deRSE24

Research Software Engineering @ Intel

Chris Dahnken
Extreme Computing Software and Systems
DCAI.DAIS.CSE.CWE.TCE.XCSS

Foreword

- The following thoughts are
  - a personal perspective
  - reflections from inside Intel (yet they are likely common in the industry)
  - Very HPC and AI centric

Agenda

- There are no RSEs at Intel .. Or are there?
- What do we do and why does it matter
- Skills
- Career
- Diversity
- Summary

There are no RSEs at Intel … or are there?

- Intel is a manufacturing company at its heart and the largest number of employees work in that area
- There are about 15000 software engineers at Intel
  - Operating Systems
  - Drivers
  - Firmware
  - Compilers
  - Libraries
  - Tools
  - Applications

None of these have an RSE job title – still I believe many would say they are RSEs

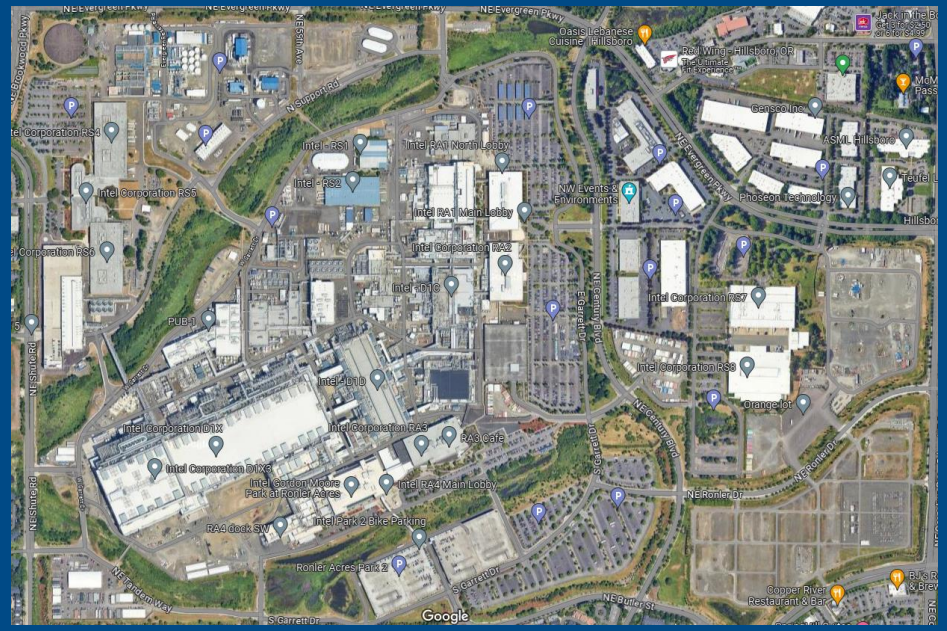

Credit: Google Maps

What do we do?

- Users/customers expect that software supports the latest processor technologies when they buy them
- This goes beyond the driver level, but requires the consideration of a whole ecosystem
- Example: CP2K dependencies
  - COSMA, ELPA, GSL, HDF5, LibINT, LibVdWXC, LibXC, LibXSMM, SPFFT, SPGLib, SuperLU
- Given the complexity of new CPU technologies, this is a job too hard for open-source communities and commercial software vendors to take on. It required a new type of engineer, working at the intersection of domain knowledge, software engineering and computer architecture, an HPC Application Engineer

CPU architects will show you this and expect that you know how to optimize your code!

