

Intel® oneAPI Base Toolkit Helps SonoScape Optimize the Performance of its S-Fetus 4.0 Obstetric Screening Assistant

“With our commitment to the independent R&D and innovation of medical equipment, SonoScape is pleased to state that our cutting-edge AI technology, powered by Intel® oneAPI architecture, has been able to realize its potential to serve medical institutions around the world.”

Feng Naizhang
Vice President, SonoScape

Obstetric screening is key to reducing maternal and perinatal mortality; however, conventional obstetric screening methods require high levels of medical expertise and are both time- and labor-intensive. To address these issues, SonoScape has launched a smart obstetric screening system based on artificial intelligence (AI) and other technologies. The system automates output of screening results through automatic structure recognition, measurement, classification, and diagnosis to significantly enhance efficiency and reduce the workload of doctors.¹

The S-Fetus 4.0 Obstetric Screening Assistant² uses deep learning to power a smart scenario-based work model that allows doctors to perform sonography without the need to manually control equipment and enables real-time dynamic acquisition of standard planes and automatic measurement of fetal biometry and growth index, an industry first. SonoScape's aim is to simplify obstetric screening workflows and make it easier for patients to get care. To enhance its performance, SonoScape used the Intel® oneAPI Base Toolkit for cross-architecture development and optimization to speed processing of multimodal data. Through a platform based on the Intel® Core™ i7 processor, performance was increased by approximately 20x³ while achieving higher price performance, cross-architecture scalability, and flexibility.

Background: Applications and Challenges of Diagnostic Ultrasound in Obstetric Examinations

Diagnostic ultrasound is a technique in which ultrasound is used to measure the data and morphology of a patient's physiology or tissue structure to discover diseases and provide medical guidance.⁴ Owing to safety, non-invasiveness, cost performance, practicality, repeatability, and broad adaptability, the market for diagnostic ultrasound equipment is growing rapidly. According to data from Fortune Business Insights, the size of the global diagnostic ultrasound equipment market was USD 7.26 billion in 2020, and is expected to reach USD 12.93 billion by the end of 2028, representing a compound annual growth rate (CAGR) of 7.8%.⁵

Though 2D ultrasound is indispensable for the diagnosis of obstetric and gynecological diseases (especially in intrauterine fetal testing), conventional ultrasonography techniques rely heavily on the expertise of the sonographer. As time-consuming and skill-intensive manual operations are required throughout the entire process, ultrasonography poses challenges to hospitals in smaller communities and less-developed areas that have limited access to medical technology.

To address these issues, SonoScape has developed a smart diagnostic ultrasound solution based on AI technologies that are capable of classification, detection, and segmentation of a variety of anatomical structures from ultrasound images through deep learning algorithms represented by convolutional neural networks (CNNs).⁶ However, the current diagnostic ultrasound solution faces several challenges:

- The equipment requires a high amount of user intervention and possesses inherent delays, such as when the operator must adapt to different operating procedures when switching between modes.
- Computing power requirements are rising as AI algorithms grow in complexity. These algorithms often use external accelerators, such as GPUs, that increase cost, use more power, and require additional testing and certification. Continuous AI optimization for the best performance and user experience has become a key challenge.

SonoScape Utilizes the Intel oneAPI Base Toolkit to Optimize the Performance of its S-Fetus 4.0 Obstetric Screening Assistant

SonoScape S-Fetus 4.0 Obstetric Screening Assistant

Based on the standardized collection and measurement of ultrasound scan sections, clinicians can use obstetric screening to detect most fetal structural abnormalities. SonoScape's proprietary S-Fetus 4.0 Obstetric Screening Assistant is the first globally available smart obstetric screening technology based on deep learning. When combined with the SonoScape P60 and S60 ultrasound platforms, the S-Fetus 4.0 is capable of real-time recognition of sections during the sonography process, automatic acquisition of standard sections, automatic measurement, and automatic feeding of results into the corresponding fetal growth sections of the medical report. Boasting the first smart obstetric screening function in the industry, the S-Fetus 4.0 significantly improves on conventional

human-computer interaction methods by providing a smart scenario-based work model that allows doctors to perform sonography without the need to manually control complex equipment, simplifying the sonogram process, improving efficiency, and reducing the workload of the sonographer. The function provides effective frontend quality control during the ultrasound process, enhances screening quality, and provides additional guiding data in real time to help both doctors and patients.



