



# Iris: Microarchitectural Event Database


5/20/2024

Nicolas Castaneda, Kevin He, Jerry Zhao

- Goal: Understand RTL Behavior
  - Debug increasingly complex cores
  - Add features to complex systems and identify tradeoffs
  - Identify performance bottlenecks
  - Current solutions either fail to capture sufficient information or fail to deliver it in a comprehensible manner

Instruction  
Commit Log

Waveform  
Viewer





# Motivation for Event Logs

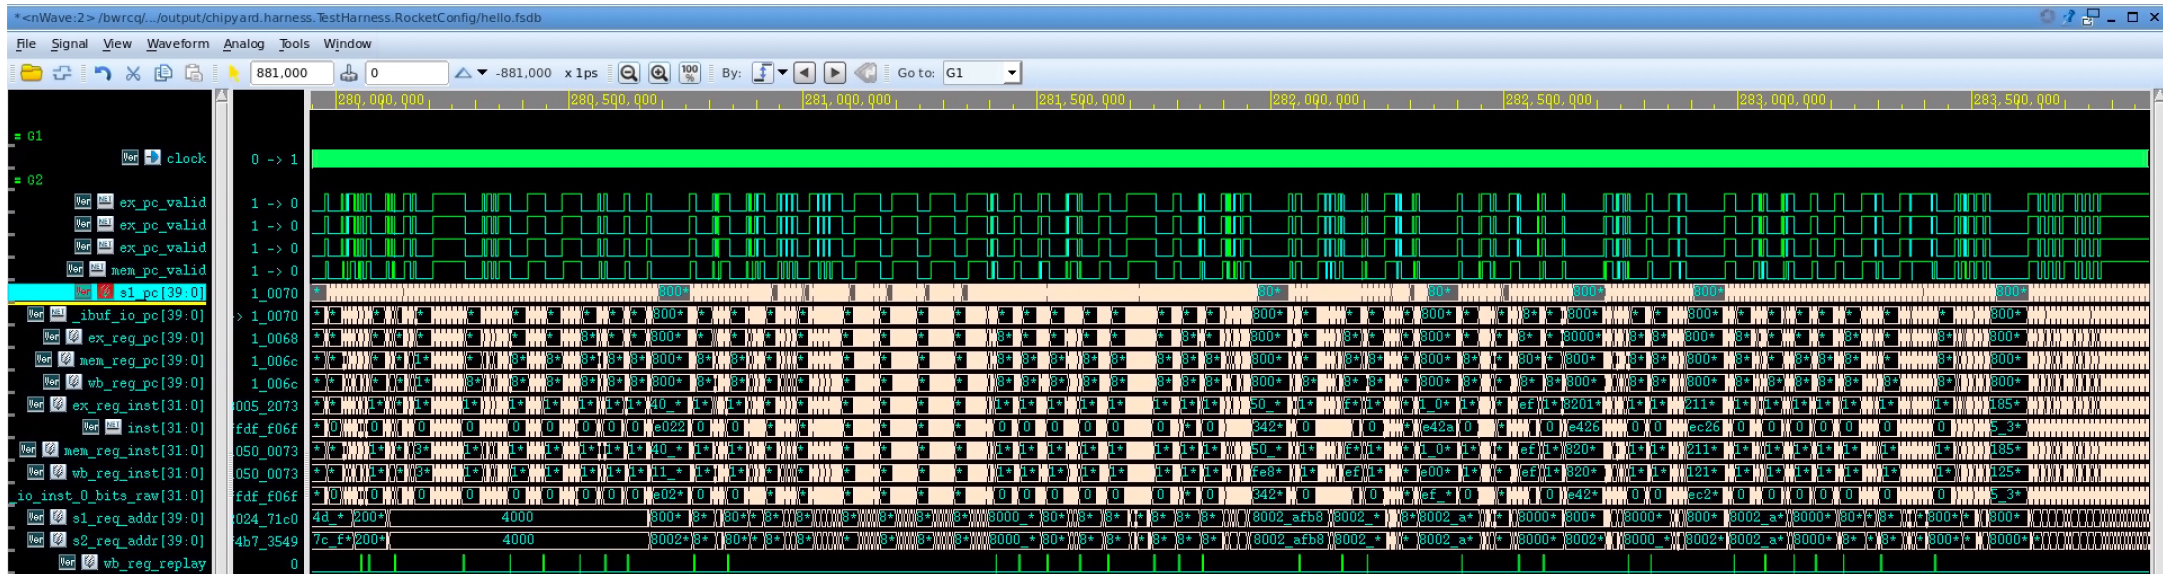
- Trace-like logs are too coarse-grained
  - No information regarding the relationship between multiple resident instructions and the microarchitecture
  - No way to reason about how instructions interact with the microarchitecture
  - Insufficient information regarding timing

```

C0: 2619 [1] pc=[0000000080001b1c] W[r 0=0000000000000000][0] R[r14=0000000000000000] R[r 0=0000000000000000] inst=[0000c701] c.beqz a4, pc + 8
C0: 2620 [1] pc=[0000000080001b24] W[r14=00000000800020e8][1] R[r 3=0000000080002678] R[r 0=0000000000000000] inst=[a7018713] addi a4, gp, -1424
C0: 2621 [1] pc=[0000000080001b28] W[r15=0000000000000000][1] R[r14=00000000800020e8] R[r 0=0000000000000000] inst=[0000471c] c.lw a5, 8(a4)
C0: 2622 [1] pc=[0000000080001b2a] W[r16=000000000000001f][1] R[r 0=0000000000000000] R[r 0=0000000000000000] inst=[0000487d] c.li a6, 31
C0: 2623 [1] pc=[0000000080001b2c] W[r10=ffffffffffffffff][1] R[r 0=0000000000000000] R[r 0=0000000000000000] inst=[0000557d] c.li a0, -1
