RISC-V SUMMIT
EUROPE 2025

RISC-V and Software

RISC-V 101

Nathan Egge

Software Readiness

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor
- Examples
  - Seagate custom SOCs in HDD [2020]

Seagate says it's designed two of its own RISC-V CPU cores – and they'll do more than just control storage drives

Processors are 'critical to future products', set to perform tasks for host systems

Chris Mellor Tue 8 Dec 2020 // 18:00 UTC

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor

- Examples
  - Seagate custom SOCs in HDD [2020]
  - Meta custom RISC-V transcoding [2023]

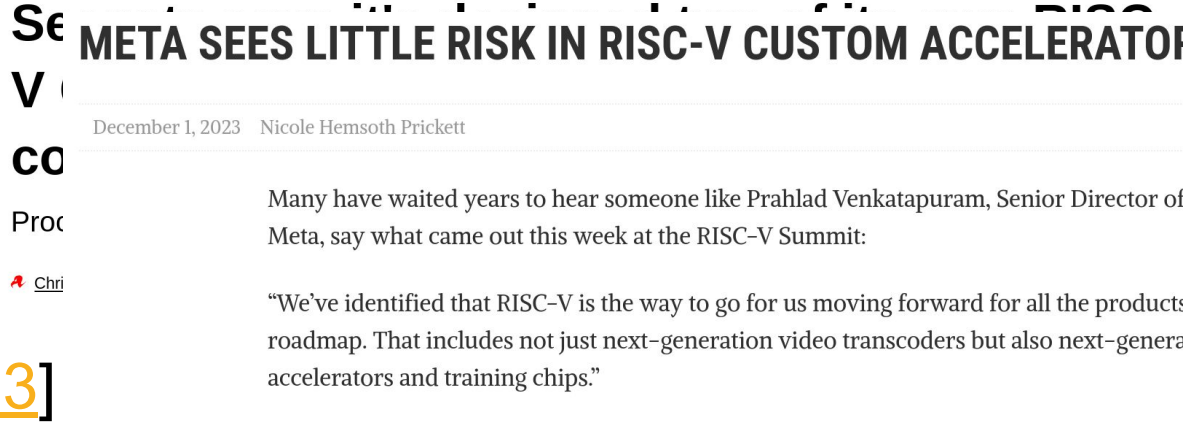

META SEES LITTLE RISK IN RISC-V CUSTOM ACCELERATOR

December 1, 2023 Nicole Hemsoth Prickett

Many have waited years to hear someone like Prahlad Venkatapuram, Senior Director of Meta, say what came out this week at the RISC-V Summit:

"We've identified that RISC-V is the way to go for us moving forward for all the products roadmap. That includes not just next-generation video transcoders but also next-genera accelerators and training chips."

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor

- Examples
  - Seagate custom SOCs in HDD [2020]
  - Meta custom RISC-V transcoding [2023]
  - Nvidia RISC-V in GPUs for 9 years! [2024]

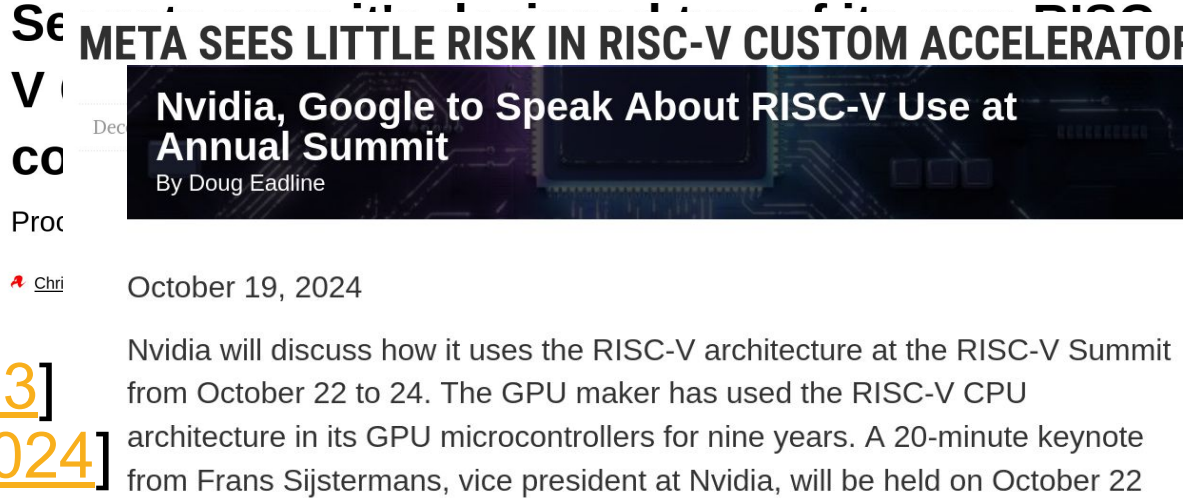

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor

- Examples
  - Seagate custom SOCs in HDD [2020]
  - Meta custom RISC-V transcoding [2023]
  - Nvidia RISC-V in GPUs for 9 years! [2024]
  - Meta testing RISC-V for AI accelerator [2025]

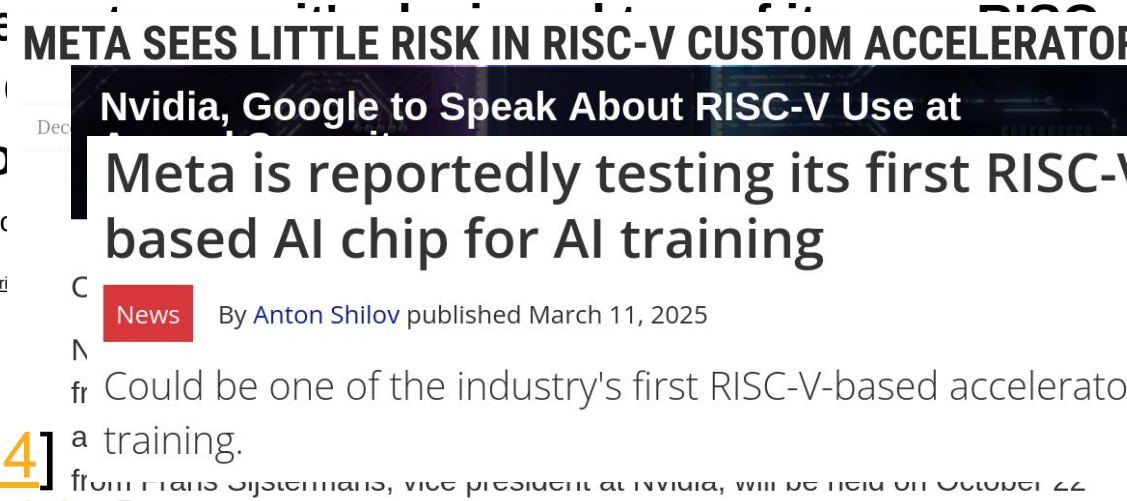

RISC-V Software Success Today

- RISC-V has good adoption in microcontrollers
  - Single purpose application
  - Limited set of standard extensions needed, custom instructions
  - RTOS or Bare Metal
  - Control often a driving factor

- Examples
  - Seagate custom SOCs in HDD [2020]
  - Meta custom RISC-V transcoding [2023]
  - Nvidia RISC-V in GPUs for 9 years! [2024]
  - Meta testing RISC-V for AI accelerator [2025]
  - Billions and billions of cores shipped! \o/

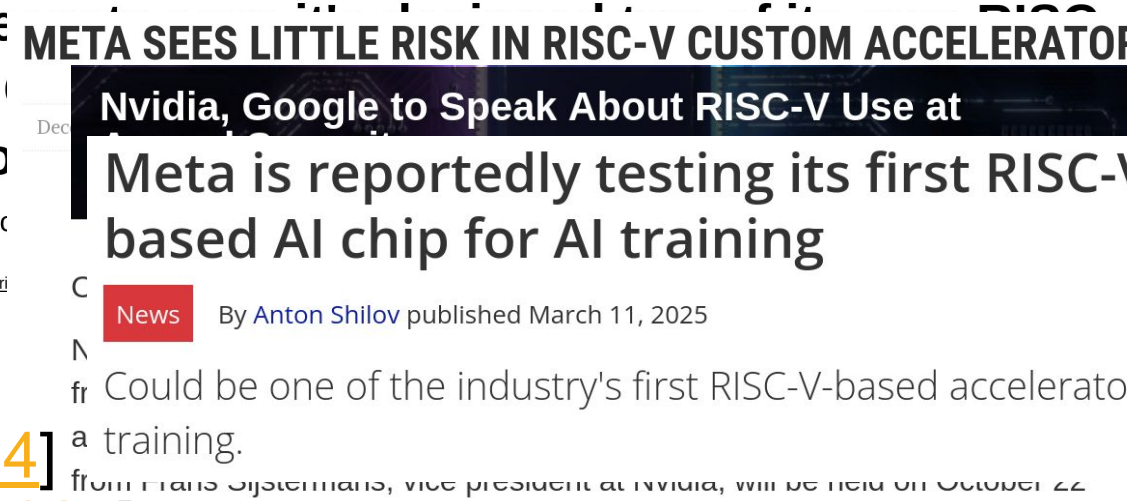

Operating Systems

- Most systems software in C/C++ without significant specialization
  - libc + syscalls good enough for POSIX support
- Examples
  - Linux: Debian, Fedora, Gentoo, etc.
  - Embedded: Yocto
  - RTOS: Zephyr, FreeRTOS
- Enterprise Linux Support is here!
  - Upstream distros like Fedora worked to enable downstream counterparts
- What % of Linux packages are enabled?

Operating Systems

- Most systems software in C/C++ without significant specialization
  - libc + syscalls good enough for POSIX support
- Examples
  - Linux: Debian, Fedora, Gentoo, etc.
  - Embedded: Yocto
  - RTOS: Zephyr, FreeRTOS
- Enterprise Linux Support is here!
  - Upstream distros like Fedora worked to enable
- What % of Linux packages are enabled?
  - As of July 2024, 97% in Debian!

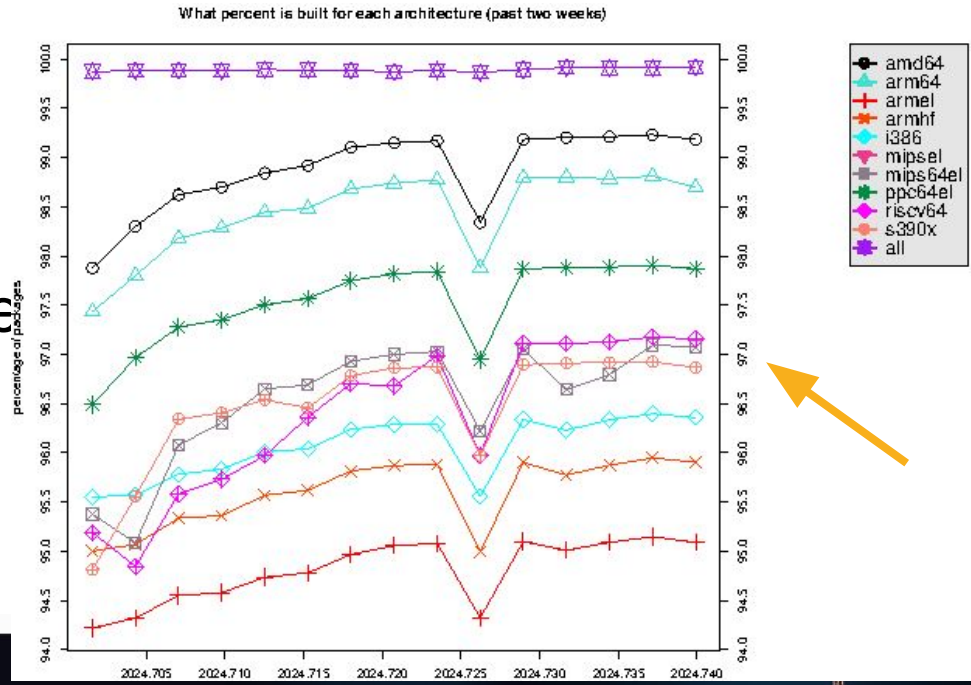

Operating Systems

- Most systems software in C/C++ without significant specialization
  - libc + syscalls good enough for POSIX support
- Examples
  - Linux: Debian, Fedora, Gentoo, etc.
  - Embedded: Yocto
  - RTOS: Zephyr, FreeRTOS
- Enterprise Linux Support is here!
  - Upstream distros like Fedora worked to enable
- What % of Linux packages are enabled?
  - ~~As of July 2024, 97% in Debian!~~
  - As of May 2025, >98% in Debian!

