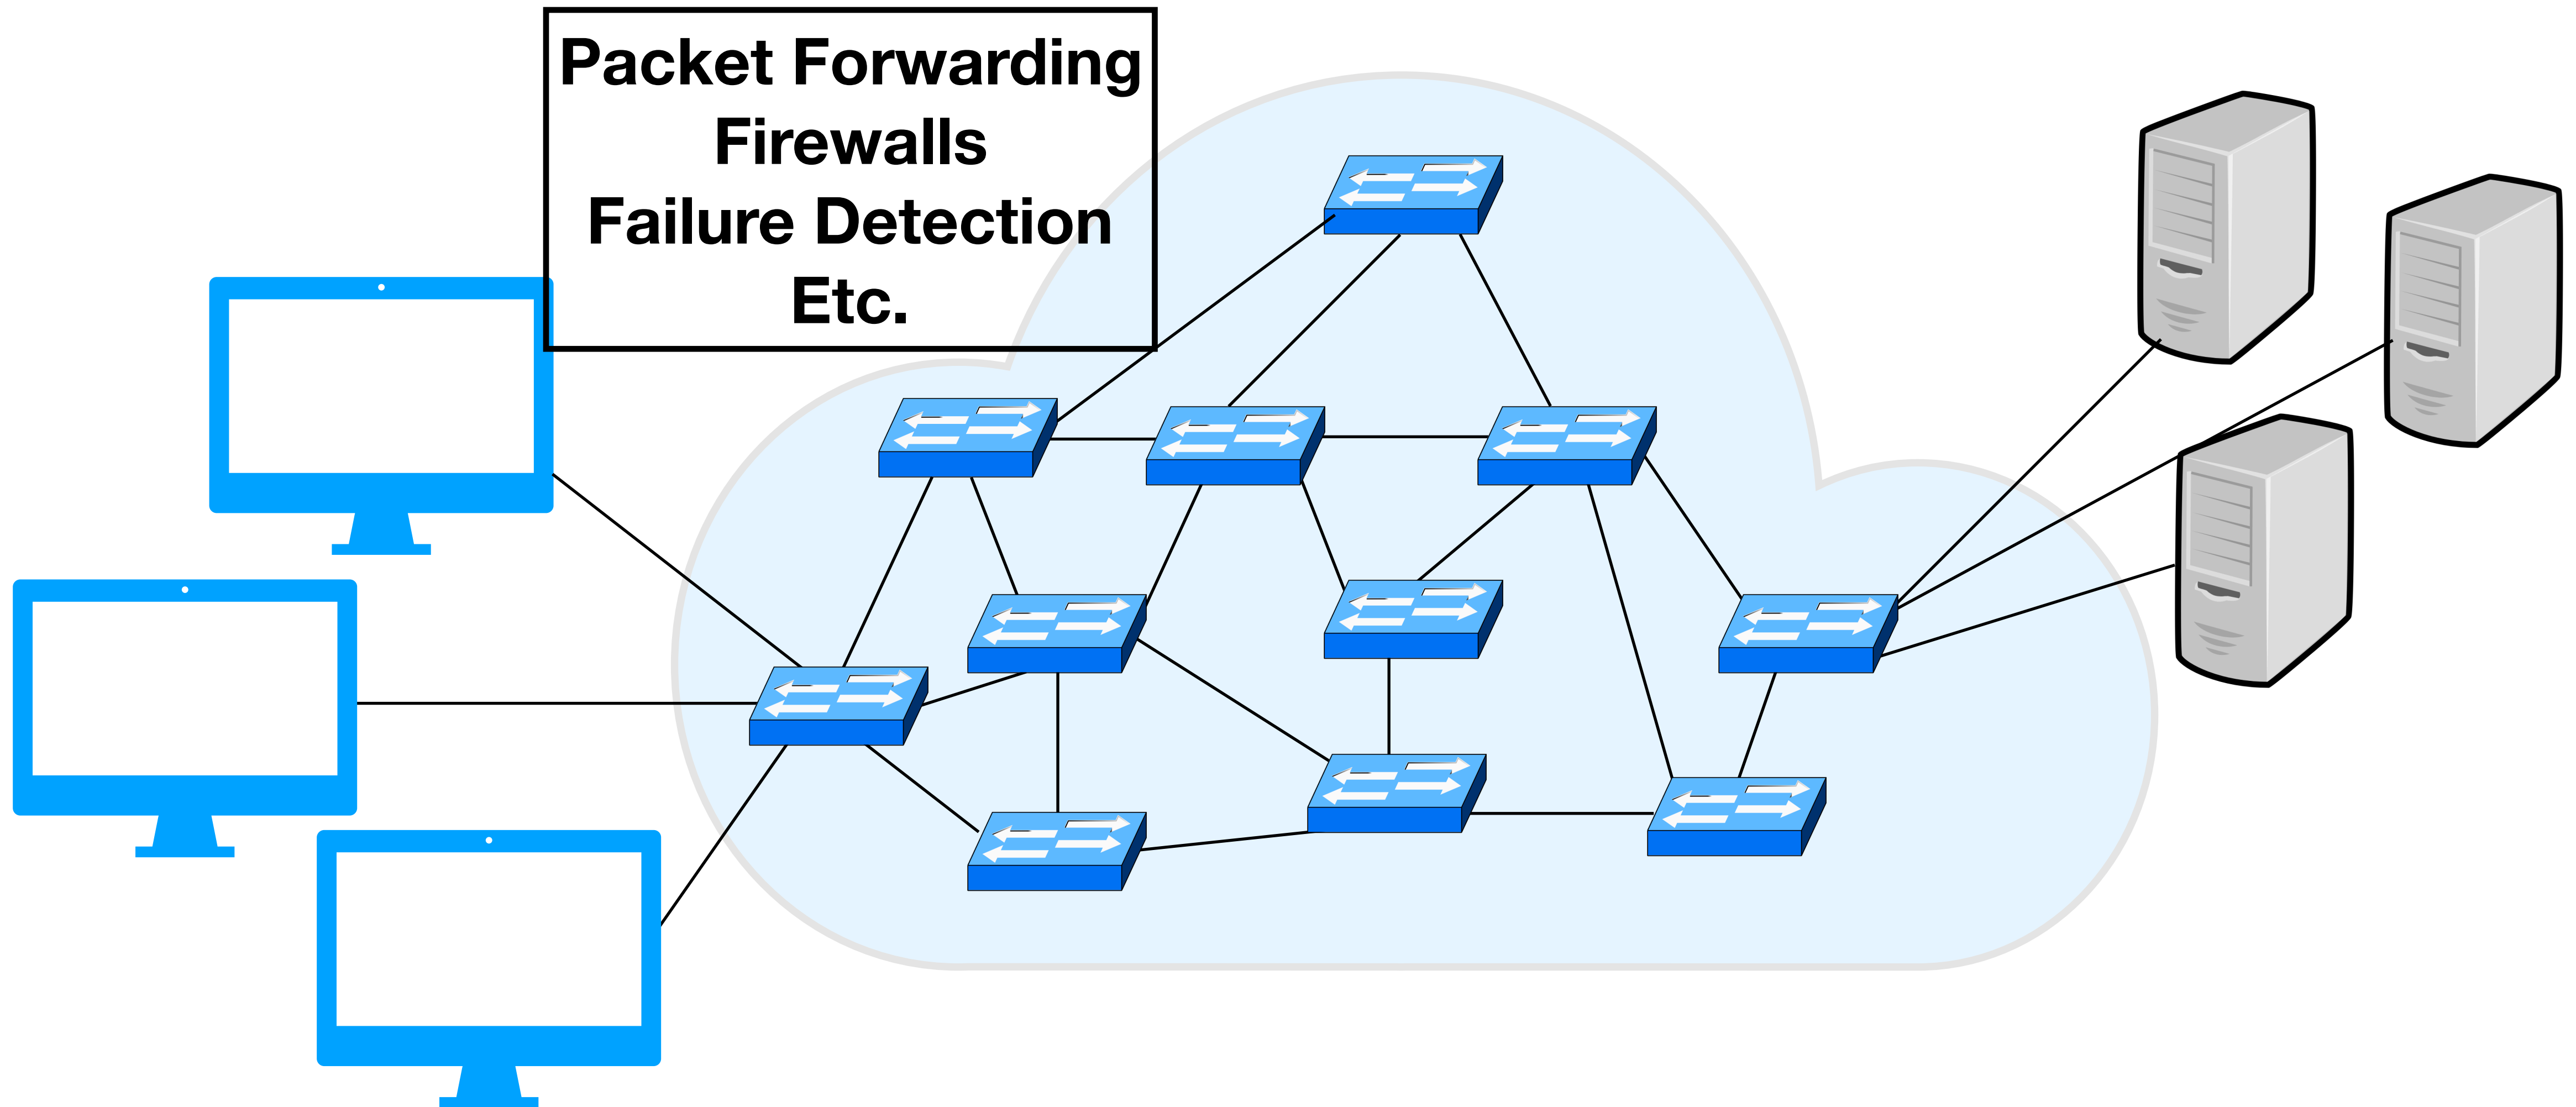


# Language Expressiveness Under Extreme Scarcity in Programmable Data Planes

Mary Hogan



# Switches must process billions of packets per second



# We need specialized hardware to process at high speeds



**Programmable switches**

- ✓ **Support network programs (firewalls, packet forwarding, etc.)**
- ✓ **Fast processing**

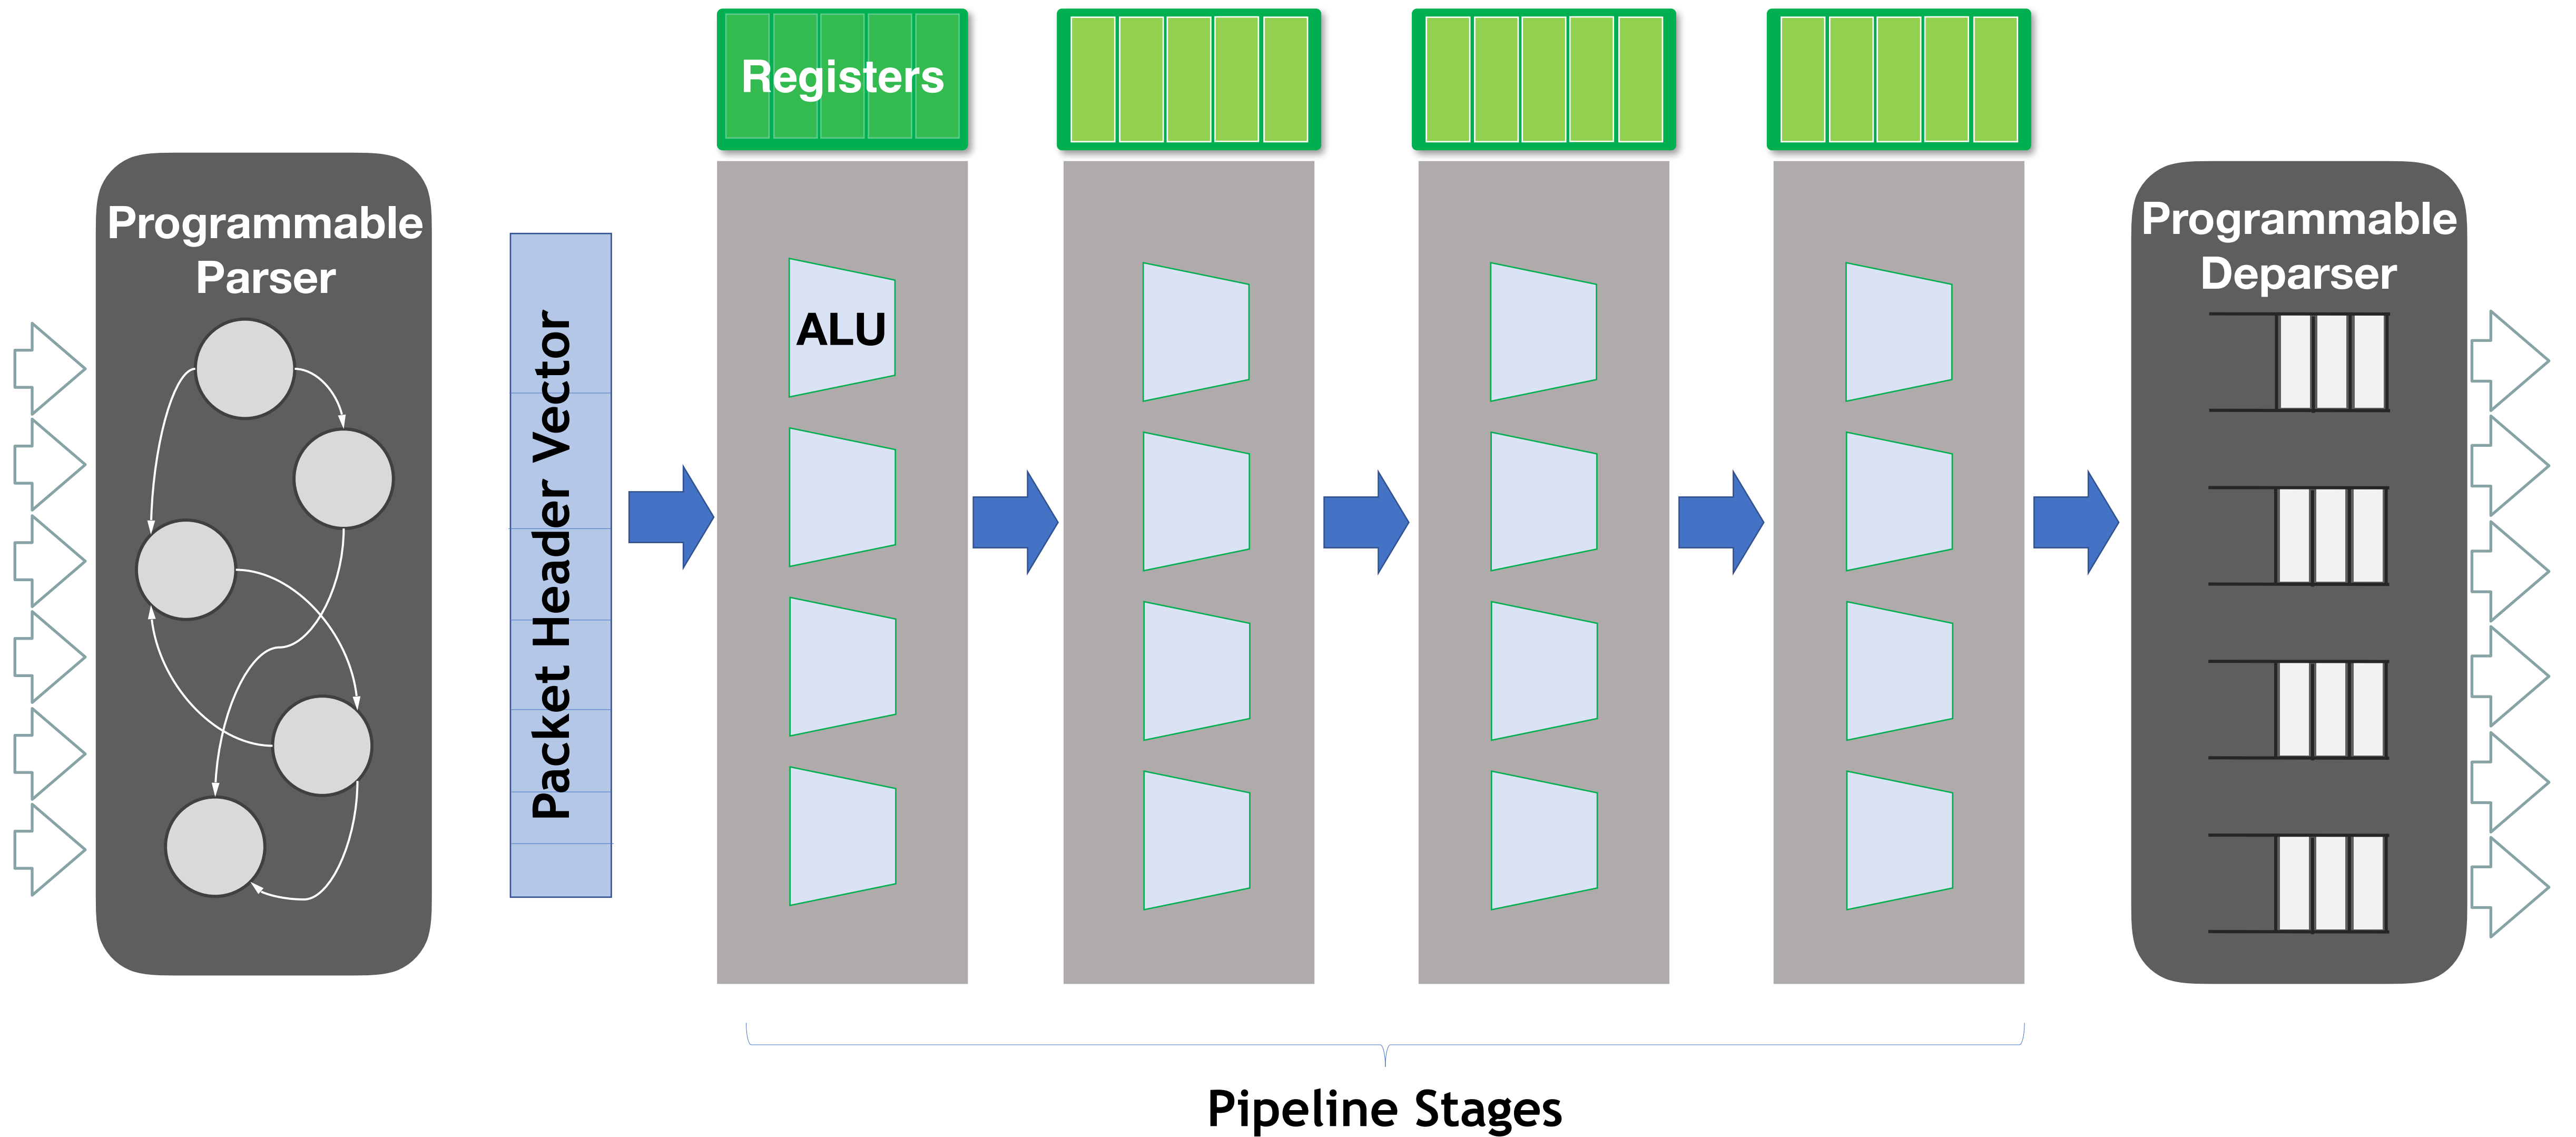
**No free lunch...**

**(1) Processing restrictions**

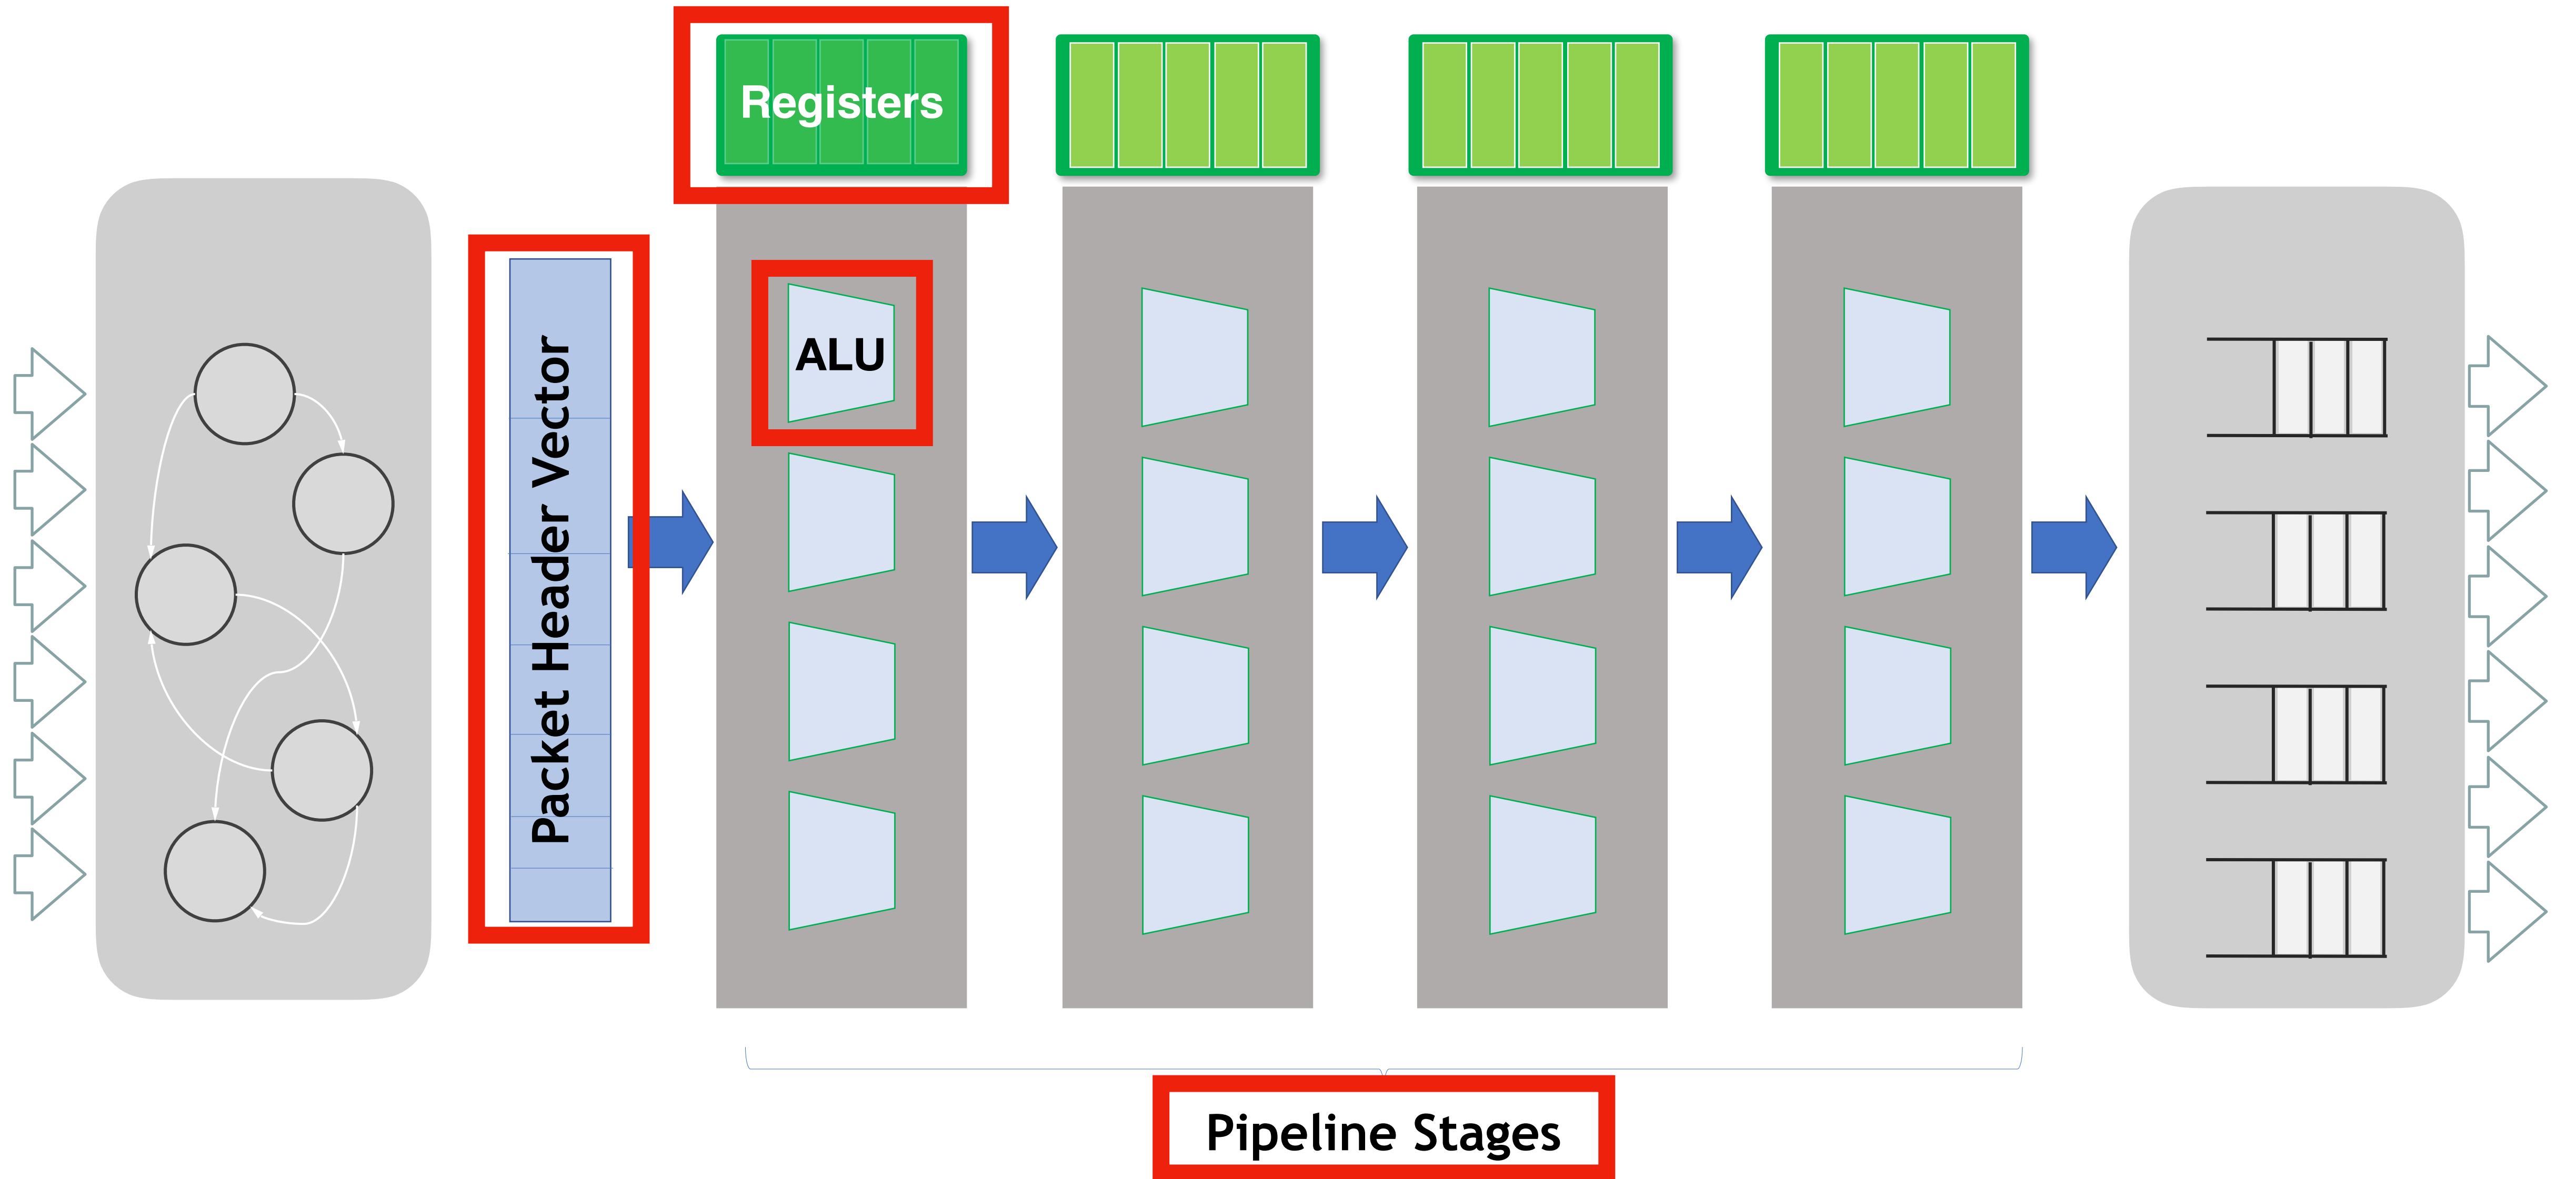
**(2) Resource restrictions**

# Protocol-Independent Switch Architecture

# PISA




# PISA



# Language Tightly Coupled with Hardware

The gap between expressiveness  
and scarcity forces programmers to  
make decisions about low-level  
details.



W  
exp  
prog  
high

have to  
y scarce  
hand

# Network Cache

Key	Value
1	A
2	B
3	C
4	D

**Key-Value Store for  
popular content**

**What if the content  
popularity changes  
over time?**


NetCache, *Jin et al.* [SOSP'17]