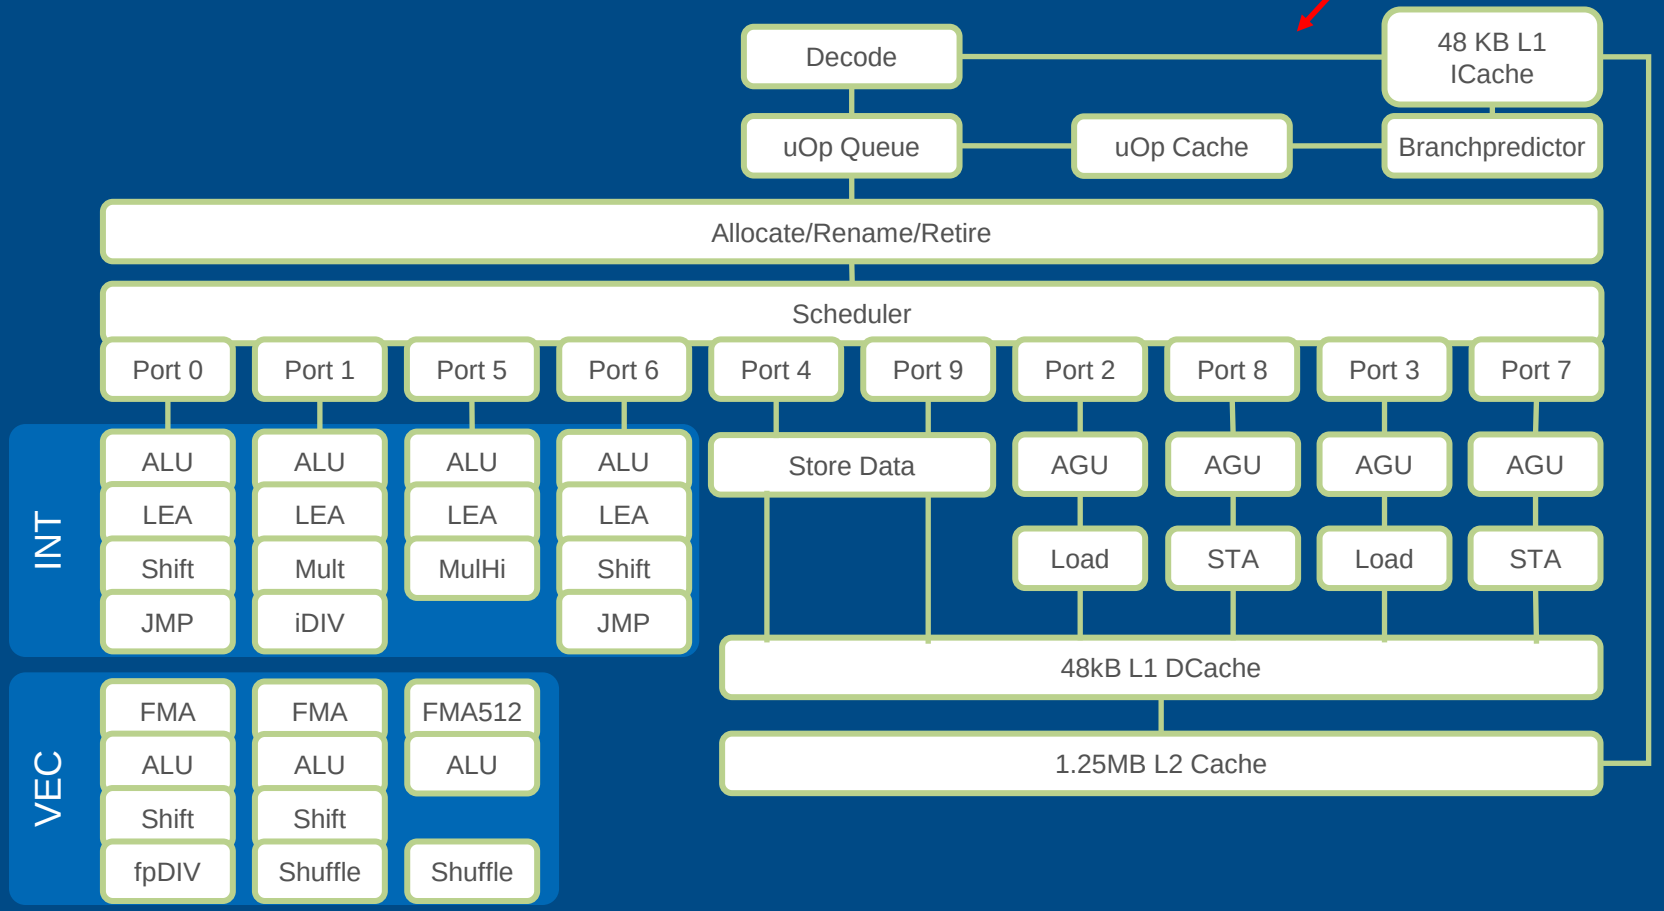

¯_(ツ)_/¯

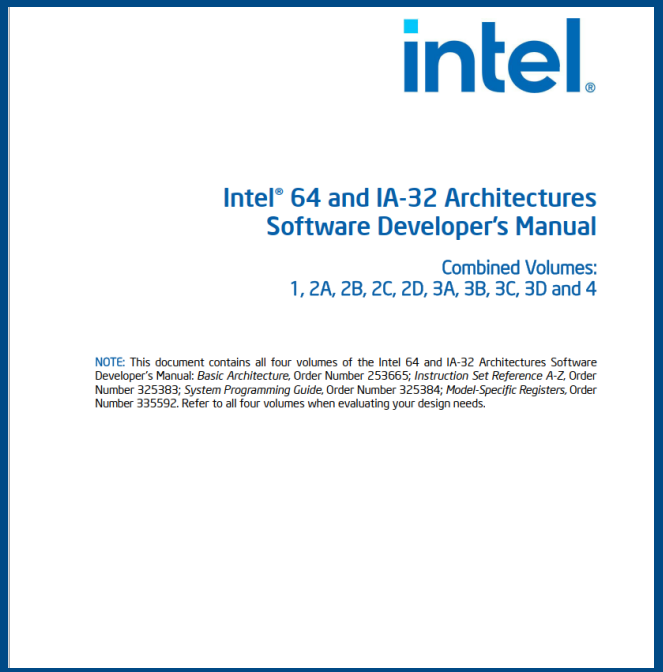

What do we do?