Figure 1. SonoScape's professional P60 obstetrics device equipped with the S-Fetus 4.0

Utilizing core algorithms, original architecture, and cross-architecture hardware, the S-Fetus 4.0 achieves a fundamental technical breakthrough that provides a smart, scenario-based, full-process, and easily adoptable solution to improve the work efficiency and consistency of doctors. Comprehensive scenario-based functions ensure that doctors do not need to switch between manual and smart modes by default throughout the entire process, and reports can be completed with the swipe of a finger.




Figure 2. Process diagram of the S-Fetus 4.0 Obstetric Screening Assistant

The front end of the S-Fetus 4.0 generates multimodal data in accordance with scenario requirements, while post processing handles reconstruction, processing, and optimization. Working on reconstructed and optimized data, the real-time AI recognition and tracking module analyzes and extracts standard surfaces. In this process, the standard surface decision-making and dispatch module follows a predefined strategy to adaptively extract quantified features, then it performs quantitative analysis and automatically integrates into subsequent operations.

During development, SonoScape and Intel engineers worked together to address several challenges:

- **Further performance optimization.** Many relevant deep learning algorithms must work in tandem to rapidly process tasks that use different data types and to optimally execute user-initiated tasks without latency. This results in higher computing power and algorithm optimization requirements for ultrasound platforms.
- **Mobile application demands.** The SonoScape diagnostic ultrasound system with S-Fetus 4.0 Obstetric Screening Assistant is a mobile system with limits on overall power consumption and system size, making it a challenge to use discrete GPUs.
- **Cross-architecture expansion for different scenarios.** The S-Fetus 4.0 Obstetric Screening Assistant needs to support migration and expansion across multiple architectures to operate in a variety of complex scenarios.

To solve these challenges, SonoScape partnered with Intel to optimize the AI performance of its obstetric screening assistant by using the Intel oneAPI Base Toolkit.

Intel oneAPI Toolkits

OneAPI is a cross-industry, open, standards-based unified programming model that delivers a common developer experience across architectures for faster application performance, more productivity, and greater innovation. The oneAPI initiative encourages collaboration on the common specifications and compatible oneAPI implementations across the ecosystem.

The model is designed to simplify the development process across multiple architectures (such as CPUs, GPUs, FPGAs, and other accelerators). With a complete set of cross-architecture libraries and tools, oneAPI helps developers develop performant code quickly and correctly across heterogeneous environments.

As shown in Figure 3, the oneAPI project aims to build on Intel's rich heritage of CPU tools and expand to XPU's. It includes a complete set of advanced compilers, libraries and porting, analysis, and debugging tools. Intel's reference implementation of oneAPI is a set of toolkits. Intel oneAPI Base Toolkit for Native Code Developers is a core set of high-performance tools for building C++, Data Parallel C++ applications, and oneAPI library-based applications.