# Count-Min Sketch

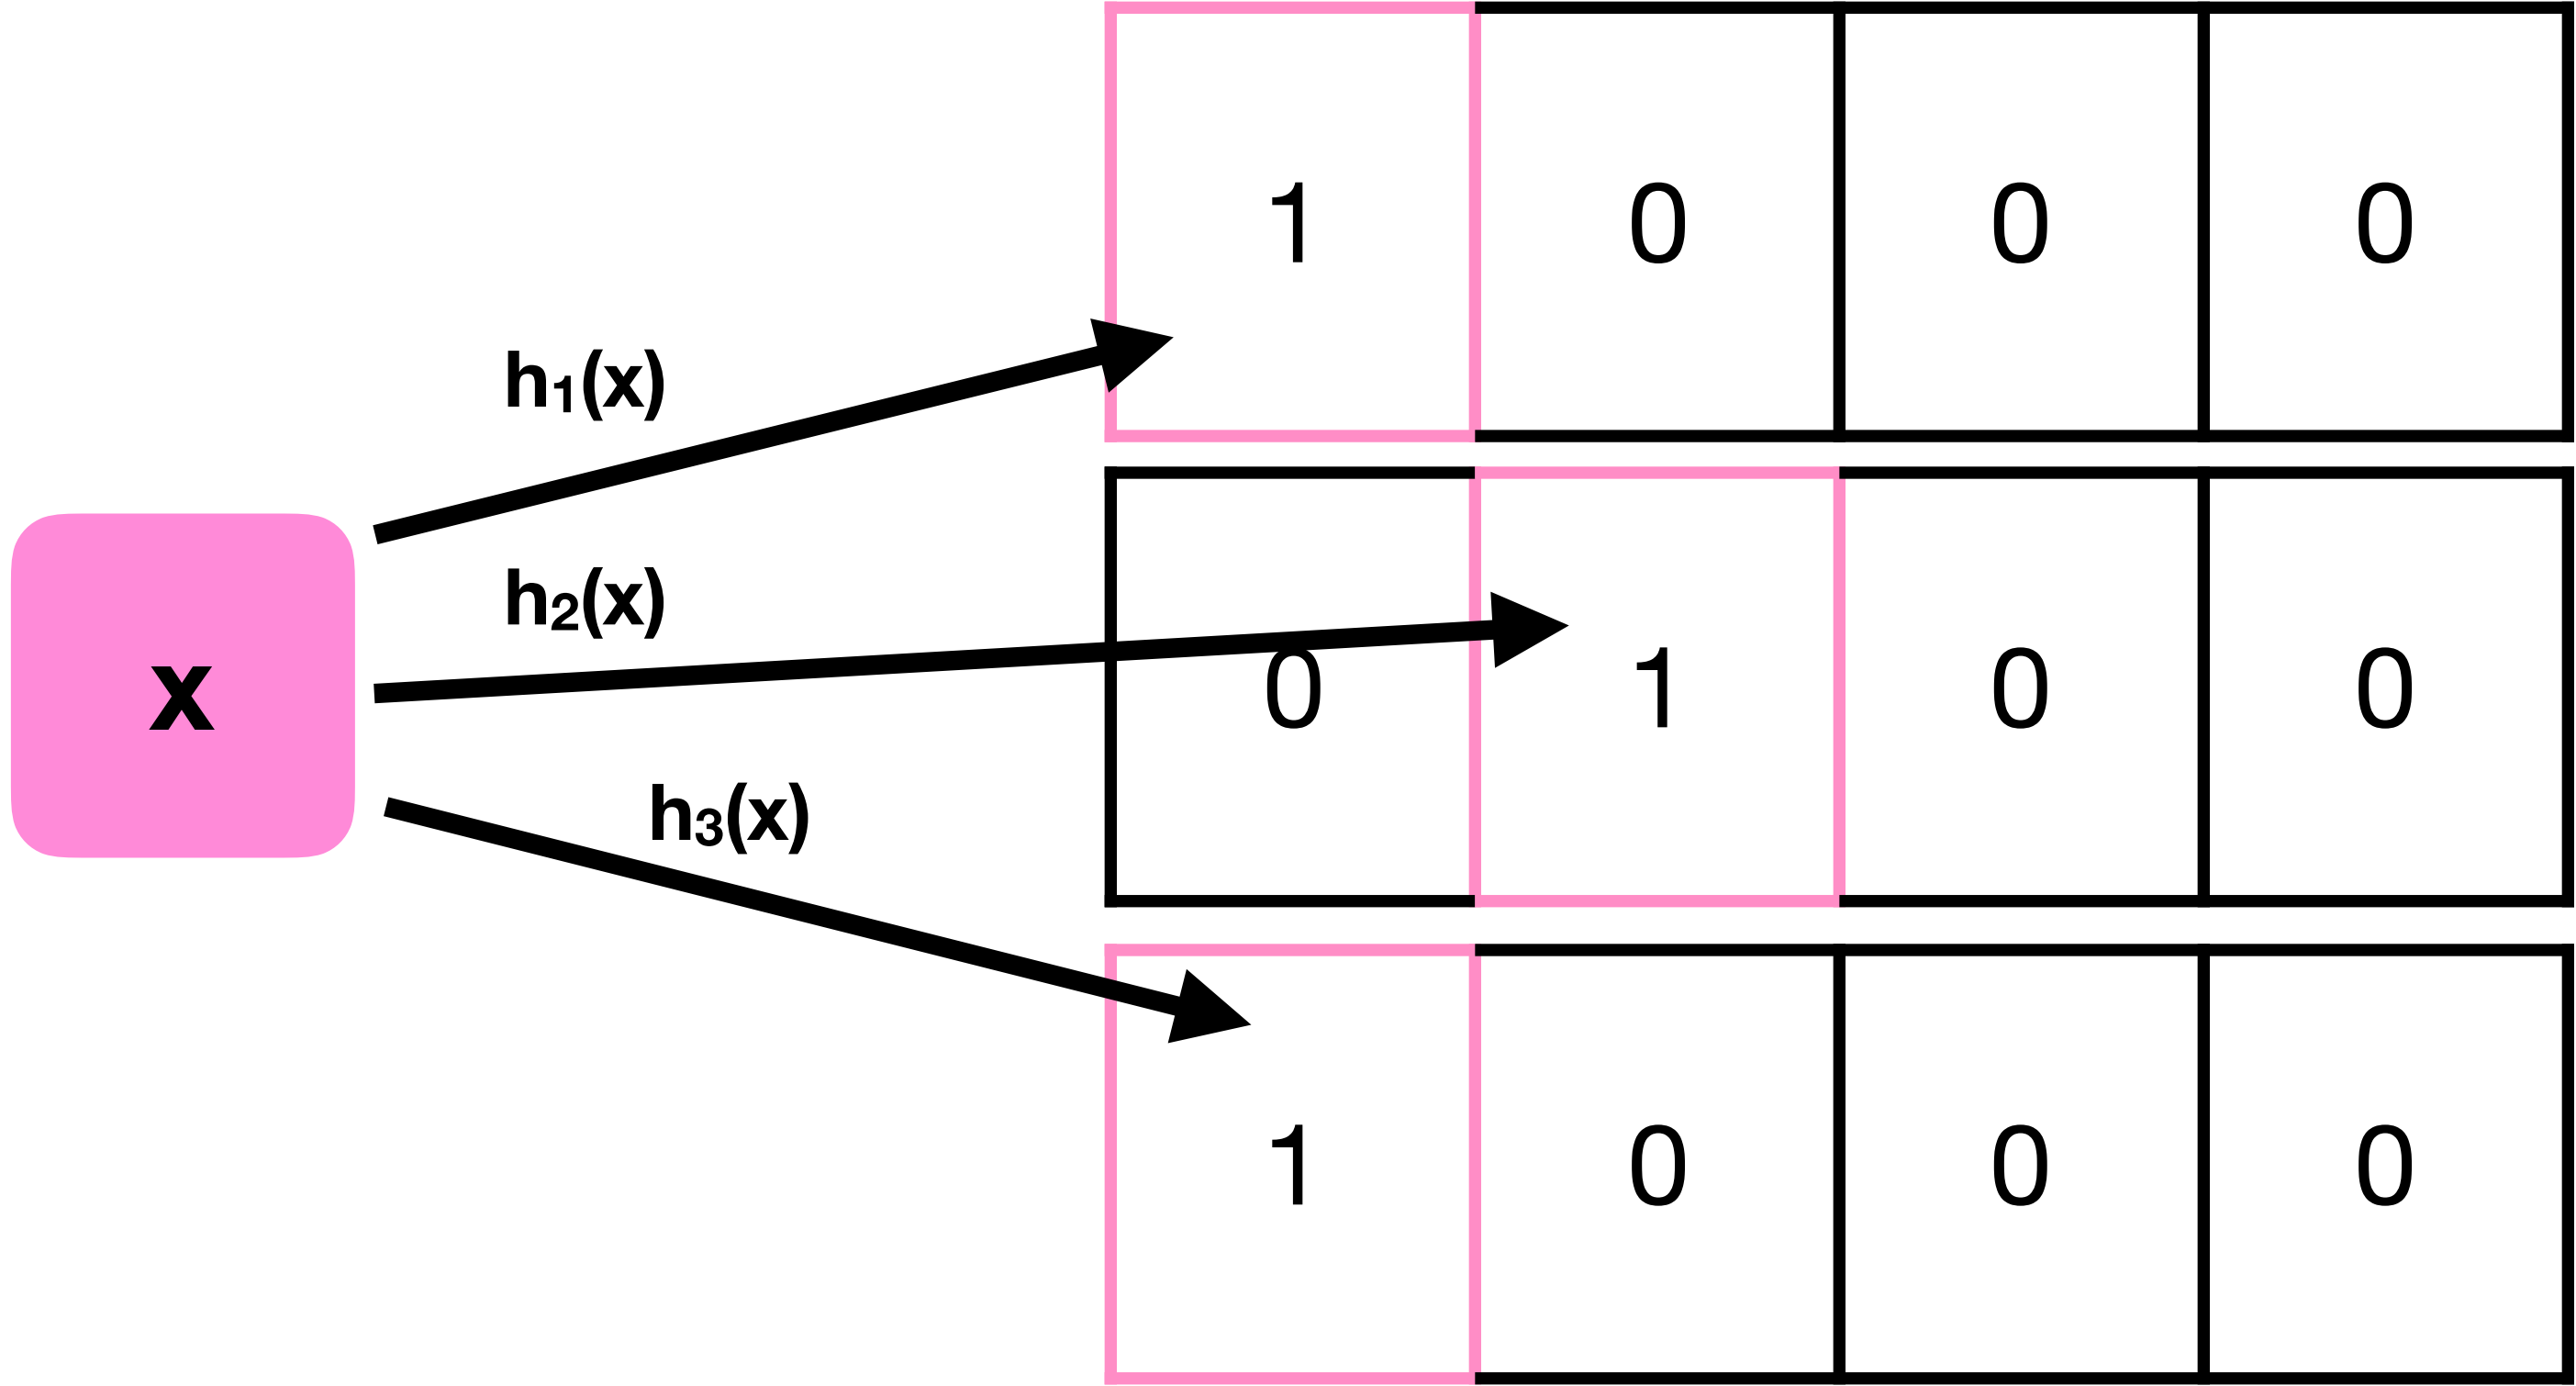
0	0	0	0
0	0	0	0
0	0	0	0

# Count-Min Sketch




0	0	0	0
0	0	0	0
0	0	0	0


# Count-Min Sketch



# Count-Min Sketch



# Count-Min Sketch



Count(x) = 1

2	0	0	0
0	1	1	0
1	0	0	1

# Building a Cache


Key	Value
1	A
2	B
3	C
4	D

**Key-Value Store for popular content**


Estimated # of Requests			
0	0	0	0
0	0	0	0
0	0	0	0

**Popularity Tracker (CMS)**

# Allocating Resources to the Cache



# Allocating Resources to the Cache



**There is a fundamental tradeoff between accuracy and resource allocation.**



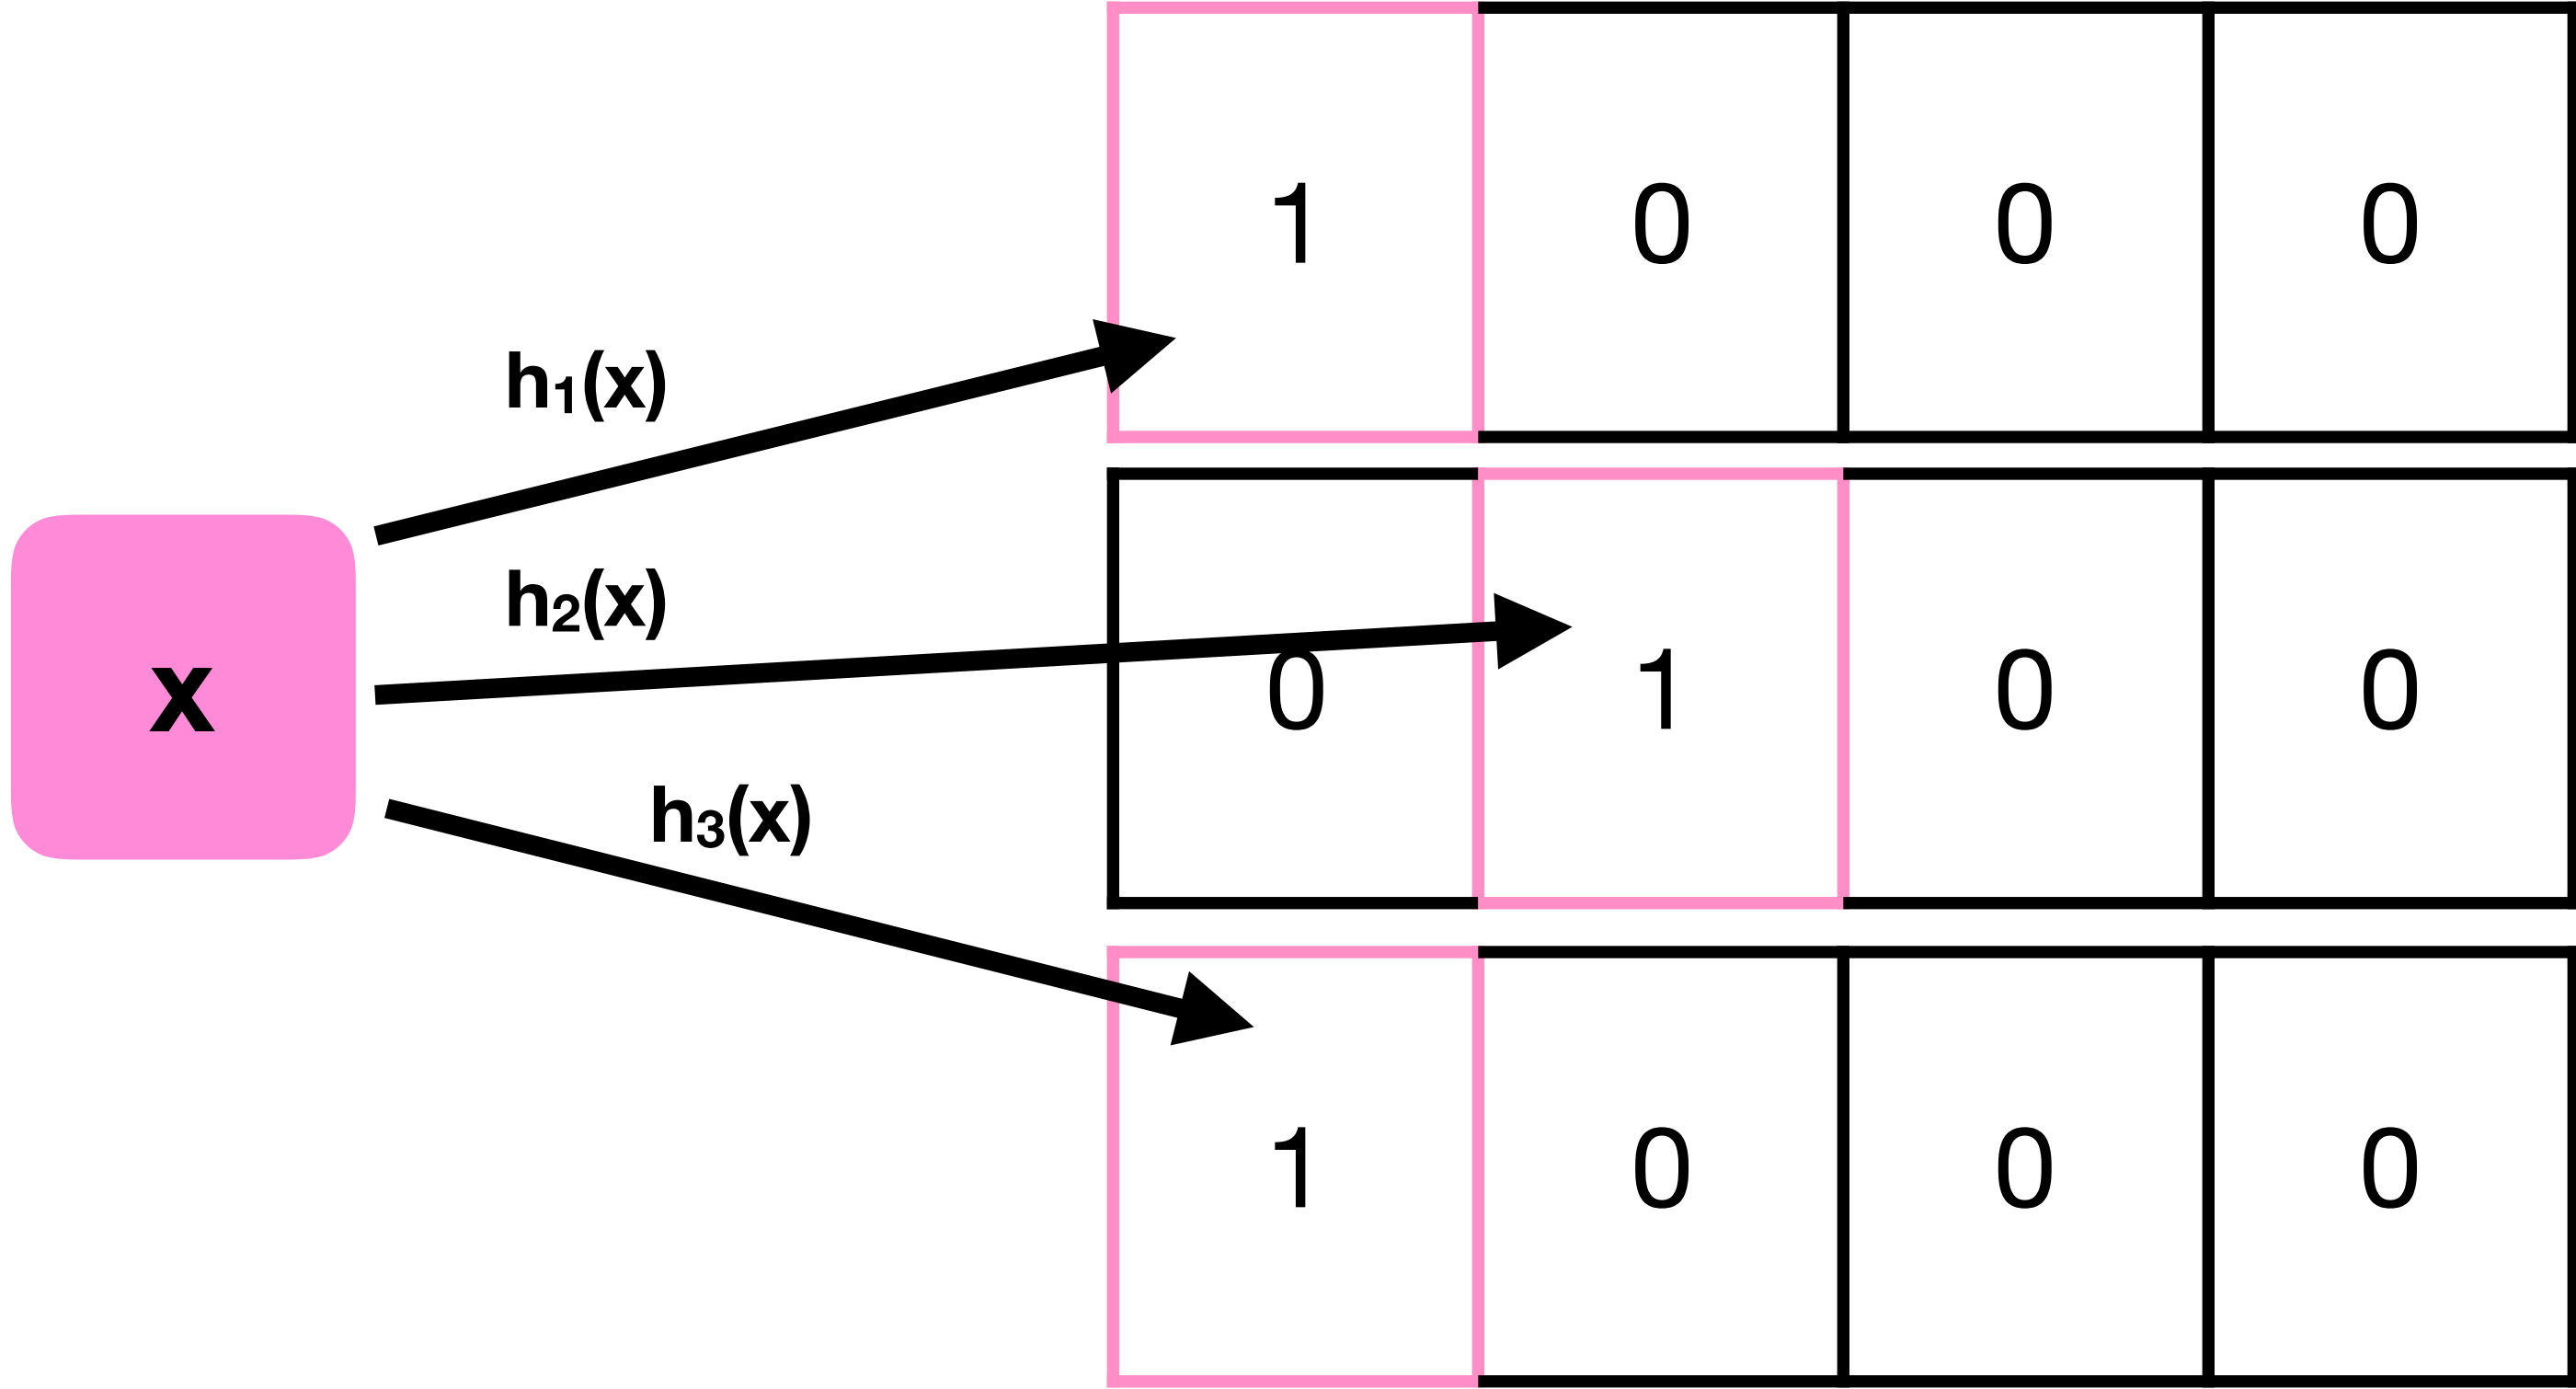
# Cache in P4

```
kv_size = 100;  
cms_cols = 50;
```

# Cache in P4

```
kv_size = 100;  
cms_cols = 50;  
register<bit<32>>(cms_cols) cms_row0;  
register<bit<32>>(cms_cols) cms_row1;  
register<bit<32>>(cms_cols) cms_row2;
```


# Count-Min Sketch



# Cache in P4

```
kv_size = 100;  
cms_cols = 50;  
register<bit<32>>(cms_cols) cms_row0;  
register<bit<32>>(cms_cols) cms_row1;  
register<bit<32>>(cms_cols) cms_row2;
```

```
def cms_hash0 () { ... }  
def cms_hash1 () { ... }  
def cms_hash2 () { ... }
```



**Low-level details  
are intertwined with  
the programming  
language.**

# The System

**Practical system** that decouples the programming language from the hardware to simplify programming network devices

**P4AI**  
[HotNets'20]

**Parasol**  
[In Submission]

Language abstraction for expressiveness


Separate low-level details from program code

**P4AI**  
[NSDI'22]

**Parasol**  
[In Submission]

Optimization with scarce resources

Automatically tailor program to environment



Compile to domain-specific code (P4)

P4 code that fits switch architecture



Programmable switches

Compile to hardware binary

Ready to run in the network!

# P4All

P4All  
[HotNets'20]

Language abstraction  
for expressiveness

Separate low-level details  
from program code

P4All  
[NSDI'22]

Optimization with  
scarce resources

Automatically tailor  
program to environment

Reusable **elastic structures** that can stretch or  
shrink to fill available resources

# Cache in P4

```
kv_size = 100;  
cms_cols = 50;  
register<bit<32>>(cms_cols) cms_row0;  
register<bit<32>>(cms_cols) cms_row1;  
register<bit<32>>(cms_cols) cms_row2;
```



**Low-level details  
are intertwined with  
the programming  
language.**


```
def cms_hash0 () { ... }  
def cms_hash1 () { ... }  
def cms_hash2 () { ... }
```

# Elastic Parameters

```
symbolic int kv_size;  
kv_size = 100;  
symbolic int cms_rows;  
cms_cols = 50;  
symbolic int cms_cols;
```


Symbolic values are placeholders.

keys = 100



KV Store

rows = 3



CMS



# Elastic Structure

```
kv_size = 100;  
cms_cols = 50;  
register<bit<32>>(cms_cols) cms_row0;  
register<bit<32>>(cms_cols) cms_row1;  
register<bit<32>>(cms_cols) cms_row2;
```

```
symbolic int kv_size;  
symbolic int cms_rows;  
symbolic int cms_cols;  
register<bit<32>>(cms_cols)[cms_rows] cms;
```

# Elastic Operations

```
kv_size = 100;
cms_cols = 50;
register<bit<32>>(cms_cols) cms_row0;
register<bit<32>>(cms_cols) cms_row1;
register<bit<32>>(cms_cols) cms_row2;
```

```
def cms_hash0 () { ... }
def cms_hash1 () { ... }
def cms_hash2 () { ... }
```

```
symbolic int kv_size;
symbolic int cms_rows;
symbolic int cms_cols;
register<bit<32>>(cms_cols)[cms_rows] cms;
```

```
def cms_hash () {
    for (i < cms_rows) {
        hash_and_increment(cms[i]);
    }
}
```

# The System

**P4AI**  
[HotNets'20]  
**Parasol**  
[In Submission]

Language abstraction  
for expressiveness

Elastic structures defined  
with symbolic values,  
arrays, bounded loops

**P4AI**  
[NSDI'22]  
**Parasol**  
[In Submission]

Optimization with  
scarce resources

Automatically tailor  
program to environment



Compile to domain-  
specific code (P4)

P4 code that fits switch  
architecture




Programmable  
switches

Compile to hardware  
binary

Ready to run in  
the network!

**P4All  
Program**






**Concrete values  
for symbolic values  
(P4 Program)**

**P4All  
Program**

**Target Specification  
(resource constraints, etc.)**



+



**Concrete values  
for symbolic values  
(P4 Program)**



Values that  
yield  
**optimal**  
performance

**P4All  
Program**

**Target Specification  
(resource constraints, etc.)**



+



**P4All Optimizer**

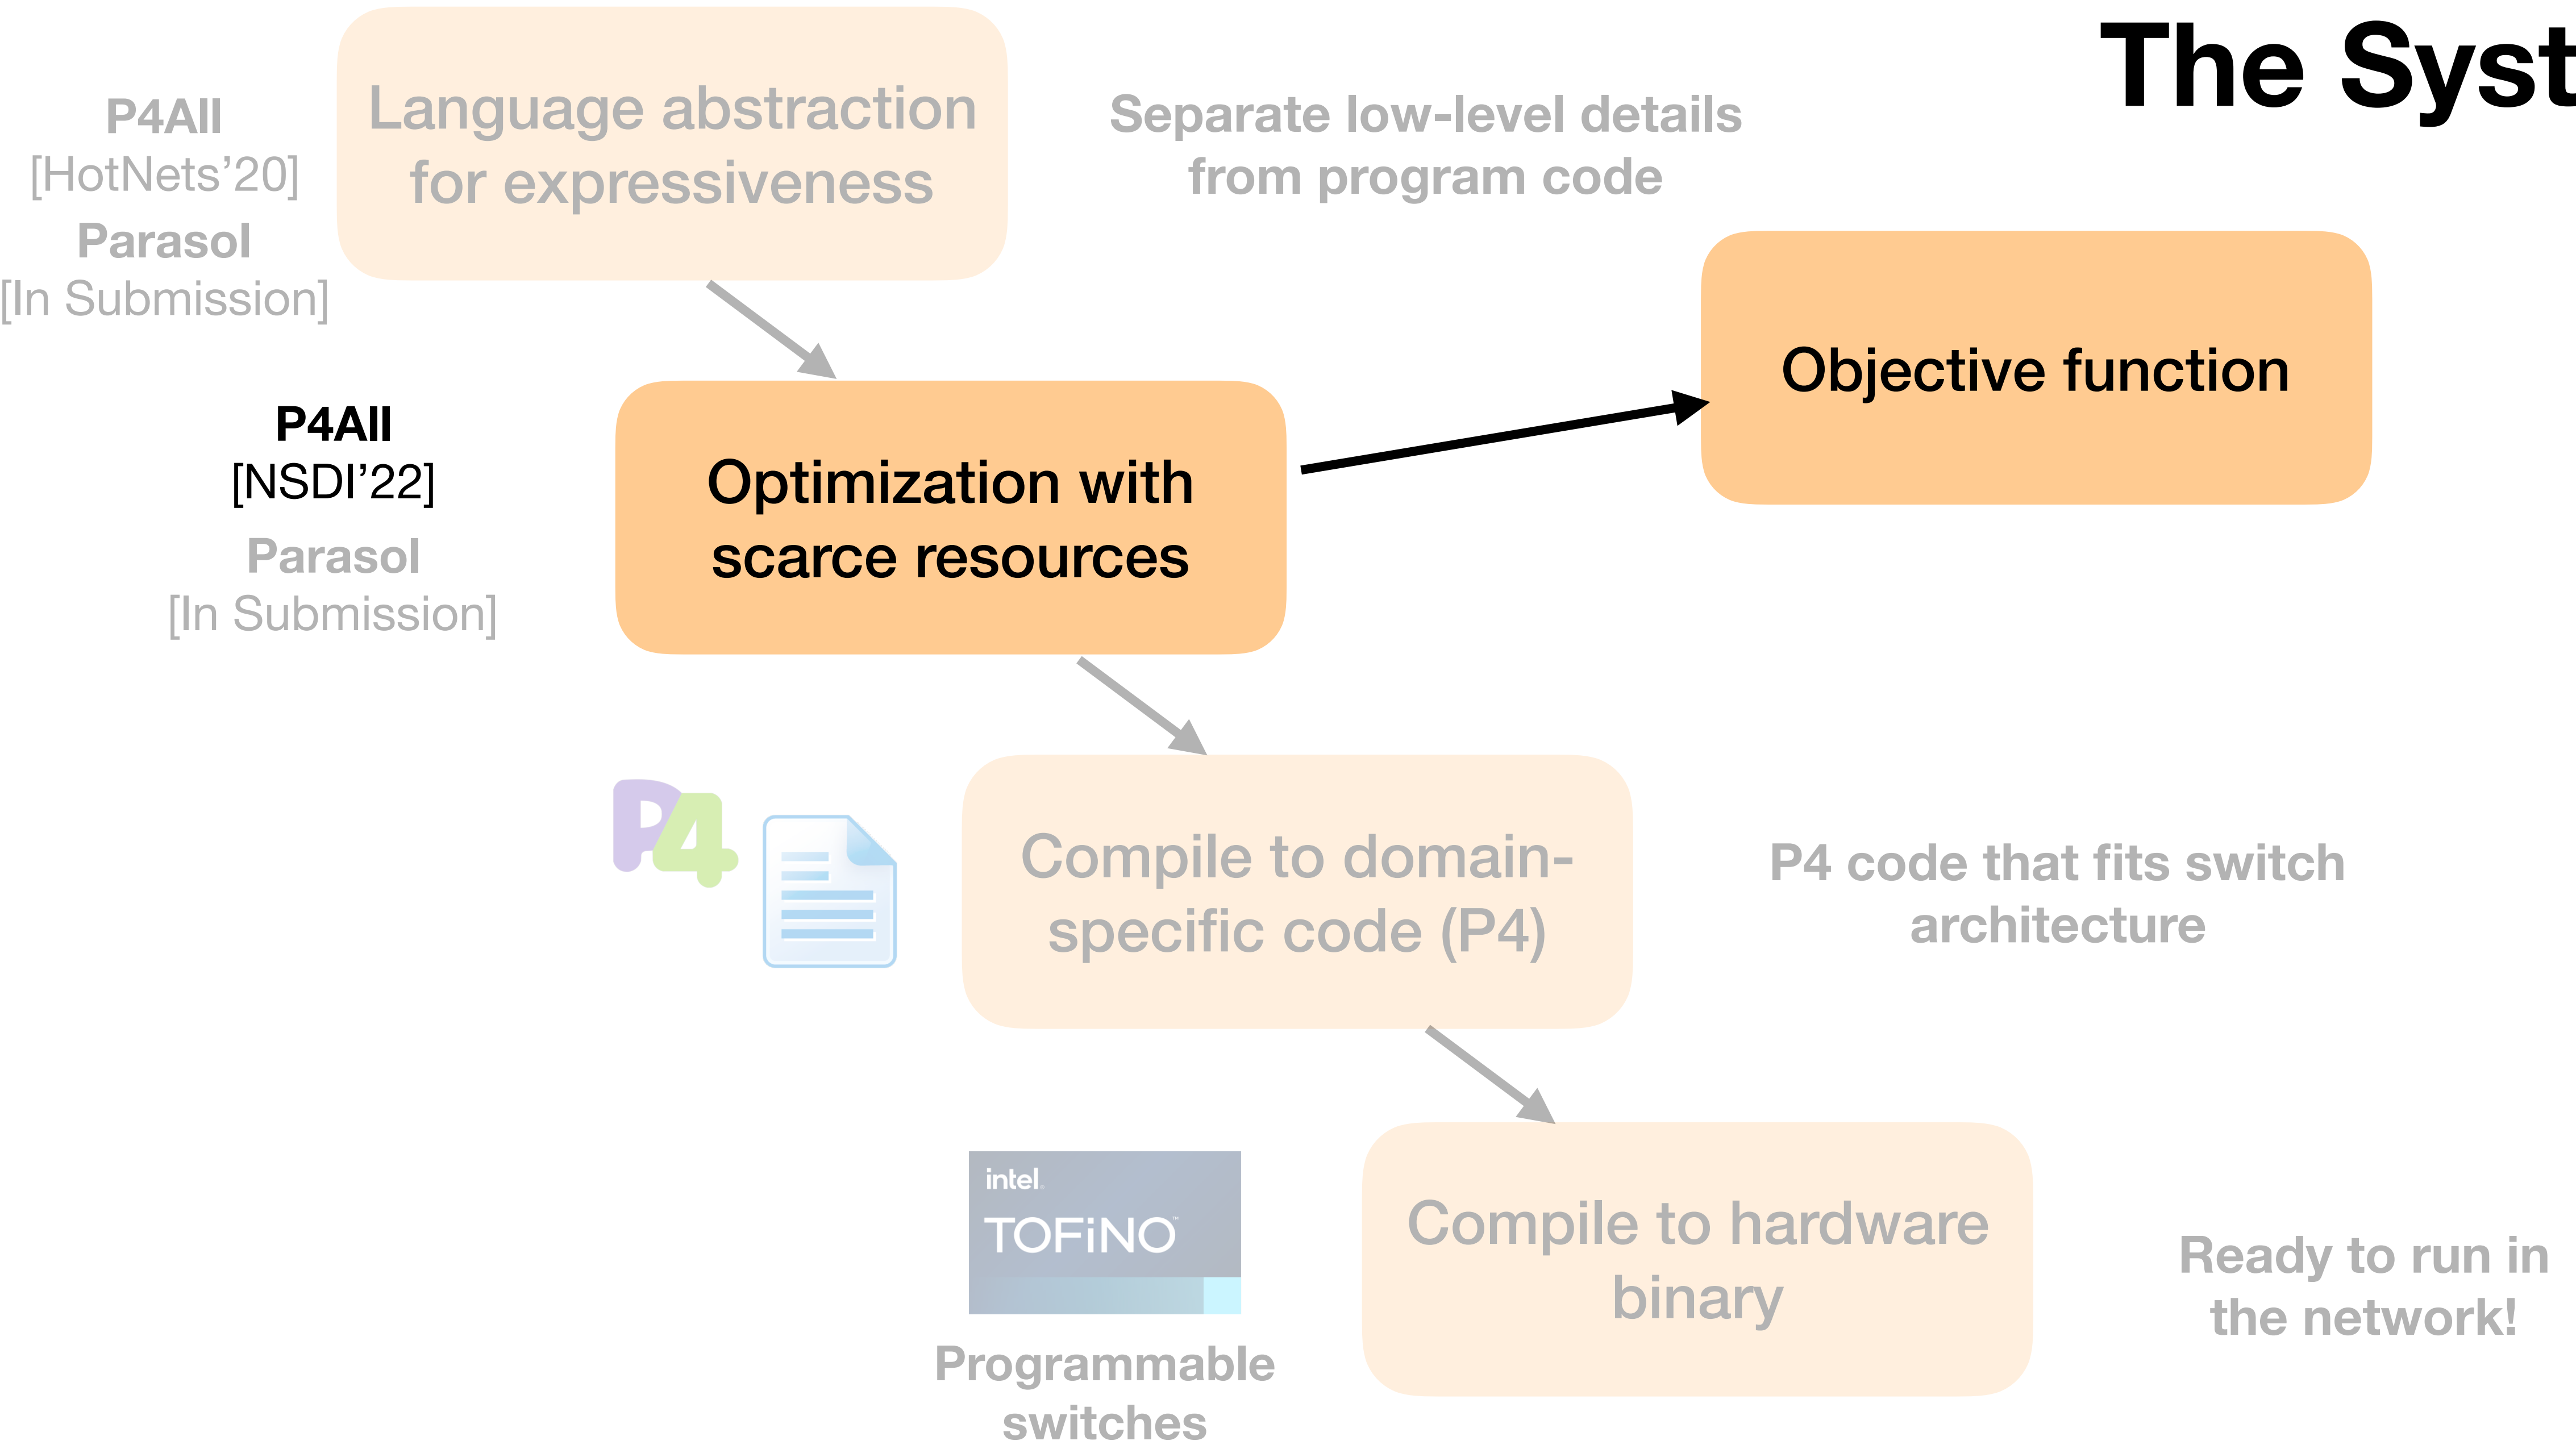
**Generate and Solve Integer-  
Linear Program (ILP)**



**Concrete values  
for symbolic values  
(P4 Program)**

Values that  
yield  
**optimal**  
performance

# The System



# Cache Performance

Key	Value
1	A
2	B
3	C
4	D

**Key-Value Store for popular content**

$$\text{Hit rate} = \sum_{i=1}^{kv\_size} \frac{1}{i^\alpha}$$

Workload-dependent parameter



# Cache Performance

Key	Value
1	A
2	B
3	C
4	D

**Key-Value Store for popular content**

Estimated # of Requests			
0	0	0	0
0	0	0	0
0	0	0	0

**Popularity Tracker (CMS)**

# Cache Performance

$$\boxed{\text{cms\_cols}} = 3 (1/\epsilon)^{1/\alpha}$$

$$\epsilon = \frac{\text{CMS error}}{3} = \frac{3}{\text{cms\_cols}}$$

Estimated # of Requests

0	0	0	0
0	0	0	0
0	0	0	0

Popularity Tracker  
(CMS)

# Cache Performance

$$\text{Hit rate} = \sum_{i=1}^{kv\_size} \frac{1}{i} - \epsilon = \frac{3}{cms\_cols}$$