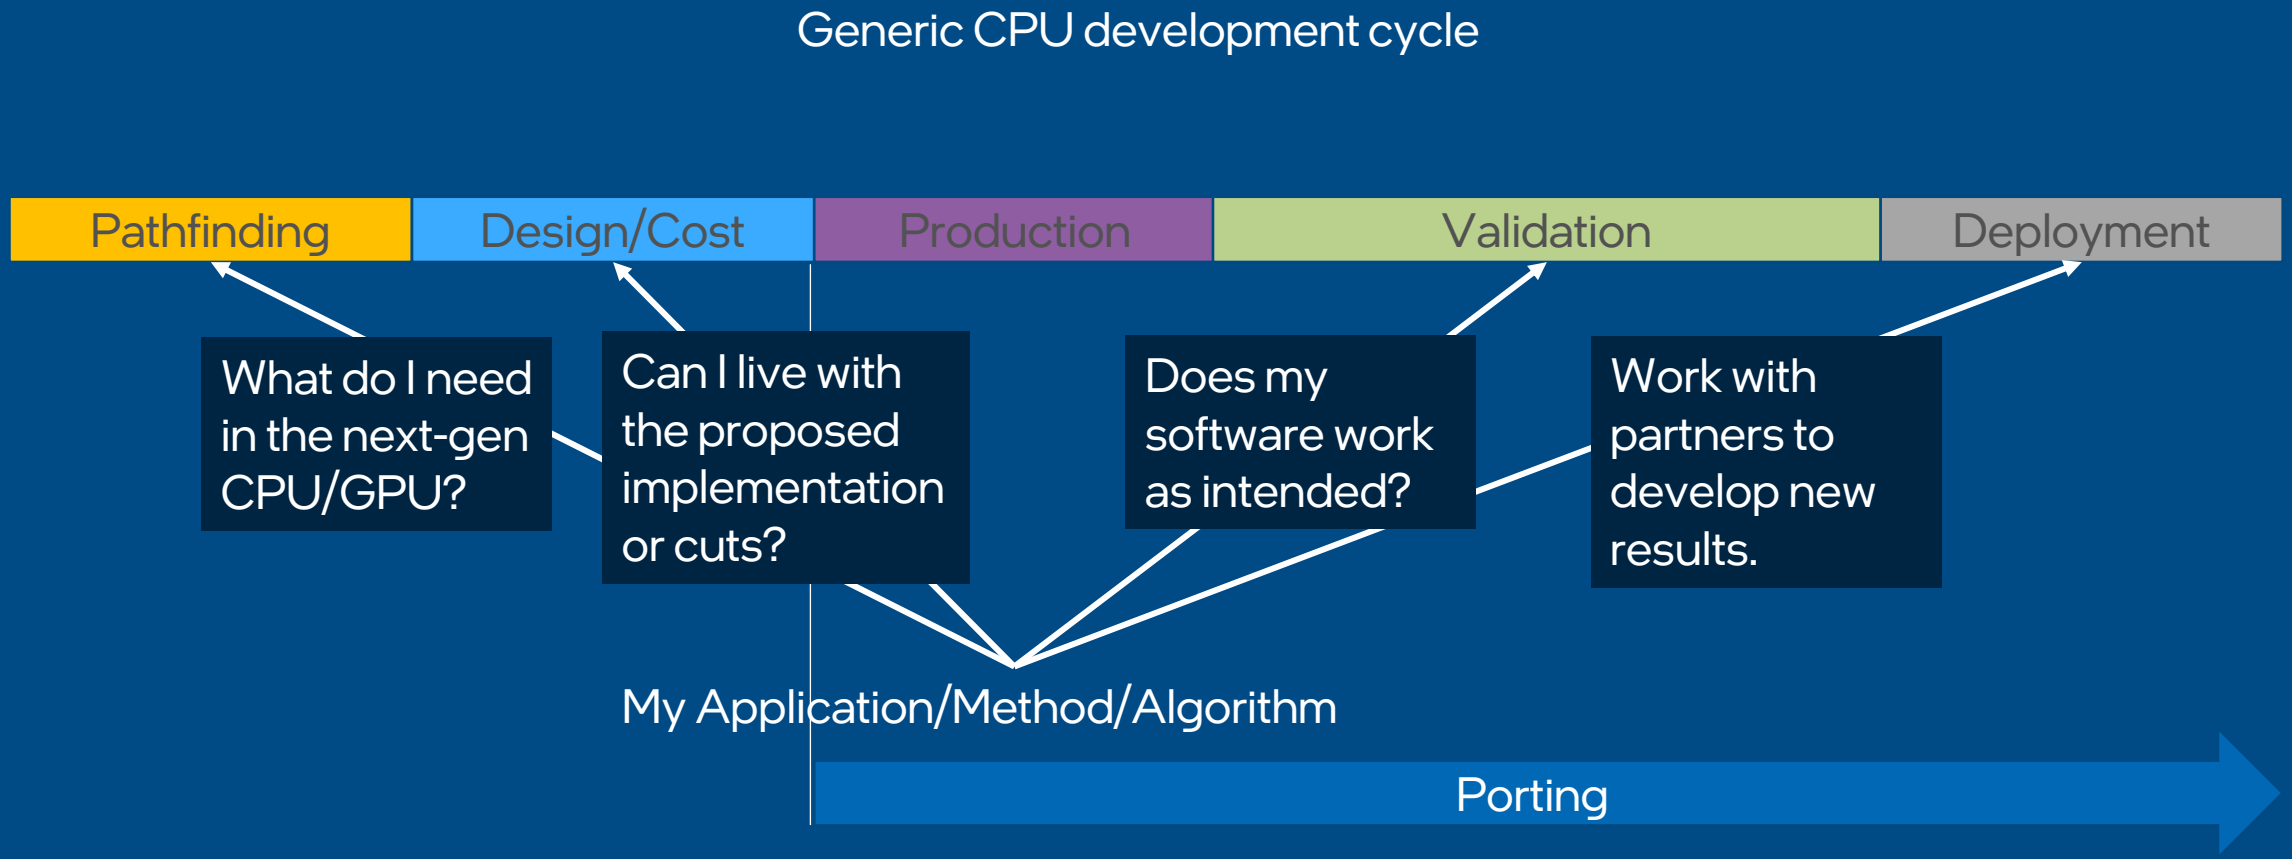

Application Engineers/RSEs work with the architecture teams across the spectrum of the CPU/GPU development cycle

What do we do?

