



The Heterogeneous Software Revolution Is Ready for You

AJ Guillon, YetiWare Inc.


Khronos Mission

PROMOTER MEMBERS

KHRONOS GROUP

Over 100 members worldwide
Any company is welcome to join

Members include: 3D Incorporated, AMOTIVE, Adobe, ATERA, amazon.com, AXELL CORPORATION, AXIS COMMUNICATIONS, BASE MARK, BINOMIAL, BLIZZARD, BROADCOM, BRENWILL, cadence, CANONICAL, CATAPULT Digital, CEVA, codeplay, Continental, Coordinate, COREAVI, DASSAULT SYSTEMES, ETRI, HARMAN, HITACHI Inspire the Next, htc, HYPREAL, igalia, Imperial College London, ITRI, KADAB, KNU, LG, Linaro, Los Alamos NATIONAL UNIVERSITY, LUNAR, matrox, MAXON, MEDIATEK, Mentor Graphics, Microsoft, MIT Lincoln Laboratory, mobica, Movidius, mozilla, NEC, Nintendo, NXP, oculus, ON, OSU, Panasonic, PIXAR, PLUTO, Qt The Qt Company, RAZER, RENESAS, Rockwell Collins, sensics, Silicon Studio, SIRU, socionext, SPREADTRUM, STREAM COMPUTING Performance Engineers, SYNOPSIS, TAKUMI, TAMPERE UNIVERSITY OF TECHNOLOGY, TU WIEN, TEXAS INSTRUMENTS, thinci, Think Silicon, tobii, TOSHIBA, umbra, unity, University of BRISTOL, TUCL, University of Windsor, Visteon, vmware, XILINX, zSpace.



Khronos is an International Industry Consortium of over 100 companies creating royalty-free, **open standard APIs** to enable software to access hardware acceleration for **3D graphics, Virtual and Augmented Reality, Parallel Computing, Neural Networks and Vision Processing**

OpenCL SC TSG




SPIR-V Ecosystem



Khronos has open sourced these tools and translators

<https://github.com/KhronosGroup/SPIRV-Tools>



SPiR-V (Dis)Assembler

SPiR-V Validator

Other Intermediate Forms

OpenCL C Front-end

OpenCL C++ Front-end

LLVM

LLVM to SPiR-V Bi-directional Translator

```

SPIR-V Magic #: 0x07230203
SPIR-V Version 99
Builder's Magic #: 0x051a00BB
<id> bound is 50
0
OpMemoryModel
Logical
GLSL450
OpEntryPoint
Fragment shader
function <id> 4
OpTypeVoid
<id> is 2
OpTypeFunction
<id> is 3
return type <id> is 2
OpFunction
Result Type <id> is 2
Result <id> is 4
0
Function Type <id> is 3
    
```

