

Towards Open-Source and Automatic Performance Characterization Hardware

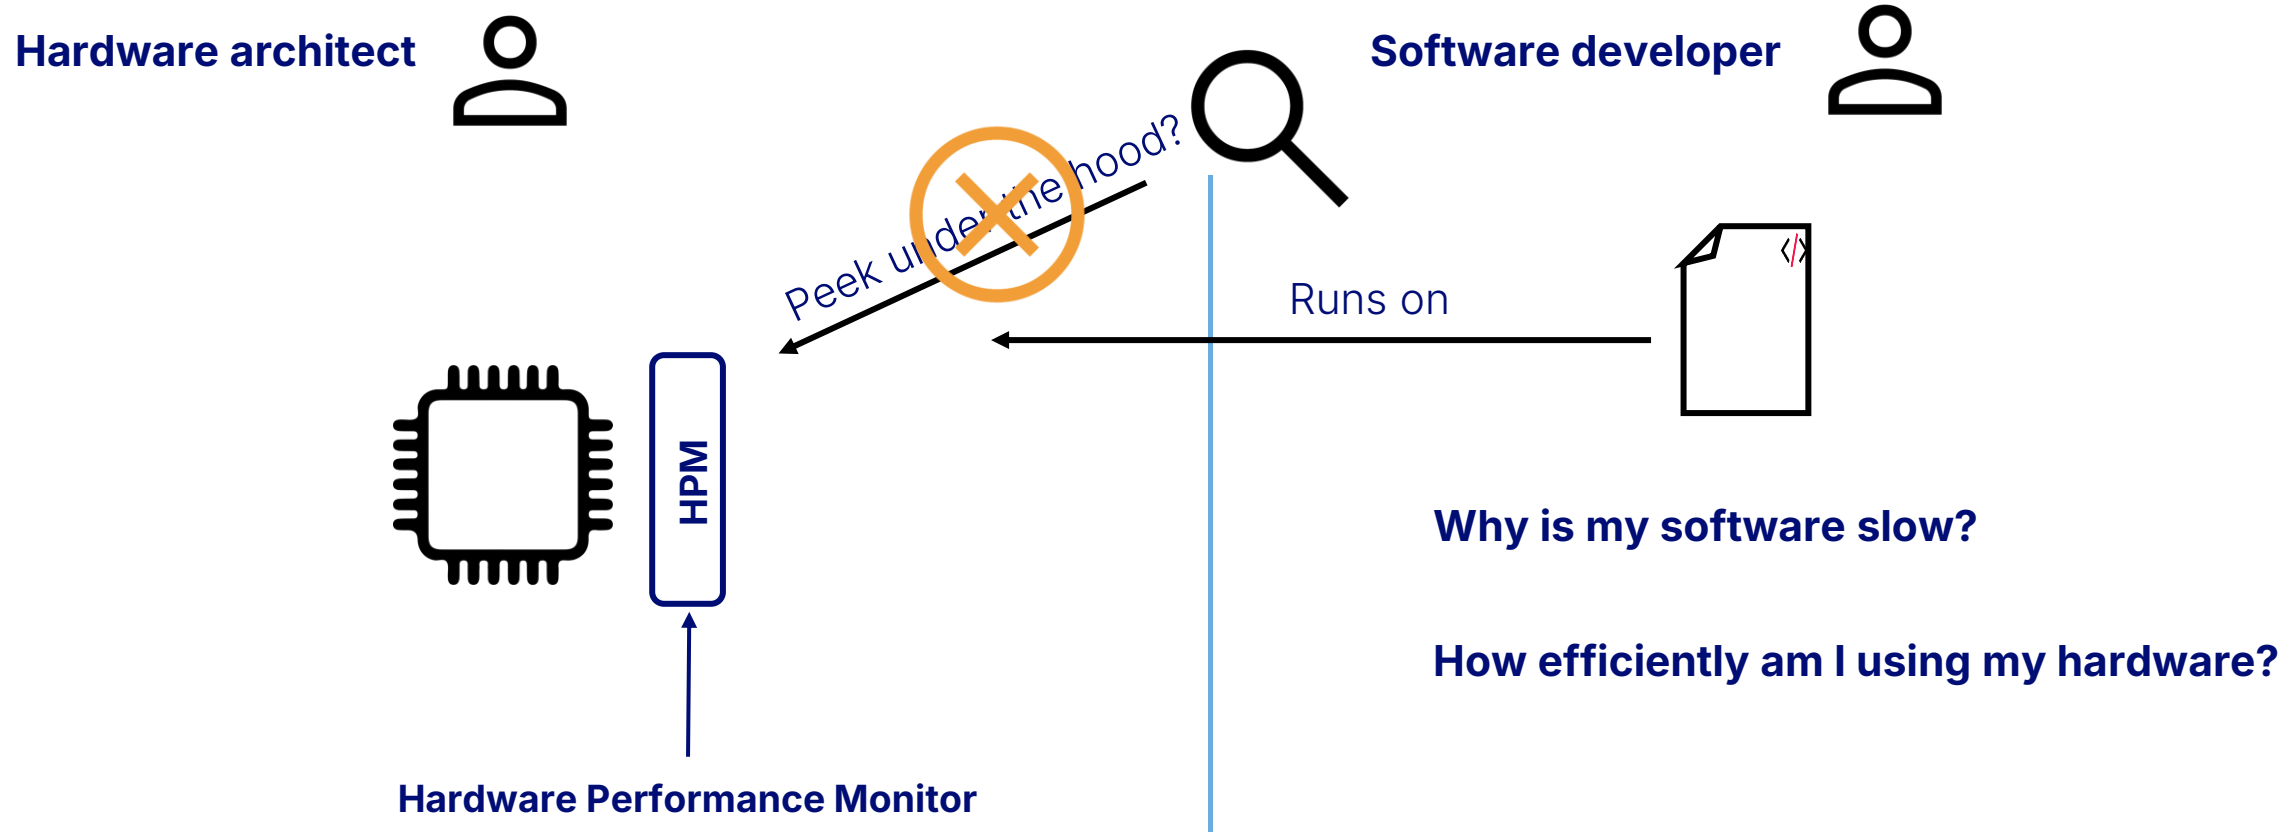
Matthew Edwin Weingarten



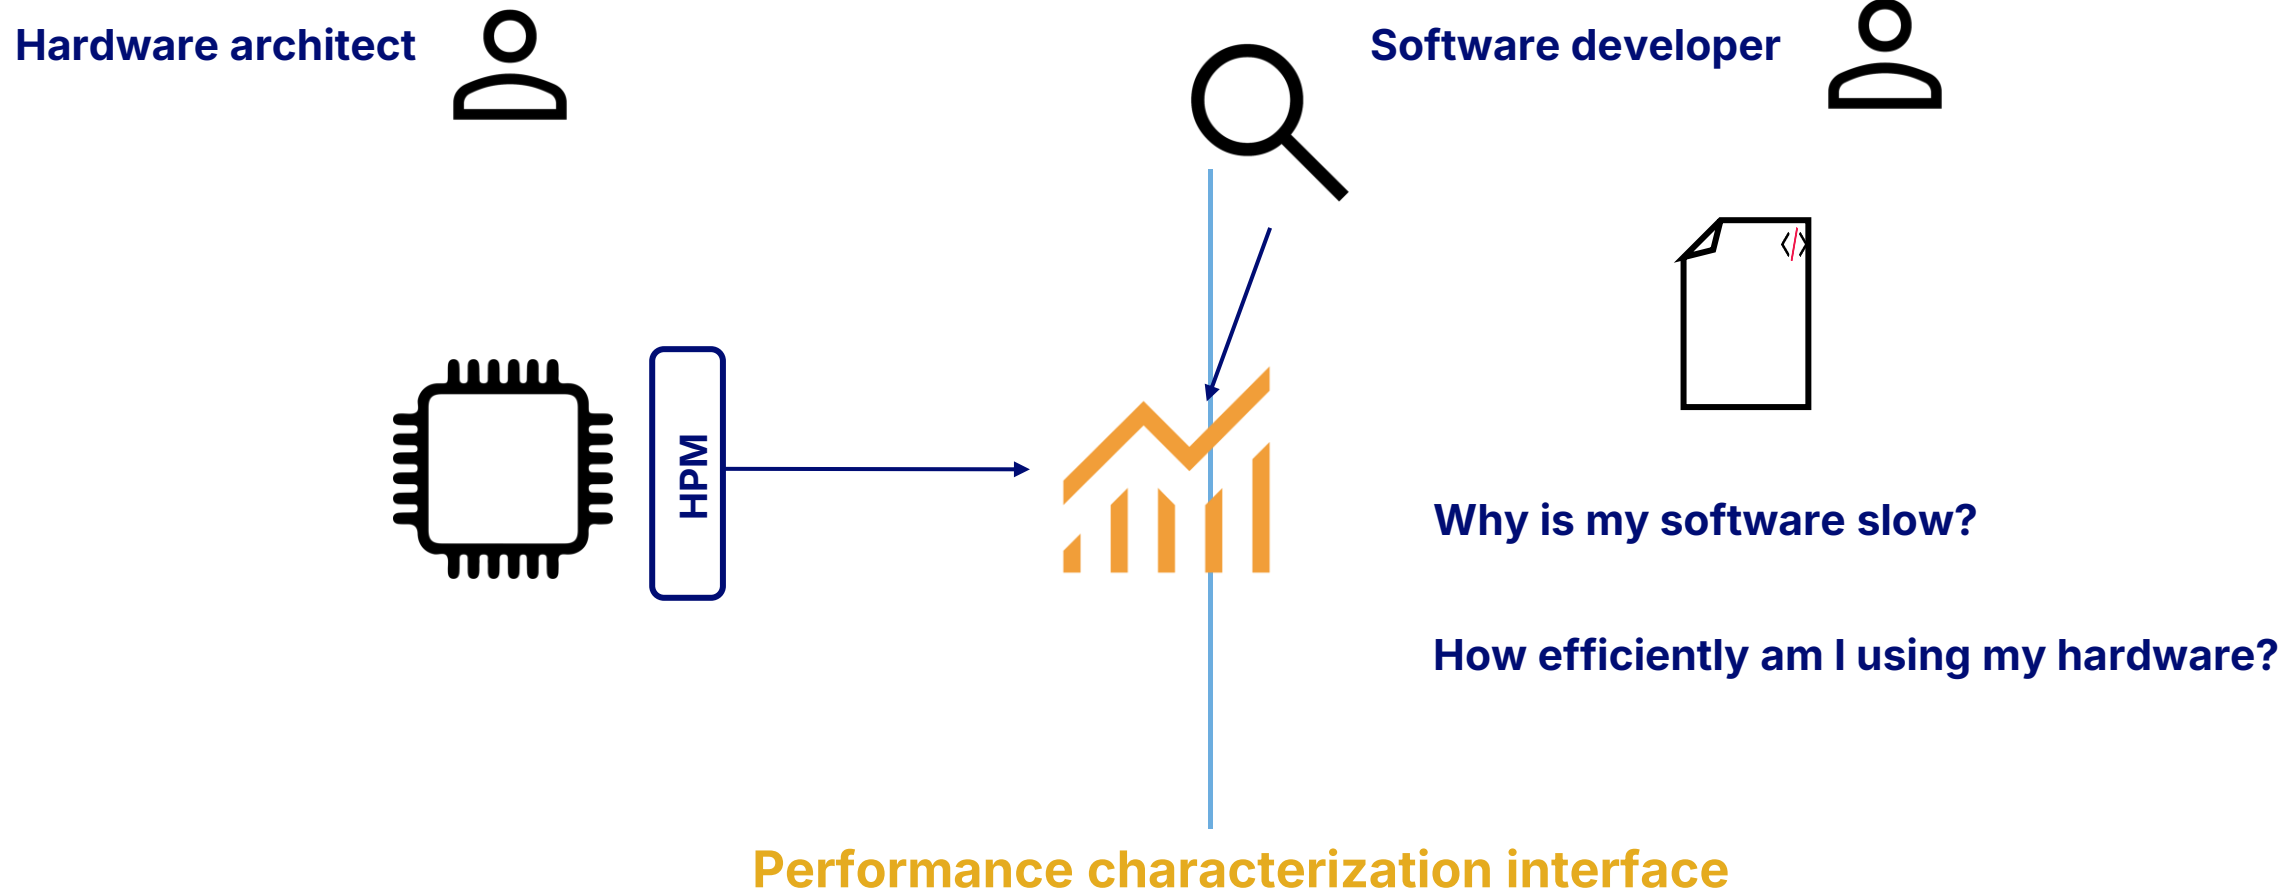