C0: 2624 [1] pc=[0000000080001b2e] W[r 0=0000000000000000][0] R[r16=000000000000001f] R[r15=0000000000000000] inst=[04f84763] blt a6, a5, pc + 78
C0: 2625 [1] pc=[0000000080001b32] W[r 0=0000000000000000][0] R[r17=0000000000000000] R[r 0=0000000000000000] inst=[02088d63] beqz a7, pc + 58
C0: 2676 [1] pc=[0000000080001b6c] W[r13=0000000000000001][1] R[r15=0000000000000000] R[r 0=0000000000000000] inst=[0017869b] addiw a3, a5, 1
C0: 2677 [1] pc=[0000000080001b70] W[r15=0000000000000002][1] R[r15=0000000000000000] R[r 0=0000000000000000] inst=[00000789] c.addi a5, 2
C0: 2678 [1] pc=[0000000080001b72] W[r15=0000000000000010][1] R[r15=0000000000000002] R[r 0=0000000000000000] inst=[0000078e] c.slli a5, 3
C0: 2679 [1] pc=[0000000080001b74] W[r 0=0000000000000000][0] R[r14=00000000800020e8] R[r13=0000000000000001] inst=[0000c714] c.sw a3, 8(a4)
C0: 2680 [1] pc=[0000000080001b76] W[r14=00000000800020f8][1] R[r14=00000000800020e8] R[r15=0000000000000010] inst=[0000973e] c.add a4, a5
C0: 2681 [1] pc=[0000000080001b78] W[r 0=0000000000000000][0] R[r14=00000000800020f8] R[r11=0000000080001ad2] inst=[0000e30c] c.sd a1, 0(a4)
C0: 2682 [1] pc=[0000000080001b7a] W[r10=0000000000000000][1] R[r 0=0000000000000000] R[r 0=0000000000000000] inst=[00004501] c.li a0, 0
C0: 2683 [1] pc=[0000000080001b7c] W[r 0=0000000080001b7e][1] R[r 1=000000008000013a] R[r 0=0000000000000000] inst=[00008082] ret

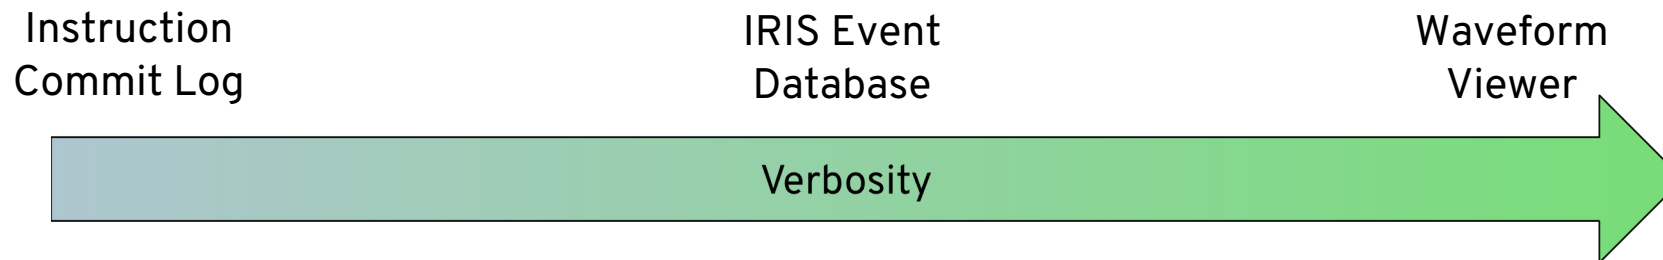
```

