



# Modeling and optimizing the performance of solid state drives (SSDs) using Intel® CoFluent™ Studio

---

## Abstract

In this paper, a team of engineers from the Intel Non-Volatile Memory Solutions Group presents an approach to explore improving the architecture of solid-state drives (SSDs). This explorative approach is based on models captured using Intel® CoFluent™ Studio. The team's goal was to identify how best to use a model-based approach to optimize the performance and cost of next generation Intel® Solid State Drives (Intel® SSDs), by making informed, data-driven decisions earlier in the design process.

Using the model-based approach, the team created an executable specification of SSD architecture in the form of an Intel CoFluent model. Using that model, the team then explored the functional behavior and performance of the SSDs for different workloads. The executable specification helped the team identify and solve architectural issues, as well as optimize the performance of the SSDs before hardware/software development began.

“The SSD model captured using Intel® CoFluent™ Studio enabled our firmware development team and system-on-a-chip architecture team to make informed, data-driven decisions earlier in the design process, and ultimately led to increased performance and reduced cost.”

– David Carlton  
Senior Media-Systems Architect  
Intel Corporation

## Authors

**David Carlton**  
Intel Corporation  
**Jérôme Lemaitre**  
Intel Corporation

## Background

Solid-state drives (SSDs) are storage devices that are usually found in consumer mobile devices (such as notebooks and tablets), desktop computers, and data centers. Unlike traditional electro-mechanical disks (such as hard disk drives, or HDDs) which contain spinning disks and moveable read/write heads, SSDs use integrated circuits. As a consequence, SSDs typically have faster

access times than HDDs. They also tend to run silently, have a smaller form factor, and are more resistant to physical shocks.

The architecture of SSD integrated circuits consists mainly of a system-on-a-chip (SoC) controller and NAND flash non-volatile memory chips. Figure 1 shows an example of an SSD with a host interface, controller, and set of NAND dies. The controller receives commands from a host computer through a host interface. The controller then processes these commands in order to write/read data to/from memory.




Figure 1. Typical SSD architecture example with a host interface, an SoC controller, and NAND dies.

In this white paper, a team of designers from the Intel Non-Volatile Memory (NVM) Solutions Group presents an approach to model and optimize the performance of next-generation Intel® SSDs. This approach is based on analyzing SSD performance via models captured using the Intel® CoFluent™ Studio system-level modeling and simulation toolset.

### Market overview

The demand for SSDs continues to grow. Right now, that growth is driven by the performance requirements for storing data, including a demand for reduced latency and lower power consumption. This is true for both consumer mobile devices and enterprise applications.

According to a market report made public by Transparency Market Research,\* the global SSD market was valued at \$15.1 billion in 2014, and is expected to reach \$229.5 billion by 2022.<sup>1</sup> The market is expected to expand at a compound annual growth rate of 40.7% from 2015 to 2022.<sup>1</sup>

### Challenges in meeting the projected SSD market growth

A significant design challenge today is specifying the architecture of an SSD, while optimizing its performance. Many potential bottlenecks must be taken into account. For example, the host interface (such as PCI Express\* or Serial ATA\*) can affect the bandwidth of the complete system. Combining a specific host interface with a protocol such as NVMe\* can also create overhead for elements such as command buffers in the controller, depending on the size of the buffers.

In addition, NAND dies have different read and write access times, and the same NAND die can be targeted by multiple commands. This can lead to contention, which can affect the time it takes to process read and write commands. SSD designers must also consider that multiple channels can be used to exchange data between the controller and the NAND dies, and that multiple dies can be accessed using a single channel.

The controller also executes firmware code to perform tasks such as block mapping, garbage collection, error detection and correction, read and write caching, power management, and other background tasks. It is not easy to predict how many cores should be allocated to run these crucial tasks, which can also generate additional transactions between the controller and the NAND dies. This is another architectural aspect that can affect the time it takes to process commands.

Finally, validating and optimizing the behavior and performance of an SSD is difficult, because designers must avoid the discovery of potential issues late in the design cycle. Such issues often appear only in the development phase in which designers begin considering complex dynamic workloads (sequences of interleaved write and read commands). Indeed, the late discovery of issues can significantly impact development phases. In turn, this can increase time-to-market and development costs for the final product.

### Objective

Our objective is to explore the architecture and predict the behavior and performance of SSDs for different workloads earlier in the development cycle, before starting hardware and software development. We want to identify and predict with a high level of confidence what the final architecture of the SSDs should be. Also, our predictions should cover specifics, such as buffer sizes, number of cores, channels, NAND dies, and so on.

### Solution overview

In this optimization study, we used the Intel CoFluent Studio toolset to create an executable specification of the architecture of an SSD device. With Intel CoFluent Studio, the user specifies the device architecture using simple block diagrams, C/C++ code, and timing annotations. The tool then automatically generates SystemC\* code that corresponds to the architecture specification.