COLUMBIA ENGINEERING




Why do we need performance characterization?



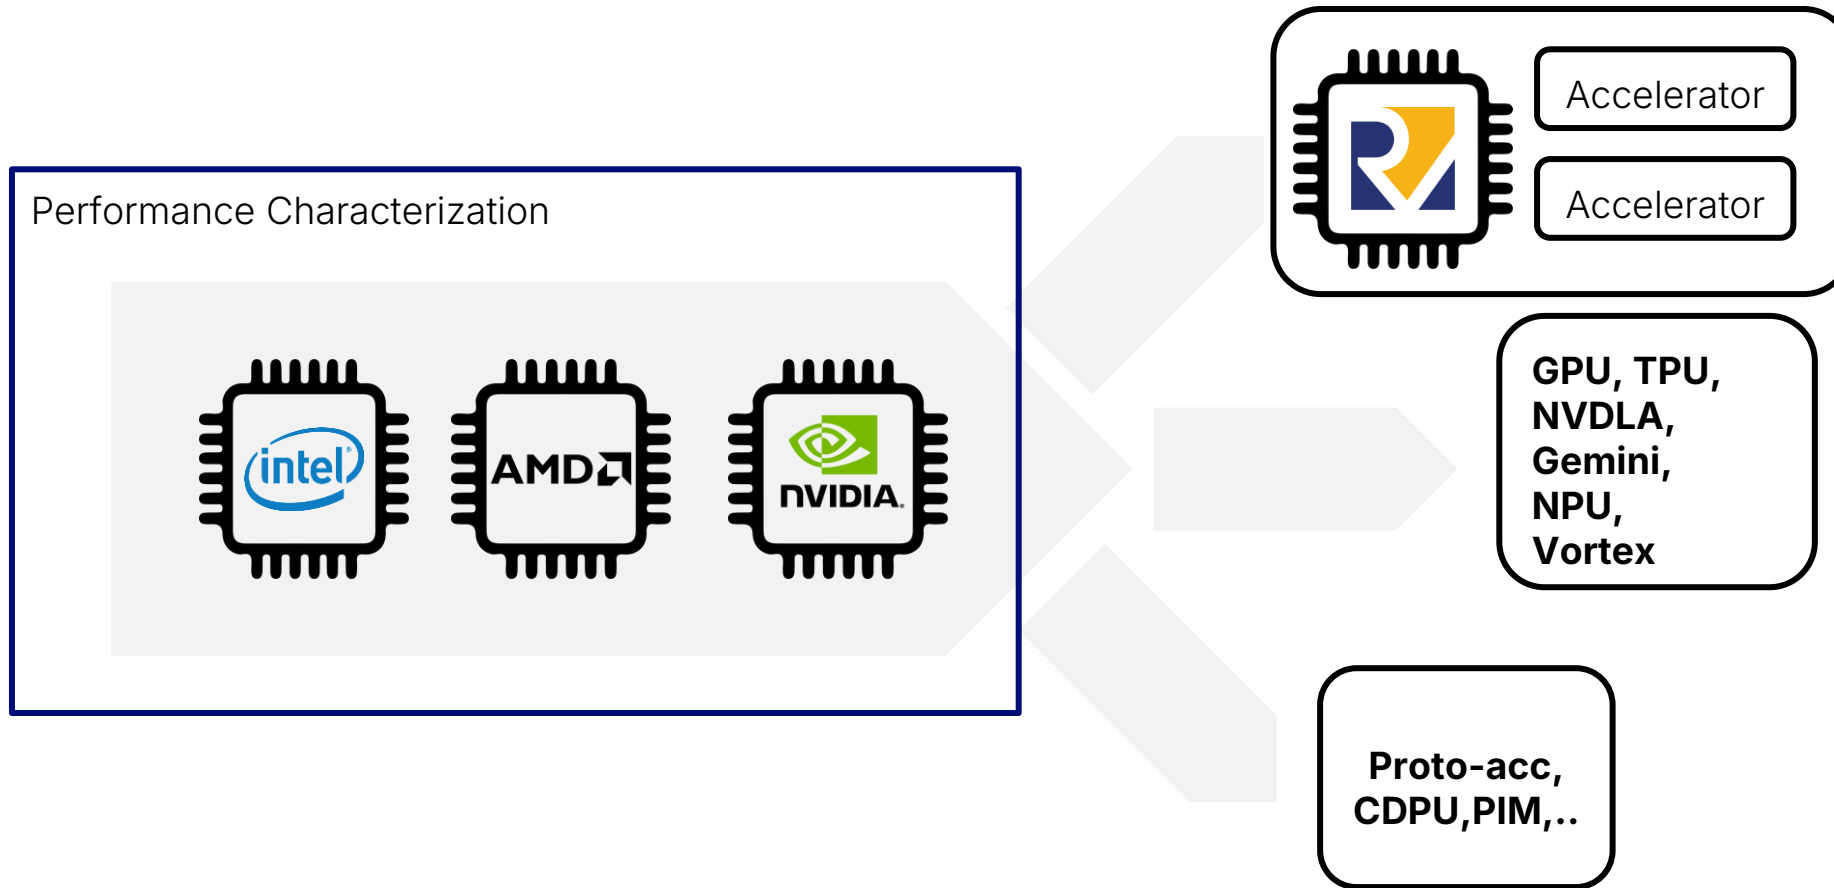
Why do we need performance characterization?