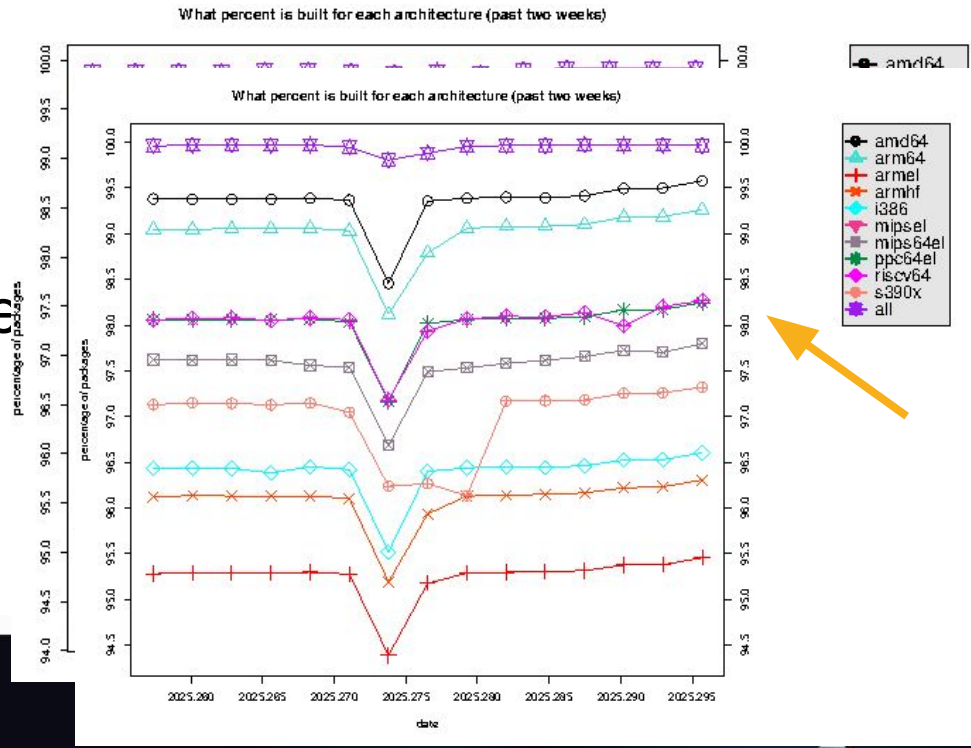

Linux Kernel

- Active work to enable RISC-V in Linux
  - Early HWCAP feature detection, but limited to 32 long bit-vector
  - RISC-V Vector 1.0 support in 6.5
  - hwprobe() syscall added in 6.6
  - PMU support, pointer masking, bitmanip, and others on-going
- SOC support still a challenge
  - Most developer boards come with a heavily modified vendor kernel
  - Requires “bring-up” to get suitable environment for development
  - No generic RISC-V kernel in Debian, can still replace rootfs

Application Software

- Must “just work” on range of heterogeneous hardware
  - Scale from IOT device, to laptop, to HEDT, to server class
  - Multimedia: IOT camera, watch streaming video, multi-channel transcoding
    - Potentially all use the same libraries
- Written in managed or interpreted languages
  - Most runtimes work, but performance limited
    - No or partial JIT, native .so not compiled for RISC-V, e.g. Python
  - Java and Go getting performance optimizations through RISE
- Variable set of workloads, performance critical execution
  - Really only one mechanism for performance, SIMD aka RISC-V Vector 1.0
  - May not contain the same extensions, or same vector length

Software Readiness

- Very much application or “domain” dependent
  - Do not need everything to be perfect, just enough to get work done

- RISC-V Software Ecosystem Dashboard [1]
  - Attempts to catalog key software components based on
    - Enabled: RISC-V base support established
    - In Progress: Active development underway
    - Optimized: Software performant on RISC-V
    - TBD: No commitment to RISC-V enablement

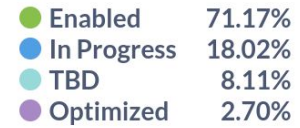

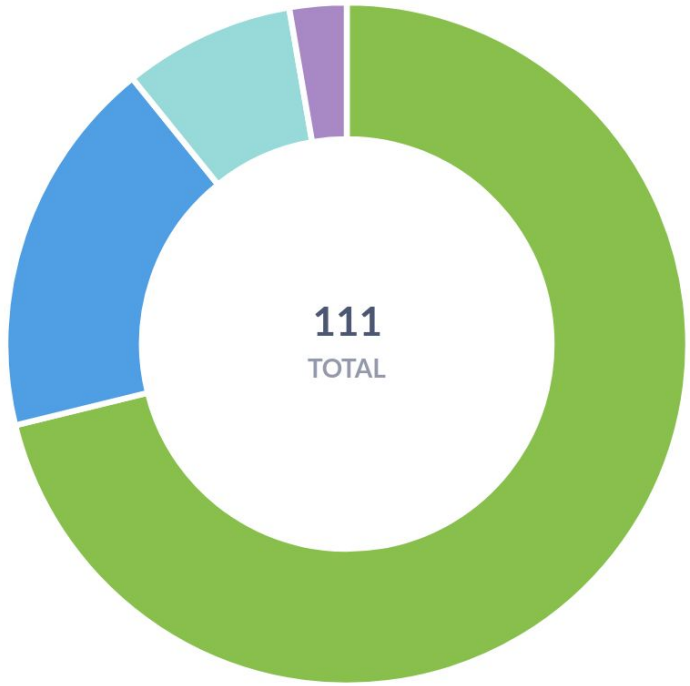

[1] https://tech.riscv.org/software-ecosystem

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)
  - Does it run

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)
  - Does it run
  - Does it run correctly (unit and integration tests)

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)
  - Does it run
  - Does it run correctly (unit and integration tests)
  - Does it run correctly on my hardware

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)
  - Does it run
  - Does it run correctly (unit and integration tests)
  - Does it run correctly on my hardware
  - Does it run correctly on my hardware with enough performance

Software Readiness

- In practice this is hard to measure, readiness can also mean
  - Does it build (configure and compile)
  - Does it run
  - Does it run correctly (unit and integration tests)
  - Does it run correctly on my hardware
  - Does it run correctly on my hardware with enough performance
- Optimized software is also nebulous
  - Performance often achieved over time through incremental improvements
  - Unclear what the lower bound is on compute
    - dav1d-1.5.1 still improving 7 years later

Languages and Runtimes

Language	Implementation	Status	Notes
C/C++	GCC, Clang	Good	RVV Intrinsics, tunings per target, autovectorization
Javascript	v8, Spidermonkey	Works	Upstreamed, v8 wiki, spidermonkey initial support Plenty of performance work ongoing
WebAssembly	v8, Cranelift	Works	Upstreamed, available, plenty of work ongoing still
Go	golang	Good	Since Go 1.16 Supports also cgo.
Rust	rustc (LLVM, Cranelift)	Works	But no RVV intrinsics yet, cpu features runtime detection coming soon
Python	CPython, pypy	Good	You can run pytorch just fine, jit backend for pypy
Java	OpenJDK	Good	Tracker, Apertus Distributes LTS for Java 11, 17, 21 and 22
Haskell	GHC	Works	Tracker, both LLVM and NCG backends are supported
Erlang	otp	Works	No JIT yet

Additional Tools

Tool	Type	Status	Notes
GDB	Debugger	Works	Does not print RVV registers yet
LLDB	Debugger	Works	Less available by default
linux-perf	Profiler	Sort-of	On some platforms only custom events are available
rr-project	Debugger	Missing	Tracker, Could work for cpu with **Zacas** support
mold	Linker	Good	Works
wild	Linker	Initial	Tracker, Works ongoing

Ways to Improve Performance

Auto-vectorization

- Pros: Compiler does all the work
  - Performance can get better with newer compilers
- Cons: Language and code have to give hints
  - Scalar code often does not map to efficient vector operations
  - Compiler support may not always be present (or bug free)

Intrinsics

- Pros: Code uses primitives present in the instruction set
  - Same language as the rest of the code, easy to reason about and debug
  - The instruction scheduling should be optimal and tuned for the target
- Cons: Compiler support may not always be present
  - Intrinsic version changes force code updates 0.11 -> 0.12

Pure Assembly

- Pros: Full control, no chance of mis-compilation
  - Overcome ABI limitations, not everything representable with intrinsics
- Cons: Must account for everything: instruction scheduling, register allocation, etc
  - Difficult to write, difficult to debug, difficult to modify

Code Size Considerations

- Can trade binary size for more specialization
  - Multiple implementations selected at runtime, even with same extensions
  - e.g., Intrinsics + Function Multi-Versioning for micro-architecture tuning
- Some deployments sensitive to binary size, no universal solution
  - Desktop application on DVD may be fine
  - Mobile applications highly sensitive to download time
  - Middleware vendors differentiate on binary size
  - Server can and often rebuild everything bespoke for hardware
- Reasonable, domain-specific tradeoffs should be made

Conclusions

- Most software “ready” in that it will build and run on Linux
- Good performance is domain specific, need to test on target HW
- Toolchain support is good and RISC-V parity steadily improving
- Many software workloads will run fine as-is and unmodified
- Top priority for RISC-V enablement is more optimizations, e.g., for V (vector), Zb{a,b,c,s} (bit manip) and Zvk (vector crypto)
- You can start today!