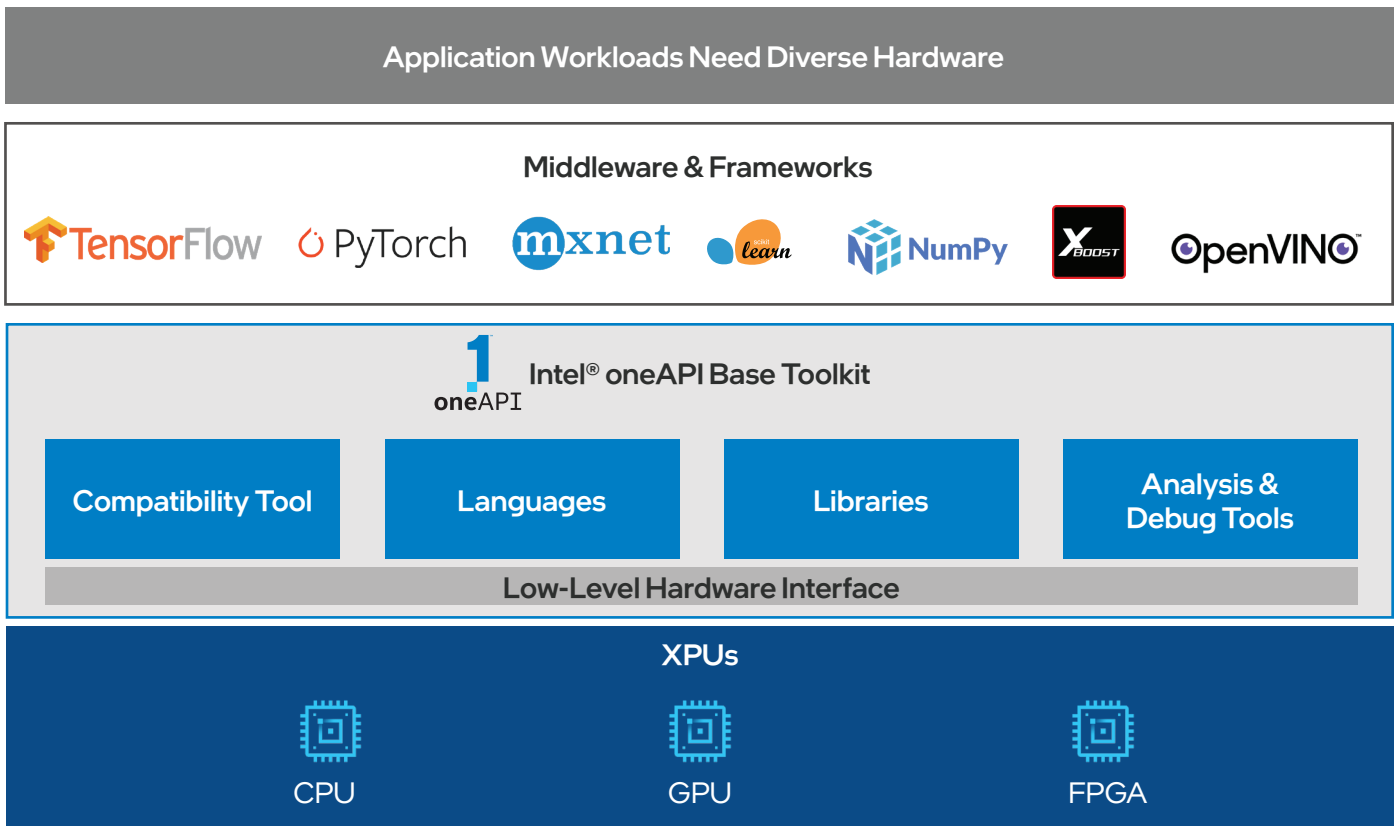


Figure 3. Intel oneAPI Base Toolkit

Intel oneAPI Base Toolkit Helps SonoScape Optimize the Performance of its Obstetric Screening Assistant

After integrating the Intel oneAPI Base Toolkit to their system, SonoScape noted several paths to optimization.

At the hardware layer, the solution utilizes a computing architecture based on an 11th Gen Intel® Core™ i7 processor that delivers enhanced execution performance, features a new core and graphics architecture, and provides AI-based optimization for excellent performance for various loads. Equipped with Intel® Deep Learning Boost (Intel® DL Boost) technology, the processor provides strong support for AI engines and enhanced performance for complex loads such as AI and data analysis.

11th Gen Intel Core processors also have integrated Intel® Iris® Xe graphics, enabling workloads to leverage this integrated GPU. It can support a rich variety of data types and features a low-power architecture.

The data processing flow of the solution is shown below (Figure 4). Equipped with cores optimized for the handling of data-intensive loads, the Intel Iris Xe graphics are responsible for real-time recognition and tracking processes and the realization of high-frequency real-time execution (each image frame must be processed or intelligently inferred).

The Intel Core i7 processor handles standard surface decision-making and dispatch; adaptive section feature extraction, quantitative analysis, and other processes; and the execution of operational logic and AI inference during downtime. Data-intensive and responsible for logical inference, the multimodal data optimization and processing module has been optimized in five key aspects through the oneAPI Toolkit. After optimization, the SonoScape obstetric screening assistant can flexibly utilize all CPU and iGPU resources, providing enhanced performance to meet operational demands and improve the patient experience.

SonoScape and Intel focused on the optimization and performance testing of the following platform:




Figure 4. Architecture of the SonoScape obstetric screening assistant

Comprehensive Performance Optimization using Intel Software Tools

Optimization #1: First, SonoScape used the Intel® VTune™ Profiler to analyze their workload. The profiler can quickly identify CPU and GPU load performance bottlenecks and provide relevant information. As shown in the figure below, vector processing makes full use of Intel's high instruction throughput and supports the parallel processing of data to rapidly improve performance over scalar operations.