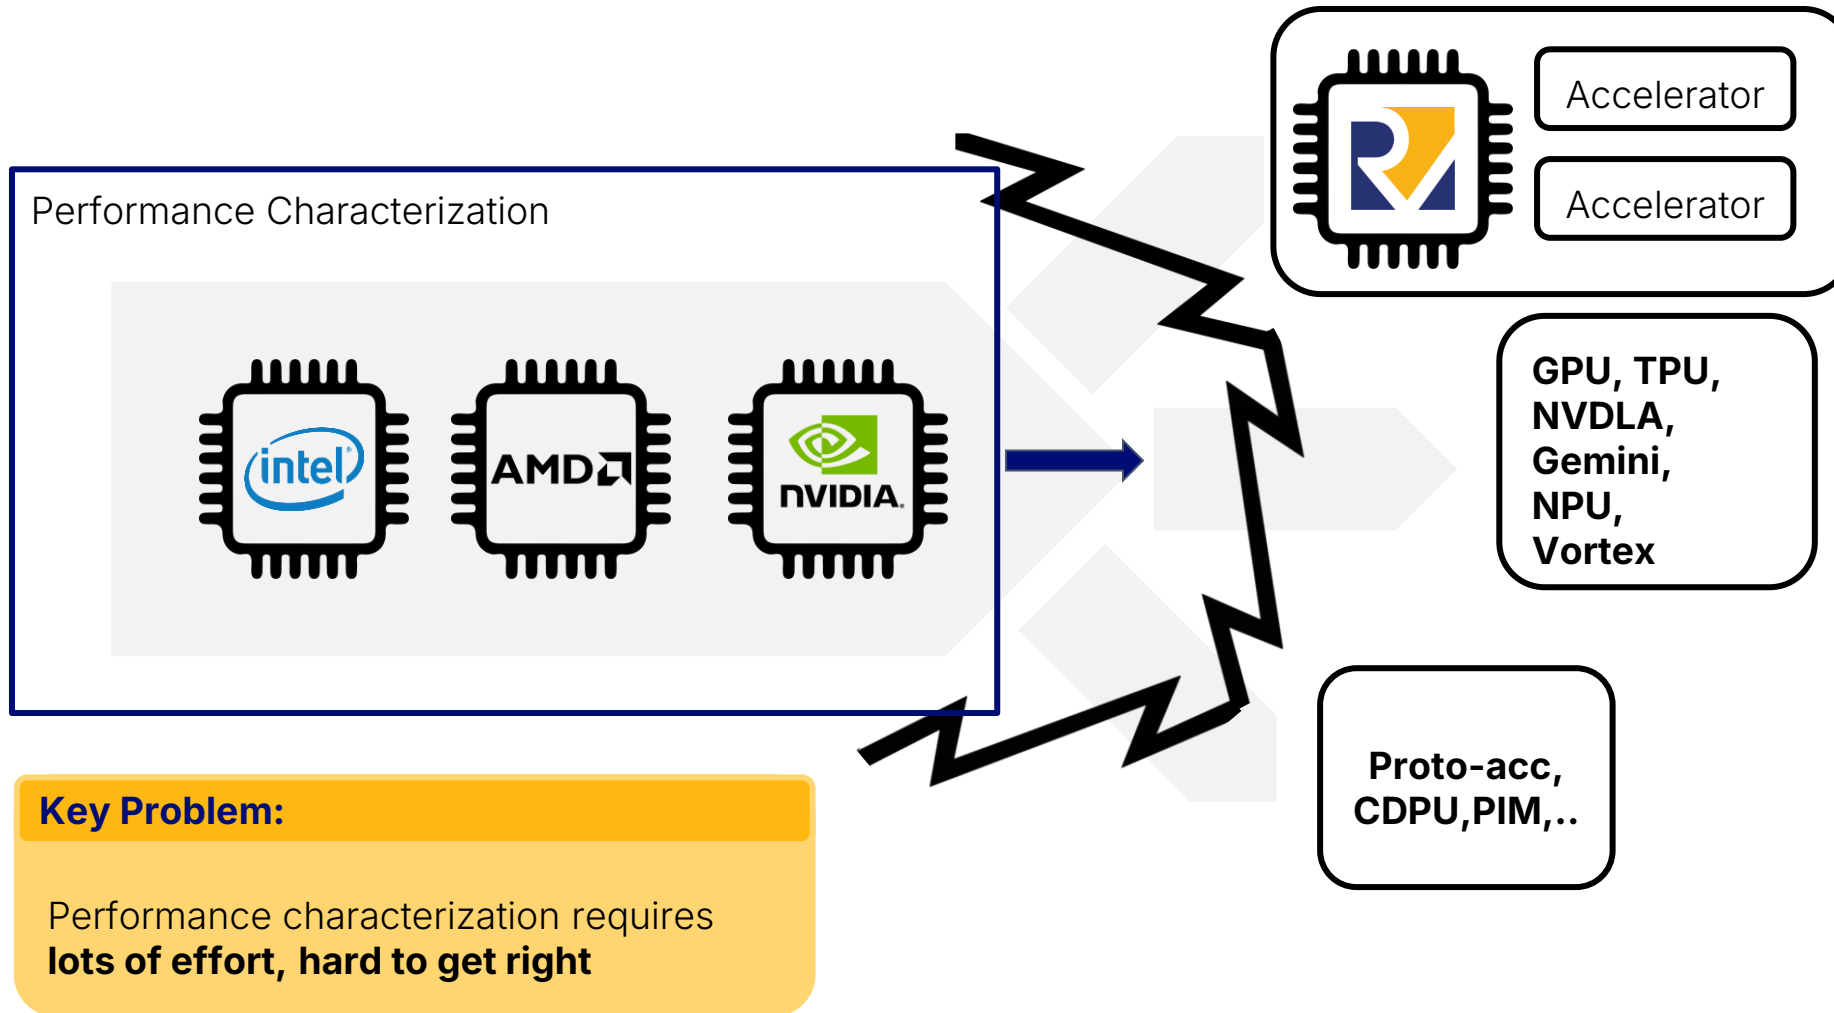
What are the trends in hardware design?