The executable system specification is used to validate complex dynamic workloads earlier. This helps to reduce the risk of finding operational issues later in the design cycle. The insights gained from the executable specification are also used by the team to revise and improve the quality of the architecture specification document.

Figure 2 gives an overview of the flow we used to explore the architecture and optimize the performance of next-generation Intel SSDs.

For this study, we modeled the architecture of the SSD device as a hierarchical composition of many modules. We captured and validated these modules in separate Intel CoFluent models. Each module is itself composed of several hierarchical blocks.

We also used Intel CoFluent Studio to capture a simple testbench around this architectural model, in order to validate the behavior and performance of the SSD device.

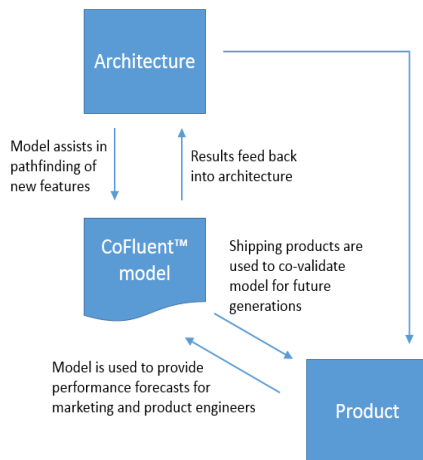


Figure 2. High level approach to optimize architecture and the performance of next-generation Intel® Solid State Drives based on Intel® CoFluent™ Studio models.

Figure 3 gives an overview of the top-level graphical representation of the CoFluent SSD device model and its testbench (the host and host interface modules).

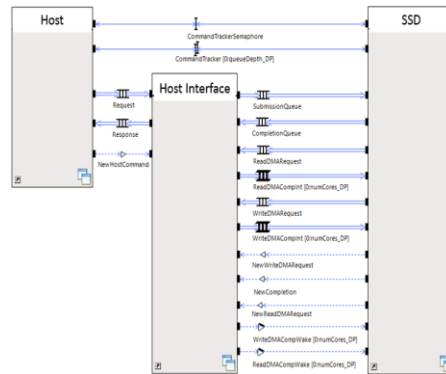


Figure 3. Top-level graphical representation of the SSD device model and its testbench, captured using Intel® CoFluent™ Studio.

The testbench stimulated the SSD device model with multiple workloads: sequences of uniquely identifiable, time-stamped write and read commands. Using this testbench, it was possible to generate random sequences of write and read commands, or set a ratio between sequences of write and read commands.

The commands were processed by the SSD device. Ultimately, the commands were consumed by the testbench, which generated latency and throughput measurements for each command. Those measurements then enabled the team to verify whether the SSD device met the performance requirements for complex workloads.

The dynamic nature of the model also allowed the team to capture transient effects and study sources of nonuniform performance. This made it possible for us to study and optimize future SSD architectures.

### Usage and results

The Intel CoFluent SSD behavioral model enabled us to make informed, data-driven decisions that increased SSD performance and reduced cost. We show here the results of several studies using the Intel CoFluent model, in order to highlight some of the architecture questions that the model helped us answer.

#### Study 1 – Sizing the transfer buffer

The SSD controller contains an SRAM-based transfer buffer for sending data to and from the NAND storage devices and the host. Creating a system with enough SRAM is crucial so that SRAM doesn't become a bottleneck for performance. At the same time, SRAM is expensive. A key business goal is to use only the amount of SRAM necessary to achieve the performance goals of the product, and use no more than that.

The Intel CoFluent model captures the behavior of traffic through the system, and accurately models the allocation and freeing of this memory. Because of that, it is possible to model the amount of SRAM needed for a given drive. This kind of modeling allows us to more clearly see how the amount of SRAM affects drive capacity, backend performance, cost and other factors.

In this study, the workload generated by the testbench is a sequence of write commands — see Figure 4. The figure shows normalized results of several simulations of a particular SSD model with different hypothetical transfer buffer sizes. These results show the point at which a greater transfer buffer yields diminishing returns on performance.




Figure 4 – Effect of SRAM transfer buffer size in the controller, on throughput in a sequential write workload.

Note that it would be difficult or impossible to obtain these types of results based on analytical models or spreadsheets. The reason is that this study relies on looking at traffic through the system and how commands line up in an actual workload for dynamically releasing and allocating the transfer buffer.

The performance analysis in this first study helped the team choose the appropriate amount of SRAM transfer buffer for this particular system design. The result was that the team was able to allocate an adequate transfer buffer to the system, and avoid costly over-provisioning.

**Study 2 – Optimizing the compute capability**

For some workloads, bottlenecks are the host interface data-transfer speed, NAND back-end performance, and transfer buffer size. In other workloads, the SoC capability to process host commands becomes the main performance bottleneck.

In the SoC case, developers generally want to know how much performance gain they would see with specific firmware improvements or different CPU architectures. They also want to know how such improvements would interact with other potential bottlenecks.