Figure 5. Scalar processing vs. Vector processing

SonoScape also made use of the DPC++ Compiler in the oneAPI toolkit to recompile its code and generate vector instructions for enhanced performance, reducing the processing speed of the workload from 141 ms to just 33 ms.⁷

Optimization #2. Once performance bottlenecks were identified by the VTune Profiler, SonoScape replaced them with APIs from Intel® Integrated Performance Primitives (Intel® IPP), a cross-platform software library of functions that include accelerators for image processing, signal processing, data compression, encryption mechanisms, and other applications. Intel IPP can be optimized for CPUs to unlock the latest features of Intel architecture platforms (such as AVX-512) to improve application performance.

For example, the `ippsCrossCorrNorm_32f` and `ippsDotProd_32f64f` functions can improve performance by removing dual-layer loop calculations and multiplication/addition loops. Through such optimization, SonoScape was able to further improve the processing speed of the workload from 33 ms to 13.787 ms.⁷

Optimization #3. Originally developed by Intel, the Open Source Computer Vision Library (OpenCV) OpenCV can be used to develop real-time image processing, computer vision, and pattern recognition programs, and supports the utilization of Intel IPP for accelerated processing.⁸ By replacing OpenCV functions in the source code with IPP functions, the solution scales well in large-scale data scenarios and performs well across all generations of Intel platforms.

Optimization #4. SonoScape's S-Fetus 4.0 obstetric screening assistant also utilizes the Intel® DPC++ Compatibility Tool to efficiently migrate existing CUDA code to DPC++, ensuring cross-architecture compatibility and minimizing the time required for migration. As shown in Figure 6, the tool provides powerful interactive functions to help developers migrate CUDA code, including kernel code

and API calls. The tool can automatically migrate 80-90 percent⁹ of the code (depending on complexity) and embeds comments to help developers complete the manual step of the migration process. In this case study, nearly 100 percent of the code was automatically migrated in a readable and usable manner.

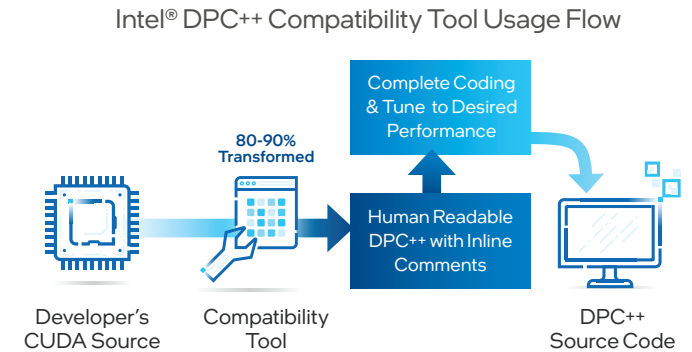


Figure 6. Workflow chart of the Intel DPC++ Compatibility Tool

After these optimizations were completed, the performance of the SonoScape S-Fetus 4.0 running on heterogeneous platform based on Intel oneAPI DPC++ was increased by nearly 20x that of the baseline performance data recorded before optimization, as shown in figure 7.⁷




Figure 7. Performance Improvement with the Intel oneAPI Base Toolkit⁷

(Baseline: Code before optimization; Optimization 1: Intel oneAPI DPC++ Compiler; Optimization 2: Intel IPP used to replace loop source code; Optimization 3: Intel IPP used to replace OpenCV functions; Optimization 4: CPU + iGPU execution after CUDA migration)

Result: Excellent Performance and Cross-Architecture Scalability

By using Intel Core i7 processors with integrated Intel Iris Xe graphics to provide underlying computing power and the Intel oneAPI heterogeneous platform for optimization, the SonoScape obstetric screening assistant was able to balance performance, cost-effectiveness, and scalability across multiple platforms.

- **Performance.** By using Intel XPU and Intel oneAPI Toolkits, the SonoScape obstetric screening assistant was able to realize up to 20x improved performance vs. unoptimized systems, laying a solid foundation for efficient obstetric diagnostic ultrasound.⁷
- **Cost savings.** By performing comprehensive optimization and using the powerful performance and flexible architecture of Intel Core i7 processor, SonoScape only needs CPU and iGPU resources to achieve its performance targets. These hardware simplifications reduce the demands for power supply, heat dissipation, and space. The solution can now be mounted on smaller diagnostic ultrasound equipment for more flexible configuration options. The integration of CPU and iGPU resources also provides longer battery life, along with higher scalability and reliability.
- **Heterogeneous Scalability.** The solution supports unified programming on heterogeneous hardware such as CPUs and iGPUs, improves the development efficiency of cross-architecture programming, and enables the flexible execution of obstetric screening assistants on different hardware configurations all while ensuring a smooth user experience.