# Optimal Allocation = Optimal Performance

**ILP:**            maximize    **key value hit rate - CMS error**  
                      subject to    **resource constraints**

$$\left( \sum_{i=1}^{kv\_size} \frac{1}{i} \right) - \frac{3}{cms\_cols}$$

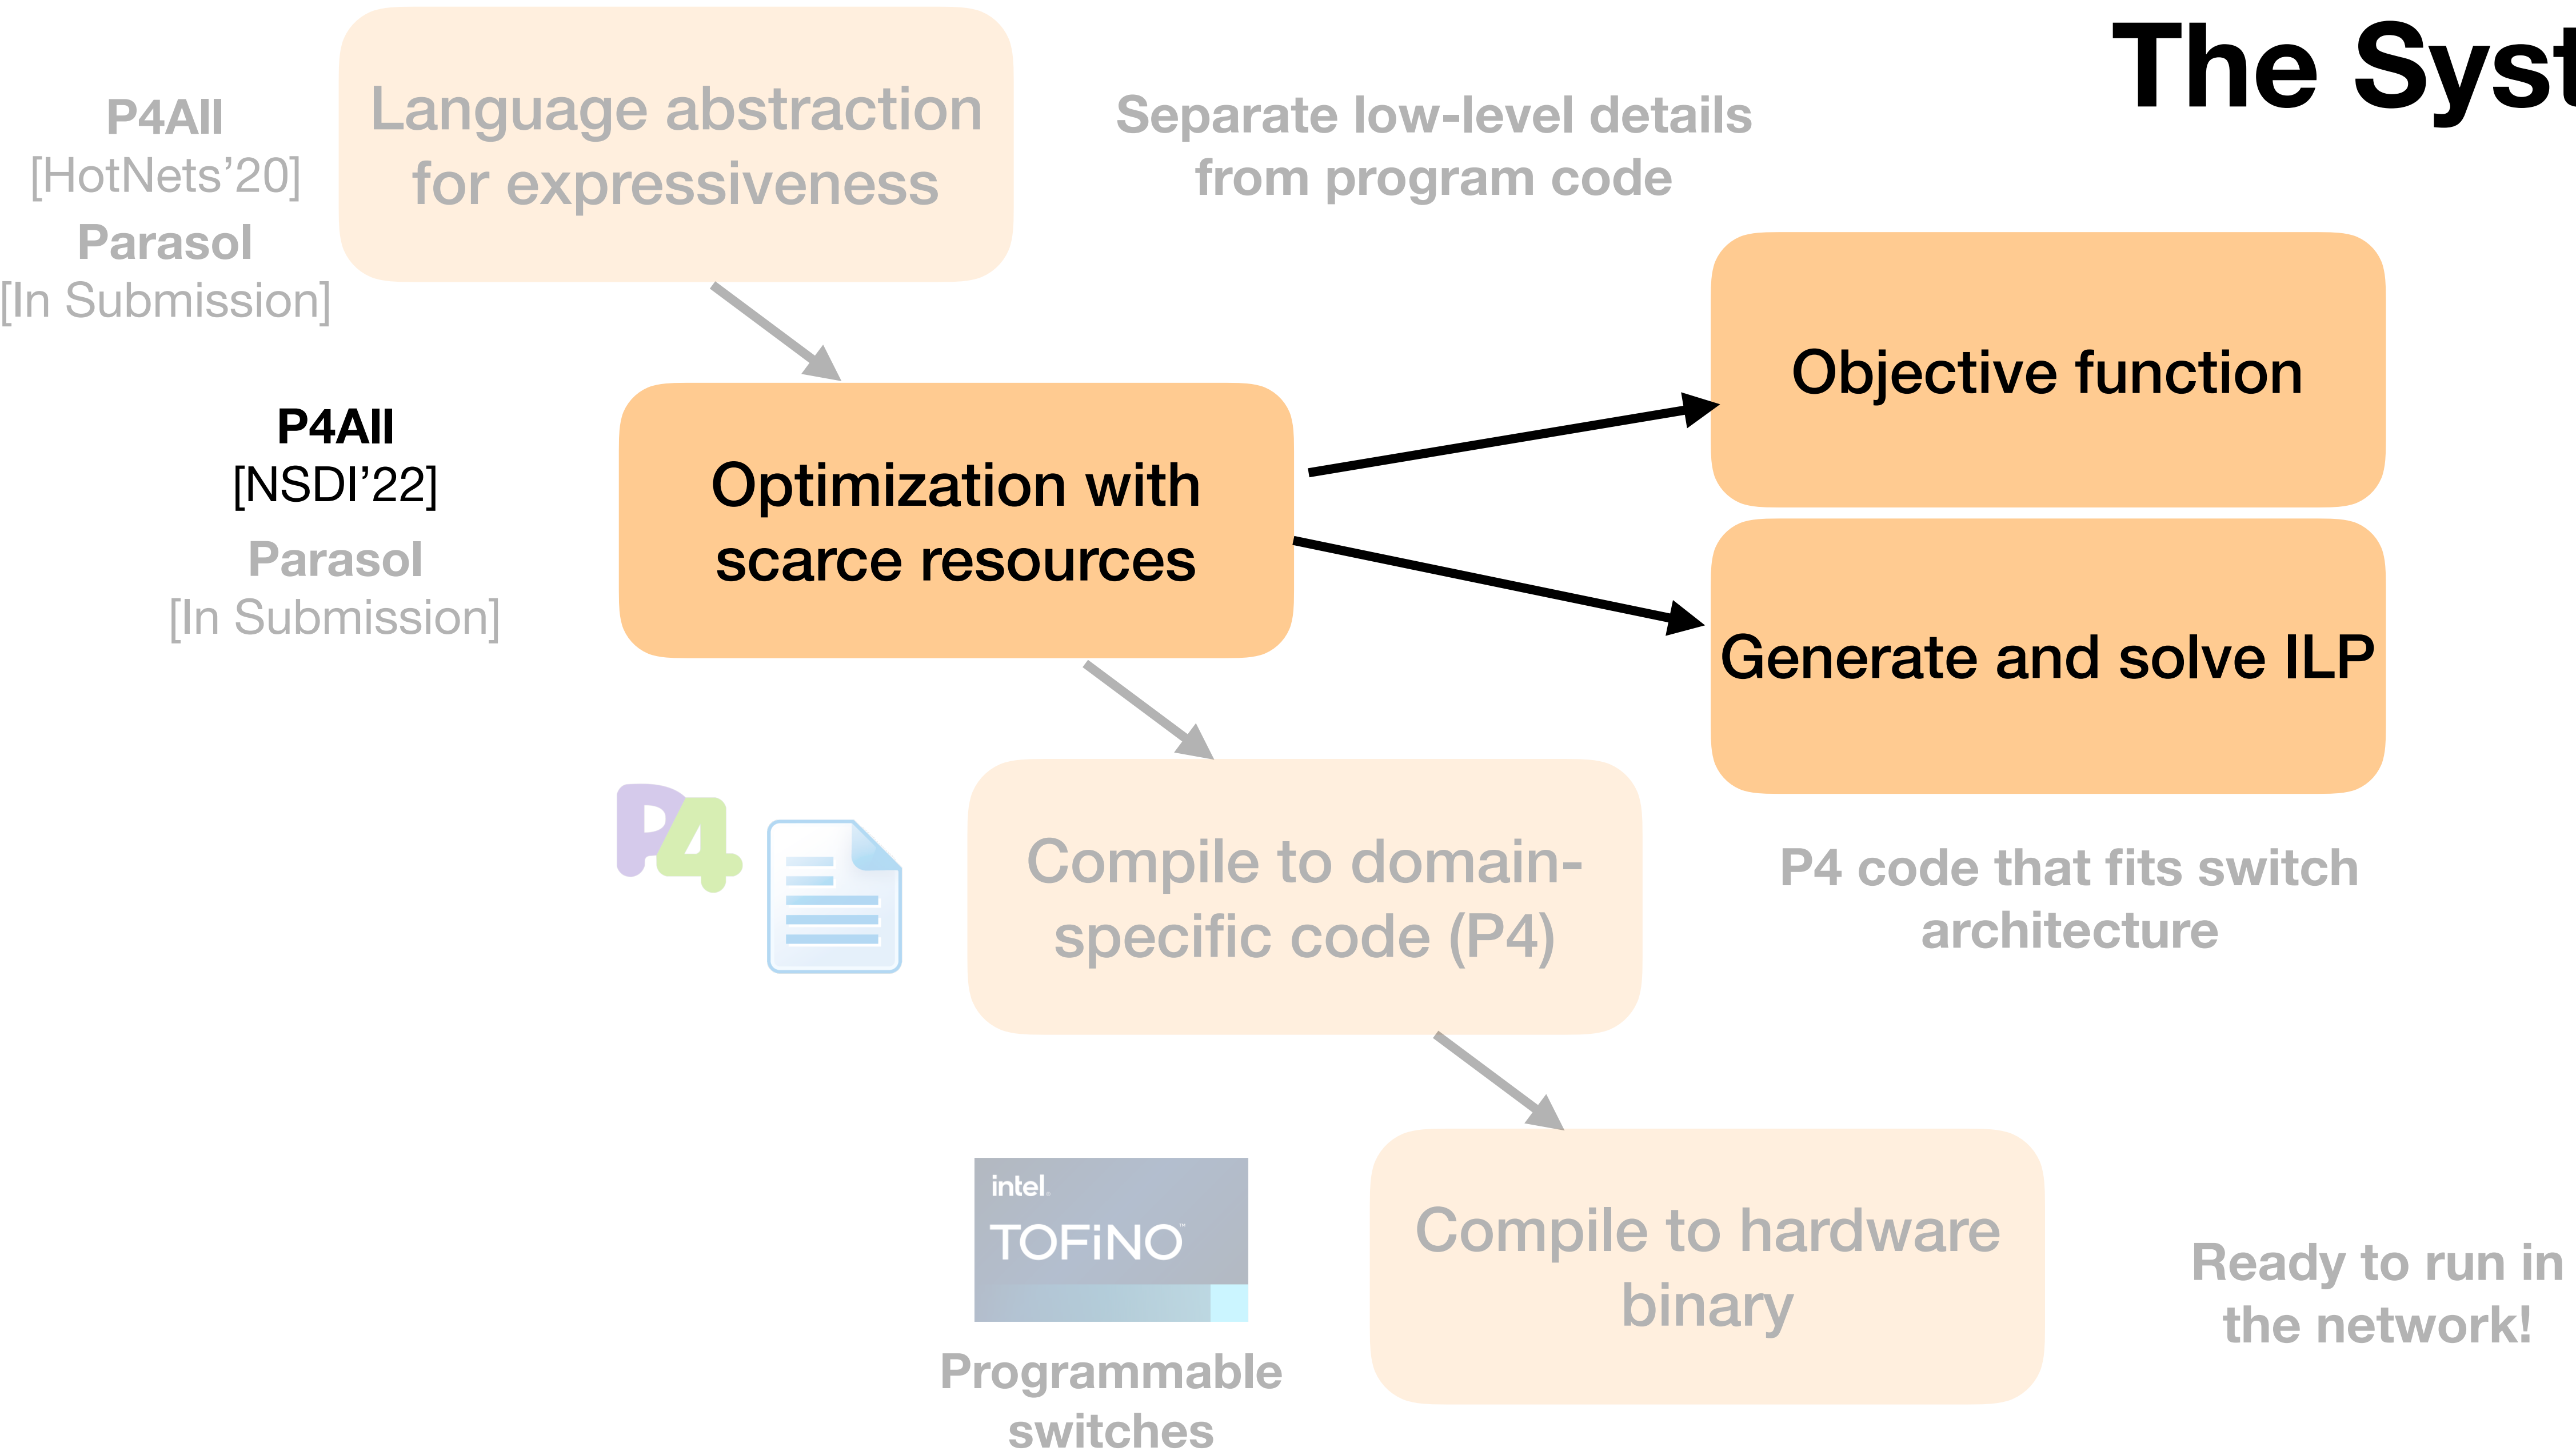
**(If the objective is not linear, we use a linear approximation)**

# Optimal Allocation = Optimal Performance

**ILP:**      maximize    **objective function**  
              subject to    **resource constraints**

**Objective function = performance as a function of structure size**

# The System




**P4All  
Program**

**Target Specification  
(resource constraints, etc.)**



+



**P4All Optimizer**

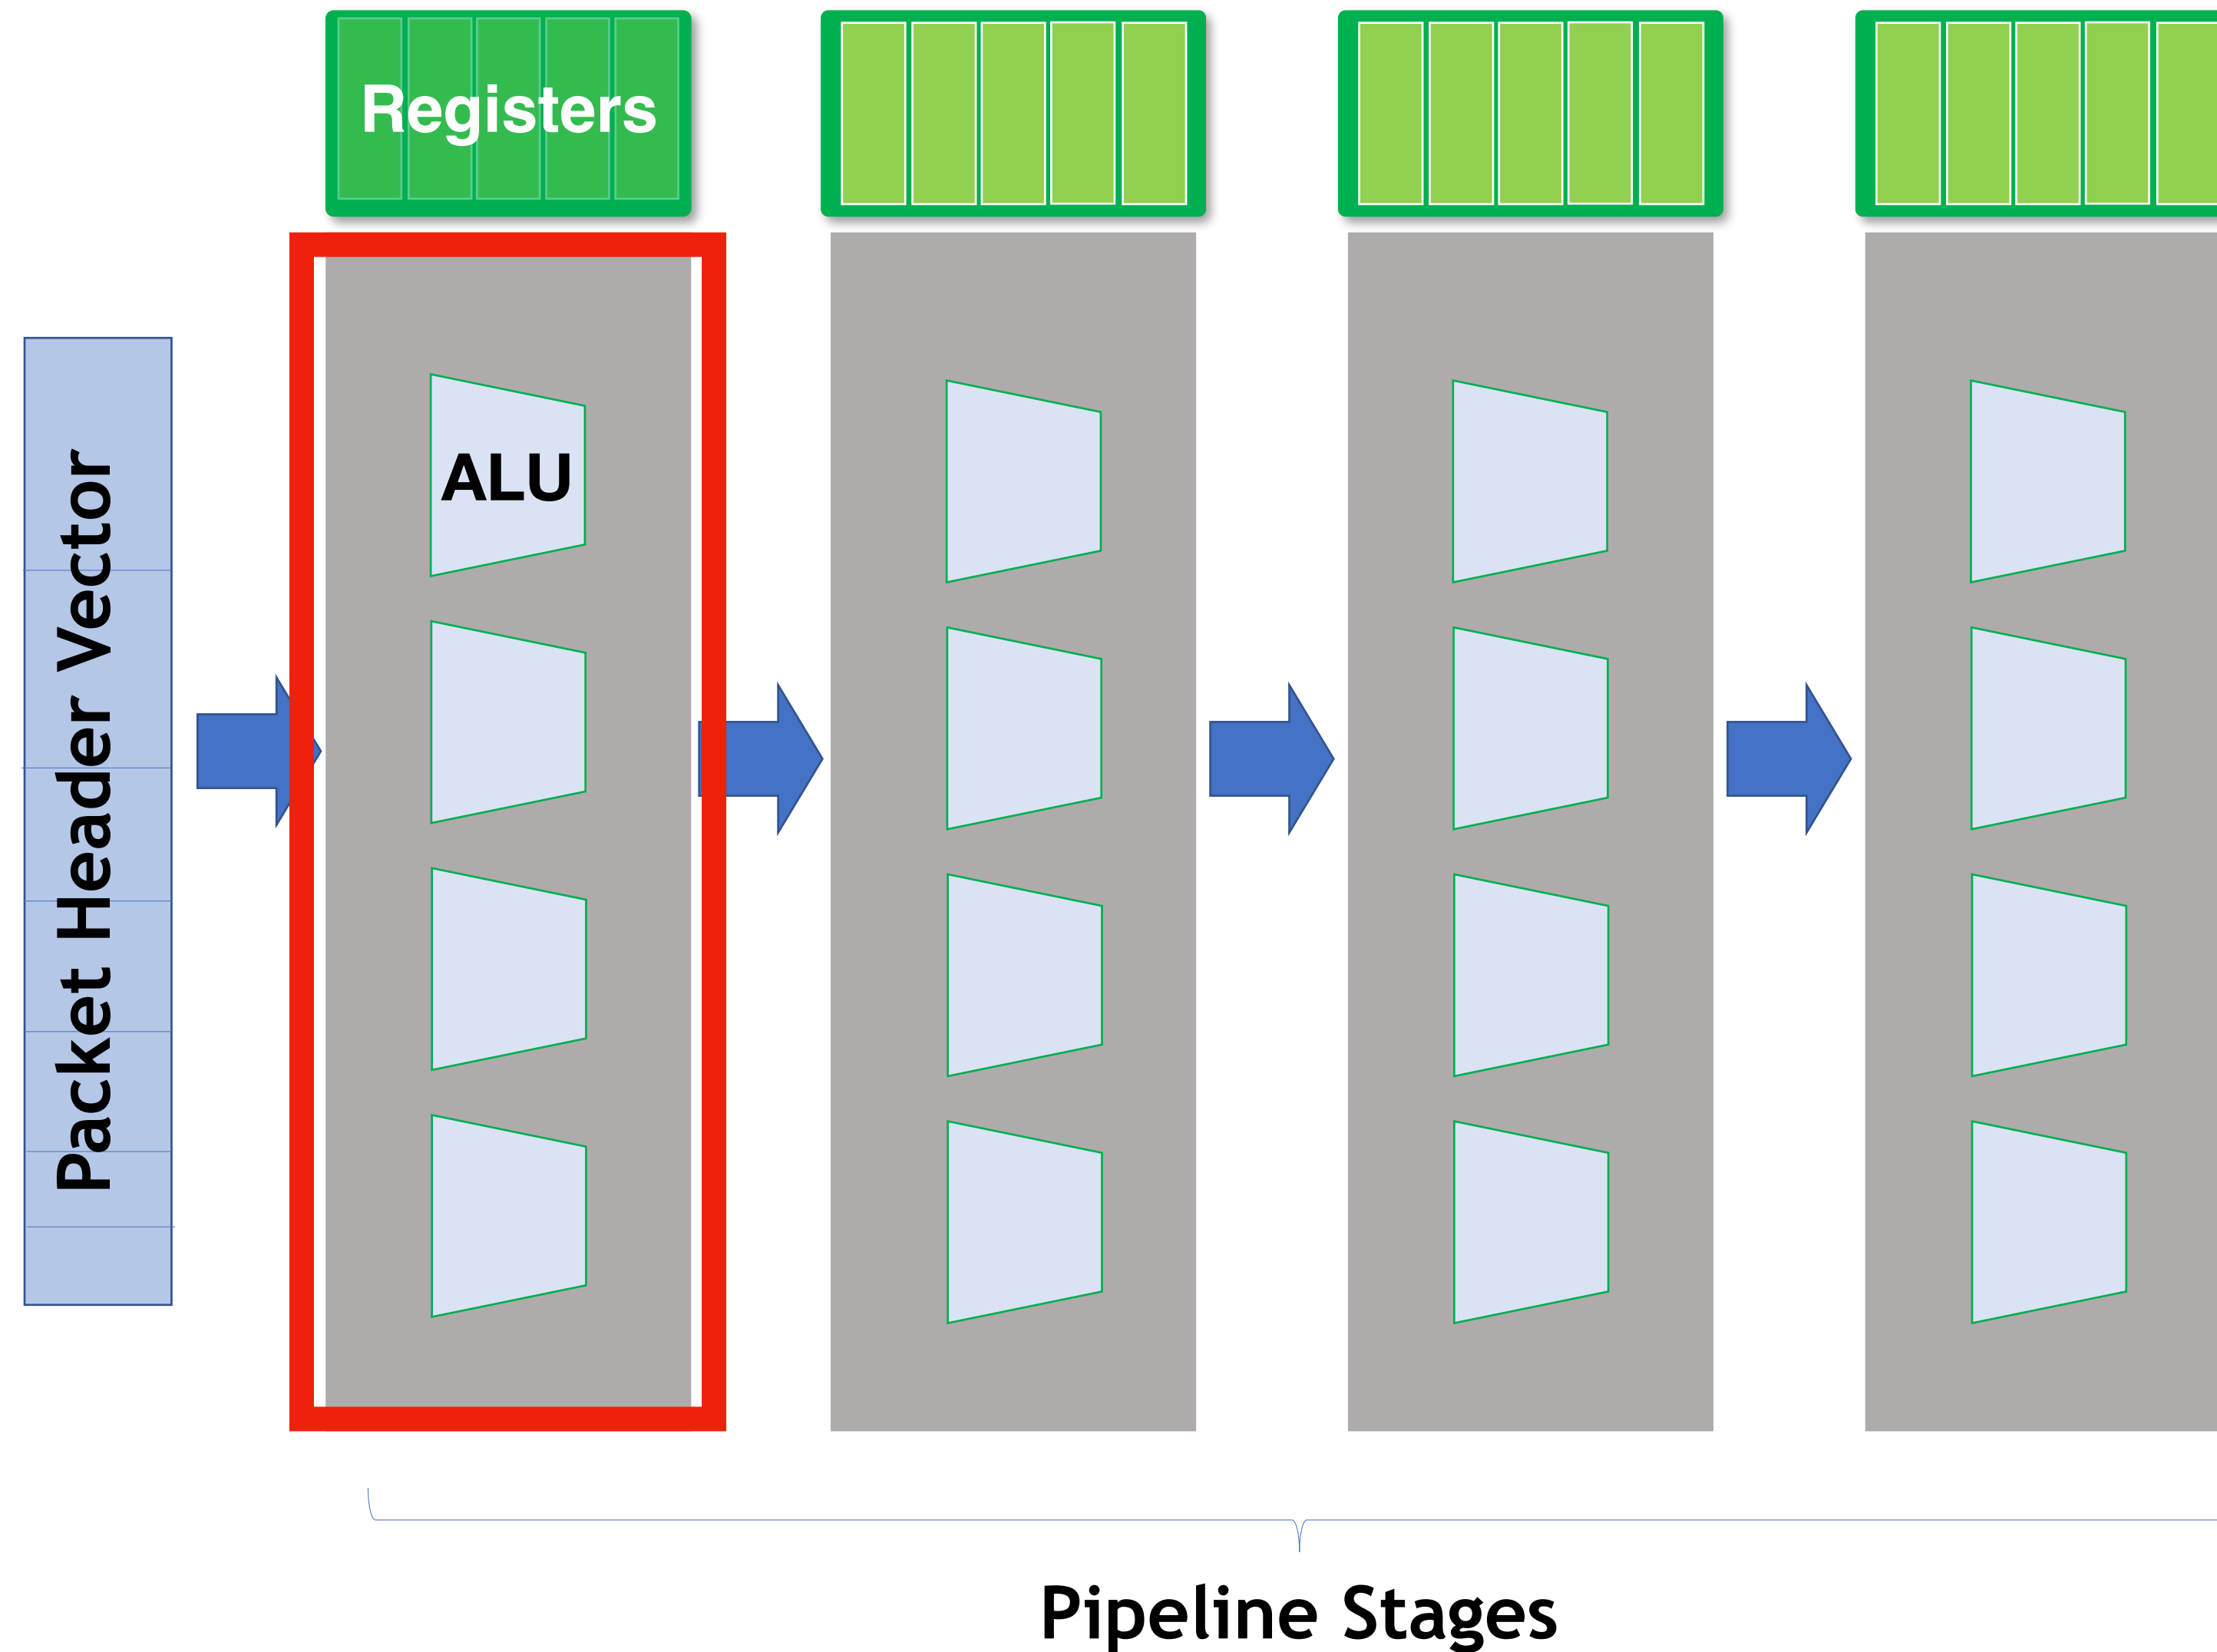
**Unroll Symbolic Loops**

**Generate and Solve Integer-  
Linear Program (ILP)**



**Concrete values  
for symbolic values  
(P4 Program)**

# Resources Determine Bound






**P4All  
Program**

**Target Specification  
(resource constraints, etc.)**



+



**P4All Optimizer**

**Parse and Generate  
Dependency Graph**

**Generate and Solve Integer-  
Linear Program (ILP)**



**Concrete values  
for symbolic values  
(P4 Program)**

# Program Dependencies

```
def cms_hash () {  
  for (i < cms_rows) {  
    hash_and_increment(cms[i]); ← Read after write dependency  
  }  
}
```

We read and write to min for each row


```
def hash_and_increment (cms_row) {  
  int index = hash();  
  int stored_val = read_val(cms_row, index);  
  if (checked_val < min) {  
    min = stored_val; }  
}
```

# Dependency Graph

```
def cms_hash () {  
  for (i < cms_rows) {  
    hash_and_increment(cms[i]);  
  }  
}
```

Walk the graph until  
**nodes visited = number of available stages**

**Stages available = 2**

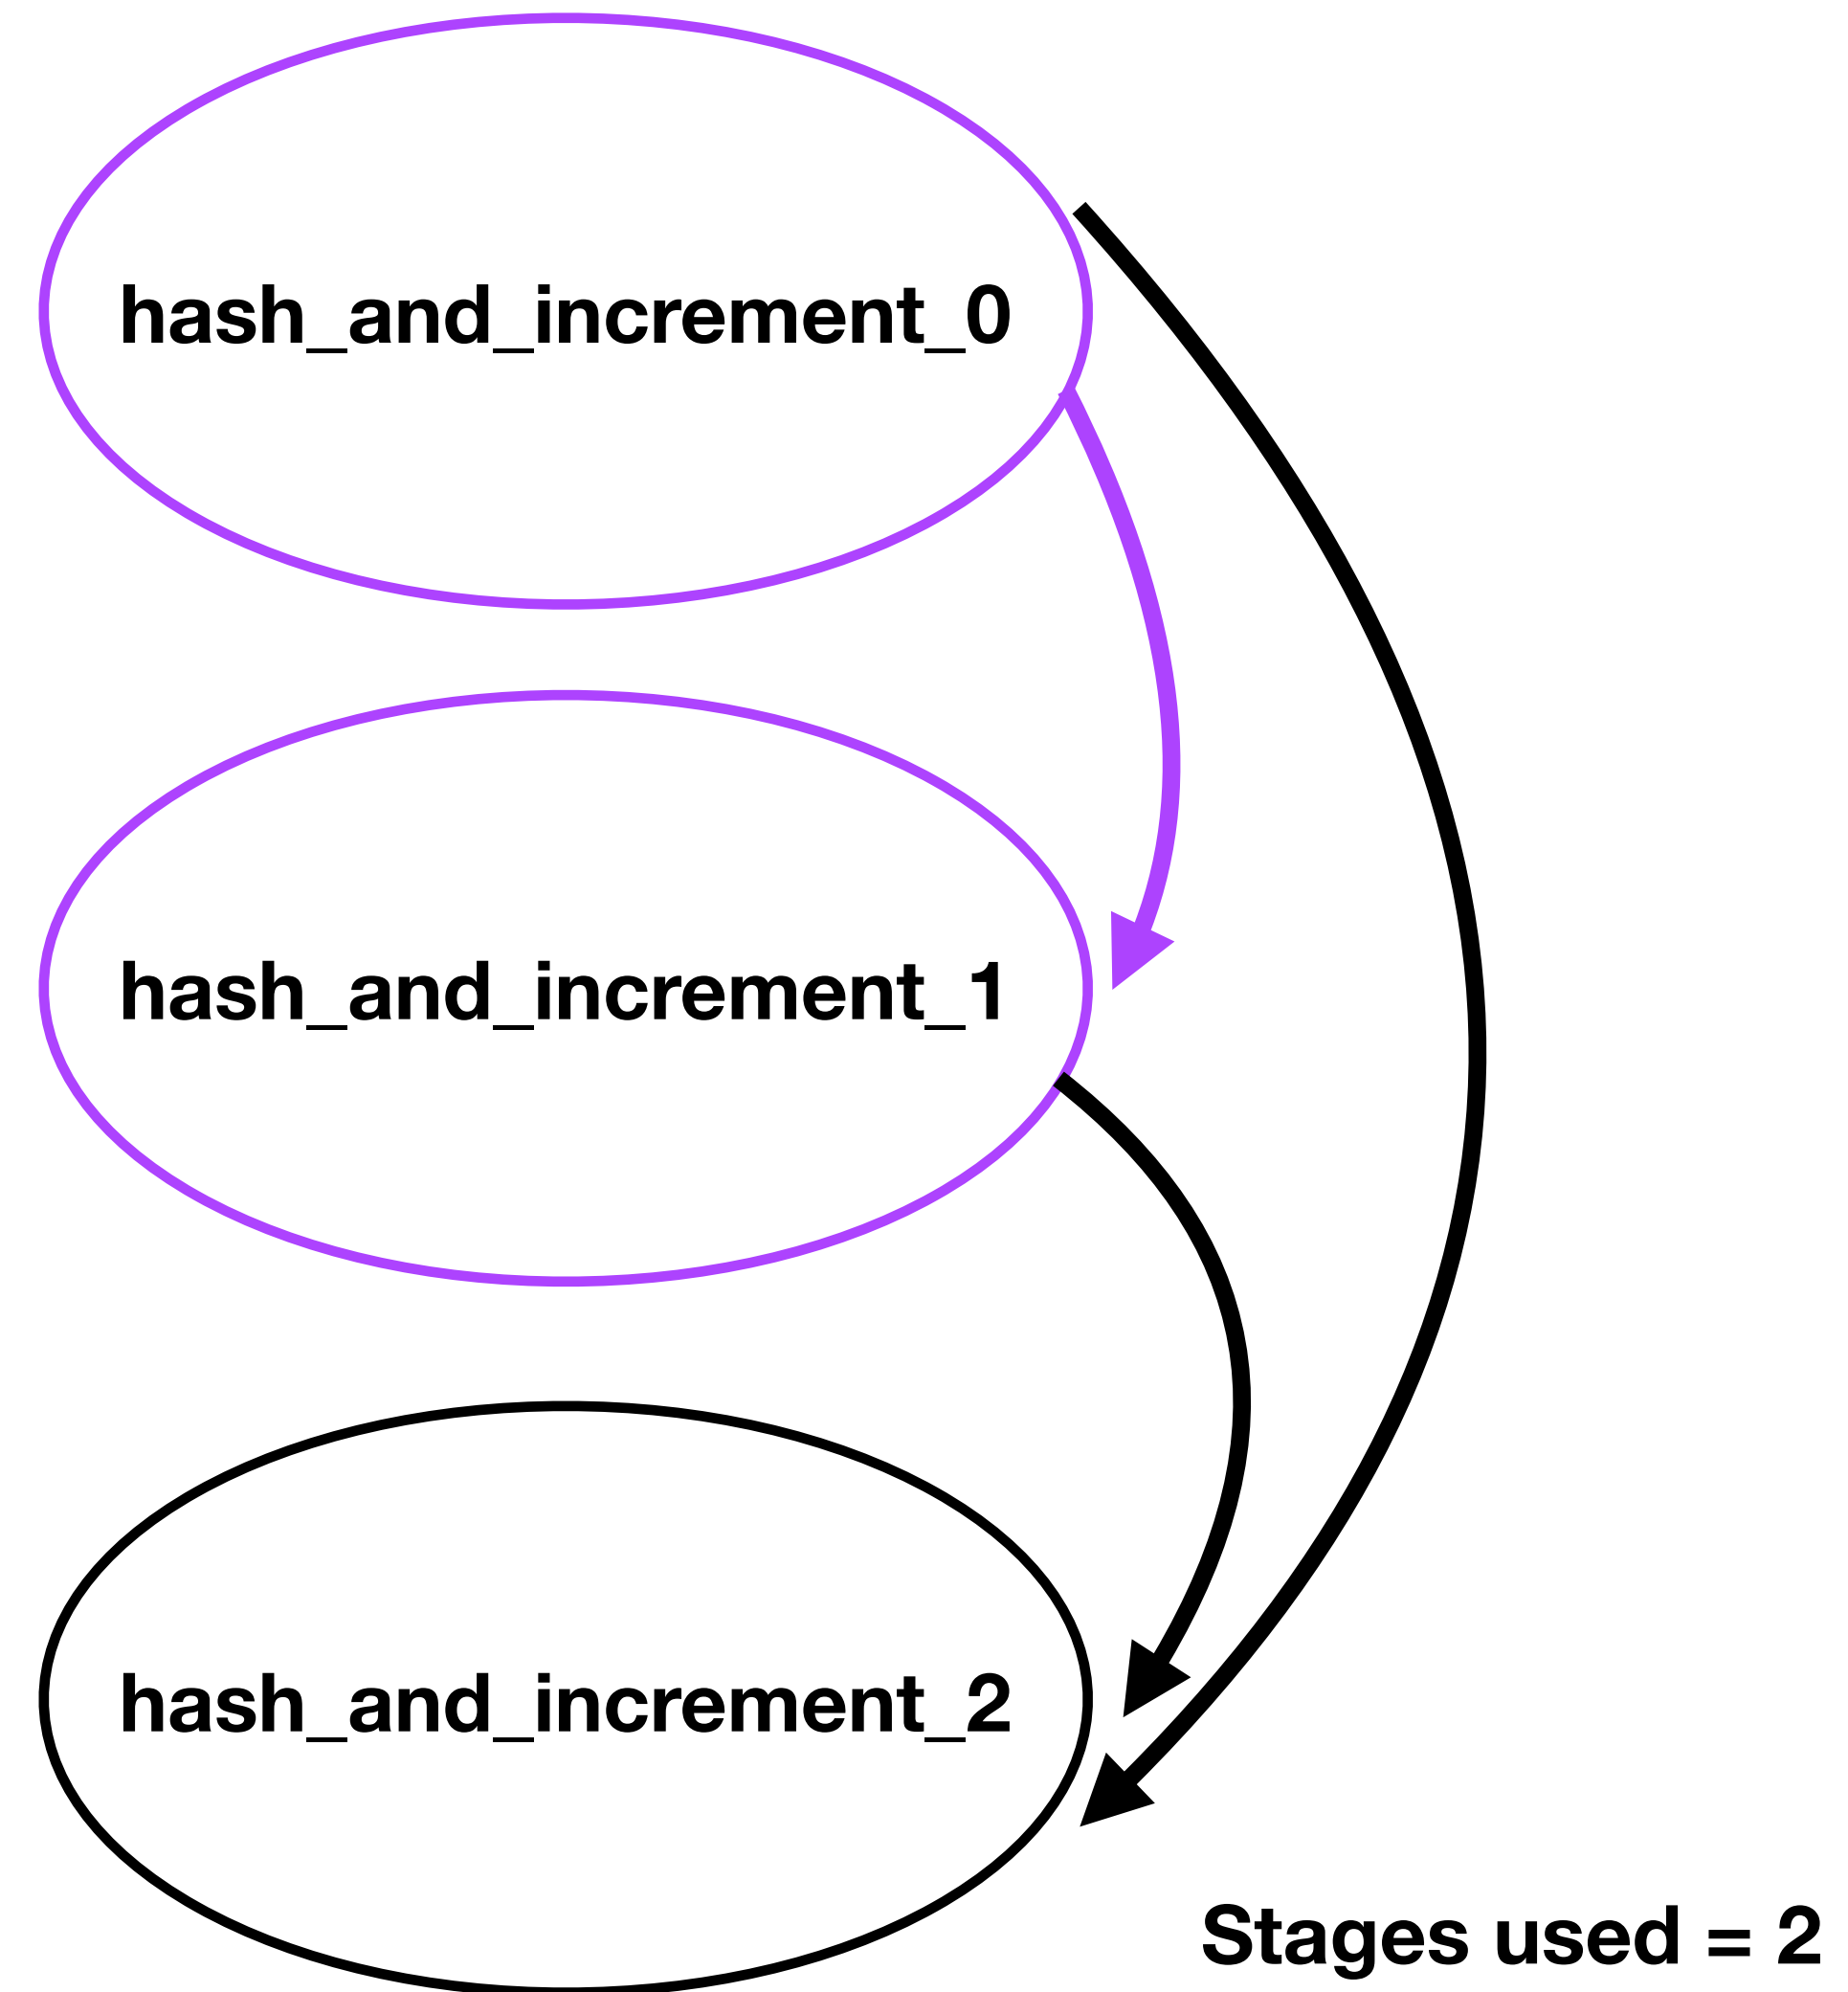


# Dependency Graph

```
def cms_hash () {  
  for (i < cms_rows) {  
    hash_and_increment(cms[i]);  
  }  
}
```

`cms_rows`  $\leq$  2

Stages available = 2





**P4All  
Program**

**Target Specification  
(resource constraints, etc.)**



+



**P4All Optimizer**

Parse and Generate  
Dependency Graph

Variables

Constraints

∩ S

Line


Objective

ILP)

