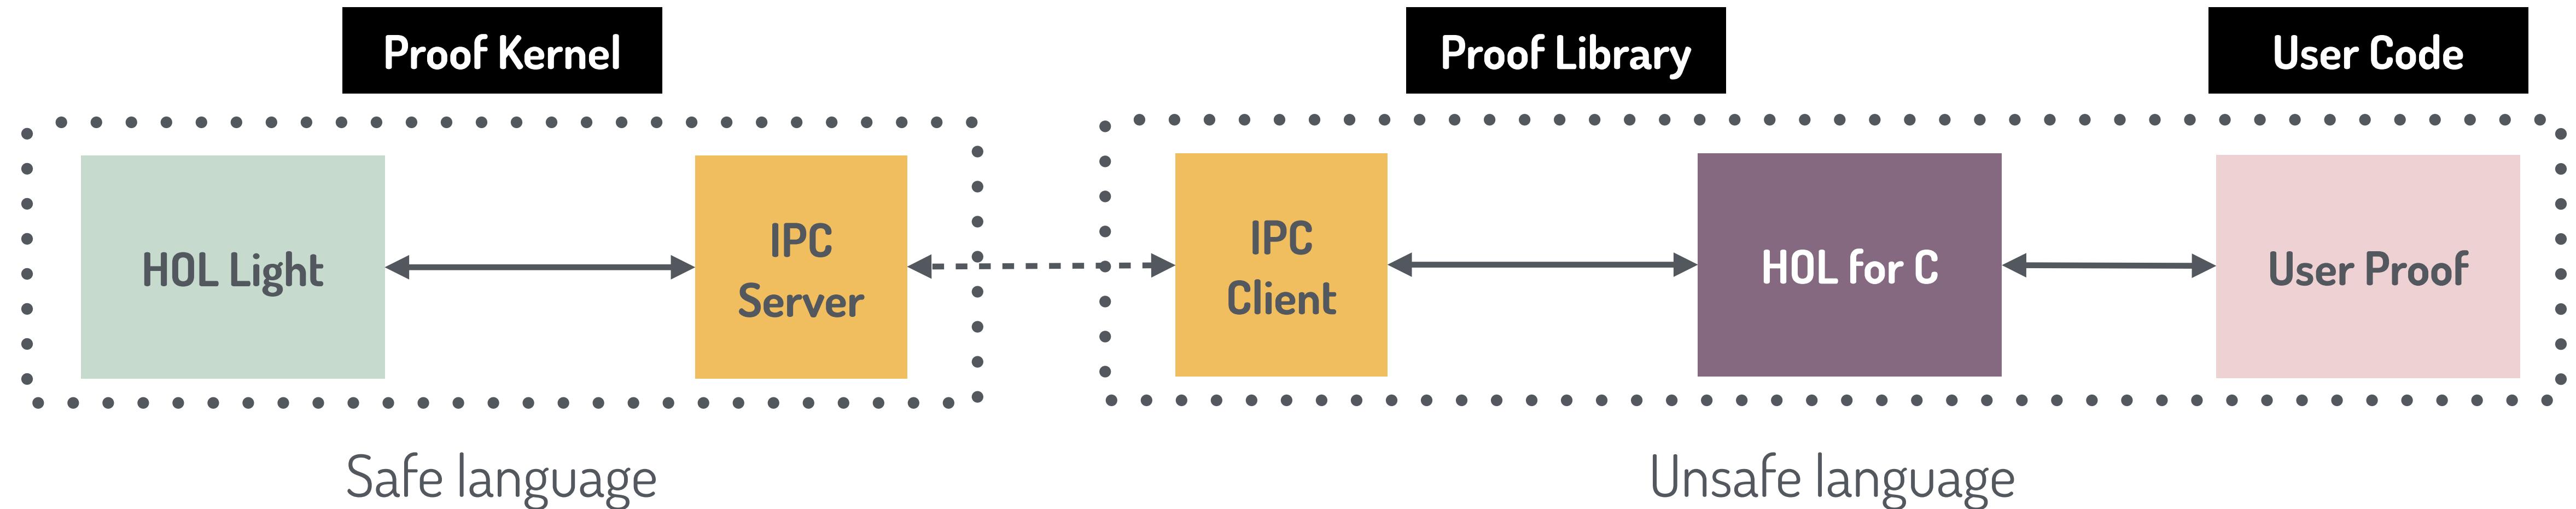
## how to ensure correct-by-construction

- the original LCF design heavily relies on **type abstraction** provided by the ML language

# The Architecture of CStar

## how to ensure correct-by-construction

- the original LCF design heavily relies on **type abstraction** provided by the ML language
- we worked out a **language-agnostic** variant of the LCF-architecture with **process separation**



# The Architecture of CStar

how to maintain the program-proof states





# The Architecture of CStar

how to maintain the program-proof states

- compute symbolic heap information via **forward symbolic execution**



# The Architecture of CStar

how to maintain the program-proof states

- compute symbolic heap information via **forward symbolic execution**
- CStar now employs (but is largely **agnostic** of) QCP's symbolic executor for **separation-logic** reasoning

## QCP: A Practical Separation Logic-based C Program Verification Tool

Xiwei Wu<sup>1</sup>, Yueyang Feng<sup>1,\*</sup>, Xiaoyang Lu<sup>1,\*</sup>, Tianchuan Lin<sup>1</sup>, Kan Liu<sup>1</sup>,  
Zhiyi Wang<sup>2</sup>, Shushu Wu<sup>1</sup>, Lihan Xie<sup>1</sup>, Chengxi Yang<sup>1</sup>, Hongyi Zhong<sup>1</sup>, Naijun  
Zhan<sup>2</sup>, Zhenjiang Hu<sup>2</sup>, and Qinxiang Cao<sup>1,†</sup>

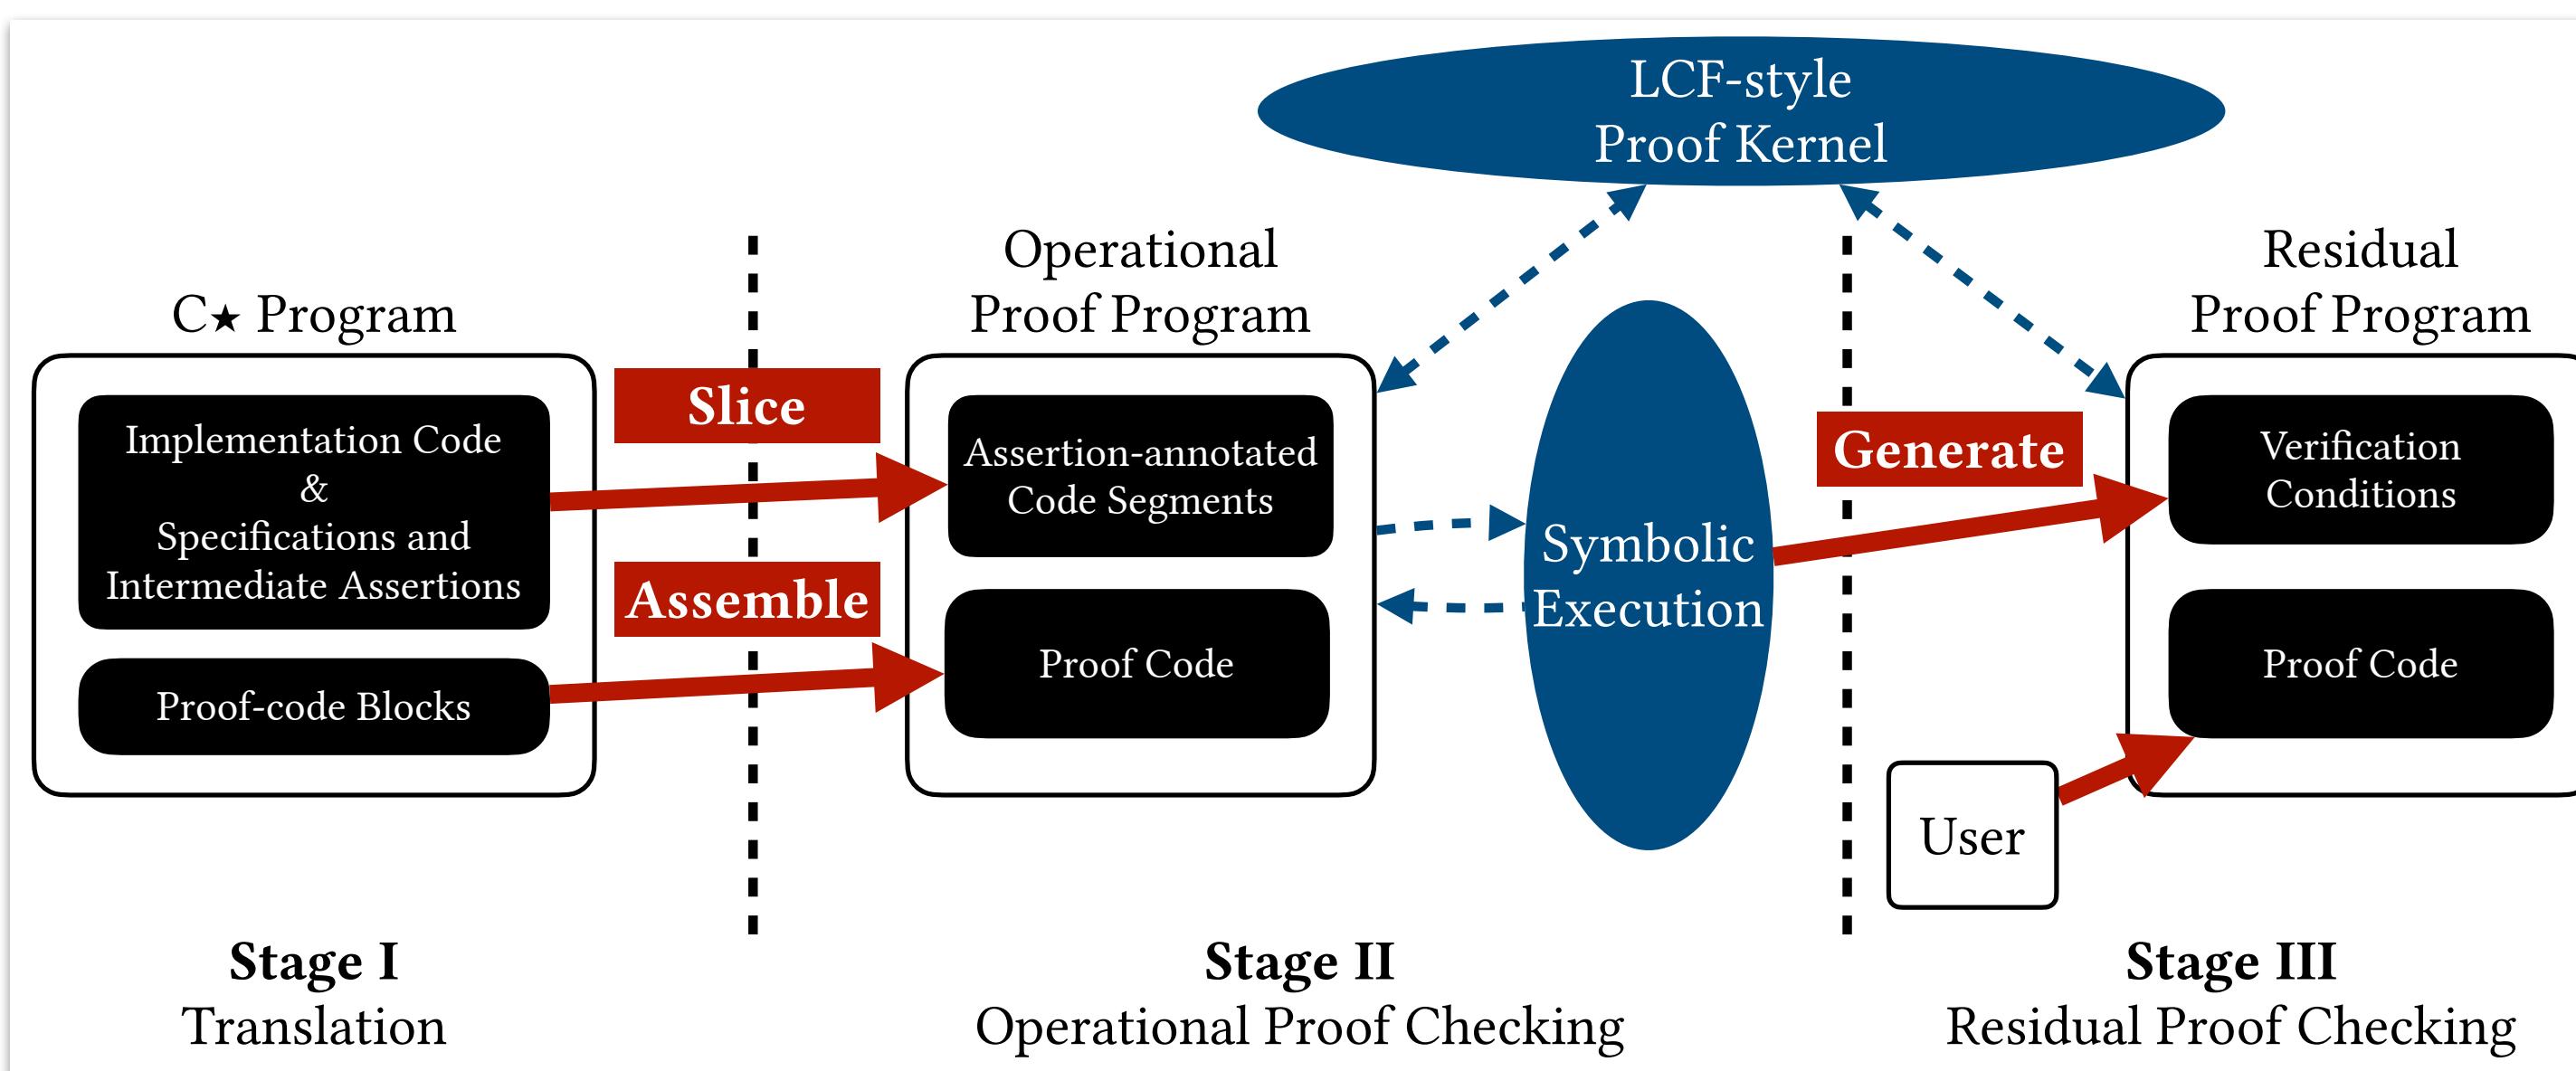
<sup>1</sup> Shanghai Jiao Tong University  
{yashen, fyyvexoben, luxy1115, caoqinxiang}@sjtu.edu.cn

<sup>2</sup> Peking University  
2301111964@stu.pku.edu.cn

# The Architecture of CStar

how to maintain the program-proof states

- compute symbolic heap information via **forward symbolic execution**
- CStar now employs (but is largely **agnostic** of) QCP's symbolic executor for **separation-logic** reasoning
- QCP evolves the program-proof states for program code; CStar evolves the program-proof states for proof code



## QCP: A Practical Separation Logic-based C Program Verification Tool

Xiwei Wu<sup>1</sup>, Yueyang Feng<sup>1,\*</sup>, Xiaoyang Lu<sup>1,\*</sup>, Tianchuan Lin<sup>1</sup>, Kan Liu<sup>1</sup>,  
Zhiyi Wang<sup>2</sup>, Shushu Wu<sup>1</sup>, Lihan Xie<sup>1</sup>, Chengxi Yang<sup>1</sup>, Hongyi Zhong<sup>1</sup>, Naijun  
Zhan<sup>2</sup>, Zhenjiang Hu<sup>2</sup>, and Qinxiang Cao<sup>1,†</sup>

<sup>1</sup> Shanghai Jiao Tong University  
{yashen, fyyvexoben, luxy1115, caoqinxiang}@sjtu.edu.cn