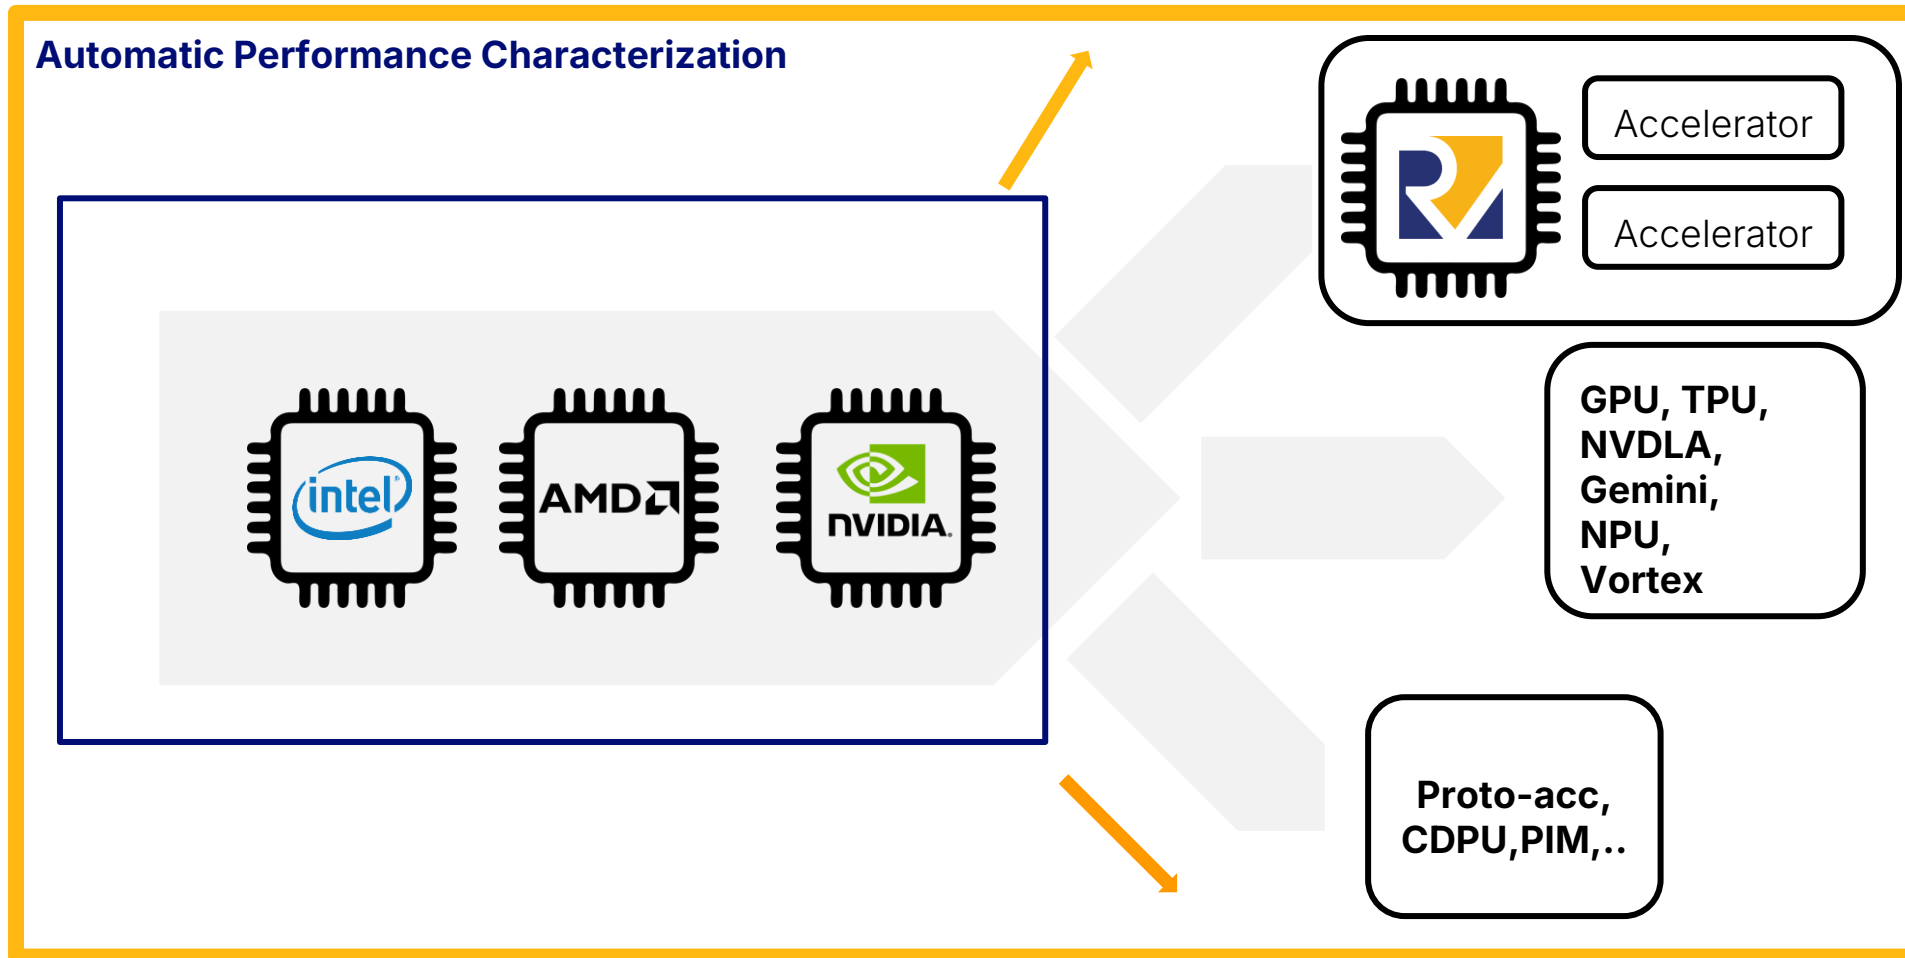
What are the trends in hardware design?



What are the trends in hardware design?






Our vision: Open-source and automatic characterization



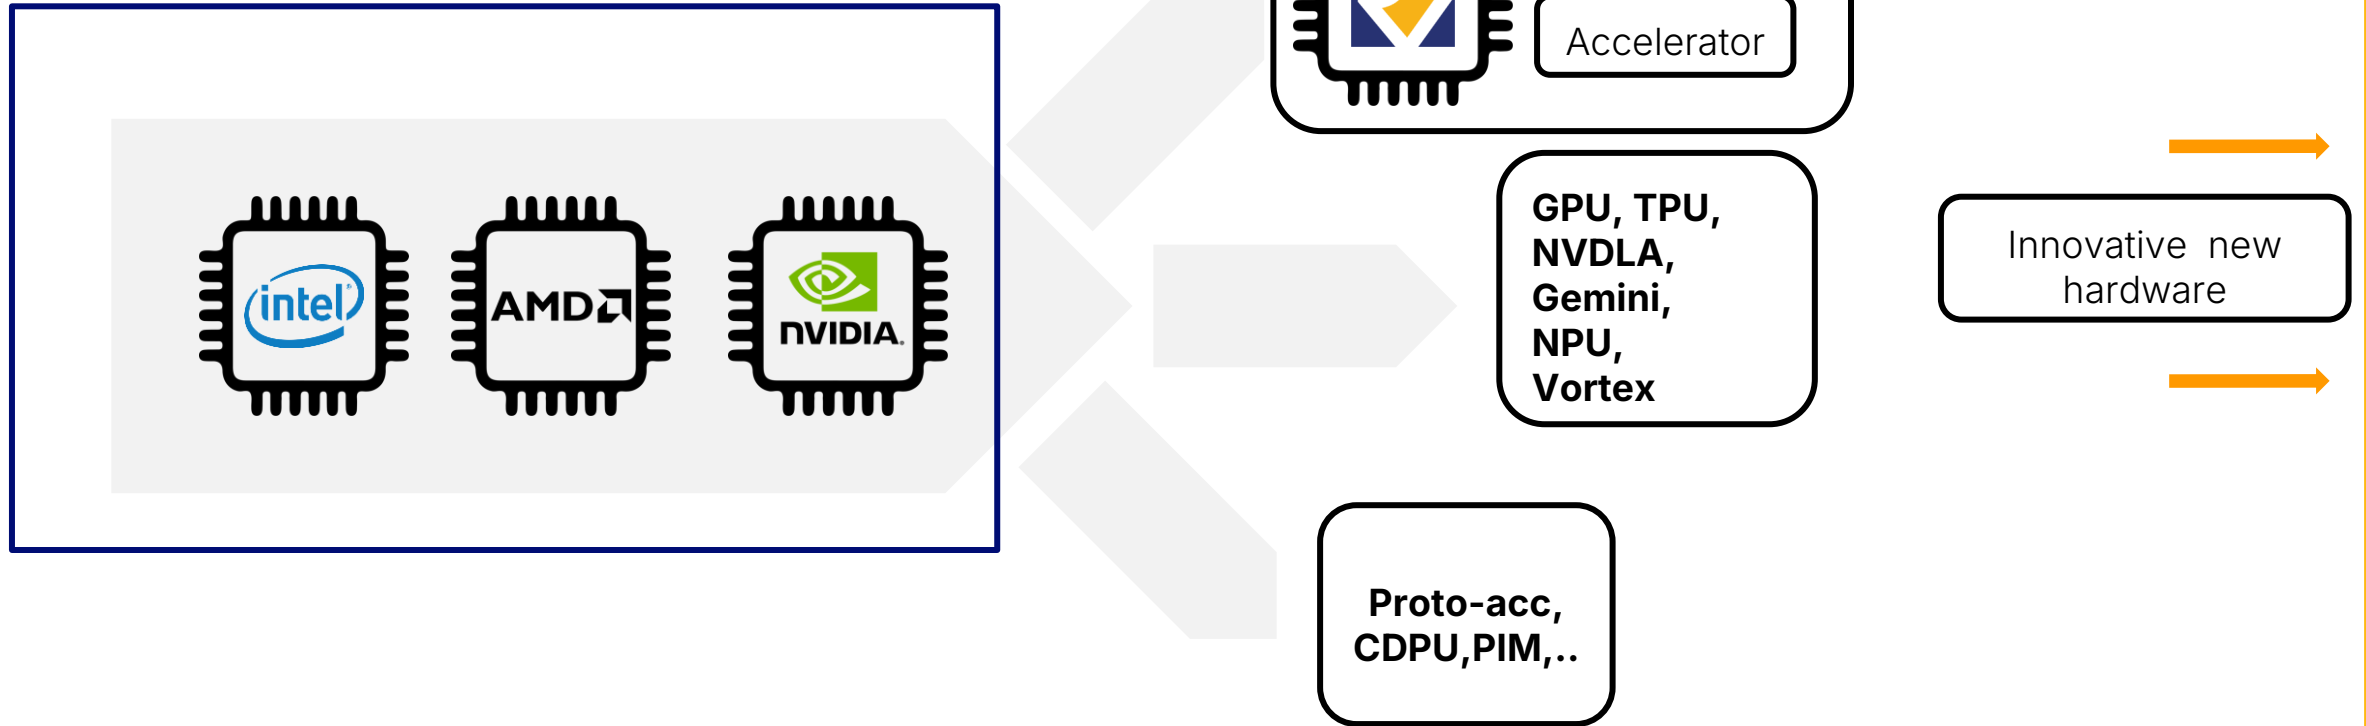
Our vision: Open-source and automatic characterization