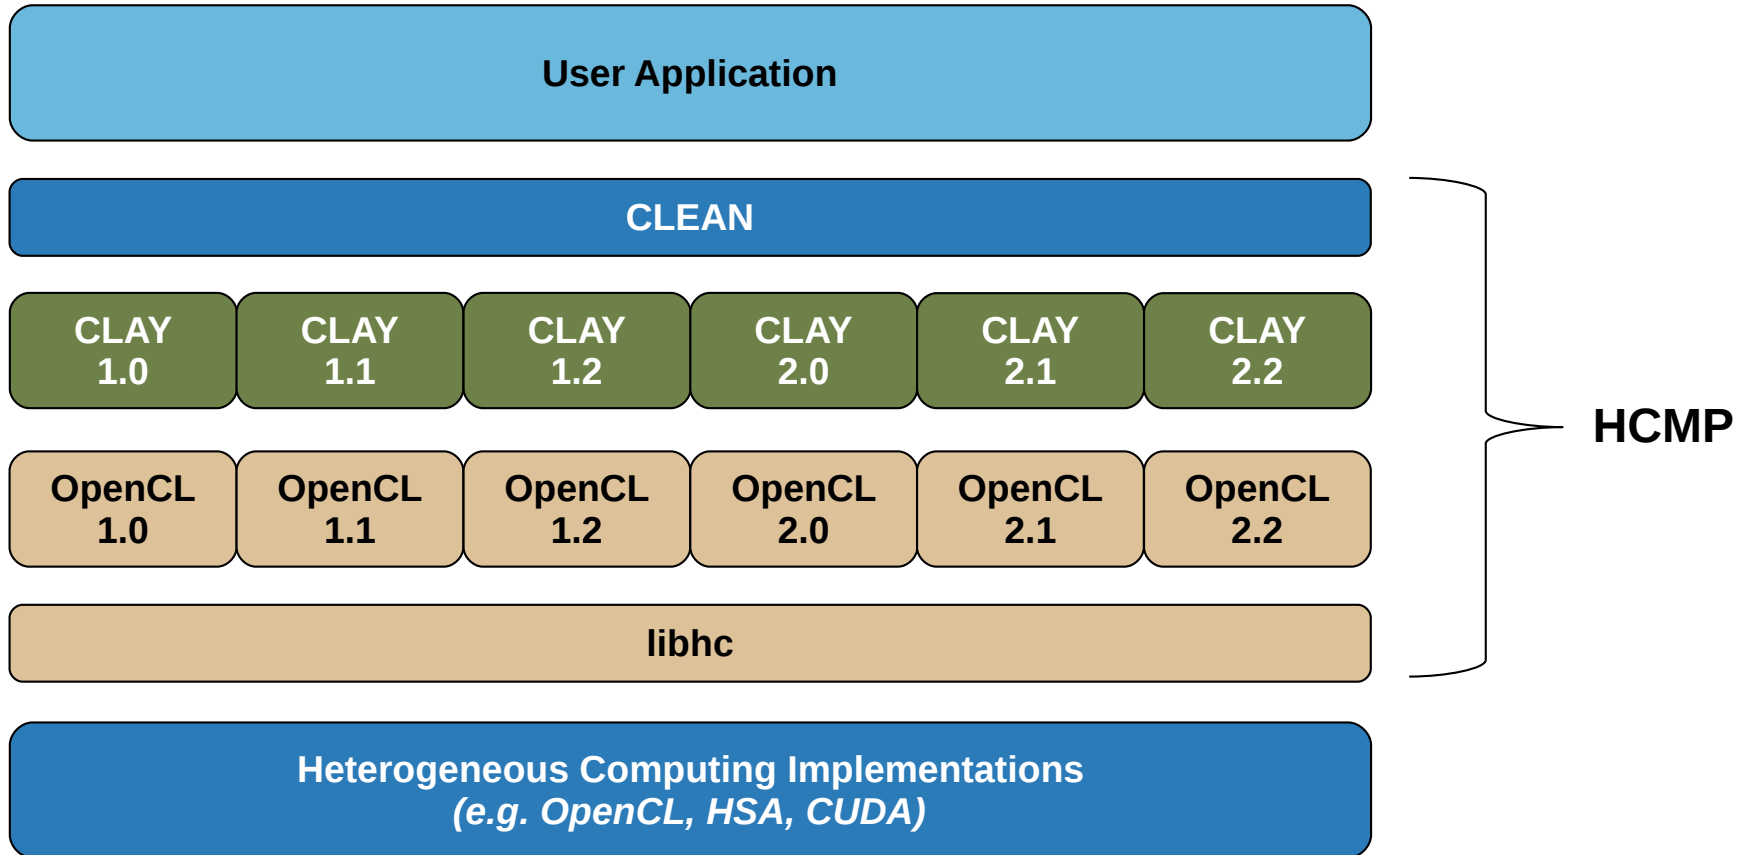
SPiR-V

- Khronos defined and controlled cross-API intermediate language
 - Native support for graphics and parallel constructs
 - 32-bit Word Stream
 - Extensible and easily parsed
- Retains data object and control flow information for effective code generation and translation