<sup>2</sup> Peking University  
2301111964@stu.pku.edu.cn



users can now use the full power of C  
to conduct proof and verification



# CStar Demo: Tactics

all implemented outside the proof kernel



# CStar Demo: SMT Integration

achieve controllable automation



# CStar Demo: ProofEncapsulation

one can implement program-proof libraries

# CStar

## Unifying Programming and Verification in C





# Unifying Programming and Verification in C

prob i: proof encapsulation





# Unifying Programming and Verification in C

prob i: proof encapsulation

prob ii: programmable proof



## Unifying Programming and Verification in C

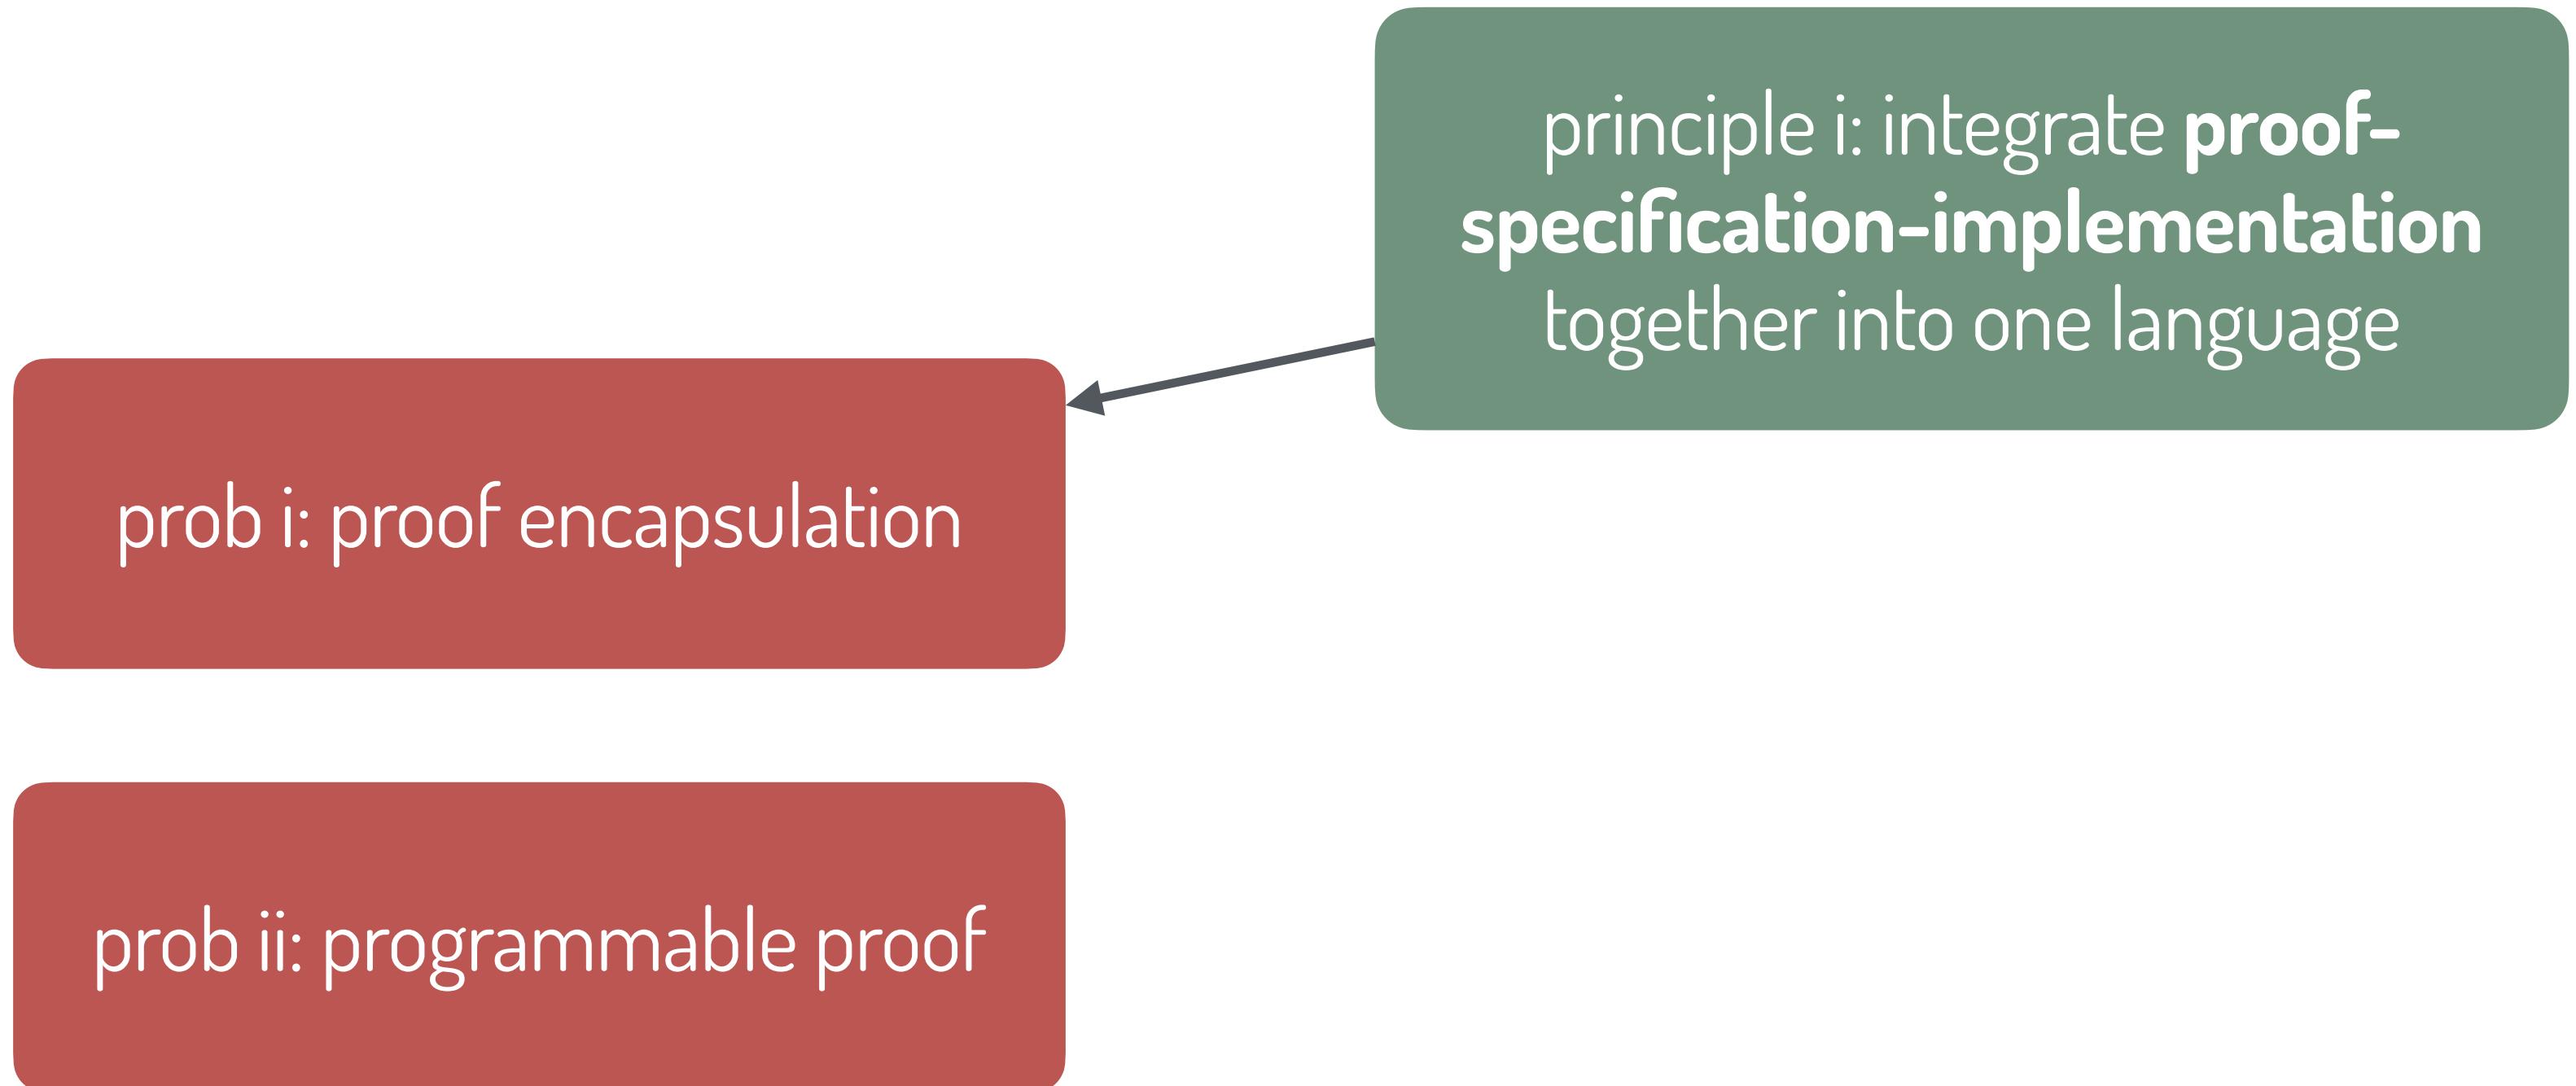
prob i: proof encapsulation

principle i: integrate **proof-specification-implementation** together into one language