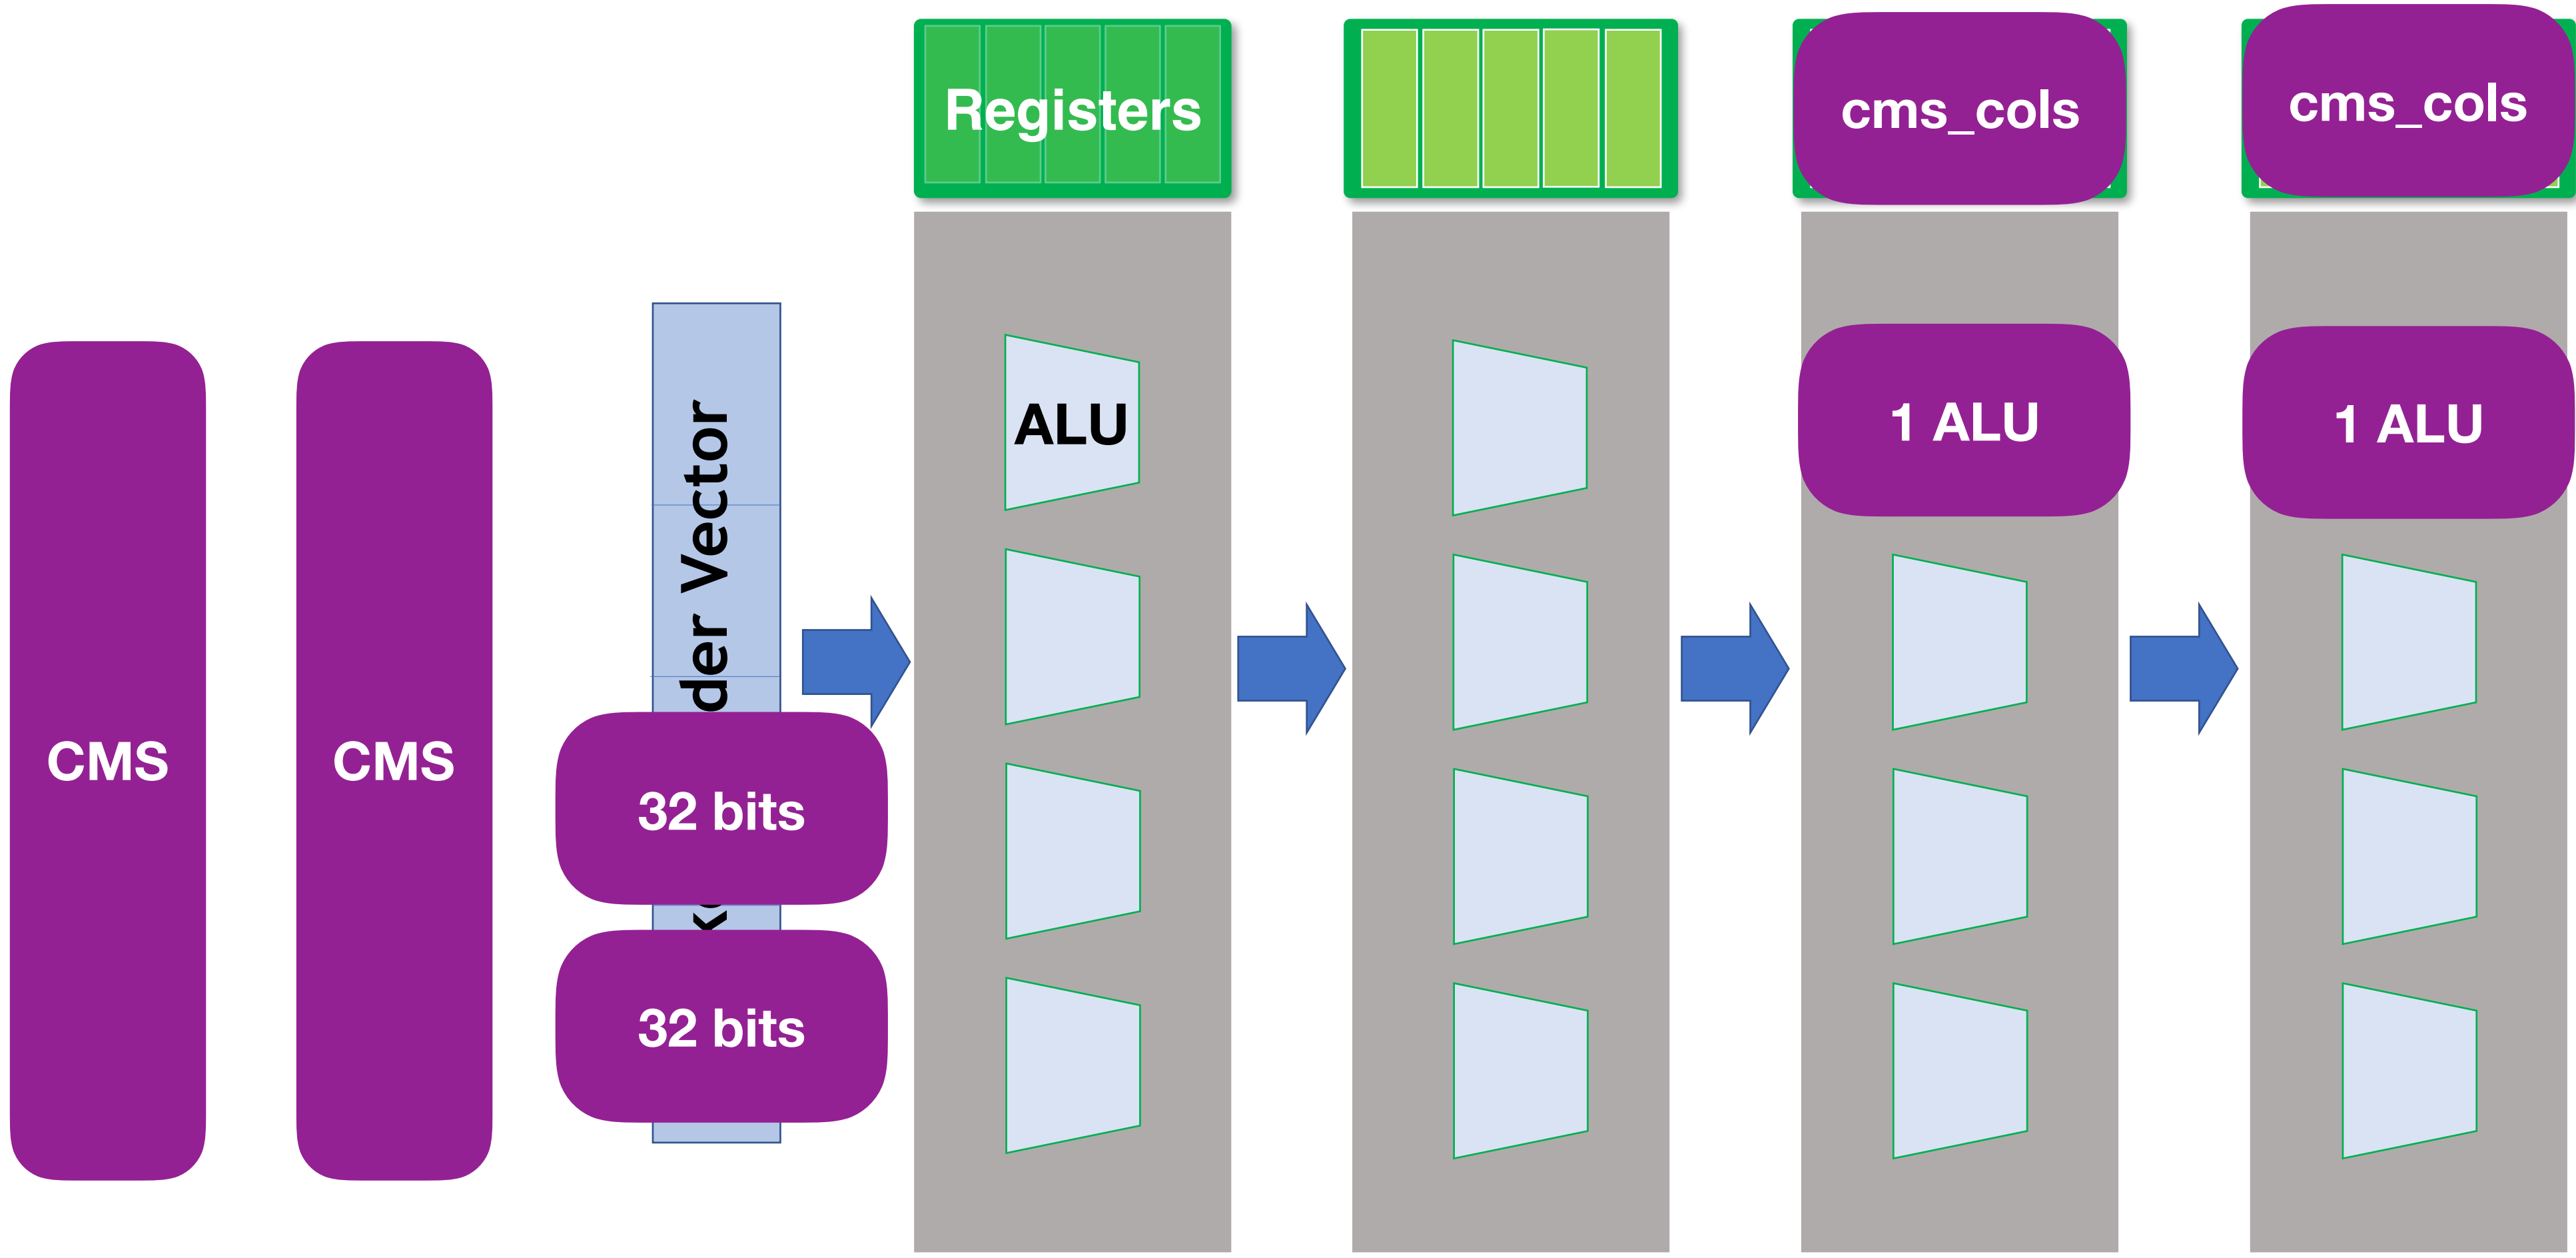


**Concrete values  
for symbolic values  
(P4 Program)**

# ILP assigns resources to program elements

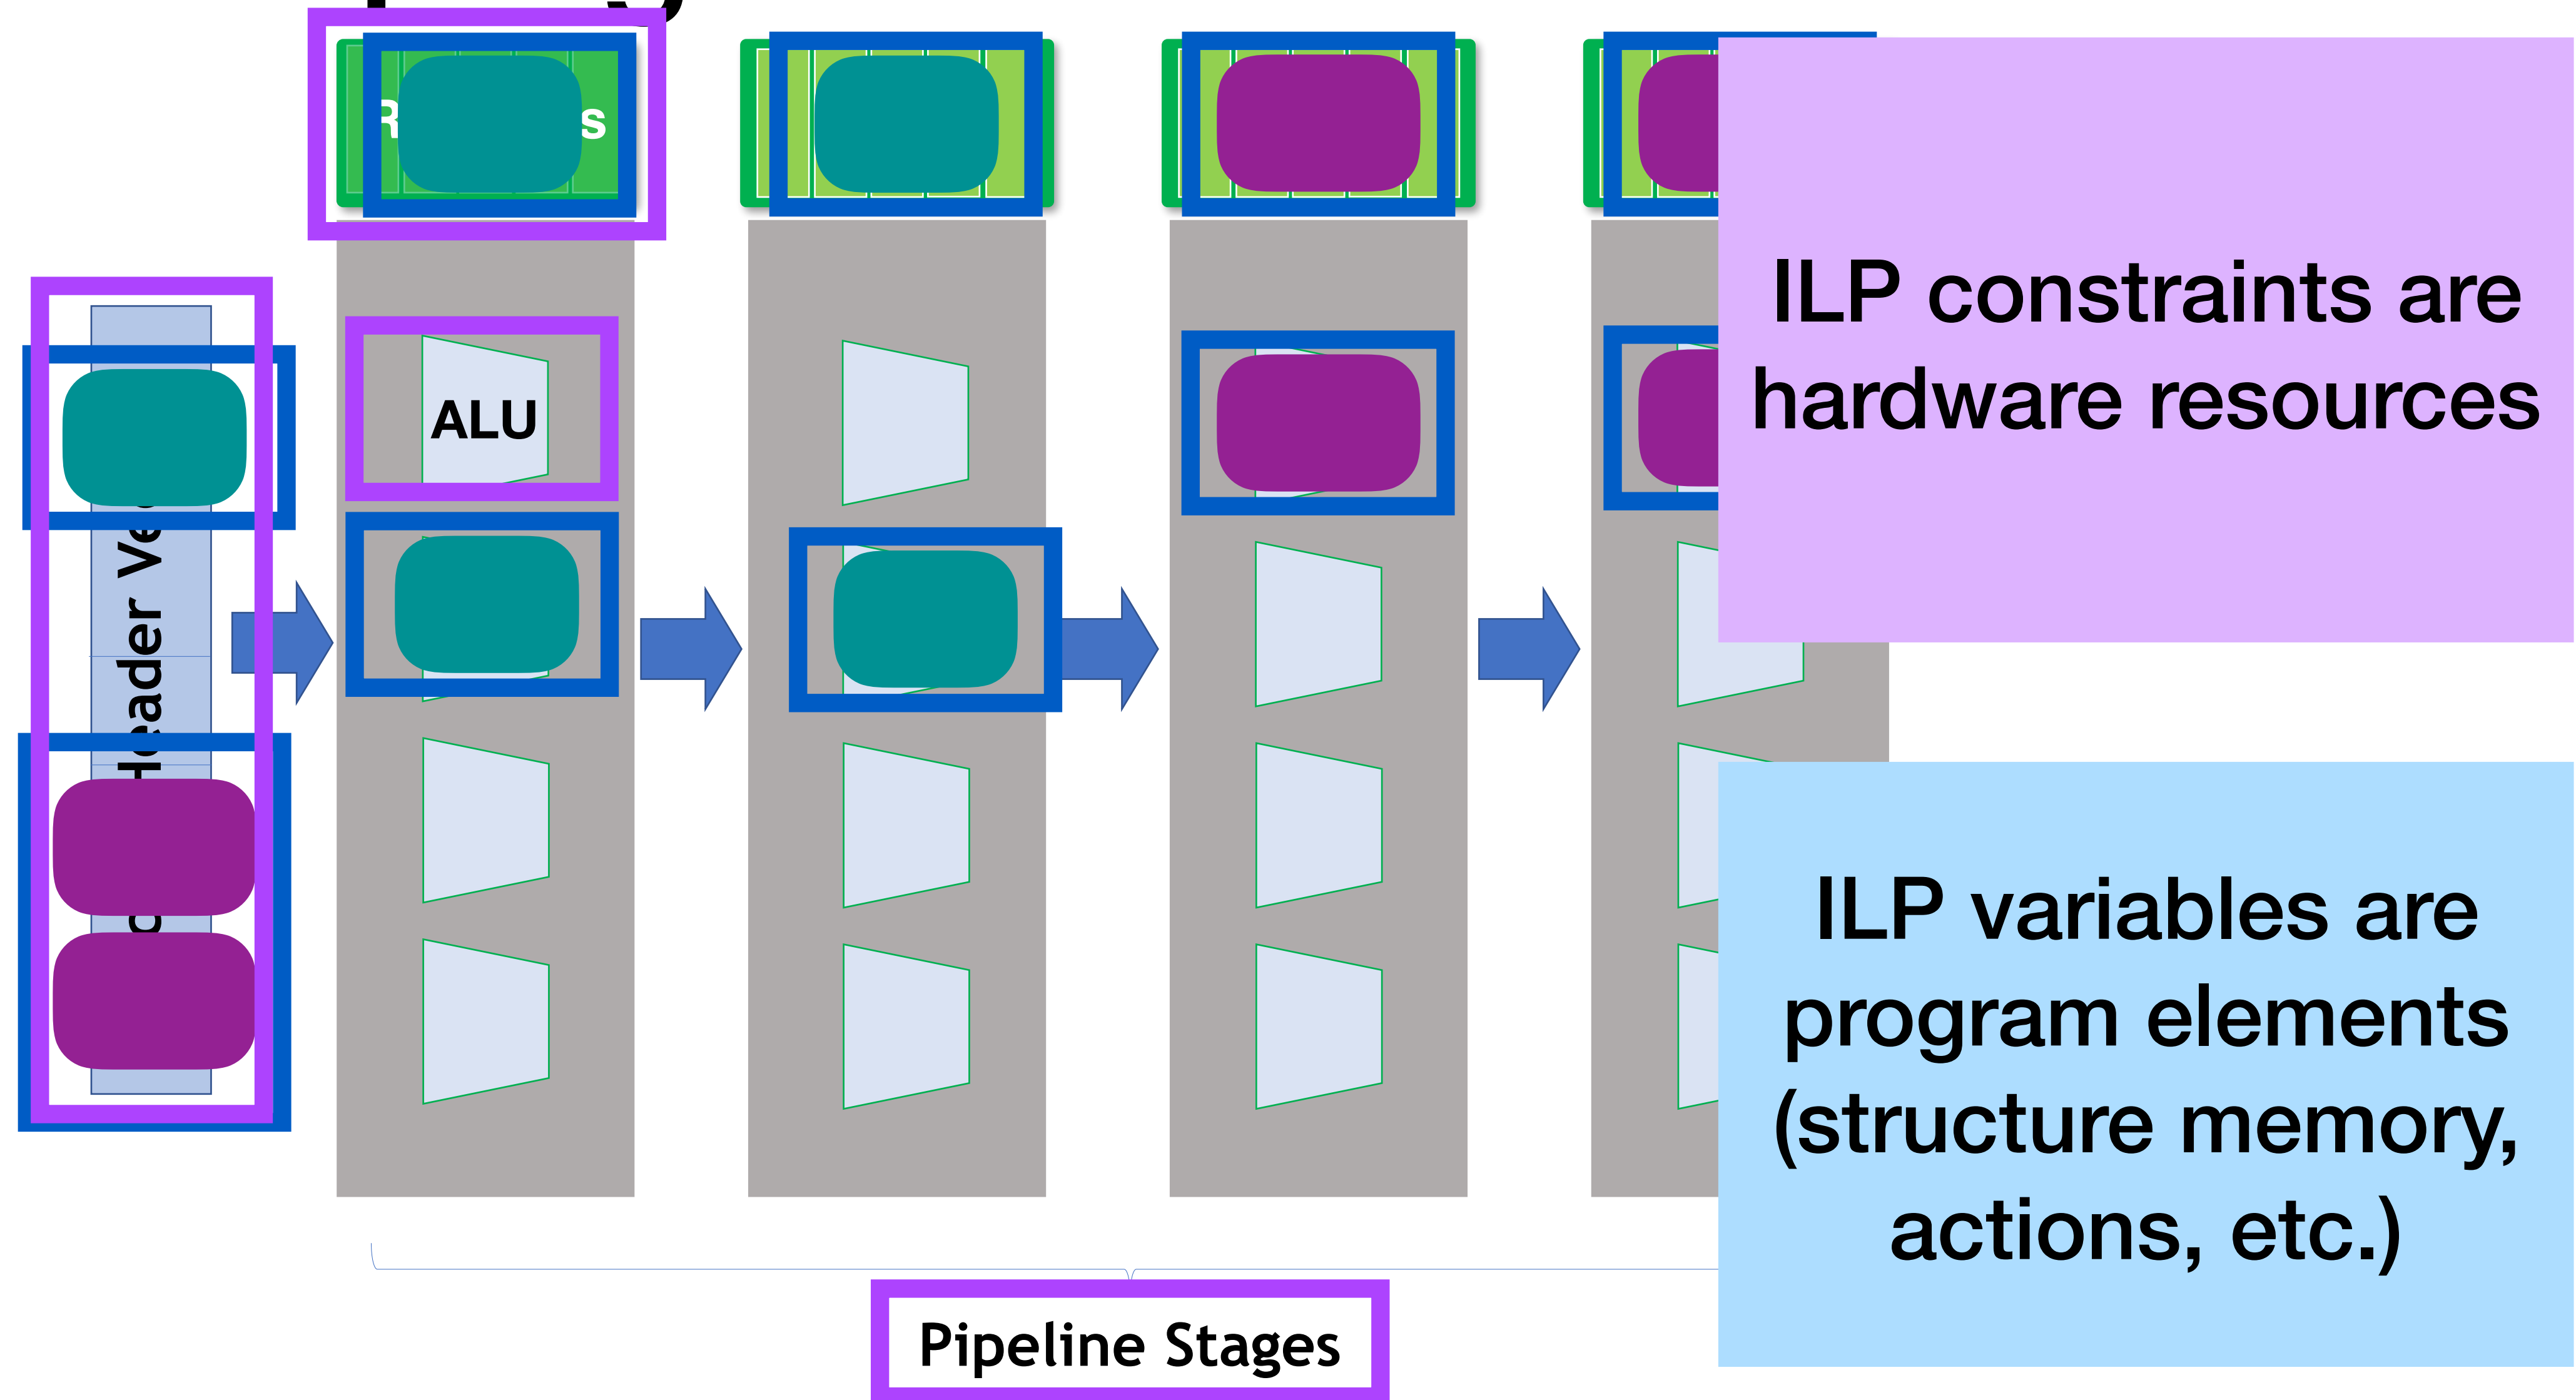


# ILP Variables



Pipeline Stages

# ILP assigns resources to program elements

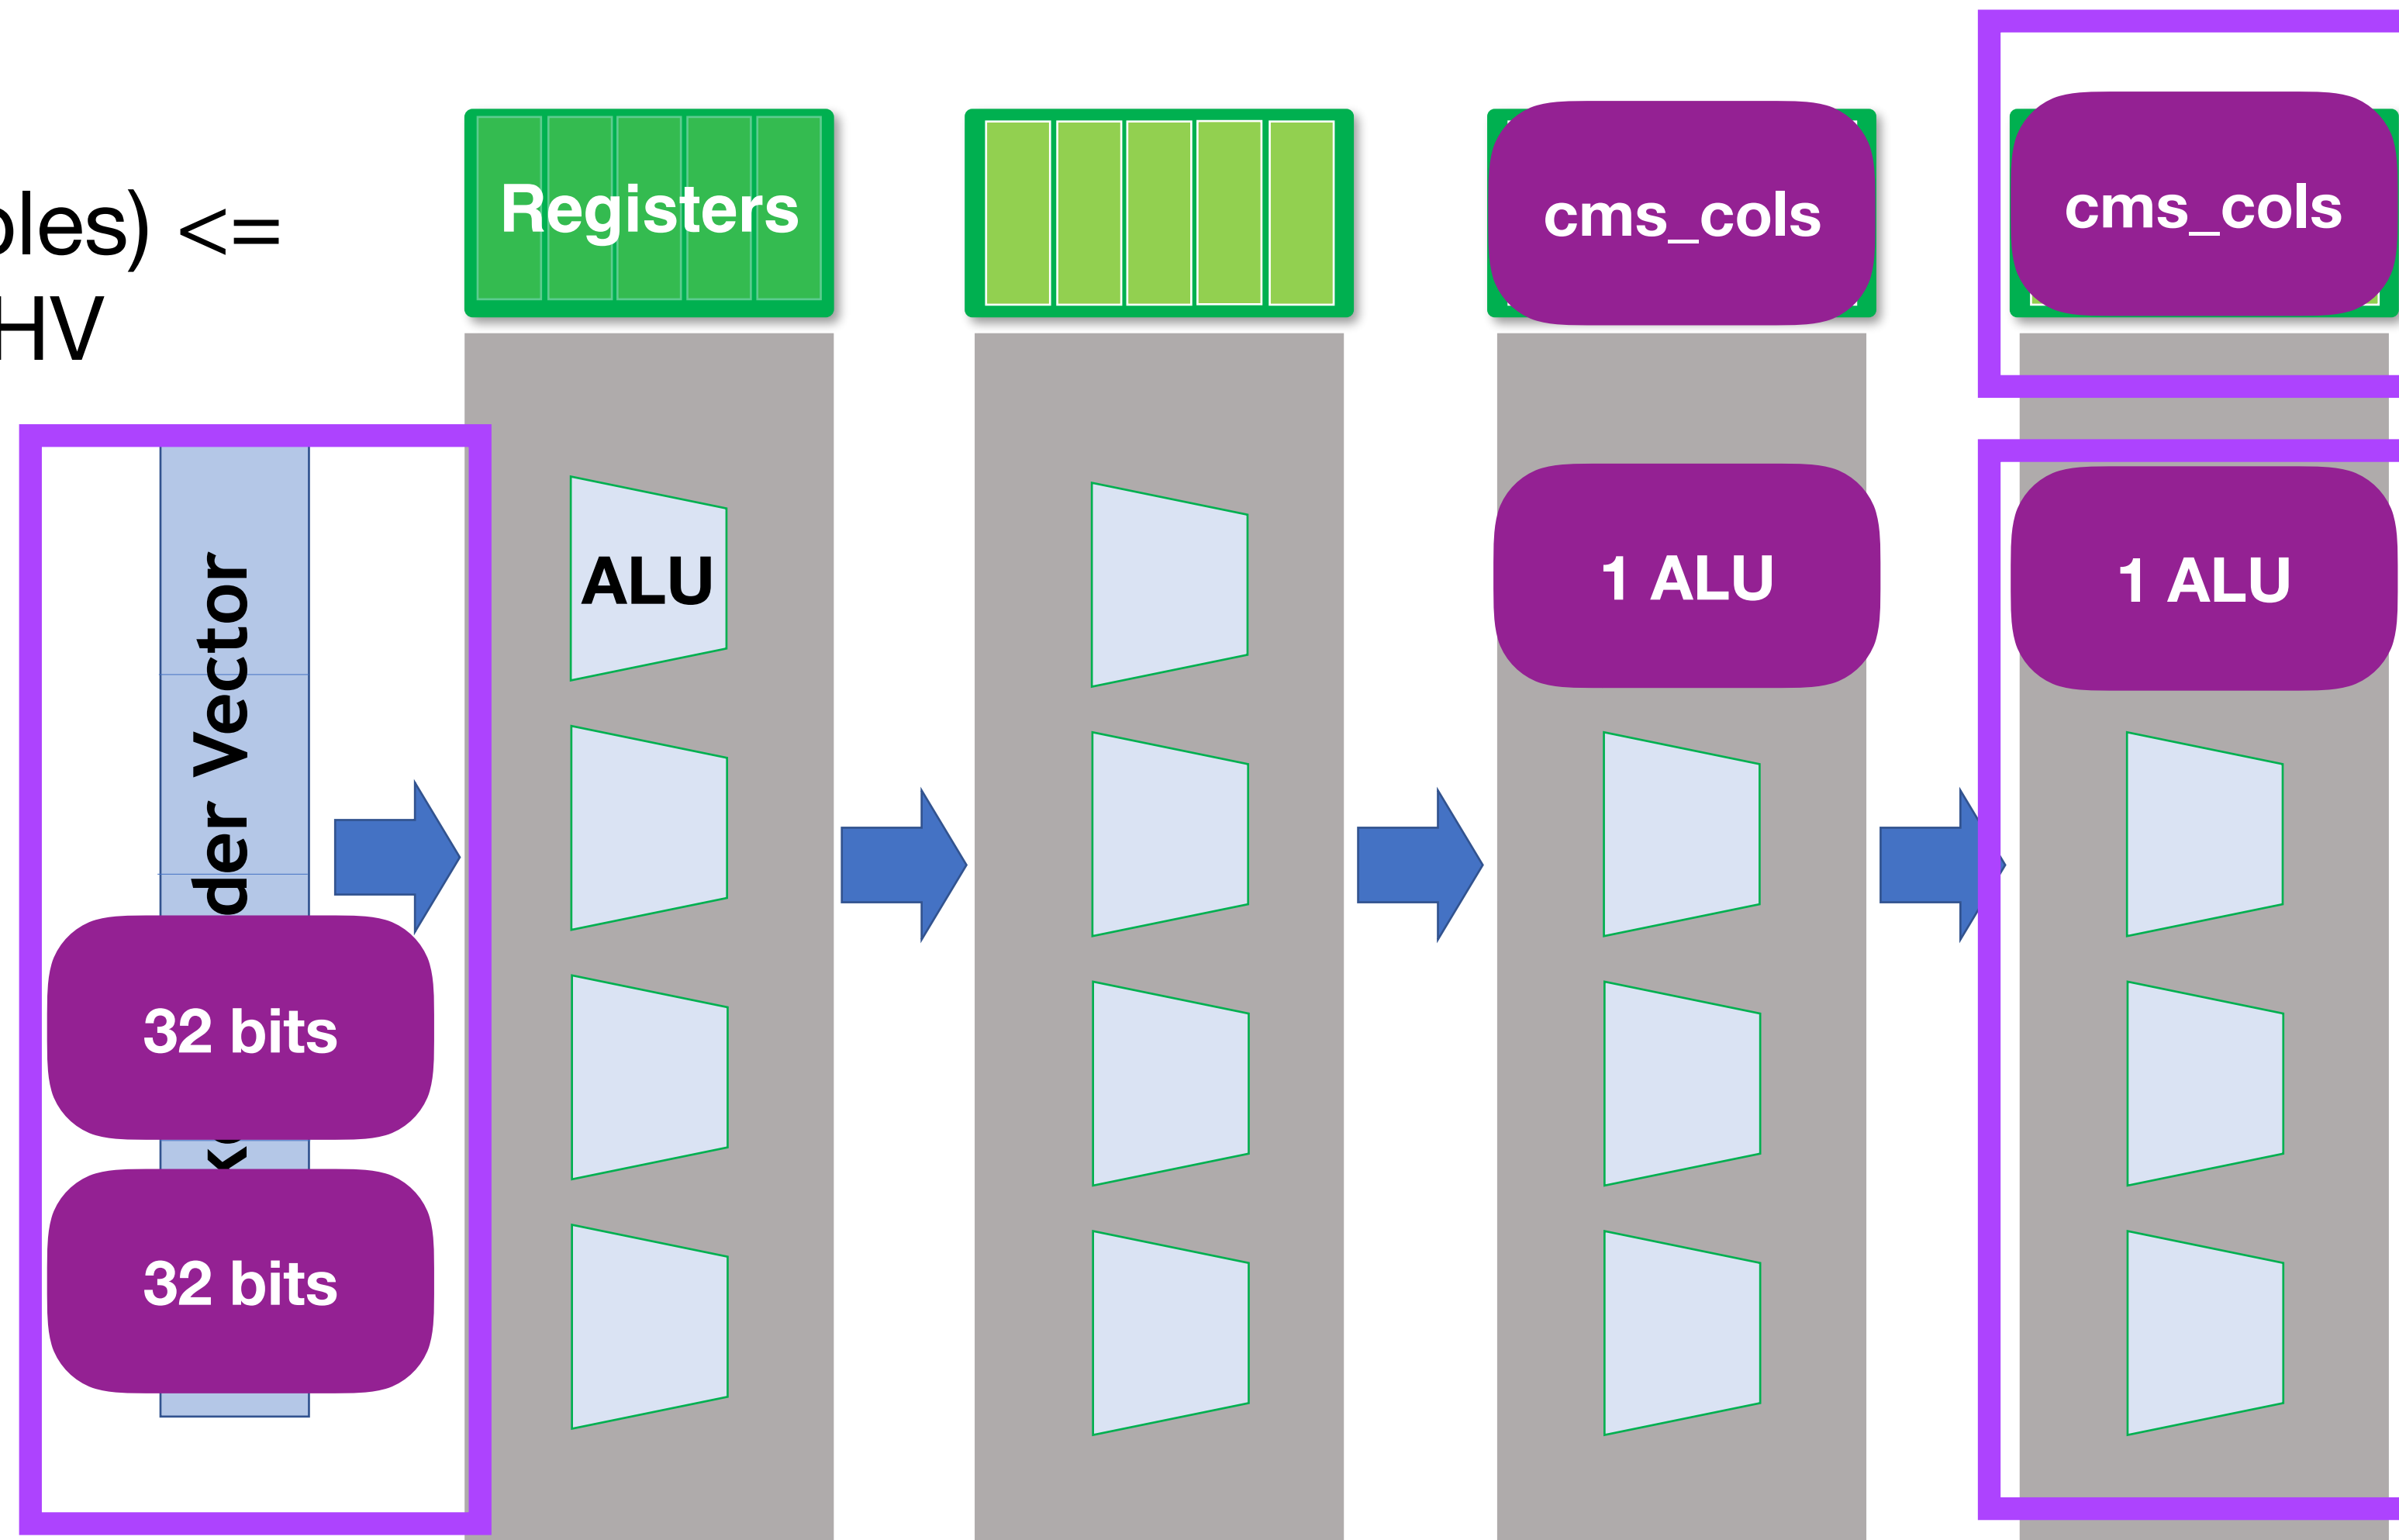




# ILP Constraints

Sum(register variables per stage)  $\leq$  Available registers per stage

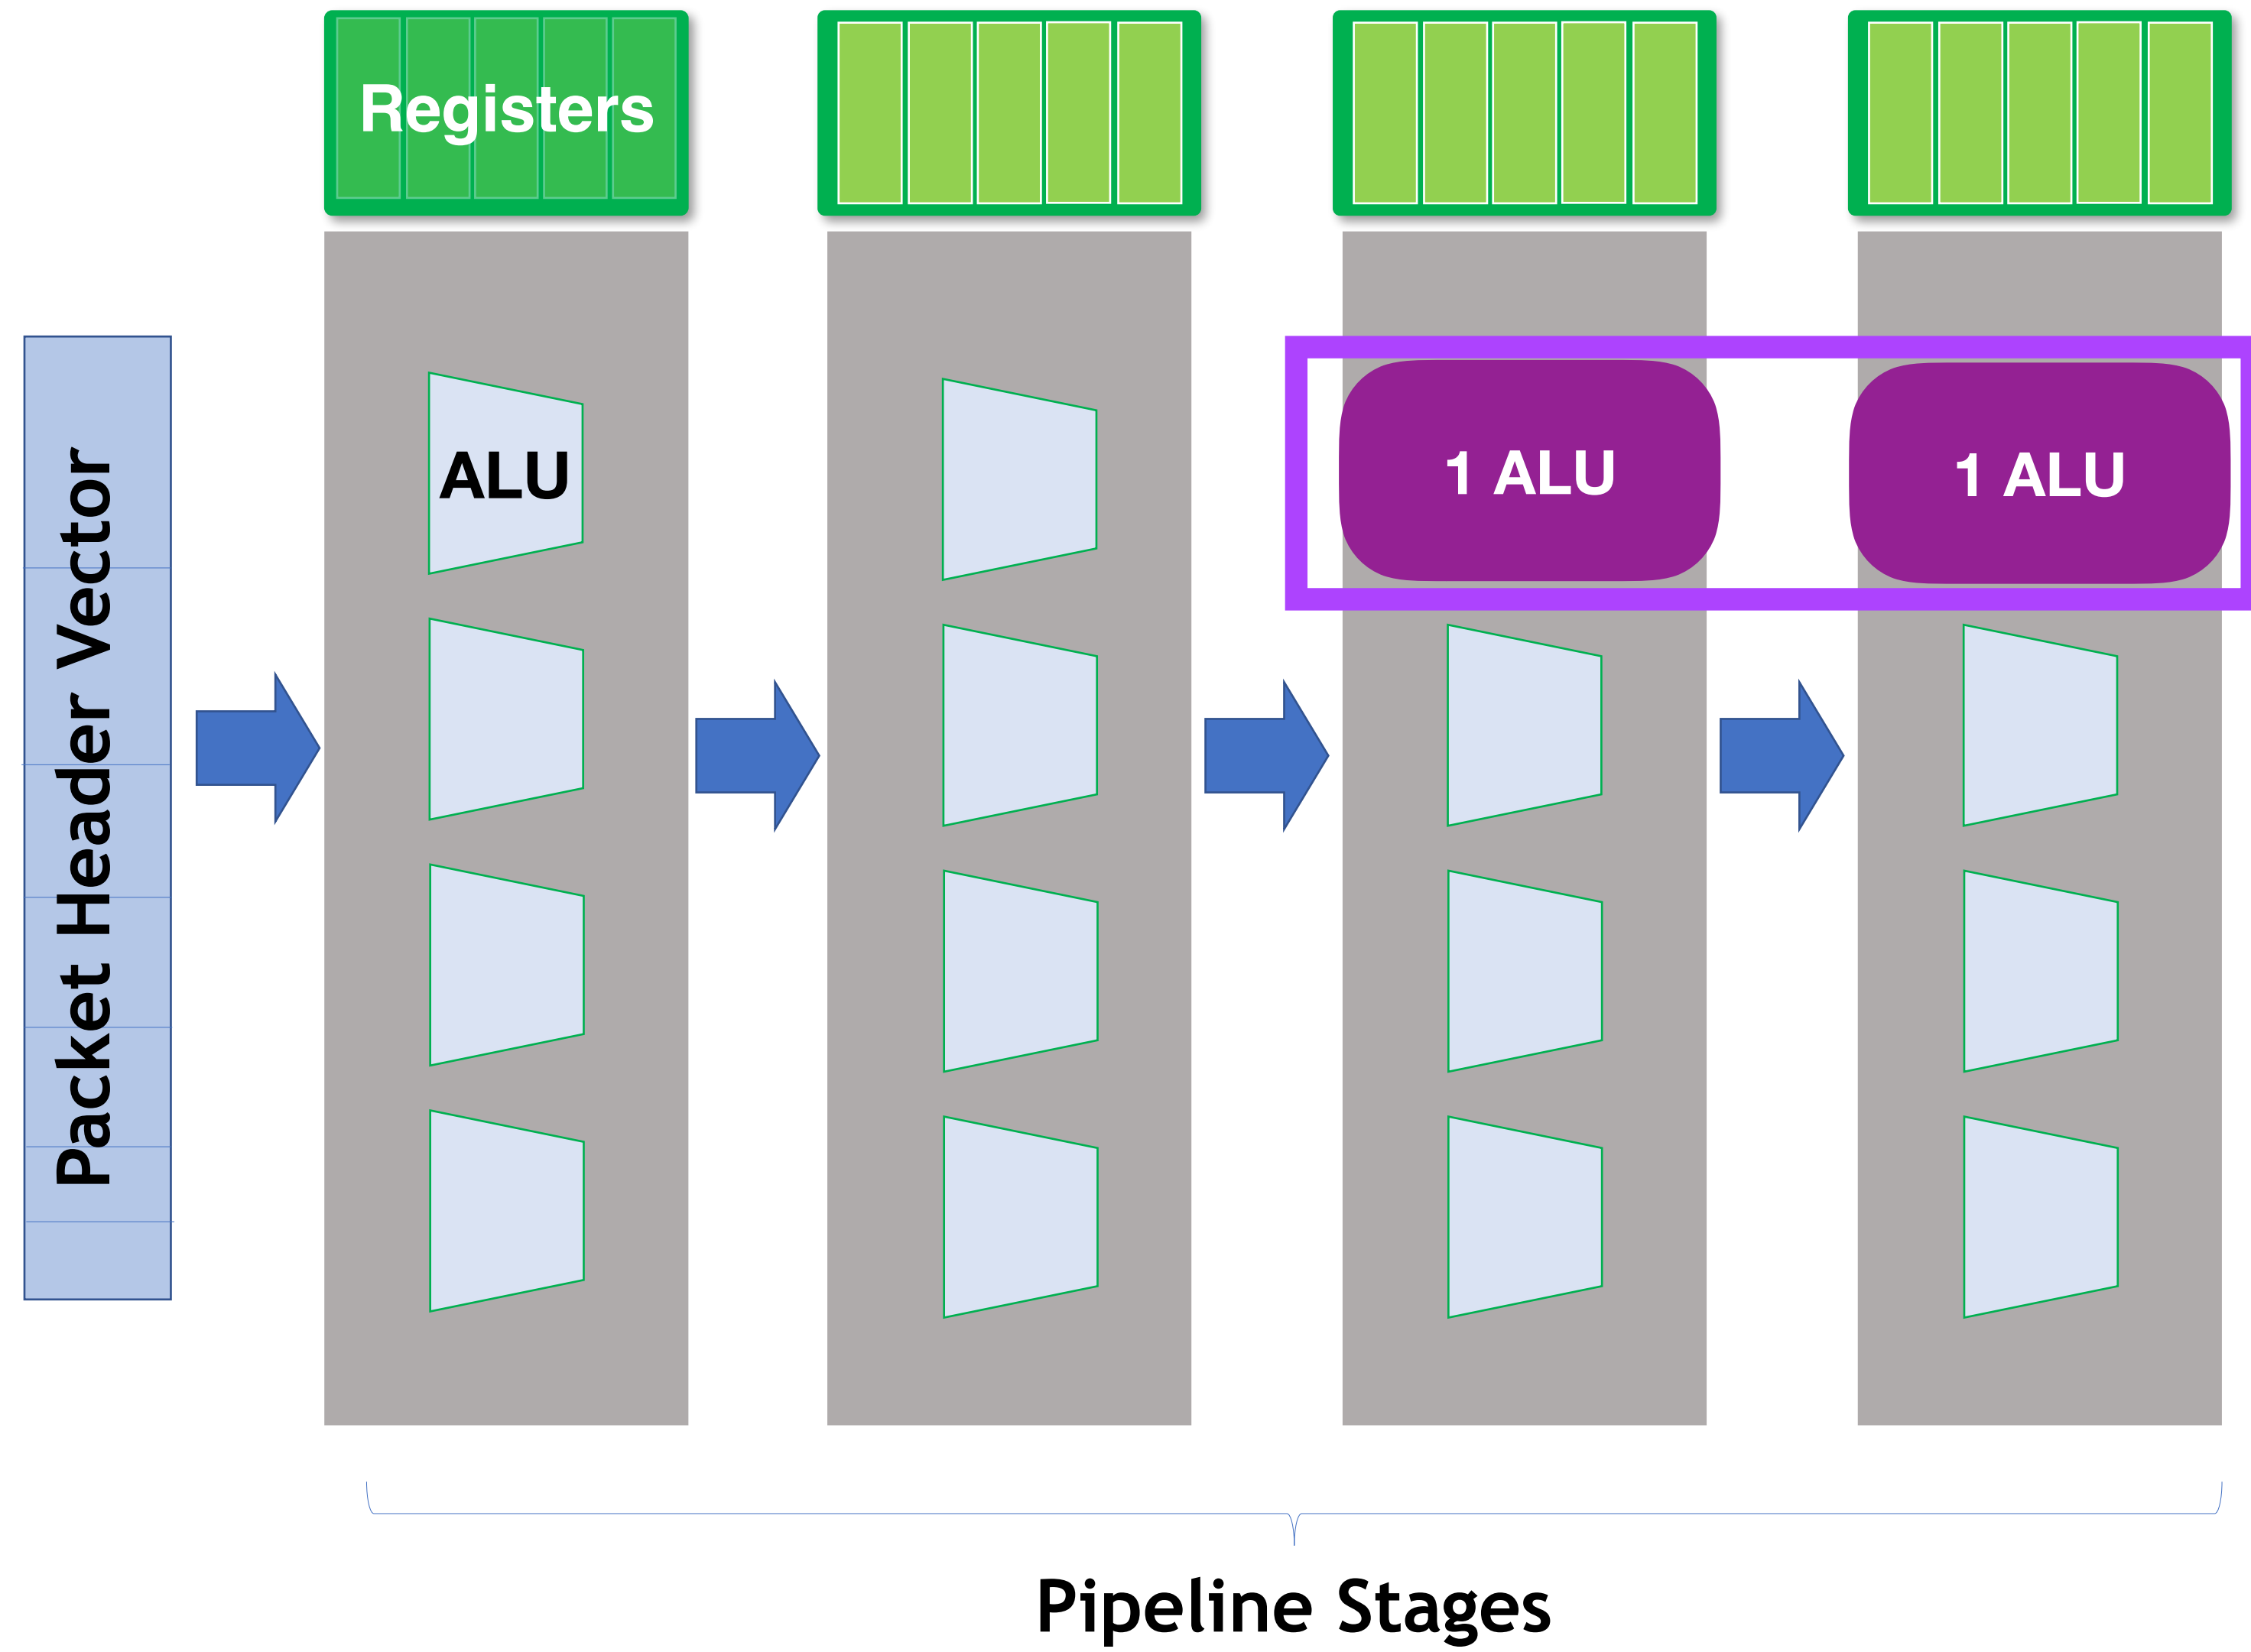
Sum(PHV variables)  $\leq$  Available PHV



Sum(action variables per stage)  $\leq$  Available ALUs per stage

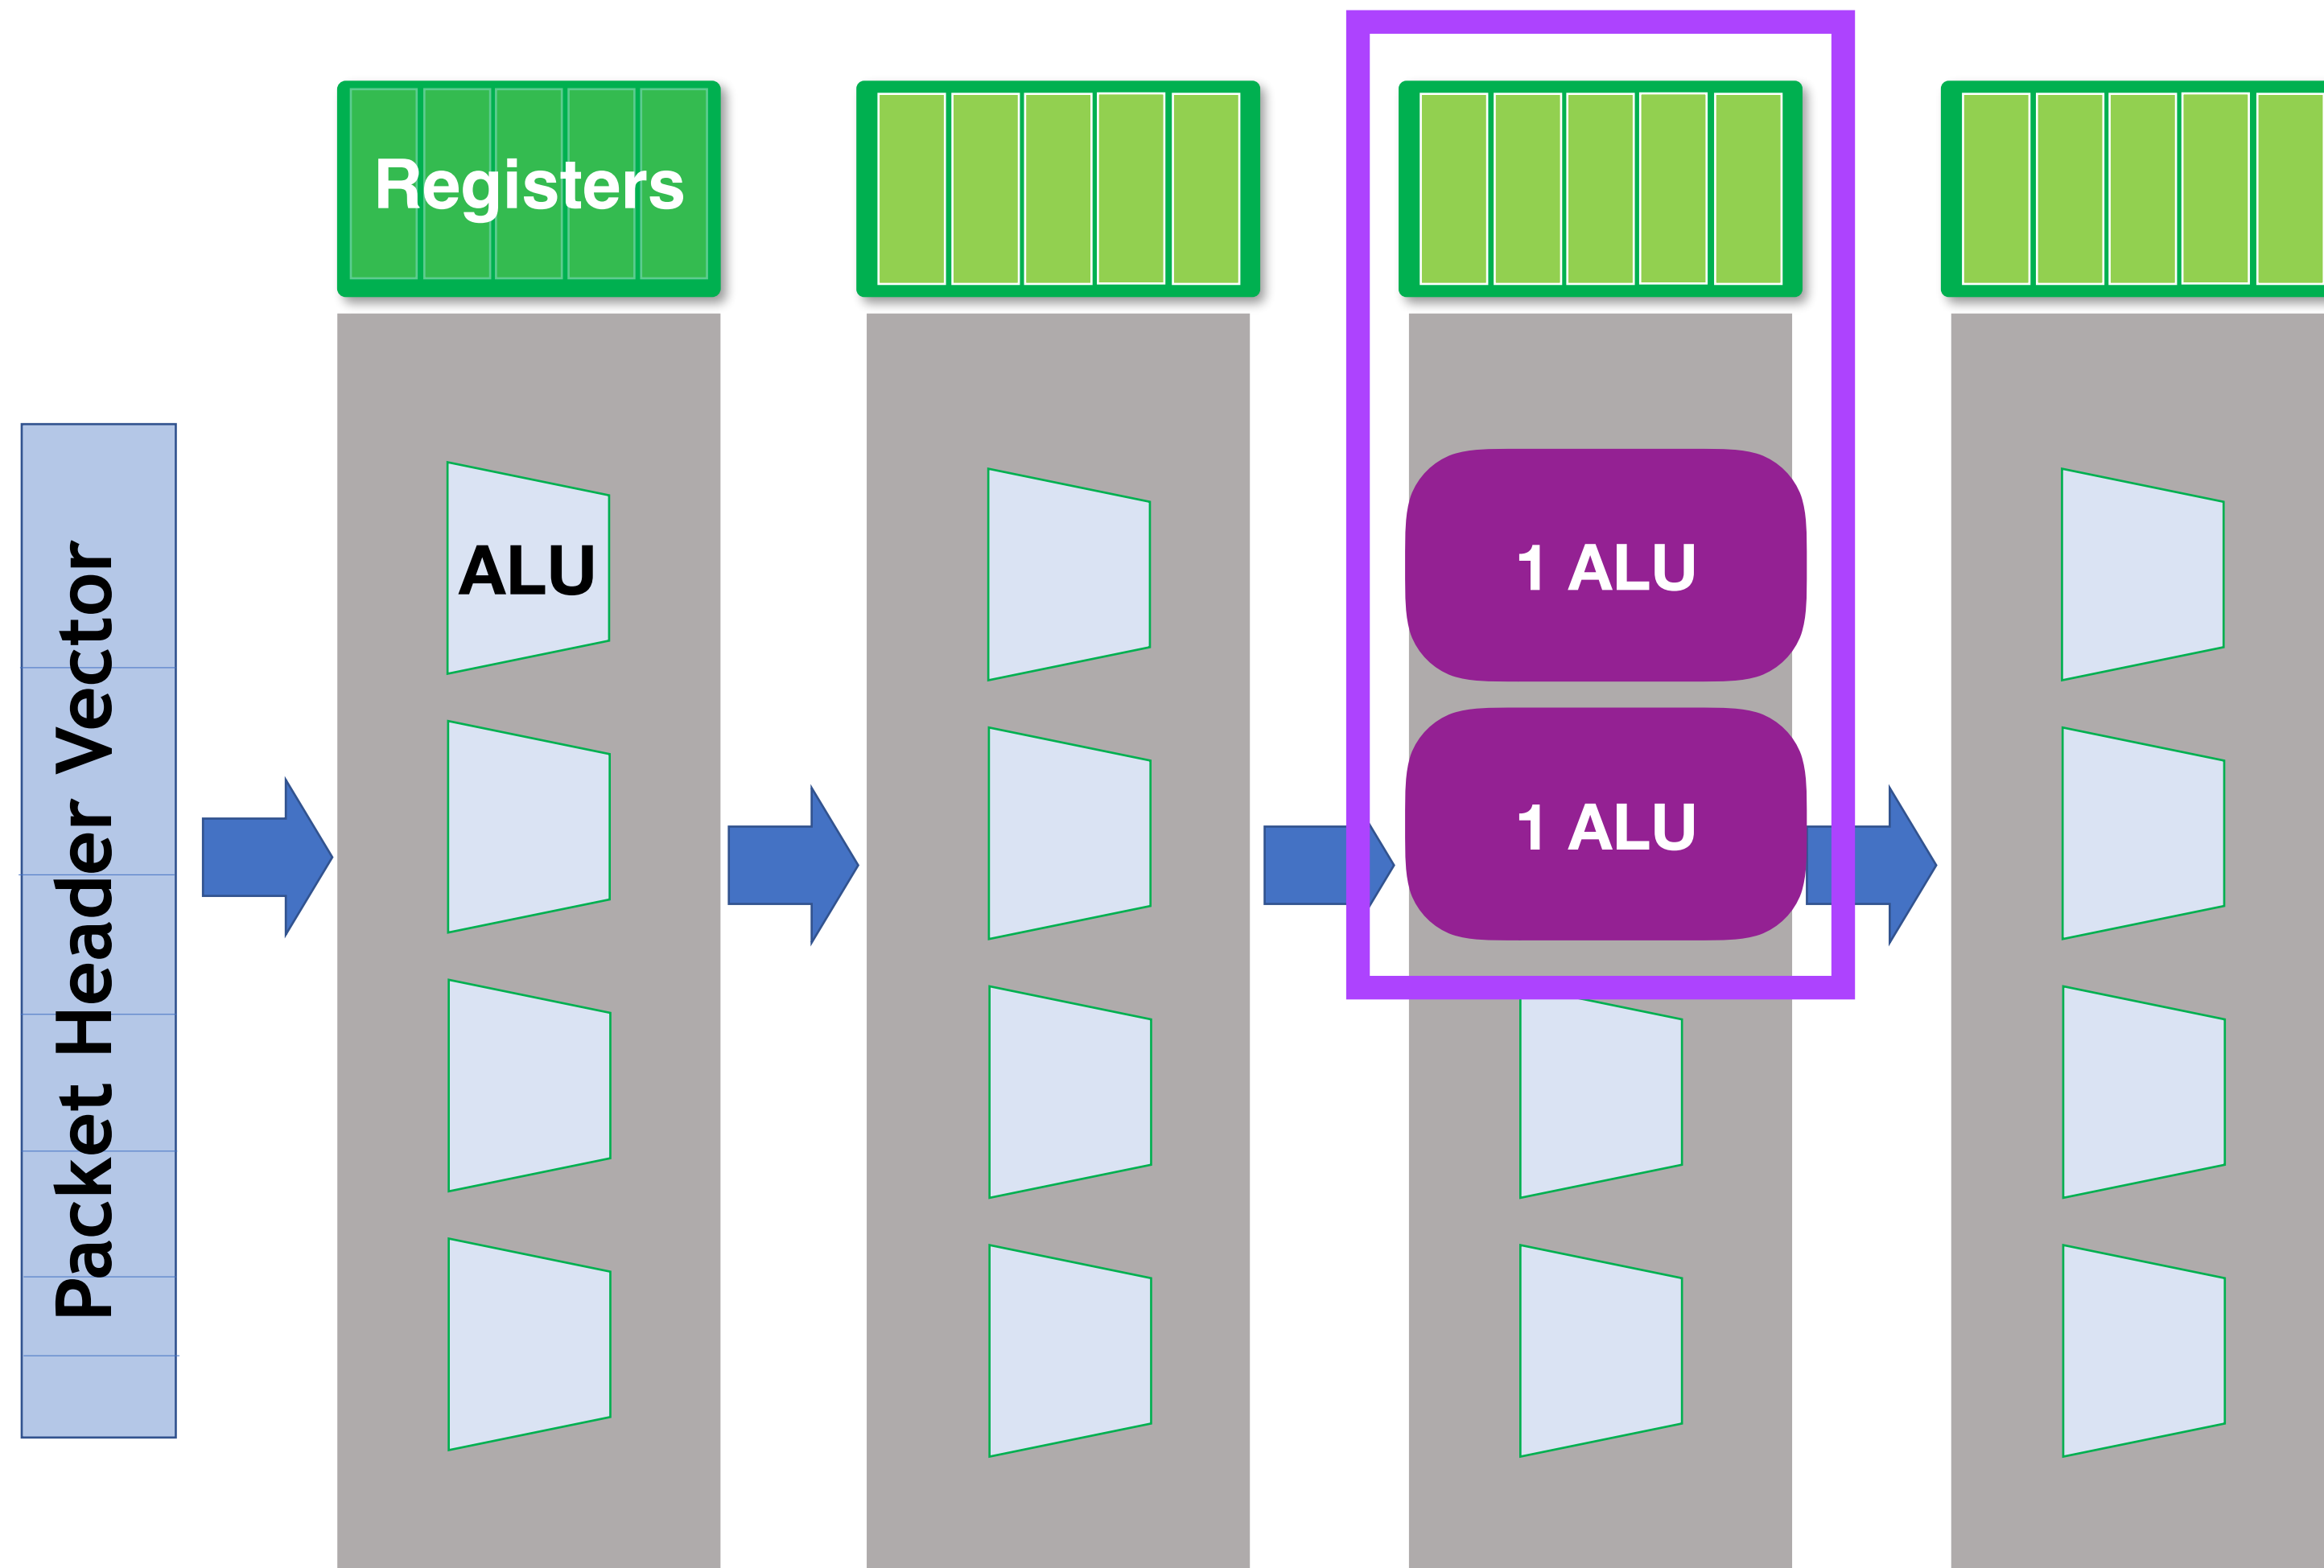
Pipeline Stages

# ILP Constraints



If action<sub>0</sub> and action<sub>1</sub> are dependent,  
Stage<sub>a0</sub> != Stage<sub>a1</sub>

# ILP Constraints

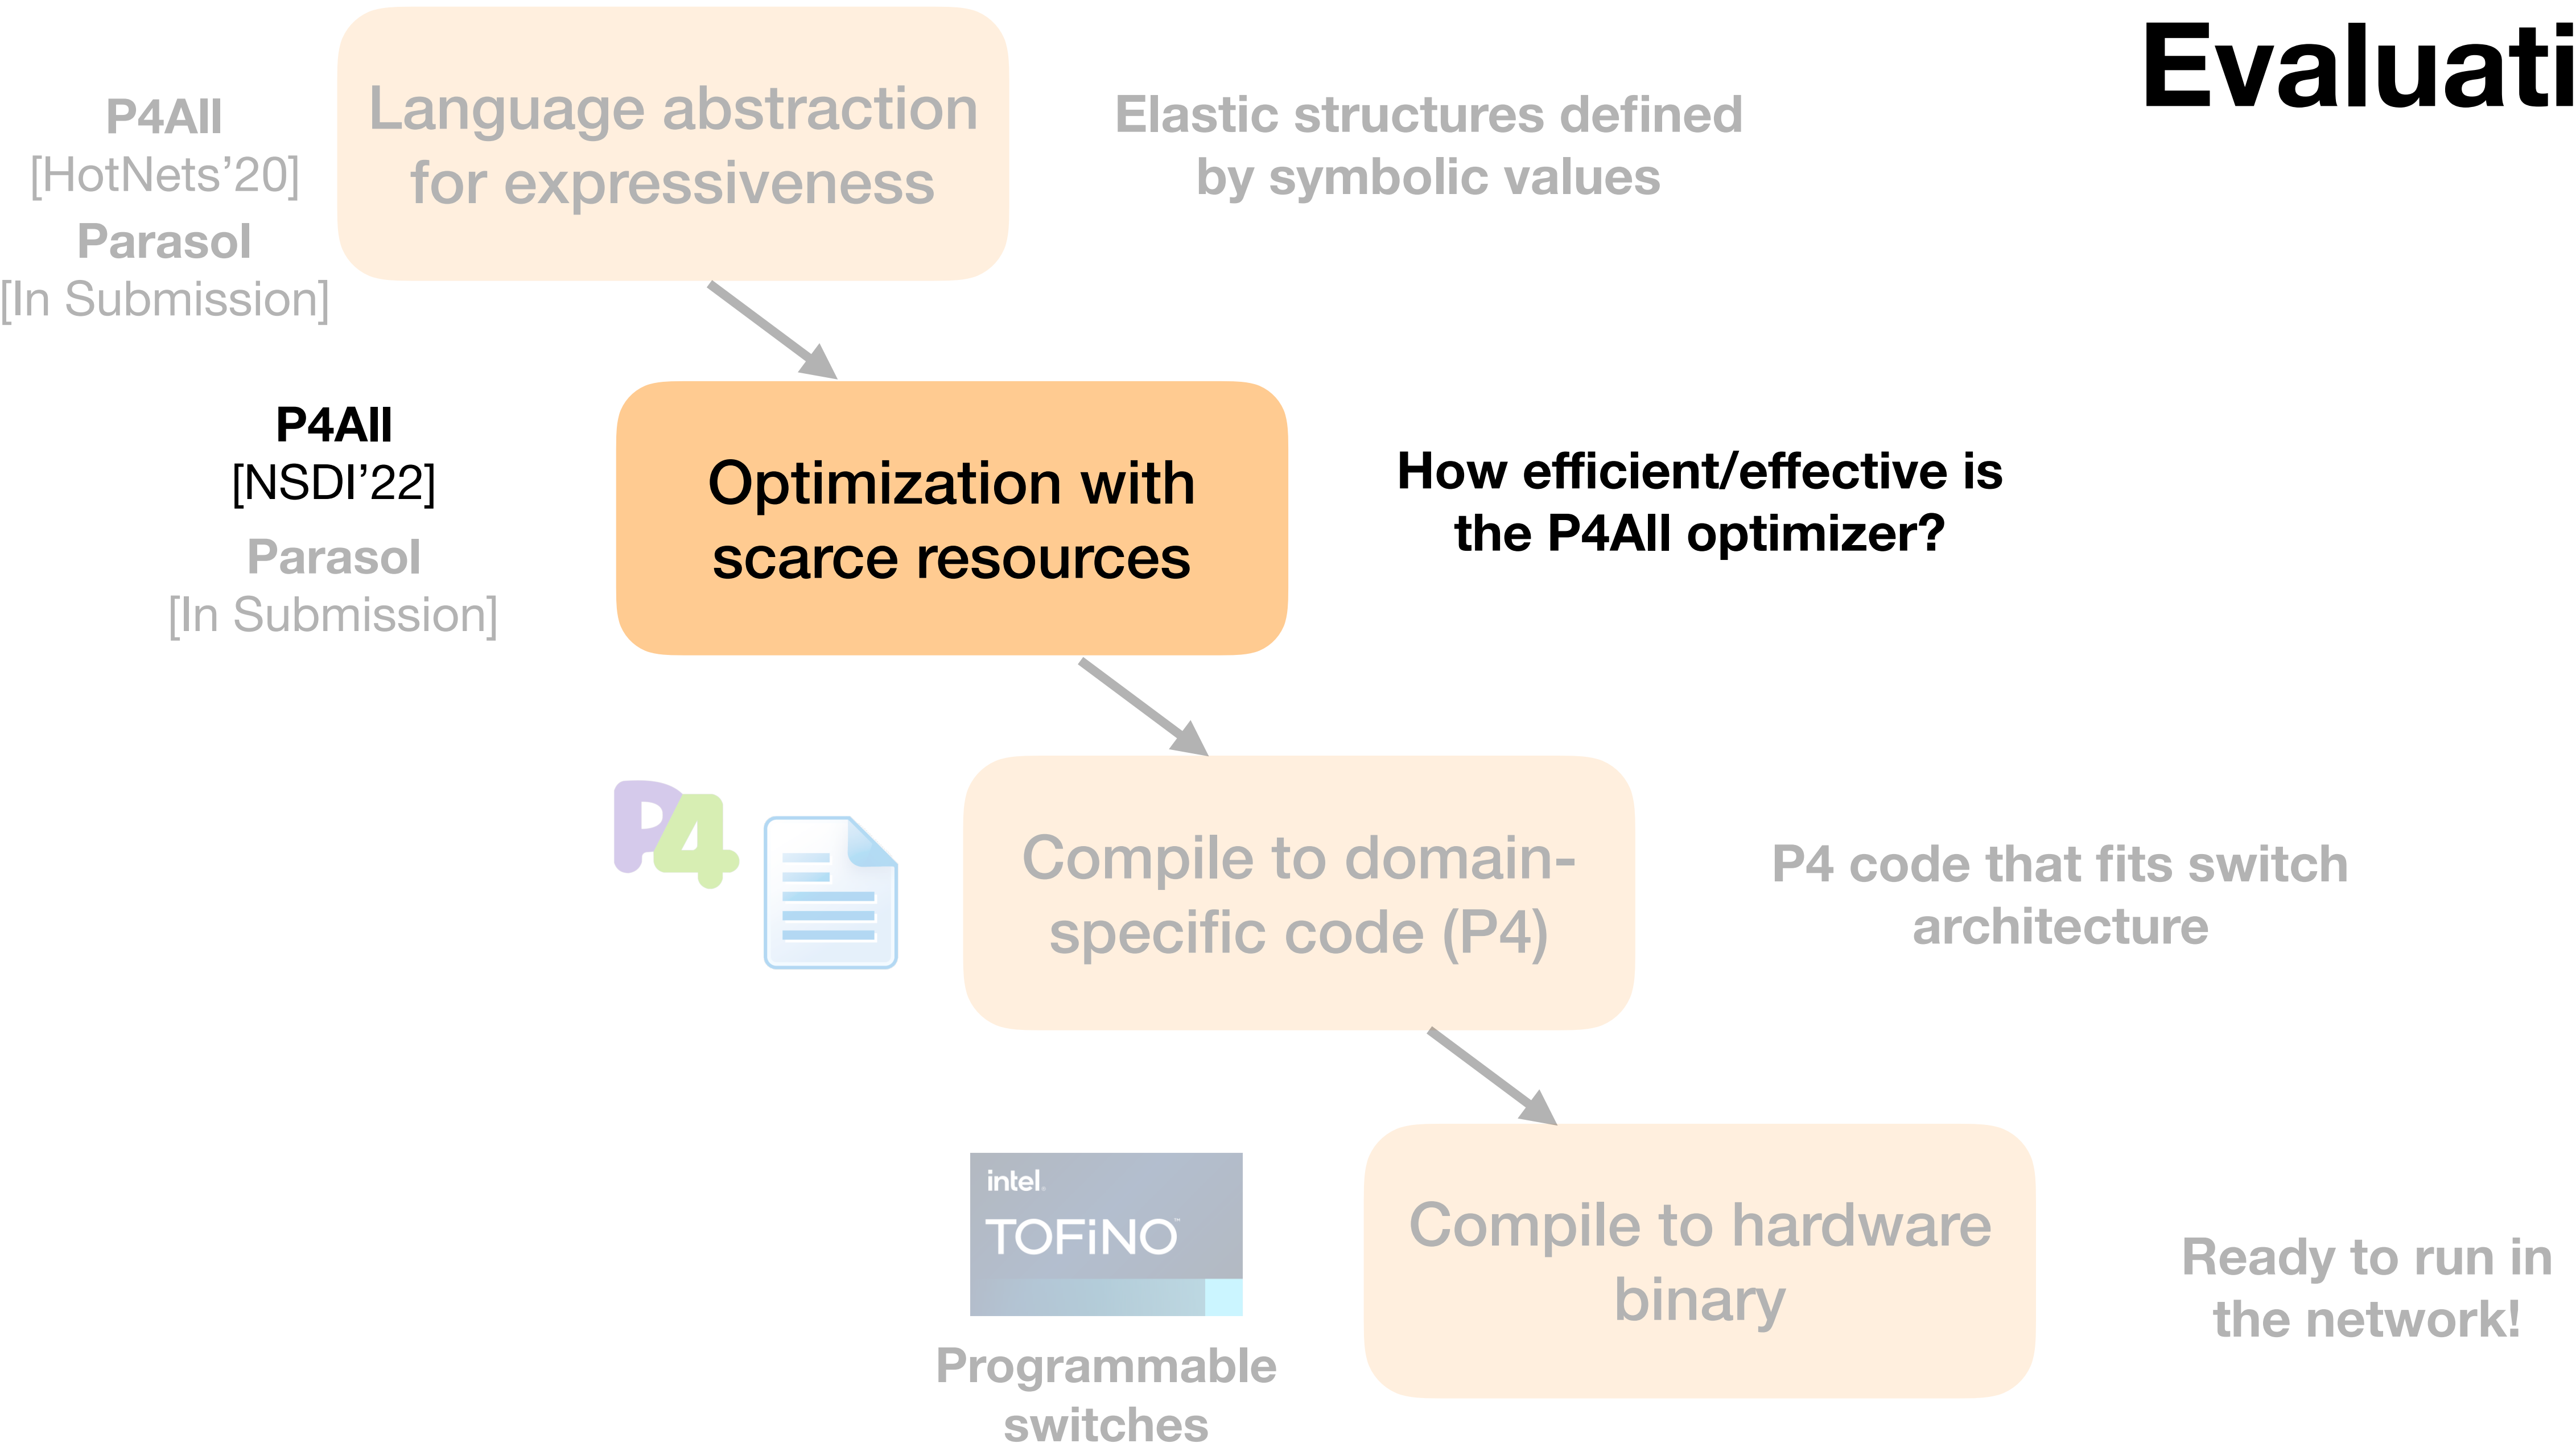


If action<sub>0</sub> and action<sub>1</sub> are dependent,  
Stage<sub>a0</sub> != Stage<sub>a1</sub>

If action<sub>0</sub> and action<sub>1</sub> access the same register arrays,  
Stage<sub>a0</sub> = Stage<sub>a1</sub>

Pipeline Stages

# Evaluation



# Optimizer Efficiency

<b>Application</b>	<b>Optimizer Time (s)</b>
CMS	
Key-value store	
Key-value store + CMS	
Switch.p4	
IP forwarding + stateful firewall	
Beaucoup	
Precision	
NetChain	
SketchLearn	
Conquest	
Bloom filter	

# Optimizer Efficiency

<b>Application</b>	<b>Optimizer Time (s)</b>
CMS	1.8
Key-value store	
Key-value store + CMS	
Switch.p4	0.2
IP forwarding + stateful firewall	0.4
Beaucoup	0.1
Precision	
NetChain	
SketchLearn	2.4
Conquest	5.8
Bloom filter	

# Optimizer Efficiency

<b>Application</b>	<b>Optimizer Time (s)</b>
CMS	1.8
Key-value store	15.4
Key-value store + CMS	27.9
Switch.p4	0.2
IP forwarding + stateful firewall	0.4
Beaucoup	0.1
Precision	25.7
NetChain	27.9
SketchLearn	2.4
Conquest	5.8
Bloom filter	513.6

# P4All vs Hand-Optimized Code

Application