Getting Started Guide

RISC-V SUMMIT
EUROPE 2025

Ways to Develop for RISC-V

- FPGA
  - Pros: Cycle accurate model of hardware
  - Cons: Whole system booting very slow, 10's of MHz
- Emulation
  - Pros: First to get RVI extensions, flexible and configurable
  - Cons: Essentially impossible to measure performance
- Hardware
  - Pros: Performance will match what ships exactly,
  - Cons: Long lead time, may overfit microarchitecture, limited vector length

RISC-V Summit EU 2024

- Presented single slide on manually prebuilt developer images
  - Since then work has focused on build automation

Prebuilt Developer Images

- Facilitate development by providing up-to-date toolchains for building and testing
  - Latest toolchain package versions
    - clang-18.1.5
    - gcc-13.2.1_p20240503
    - rust-1.77.1
    - binutils-2.42
    - cmake-3.29.3
    - python-3.12.3
    - perl-5.38.2
    - git-2.45.1
    - subversion-1.14.3
  - Kendryte K230 and Banana Pi BPI-F3

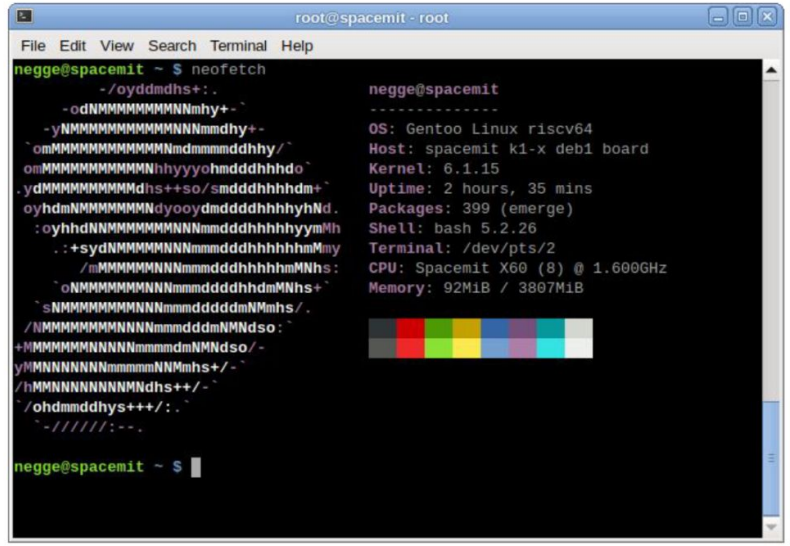

[1] https://people.videolan.org/~negge/canaan-3G-2024-04-08.img.xz
[2] https://people.videolan.org/~negge/spacemit-4G-2024-05-15.img.xz

ROMA II image coming soon!!

RISE

RISC-V SUMMIT EUROPE 2025

Gentoo Developer Images

- Project Goal
  - Fastest way to create bootable images with **up-to-date** toolchains!
- *Key Idea*: Automate developer image building
  - Now takes only ~300 minutes (!) to cross compile bootable image
  - U-Boot + Kernel + ~330 software packages
  - Bespoke CFLAGS possible, testing surfaced several gcc autovector issues [1]
- Partnered with Luca Barbato, RISC-V Gentoo developer
  - Fixed multiple issues unblocking full cross compilation **<-- no other distro has this**
  - Right now BPI-F3, ROMA-II, Orange Pi RV2D and others based on SpacemiT K1
  - Now building *everything* with **-O3 -march=rv64gcv_zvl256b**

New with gcc-15

[1] GCC Bug 116242 - [meta-bug] Tracker for zvl issues in RISC-V

Developing on ROMA-II with RVV 1.0

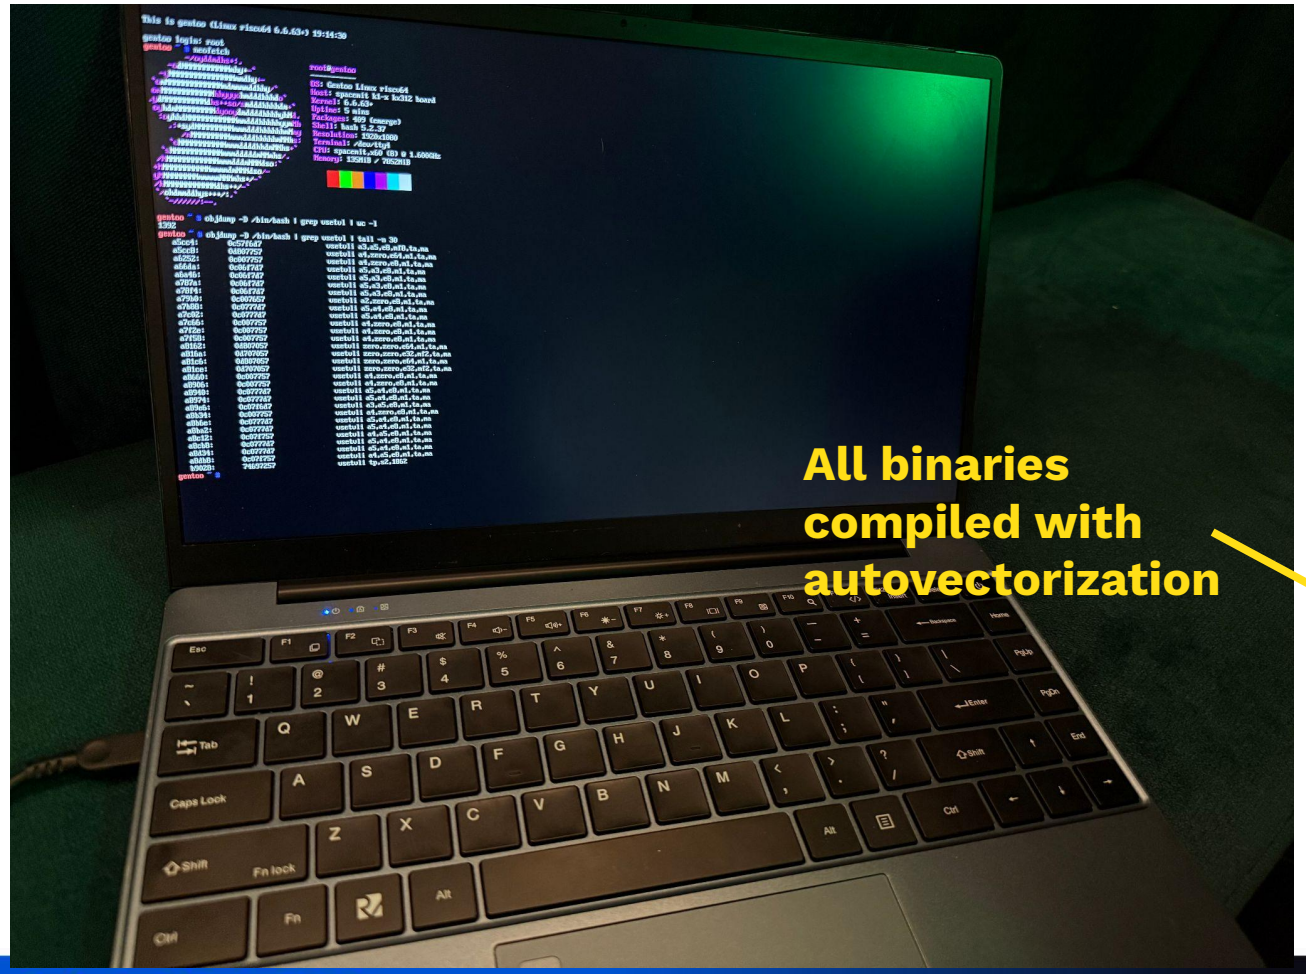

All binaries compiled with autovectorization

```
This is gentoo (Linux riscv64 6.6.63+) 19:14:30

gentoo login: root
gentoo ~ # neofetch
                                    root@gentoo
                                    -----------
                                    OS: Gentoo Linux riscv64
                                    Host: spacenit k1-x k3i2 board
                                    Kernel: 6.6.63+
                                    Uptime: 5 mins
                                    Packages: 409 (emerge)
                                    Shell: bash 5.2.37
                                    Resolution: 1920x1080
                                    Terminal: /dev/tty4
                                    CPU: spacenit,x60 (8) @ 1.600GHz
                                    Memory: 135MiB / 7852MiB

gentoo ~ # objdump -D /bin/bash | grep vsetvl | wc -l
1392
gentoo ~ # objdump -D /bin/bash | grep vsetvl | tail -n 30
    a5cc4:      0c57f6d7                vsetvli a3,a5,e8,mf8,ta,ma
    a5cc8:      0d807757                vsetvli a4,zero,e64,m1,ta,ma
    a6252:      0c007757                vsetvli a4,zero,e8,m1,ta,ma
    a66da:      0c06f7d7                vsetvli a5,a3,e8,m1,ta,ma
    a6a16:      0c06f7d7                vsetvli a5,a3,e8,m1,ta,ma
    a787a:      0c06f7d7                vsetvli a5,a3,e8,m1,ta,ma
    a78f4:      0c06f7d7                vsetvli a5,a3,e8,m1,ta,ma
    a79b0:      0c007657                vsetvli a2,zero,e8,m1,ta,ma
    a7b88:      0c0777d7                vsetvli a5,a4,e8,m1,ta,ma
    a7c02:      0c0777d7                vsetvli a5,a4,e8,m1,ta,ma
    a7c66:      0c007757                vsetvli a4,zero,e8,m1,ta,ma
    a7f2e:      0c007757                vsetvli a4,zero,e8,m1,ta,ma
    a7f58:      0c007757                vsetvli a4,zero,e8,m1,ta,ma
    a8162:      0d807057                vsetvli zero,zero,e64,m1,ta,ma
    a816a:      0d707057                vsetvli zero,zero,e32,mf2,ta,ma
    a81c6:      0d807057                vsetvli zero,zero,e64,m1,ta,ma
```

Developing on ROMA-II with RVV 1.0

- But is it faster…
- Greg Sterling of RISC-V International compared images with and without RVV
- Built using gcc-15.1 with -march=rv64gcv_zvl256b
  - Good speed-ups for C only code:
    - qwen:0.5b model 302% faster than scalar
    - deepseek-r1:1.5b model 368% faster than scalar
- ISCAS landing 128-bit RVV in llama.cpp (used by ollama demo)
  - Good improvements gcc-14 -> gcc-15
  - Need to test on other HW with gcc-15 autovectorizer enabled

Installing the Image [1] (from Aug-15)

```
$ dd if=gentoo-linux-k1_dev-sdcard-2.0rc3.img of=/dev/sdc conv=sync status=progress
```