prob ii: programmable proof



## Unifying Programming and Verification in C



## Unifying Programming and Verification in C



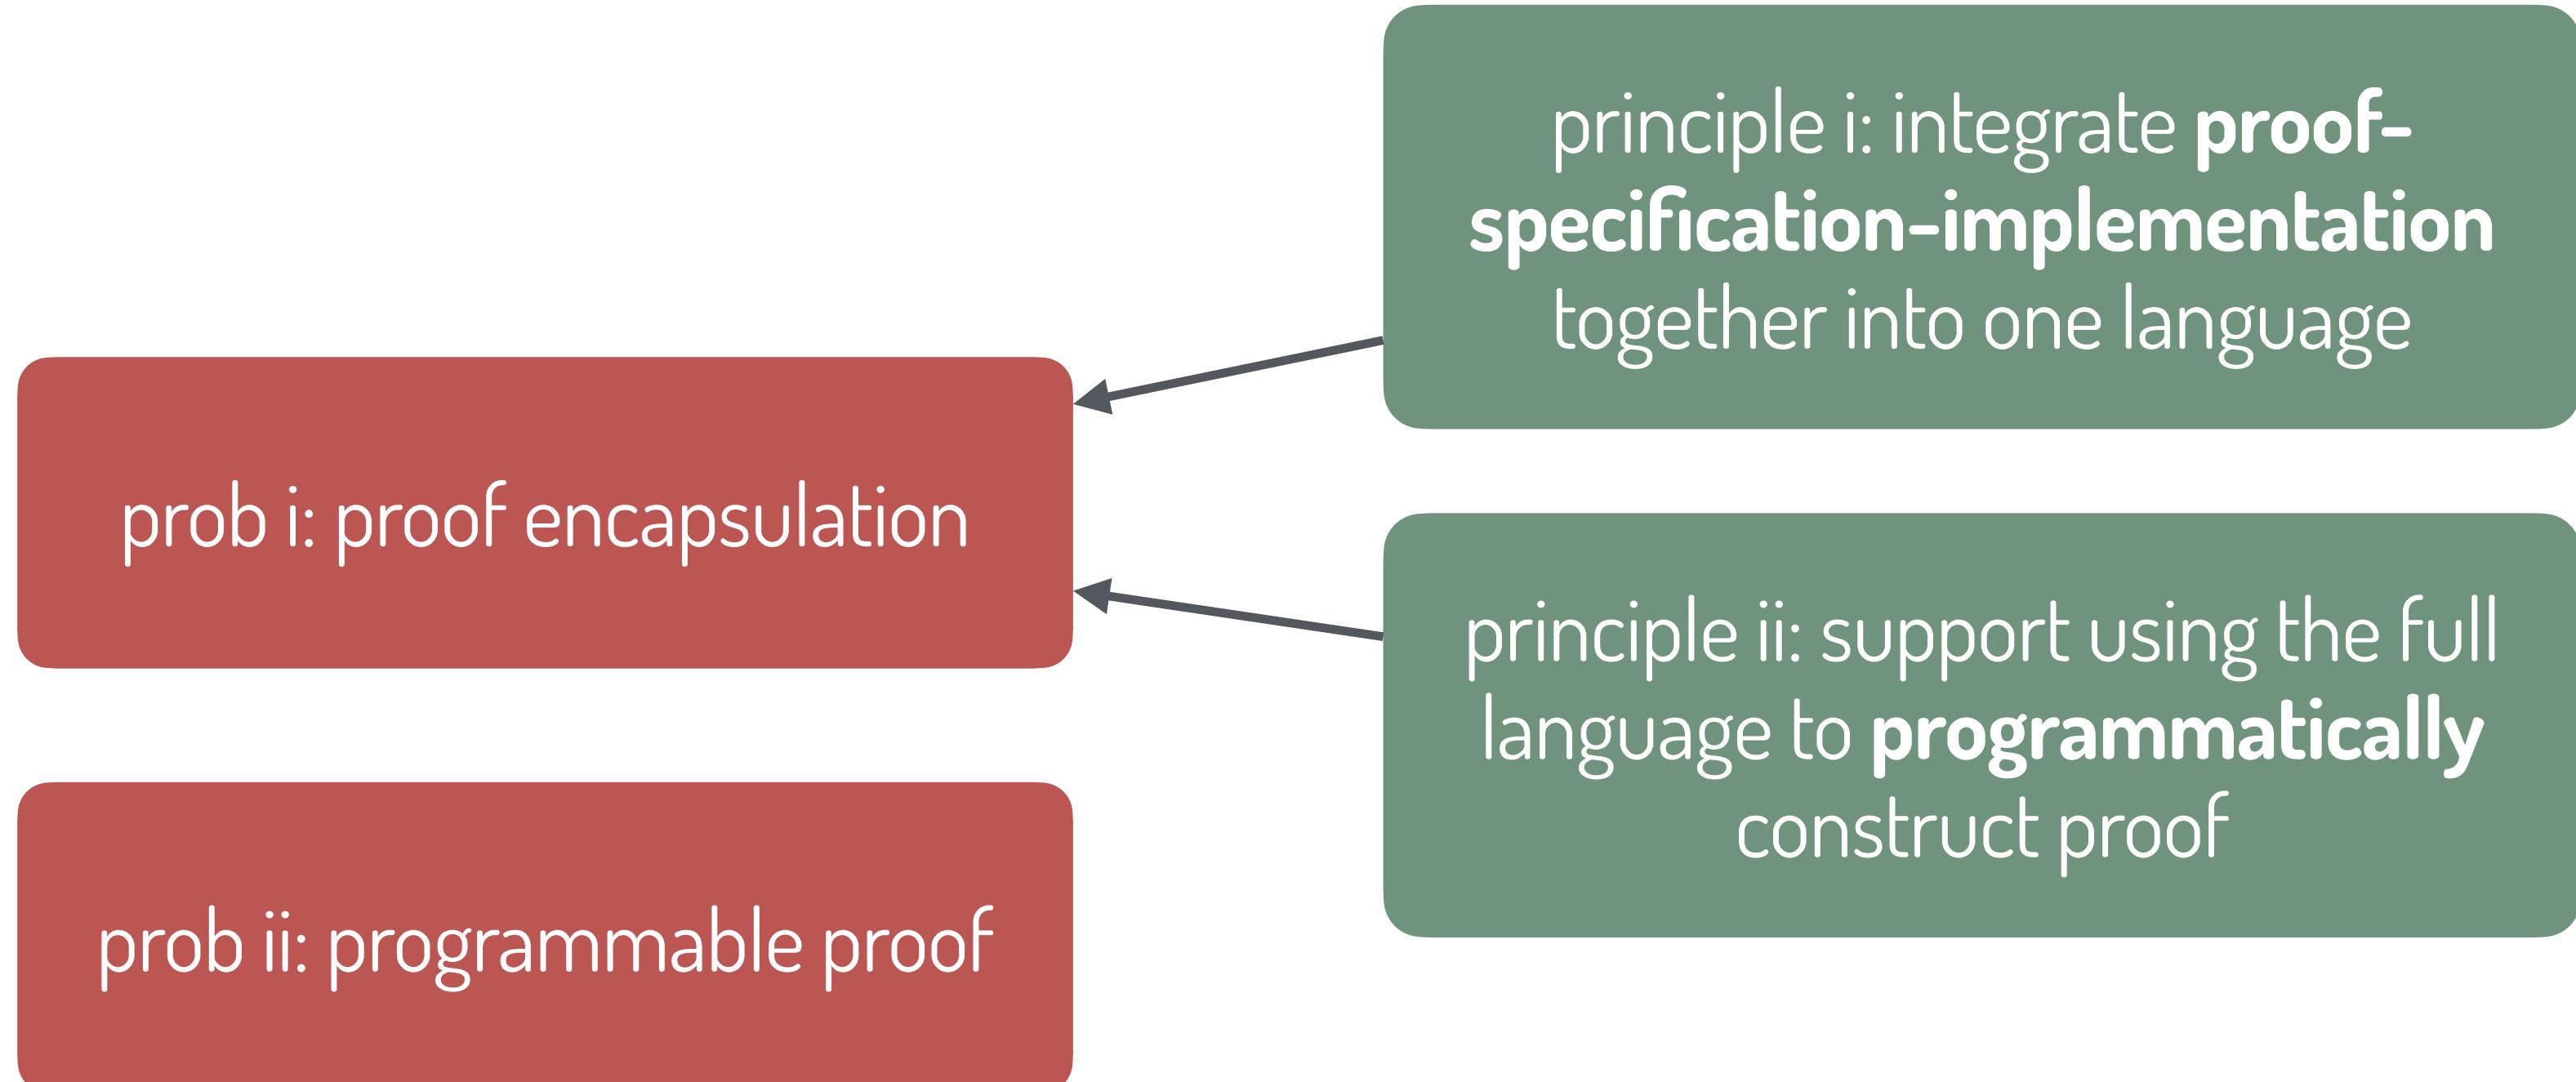
prob i: proof encapsulation

prob ii: programmable proof

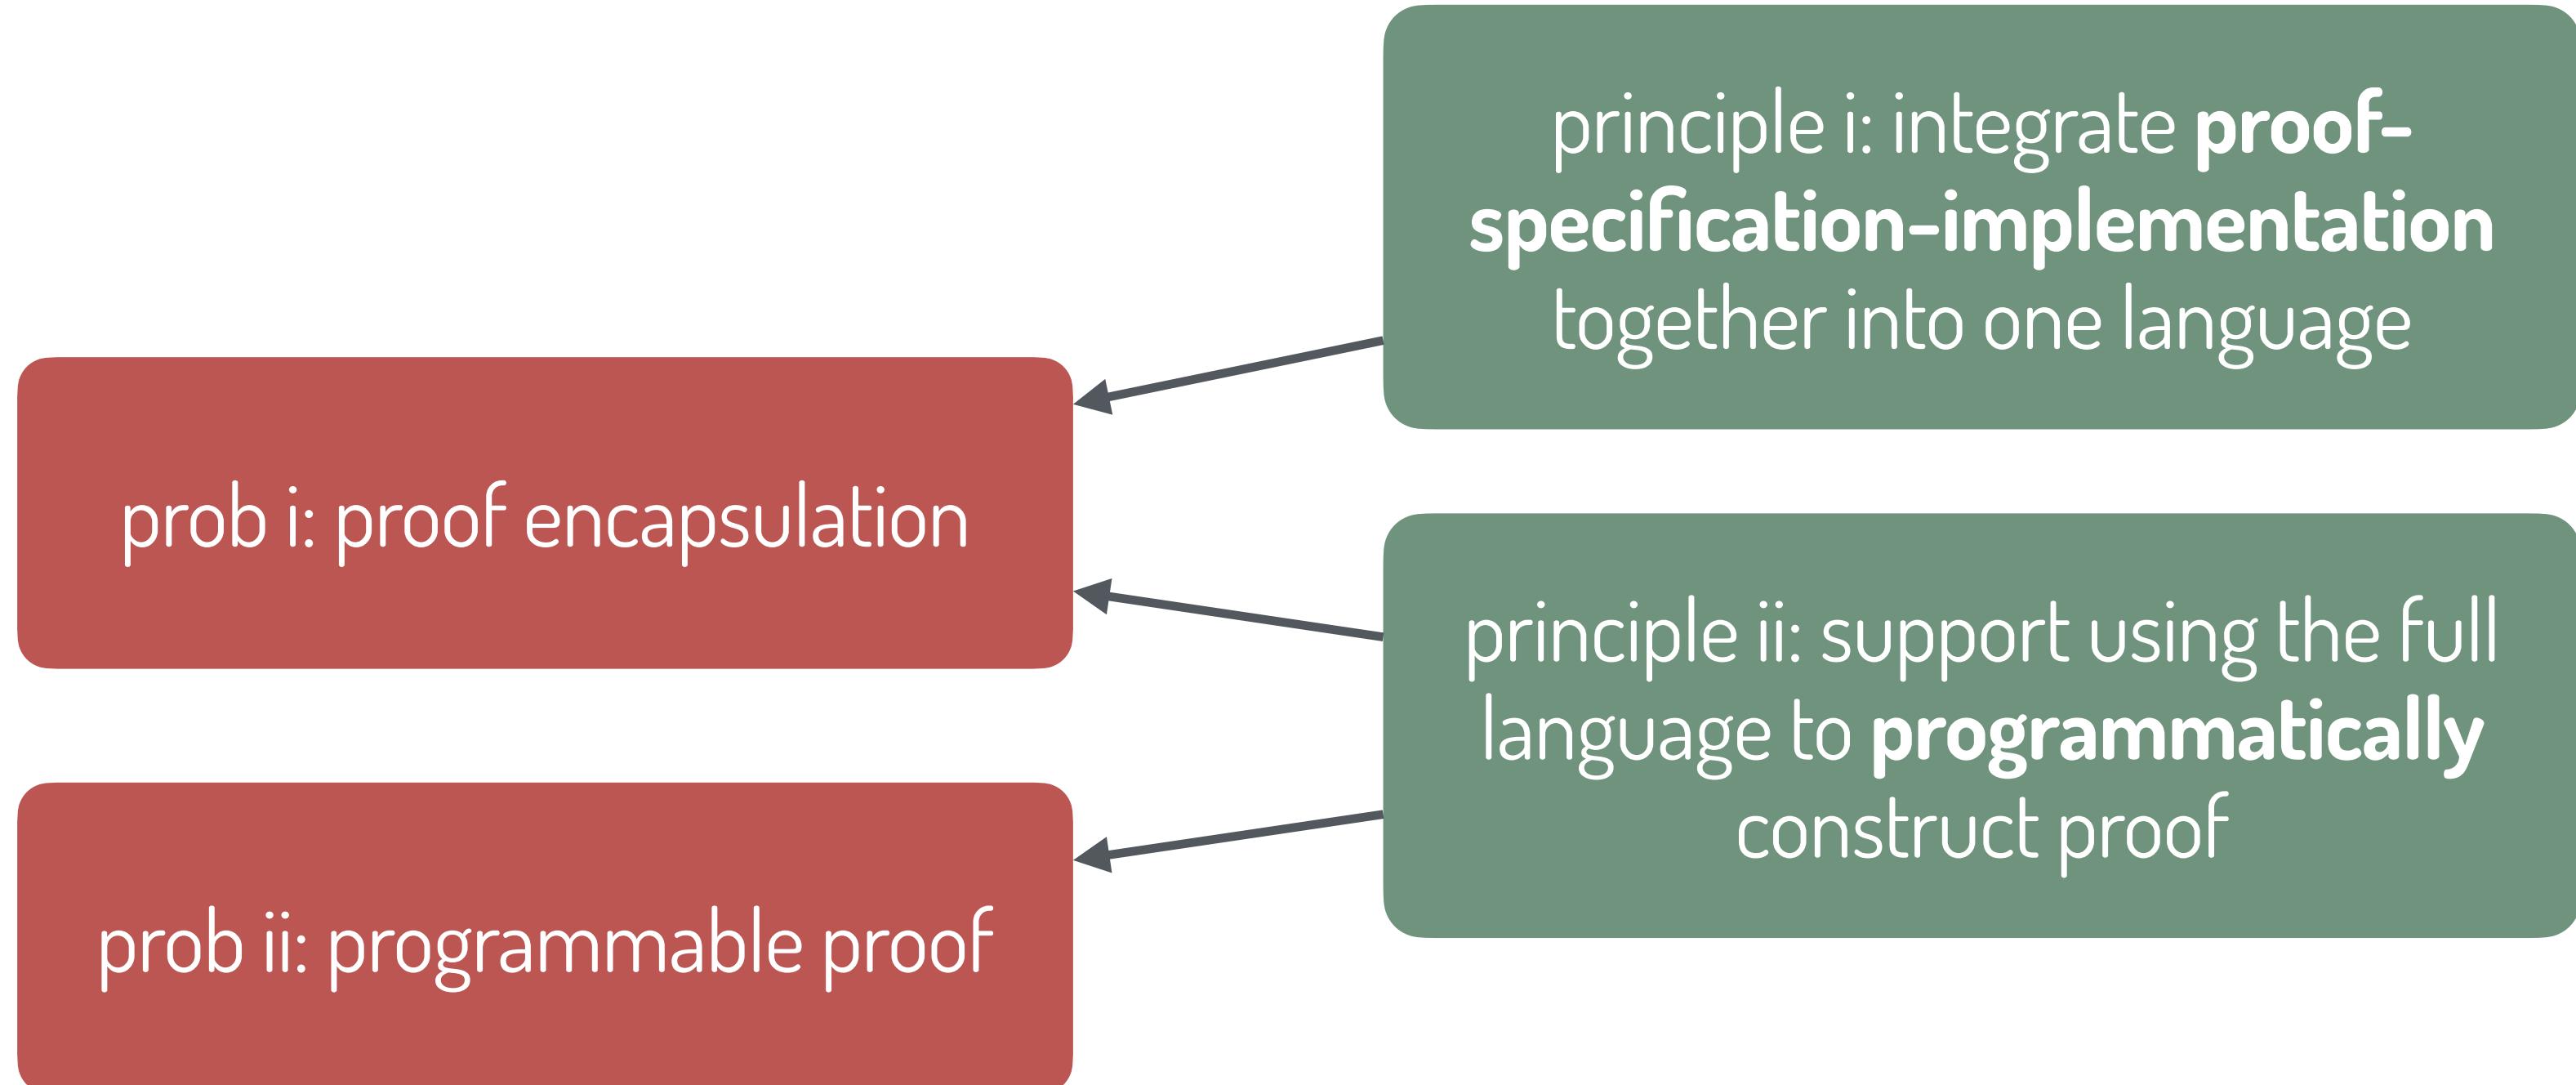
principle i: integrate **proof-specification-implementation** together into one language

principle ii: support using the full language to **programmatically** construct proof