CMS
Key-value store
Key-value store + CMS
Switch.p4
IP forwarding + stateful firewall
Beaucoup
Precision
NetChain
SketchLearn
Conquest
Bloom filter



**P4All generates same resource usage as hand-optimized code**



**P4All solution achieves marginally higher accuracy than hand-optimized code**



**P4All solution requires more resources**



# Does P4All Generalize?

P4All  
[NSDI'22]

**Objective function = performance as a function of structure size**



**Analytical functions provide bounds on worst-case performance**



**Parameters only relate to resource usage**



**ILP optimizes resource usage**



**Workload dependence information is limited**

**P4All works well when we have an analytical function and are only concerned with resource usage**



**Analytical functions not always possible**

# Optimization Objective

**P4All**  
[NSDI'22]

**Objective function = performance as a function of structure size**

**Parasol**  
[In Submission]

**Objective function = run-time performance measurements**

# Parasol

Parasol  
[In Submission]

Language abstraction  
for expressiveness

Separates what to believe about details  
from any parameter code

Parasol  
[In Submission]

Optimization with  
scarce resources

Simulation-based optimization  
Automatically tailors  
finds parameter values for  
program to environment  
optimal performance

Optimize any program parameter using  
empirical performance objectives

# Extended Symbolics

$$\text{Objective (P4All)} = \sum_{i=1}^{kv\_size} \frac{1}{i} - \frac{3}{cms\_cols}$$

## Cache.p4

```
symbolic int kv_size;  
symbolic int cms_rows;  
symbolic int cms_cols;  
def cms_hash () {  
  for (i < cms_rows) {  
    hash_to_row(i);  
  }  
}
```

When should cache entries expire?

When is a key popular enough to go into the cache?

Is CMS the right data structure?