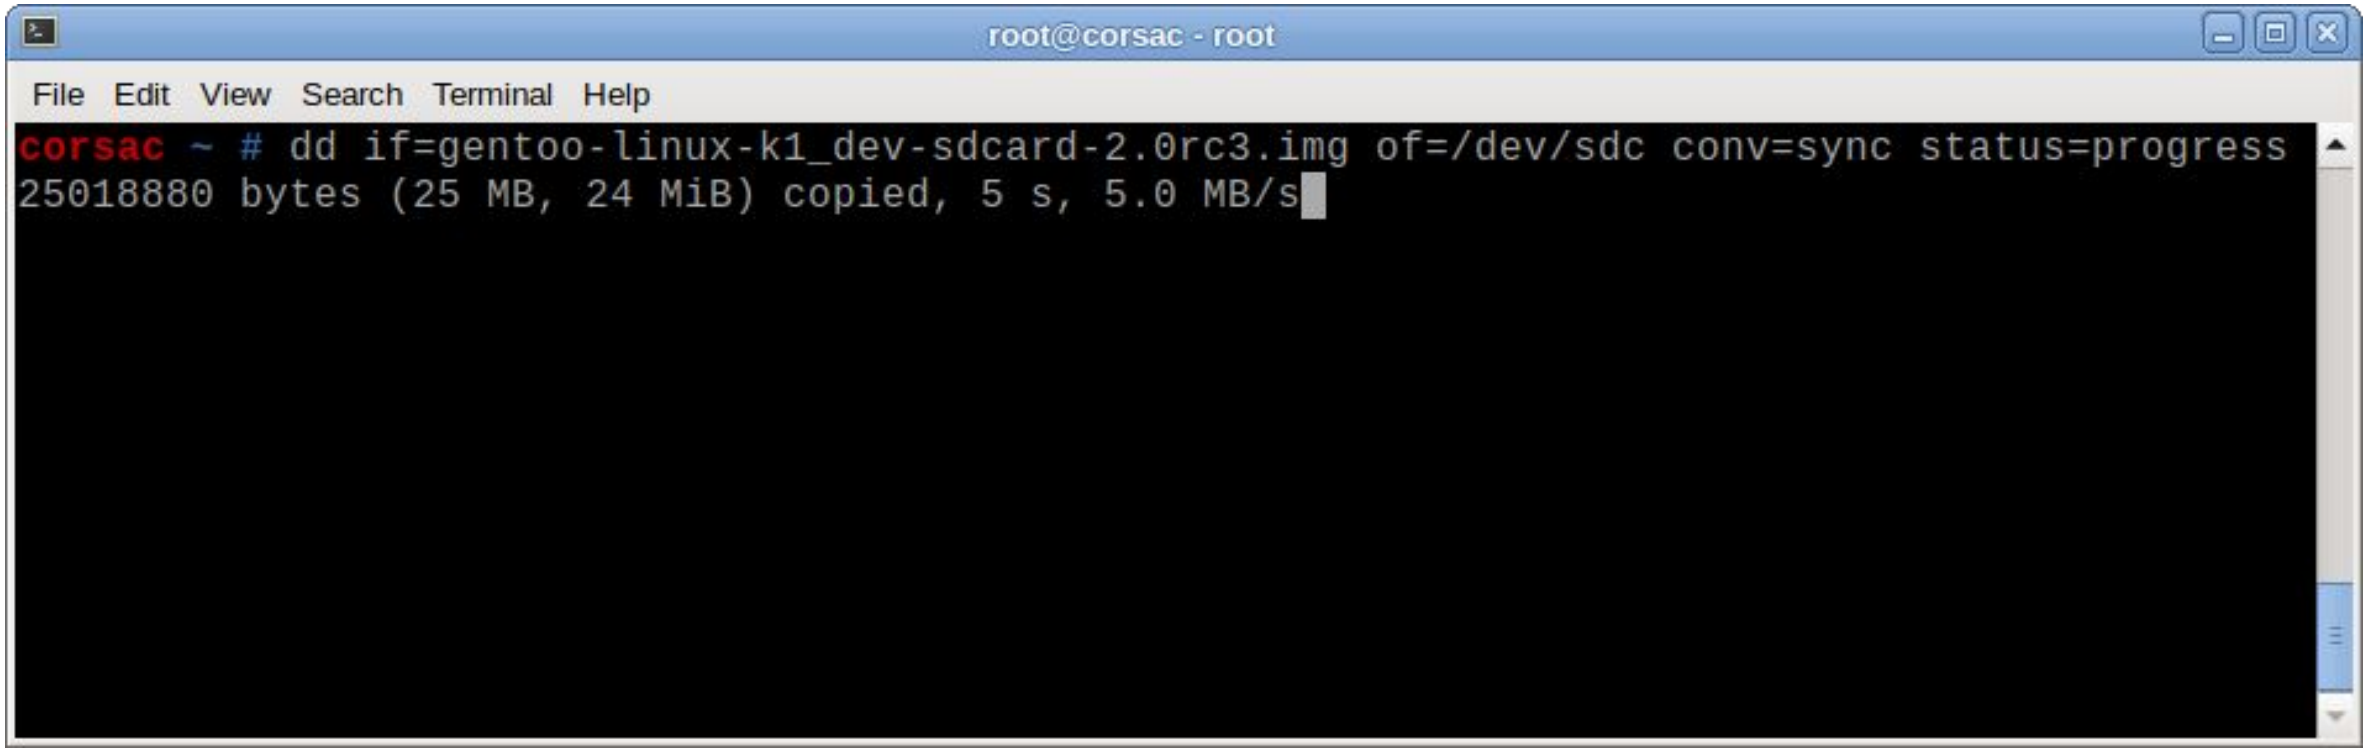

[1] https://dev.gentoo.org/~lu_zero/gentoo-linux-k1_dev-sdcard-2.0rc3.img.xz

Boot Process

```
U-Boot SPL 2022.10spacemit (Aug 14 2024 - 20:15:22 -0000)
DDR type LPDDR4X
lpddr4_silicon_init consume 11ms
Change DDR data rate to 2400MT/s
Boot from fit configuration k1-x_deb1
## Checking hash(es) for config conf_2 ... OK
## Checking hash(es) for Image uboot ... crc32+ OK
## Checking hash(es) for Image fdt_2 ... crc32+ OK
## Checking hash(es) for config config_1 ... OK
## Checking hash(es) for Image opensbi ... crc32+ OK


U-Boot 2022.10spacemit (Aug 14 2024 - 20:15:22 -0000)

CPU:   rv64imafdcv
Model: spacemit k1-x deb1 board
DRAM:  DDR size = 4096 MB
DDR size = 4096 MB
DDR size = 4096 MB
```

```
## Loading kernel from FIT Image at 11000000 ...
   Using 'conf-default' configuration
   Verifying Hash Integrity ... OK
   Trying 'kernel' kernel subimage
     Description:  Linux 6.6.36+
     Type:         Kernel Image
     Compression:  gzip compressed
     Data Start:   0x110000bc
     Data Size:    14255955 Bytes = 13.6 MiB
     Architecture: RISC-V
     OS:           Linux
     Load Address: 0x00200000
     Entry Point:  0x00200000
     Hash algo:    crc32
     Hash value:   7c3065e0
   Verifying Hash Integrity ... crc32+ OK
## Flattened Device Tree blob at 31000000
   Booting using the fdt blob at 0x31000000
   Uncompressing Kernel Image
```


Boot Process (Con’t)

```
   OpenRC 0.54.2 is starting up Gentoo Linux (riscv64)

 * Mounting /proc ...                          * Create Volatile Files and Directories ...
 [ ok ]                                        [ ok ]
 * Mounting /run ...                          INIT: Entering runlevel: 3
 [ ok ]                                        * Starting metalog ...
 * /run/openrc: creating directory             [ ok ]
 * /run/lock: creating directory               * Starting DHCP Client Daemon ...
 * /run/lock: correcting owner                dhcp_vendor: No such process
 * Caching service dependencies ...            [ ok ]
[    5.445256] usb 2-1.5: new high-speed       * Mounting network filesystems ...
 [ ok ]                                        [ ok ]
 * Mounting /sys ...                           * Starting sshd ...
 [ ok ]                                        [ ok ]
 * Mounting debug filesystem ...               * Starting local ...
 [ ok ]                                        [ ok ]
 * Mounting config filesystem ...
 [ ok ]
 * Mounting fuse control filesystem ...       This is localhost (Linux riscv64 6.6.36+) 21:56:52

                                              localhost login:
```