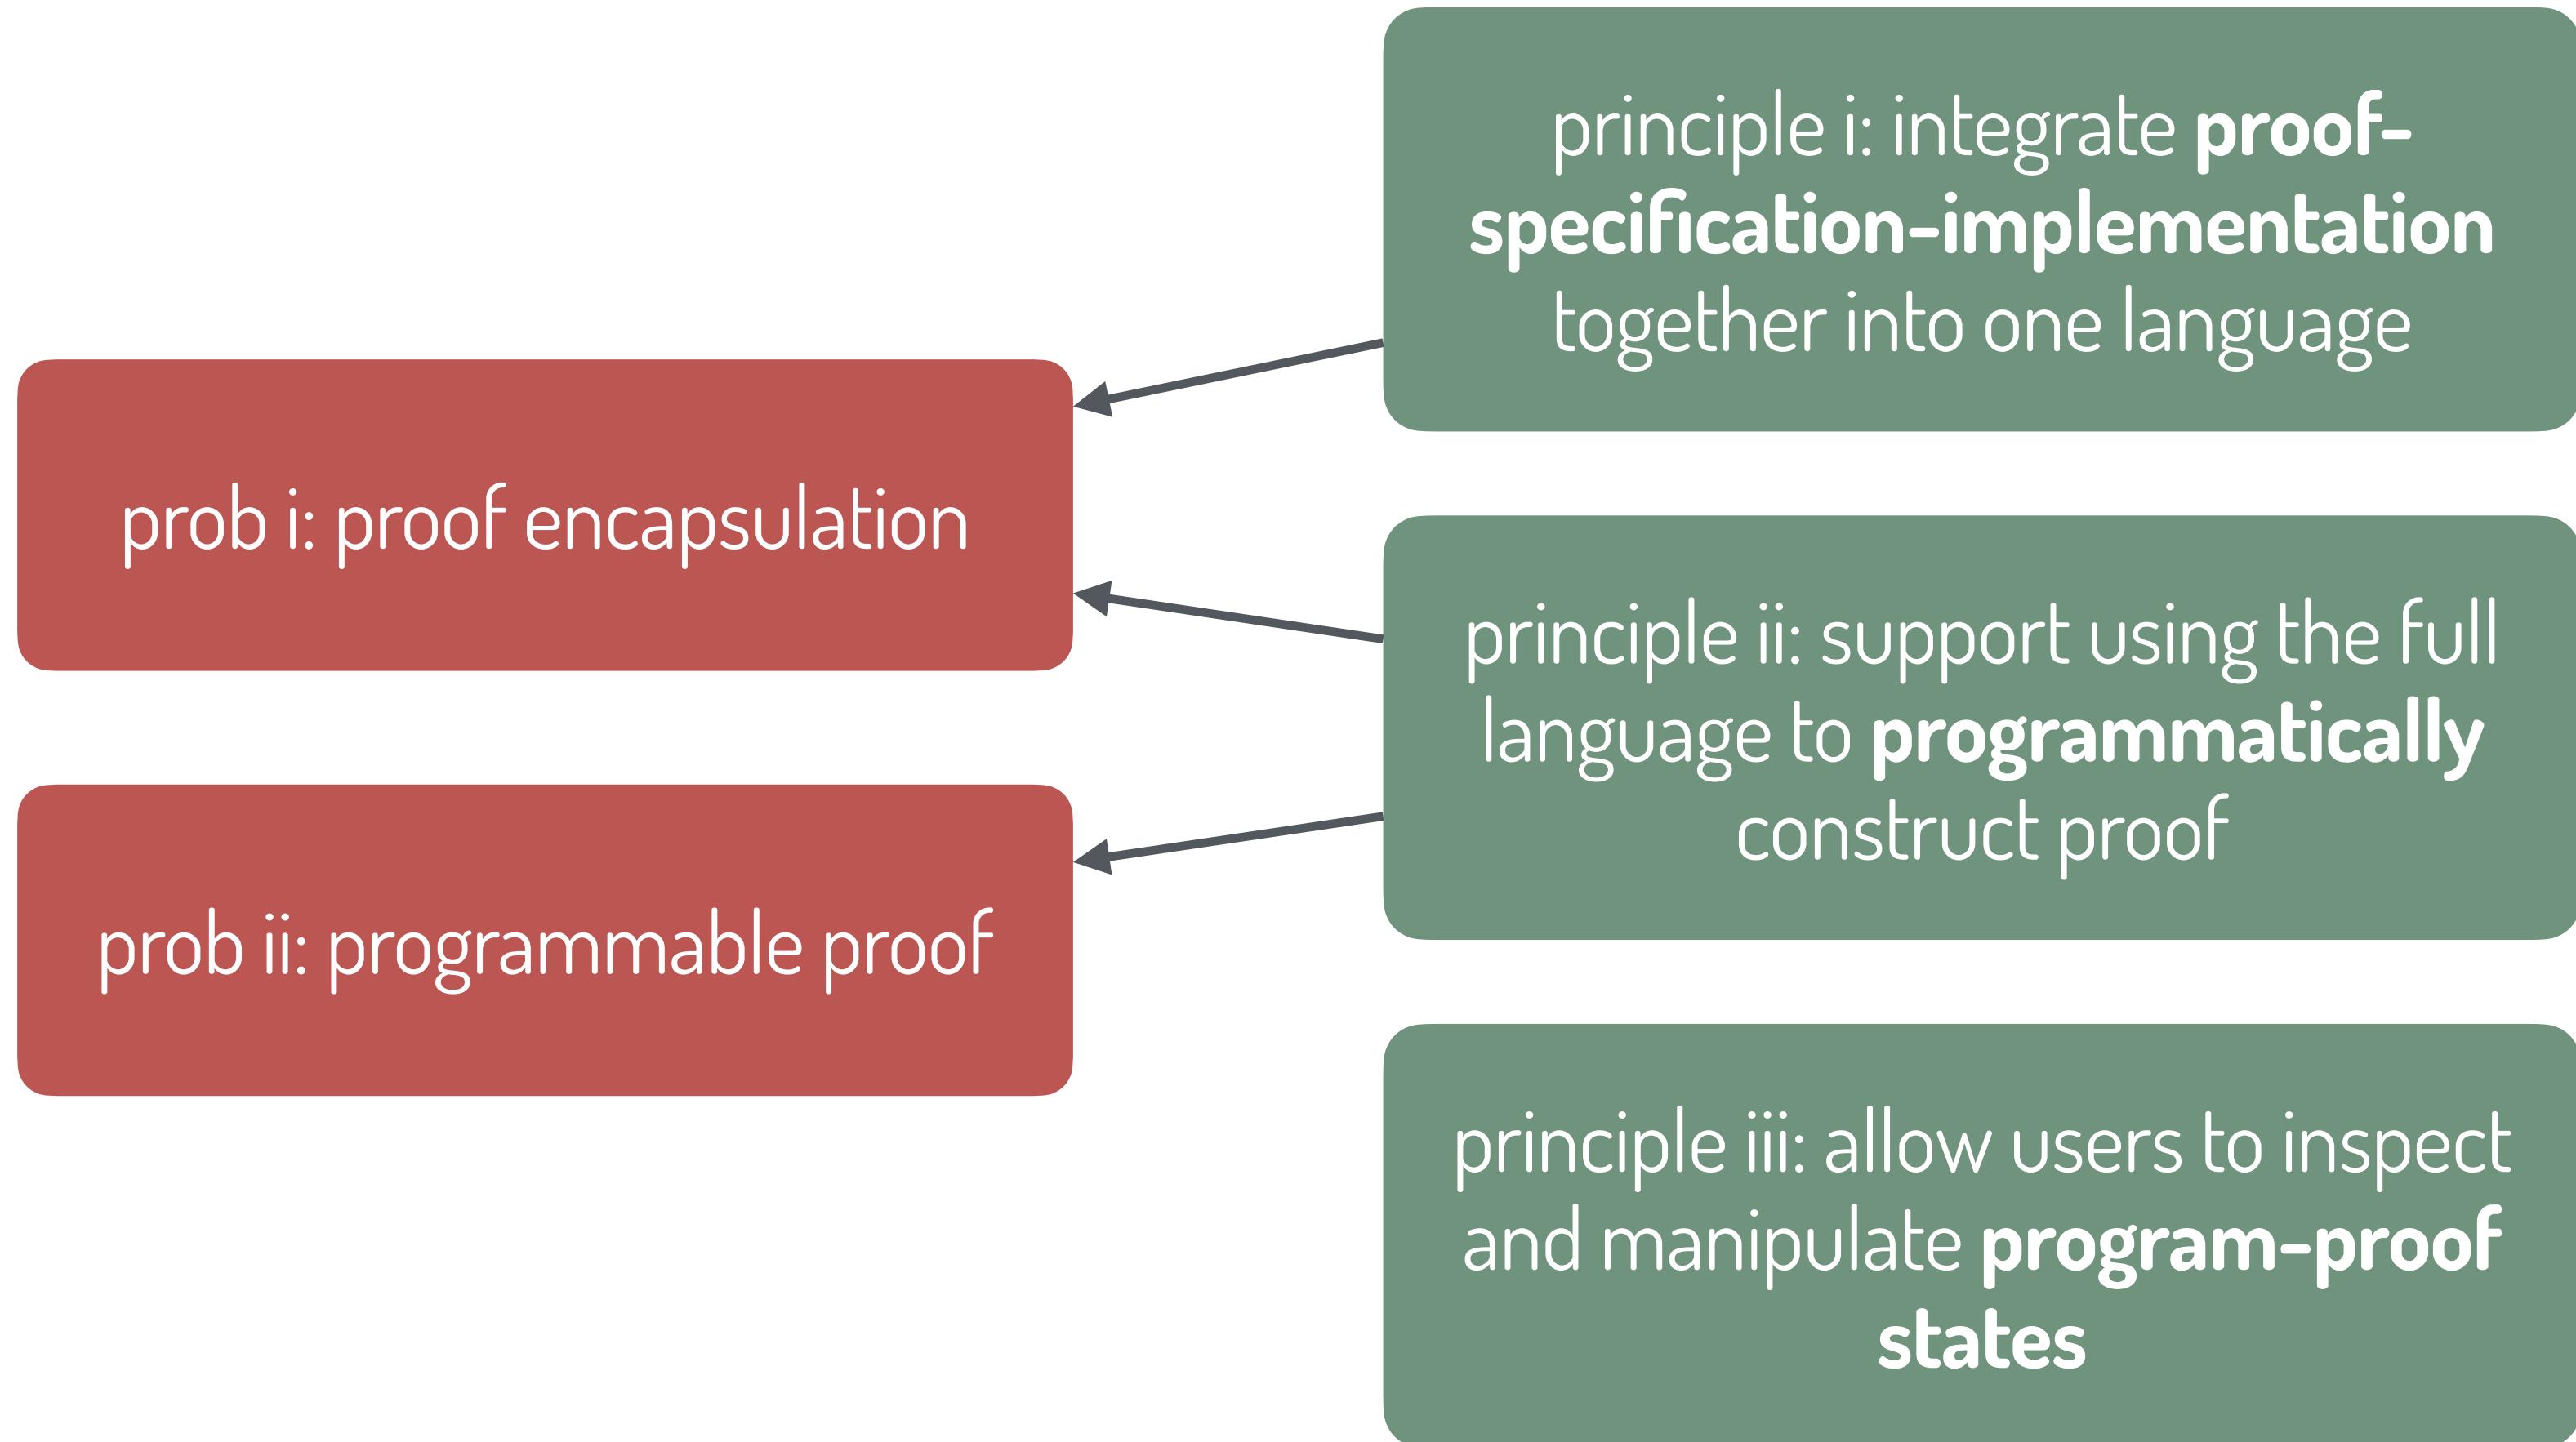
## Unifying Programming and Verification in C



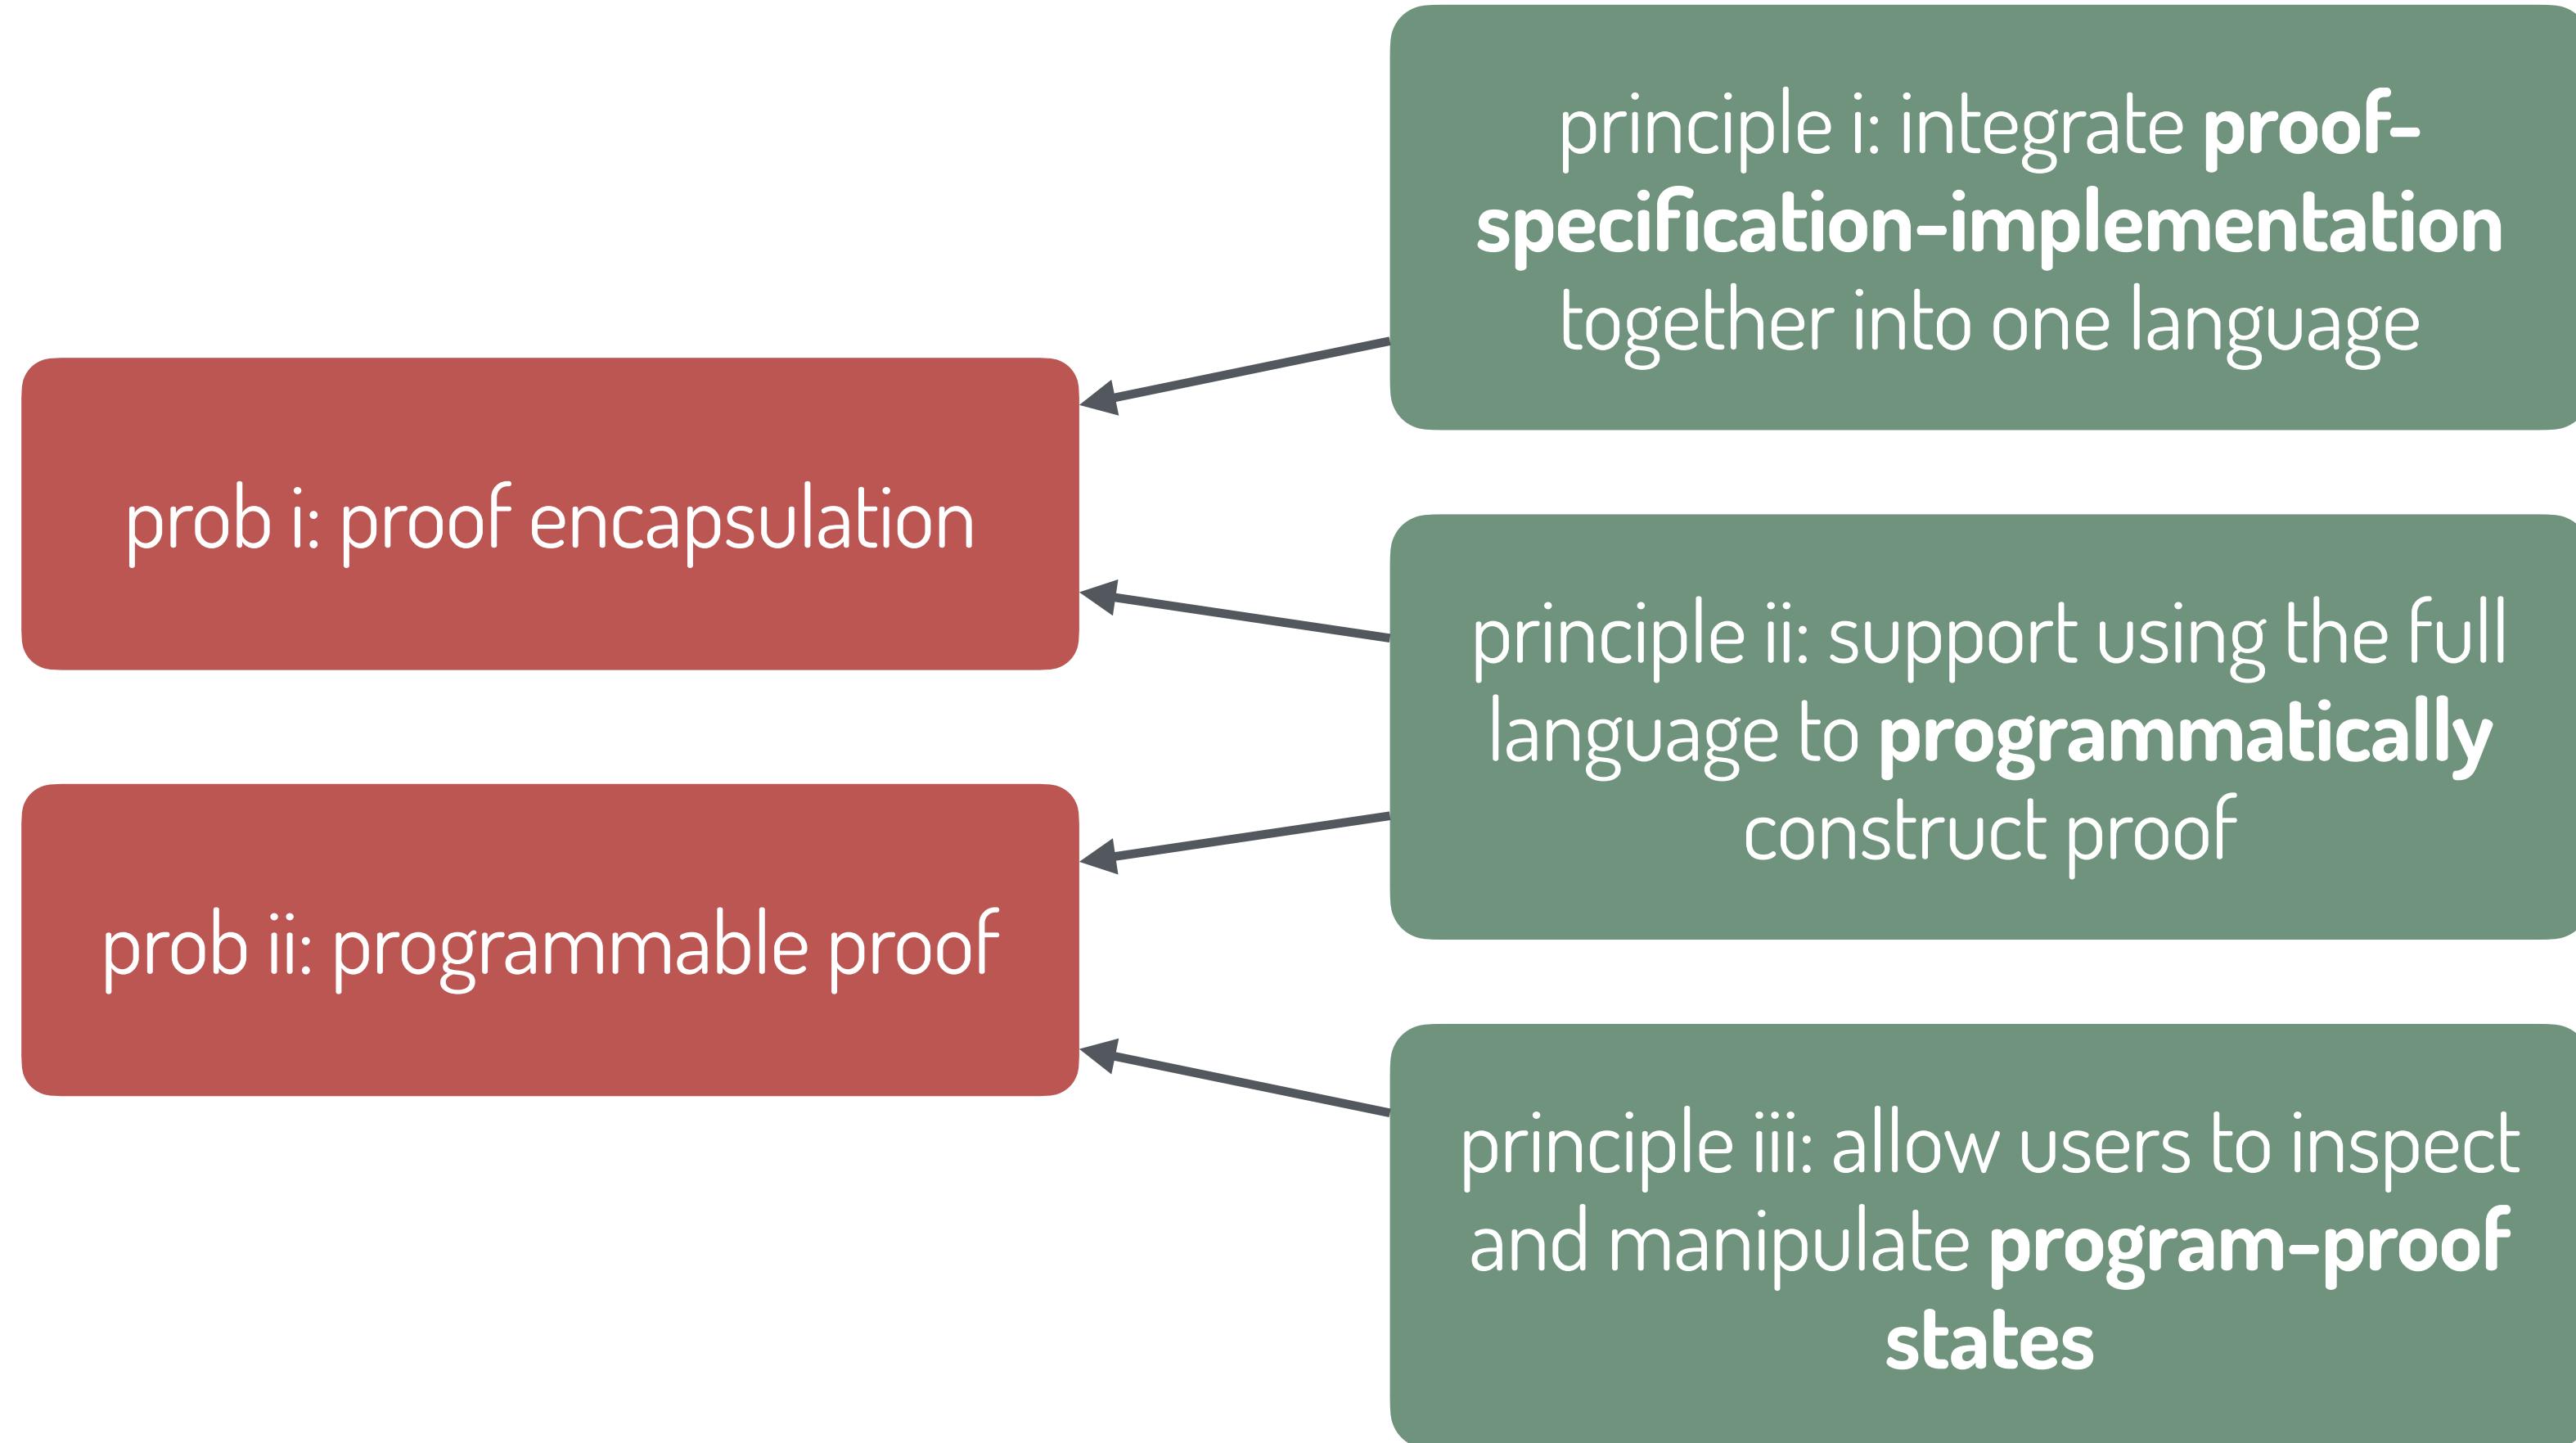
## Unifying Programming and Verification in C