- Waveforms contain too much information
  - Value for every bit in the RTL design over millions of cycles
  - Requires extensive knowledge of microarchitecture signals
  - Time consuming for initial debugging
  - Dependency chains between signals aren't captured

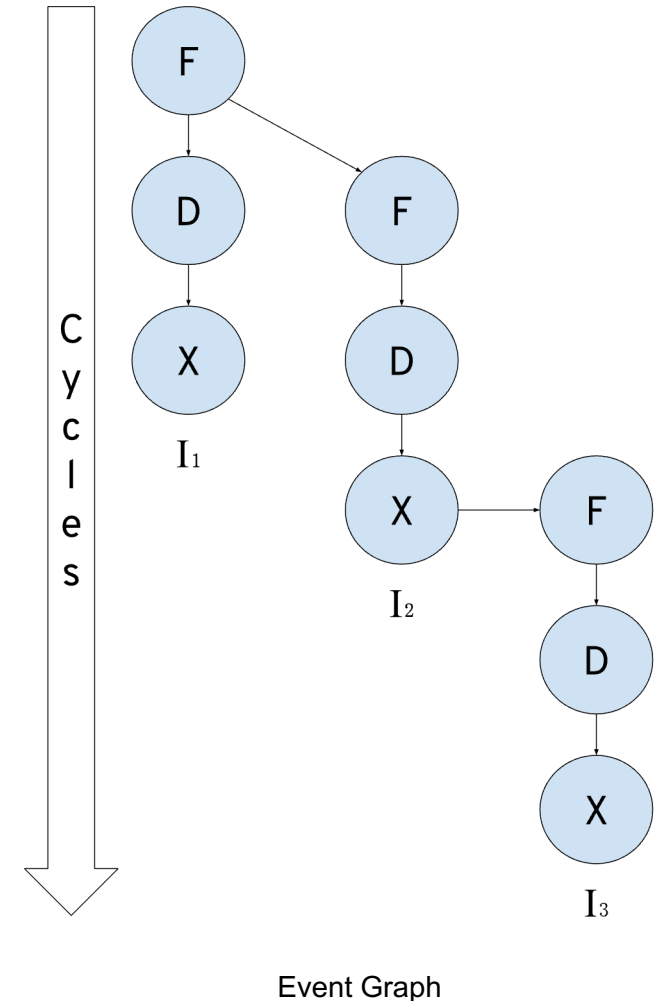


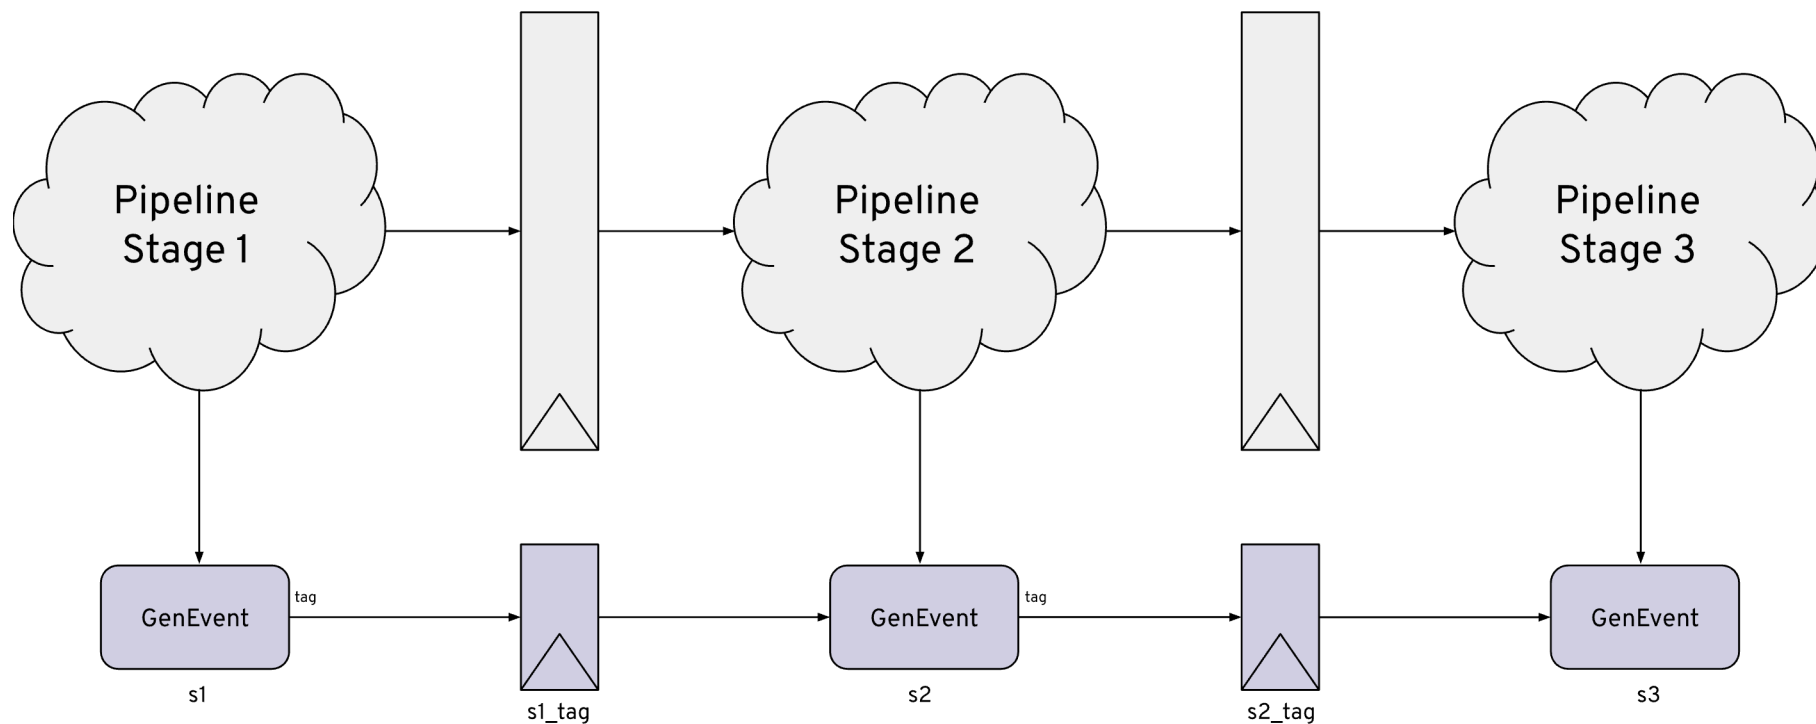
Rocket Waveforms

- Our solution: IrisDB
  - Flexible API for extracting microarchitectural events and data in RTL
  - Middle ground between waveforms and instruction commit logs
    - Provides a good starting point for debugging
  - Outputs event log for post-processing or analysis



- Represents the microarchitectural state as a sequence of dependent events
  - Microarchitecture agnostic representation
  - Nodes represent single cycle events
  - Edges represent the resolution of a hazard, allowing the subsequent event to occur
- Easily configured and analyzed
  - Exposes a standard DB schema
  - Can easily change resolution of events
  - Can easily query graph for specific information (i.e. an instruction trace)