Full Gentoo Linux System

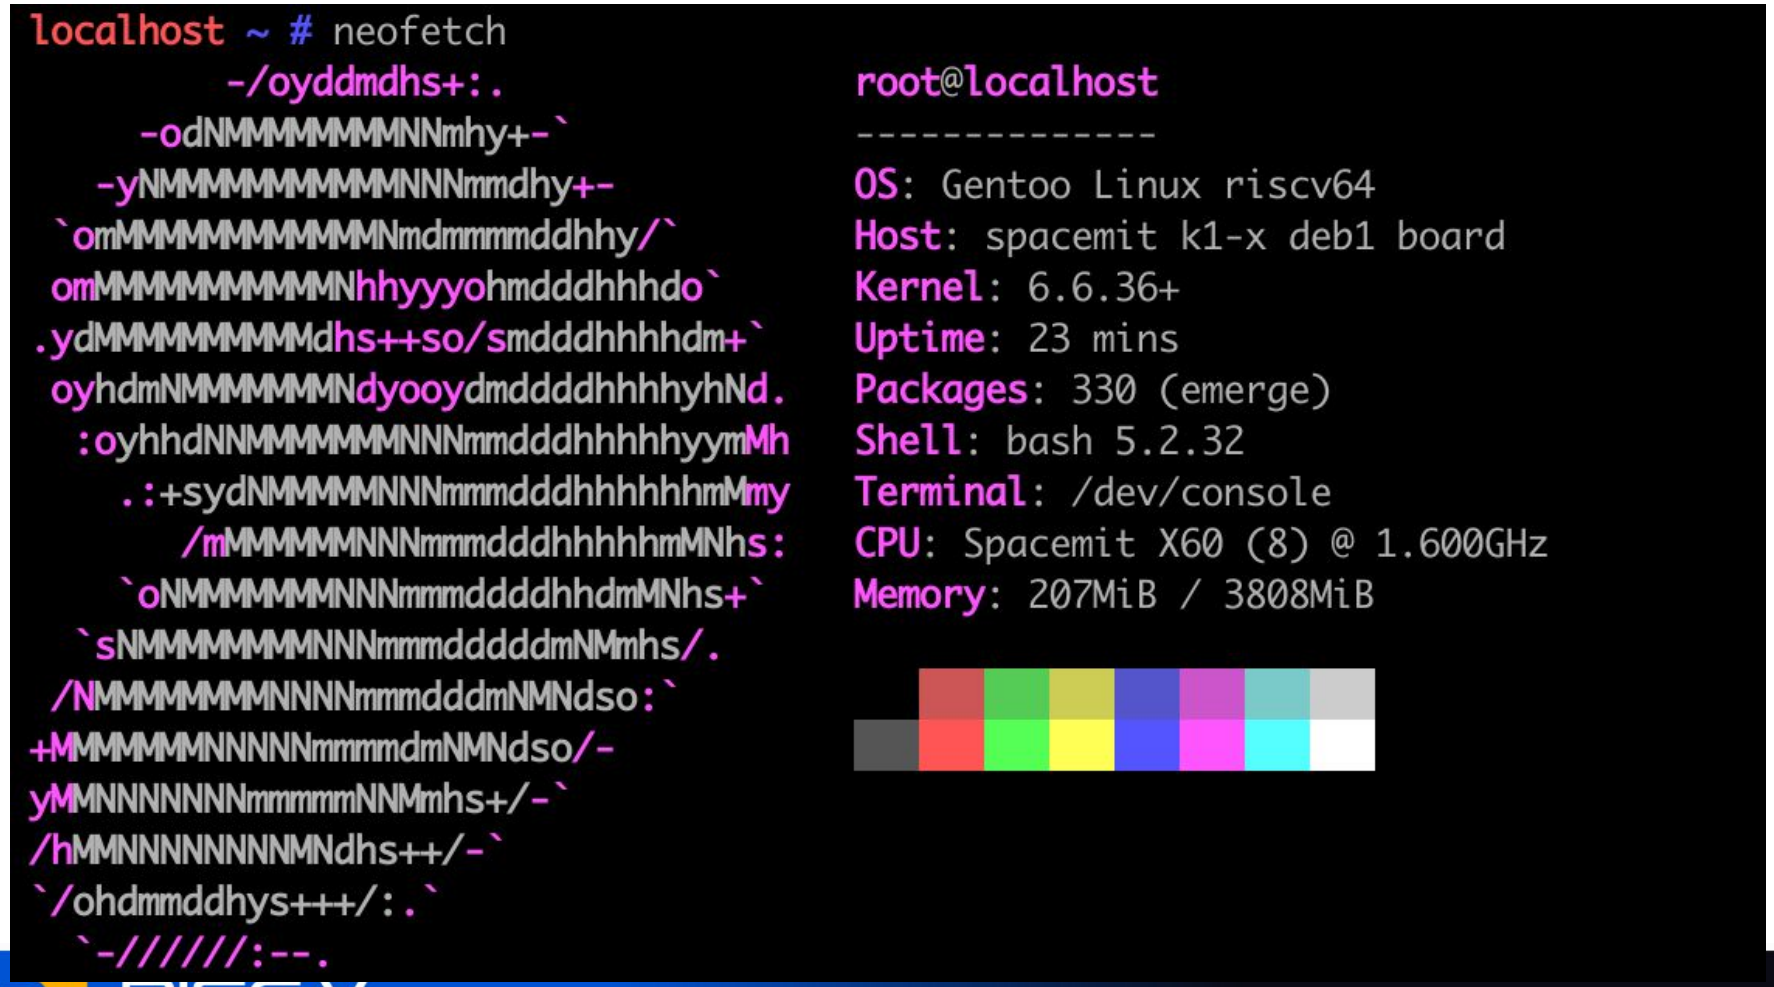

Up-to-date Toolchains!

```
localhost ~ # clang --version
clang version 18.1.8
Target: riscv64-unknown-linux-gnu
Thread model: posix
InstalledDir: /usr/lib/llvm/18/bin
Configuration file: /etc/clang/riscv64-unknown-linux-gnu-clang.cfg
localhost ~ # gcc --version
gcc (Gentoo 14.2.0 p4) 14.2.0
Copyright (C) 2024 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.  There is NO
warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

localhost ~ # /usr/lib/ld-linux-riscv64-lp64d.so.1 --version
ld.so (Gentoo 2.40 (patchset 1)) stable release version 2.40.
Copyright (C) 2024 Free Software Foundation, Inc.
This is free software; see the source for copying conditions.
There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A
PARTICULAR PURPOSE.
localhost ~ #
```

2024 Jun 20

2024 Aug 1

2024 Jul 22

Wifi and ethernet just work out of the box

```
localhost ~ # modprobe 8852bs
localhost ~ # /etc/init.d/wpa_supplicant start
 * Starting WPA Supplicant Daemon ...
Successfully initialized wpa_supplicant                          [ ok ]
localhost ~ # ifconfig wlan0
wlan0: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 1500
        inet 192.168.9.202  netmask 255.255.255.0  broadcast 192.168.9.255
        inet6 fd1a:637d:f215:0:d7e5:b531:88fc:b12e  prefixlen 64  scopeid 0x0<gl
obal>
        inet6 fd95:b4c7:7c8b:0:300c:64b3:c757:b96f  prefixlen 64  scopeid 0x0<gl
obal>
        inet6 fd8d:88cb:94f4:0:b5d0:53e4:4cf:2073  prefixlen 64  scopeid 0x0<glo
bal>
        inet6 fe80::b7a2:5296:db95:fe64  prefixlen 64  scopeid 0x20<link>
        inet6 fd1a:637d:f215::93c  prefixlen 128  scopeid 0x0<global>
        ether c0:4b:24:36:6b:af  txqueuelen 1000  (Ethernet)
        RX packets 12326  bytes 51588926 (49.1 MiB)
        RX errors 0  dropped 0  overruns 0  frame 0
        TX packets 12397  bytes 696222 (679.9 KiB)
```

Emerge (install) new packages

```
localhost ~ # ldconfig
localhost ~ # emerge-webrsync -q
 * Latest snapshot date: 20240814
 *
 * Approximate snapshot timestamp: 1723682700
 *        Current local timestamp: 1723682400
 *
 * The current local timestamp is possibly identical to the
 * timestamp of the latest snapshot. In order to force sync, use
 * the --revert option or remove the timestamp file located at
 * '/var/db/repos/gentoo/metadata/timestamp.x'.
localhost ~ # getuto && emerge -g neofetch

 * IMPORTANT: 18 news items need reading for repository 'gentoo'.
 * Use eselect news read to view new items.


Local copy of remote index is up-to-date and will be used.
Calculating dependencies... done!
```

Future Work

- Experiment with alternate whole system build configs
  - Crossdev already supports `riscv64-unknown-linux-musl` as target
  - Paves the way to build the whole system using clang
- Improve the overall cross-building experience
  - This project already found many bugs
    - Few packages (e.g. perl) already got some fixes
    - Crossdev has a pending patch to make it profile-aware
- Automate image generation on monthly or weekly basis
  - Latest images:

[1] Canaan K230 https://people.videolan.org/~negge/canaan-3G-2024-04-08.img.xz
[2] BPI F3 / ROMA II https://dev.gentoo.org/~lu_zero/riscv/gentoo-linux-k1_dev-sdcard-2.1.0-20250411-rvv.img.xz
[3] Orange Pi RV2 https://people.videolan.org/~negge/gentoo-linux-x1_dev-sdcard-20250422-rvv.img.xz

Fedora Images Available

- Community built as rv64gc [1]
  - Images for many SOCs and devices
- Official Fedora 41 and 42 builds [2]
  - Supporting:
    - StarFive VisionFive 2
    - SiFive HiFive Unmatched
    - SiFive HiFive Premier P550

[1] https://images.fedoravforce.org
[2] https://fedoraproject.org/wiki/Architectures/RISC-V