Non-resource parameters

# Extended Symbolics

$$\text{Objective (Parasol)} = \frac{\text{number of cache hits}}{\text{number of cache requests}}$$

When should cache entries expire?  
Cache.p4

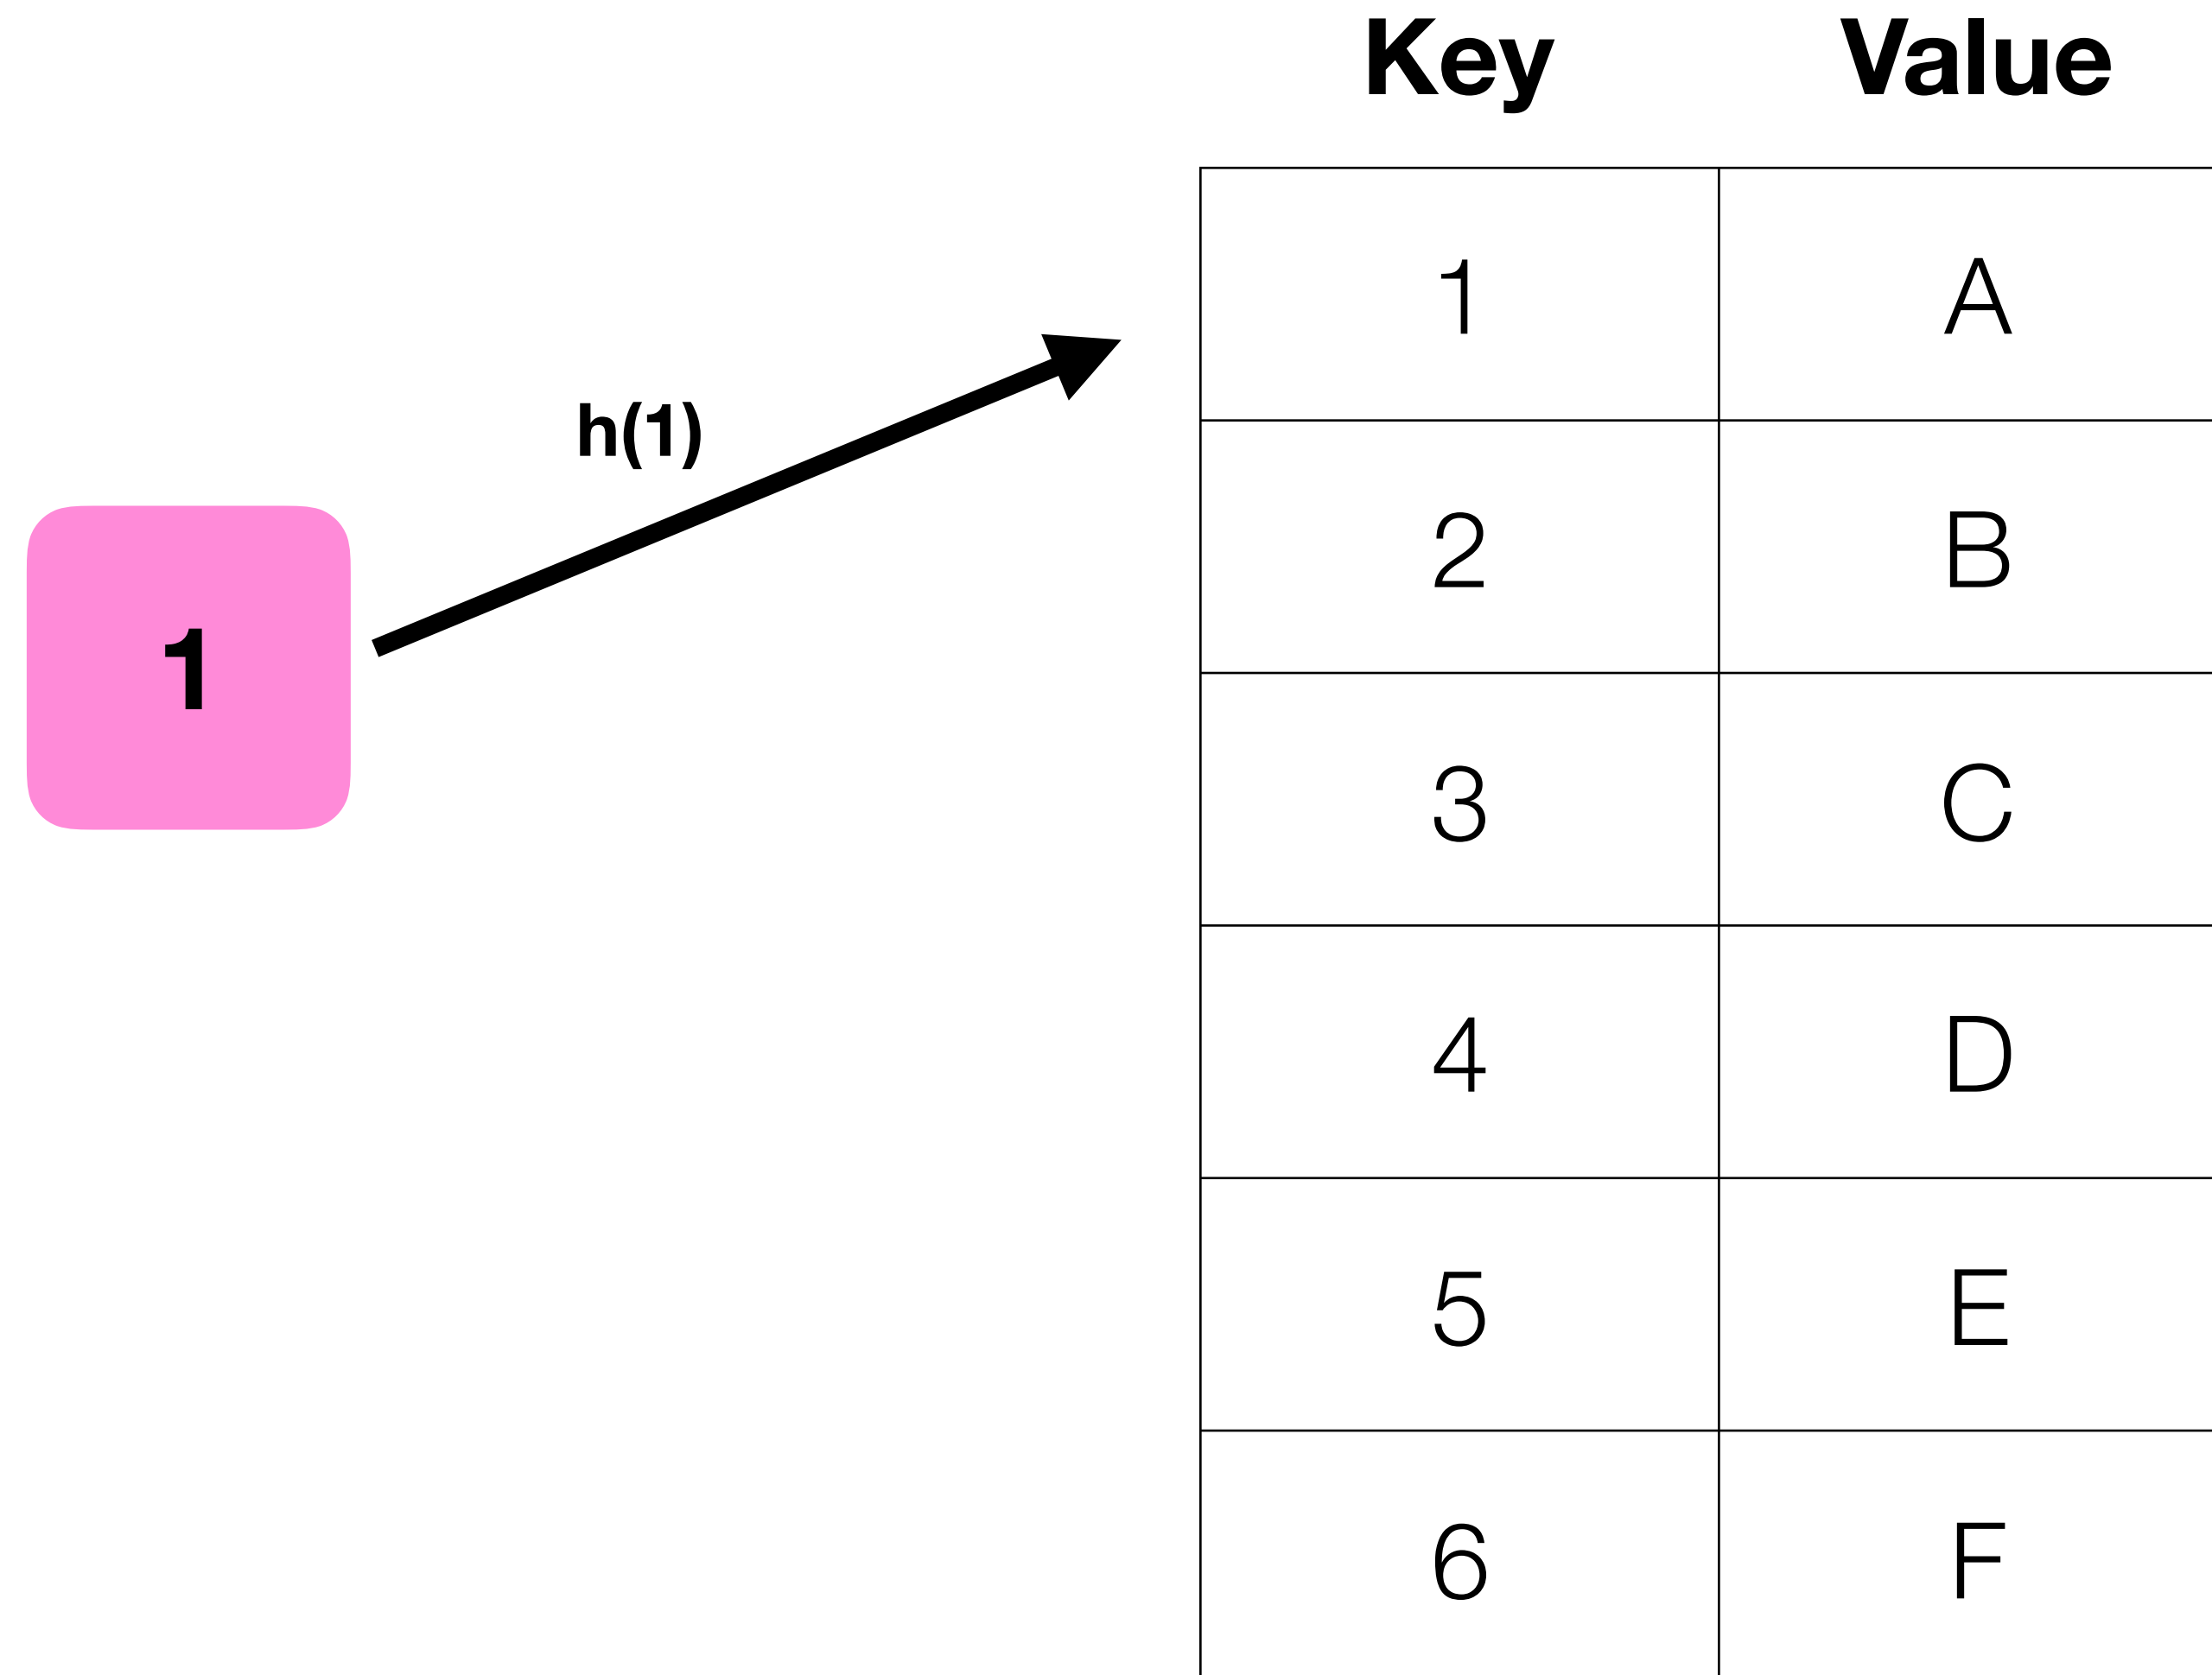
```
symbolic int kv_size;  
symbolic int cms_rows;  
symbolic int cms_cols;  
def cms_hash () {  
  for (i < cms_rows) {  
    hash_to_row(i);  
  }  
}
```

When is a key popular enough to go into the cache?  
Is CMS the right data structure?

Cache.p4 (Cache.dpt)

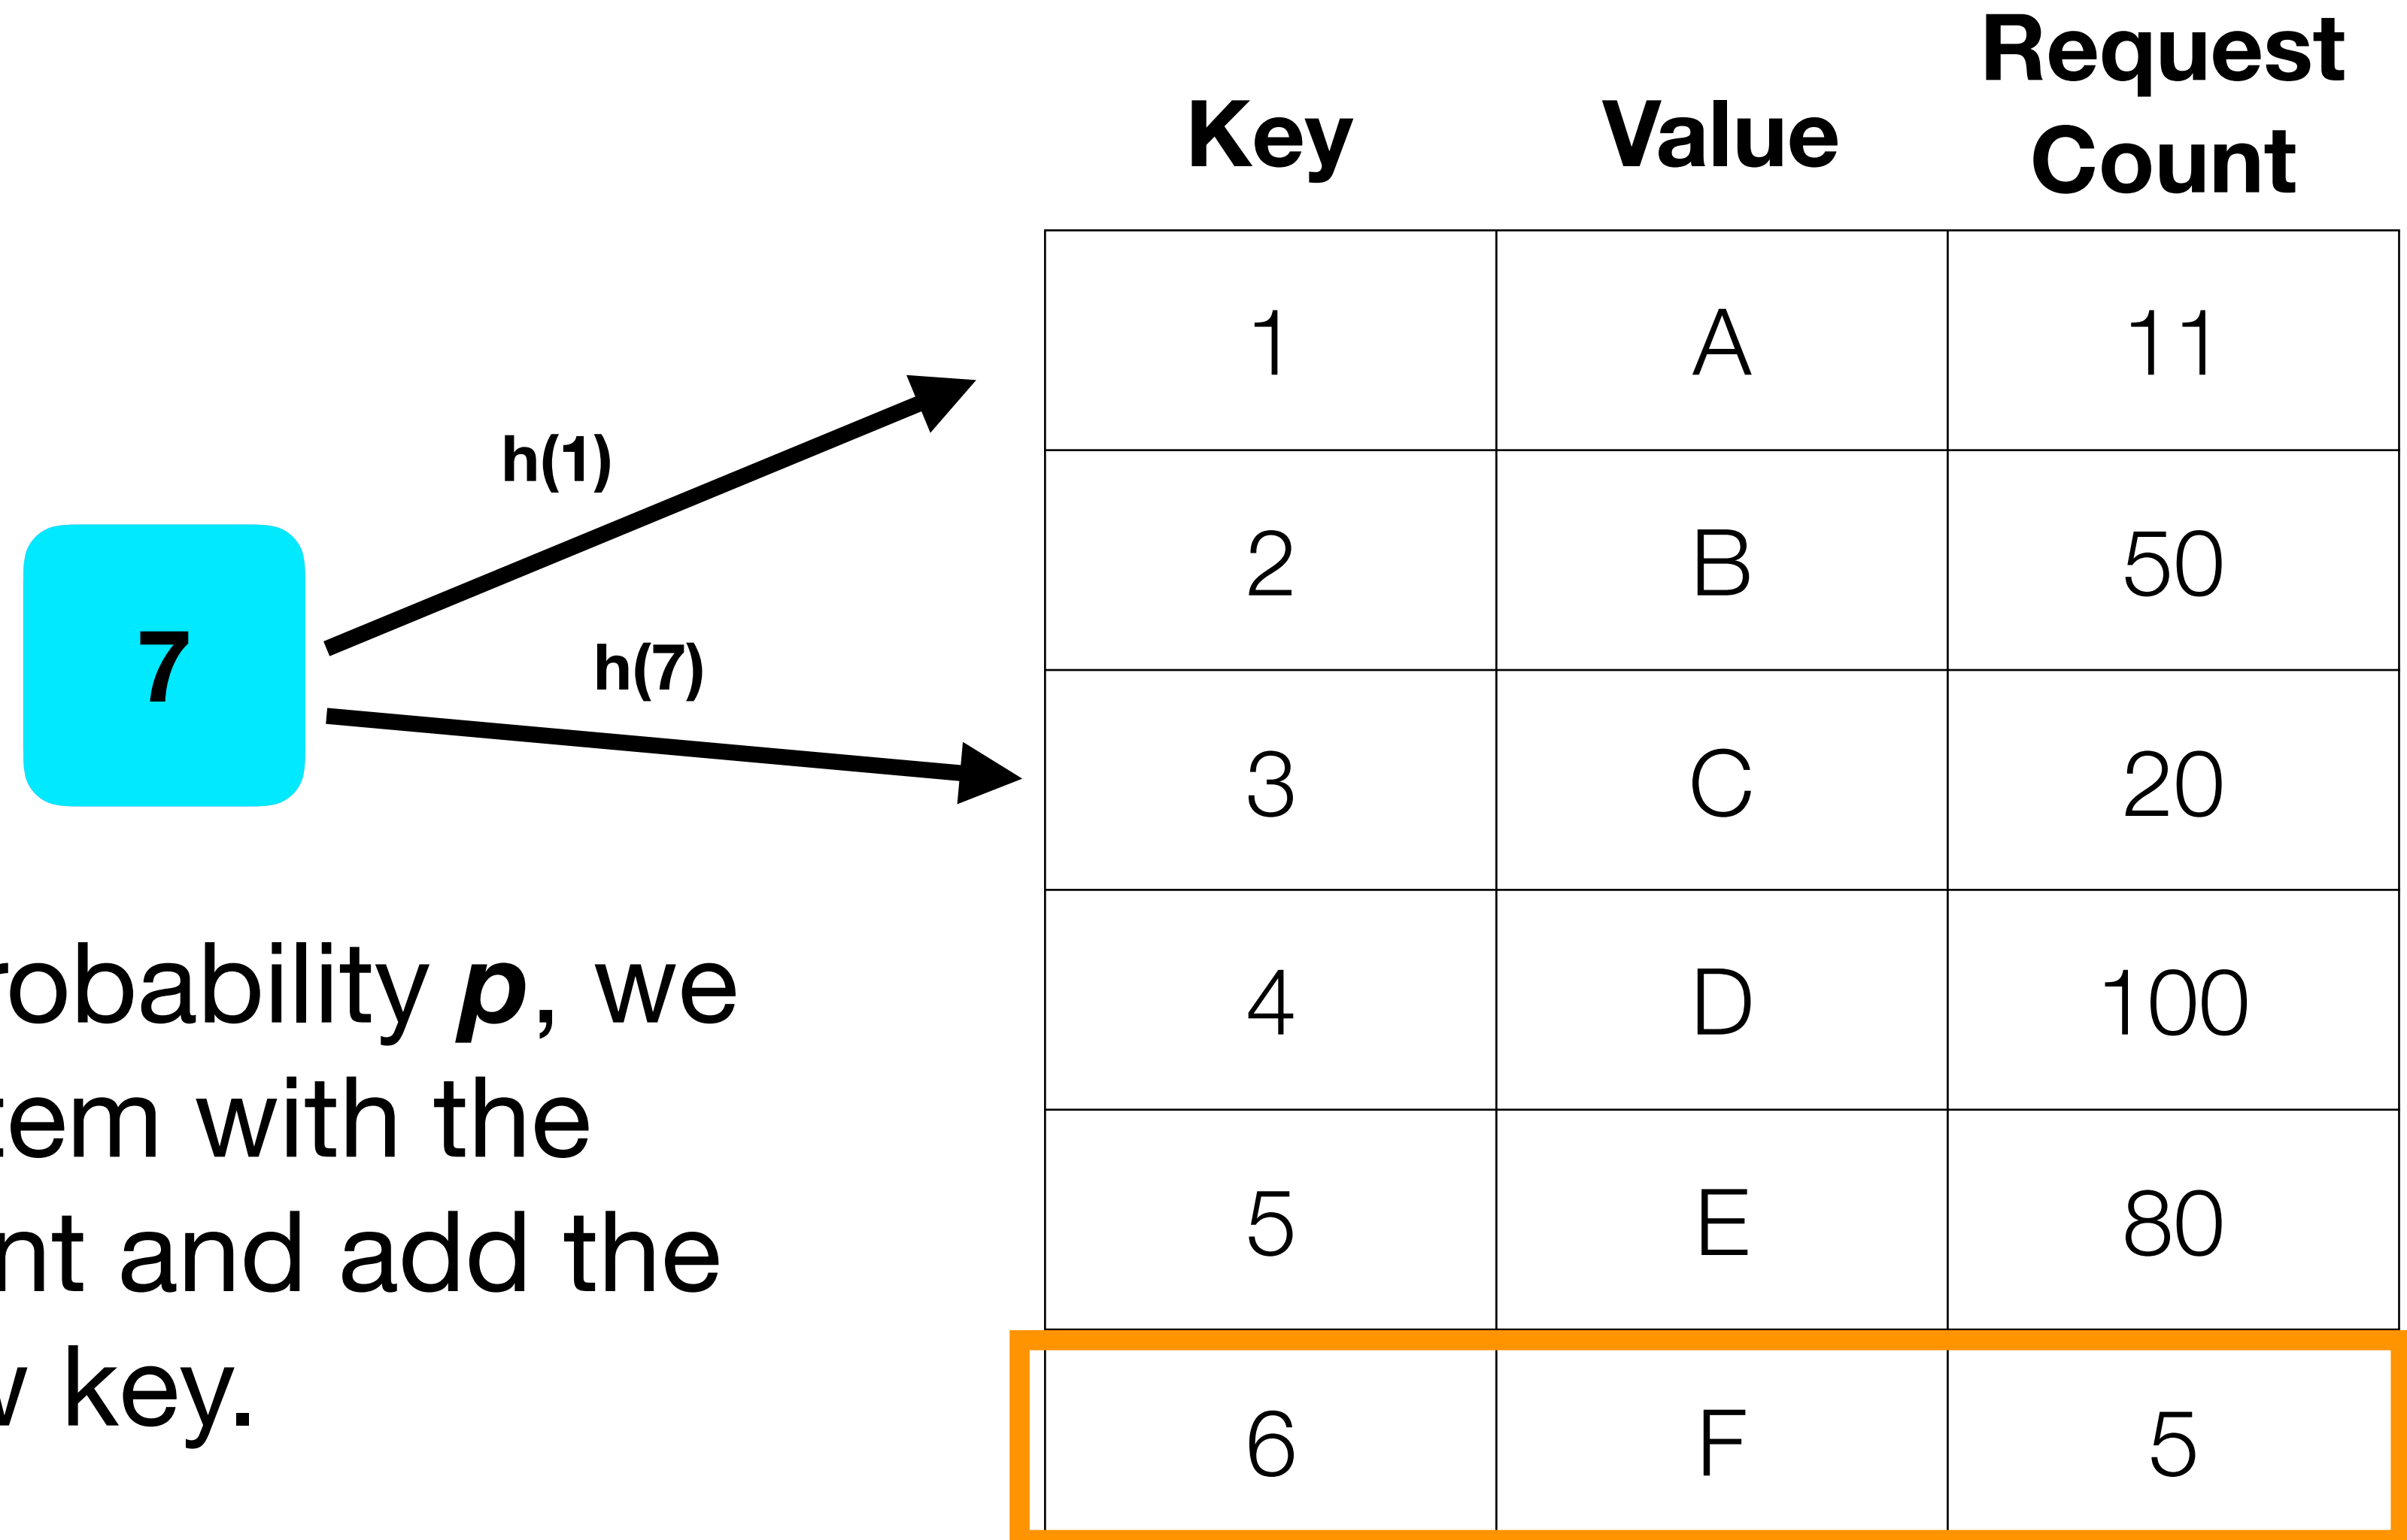
```
symbolic int kv_size;  
symbolic int cms_rows;  
symbolic int cms_cols;  
symbolic int timeout;  
symbolic int threshold;  
def cms_hash () {  
  for (i < cms_rows) {  
    hash_to_row(i);  
  }  
}
```

# Precision Cache



Efficient Measurement on Programmable Switches, *Basat et al.* [ICNP'18]

# Precision Cache



With some probability  $p$ , we evict the item with the smallest count and add the new key.

# Choice of Data Structure

Is CMS the right data structure?

Cache.p4

```
symbolic int kv_size;
symbolic int tracker_rows;
symbolic int tracker_cols;
symbolic int timeout;
symbolic int threshold;
symbolic bool cms_tracker;

module CMS {...}
module Precision {...}
tracker = CMS if cms_tracker else Precision;
```



# The System

P4All  
[HotNets'20]  
Parasol  
[In Submission]

Language abstraction  
for expressiveness

Extend symbolic values to  
*any* parameter

P4All  
[NSDI'22]  
Parasol  
[In Submission]

Optimization with  
scarce resources

Simulate and measure  
performance



Compile to domain-  
specific code (P4)

P4 code that fits switch  
architecture



Compile to hardware  
binary

Ready to run in  
the network!

# Objective Functions

$$\text{Hit rate} = \frac{\text{number of cache hits}}{\text{number of cache requests}}$$

# Extern Functions

$$\text{Hit rate} = \frac{\text{number of cache hits}}{\text{number of cache requests}}$$

## Cache.p4

```
if (key in cache) {  
    LogCacheHit();  
}  
LogCacheRequest();
```

Externs are like “ghost code” that run ONLY during simulation

# Extern Functions

$$\text{Hit rate} = \frac{\text{number of cache hits}}{\text{number of cache requests}}$$

## Cache.p4

```
if (key in cache) {  
    LogCacheHit();  
}  
LogCacheRequest();
```

## Measurements.py

```
requests = 0  
hits = 0  
def LogCacheHit():  
    hits += 1  
def LogCacheRequest():  
    requests += 1
```

# Scoring Performance

## Cache.p4

```
if (key in cache) {  
    LogCacheHit();  
}  
LogCacheRequest();
```

## HitRate.py

```
def CacheScore(hits, requests):  
    return hits / requests
```

## Measurements.py

```
requests = 0  
hits = 0  
def LogCacheHit():  
    hits += 1  
def LogCacheRequest():  
    requests += 1
```

vs

$$\sum_{i=1}^{kv\_size} \frac{1}{i^\alpha}$$

# The System

P4All  
[HotNets'20]

Language abstraction  
for expressiveness

Separate low-level details  
from program code

Choose parameters to  
simulate

Simulate and measure  
performance

P4 code that fits switch  
architecture

Ready to run in  
the network!


P4All  
[NSDI'22]

Parasol  
[In Submission]

Optimization with  
scarce resources

Compile to domain-  
specific code (P4)

Compile to hardware  
binary

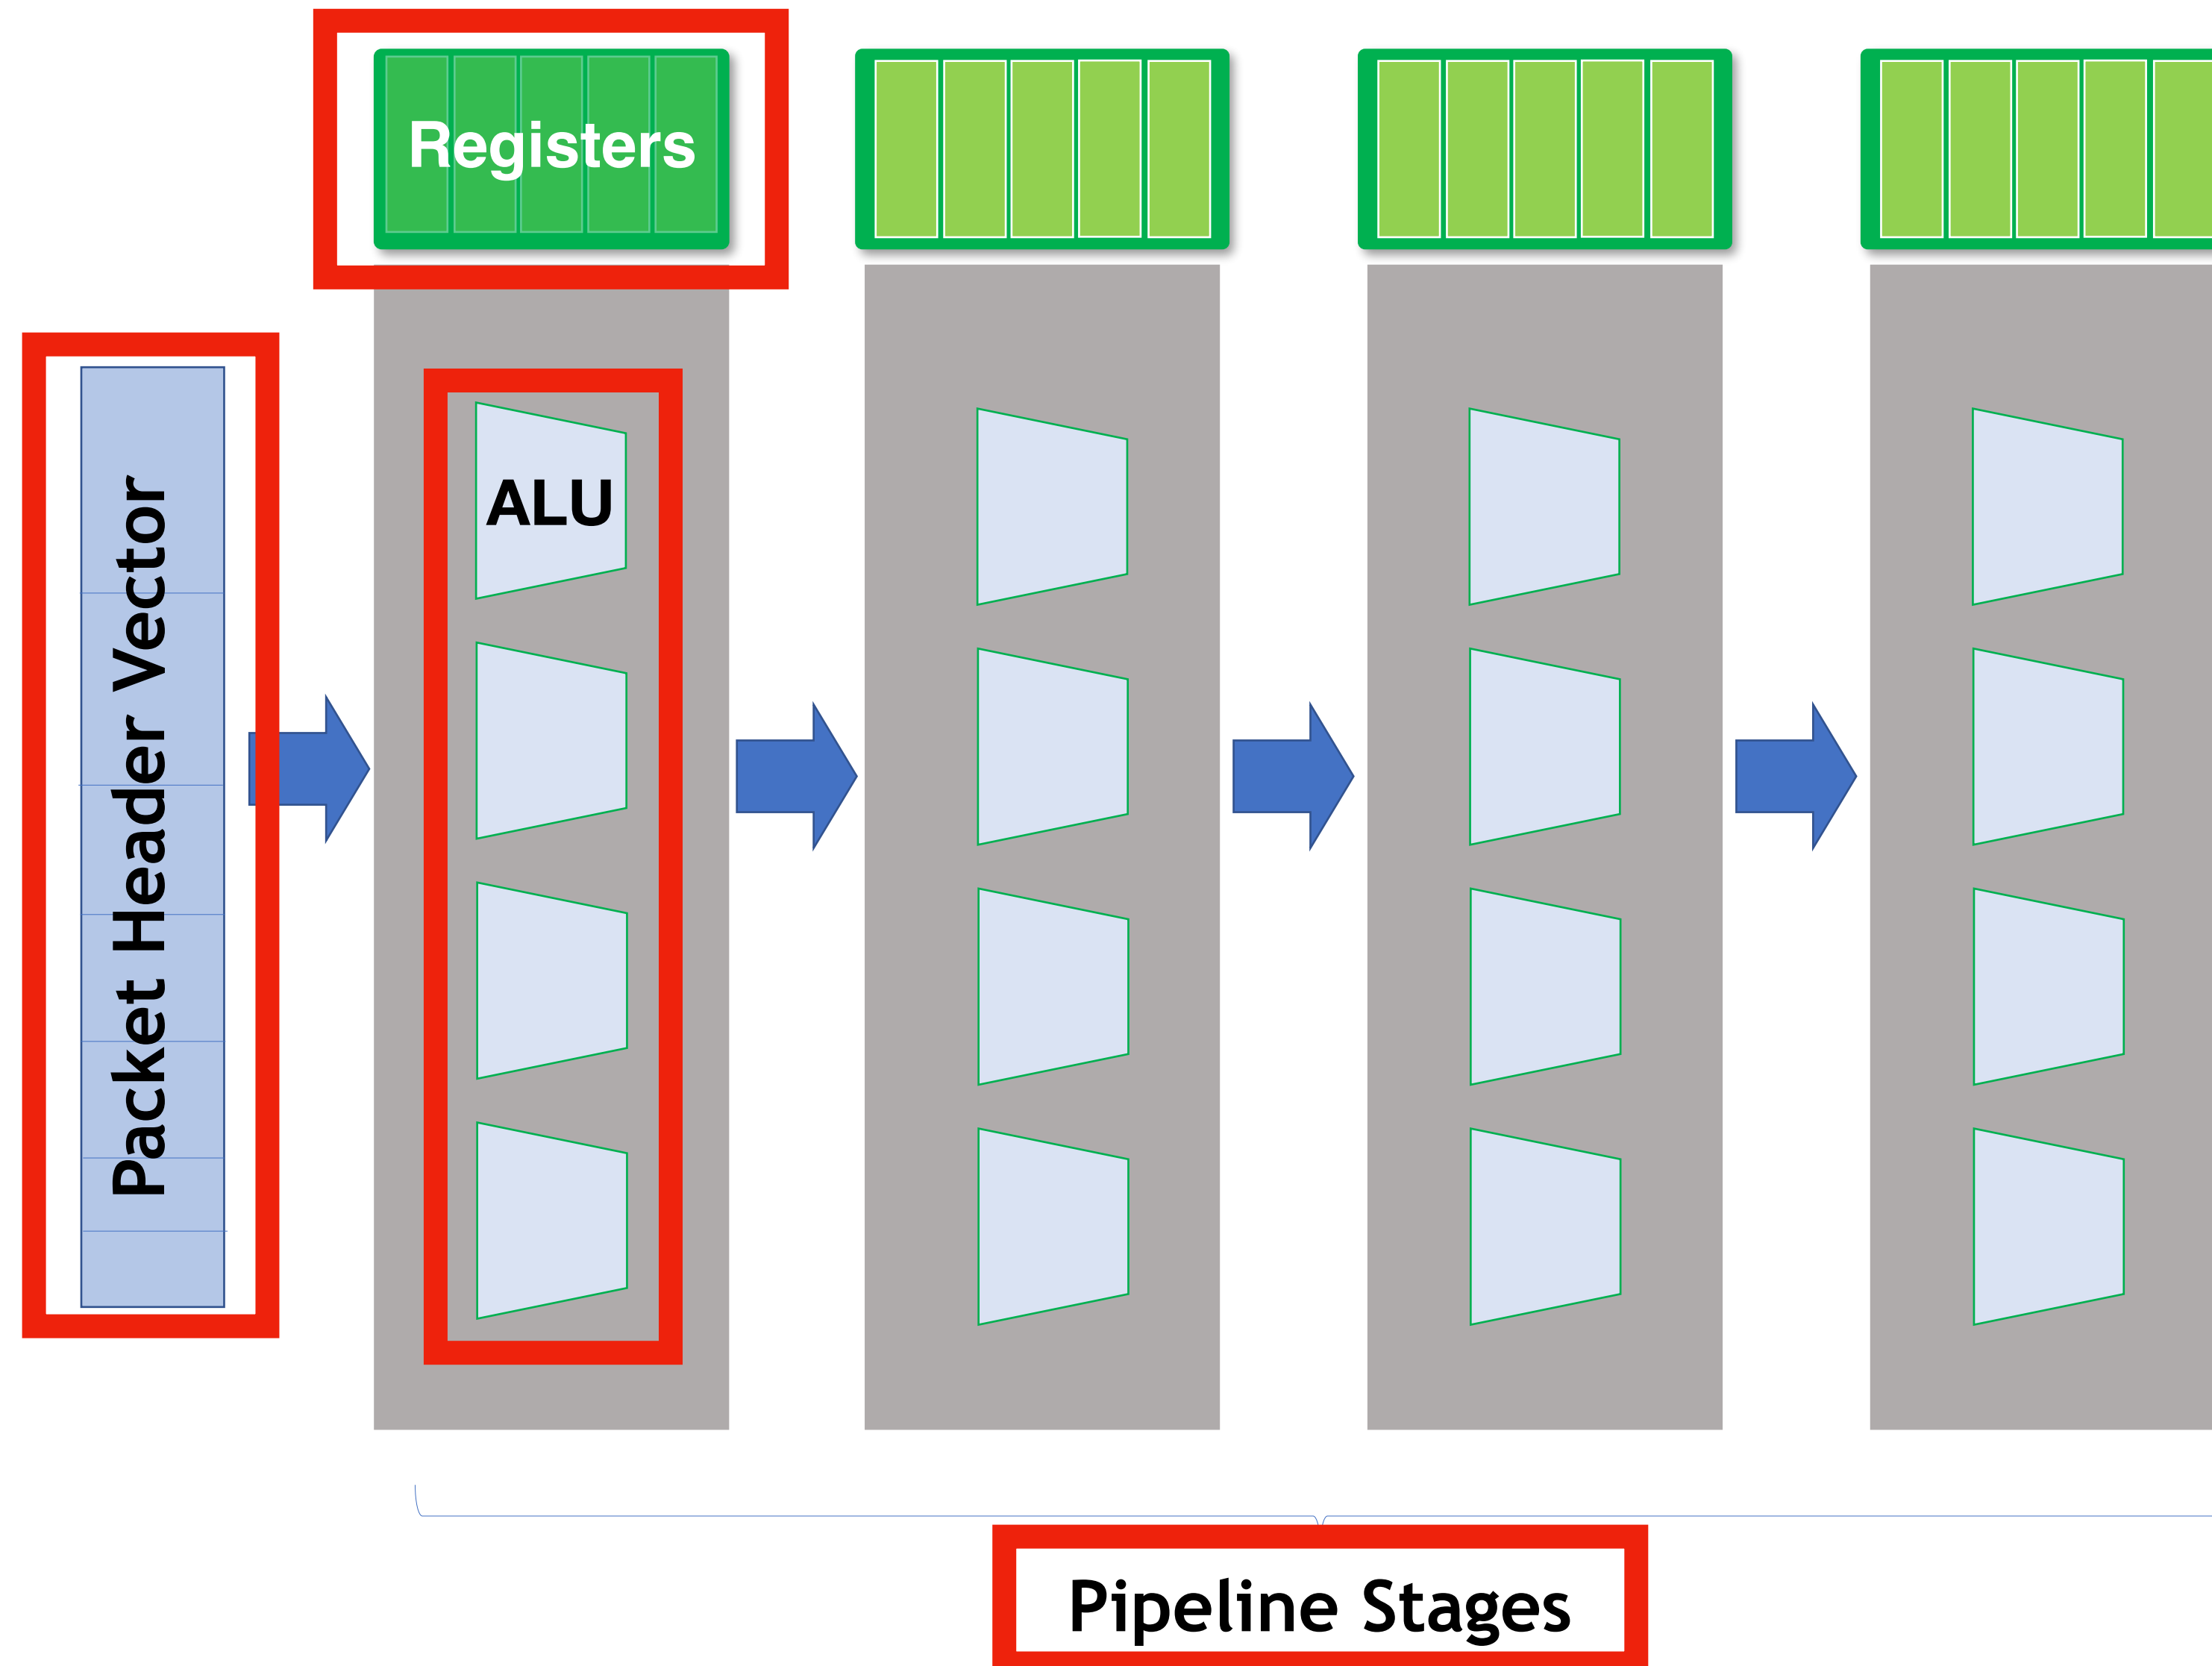


Programmable  
switches

# Choosing Parameter Values

```
symbolic int kv_size;  
symbolic int tracker_rows;  
symbolic int tracker_cols;  
symbolic int timeout;  
symbolic int threshold;  
symbolic bool cms_tracker;
```

# Resource Constraints






# Resources Limit Parameter Values

```
symbolic int kv_size;  
symbolic int tracker_rows;  
symbolic int tracker_cols;
```

We use heuristics to estimate resource usage



Compilation fails if program uses too many resources



We can compile\* the program (with chosen allocation) to see if it fits within resource constraints.