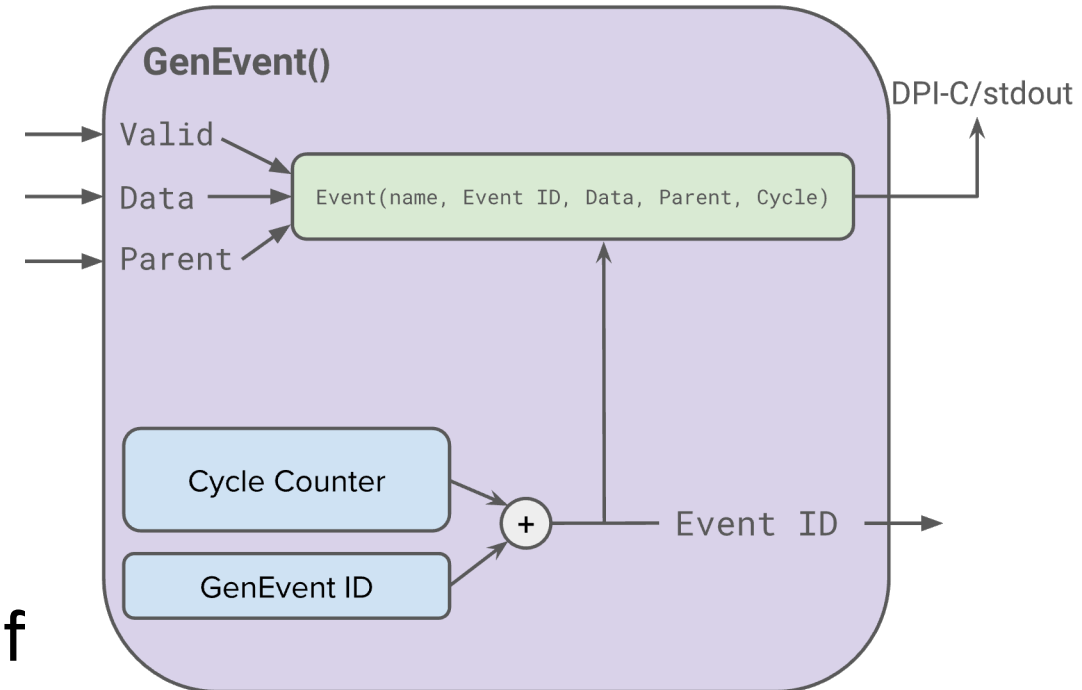


```
s1_tag := GenEvent("s1", s1_valid, s1_data, None)
```

```
s2_tag := GenEvent("s2", s2_valid, s2_data, s1_tag)
```

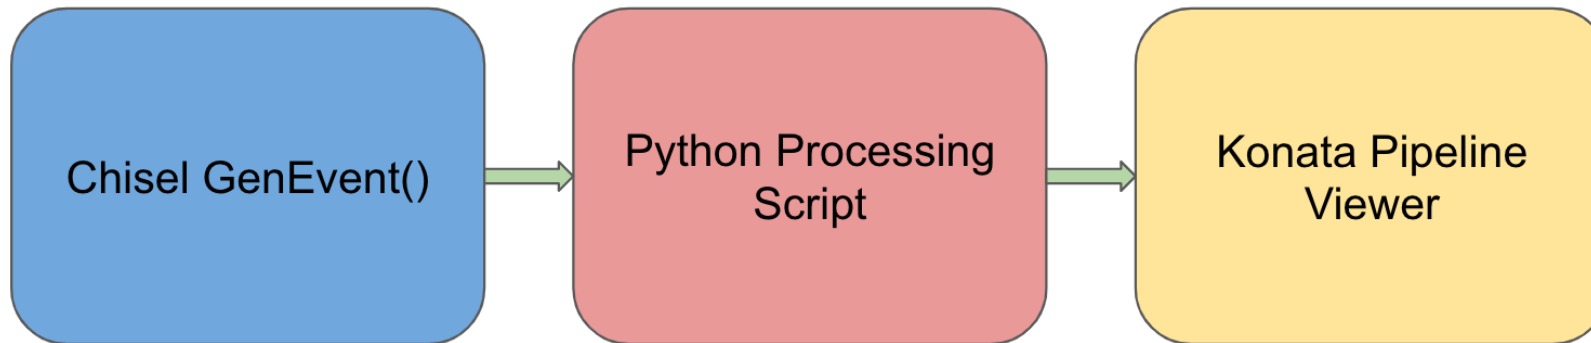
```
GenEvent("s3", s3_valid, s3_data, s2_tag)
```