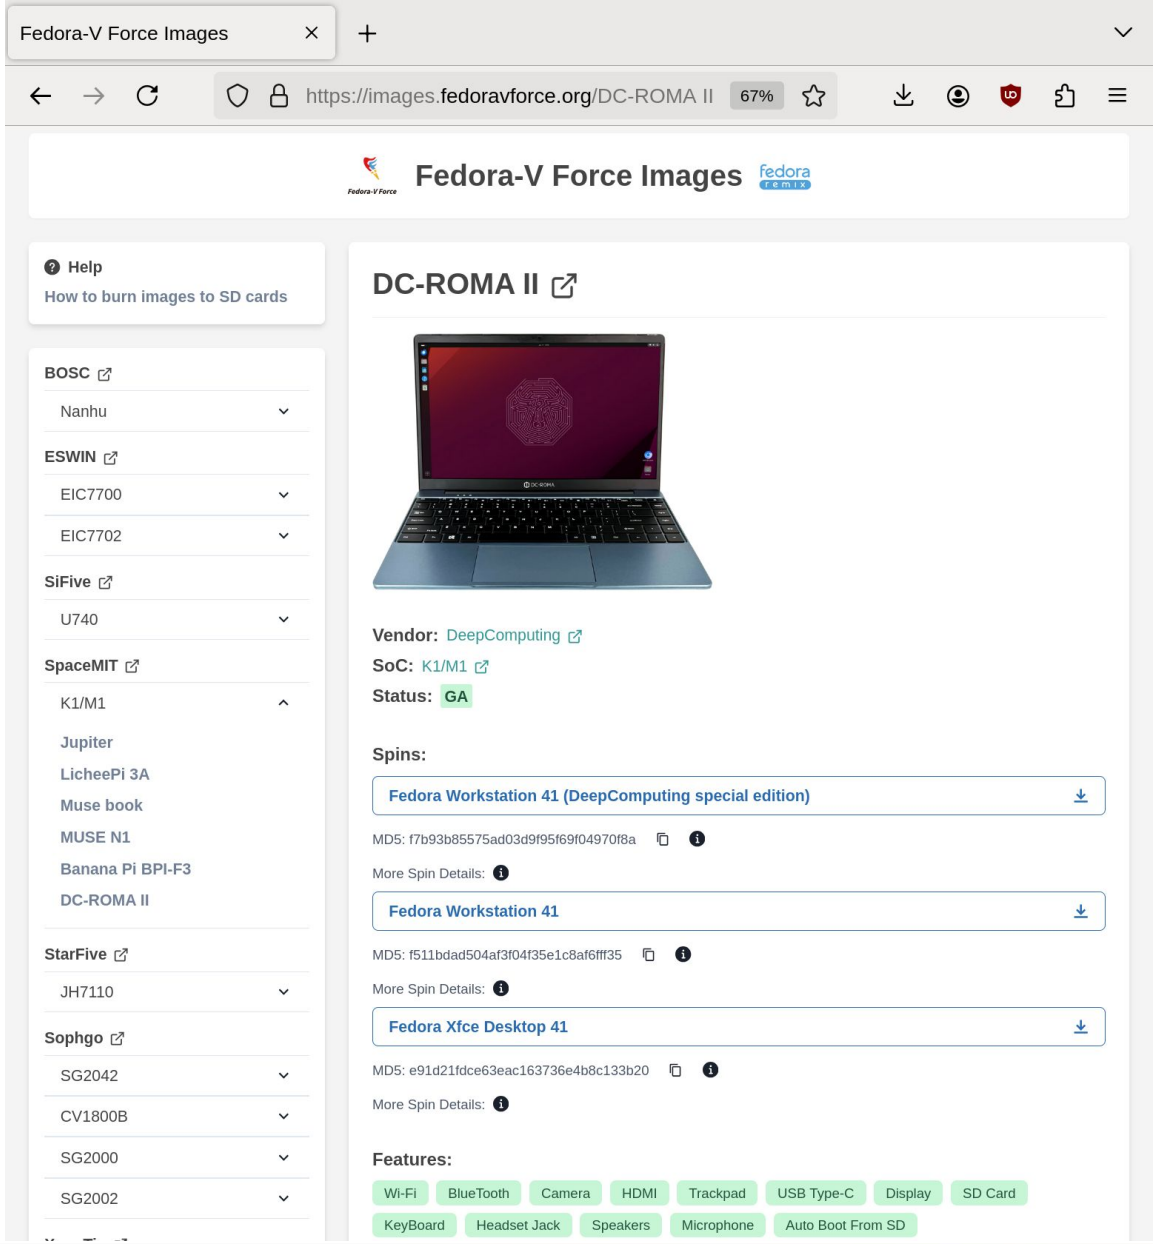

Demo in RISE Lounge

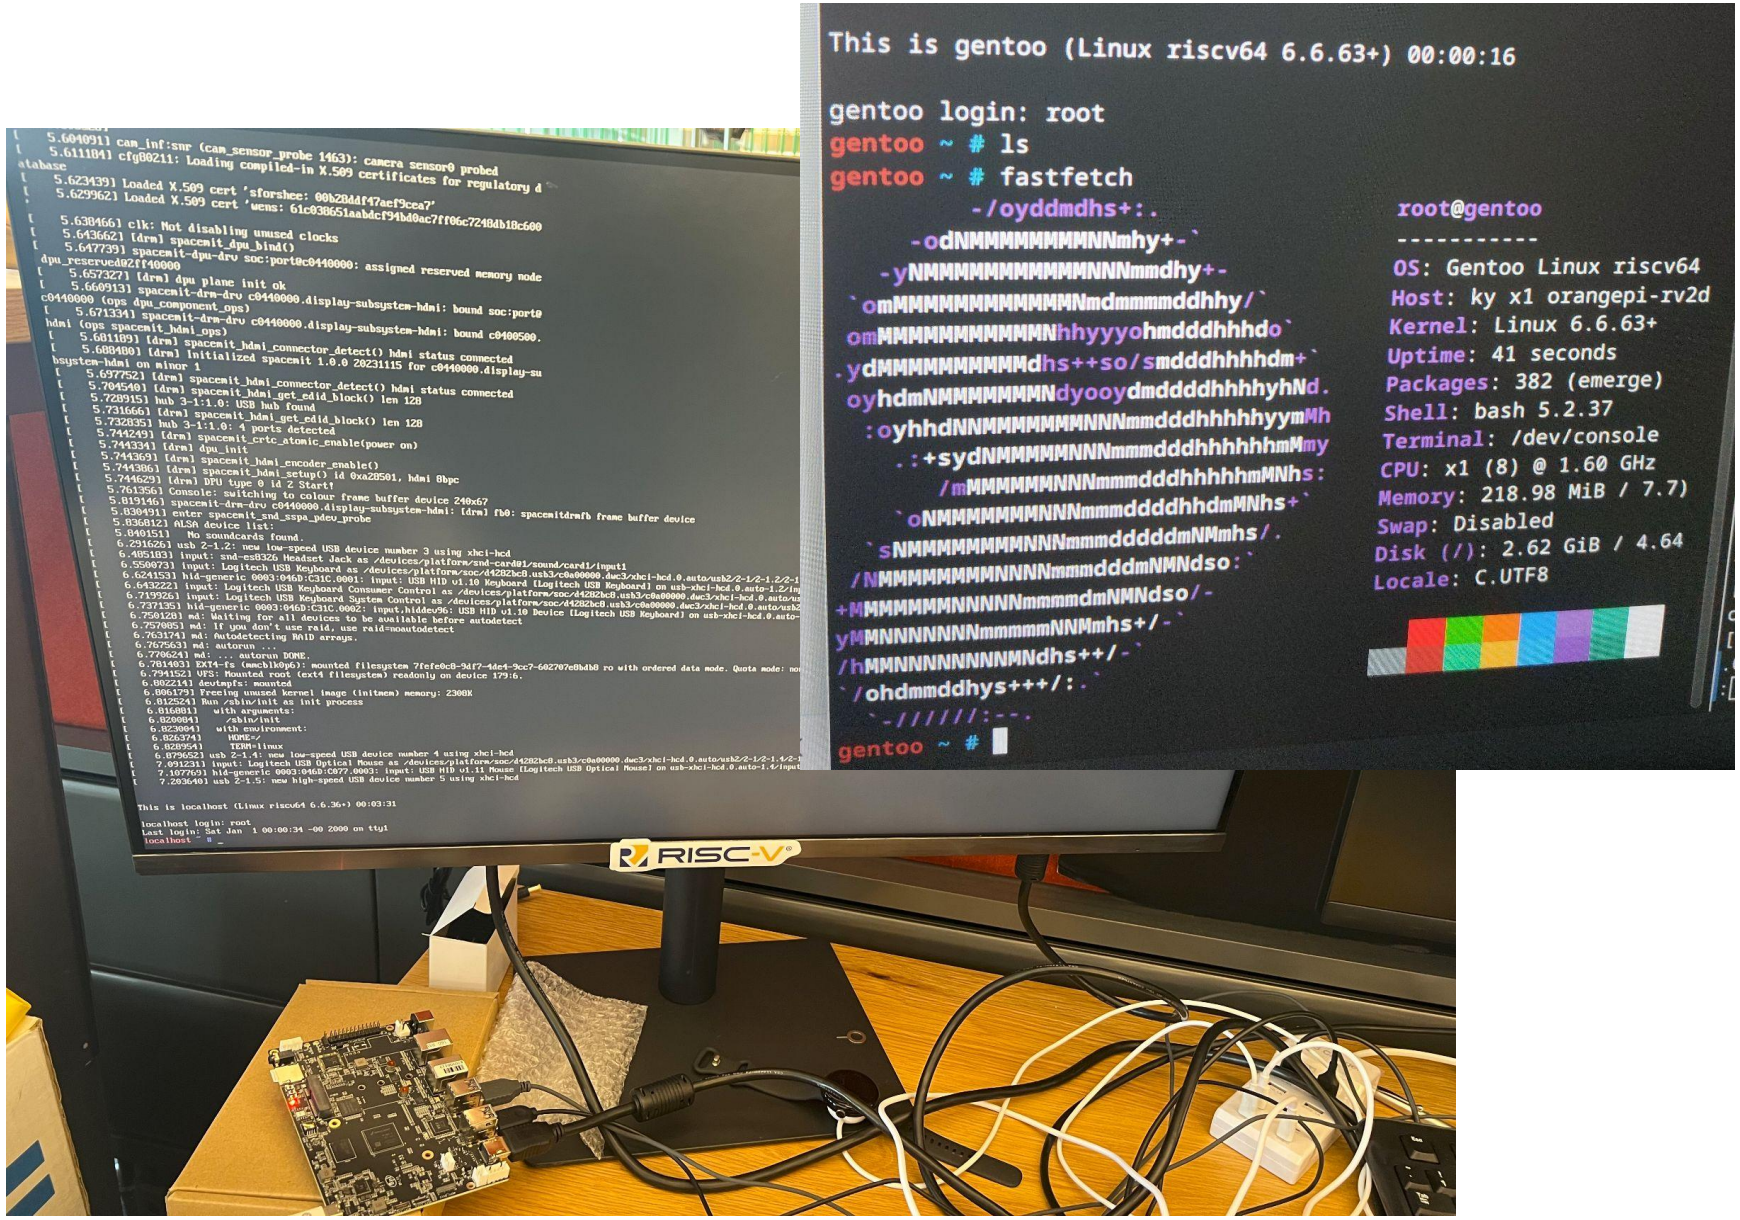

Example RVV 1.0 Code #1 - Get vector length

```
.global _start

_start:
  csrr a0, vlenb
  addi a7, x0, 93
  ecall
```

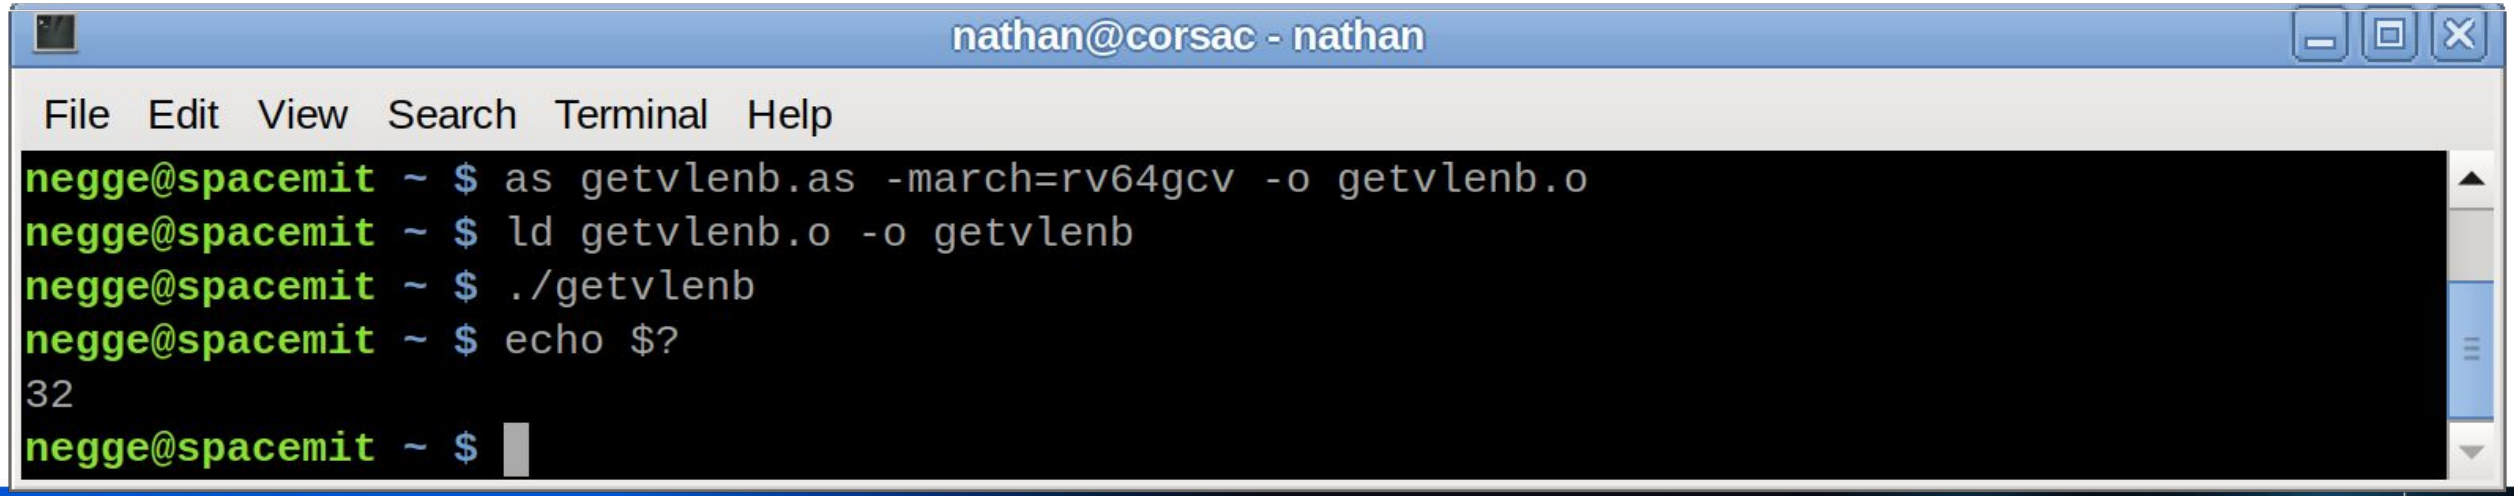

Example RVV 1.0 Code #2 - Run-time detect

```c
#include <sys/auxv.h>
#include <stdio.h>

#define ISA_V_HWCAP (1 << ('v' - 'a'))

void main() {
  unsigned long hw_cap = getauxval(AT_HWCAP);
  printf("RVV %s\n", hw_cap & ISA_V_HWCAP ? "detected" : "not found");
}
```