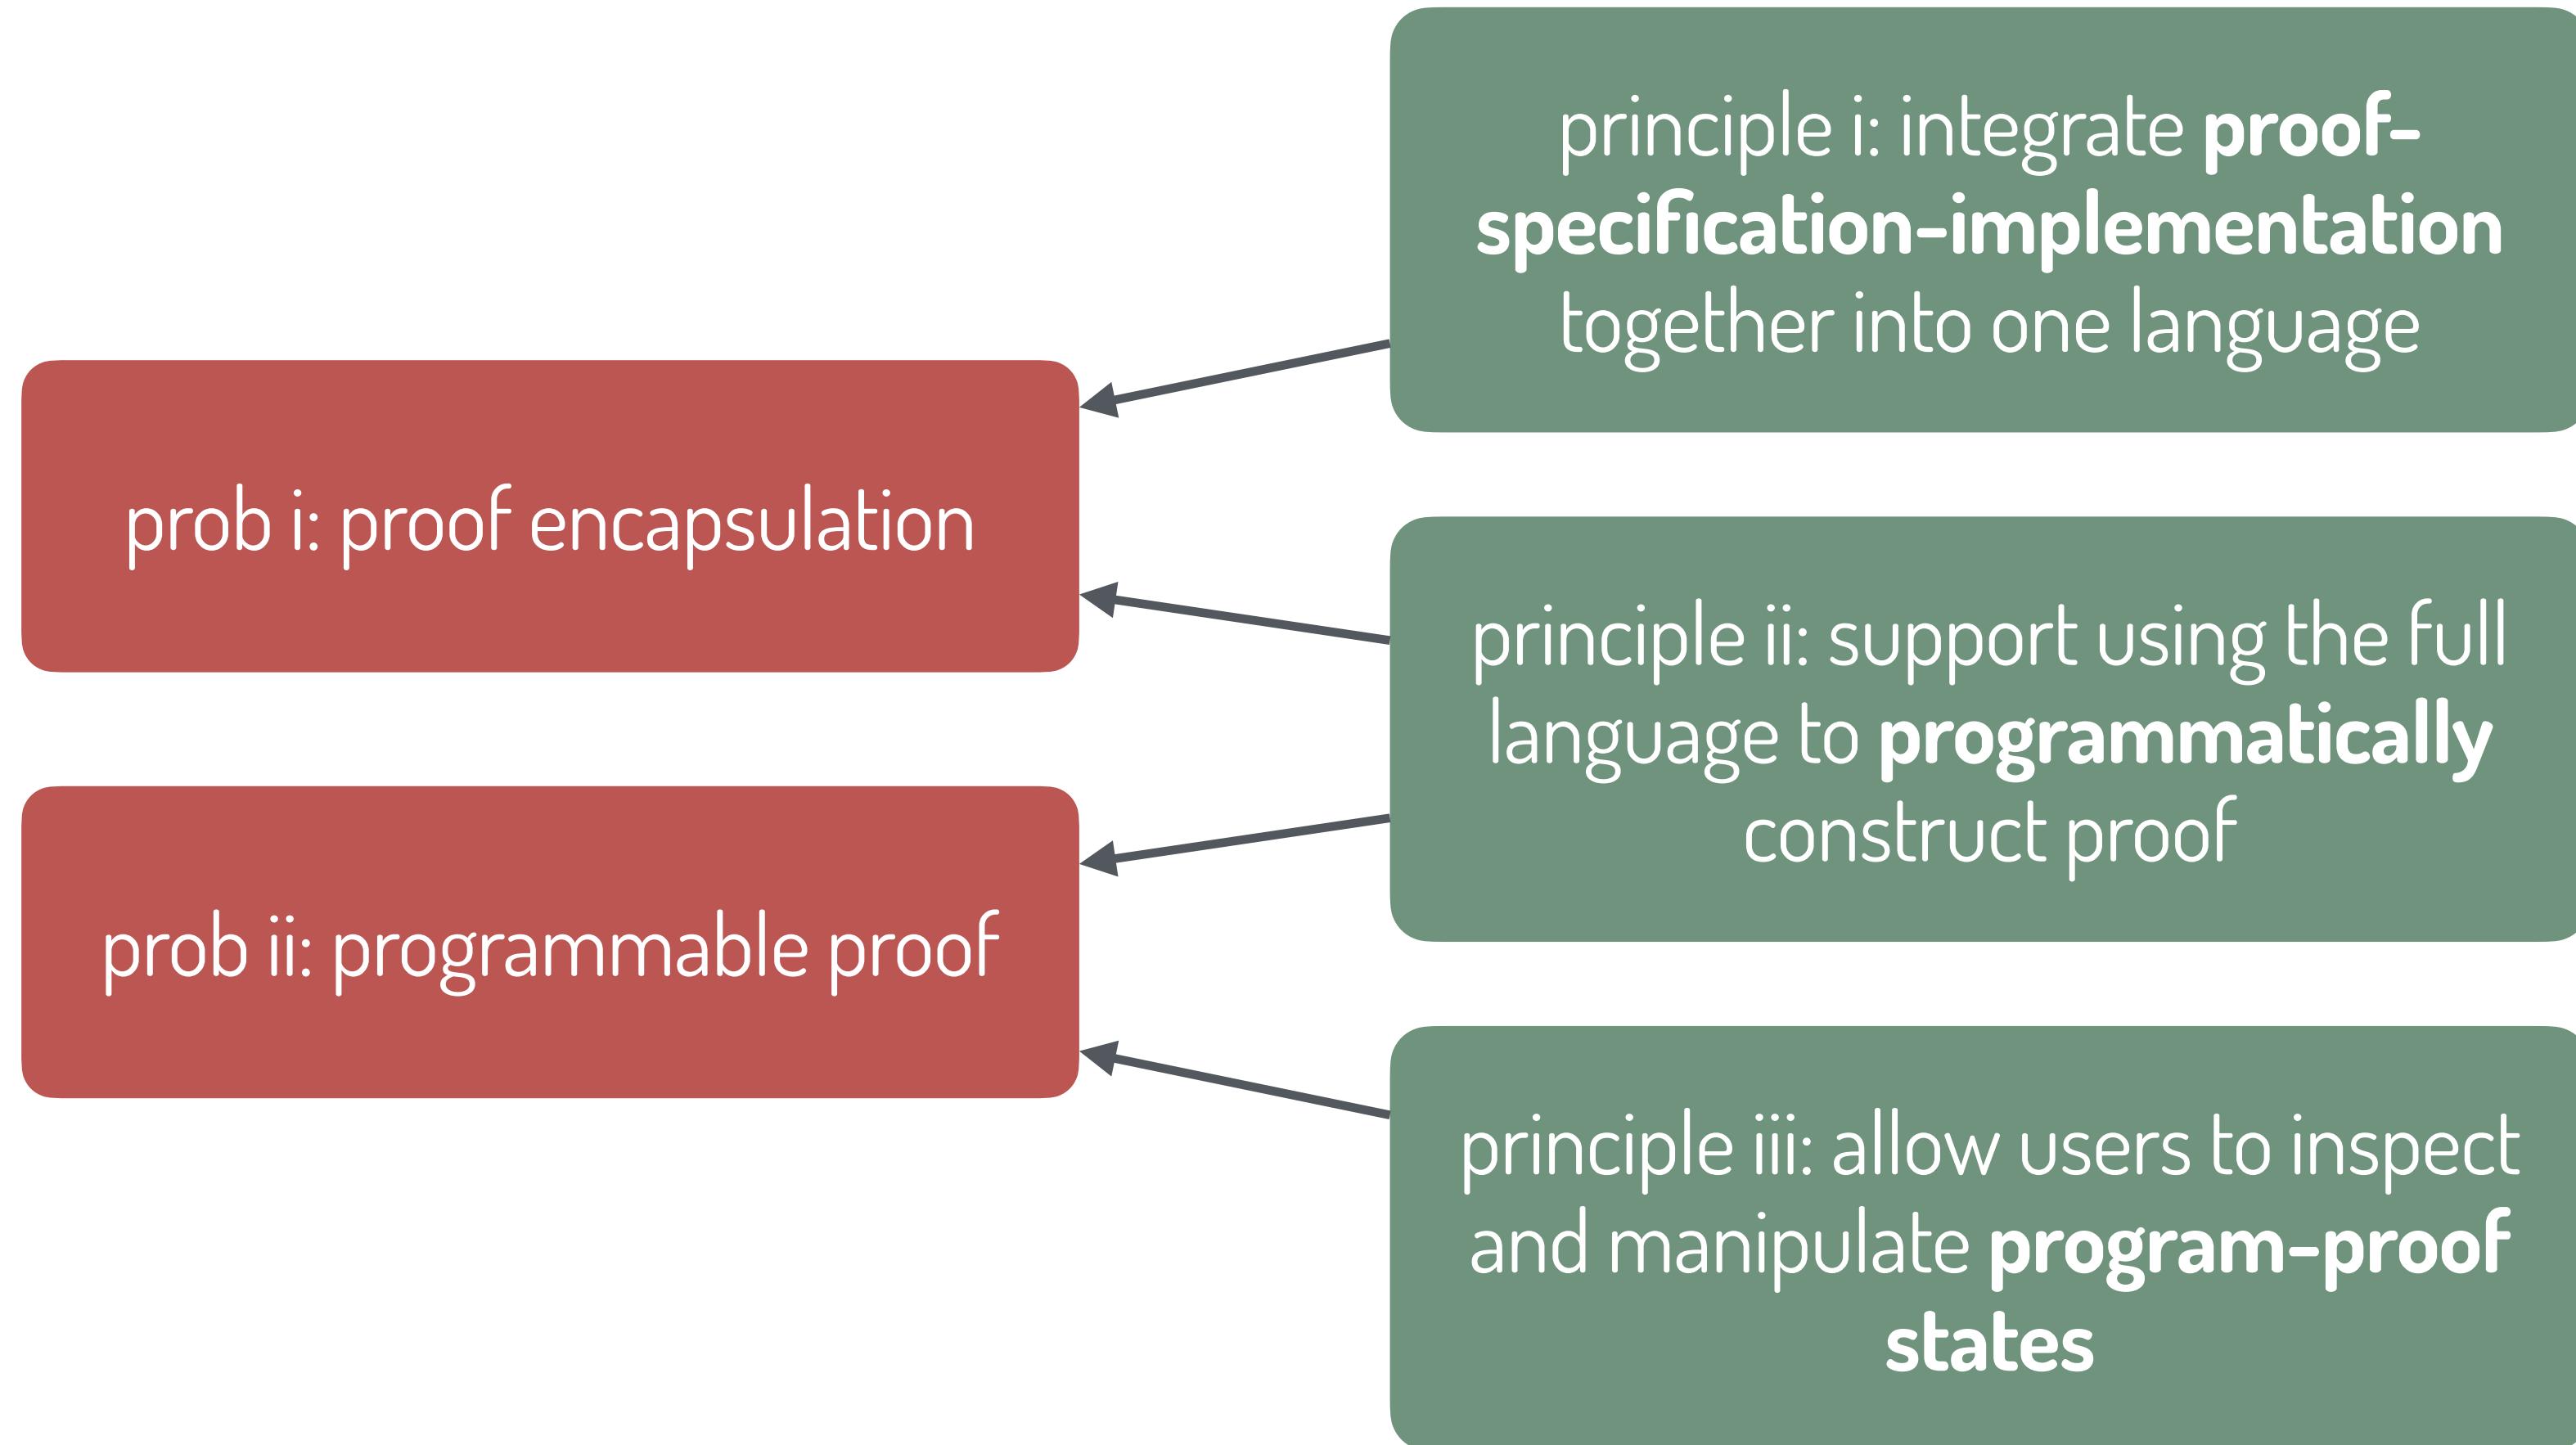
## Unifying Programming and Verification in C