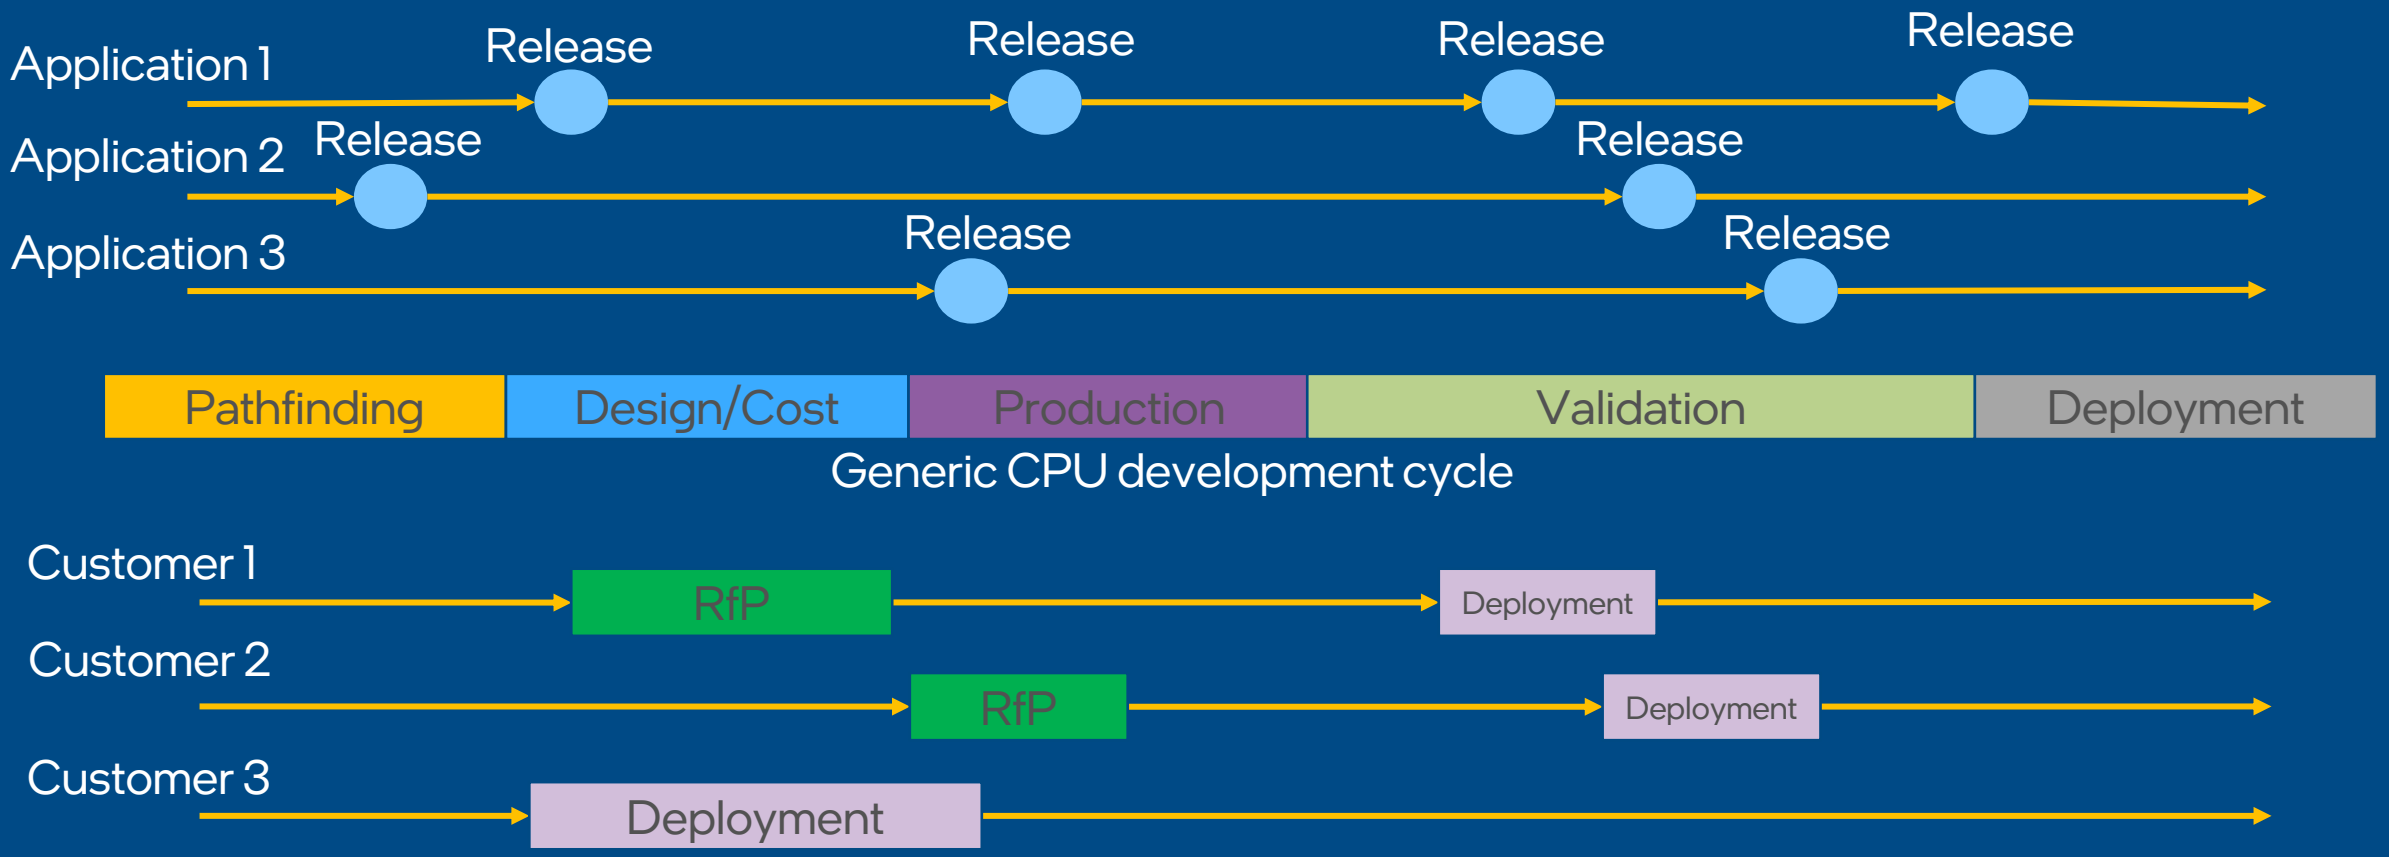

Application enabling does interfere with the application release cycles as well as with customer sales decisions. An application that is required for in a sales decision in a Request for Procurement (RfP) might need to be simulated or projected, but needs to be available at deployment for acceptance and operation.

Skills – Minimal

Numerical Methods	Programming /SW	Hardware Comp. Arch.	Domain Knowledge (MSc)
BLAS LAPACK FFT	C/C++ FORTRAN OpenMP MPI Linux Bash SLURM	SIMD	CFD QCD AI/ML Seismic/O&G Quantum Chemistry Molecular Dyn. Financial Cosmology Fusion Numerical Weather Compiler tech. Electrical Eng.

Skills - Desired

Numerical Methods	Programming /SW	Hardware Comp. Arch.	Domain Knowledge (PhD)
BLAS	C/C++	Pipelines	CFD
LAPACK	FORTRAN	Caches	QCD
FFT	OpenMP	OOO Execution	AI/ML
Finite Difference	MPI	SIMD	Seismic/O&G
Finite Elements	Assembly	SIMT	Quantum Chemistry
Finite Volume	OpenCL	Perf. Monitoring	Molecular Dyn.
Lattice Boltzmann	SYCL	Networking	Financial
Monte Carlo	CUDA	Infiniband	Cosmology
	Pthreads	Tracing	Fusion
	Linux	Simulation	Numerical Weather
	Bash		Compiler tech.
	SLURM		Electrical Eng.
	EMACS!		