Automatic Performance Characterization

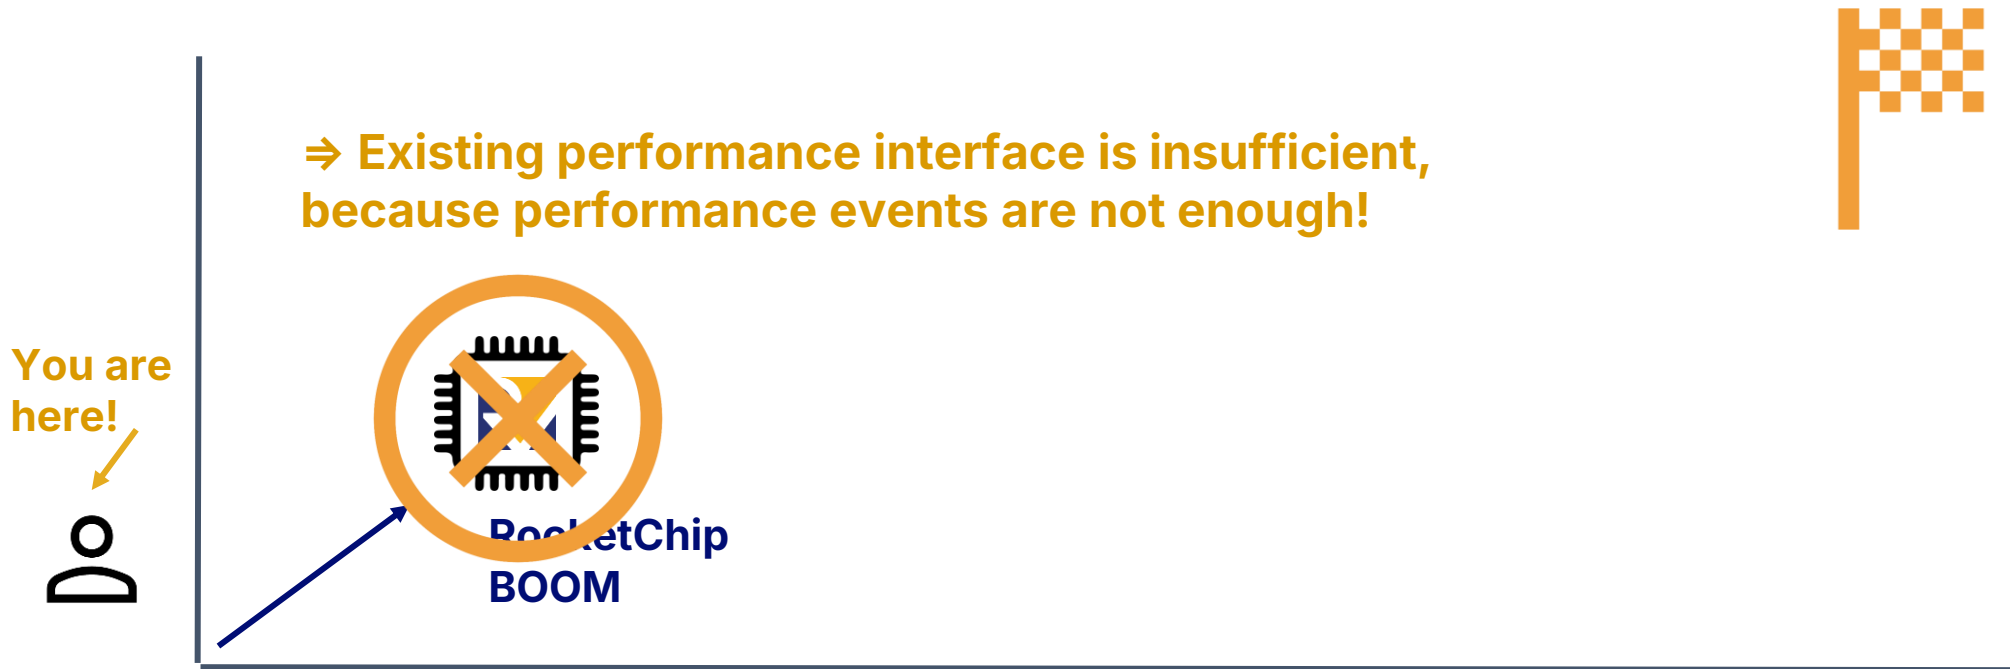


Our vision: Open-source and automatic characterization

Automatic Performance Characterization



Our vision: Open-source and automatic characterization



Key Vision:

Fully **automatic** and **generalized** performance characterization


Key Problem:

Performance characterization requires **lots of effort, hard to get right**

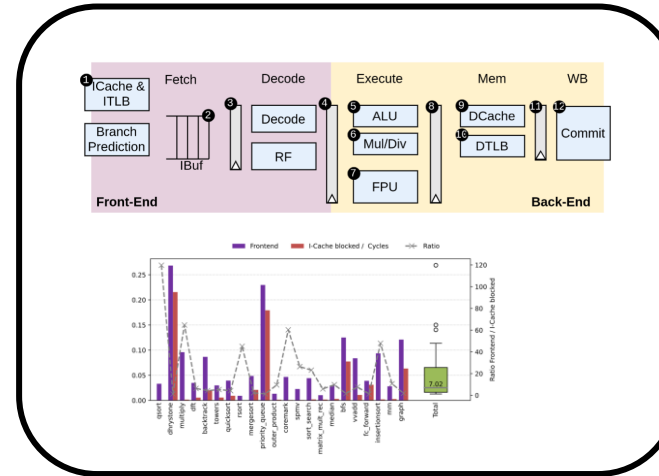
Key Contribution:

Performance characterization HPM events for **RocketCore** and **BOOM**

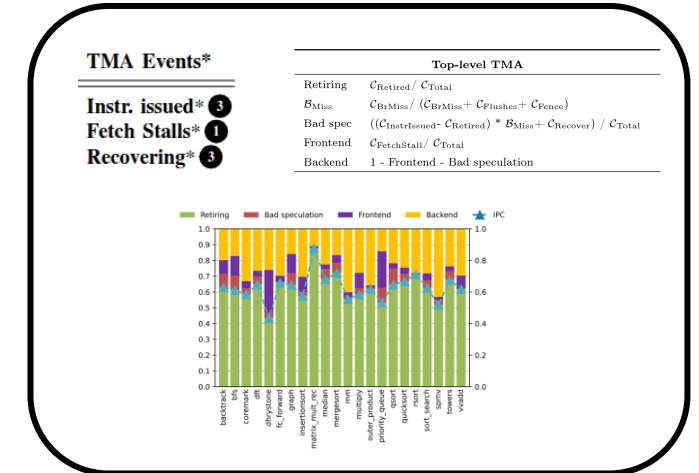
Outline



Introduction & motivation

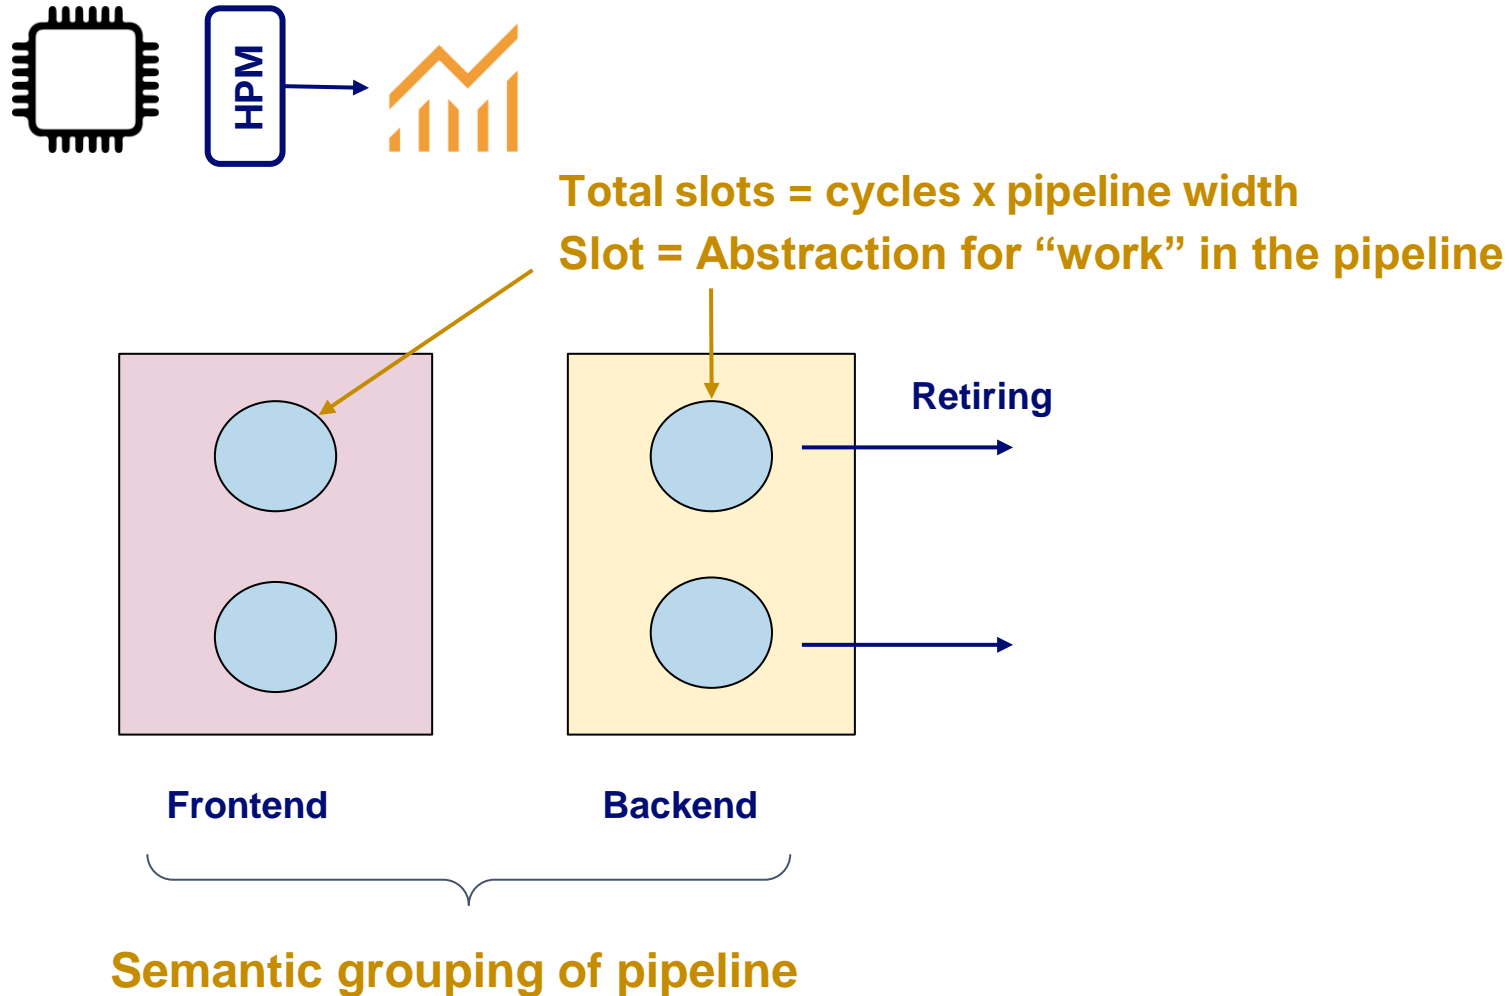


Performance characterization on RISC-V




Methodology & results


What is Top-Down Microarchitectural Analysis (TMA)?



What is Top-Down Microarchitectural Analysis (TMA)?




"Bad" slot



Frontend

"Good" slot





Backend


TMA

⇒ Categorize slots

What is Top-Down Microarchitectural Analysis (TMA)?



Frontend



Backend



"Bad" slot

"Good" slot


TMA

⇒ Categorize slots


What is Top-Down Microarchitectural Analysis (TMA)?




Frontend



Backend



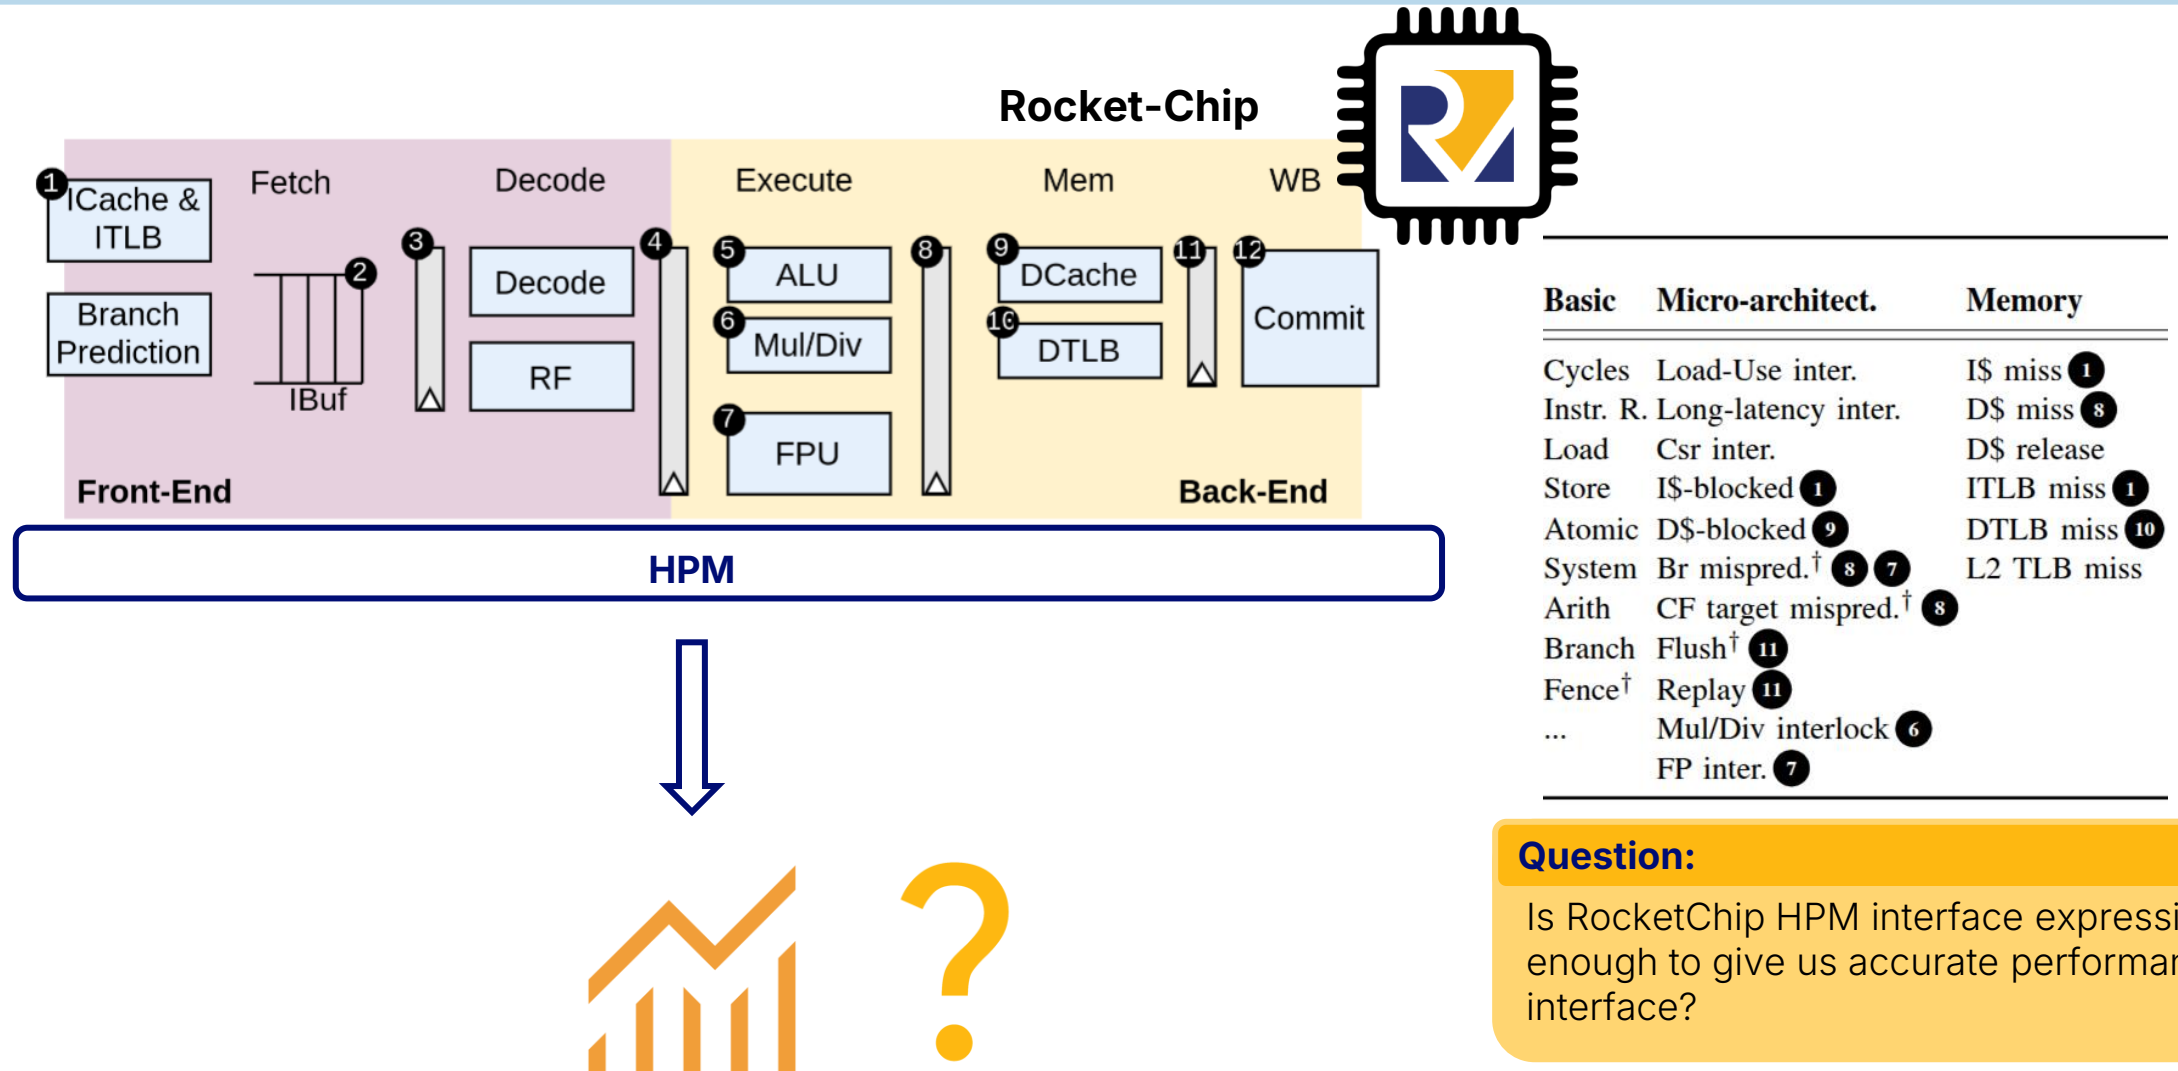
"**Good**" slot turns into "**Bad**"
slot due to bad speculation
(e.g., branch misprediction)