## Unifying Programming and Verification in C

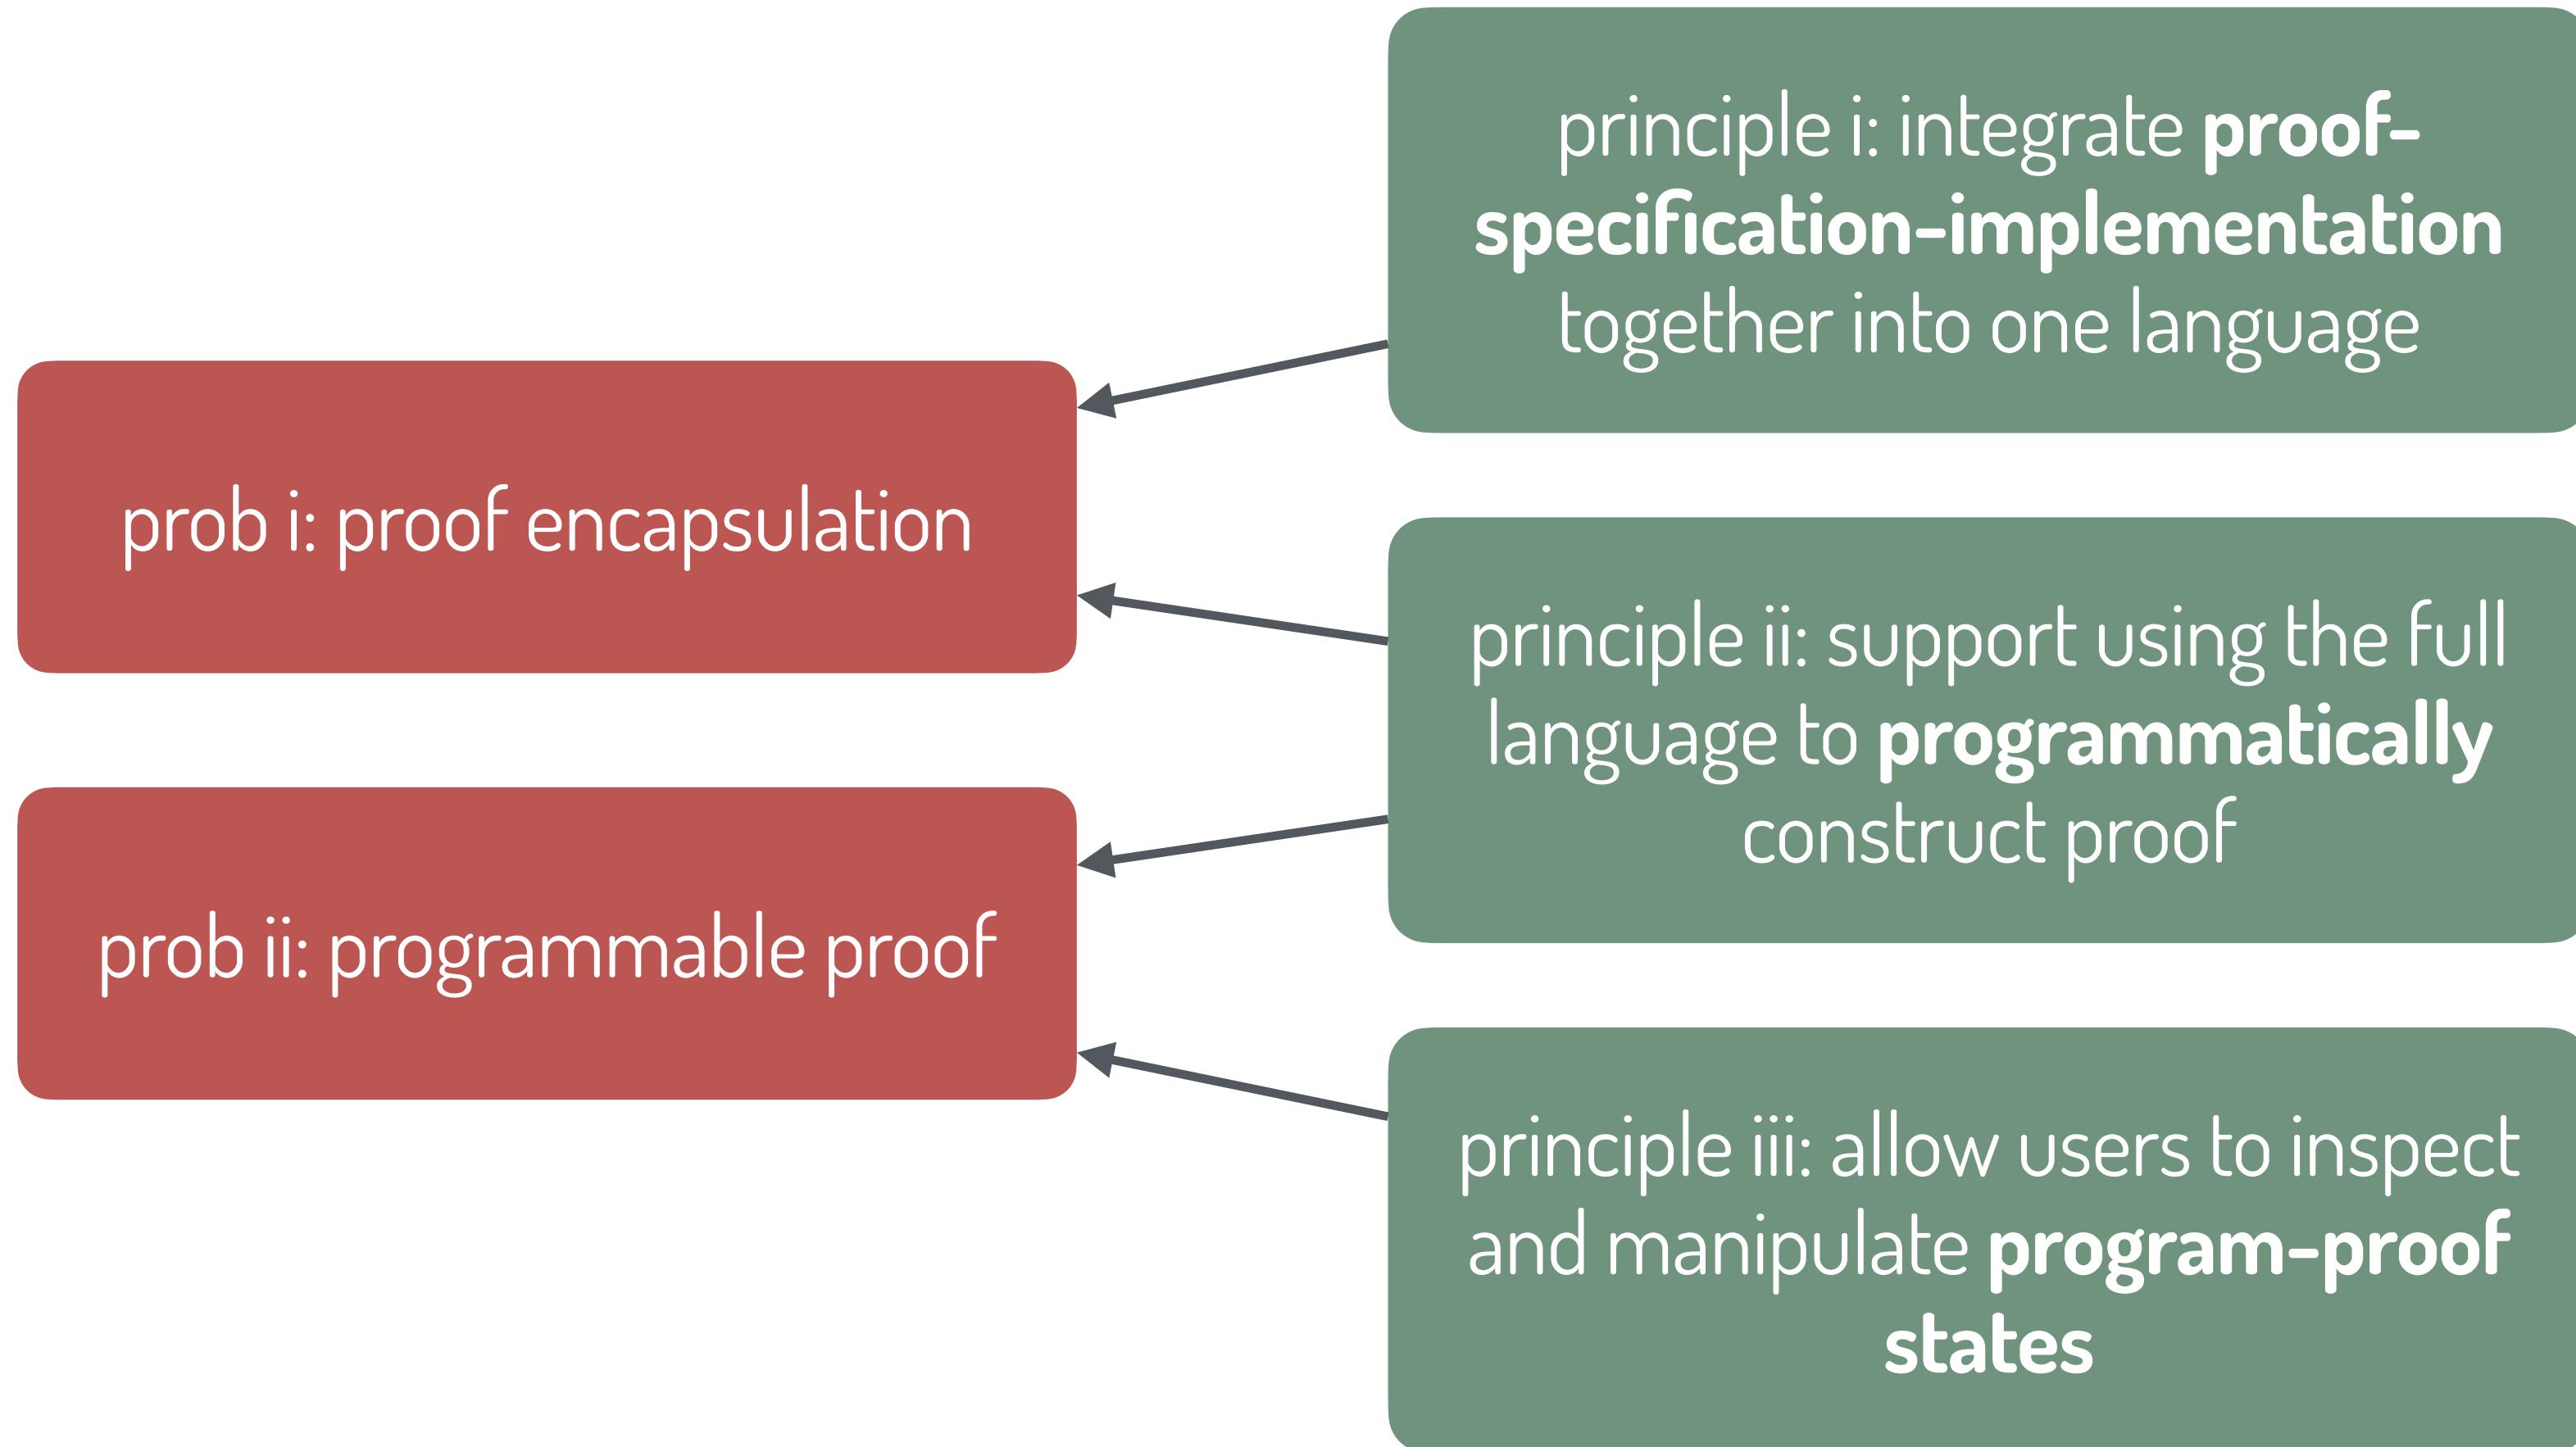


## Unifying Programming and Verification in C



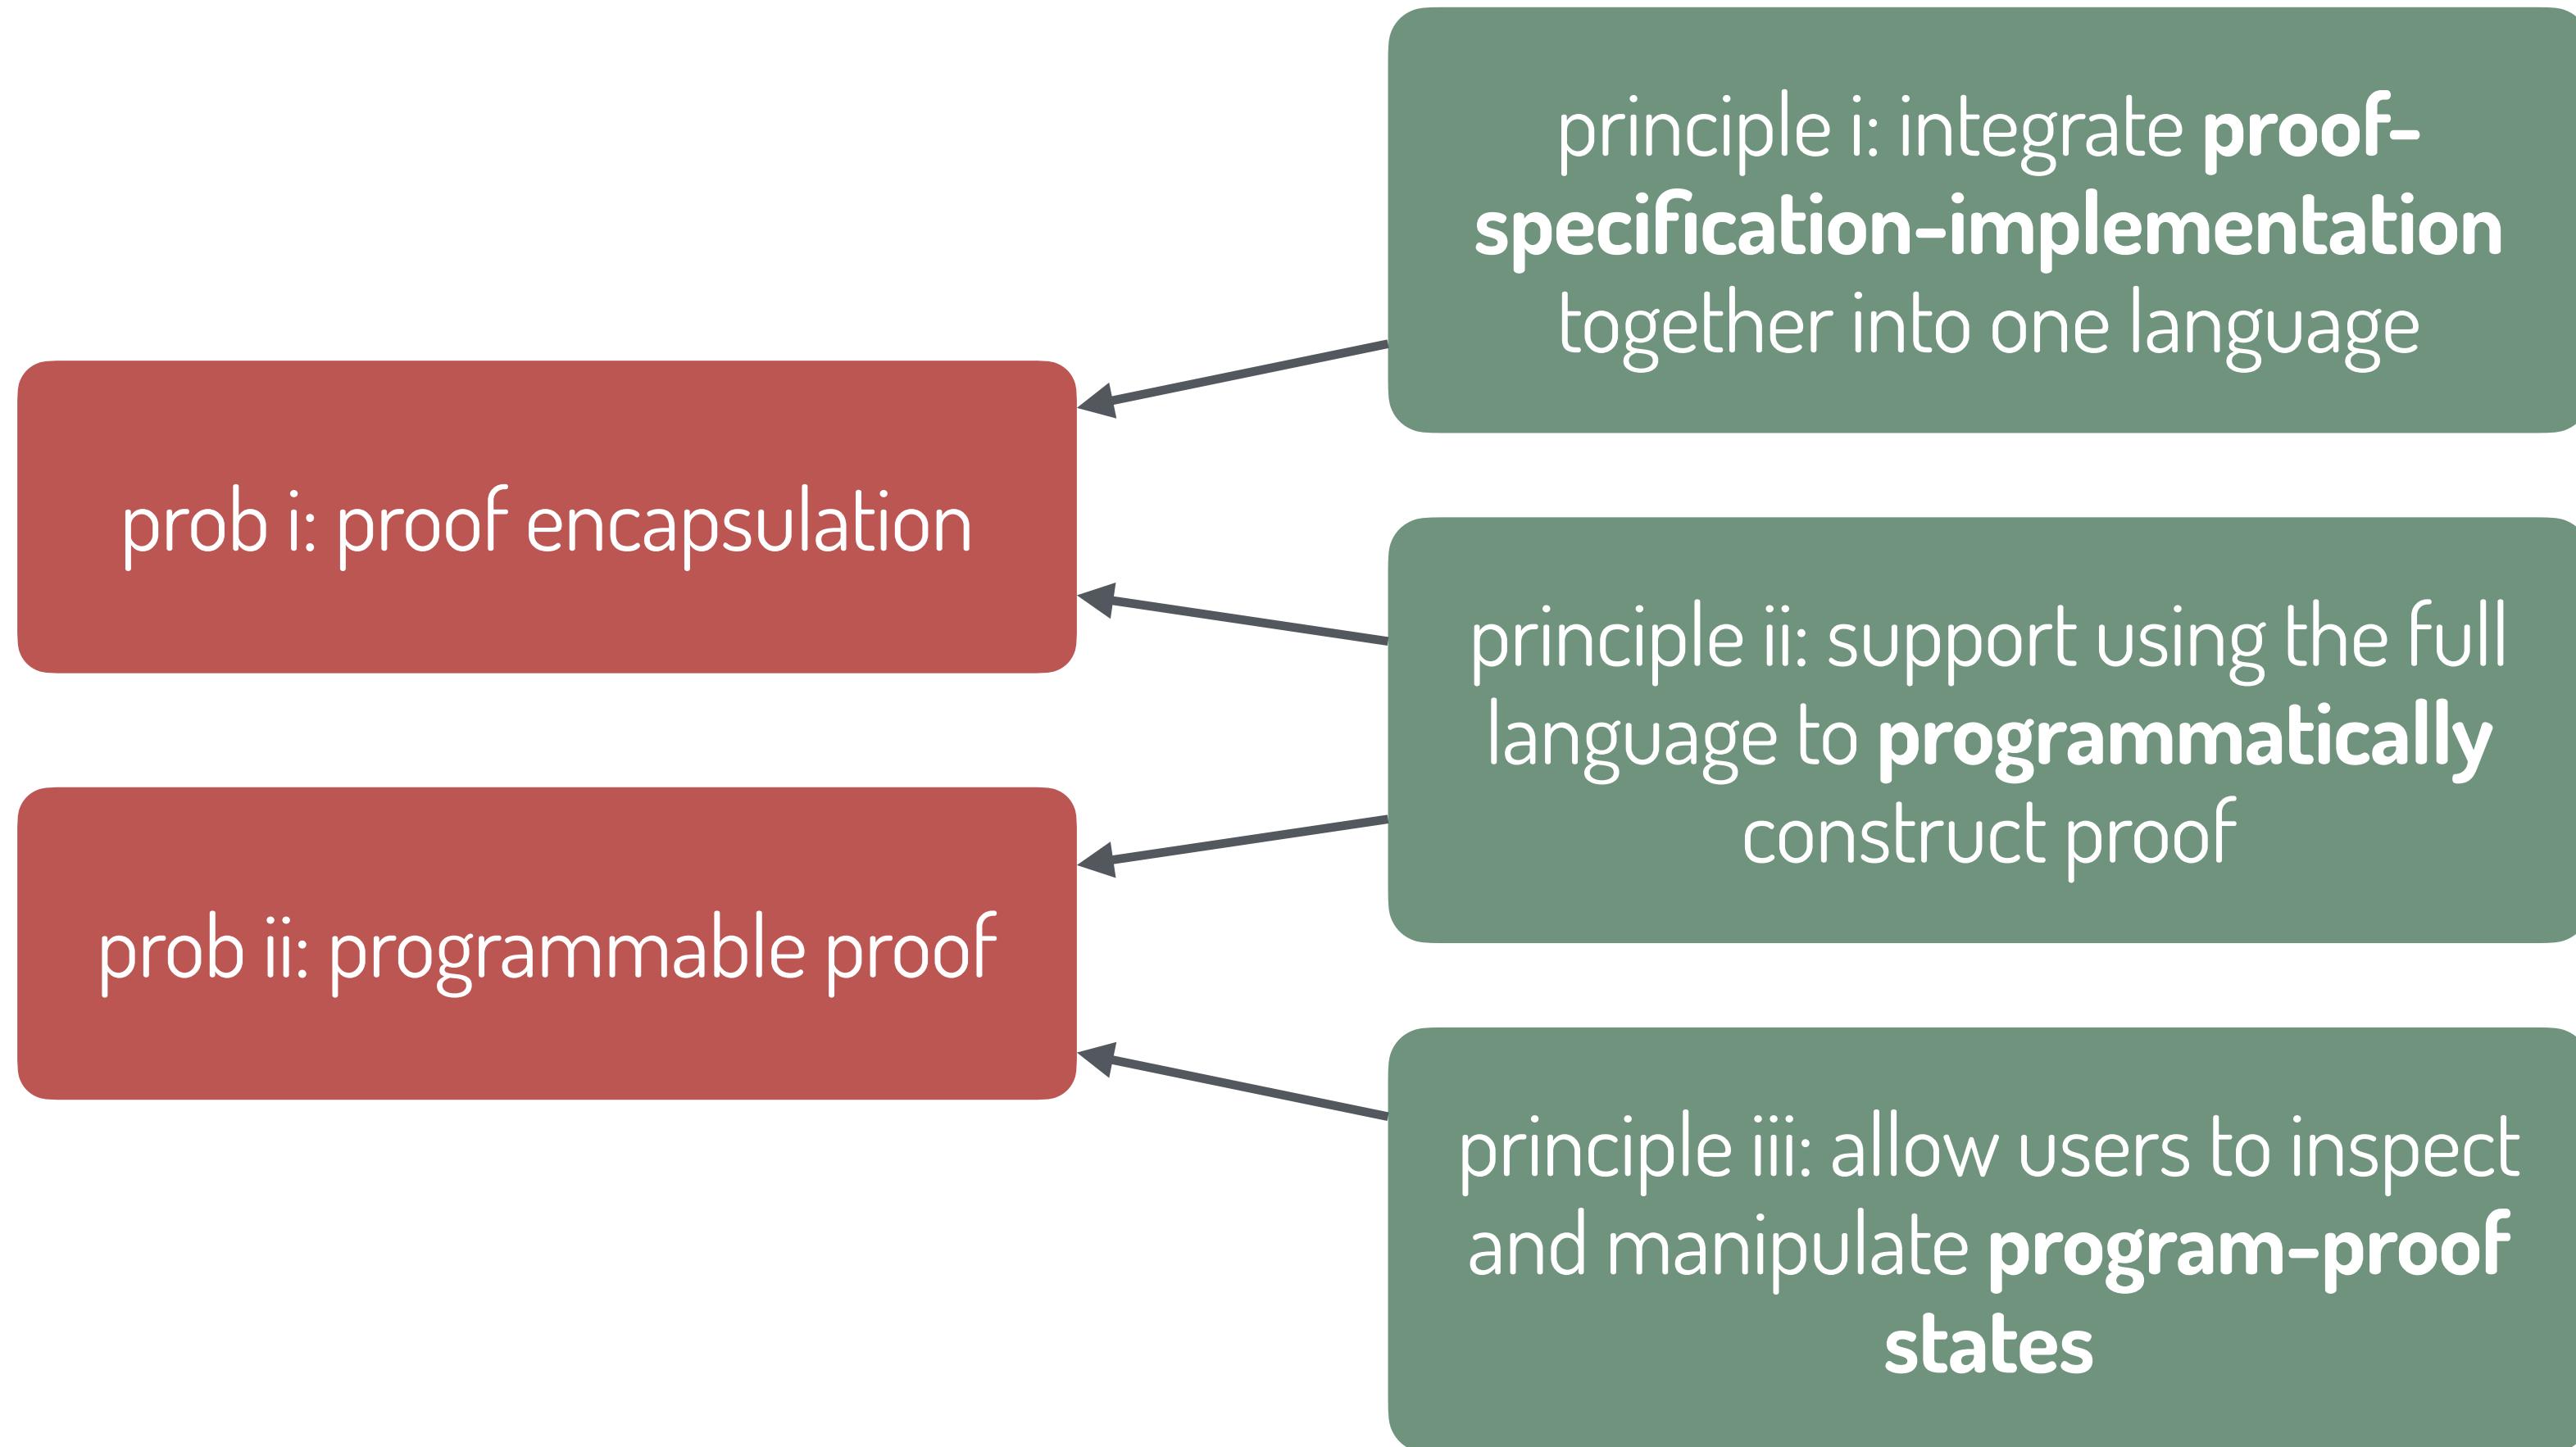
- CStar reuses clang's toolchain
  - use clang for compilation
