RISC-V Unified DB

Derek R. Hower (1) Afonso Oliveira (2)

(1) Qualcomm, Inc. (2) Synopsys, Inc.

Abstract

This work presents the RISC-V Unified Database (UDB), a step towards a unified, machine-readable source of truth for the RISC-V Specification. Currently, the RISC-V specification (including both ratified and de-facto parts) is scattered across multiple repositories and cloud storage files. Little of it is machine-readable, and in many cases information is duplicated in error-prone ways. This presents a barrier to further RISC-V growth as it is increasingly difficult to understand (and especially **verify**) the complex and growing specification. By gathering the specification in a single database, UDB provides the means to quickly find information and to generate artifacts directly from an authority. Towards that end, UDB currently includes over ten generators, including ones that produce ISA manuals, instruction and CSR indices, and an Instruction Set Simulator (ISS) - all from the same source. Though tremendous progress has been made in the last year, UDB is still a work in progress, and we are actively looking for increased community participation.

Specification Challenges

- Multiple disconnected documents and repositories
- Prose-only sources with no machine-readable schema
- Extensions unversioned in releases
- No single source of instructions, CSRs or pseudo-ops
- Conflicting names/encodings across manuals

Key Implications

- Custom docs for each CPU require manual refactoring
  - Error-prone and not verifiable
- Testing gaps: no auto-gen of compliance suites

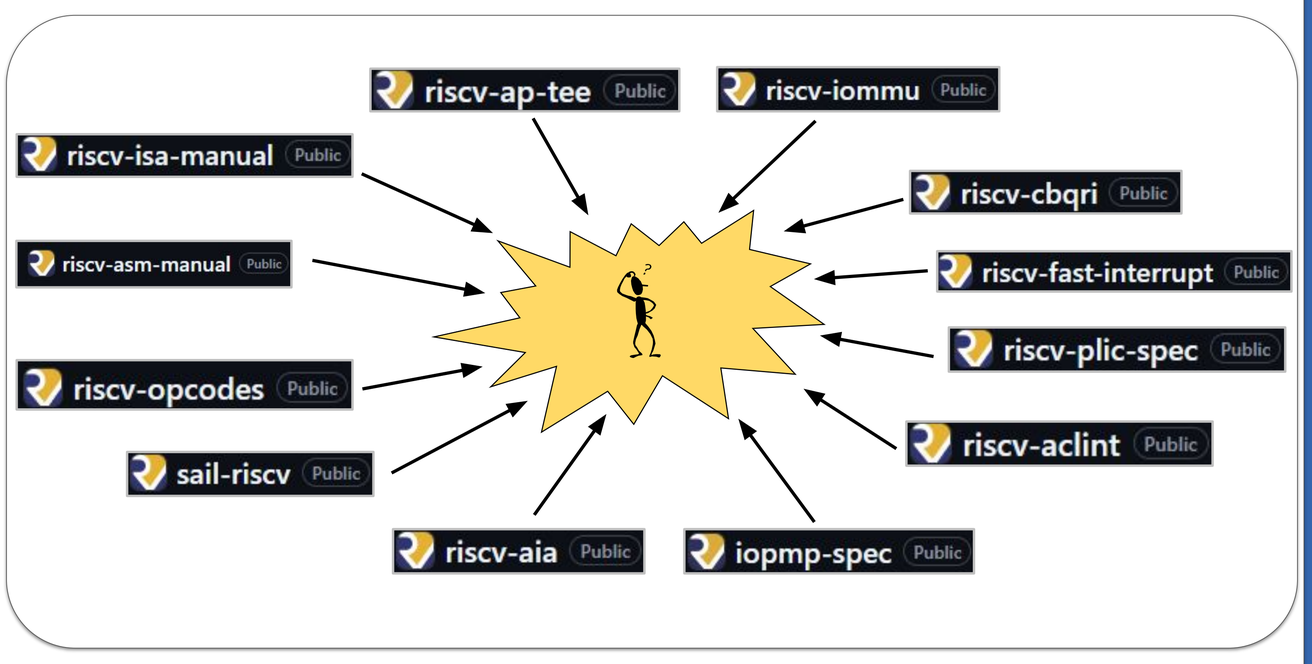

RISC-V Unified DB Overview

- Unified, machine-readable source-of-truth
- Schema-enforced artifacts (instructions, CSRs & profiles)
- Split information and tools - allows downstream usage
- Overlay support for custom extensions
- Cross-validated against established sources (e.g. LLVM)

RISC-V Unified DB Outputs

- Documentation - ISA, extension and profile manuals, Instruction and CSR appendix
- Code - assembler definitions, test suites, decoder constants
- ISS - ISS hart integrated with Renode

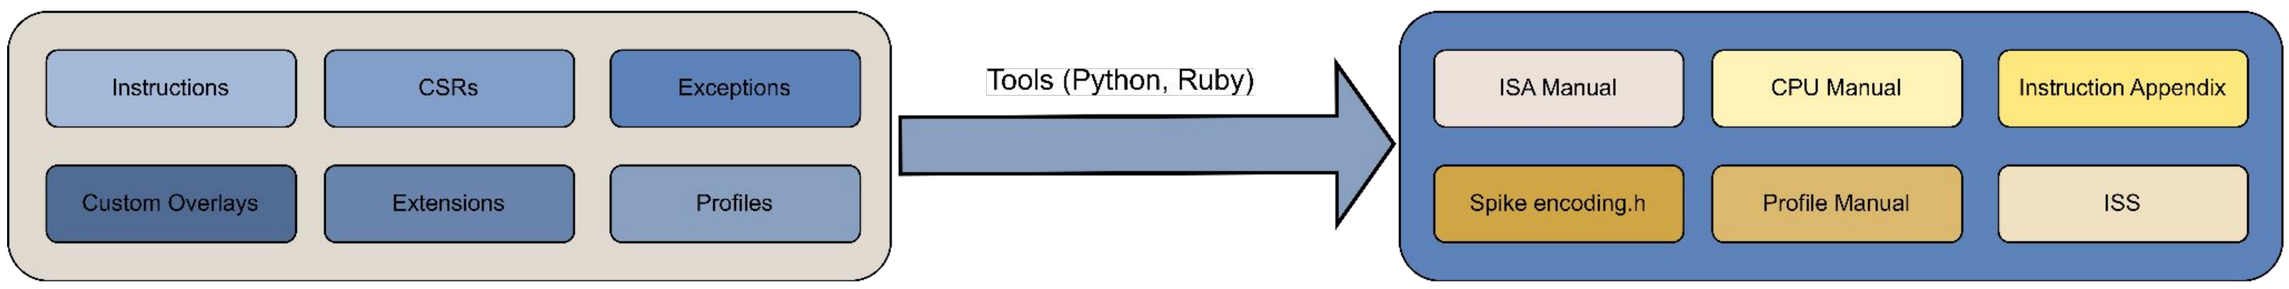

Use Cases

- Documentation - Qualcomm's Xqci, Synopsys' ARC-V
- Substituting RISC-V opcodes: ongoing for Golang, ACT, Spike, OpenHW CVW verification
- New tools can use the structured info to build upon
- UDB integrated Instruction Set Simulator

Community Endorsement

- Official RISC-V SIG meeting weekly for 10 months - six members regularly represented
- Collaborating with multiple SIGs: CSC, Doc, and Profiles
- BSD-3-licensed GitHub repo with active issues & discussions
- Weekly public RVI meetings open to new contributors

JOIN US!

RISC-V® QUALCOMM® SYNOPSYS®