TMA


⇒ Categorize slots

What is the current state of performance characterization?



Is the current HPM implementation enough?

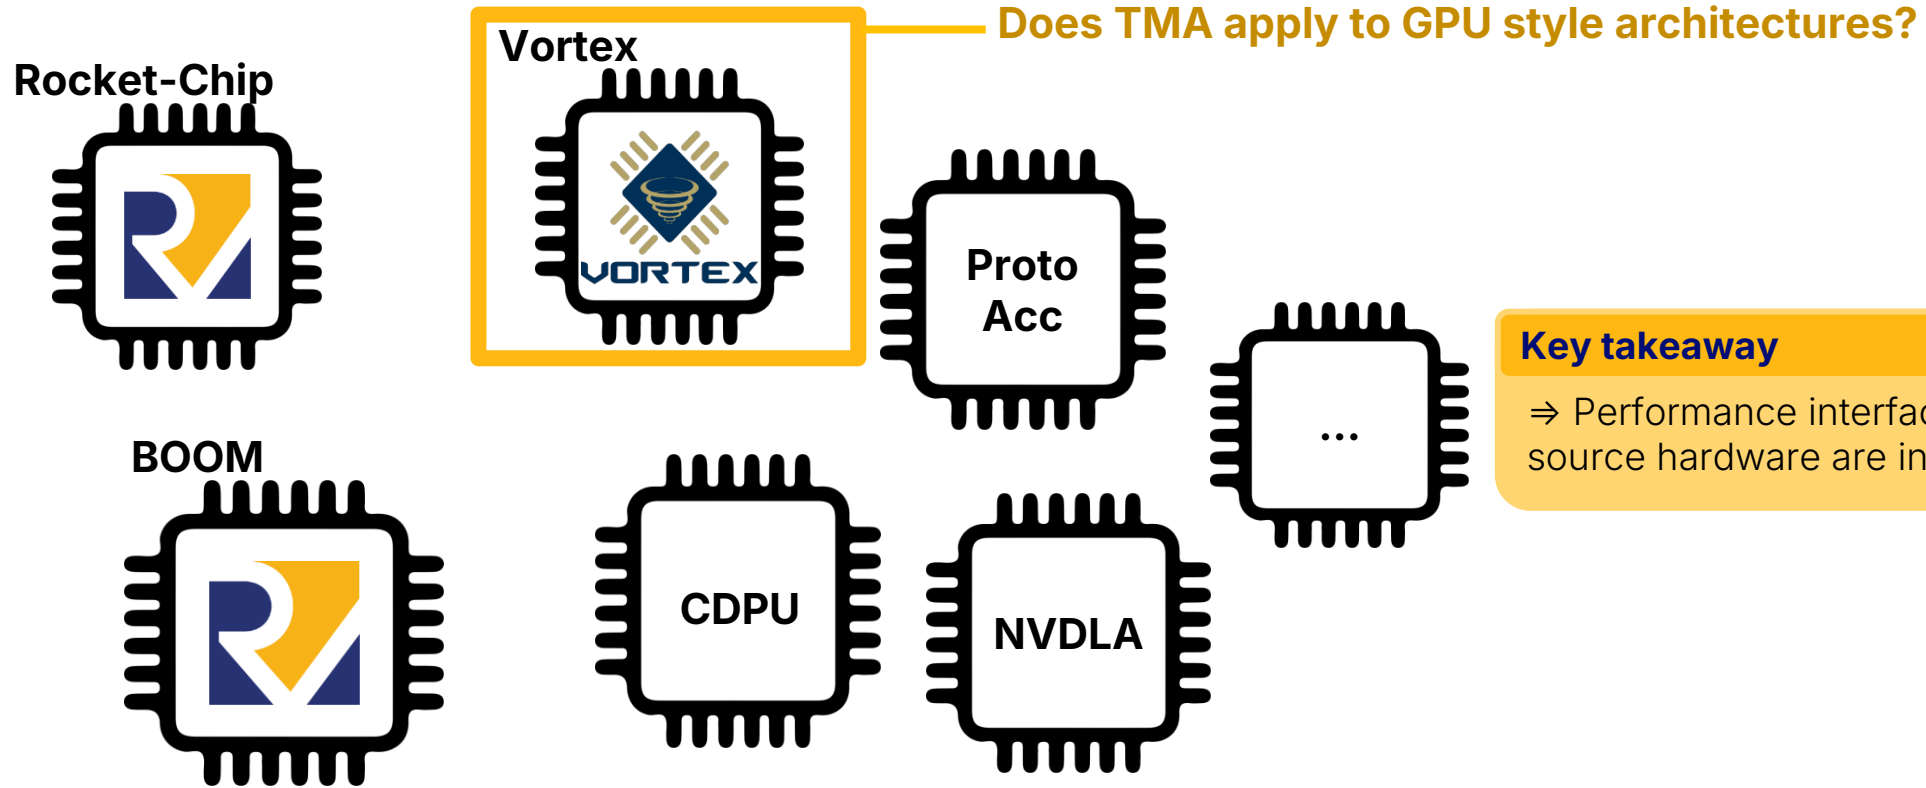
■ % slots lost in **frontend** ■ #cycles / **IS-blocked**



Key takeaway

⇒ Current frontend HPM events do not correlate with the "true" model


Is this an isolated incident?