  - use clangd for language server

## Unifying Programming and Verification in C



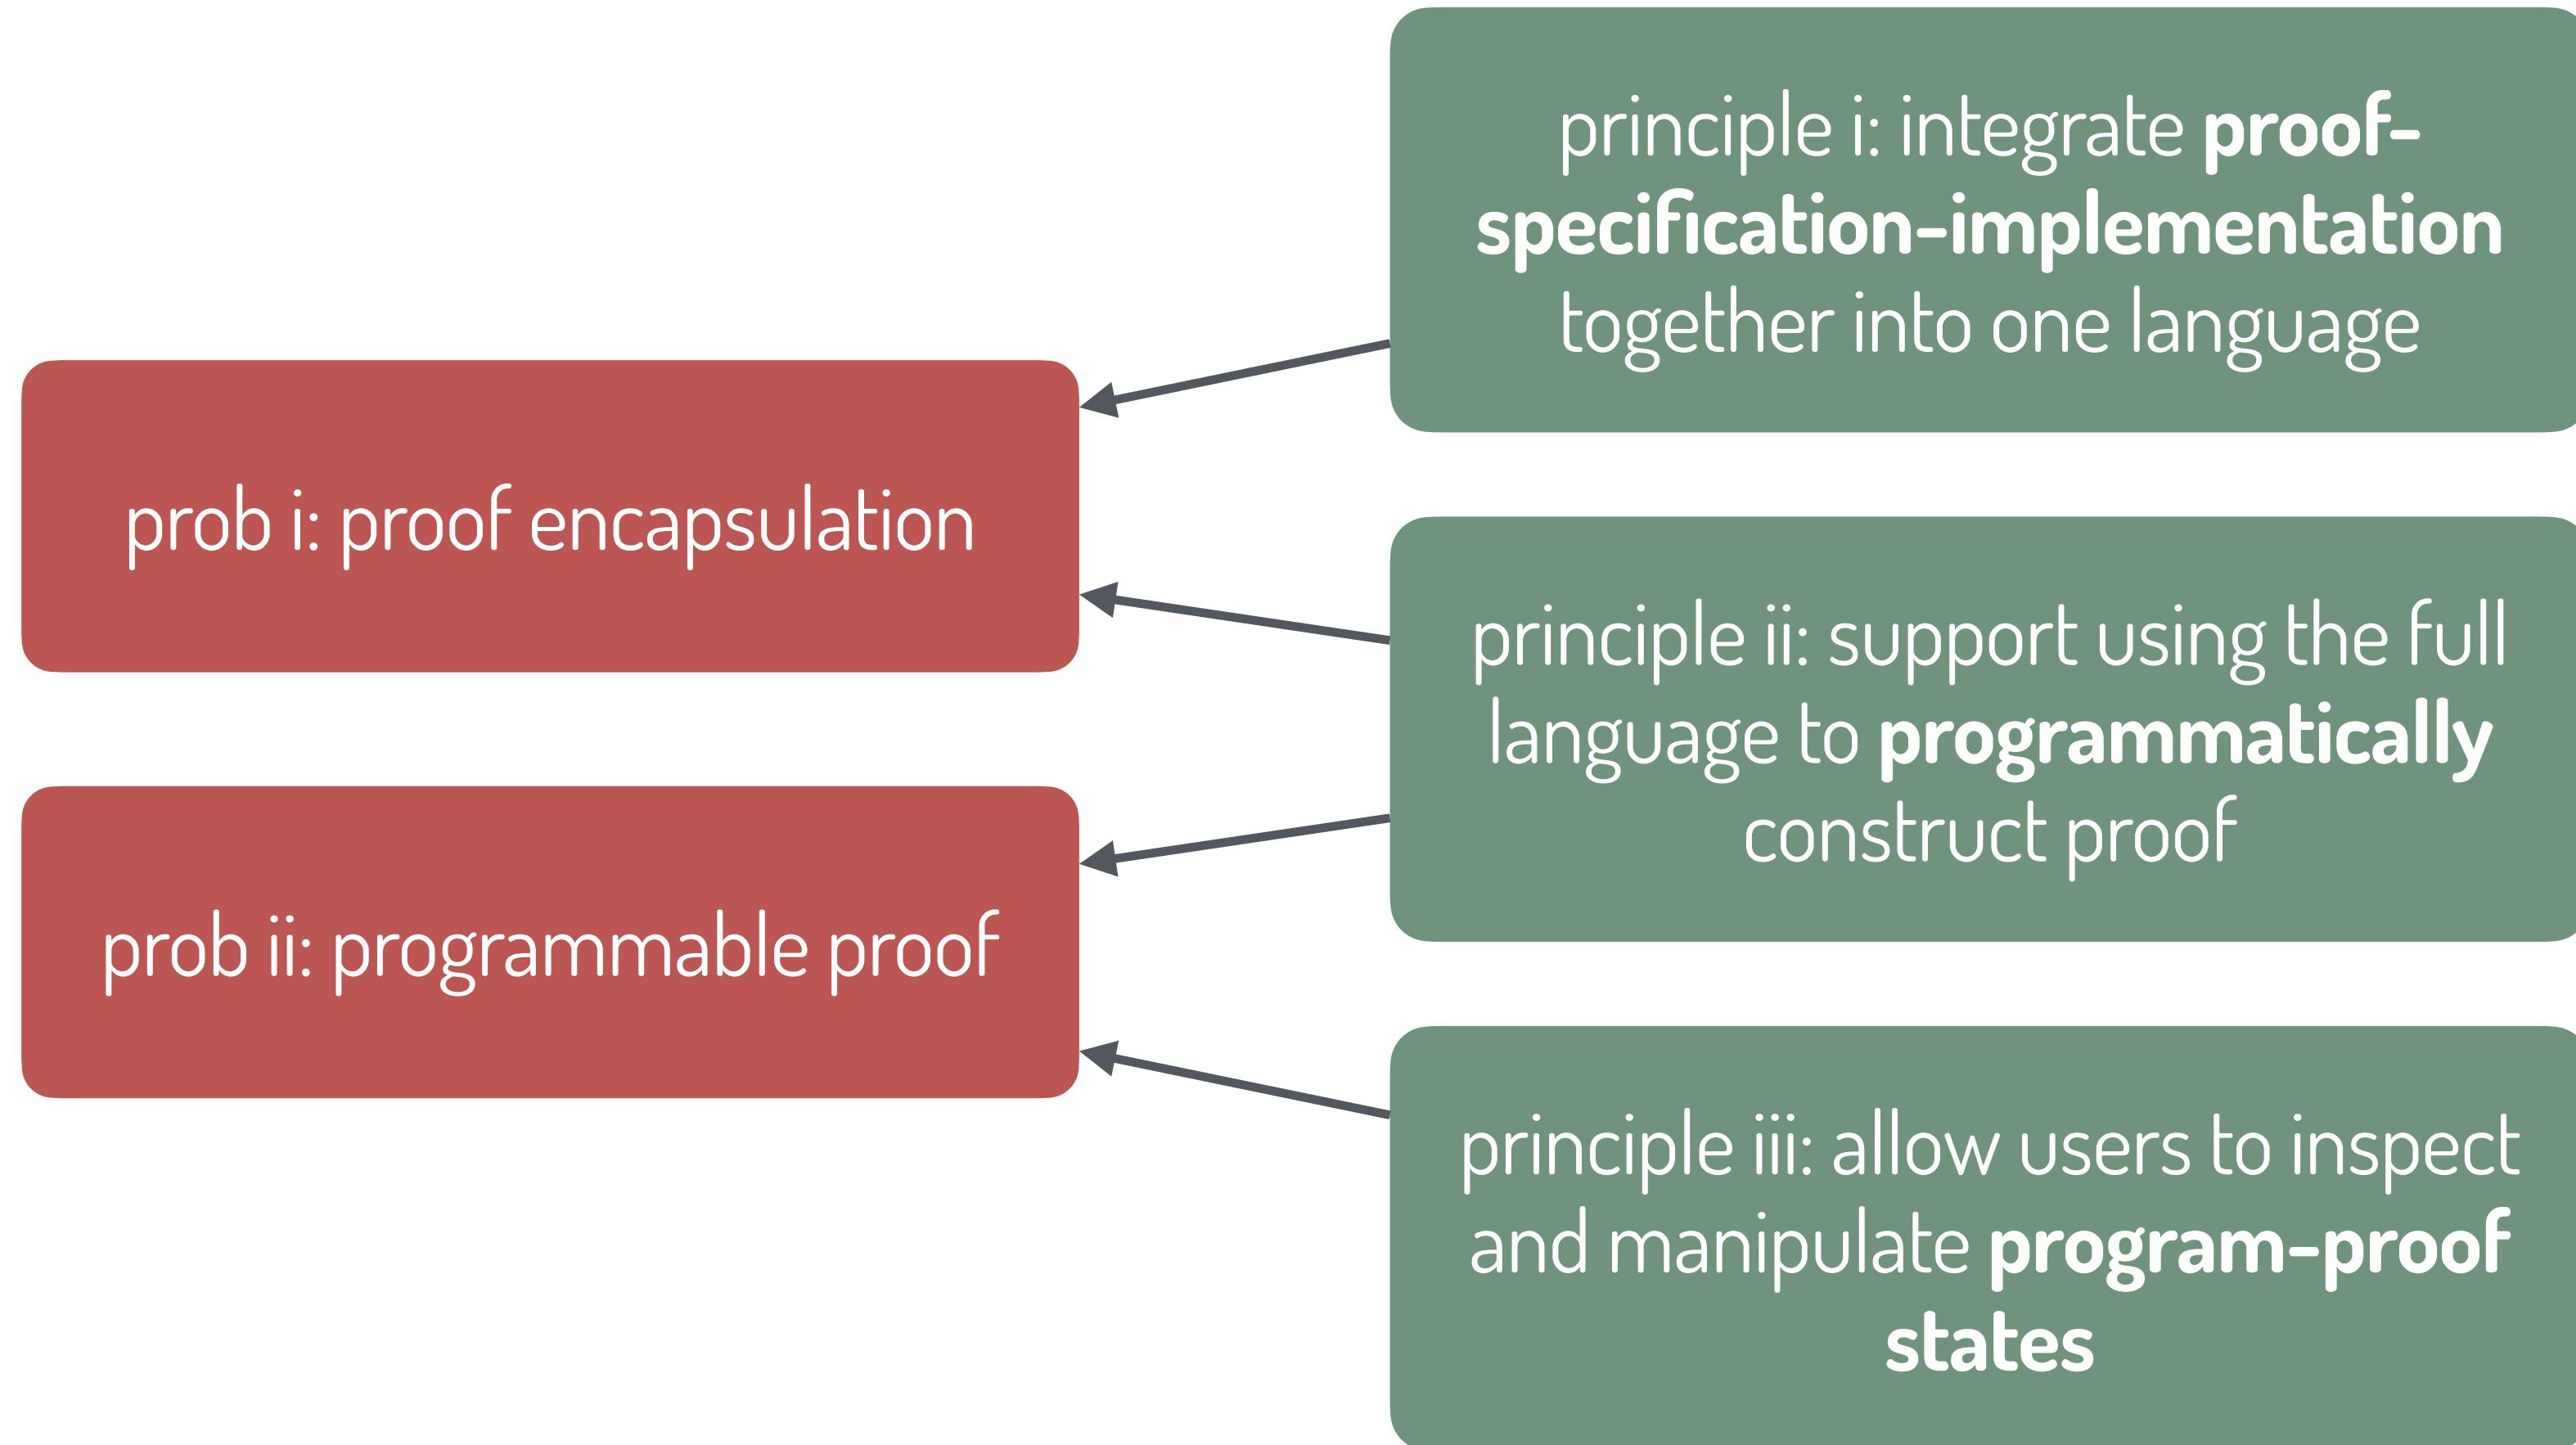
- CStar reuses clang's toolchain
  - use clang for compilation
  - use clangd for language server
- its design is language-agnostic