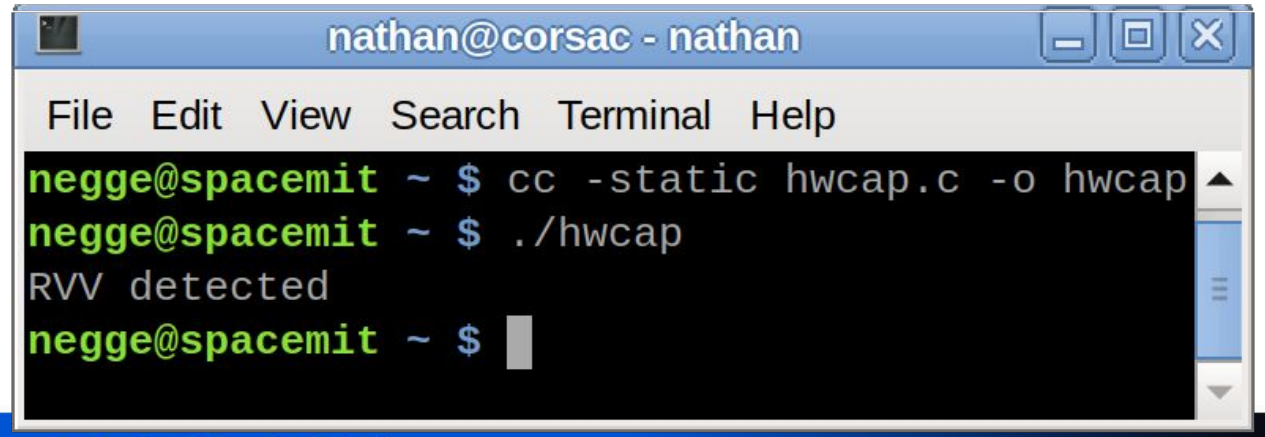

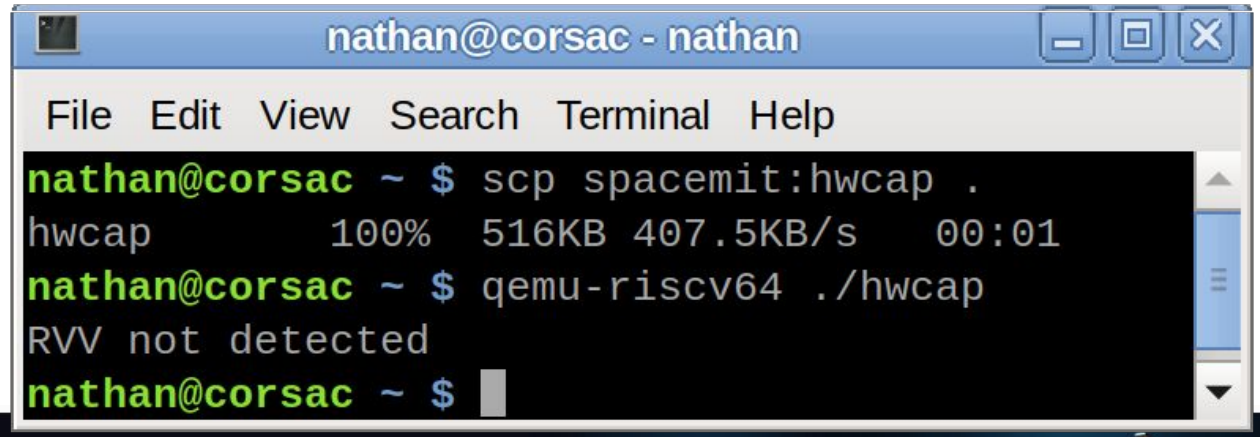

Example RVV 1.0 #3 - Application Profiling

```
$ perf record -e u_mode_cycle ./dav1d -i Bosphorus_1080p_8bit.ivf -o /dev/null
dav1d 1.5.0-3-g55fb943 - by VideoLAN
Decoded 600/600 frames (100.0%) - 10.04/30.00 fps (0.33x)
[ perf record: Woken up 84 times to write data ]
[ perf record: Captured and wrote 22.223 MB perf.data (582464 samples) ]
$ perf report
```

```
# To display the perf.data header info, please use --header/--header-only options.
#
#
# Total Lost Samples: 0
#
# Samples: 578K of event 'u_mode_cycle:u'
# Event count (approx.): 218918769644
#
# Overhead  Command       Shared Object                Symbol
# ........  ............  ...........................  ..................................................................
#
    67.59%  dav1d-worker  libdav1d.so.7.0.0            [.] prep_8tap_c
     3.68%  dav1d-worker  libdav1d.so.7.0.0            [.] put_8tap_c
     3.58%  dav1d-worker  libdav1d.so.7.0.0            [.] $xrv64i2p1_m2p0_a2p1_f2p2_d2p2_c2p0_zicsr2p0_zifencei2p0_zmmul1p0
     2.45%  dav1d-worker  libdav1d.so.7.0.0            [.] wiener_c
     2.22%  dav1d-worker  libdav1d.so.7.0.0            [.] $xrv64i2p1_m2p0_a2p1_f2p2_d2p2_c2p0_zicsr2p0_zifencei2p0_zmmul1p0
     2.11%  dav1d-worker  libdav1d.so.7.0.0            [.] prep_8tap_smooth_sharp_c
     1.44%  dav1d-worker  libdav1d.so.7.0.0            [.] prep_8tap_smooth_regular_c
     1.29%  dav1d-worker  libdav1d.so.7.0.0            [.] dav1d_mask_8bpc_rvv
     0.87%  dav1d-worker  libdav1d.so.7.0.0            [.] $xrv64i2p1_m2p0_a2p1_f2p2_d2p2_c2p0_zicsr2p0_zifencei2p0_zmmul1p0
     0.80%  dav1d-worker  libdav1d.so.7.0.0            [.] load_tmvs_c
     0.70%  dav1d-worker  libdav1d.so.7.0.0            [.] prep_8tap_sharp_c
     0.69%  dav1d-worker  libdav1d.so.7.0.0            [.] prep_8tap_smooth_c
     0.68%  dav1d-worker  libdav1d.so.7.0.0            [.] decode_b
     0.62%  dav1d-worker  libdav1d.so.7.0.0            [.] dav1d_create_lf_mask_inter
     0.58%  dav1d-worker  libdav1d.so.7.0.0            [.] put_8tap_scaled_c
     0.54%  dav1d-worker  libc.so.6                    [.] __strxfrm_l
```

RISC-V Hacking!

- Kaleidoscopico by lft [1]
  - Bare-metal demoscene art
  - 2nd place wild compo at Revision [2]
  - RP2350 Pico 2 @ 130MHz
  - 17,000 lines RISC-V asm
  - Full technical write-up

Linus Åkesson's Kaleidoscopico Turns a Raspberry Pi Pico 2 Into a Powerful Demoscene Machine

Commodore Amiga-inspired project, totaling 17,000 lines of RISC-V assembler, shows just what the RP2350 can do.

Gareth Halfacree

14 days ago • Art / Music / HW101

Follow

[1] https://www.pouet.net/prod.php?which=104053
[2] https://2025.revision-party.net/
[3] https://youtu.be/KhhLoVBpg48

Example RVV 1.0 #4 - dav1d 8bpc mc blend

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

Pseudo-code for RVV, note w can be 4, 8, 16 or 32 only

Set VL based on w
Load vectors for dst, tmp and mask
Scratch vector for widening multiply, followed by widening multiply accumulate
Narrowing shift with rounding
Store back into dst

Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
```

```
  ret
endfunc
```

Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma
```

```
  ret
endfunc
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma
  li t1, 64                 // t1 = 64;
```

```
  ret
endfunc
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma
  li t1, 64                 // t1 = 64;
1:                          // do {
  addi a4, a4, -1           //   h = h - 1;
  ...



  bnez a4, 1b               // } while (h != 0)
  ret
endfunc
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma
  li t1, 64                 // t1 = 64;
1:                          // do {
  addi a4, a4, -1           //   h = h - 1;
  ...
  add a0, a0, a1            // dst += dst_stride
  add a2, a2, a3            // tmp += w;
  add a5, a5, a3            // mask += w;
  bnez a4, 1b               // } while (h != 0)
  ret
endfunc
```

Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
  vrsub.vx v8, v8, t1      // v8 = 64 - v8;
```

Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
  vrsub.vx v8, v8, t1      // v8 = 64 - v8;
  vwmaccu.vv v16, v0, v8   // v16 = v16 + v0*v8;
```

Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
  vrsub.vx v8, v8, t1      // v8 = 64 - v8;
  vwmaccu.vv v16, v0, v8   // v16 = v16 + v0*v8;
  vnsra.wi v0, v16, 6      // v0 = (v16 + 32) >> 6;
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
  vrsub.vx v8, v8, t1      // v8 = 64 - v8;
  vwmaccu.vv v16, v0, v8   // v16 = v16 + v0*v8;
  vnsra.wi v0, v16, 6      // v0 = (v16 + 32) >> 6; <-- does not work due to signed-signed operation
  vnclipu.wi v0, v16, 6    // v0 = MAX(0, MIN(256, (v16 + 32) >> 6));
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (inner loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h --> 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
  vle8.v v0, (a0)          // v0 = *dst;
  vle8.v v4, (a2)          // v4 = *tmp;
  vle8.v v8, (a5)          // v8 = *mask;
  vwmulu.vv v16, v4, v8    // v16 = v4*v8;
  vrsub.vx v8, v8, t1      // v8 = 64 - v8;
  vwmaccu.vv v16, v0, v8   // v16 = v16 + v0*v8;
  vnsra.wi v0, v16, 6      // v0 = (v16 + 32) >> 6; <-- does not work due to signed-signed operation
  vnclipu.wi v0, v16, 6    // v0 = MAX(0, MIN(256, (v16 + 32) >> 6));
  vse8.v v0, (a0)          // *dst = v0;
```

Example RVV 1.0 #4 - dav1d 8bpc mc blend (all together)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}
```

```
function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma
  csrw vxrm, zero
  li t1, 64                 // t1 = 64;
1:                          // do {
  addi a4, a4, -1           //   h = h - 1;

  vle8.v v0, (a0)           // v0 = *dst;
  vle8.v v4, (a2)           // v4 = *tmp;
  vle8.v v8, (a5)           // v8 = *mask;

  vwmulu.vv v16, v4, v8     // v16 = v4*v8;
  vrsub.vx v8, v8, t1       // v8 = 64 - v8;
  vwmaccu.vv v16, v0, v8    // v16 = v16 + v0*v8;
  vnclipu.wi v0, v16, 6     // v0 = MAX(0, MIN(256, (v16 + 32) >> 6));
  vse8.v v0, (a0)           // *dst = v0;

  add a0, a0, a1            // dst += dst_stride
  add a2, a2, a3            // tmp += w;
  add a5, a5, a3            // mask += w;

  bnez a4, 1b               // } while (h != 0)
  ret
endfunc
```


Example RVV 1.0 #4 - dav1d 8bpc mc blend (checkasm)

- Run checkasm to verify correctness
  - Passes at width of 4, 8 and 16 but fails when w = 32
  - What is going on?