Outlook: Accelerated Integration of AI and Medical Applications

Smart diagnostic ultrasound is a key application of the integration of AI and medical technologies that helps to reduce doctor workloads and improve the speed of medical processes.¹⁰ To facilitate the use of AI and medical applications, Intel is working with partners such as SonoScape to accelerate digital innovation through XPU architecture made up of CPUs, iGPUs, dedicated accelerators, FPGAs, and software and hardware products such as the oneAPI programming model in the medical industry.

“The Intel® oneAPI Base Toolkit helped us optimize key modules in an efficient manner, realizing a 20x⁷ increase in performance and unified development on cross-architecture XPU platforms. Through Intel technologies, our obstetric screening assistant has achieved breakthroughs in terms of performance and scalability and can now provide a more efficient means of smart obstetric diagnosis to help medical institutions transition from conventional ultrasound to smart ultrasound and assist doctors in precise and efficient work to improve patient outcomes.”

Zhou Guoyi
Head of SonoScape Medical Innovation Research Center

About SonoScape

Founded in 2002 in Shenzhen, China, SonoScape has committed itself to “Caring for Life through Innovation” by providing ultrasound and endoscopy solutions. With seamless support, SonoScape provides worldwide sales and service in more than 130 countries, benefiting local hospitals and doctors with comprehensive imaging diagnostic evidence and technical support. Investing 20 percent of total revenue into R&D annually, SonoScape has continuously introduced new medical products and technologies into the market each year. It now expands into seven R&D centers in Shenzhen, Shanghai, Harbin, Wuhan, Tokyo, Seattle, and Silicon Valley. For more information, please visit our official website www.sonoscape.com.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge, and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.

Solution provided by:



¹ Efficiency increase claim of 50% is based on assessment data after clinical evaluation from 18 doctors of intermediate & senior experience in 5 medical facilities after 1 month period. Reduction in workload claim of 70% based on evaluation of necessary steps to complete medical checkup using standard operation procedures vs. S-Fetus.

² For more information about S-Fetus 4.0 Obstetric Screening Assistant, please visit https://www.sonoscape.com/html/2020/exceed_0715/113.html

³ Test results provided by SonoScape. Test configuration: Intel® Core™ i7-1185GRE processor @ 2.80GHz, Intel Iris® Xe graphics @ 1.35 GHz, 96EU, Ubuntu 20.04, Intel® oneAPI DPC++/C++ Compiler, Intel® DPC++ Compatibility Tool, Intel® oneAPI DPC++ Library, Intel® Integrated Performance Primitives, Intel® VTune™ Profiler

⁴ Wells, P. N. T., "Physical Principles of Ultrasonic Diagnosis." *Medical and Biological Engineering* 8, No. 2 (1970): 219–219.

⁵ <https://www.fortunebusinessinsights.com/industry-reports/ultrasound-equipment-market-100515>

⁶ Shengfeng Liu, et al., "Deep Learning in Medical Ultrasound Analysis: A Review." *Engineering* 5, No. 2 (2019): 261–275

⁷ Test results provided by SonoScape. See backup for testing configurations.

⁸ <https://en.wikipedia.org/wiki/OpenCV>

⁹ <https://www.intel.com/content/www/us/en/developer/articles/technical/heterogeneous-programming-using-oneapi.html>

¹⁰ Luo, Dandan, et al., "A Prenatal Ultrasound Scanning Approach: One-Touch Technique in Second and Third Trimesters." *Ultrasound Med Biol.* 47, No. 8 (2021): 2258–2265. https://www.researchgate.net/publication/351951854_A_Prenatal_Ultrasound_Scanning_Approach_One-Touch_Technique_in_Second_and_Third_Trimesters

Backup

Testing by SonoScape as of Sep 3, 2021. Test configuration: Intel® Core™ i7-1185GRE processor @ 2.80GHz, with or without Intel Iris® Xe graphics @ 1.35 GHz, 96EU, Ubuntu 20.04, Intel® oneAPI DPC++/C++ Compiler, Intel® DPC++ Compatibility Tool, Intel® oneAPI DPC++ Library, Intel® Integrated Performance Primitives, Intel® VTune™ Profiler

Notices and Disclaimers

Performance varies by use, configuration, and other factors. Learn more at www.Intel.com/PerformanceIndex

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software, or service activation.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

0422/E0H/MESH/PDF 350912-001US