In this study, we considered a large-queue-depth random-read workload. Figure 5 shows how, in this study, random read input/output operations per second (IOPS) are improved as a function of overall system compute capability. The normalized results in the figure are for two potential SoC device architectures of a future SSD. This highlights the fact that




Figure 5 – Impact of the controller’s command processing capabilities (which in turn are related to CPU compute power and firmware efficiency) on random read input/output operations per second (IOPS).

increasing the capability of other resources can affect the compute requirements of the system.

Again, note that it would have been difficult to obtain these results using an analytical model or a spreadsheet because we wanted to analyze how secondary bottlenecks interfere with the primary CPU bottleneck. Using an Intel CoFluent model for this type of study allowed us to analyze the dynamic evolution of the CPU compute power and firmware efficiency, as other bottlenecks became more prominent.

The results of this study provided valuable feedback to the firmware development team and SoC architecture team. It also improved the communication between the two teams. This is important because, in a typical development environment, firmware and architecture development teams must work together to design and build a system that meets the SSD performance goals of future products.

Our study here showed that using Intel CoFluent Studio allows us to test different controller architectures earlier in the design process. In turn, that allows teams to predict with a high level of confidence how much compute capability is needed to reach targets, and what trade-offs are possible and might be most practical.

### Study 3 – Optimizing the multi-bit programming algorithm

In this study, we examined how best to optimize the multi-bit programming algorithm.

Current state-of-the-art SSDs are able to store multiple bits of data (3 bits in triple-level cells “TLC” NAND) in a single flash memory cell. The extra bits that need to be programmed in each cell introduce new complexity into the algorithms in the controller that manages how this data is stored. For example, all 3 bits could be written to a cell at once; or the bits could instead be written to a cell one at a time, alternating between different cells.

Each programming scheme uses system resources differently and can have very different performance characteristics relative to one another. Furthermore these performance characteristics can change from one SSD generation to another, based on changes in NAND timings and how these interact with other system bottlenecks (which also vary from generation to generation).


In Figure 6, we show the normalized results from simulations of the SSD model for some trial programming algorithms for a TLC-based SSD.

In this study, the workload generated by the testbench is a sequence of write commands. By modeling the behavior of the controller’s programming algorithm, we can obtain a highly accurate estimate of the system performance in each case.

Without an executable specification, it would have been difficult to predict how different programming algorithms (which have nontrivial system dependencies) end up affecting performance. In particular, the granularity of our Intel CoFluent model enabled us to understand how transitions between different phases of a given programming algorithm affect resource utilization and performance.

The Intel CoFluent model also allows for quick experimentation of novel algorithms. Both of these use cases

reduce the need for costly firmware development, since only the best algorithm needs to be realized.



*Figure 6 – Performance of various multi-bit programming algorithms designed to optimize device performance. By simulating the high-level behavior of different algorithms, we gain enough information to choose the best algorithm for our design, and help reduce costly firmware development.*

### Conclusion

Intel engineers were looking for effective approaches to optimizing the performance of next-generation Intel SSDs. In this study, they explored the effectiveness of modeling and simulating SSD devices using Intel CoFluent Studio.

Using an innovative model-based approach with Intel CoFluent Studio, the team created an executable specification of the architecture of the SSDs. The executable specification was then used to validate functional behavior, and then predict the performance of the SSDs for different workloads. The executable Intel CoFluent model was highly effective, and enabled the team to identify and solve architectural issues before beginning the hardware/software development phase.

### Key results

- Using the Intel® CoFluent™ Studio model, the team simulated and identified the amount of SRAM required to implement a transfer buffer, while optimizing the overall performance and cost of a new SSD architecture.
- The team explored different models and simulations of controller architectures early in the design process. This allowed the team to more accurately predict how much compute capability would be required to reach targets, and what trade-offs would be both possible and most effective.
- The team compared different schemes to program multi-bit flash memory cells in the controller. The team then identified new algorithms that would optimize the performance of next-generation TLC-based SSDs.
- The Intel CoFluent model helped to formalize communication between the firmware development and SoC architecture teams, who must typically work together to make informed, data-driven design decisions.





For more information on Intel® CoFluent™ Studio, visit  
[www.cofluent.intel.com](http://www.cofluent.intel.com)

<sup>1</sup> Transparency Market Research, "Solid State Drive Market – Global Industry Analysis, Size, Share, Growth, Trends and Forecast 2015-2022", 2015-09-18, <http://www.transparencymarketresearch.com/solid-state-drive-market.html>

Information in this document is provided in connection with Intel® products. No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

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.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting [www.intel.com/design/literature.htm](http://www.intel.com/design/literature.htm).

Intel, the Intel logo, and Intel CoFluent are trademarks of Intel Corporation in the U.S. and/or other countries.

Copyright © 2017 Intel Corporation. All rights reserved.

\*Other names and brands may be claimed as the property of others.