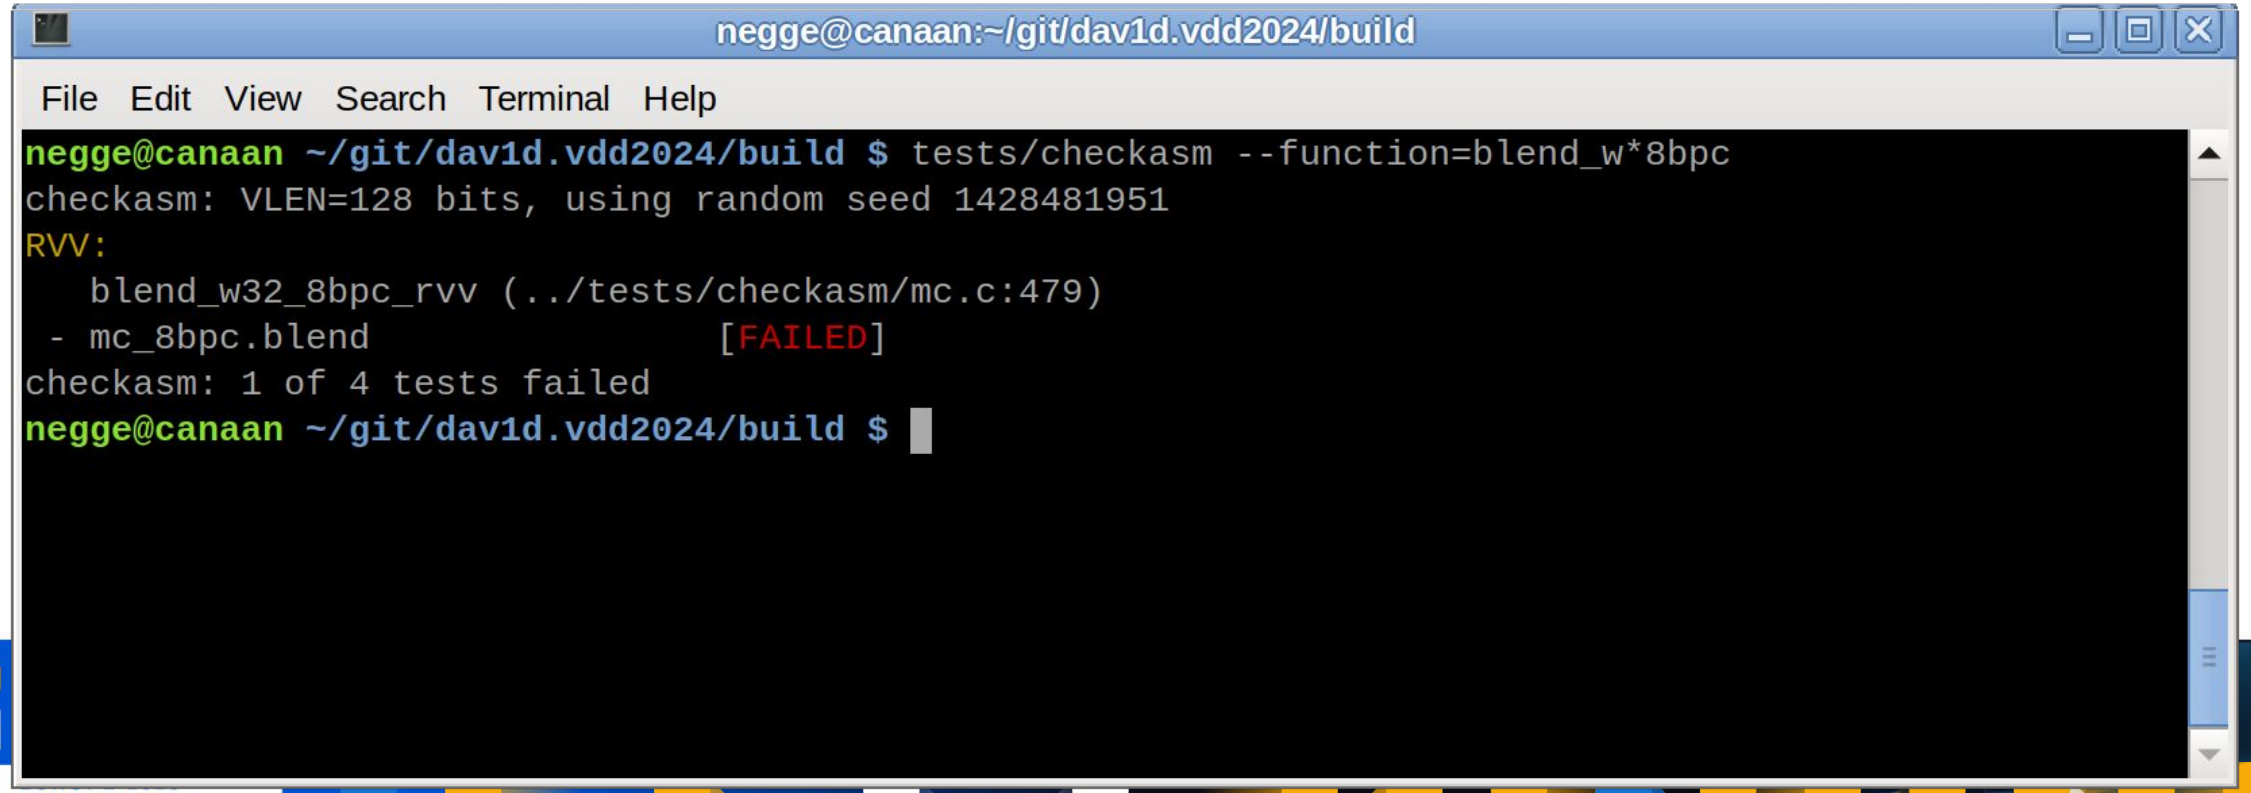

Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}

function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m1, ta, ma

  ret
endfunc
```

Canaan K230 has VLEN = 128
LMUL = m1 is not large enough for
8bpc * 32 = 256 bits

Example RVV 1.0 #4 - dav1d 8bpc mc blend (outer loop)

```
#define blend_px(a, b, m) (((a * (64 - m) + b * m) + 32) >> 6)
static void blend_c(uint8_t *dst, const ptrdiff_t dst_stride, const uint8_t *tmp,
                    const int w, int h, const uint8_t *mask) {
    for (; h-- > 0; dst += dst_stride, tmp += w, mask += w)
        for (int x = 0; x < w; x++) dst[x] = blend_px(dst[x], tmp[x], mask[x]);
}

function blend_8bpc_rvv, export=1, ext="v"
  vsetvli zero, a3, e8, m2, ta, ma

  ret
endfunc
```

Canaan K230 has VLEN = 128
LMUL = m1 is not large enough for
8bpc * 32 = 256 bits

Example RVV 1.0 #4 - dav1d 8bpc mc blend (checkasm)

- Run checkasm to verify correctness
  - Passes at all widths, of 4, 8, 16 and 32

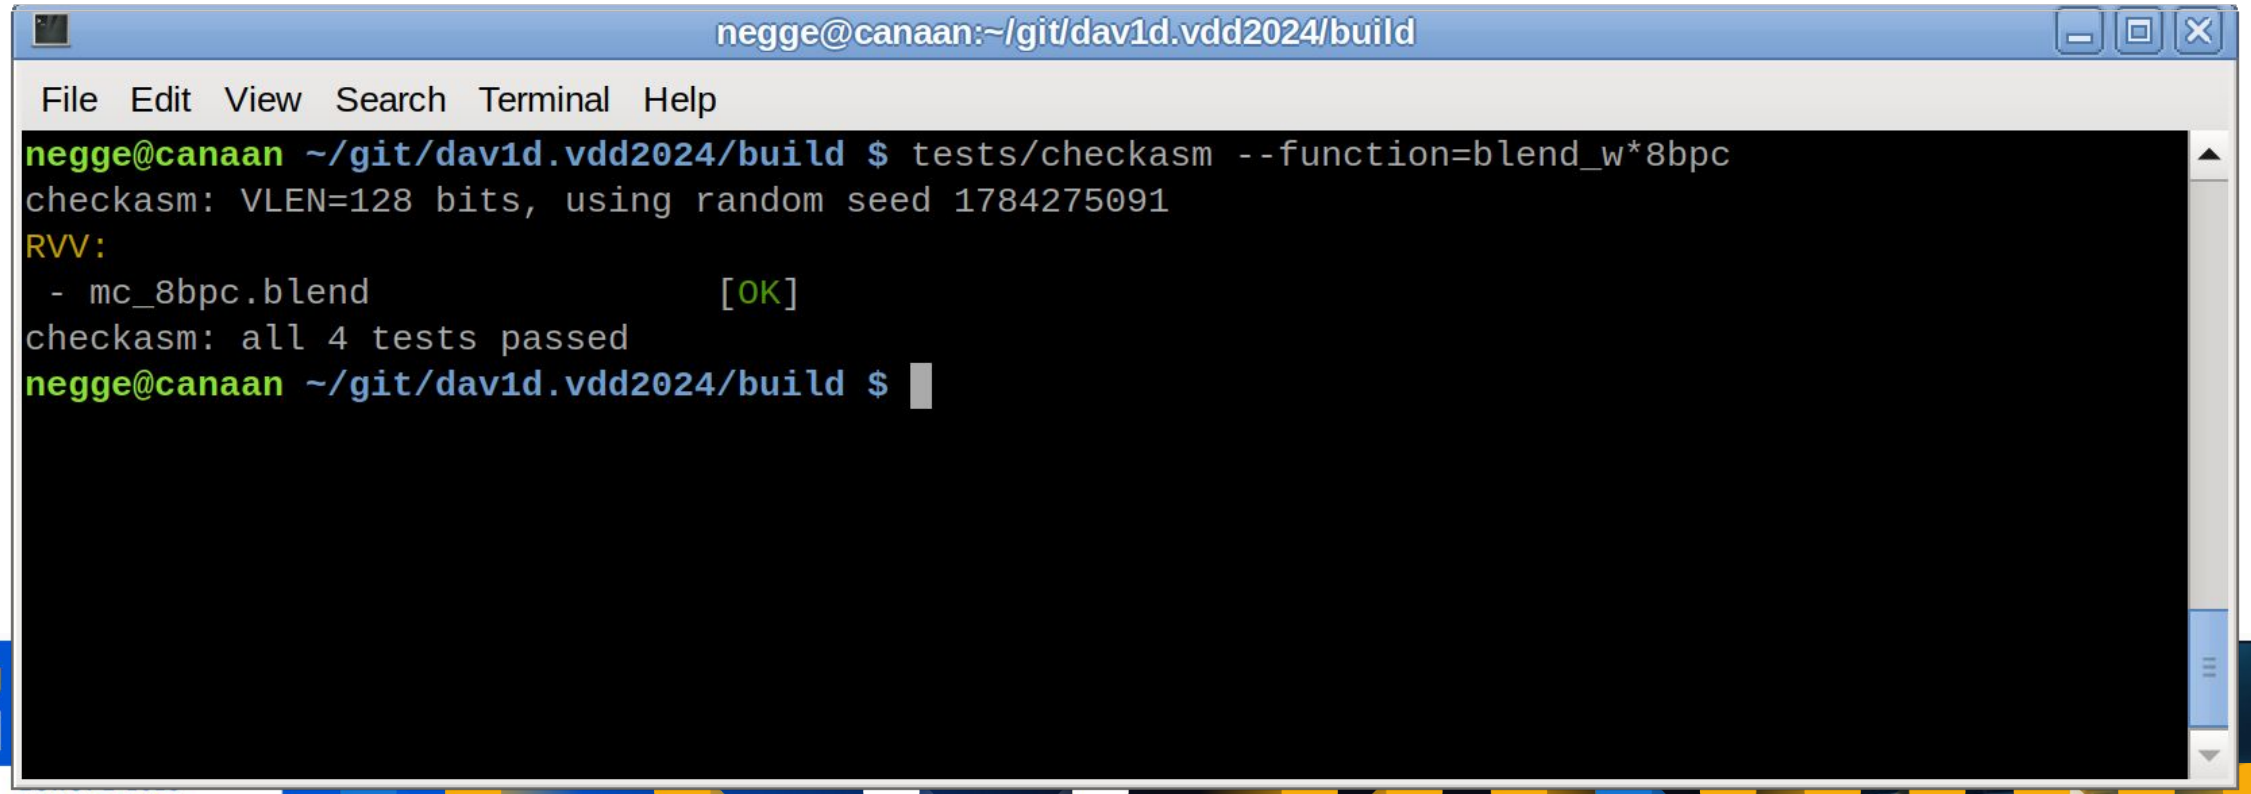

RISC-V SUMMIT
EUROPE 2025

Questions?