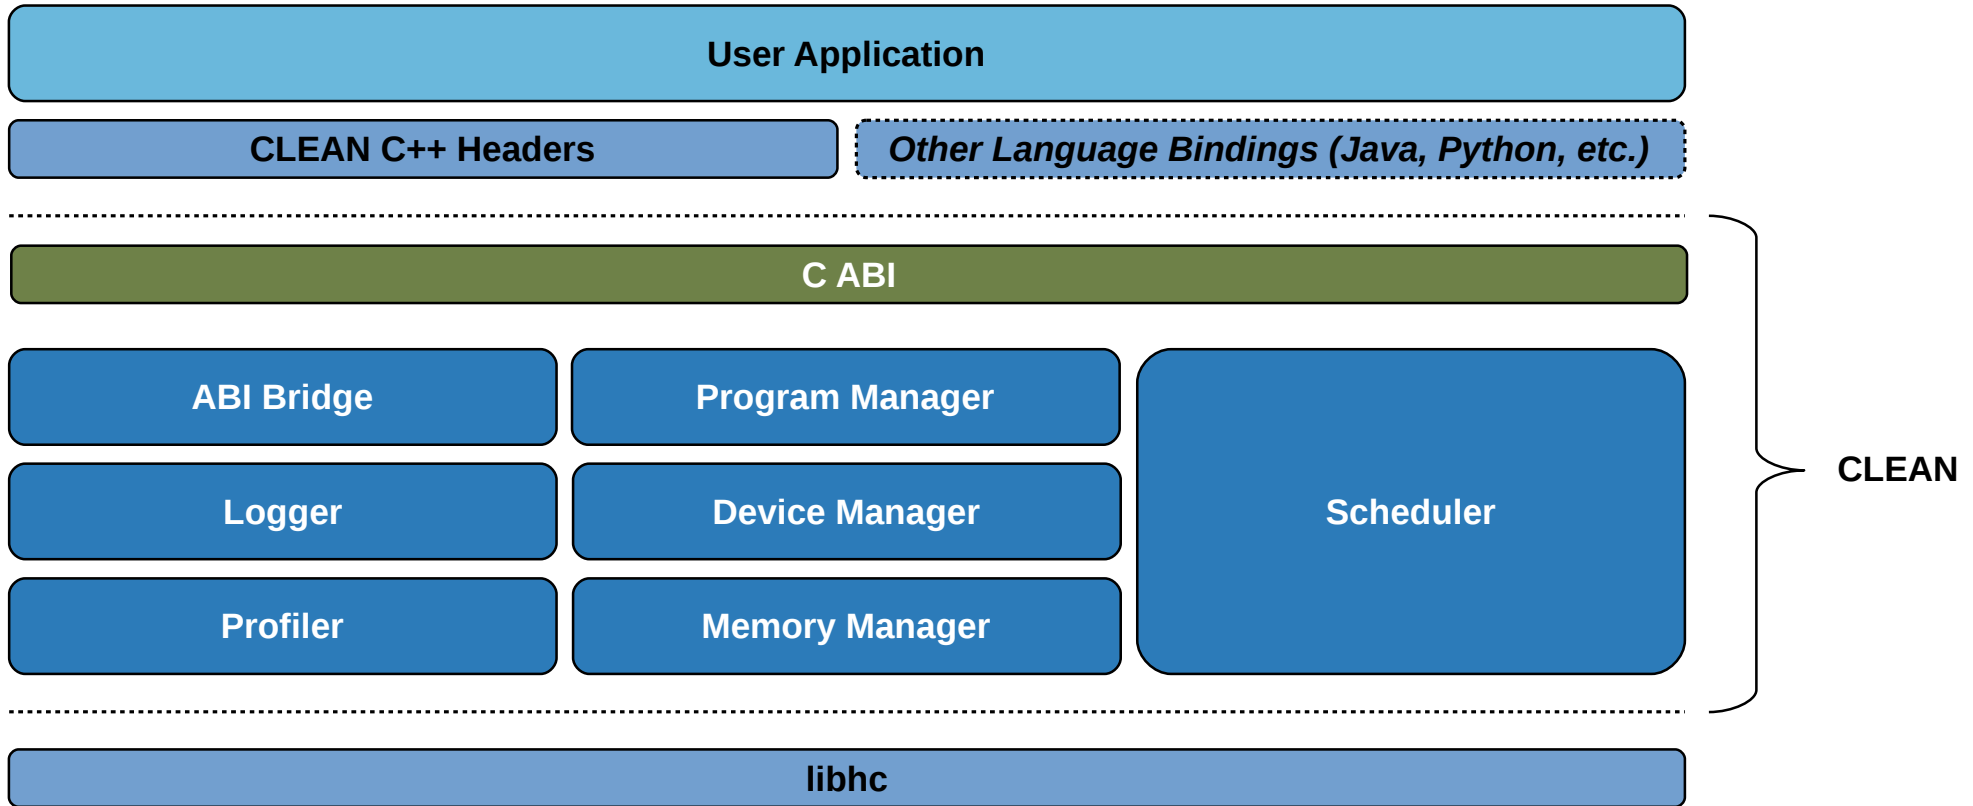
IHV Driver Runtimes



HCMP Architecture

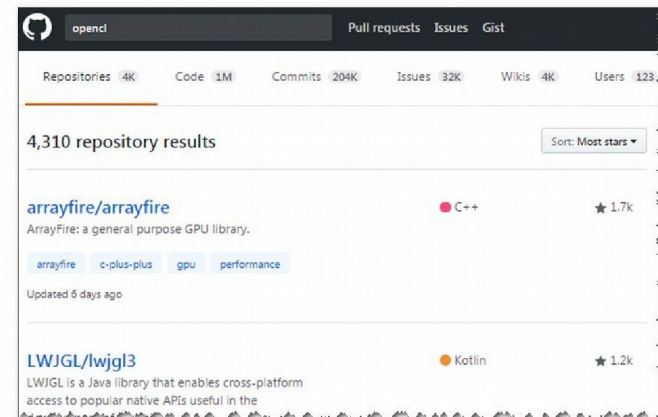


CLEAN Architecture



Lessons from OpenCL's First Eight Years

- **OpenCL is really needed**
 - Almost every silicon vendor shipping OpenCL implementations
 - 100s of software products relying OpenCL acceleration
 - Over 4,000 GitHub repositories building on OpenCL
- **A layered ecosystem is the way to go**
 - Low-level should provide explicit hardware access
 - Complex drivers get in the way - keep the lowest level as simple as possible
 - 'Ease-of-use' can be provided by higher layers
- **Performance Portability can't be done by magic**
 - Responsibility of higher levels with system/domain context
- **Users love language flexibility**
 - SPIR-V is awesome - choice of language front-ends
- **Market segments need deployment flexibility**
 - DSP inferencing should not be forced to ship IEEE FP32
 - Feature sets - toggle capabilities within a coherent framework
- **Industry adoption paths matter**
 - Apple, Google, NVIDIA...

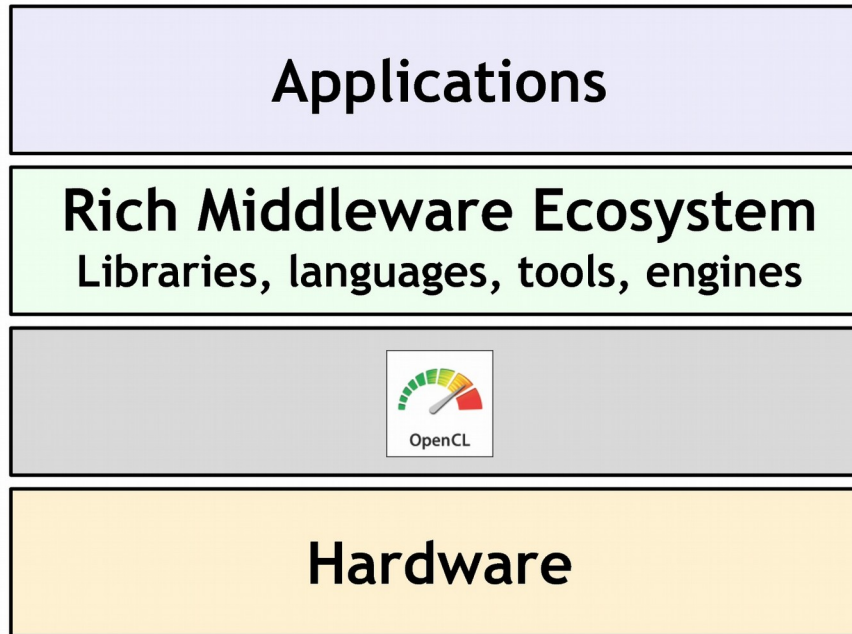


Embrace the Layered Ecosystem

OpenCL mixes providing low-level hardware access with 'ease-of-use'

Didn't make it clear that low-level performance portability is impossible

Did not focus on rapidly porting efficient libraries



Middleware just needs direct access to hardware. Driver should 'get out of the way'

Middleware can provide ease of use


Middleware has the system/domain context to try to provide performance portability