Career

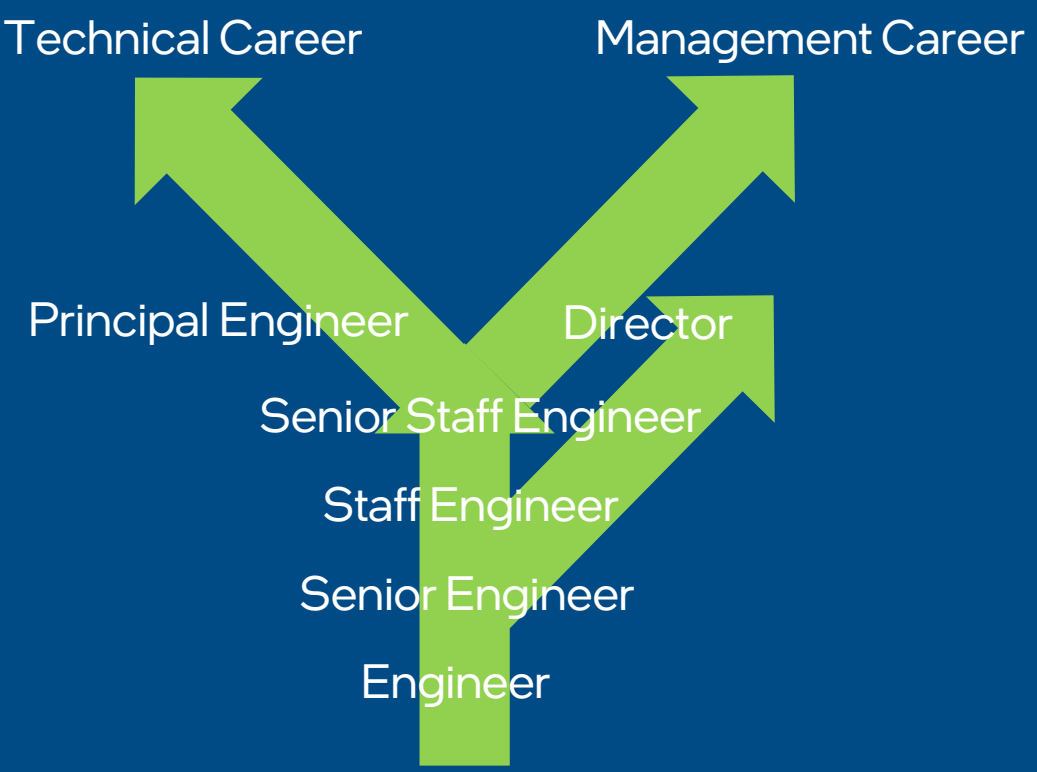

- Application Engineers/RSEs are generally on the technical career path
- Becoming a manager is common from Senior Engineer level on
- It is totally acceptable to manage more senior people
- Moving to Principal Engineer level requires approval from a committee
- About 10% of the engineering workforce make to this level

I won't talk about salaries here, but there are sources on the web that give precise numbers, e.g. www.levels.fyi.

Diversity (focusing on F/M/D here)

- Diversity is a company priority and a corporate value
  - “We value diversity and embrace difference.”
- A lot of senior leaders (Senior/Executive Vice Presidents) are female, but still the minority (F/M=6/10) in the top leadership. This changes frequently and it has during the creation of this talk.
- Personally, my results are mixed. I have a hard time to maintain >10% females on the team.
  - Too HPC focused? Too domain focused?
  - In many core development roles, e.g. compiler and libraries, the number of females is higher!

Summary

- Application Engineers (AE) at Intel are an industrial instantiation of an RSE
- AEs are responsible for
  - Porting to and optimizing for new technologies
  - Contributing to compiler and library development
  - Providing insight into application requirements for CPU/GPU architecture
  - Supporting sales
- AEs are on the technical career path, also serve as technical managers


intel

What do we do – Examples – LibXSMM

- LibXSMM (https://github.com/libxsmm/libxsmm) is a library for specialized dense and sparse matrix operations as well as for deep learning primitives such as small convolutions. The library is targeting Intel Architecture
with Intel SSE, Intel AVX, Intel AVX2, Intel AVX-512 (with VNNI and Bfloat16), and Intel AMX (Advanced Matrix Extensions) supported by future Intel processor code-named Sapphire Rapids. Code generation is mainly based on Just-In-Time (JIT) code specialization for compiler-independent performance (matrix multiplications, matrix transpose/copy, sparse functionality, and deep learning). LIBXSMM is suitable for "build once and deploy everywhere", i.e., no special target flags are needed to exploit the available performance. Supported GEMM datatypes are: FP64, FP32, bfloat16, int16, and int8.

What do we do – Examples – LibXSMM

- Original problem statement:

$$C = C + \sum_i A_i \cdot B_i$$

where A and B are 23x23 matrices (which is the number of basis functions required to represent water in CP2K).

- Standard libraries BLAS(MKL) are designed to work well with large matrix sizes, not small ones.
- Used in CP2K, SeiSol, NekBox,Nek5000,PyFR, and many more (https://libxsmm.readthedocs.io/en/latest/#performance)