

Non-Volatile Memory and Java

[Mario Wolczko](#)

A series of short articles about the impact of non-volatile memory (NVM) on the Java platform.

Part 1: Introducing NVM

Non-volatile Random Access Memory (NVRAM) has [arrived into the computing mainstream](#). This development is likely to be highly disruptive: it will change the economics of the memory hierarchy by providing a new, intermediate level between DRAM (Dynamic Random Access Memory) and flash storage, but fully exploiting the new technology will require widespread changes in how we architect and write software. Despite this, there is surprisingly little awareness on the part of programmers (and their management) of the technology and its likely impact, and relatively little activity in academia (compared to the magnitude of the paradigm shift) in developing techniques and tools which programmers will need to respond to the change.

In this series I will discuss the possible impact of NVRAM on the Java ecosystem. Java is the most widely used programming language: there are millions of Java developers and billions of lines of Java code in daily use.

I'll start by describing the released Intel hardware and its software interface. Next I'll go into the consequences of this combination on software, entailing both opportunities and challenges. After that I'll get into the specifics of how Java may be impacted, the choices facing the platform in addressing the technology and how those choices have been resolved in proposals to date. Finally, I'll describe some ideas I've been pursuing at Oracle Labs and explain the underlying rationale.

I'll try to keep this accessible to typical Java programmers (and their

managers) by explaining concepts and terms as necessary.

Introduction

For as long as most people can remember, computers have lost the content of main memory when power is removed. (Before semiconductor memory there was ferrite core storage, which was also non-volatile, although I don't know if anyone exploited this property, or, if so, how.) However, that is about to change. Intel has just released Non-Volatile DIMMs (NV-DIMM), under the brand [Optane DC Persistent Memory](#). The underlying circuit technology is known by the brand [3D XPoint](#). SSDs using 3D XPoint chips have been shipping since mid-2017, and now computers are available with these devices attached as NVRAM. Many other companies are also working towards similar goals and we can expect there to be a variety of competing technologies and products appearing in the coming years. However, Intel's entry into the field has given it new legitimacy and importance.

What we know about Intel's hardware

Many details about the technology developed by Intel and [Micron](#) (Intel's technology partner) are yet to emerge. As of this writing, the details below are based on Intel statements, press articles and a [UCSD paper](#); I have not yet obtained access to a working system.

In addition to byte-addressability and non-volatility:

1. *Memory density* is significantly higher than DRAM. The first Persistent Memory Module (PMM) sizes are 128, 256 and 512GB using [128Gb parts](#); in comparison, the largest DRAM DIMM is currently [256GB using 16Gb parts](#).
2. *Cost/bit* is significantly lower than DRAM, although significantly higher than flash. [Early pricing/bit](#) appears to come in around 20% the price of DRAM (up to 40%, for the high-capacity modules, although this will likely drop over time). DRAM prices are also currently falling from a

[recent high](#) so the ratio will be in a flux for a while.

3. *Read and write latencies* are higher than DRAM, but within an order of magnitude and hence much lower (by several orders of magnitude) than flash. I haven't seen any figures from Intel, but a [study from UCSD](#) confirms this: they measure random read latency at around 300ns (compared to around 80ns for DRAM) and sequential read latency of around 170ns. Write latencies are closer: 94ns for Optane vs 86ns for DRAM (more on this later).
4. *Read bandwidth* peaks at about one-third that of DRAM, while write bandwidth is about a sixth that of DRAM (from the [UCSD paper](#)).
5. *Endurance* is expected to be much higher than flash, although how close to DRAM we don't yet know — Intel has not released any official endurance figures. [Early marketing](#) claimed Optane would have 1000 times the speed (which I take to mean reciprocal media latency) and endurance of flash, at 10x DRAM density. That would imply an endurance on the order of 10^8 write cycles per cell (compared to 10^5 for flash, and $>10^{15}$ for DRAM). A quick calculation confirms that this would suffice: the UCSD paper reports a single-DIMM peak write bandwidth of 1.5GB/s for 256B blocks (the Optane native block size). This translates to about 6M writes/second, or up to about 1P writes in the warrantied 5 year life, or 2M writes/block in a 128GB DIMM, assuming perfect wear leveling.

Because of the higher latency (and perhaps lesser endurance, to be confirmed) than DRAM, systems using Optane will still contain DRAM. New systems capable of supporting Optane have also been released. In Intel's new Cascade Lake architecture each CPU socket has up to [six memory modules attached via two memory controllers](#), for a maximum capacity of 3.75TB per socket (6x0.5TB PMMs and 6x128GB DIMMs). Each PMM contains an internal indirection table which is used to map CPU physical addresses to internal addresses; this allows wear leveling and bad block management to be done internally to each module (from the [UCSD paper](#), §2.1.1) and should improve practical endurance. A module or module

partition can be configured in one of two modes: in *Memory* mode the DRAM attached to the same controller acts as a cache, while in *AppDirect* mode the persistent memory is accessed by the CPU directly. Memory mode does not support persistence — there is no way to force writebacks and therefore no way to know when updates become durable — and so I will not consider this mode any further.

To bring a technology like this to market requires a wide variety of developments involving many aspects of system architecture, some of which I shall describe in later parts, and which have taken many years to complete. It's not "just" a matter of inventing a non-volatile circuit technology and building pin-compatible DIMMs. This gives Intel an advantage, in that there aren't many industry players in control of all the necessary technologies. There's a downside too, however: Intel's PMMs won't work in other manufacturer's systems; they won't even work in older Intel hardware. An enterprise customer wanting multiple sources of non-volatile memory technology before making mission-critical changes to its IT architecture will have to wait for competitors to catch up. To help minimize divergence, an industry consortium (SNIA — the Storage Networking Industry Association) has been authoring a roadmap under its [Solid State Storage Initiative](#).

There are many other companies developing NVRAM technologies, so I am confident that multiple sources will emerge, eventually, although the timescale is difficult to predict. Some, e.g., [Nantero](#), claim that their technology will be as fast as DRAM. If speed and endurance eventually match or improve upon DRAM, then it becomes conceivable to replace *all* DRAM with NVRAM, but this does not appear imminent. However, if DRAM is eventually replaced by NVM it would be nice if we could avoid having to go through *two* rounds of disruption in the way we build software, and plan for that eventuality.

In the [next part](#) I will discuss software issues in using Optane PMMs.