## Unifying Programming and Verification in C



- CStar reuses clang's toolchain
  - use clang for compilation
  - use clangd for language server
- its design is language-agnostic
  - RustStar?

## Unifying Programming and Verification in C



- **CStar reuses clang's toolchain**
  - use clang for compilation
  - use clangd for language server
- **its design is language-agnostic**
  - RustStar?
  - you own language-Star?



# What are Still Missing

there are two looming "clouds"



# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style



# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
- the logical core is simply-typed lambda calculus, which is in many ways simpler than C



# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
  - the logical core is simply-typed lambda calculus, which is in many ways simpler than C
- **major one: program-proof states are non-local**



# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
  - the logical core is simply-typed lambda calculus, which is in many ways simpler than C
- **major one: program-proof states are non-local**
  - at every program point, there is **a single global state**

# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
- the logical core is simply-typed lambda calculus, which is in many ways simpler than C
- major one: program-proof states are non-local
  - at every program point, there is a single global state

```

for ( ; (order + 1) < pool->max_order; order++)
/*@ inv let p_i2 = (((integer) p) - __hyp_vmemmap) / 32 @*/
/*@ inv let Z = ZeroPage((pointer) ((p_i2 * 4096) - hyp_physvirt_offset), 1, order) @*/
/*@ inv let OP = Owned(pool) @*/
/*@ inv let hyp_vmemmap = (pointer) __hyp_vmemmap @*/
/*@ inv let start_i2 = (*pool).range_start / 4096 @*/
/*@ inv let end_i2 = (*pool).range_end / 4096 @*/
/*@ inv let off_i = hyp_physvirt_offset / 4096 @*/
/*@ inv let V2 = each (integer i; start_i2 <= i && i < end_i2)
           {Owned<struct hyp_page>(hyp_vmemmap+(i*32)) } @*/
/*@ inv let p_page = V2.value[p_i2] @*/
/*@ inv let p_page_tweaked2 = (p_page){.order = order} @*/
/*@ inv each(integer i; start_i2 <= i && i < end_i2)
           {vmmmap_b_wf(i, hyp_vmemmap, V2.value[p_i2 = p_page_tweaked2], pool, *pool)} @*/
/*@ inv each(integer i; 0 <= i && i < ((*pool).max_order)
           {freeArea_cell_wf(i, hyp_vmemmap, V2.value, pool, (*pool))} @*/
/*@ inv hyp_pool_wf(pool, *pool, hyp_vmemmap, hyp_physvirt_offset) @*/
/*@ inv let R = each(integer i; start_i2 <= (i + off_i) && (i + off_i) < end_i2
           & (V2.value[i + off_i]).refcount == 0
           & (V2.value[i + off_i]).order != (hyp_no_order ())
           { ZeroPage((pointer) 0) + (i*4096), 1, (V2.value[i+off_i]).order } @*/
/*@ inv 0 <= order; order+1 <= (*pool).max_order @*/
/*@ inv cellPointer(hyp_vmemmap,32,start_i2,end_i2,p) @*/
/*@ inv (p_page.refcount) == 0; (p_page.order) == (hyp_no_order ()); (p_page.pool) == pool @*/
/*@ inv (p_page.node.next) == &(p->node); (p_page.node.prev) == &(p->node) @*/
/*@ inv order_aligned(p_i2,order) @*/
/*@ inv (p_i2 * 4096) + (page_size_of_order(order)) <= (*pool).range_end @*/
/*@ inv each(integer i; {p_i}@start < i && i < end_i2)
           {{{V.value[i]}@start}.refcount == 0} || ((V2.value[i]) == {V.value[i]}@start) @*/
/*@ inv {__hyp_vmemmap} unchanged; {hyp_physvirt_offset} unchanged; {pool} unchanged @*/
/*@ inv ((*pool)@start){.free_area} = (*pool).free_area == *pool @*/
{
  buddy = __find_buddy_avail(pool, p, order);
  if (!buddy)
    break;
/*CN*/ instantiate vmmmap_b_wf, hyp_page_to_pfn(buddy);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(p), 1, order);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(buddy), 1, order);
/*CN*/ lemma_attach_inc_loop(*pool, p, order);
/*CN*/ lemma2(hyp_page_to_pfn(p), order);
/*CN*/ lemma_page_size_of_order_inc(order);
/*CN*/ if ((buddy->node).next != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).next, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).prev, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != (buddy->node).next);
  list_del_init(&buddy->node);
  buddy->order = HYP_NO_ORDER;
  p = min(p, buddy);
/*CN*/ pack ZeroPage (hyp_page_to_virt(p), 1, order + 1);
}

```

excerpts from the verified attach\_page function in CN

# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
- the logical core is simply-typed lambda calculus, which is in many ways simpler than C
- major one: program-proof states are non-local
  - at every program point, there is a single global state
  - it breaks modularity and readability of reasoning

```

for ( ; (order + 1) < pool->max_order; order++)
/*@ inv let p_i2 = (((integer) p) - __hyp_vmemmap) / 32 @*/
/*@ inv let Z = ZeroPage((pointer) ((p_i2 * 4096) - hyp_physvirt_offset), 1, order) @*/
/*@ inv let OP = Owned(pool) @*/
/*@ inv let hyp_vmemmap = (pointer) __hyp_vmemmap @*/
/*@ inv let start_i2 = (*pool).range_start / 4096 @*/
/*@ inv let end_i2 = (*pool).range_end / 4096 @*/
/*@ inv let off_i = hyp_physvirt_offset / 4096 @*/
/*@ inv let V2 = each (integer i; start_i2 <= i && i < end_i2)
           {Owned<struct hyp_page>(hyp_vmemmap+(i*32)) } @*/
/*@ inv let p_page = V2.value[p_i2] @*/
/*@ inv let p_page_tweaked2 = (p_page){.order = order} @*/
/*@ inv each(integer i; start_i2 <= i && i < end_i2)
           {vmmmap_b_wf(i, hyp_vmemmap, V2.value[p_i2 = p_page_tweaked2], pool, *pool)} @*/
/*@ inv each(integer i; 0 <= i && i < ((*pool).max_order)
           {freeArea_cell_wf(i, hyp_vmemmap, V2.value, pool, (*pool))} @*/
/*@ inv hyp_pool_wf(pool, *pool, hyp_vmemmap, hyp_physvirt_offset) @*/
/*@ inv let R = each(integer i; start_i2 <= (i + off_i) && (i + off_i) < end_i2
           & (V2.value[i + off_i]).refcount == 0
           & (V2.value[i + off_i]).order != (hyp_no_order ())
           { ZeroPage((pointer) 0) + (i*4096), 1, (V2.value[i+off_i]).order } @*/
/*@ inv 0 <= order; order+1 <= (*pool).max_order @*/
/*@ inv cellPointer(hyp_vmemmap,32,start_i2,end_i2,p) @*/
/*@ inv (p_page.refcount) == 0; (p_page.order) == (hyp_no_order ()); (p_page.pool) == pool @*/
/*@ inv (p_page.node.next) == &(p->node); (p_page.node.prev) == &(p->node) @*/
/*@ inv order_aligned(p_i2,order) @*/
/*@ inv (p_i2 * 4096) + (page_size_of_order(order)) <= (*pool).range_end @*/
/*@ inv each(integer i; {p_i}@start < i && i < end_i2)
           {{(V.value[i])@start}.refcount == 0} || ((V2.value[i]) == {V.value[i]}@start) @*/
/*@ inv {__hyp_vmemmap} unchanged; {hyp_physvirt_offset} unchanged; {pool} unchanged @*/
/*@ inv ((*pool)@start){.free_area} = (*pool).free_area == *pool @*/
{
  buddy = __find_buddy_avail(pool, p, order);
  if (!buddy)
    break;
/*CN*/ instantiate vmmmap_b_wf, hyp_page_to_pfn(buddy);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(p), 1, order);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(buddy), 1, order);
/*CN*/ lemma_attach_inc_loop(*pool, p, order);
/*CN*/ lemma2(hyp_page_to_pfn(p), order);
/*CN*/ lemma_page_size_of_order_inc(order);
/*CN*/ if ((buddy->node).next != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).next, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).prev, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != (buddy->node).next);
  list_del_init(&buddy->node);
  buddy->order = HYP_NO_ORDER;
  p = min(p, buddy);
/*CN*/ pack ZeroPage (hyp_page_to_virt(p), 1, order + 1);
}

```

excerpts from the verified attach\_page function in CN

# What are Still Missing

there are two looming "clouds"

- minor one: the logical layer is functional, which is not C-style
- the logical core is simply-typed lambda calculus, which is in many ways simpler than C
- major one: program-proof states are non-local
  - at every program point, there is a single global state
  - it breaks modularity and readability of reasoning
  - ongoing work: program-proof states reified as **local capabilities**

```

for ( ; (order + 1) < pool->max_order; order++)
/*@ inv let p_i2 = (((integer) p) - __hyp_vmemmap) / 32 @*/
/*@ inv let Z = ZeroPage((pointer) ((p_i2 * 4096) - hyp_physvirt_offset), 1, order) @*/
/*@ inv let OP = Owned(pool) @*/
/*@ inv let hyp_vmemmap = (pointer) __hyp_vmemmap @*/
/*@ inv let start_i2 = (*pool).range_start / 4096 @*/
/*@ inv let end_i2 = (*pool).range_end / 4096 @*/
/*@ inv let off_i = hyp_physvirt_offset / 4096 @*/
/*@ inv let V2 = each (integer i; start_i2 <= i && i < end_i2)
           {Owned<struct hyp_page>(hyp_vmemmap+(i*32)) } @*/
/*@ inv let p_page = V2.value[p_i2] @*/
/*@ inv let p_page_tweaked2 = (p_page){.order = order} @*/
/*@ inv each(integer i; start_i2 <= i && i < end_i2)
           {vmmmap_b_wf(i, hyp_vmemmap, V2.value[p_i2 = p_page_tweaked2], pool, *pool)} @*/
/*@ inv each(integer i; 0 <= i && i < (*pool).max_order)
           {freeArea_cell_wf(i, hyp_vmemmap, V2.value, pool, (*pool))} @*/
/*@ inv hyp_pool_wf(pool, *pool, hyp_vmemmap, hyp_physvirt_offset) @*/
/*@ inv let R = each(integer i; start_i2 <= (i + off_i) && (i + off_i) < end_i2
           & (V2.value[i + off_i]).refcount == 0
           & (V2.value[i + off_i]).order != (hyp_no_order ())
           { ZeroPage((pointer) 0) + (i*4096), 1, (V2.value[i+off_i]).order } @*/
/*@ inv 0 <= order; order+1 <= (*pool).max_order @*/
/*@ inv cellPointer(hyp_vmemmap,32,start_i2,end_i2,p) @*/
/*@ inv (p_page.refcount) == 0; (p_page.order) == (hyp_no_order ()); (p_page.pool) == pool @*/
/*@ inv (p_page.node.next) == &(p->node); (p_page.node.prev) == &(p->node) @*/
/*@ inv order_aligned(p_i2,order) @*/
/*@ inv (p_i2 * 4096) + (page_size_of_order(order)) <= (*pool).range_end @*/
/*@ inv each(integer i; {p_i}@start < i && i < end_i2)
           {{(V.value[i])@start}.refcount == 0} || ((V2.value[i]) == {V.value[i]}@start)} @*/
/*@ inv {__hyp_vmemmap} unchanged; {hyp_physvirt_offset} unchanged; {pool} unchanged @*/
/*@ inv ((*pool)@start){.free_area} = (*pool).free_area == *pool @*/
{
  buddy = __find_buddy_avail(pool, p, order);
  if (!buddy)
    break;
/*CN*/ instantiate vmmmap_b_wf, hyp_page_to_pfn(buddy);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(p), 1, order);
/*CN*/ unpack ZeroPage (hyp_page_to_virt(buddy), 1, order);
/*CN*/ lemma_attach_inc_loop(*pool, p, order);
/*CN*/ lemma2(hyp_page_to_pfn(p), order);
/*CN*/ lemma_page_size_of_order_inc(order);
/*CN*/ if ((buddy->node).next != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).next, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != &pool->free_area[order])
/*CN*/   instantiate vmmmap_b_wf,
           hyp_page_to_pfn(container_of((buddy->node).prev, struct hyp_page, node));
/*CN*/ if ((buddy->node).prev != (buddy->node).next);
  list_del_init(&buddy->node);
  buddy->order = HYP_NO_ORDER;
  p = min(p, buddy);
/*CN*/ pack ZeroPage (hyp_page_to_virt(p), 1, order + 1);
}

```

excerpts from the verified attach\_page function in CN

efficiency