- Each GenEvent Module has a:
  - Event Name
  - Valid input
  - Data input
  - Optional Parent ID input
- When valid, GenEvent logs event name, cycle, inputs, and generates a unique Event ID Tag
- Event ID Tags are the primary keys of the event database
- Module outputs Event ID tag in RTL





- Run RTL simulation with GenEvent annotated architecture
- Reconstruct the event graph using NetworkX
- Perform depth-first-search to construct the instruction traces
- Format sequences into a Konata log file
- Use Konata application for waterfall visualization of instruction execution



GenEvent API to Konata visualizer flow



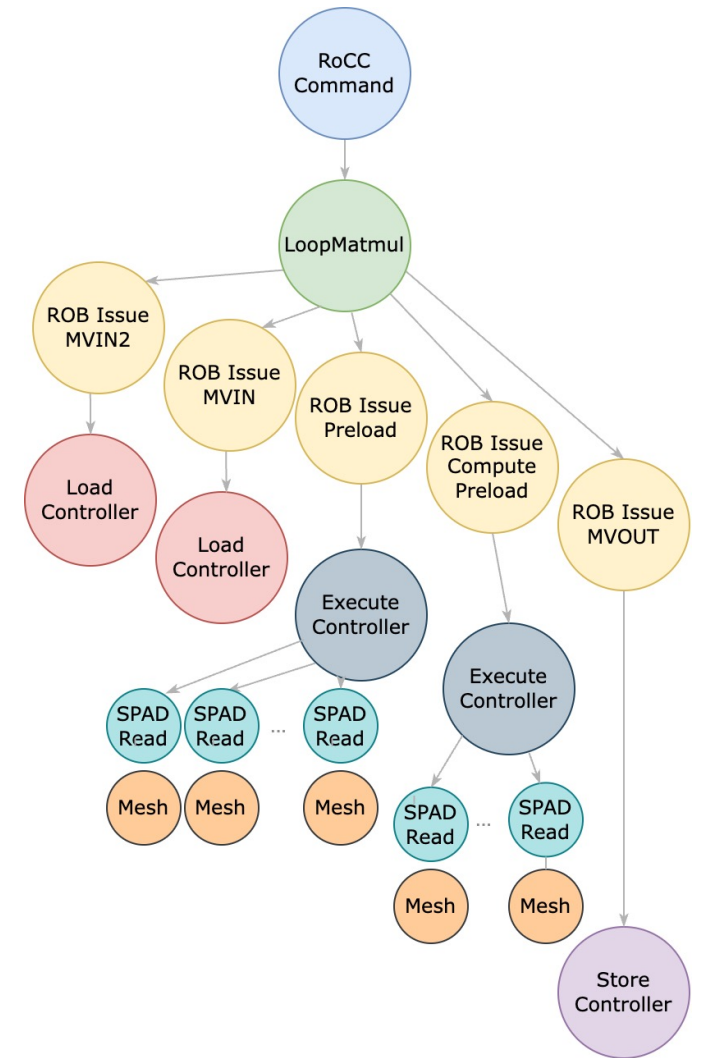
# Application: Konata Pipeline Viewer



[549320, 3575]

Rocket Konata Pipeline Visualization

- Microarchitecture agnostic:
- Sodor Educational Cores
  - 1 stage, 2 stage, 5 stage, and microcoded cores annotated
- Rocket In-Order Core
  - Integer, mul/div, and cache request/response pipelines annotated
- Gemmini Accelerator
  - Load, Store, and Execution controllers, LoopMatmul and LoopConv FSMs, scratchpad reads/writes, mesh



Gemmini Event Graph  
(Y-axis is time)



## Conclusion

- Extensible RTL event logging API
- Flexible graph event representation
- Implemented in Chisel with GenEvent
- Graph to visualization flow with Konata
  
- Questions?



- Demo!



# Acknowledgements

- Thank you, Vighnesh Iyer, Joonho Whangbo, and Ethan Gao for your help!