

James Abel

Email: j@abel.co

Experience **Nocimed, Inc., San Mateo, CA (<https://nocimed.com/>)** **2017-2019**
Vice-President of Engineering

- Architect, develop and deploy Nociscan-LS.
- Nocimed's products aide medical professionals (e.g. spine surgeons) in the diagnosis of degenerative disc disease (DDD) pain by applying Signal Processing, Machine Learning and other algorithmic and visualization techniques to MRI (Spectroscopy) scans.
- Responsible for the entire product stack: e.g. acquisition ingestion from medical facilities (MRIs), cloud-based signal processing and machine learning/algorithms (Python, MATLAB), and a web-based portal (e.g. user report access, administration, billing).
- Managed a small team of engineers while also doing technical work (architecture, development, automated testing, code reviews). Interfaced to stakeholders: IRBs, QA, Regulatory, and key customers/development partners (spine surgeons).
- Primary languages/infrastructures/techniques: Python (including PyQt, Django, pytest, packaging), CI/CD, AWS (various cloud-native services), github, OKRs, Agile/Scrum.

Hardware/Software Consultant, Redwood City, CA **2017**

- Architect, develop and test applications, primarily for the biomedical industry. Utilize Signal Processing and Machine Learning techniques.
- Primary languages: C/C++, Python
- Contributor to several open source projects (see <https://github.com/jamesabel>).

Intel Corporation, Santa Clara, CA **1988-2016**
Principal Engineer – Software and Services Group (Santa Clara, CA) **2012-2016**

- Chief Tracing Architect responsible for the architecture and design of application workload 'Tracers'. These Tracers take a software application workload and create a set of files that represent relevant execution information that is fed into CPU and full platform simulators for performance and power analysis.
- Utilized machine learning techniques for microprocessor and system design. Developed data cleaning and defined ML vector spaces. Used clustering techniques to select the most relevant aspects of the workloads to reduce computation requirements while maintaining analysis fidelity. Written mainly in Python using scientific computing libraries (e.g. scipy, numpy, matplotlib), along with some C/C++.
- Established an Intel-wide trace repository with over 50K traces that were utilized to create 50M datasets per year.
- Created a system and database (> 100M entries) to monitor and create reports on the usage of these trace files.
- Lead cross-Intel initiatives on Tracing and Simulation
- Lead and mentor a technical team of ~20 engineers in the areas of

hardware and software architecture and design.

- Established engineering quality standards, coding standards, and test requirements (unit, system, regression). Conducted code and design reviews

Principal Engineer – Atom Core Architecture Group (Austin, TX)

2011-2012

- Responsible for identification, acquisition and use of emerging workloads to improve the performance and/or power of future Atom cores, focusing on 2015+ product lines. Performed deep architecture and micro-architecture analysis both pre-silicon and post-silicon to identify improvements. Did path-finding and technical readiness analysis. Focus areas: media and DSP workloads.
- Analyzed industry standard benchmarks such as SPEC CPU, EEMBC FP and SysMark. Analyzed ISV applications and workloads.
- Created proposals for these improvements, evaluated the technical feasibility and negotiated with design, planning and software teams to get product buy-in. Established programming models associated with these features. Wrote detailed technical architecture specifications.
- Utilized proprietary tracing, simulation and analysis environments to model and analyze improvements such as new instructions and/or micro-architectural features. Utilized core and platform behavioral models. Specified and wrote new tools required to perform these evaluations. Created new analysis techniques and best-known-methods for performance and power analysis.
- Served on corporate-wide technical committees to steer Intel's architectural directions. Also served on committees to direct and select papers for internal technical conferences.

Principal Engineer, Software and Services Group (Chandler, AZ)

2002-2011

Sr. Software Engineer, Software and Solutions Group

1997-2002

- Responsible for performance optimizations of key software applications, and used these applications to improve Intel's future processor and platform products. Worked as a virtual member of the architecture and product groups, infusing software knowledge into Intel products.

Delivered workload characterization, tracing and analysis to Intel's processor and platform architecture groups. Performed pre-Si analysis in a simulation-based environment using both high level behavioral and RTL models. Gathered, root-caused and dispositioned sightings of performance anomalies in key applications. Developed point tools for application characterization, tracing, simulation and analysis (both pre- and post-Si).

Drove key performance and performance/power improvements into Intel's processor and platform products (e.g. Intel® Core™ processor family and Intel® Core™ Atom processor). This included co-owning and driving new instruction definition for HPC/multimedia/imaging/3D into instruction sets such as SSSE3-SSE4.1, AVX, as well as general micro-architectural performance improvements. Also drove performance

improvements into Intel compilers.

Selection and development of kernels, tracing, modeling new instructions in a proprietary simulator, participating in performance ‘dungeons’, and analysis and performance estimations of kernels/workloads. Provided speedup estimates of new instructions on kernels and workloads as well as input into the micro-architecture to improve the performance (and performance/power) of particular implementations.

Also worked with strategic ISVs to optimize their code for Intel Architecture as well as used this code for tech readiness.

Sr. Software Engineer, Intel Architecture Labs *1994-1997*

- Responsible for optimization of multi-media and signal processing applications. Major projects included implementation of a Dolby Digital decoder using MMX™ Technology (<http://developer.intel.com/technology/itj/q31997.htm>) and speech recognition programs. Was also a developer on Native Signal Processing/IA-SPOX project (signal processing OS that resided in a Windows driver).

Sr. Design Engineer, Intel Communications Division *1991-1994*

- Design engineer on a single-chip communications product. Responsible for interconnect design, as well as full-chip RTL models. This product targeted wireless communications, utilizing two DSPs and one microcontroller (80186 core).

CAD Engineer, ASIC/Embedded Processor Division *1988-1991*

- Responsible for behavioral models of embedded cores (8051, 80186) and other standard cells. Also responsible for software for automated test.

GenRad Corporation, Phoenix AZ **1985-1987**
Hardware Design Engineer

- Designer of automated test systems. Targets included micro-processor based adapter boards, systems and SCSI disks.

Dynacomp, Phoenix AZ **1983-1984**
Hardware Design Engineer

- Designer of single-board microcomputer systems based on the 68000 processor and Multi-bus for automated news services. Wrote low level networking drivers.

Holder of 6 patents:

7,743,233 [Sequencer address management](#)

7,689,641 [SIMD integer multiply high with round and shift](#)

7,630,585 [Image processing using unaligned memory load instructions](#)

7,043,719 [Method and system for automatically prioritizing and analyzing performance data for one or more system configurations](#)

5,991,787 [Reducing peak spectral error in inverse Fast Fourier Transform using MMX.TM. technology](#)
5,752,071 [Function coprocessor](#)

Education **MSCS, Arizona State University, Tempe, AZ - 1991**
 BSEE, Bradley University, Peoria, IL - 1983

Technical Skills **Major Languages/Programming skills: C/C++, assembler/intrinsics, Python, OOP, threading, SIMD programming, performance and power analysis**
 Multimedia algorithms, signal processing, machine learning, IA architectures and micro-architectures