**We get upper bounds on all resource-related symbolic values.**

# The System

P4AI  
[HotNets'20]

Language abstraction  
for expressiveness

Separate low-level details  
from program code

Remove invalid parameter  
choices

P4AI  
[NSDI'22]  
**Parasol**  
[In Submission]

Optimization with  
scarce resources

Choose parameters to  
simulate

Simulate and measure  
performance



Compile to domain-  
specific code (P4)

P4 code that fits switch  
architecture



Compile to hardware  
binary

Ready to run in  
the network!

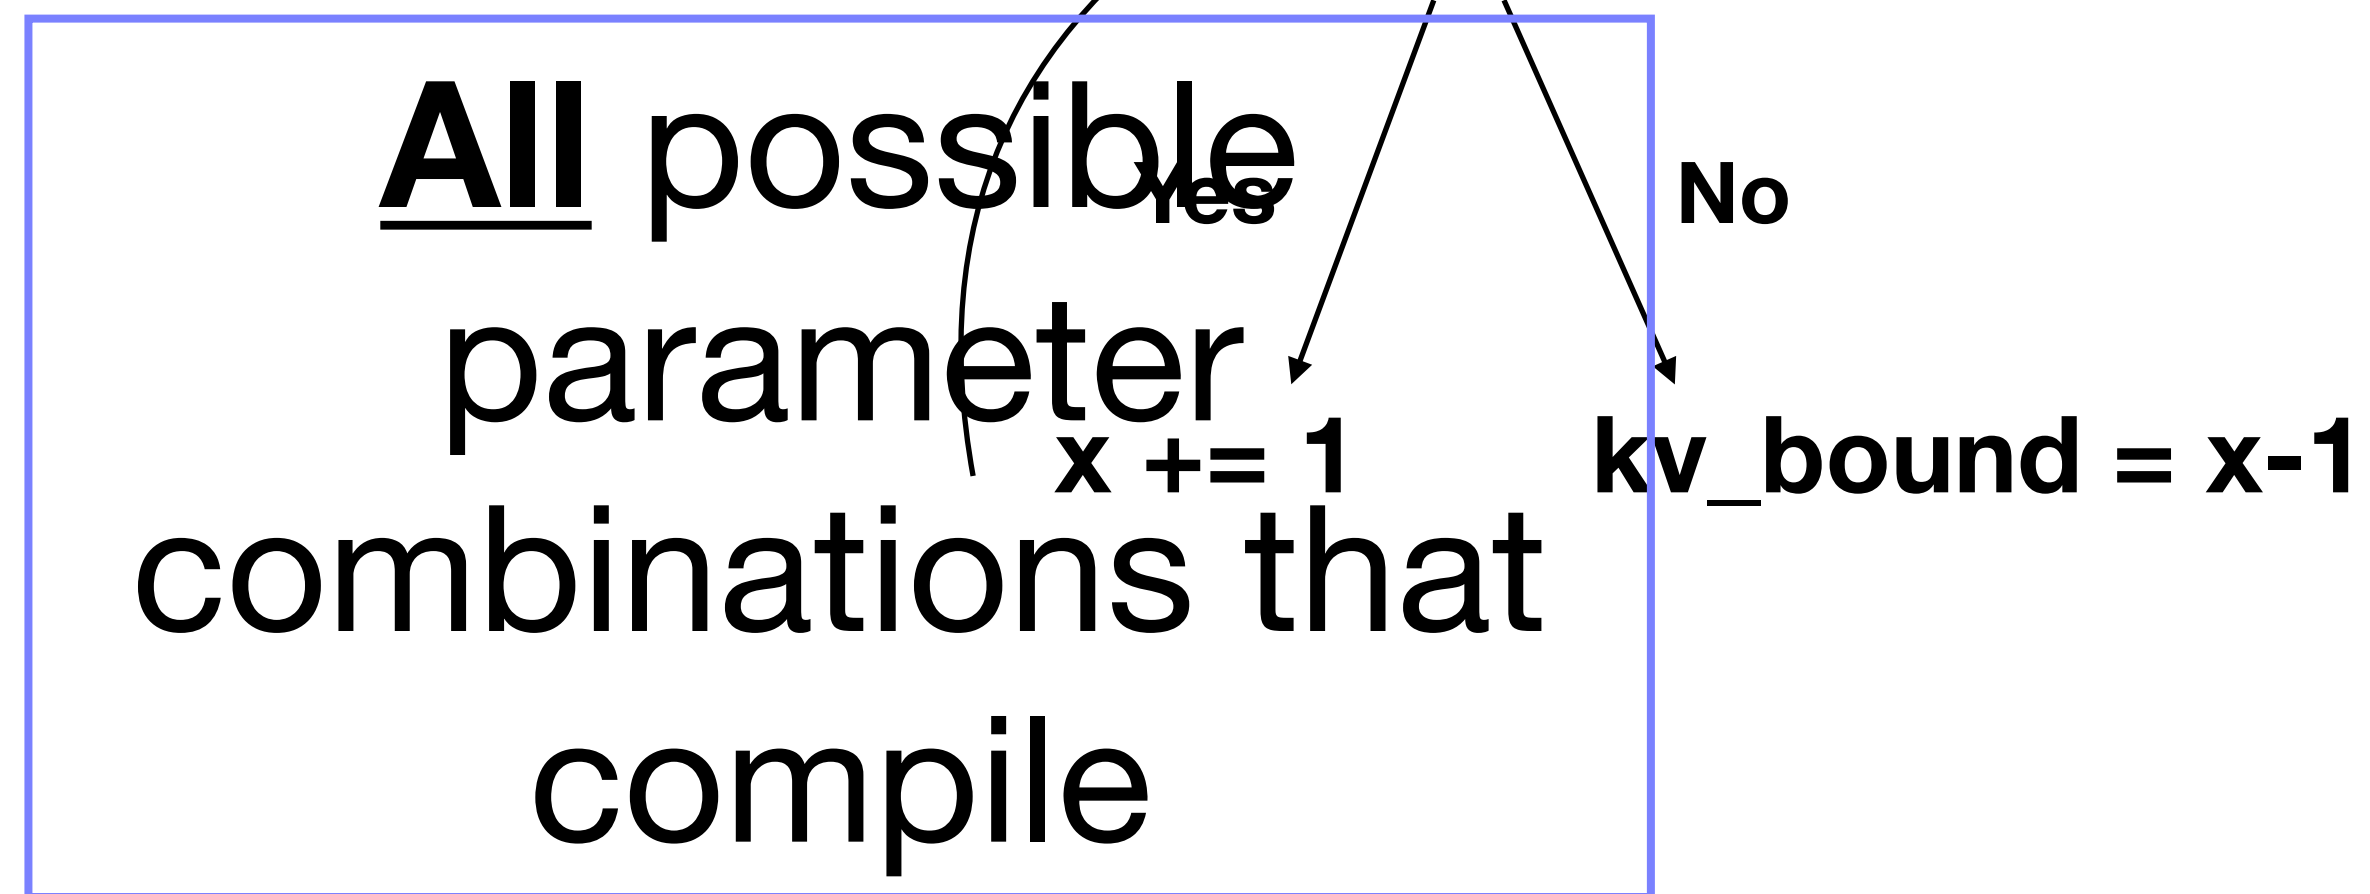
# Preprocessing

```
symbolic int kv_size;  
symbolic int tracker_rows;  
symbolic int tracker_cols;
```

kv\_size = x

tracker\_rows = 1  
tracker\_cols = 1

compile?



kv\_size : {1, ..., kv\_bound}

tracker\_rows : {1, ..., rows\_bound}

tracker\_cols : {1, ..., cols\_bound}

# The System

P4AI  
[HotNets'20]

Language abstraction  
for expressiveness

Separate low-level details  
from program code

Remove invalid parameter  
choices

P4AI  
[NSDI'22]  
**Parasol**  
[In Submission]

Optimization with  
scarce resources

Choose parameters to  
simulate

Simulate and measure  
performance



Compile to domain-  
specific code (P4)

P4 code that fits switch  
architecture



Compile to hardware  
binary

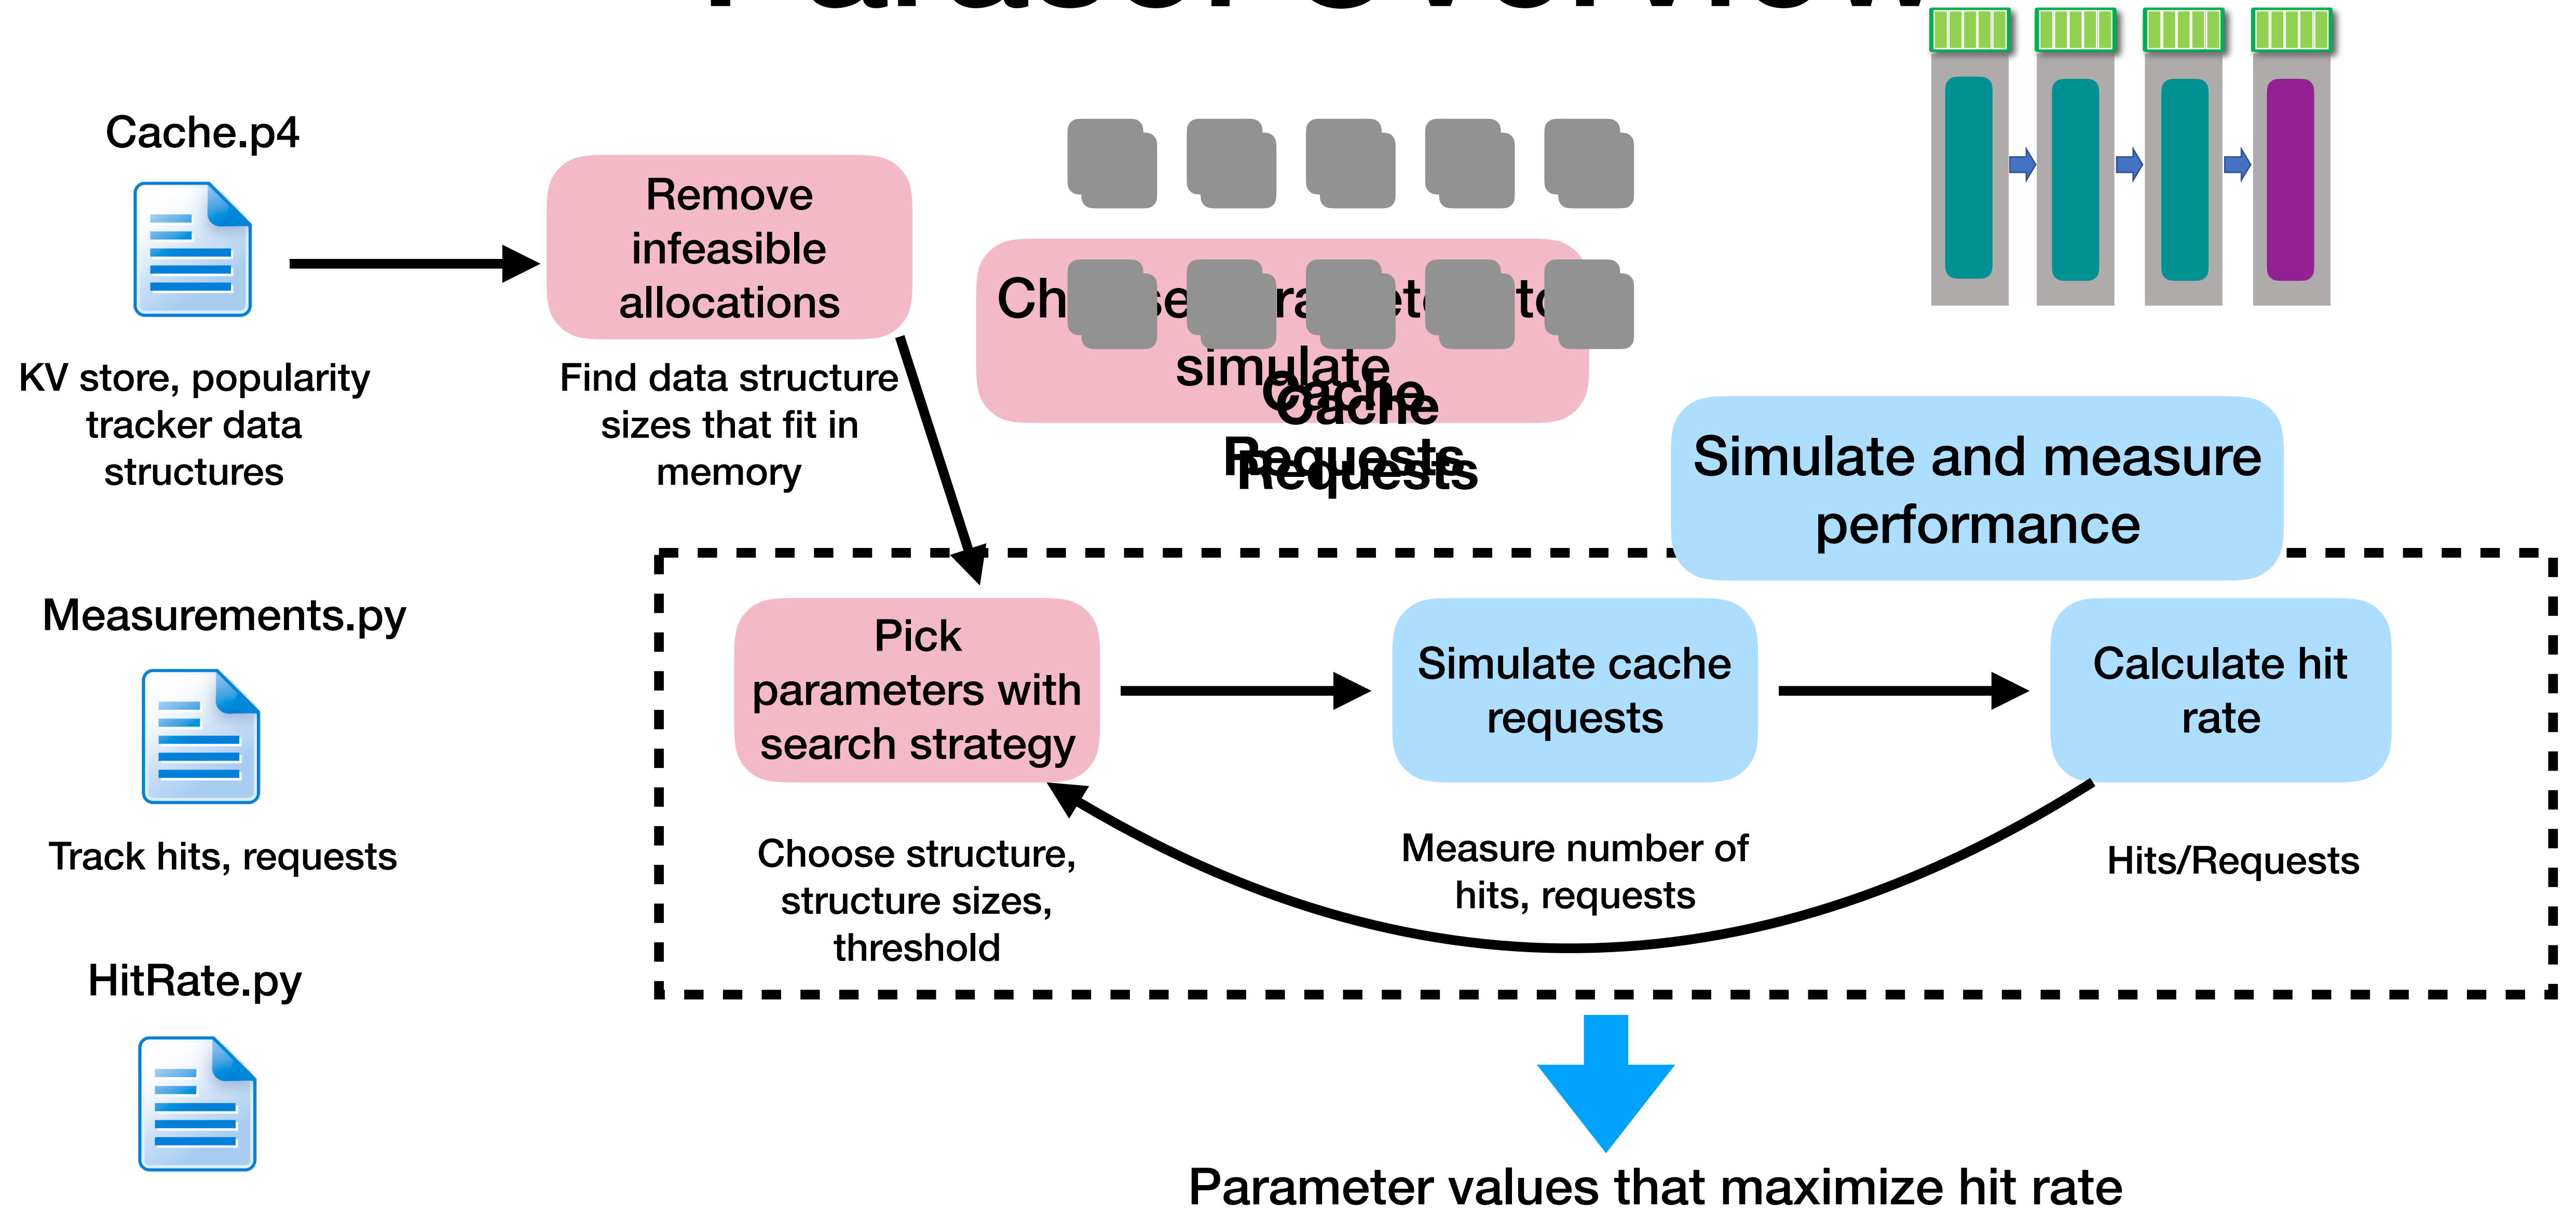
Ready to run in  
the network!

# Search Strategies

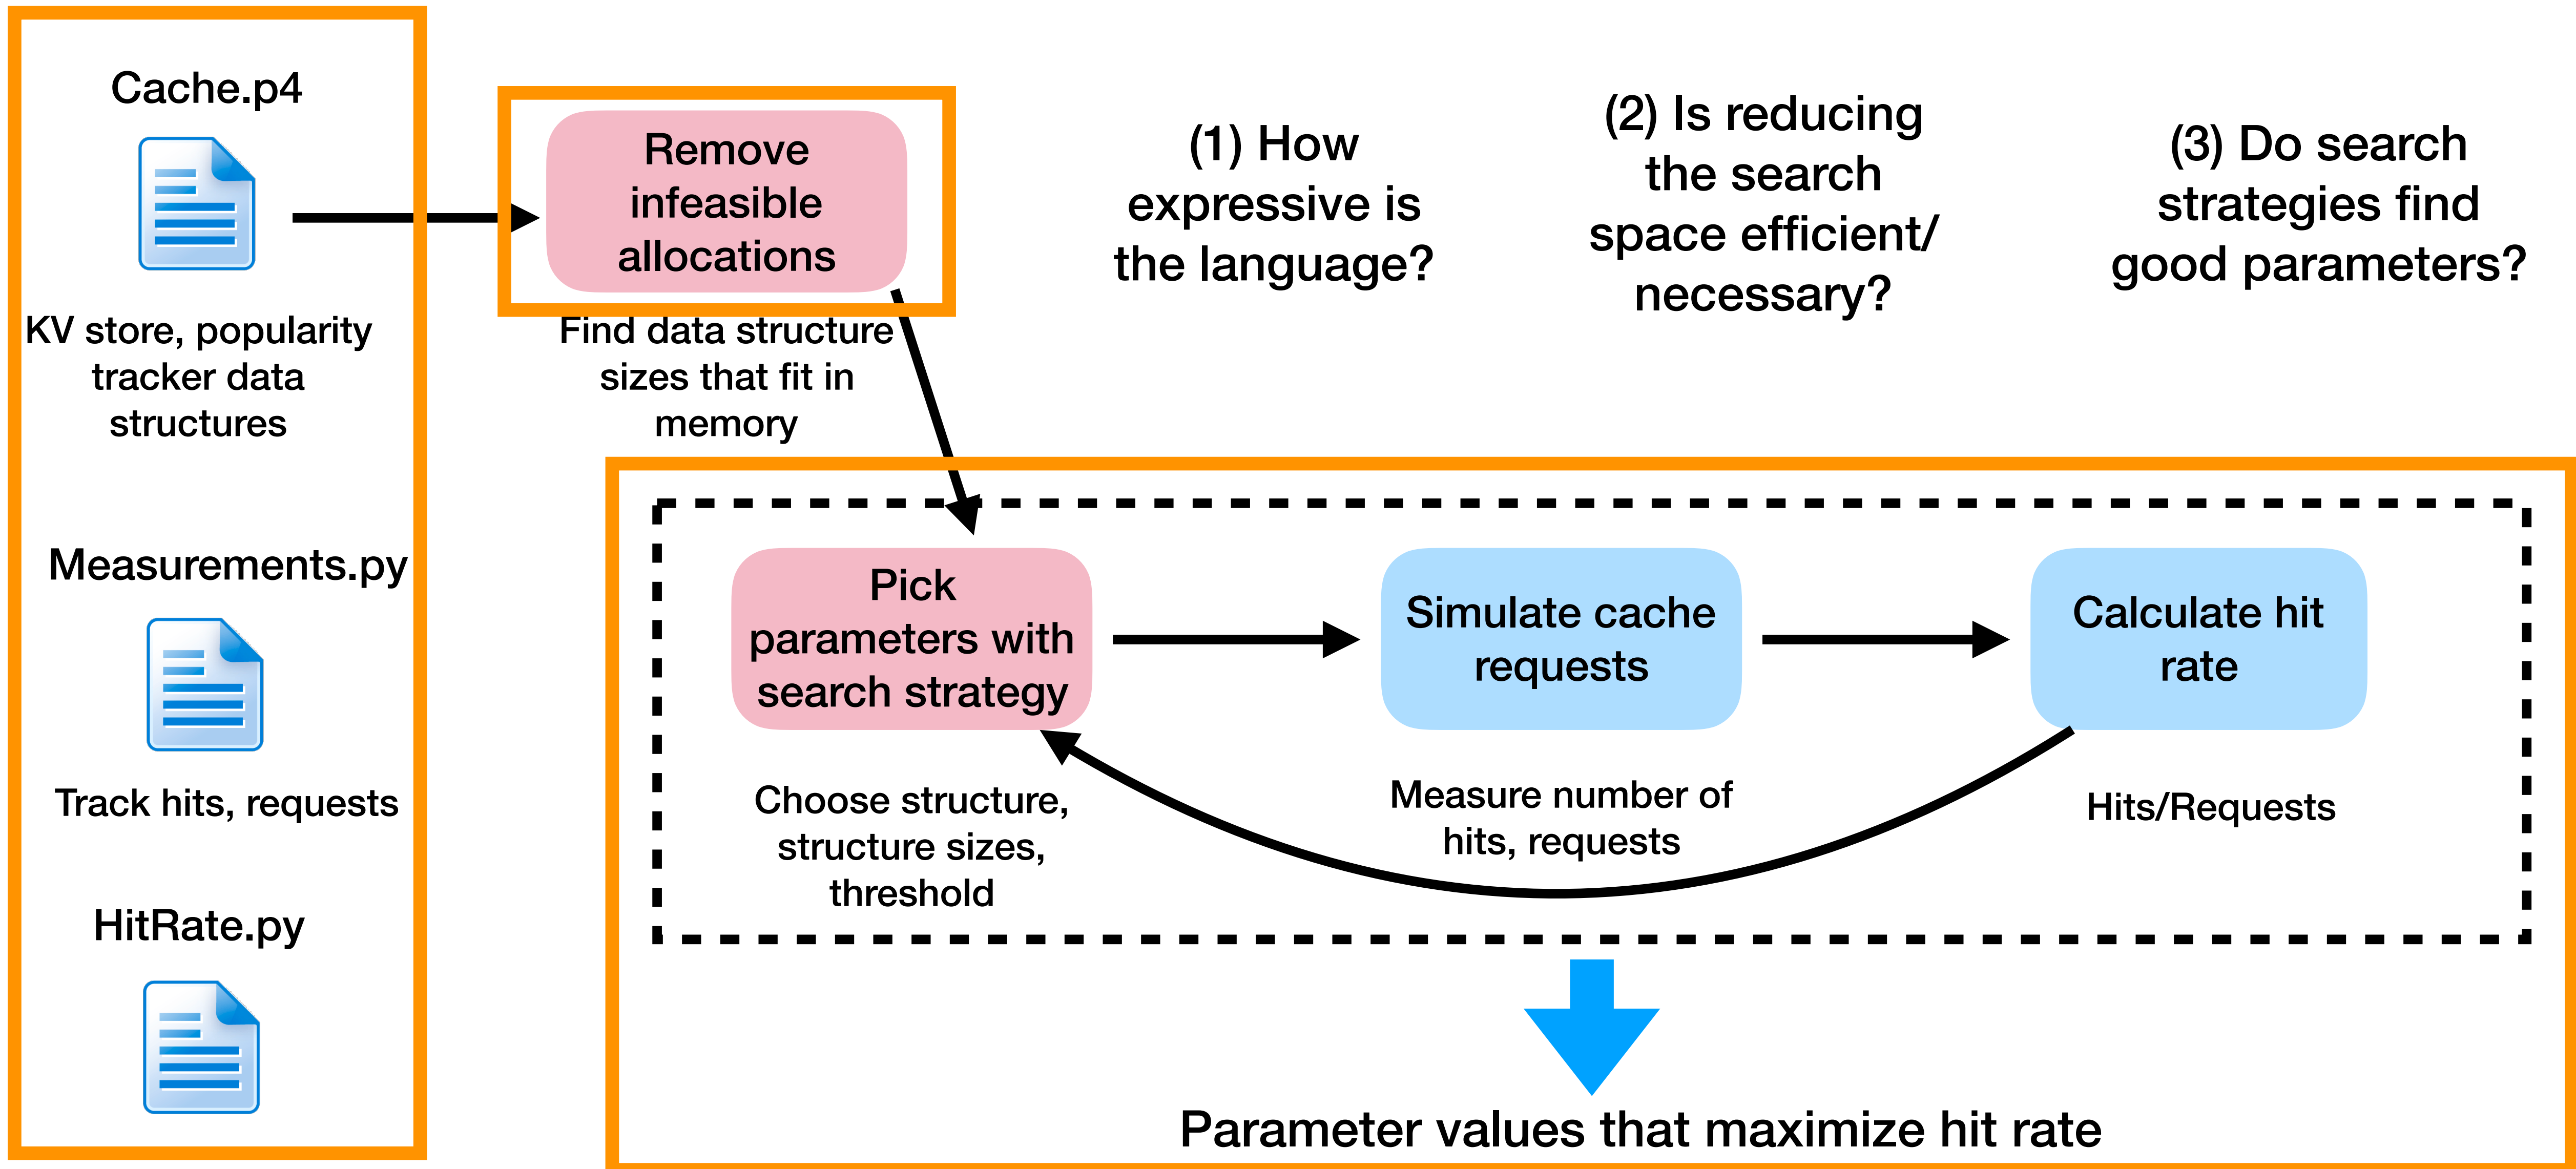
- (0) Exhaustive Search
- (1) Nelder-Mead Simplex
- (2) Simulated Annealing
- (3) Bayesian Optimization

Gradient Descent  
Use the history of  
Genetic Algorithm  
previous performance  
scores to guide the  
search

# Parasol Overview



# Parasol Evaluation



# Parasol Evaluation

<b>Application</b>
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement



# P4All vs Parasol Language

<b>Application</b>
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement

# P4All vs Parasol Language

Application	Params in P4All?
Cache	X
Hash Table	
Stateful Firewall	X
Load Balancer	X
Count-Min Sketch	
Precision	
Measurement Queries	
Congestion Detection	
(Short) RTT Measurement	X
(Long) RTT Measurement	X

# P4All vs Parasol Language

Application	Params in P4All?	Objective in P4All?
Cache	X	X
Hash Table		
Stateful Firewall	X	X
Load Balancer	X	X
Count-Min Sketch		
Precision		X
Measurement Queries		X
Congestion Detection		X
(Short) RTT Measurement	X	X
(Long) RTT Measurement	X	

# Preprocessing Evaluation

Application
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement

**How long did it take to  
reduce the search space?**

# Preprocessing Evaluation

Application	Reduction Time
Cache	2 hours
Hash Table	15 seconds
Stateful Firewall	30 seconds
Load Balancer	2 seconds
Count-Min Sketch	16 seconds
Precision	32 minutes
Measurement Queries	1.5 hours
Congestion Detection	15 seconds
(Short) RTT Measurement	23 seconds
(Long) RTT Measurement	3 seconds

# Preprocessing Evaluation

Application	Reduction Time
Cache	2 hours
Hash Table	<b>15 seconds</b>
Stateful Firewall	<b>30 seconds</b>
Load Balancer	<b>2 seconds</b>
Count-Min Sketch	<b>16 seconds</b>
Precision	32 minutes
Measurement Queries	1.5 hours
Congestion Detection	<b>15 seconds</b>
(Short) RTT Measurement	<b>23 seconds</b>
(Long) RTT Measurement	<b>3 seconds</b>

# Preprocessing Evaluation

Application	Reduction Time
Cache	<b>2 hours</b>
Hash Table	15 seconds
Stateful Firewall	30 seconds
Load Balancer	2 seconds
Count-Min Sketch	16 seconds
Precision	32 minutes
Measurement Queries	<b>1.5 hours</b>
Congestion Detection	15 seconds
(Short) RTT Measurement	23 seconds
(Long) RTT Measurement	3 seconds

# Preprocessing Evaluation

**Optimizer  
found better  
solutions  
with  
reduced  
search  
space!**

Application	Reduction Time
<b>Cache</b>	<b>2 hours</b>
Hash Table	15 seconds
<b>Stateful Firewall</b>	<b>30 seconds</b>
Load Balancer	2 seconds
Count-Min Sketch	16 seconds
Precision	32 minutes
<b>Measurement Queries</b>	<b>1.5 hours</b>
<b>Congestion Detection</b>	<b>15 seconds</b>
(Short) RTT Measurement	23 seconds
(Long) RTT Measurement	3 seconds



# Search Strategy Evaluation

Application
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement

**Do search strategies find parameter values comparable to hand-optimized code?**

# Parasol vs Hand-Optimized Code

Application
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
<b>Measurement Queries</b>
<b>Congestion Detection</b>
(Short) RTT Measurement
<b>(Long) RTT Measurement</b>




**Parasol programs had (nearly) identical performance to hand-optimized code**

# P4All vs Parasol

```
symbolic int kv_size;  
symbolic int cms_rows;  
symbolic int cms_cols;  
symbolic int timeout;  
symbolic int threshold;
```


**P4All and Parasol generate  
the same amount of packets  
forwarded to  
storage server!**

**Parasol threshold (300)**



**Miss rate = 66%**

**Suboptimal threshold (5000)**



**Miss rate = 68%**

# P4All vs Parasol

```
symbolic int kv_size;  
symbolic int tracker_rows;  
symbolic int tracker_cols;  
symbolic int timeout;  
symbolic int threshold;  
symbolic bool cms_tracker;
```

Distribution	CMS Miss Rate	Precision Miss Rate
High skew		
Moderate skew		

# Summary

**Bridge the gap between expressiveness needed for complex programs and resource scarcity**

**P4All**  
[HotNets'20]

**Parasol**  
[In Submission]

Language abstraction  
for expressiveness


Separate low-level details  
from program code

**P4All**  
[NSDI'22]

**Parasol**  
[In Submission]

Optimization with  
scarce resources

Automatically tailor  
program to environment



Compile to domain-  
specific code (p4)