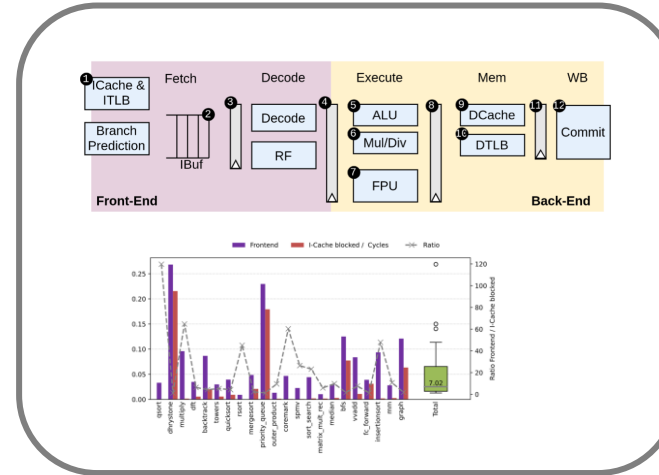
Key takeaway

⇒ Performance interfaces for most open source hardware are insufficient!

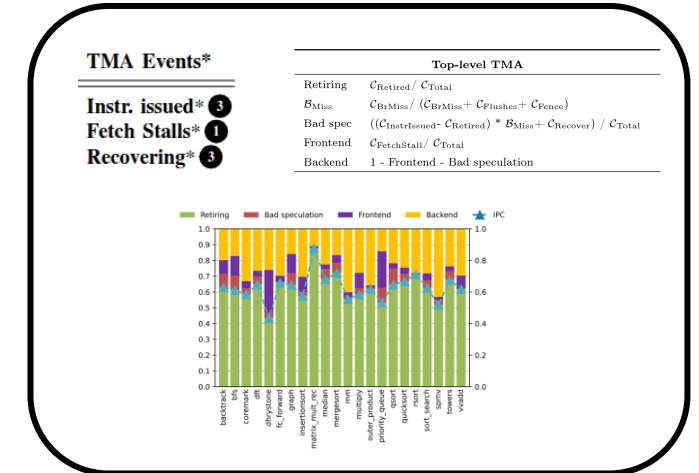
Outline



Introduction & motivation

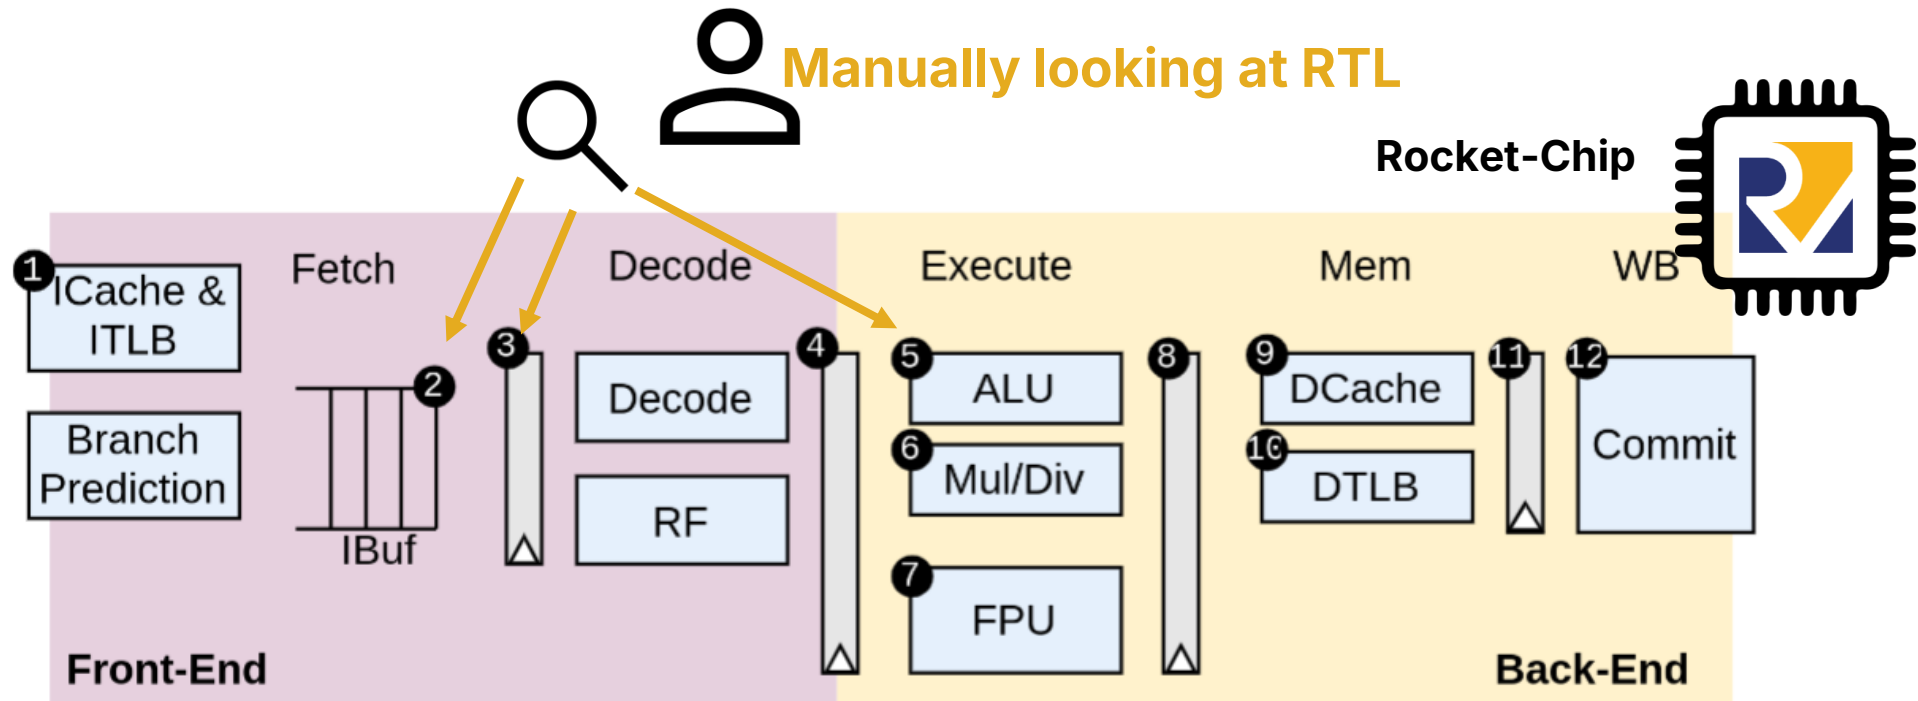


Performance characterization on RISC-V



Methodology & results

What is our current methodology?



TMA Events*

Instr. issued* ③
Fetch Stalls* ①
Recovering* ③

Top-level TMA

Retiring	$C_{\text{Retired}} / C_{\text{Total}}$
B_{Miss}	$C_{\text{BrMiss}} / (C_{\text{BrMiss}} + C_{\text{Flushes}} + C_{\text{Fence}})$
Bad spec	$((C_{\text{InstrIssued}} - C_{\text{Retired}}) * B_{\text{Miss}} + C_{\text{Recover}}) / C_{\text{Total}}$
Frontend	$C_{\text{FetchStall}} / C_{\text{Total}}$
Backend	1 - Frontend - Bad speculation

Key takeaway

⇒ This process takes multiple 2-3 months of work

What methodology is used to evaluate correctness?

⇒ Case studies to show:

(1) Can we accurately identify bottlenecks?

(1) Can we accurately predict speedups after applying optimizations?


Synthetic benchmark

```
2 beq a5,zero,.L1
3 nop
4 j .L2
5 nop
6 j .L2
7 .L2:
8 beq a5,zero,.L3
```



Always
mispredict

Always predict
correctly



Retiring Frontend Bad spec Backend



What are the results? —> This is the Baseline!



Rocket-Chip





BOOM

Key takeaway

- ⇒ (1) Creating a baseline
- => (2) Leverage insights from this design

Call for feedback: Automatic performance characterization

Manual implementations as evaluation reference



Key Problem:

Performance characterization requires **lots of effort, hard to get right**

Key Vision:

Fully **automatic** and **generalized** performance characterization

Key Contribution:

Performance characterization HPM events for **RocketCore** and **BOOM**

Contact: matthew.weingarten@columbia.edu

mattweingarten.github.io/