Run-time abstraction to compute hardware IS needed:

- Software vendors can't afford to port to every type/generation hardware
- Hardware vendors want to keep innovating under an abstraction

OpenCL Next

- Enable diverse, heterogenous computing architectures
 - Not just GPUs
- Enable vendors to ship targeted functionality for their customers/markets
 - While still being formally conformant
- Empower middleware + libraries to emulate features
 - Reduces implementation cost
 - Enables portability without vendors implementing all features
- More features in OpenCL become optional - e.g. floating point
 - Enable smaller devices to support recent OpenCL versions
 - Dependency tree between features
- Enhanced query mechanisms
 - Application queries which OpenCL features are supported by a device




OpenCL Next Feature Sets

- **Optional Feature Sets to reduce market fragmentation**
 - Opportunity for industry to flexibly agree on commonly supported functionality
 - Features optional in both language and API
- **Vendor can support ANY combination of features**
 - Not necessary to support any Feature Set - as long as supported features are conformant
 - If features include 100% of a Feature Set - can claim conformance to that Feature Set
 - Supporting popular Feature Sets may help drive sales
 - An Implementation may support multiple Feature Sets

OpenCL 2.2 Functionality

■ = queryable, optional feature



aj.guillon@yetiware.com

YouTube: <http://goo.gl/6HY0bS>