P4 code that fits switch  
architecture



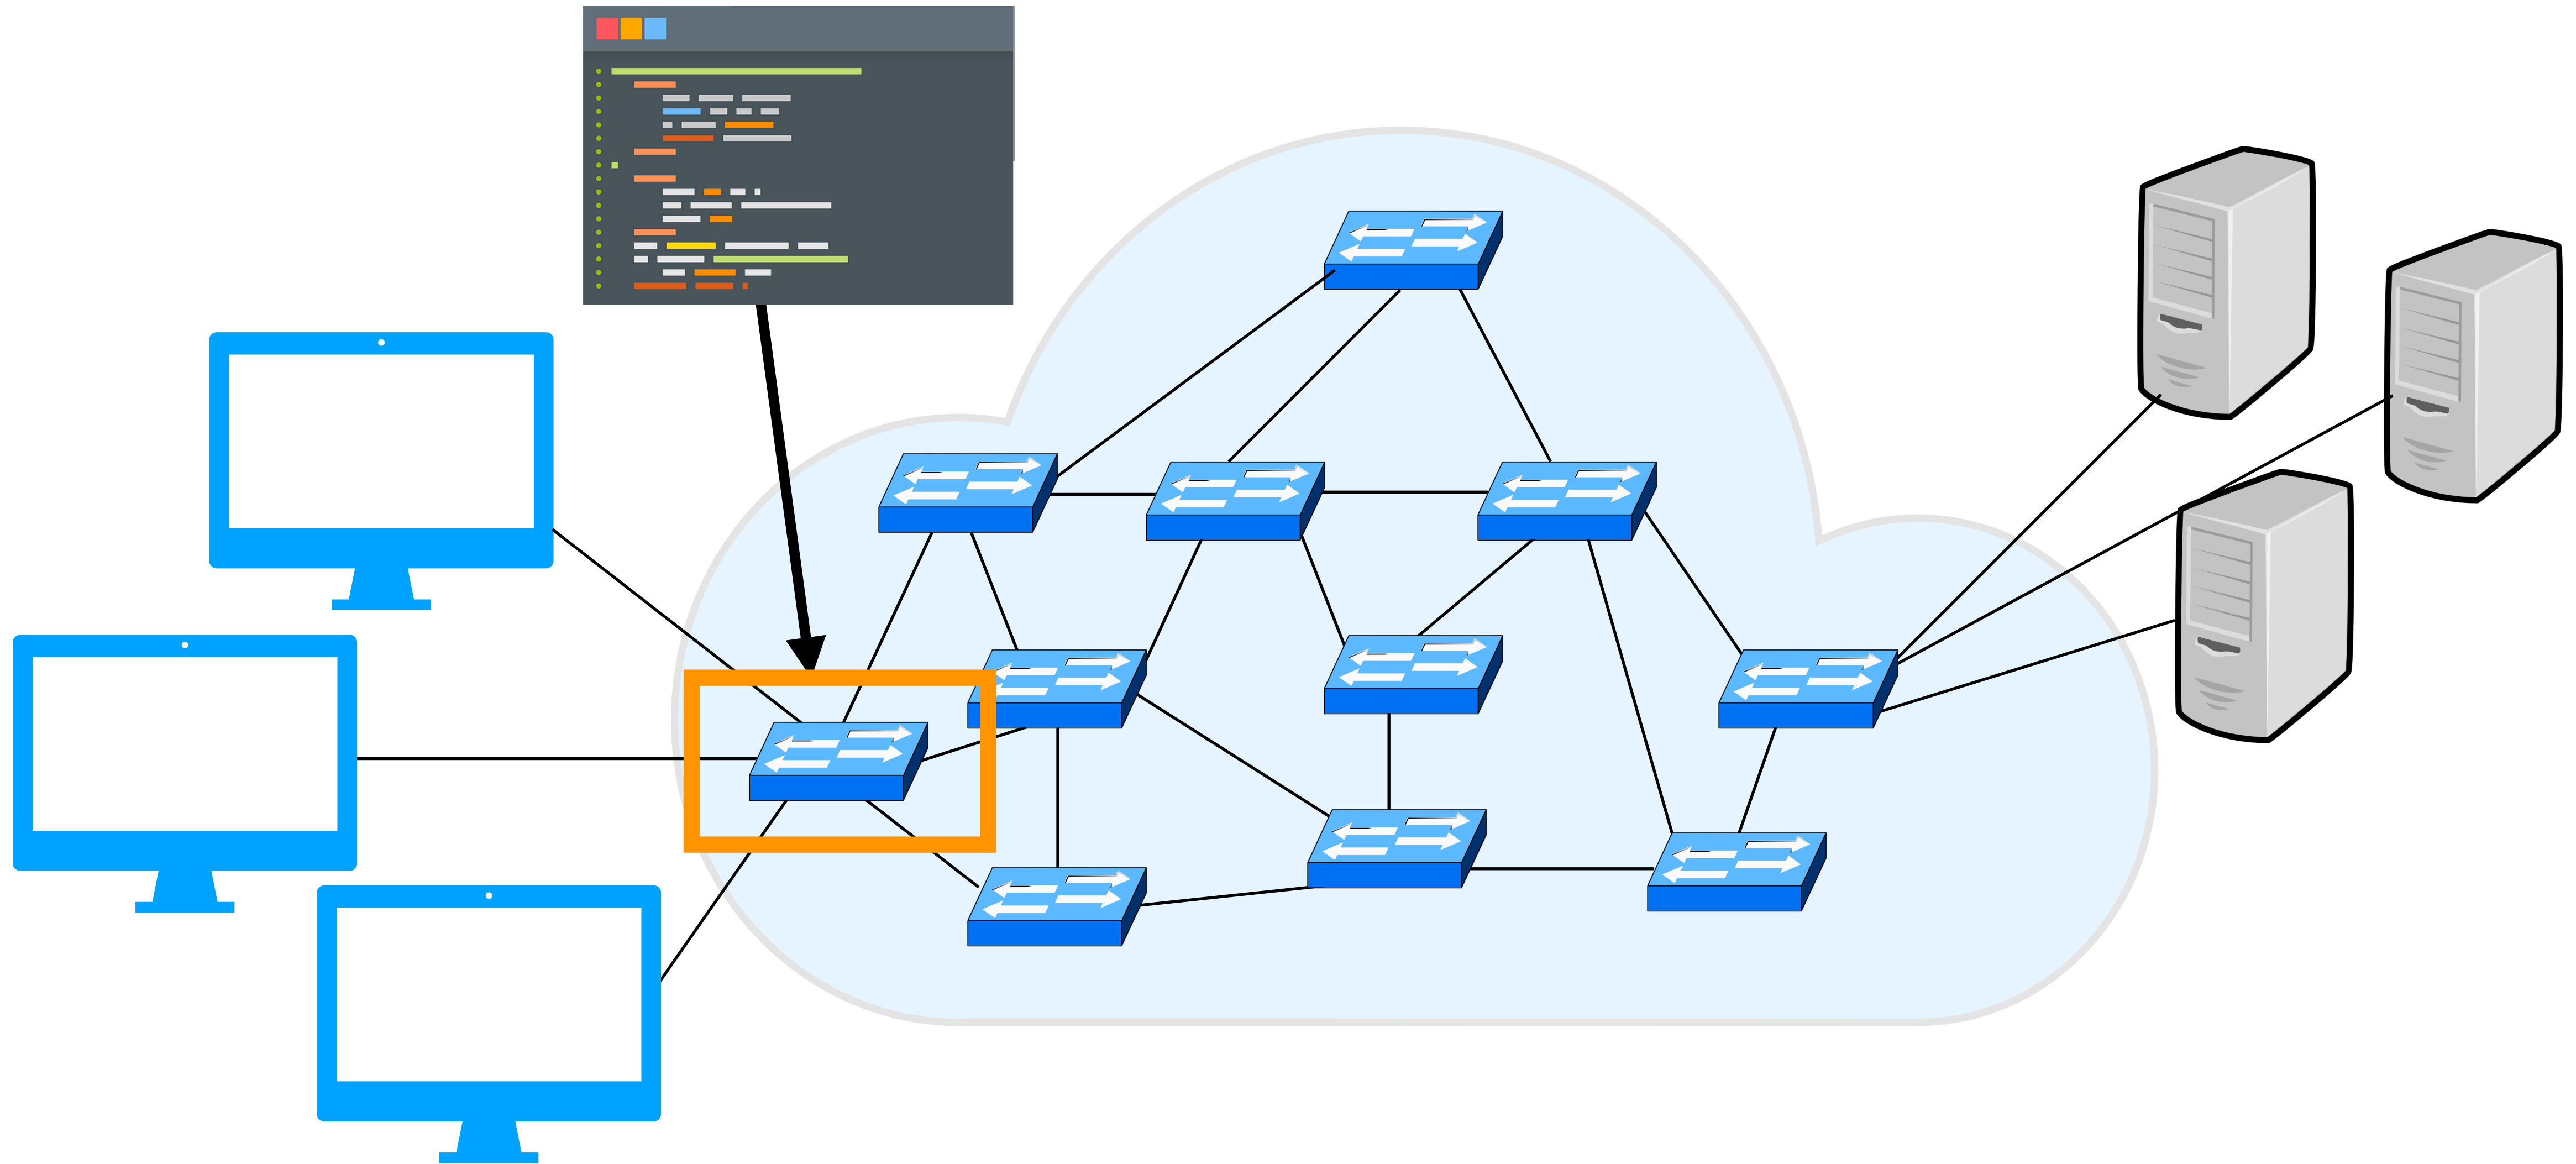
Programmable  
switches

Compile to hardware  
binary

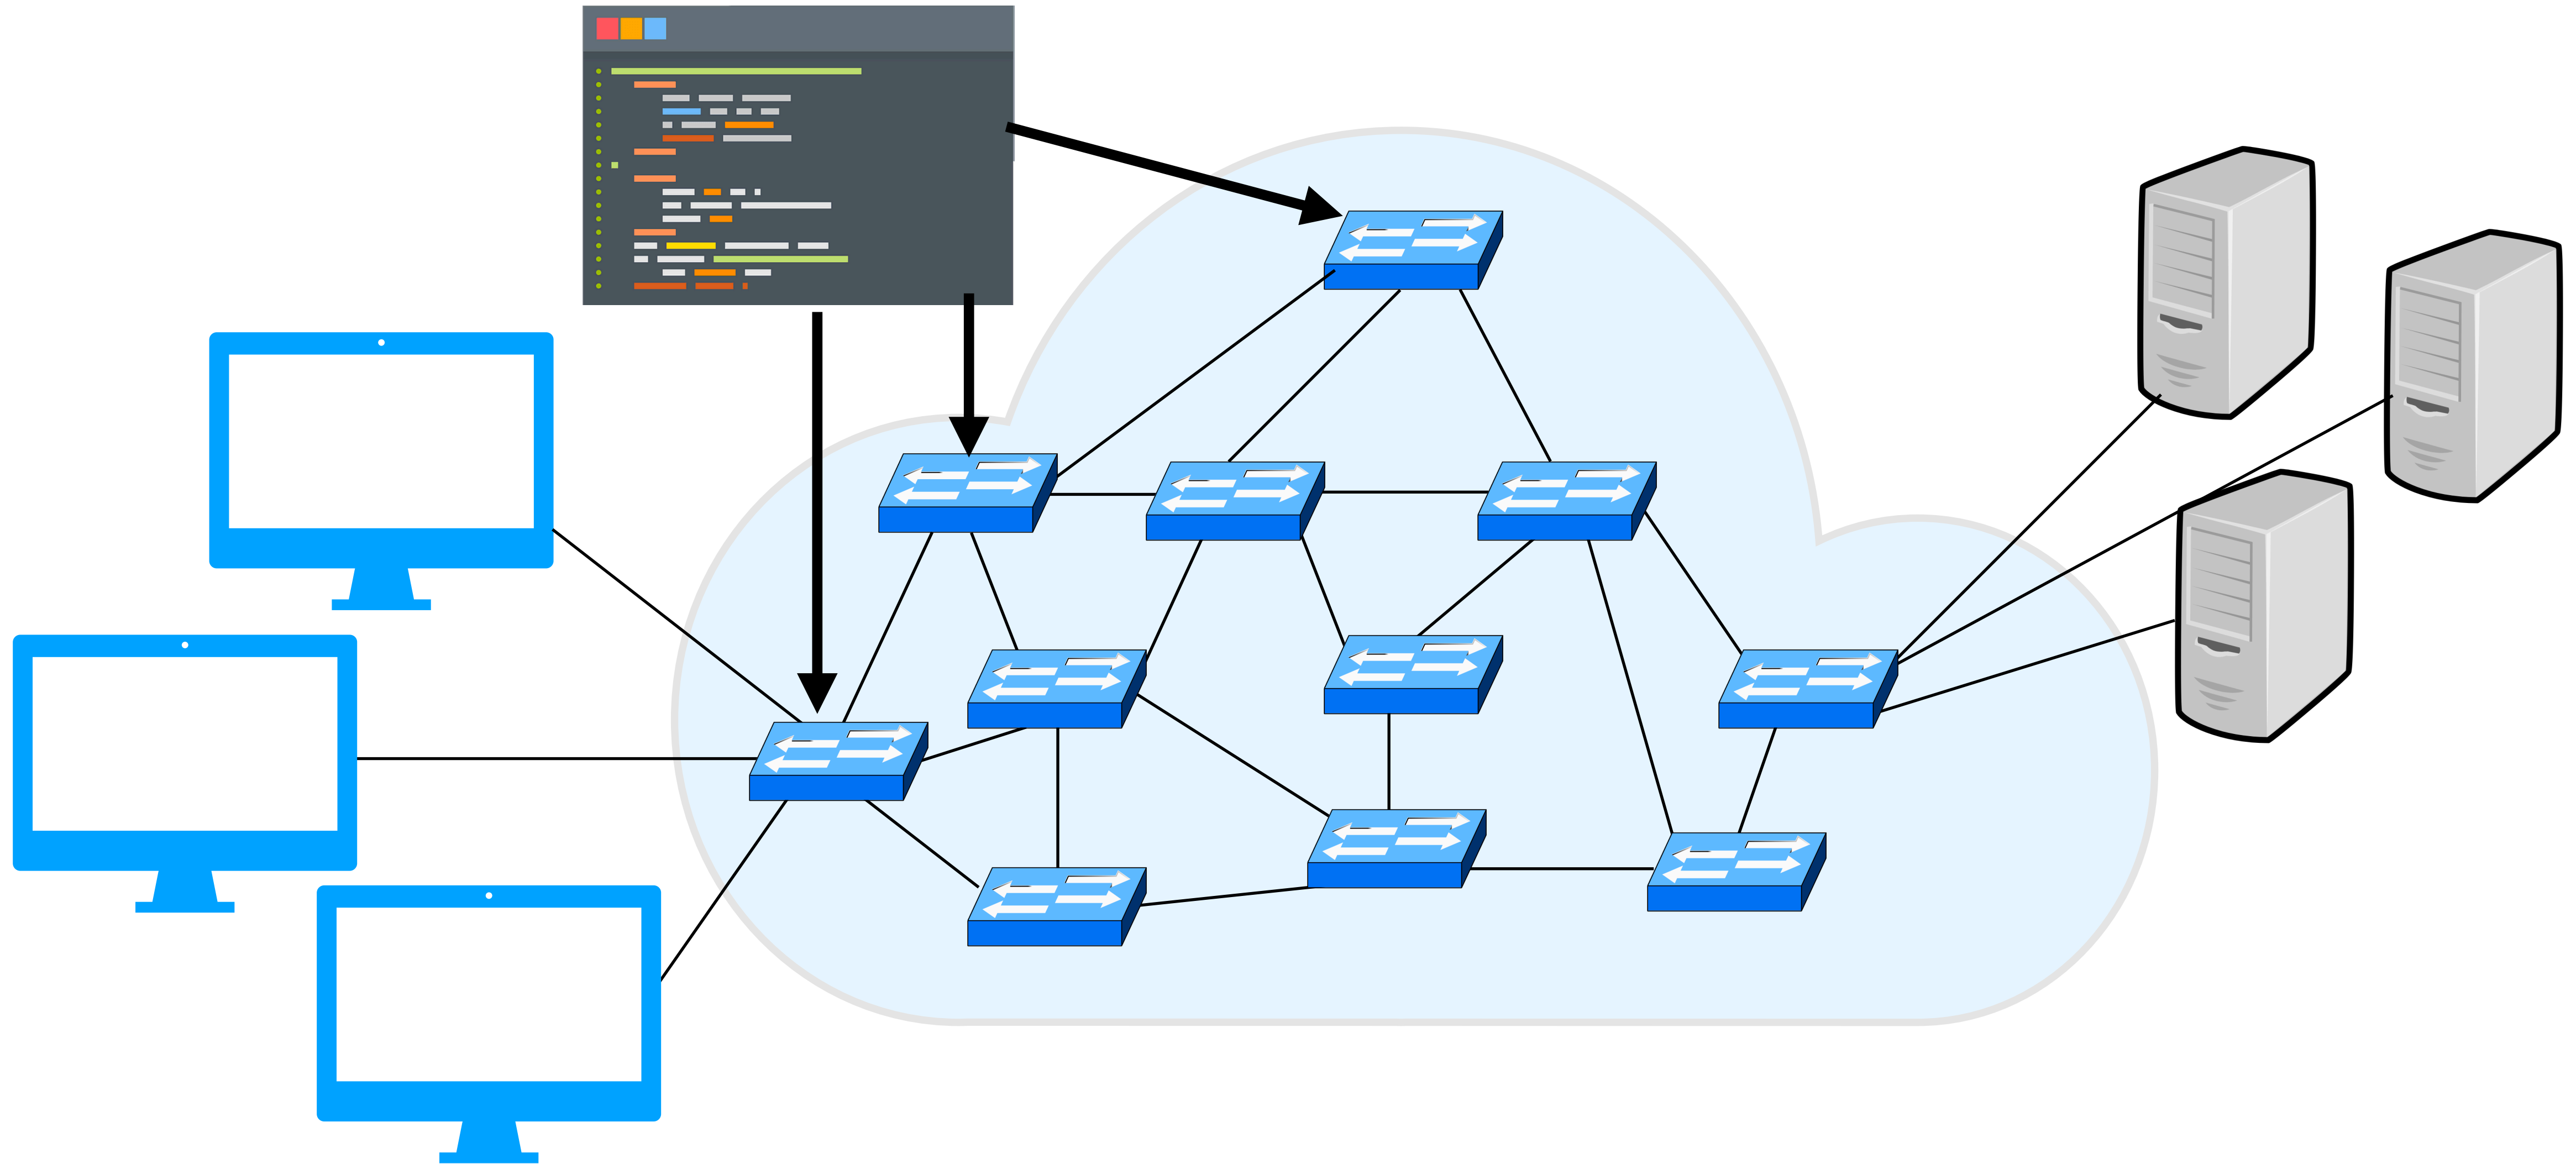
Ready to run in  
the network!

# Future Work: Programming the Entire Network

# Single-Device Programming



# Single-Device Programming





# Future Work: Programming the Entire Network

A system that:

(1) Decides where to put the program

(2) Finds the optimal resource allocation across multiple network devices

NOT only:



Programmable switches



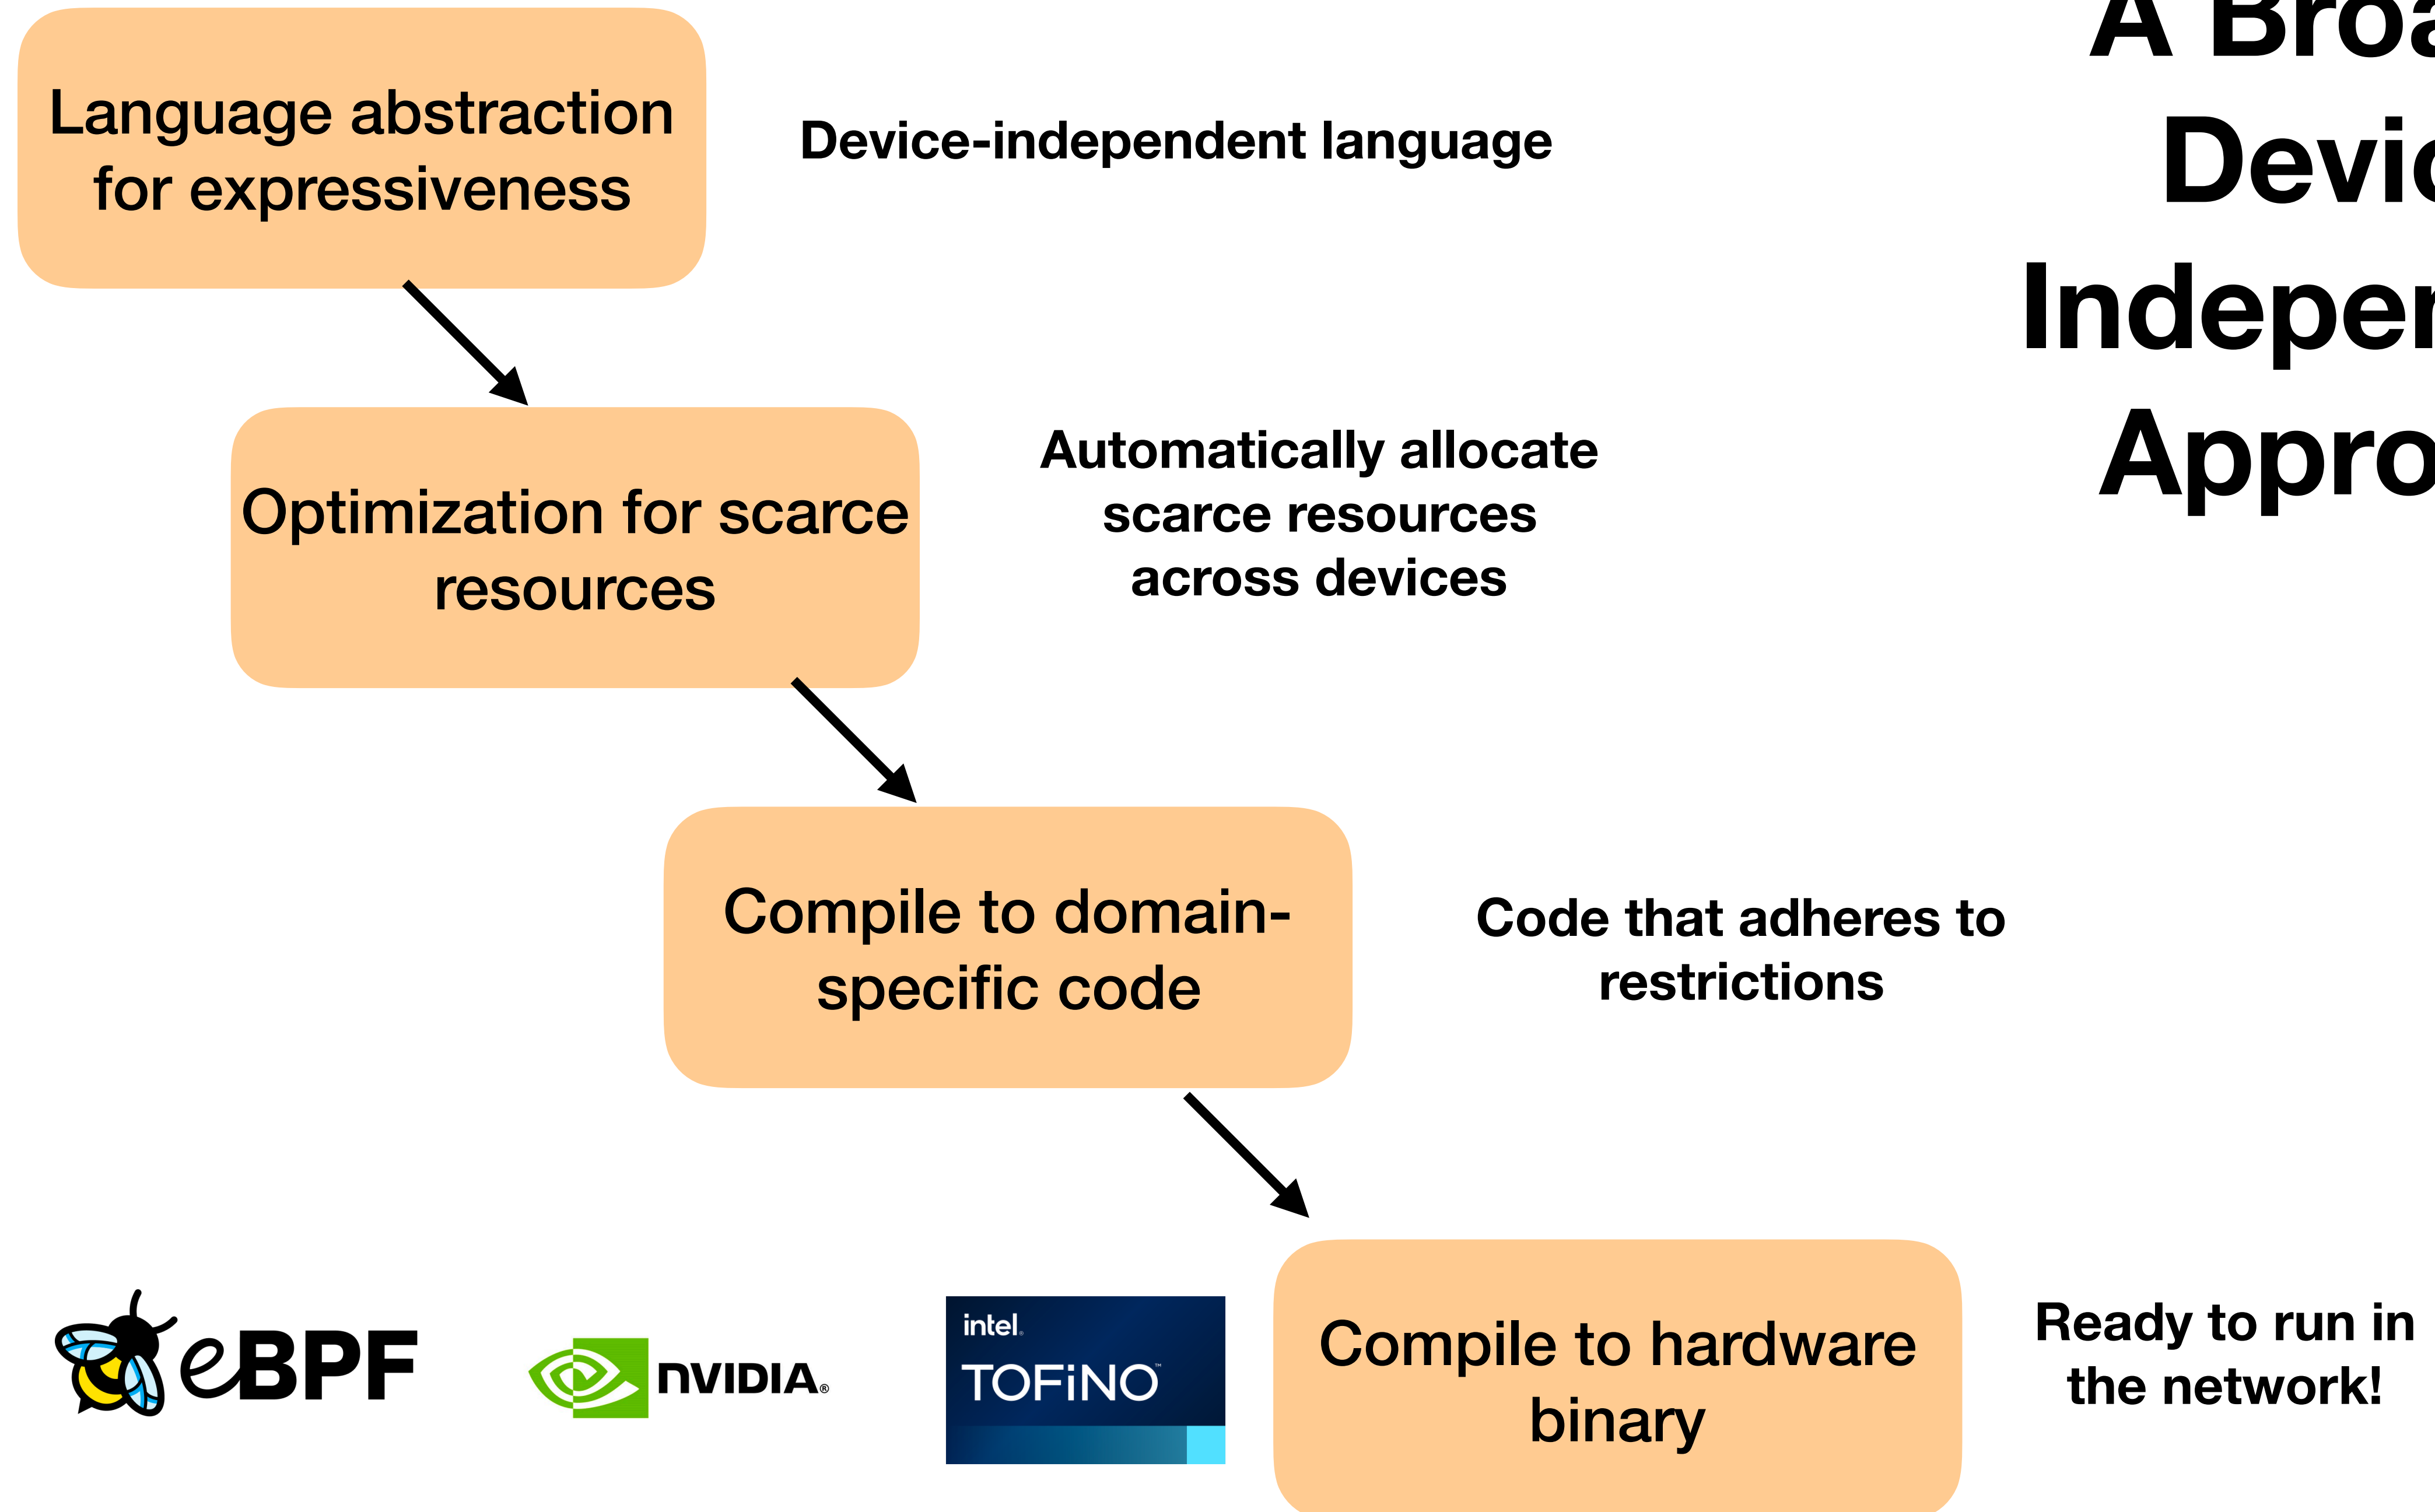
Linux kernel

Goal: Programmers can easily describe how network-wide applications should perform without grappling with low-level details.



SmartNICs

# A Broader Device- Independent Approach



# Thank you!



Jen Rexford



Dave Walker



Shir Landau Feibish




John Sonchack



Devon Loehr



Mina Tahmasbi  
Arashloo



Rob Harrison



Maria Apostolaki



Wei Luo



Sharad Agarwal



Rachee Singh



Ryan Beckett



Gerry Wan



Yiming Qiu

# Language Expressiveness Under Extreme Scarcity in Programmable Data Planes

Mary Hogan



**Backup slides**

# P4AI ILP Constraints

Variables	
Actions	#1 $\{x_{a_i,s} \mid 0 \leq s < S\}$
Registers	#2 $\{m_{r_i,s} \mid 0 \leq s < S\}$
Match-Action Tables	#3 $\{tm_{t_i,s} \mid 0 \leq s < S\}$
Metadata	#4 $\{d_i \mid i \leq U_v\}$
Constraints	
<b>Dependencies</b>	
Same-Stage	#5 $x_{a_i,s} = x_{b_i,s} \quad s < S$
Exclusion	#6 $x_{a_i,s} \leq 1 - x_{b_i,s}$ $s < S$
Precedence	#7 $x_{b_i,y} \leq 1 - x_{a_i,z}$ $y, z < S, y \leq z$
Conditional	#8 $\sum_{0 \leq s < S} x_{a_i,s} = \sum_{0 \leq s < S} x_{b_i,s}$ $0 \leq i \leq U_v$
<b>Resources</b>	
Memory	#9 $\sum_i m_{r_i,s} \cdot w_{r_i} \leq M \quad \forall s < S$ #10 $m_{r_i,s} \leq x_{a_i,s} \cdot M \quad 0 \leq s < S$ #11 $m_{r_i,s} \cdot w_0 = m_{0,s} \cdot w_{r_i}$ $\forall s < S, r \geq 1$
TCAM	#12 $\sum_i tm_{t_i,s} \cdot tw_{t_i} \leq T \quad \forall s < S$
Stateful ALUs	#13 $\sum_i H_f(a_i) \cdot x_{a_i,s} \leq F$ $\forall 0 \leq s < S$
Stateless ALUs	#14 $\sum_i H_l(a_i) \cdot x_{a_i,s} \leq L$ $\forall 0 \leq s < S$
PHV	#15 $\sum_i d_i \cdot bits_d \leq P - P_{fixed}$ #16 $d_i = \sum_{0 \leq s < S} x_{a_i,s}$ if $accesses(a,d)$
Hash Functions	#17 $\sum_i h_{ha_i,s} \leq N \quad \forall s < S$
<b>Others</b>	
At Most Once	#18 $\sum_{0 \leq s < S} x_{a_i,s} \leq 1$
Inelastic Actions	#19 $\sum_{0 < s < S} x_{a_{ne},s} = 1$

# The System

**P4All**  
[HotNets'20]  
**Parasol**  
[In Submission]


Language abstraction  
for expressiveness

Did P4All reduce code  
repetition?

**P4All**  
[NSDI'22]  
**Parasol**  
[In Submission]

Optimization with  
scarce resources

Automatically tailor  
program to environment



Compile to domain-  
specific code (P4)

P4 code that fits switch  
architecture



Compile to hardware  
binary

Ready to run in  
the network!

# Repetition Reduction

<b>Application</b>	<b>P4 LoC</b>	<b>P4All LoC</b>
CMS		
Key-value store		
Key-value store + CMS		
Switch.p4		
IP forwarding + stateful firewall		
Beaucoup		
Precision		
NetChain		
SketchLearn		
Conquest		
Bloom filter		




# Repetition Reduction

<b>Application</b>	<b>P4 LoC</b>	<b>P4All LoC</b>
CMS	207	179
Key-value store		
Key-value store + CMS	216	170
Switch.p4		
IP forwarding + stateful firewall		
Beaucoup		
Precision	366	273
NetChain		
SketchLearn		
Conquest		
Bloom filter	179	70

# Repetition Reduction

<b>Application</b>	<b>P4 LoC</b>	<b>P4All LoC</b>
CMS	207	179
Key-value store	127	127
Key-value store + CMS	216	170
Switch.p4	804	804
IP forwarding + stateful firewall	282	282
Beaucoup	541	541
Precision	366	273
NetChain	242	242
SketchLearn	445	445
Conquest	869	869
Bloom filter	179	70

# ILP Overhead



# P4All vs Parasol Language

	<b>P4All</b>	<b>Parasol</b>
<b>Base language</b>	P4	Lucid
<b>Parameters</b>	Symbolic integers	Symbolic values (integers, booleans, etc.)
<b>Objectives</b>	Closed-form objective functions	Python objective functions
<b>Optimization</b>	ILP	Simulation

# Compilation Strategies

<b>Strategy</b>	<b>Resources Considered</b>	<b>Average compile time</b>
Full Compilation	All switch resources	3 min

# Compilation Strategies

Strategy	Resources Considered	Average compile time
Full Compilation	All switch resources	3 min
Dependency Graph (P4All)	Action dependencies, pipeline stages	51 s

# Compilation Strategies

<b>Strategy</b>	<b>Resources Considered</b>	<b>Average compile time</b>
Full Compilation	All switch resources	3 min
Parasol to P4	All resources except PHV	1.5 min
Dependency Graph (P4All)	Action dependencies, stages	51 s

# Compilation Strategies

Strategy	Resources Considered	Average compile time
Full Compilation	All switch resources	3 min
Parasol to P4	All resources except PHV	1.5 min
Greedy Layout	Action dependencies, stages, memory, hash units, array accesses, ALUs	51 s
Dependency Graph (P4All)	Action dependencies, stages	51 s



# Compilation Times Per App

Application	Full Compilation	Parasol to P4	Greedy Layout	Dataflow Graph
Cache	43 s	4.5 min	3.5 min	3.5 min
Hash Table	23 s	3 s	1 s	1 s
Stateful Firewall	23 s	1.5 min	6 s	6 s
Load Balancer	9 s	1 s	1 s	1 s
Count-Min Sketch	56 s	4 s	1 s	1 s
Precision	1 min	3.5 min	4 min	4 min
Measurement Queries	1.5 min	9 s	7 s	7 s
Congestion Detection	5 min	3 min	2 s	2 s
(Short) RTT Measurement	16 s	49 s	3 s	3 s
(Long) RTT Measurement	24 s	50 s	3 s	3 s

# Solutions Per App

Application	Full Compilation	Parasol to P4	Greedy Layout	Dataflow Graph
Cache	1167	1167	1179	1573
Hash Table	60	60	60	65
Stateful Firewall	24	25	25	26
Load Balancer	13	13	13	13
Count-Min Sketch	142	143	143	143
Precision	36	36	36	39
Measurement Queries	5000	5065	5065	10816
Congestion Detection	17	25	25	26
(Short) RTT Measurement	24	24	24	26
(Long) RTT Measurement	13	13	13	13

# Search Strategy Evaluation

Application
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement

**Do search strategies find the optimal parameter values?**

# Search Strategy Evaluation

Application
Cache
Hash Table
Stateful Firewall
Load Balancer
Count-Min Sketch
Precision
Measurement Queries
Congestion Detection
(Short) RTT Measurement
(Long) RTT Measurement

Do search strategies find the optimal parameter values?

**Optimal parameters:  
simulate every set of  
parameters in reduced  
search space**

# Search Strategy Evaluation

Application	Optimal Parameters?
Cache	
Hash Table	✓
Stateful Firewall	✓
Load Balancer	
Count-Min Sketch	✓
Precision	✓
Measurement Queries	✓
Congestion Detection	
(Short) RTT Measurement	✓
(Long) RTT Measurement	

# Search Strategy Evaluation

Application	Optimal Parameters?
<b>Cache</b>	
Hash Table	✓
Stateful Firewall	✓
<b>Load Balancer</b>	
Count-Min Sketch	✓
Precision	✓
Measurement Queries	✓
<b>Congestion Detection</b>	
(Short) RTT Measurement	✓
<b>(Long) RTT Measurement</b>	

# Search Strategy Evaluation

Application
<b>Cache</b>
Hash Table
Stateful Firewall
<b>Load Balancer</b>
Count-Min Sketch
Precision
Measurement Queries
<b>Congestion Detection</b>
(Short) RTT Measurement
<b>(Long) RTT Measurement</b>

$$\frac{Opt\_score - Strategy\_score}{Opt\_score}$$

# Search Strategy Evaluation


Application
<b>Cache</b>
Hash Table
Stateful Firewall
<b>Load Balancer</b>
Count-Min Sketch
Precision
Measurement Queries
<b>Congestion Detection</b>
(Short) RTT Measurement
<b>(Long) RTT Measurement</b>

$$\frac{Opt\_score - Strategy\_score}{Opt\_score} \leq 12\%$$



# Parasol vs Hand-Optimized Code

## RTT Measurement



RTT = response timestamp  
- request timestamp

$M$ : size of data structure

$p$ : probability of adding new  
request to data structure

Parasol optimized  
parameters

**Max percentile error = 28%**

Hand-optimized parameters

**Max percentile error = 31%**