

Kontakt Patch Loads: NVMe vs SATA SSD (Windows)

tl;dr

I compared a **1TB Samsung 960 PRO NVMe** with a **2TB Samsung 850 EVO SATA SSD** and measured patch load times of a multi consisting of all section patches from Cinematic Studio Strings.

The goal was to answer the question: will the added performance of NVMe help in my project load times where most of the time seems to be spent waiting for Kontakt to load patches?

Consequently, the test is fairly limited in that only the process of patch loading was measured. (DFD streaming was not a real pain point for me on even modest SATA SSDs so I didn't bother to compare it here.)

Observations:

- The **initial patch load times** (that delay until the GUI becomes responsive) is entirely **CPU bound**. There was no difference between the two drives.
- Once the GUI was ready, **samples loaded** from disk into memory **substantially faster with NVMe**.
- If you run a purged template and rely on streaming, there probably isn't much difference in practice between these drives.
- If you prefer to fully load your samples, NVMe will do it much faster (**2x faster** in my comparison).
 - But I haven't found this necessary as even modest SATA SSDs have been able to keep up with streaming even from purged patches for me. (Note that my ASIO buffer is 512 samples. Smaller buffers might benefit from the lower latency of NVMe. A test for another day.)
- Kontakt is not as efficient as it might be in loading samples, including some bizarre pathological filesystem behaviour on Windows at least.
 - But it's also not quite as bad as I'd previously thought (and claimed on forums), so, mea culpa.

Read on for the exciting nerdy details.

Test Procedure

I wanted to do a deep-dive on a small, manageable test case, so I created a simple Kontakt multi consisting of:

- CSS 1st Violins (426.14 MB)
- CSS 2nd Violins (442.07 MB)
- CSS Violas (405.31 MB)
- CSS Cellos (450.72 MB)
- CSS Basses (326.76 MB)

This is a total sample memory consumption (as reported by Kontakt) of **2051 MB**.

That's with a **preload buffer size of 24KB**, which is my personal setting. (The default preload buffer will result in about 1.8x more memory used.)

Before taking any measurements, I would:

- Reset multi within Kontakt
- Use [RamMap](#) to flush the filesystem cache (Empty | Empty Standby List)
- Wait 10 seconds to ensure the system has quiesced

During each measurement, Windows' Performance Monitor would log the following metrics:

- CPU % Idle Time for all logical processors
 - Idle is a convenient proxy for real system utilization even though it does require reading the graphs inverted (an acquired skill).
- Disk Read Bytes/sec
- Disk Reads/sec
- Disk % Idle Time
- Average Disk sec/Read (i.e. read latency)

Drive Benchmarks

In order to meaningfully benchmark the drives, I needed to know how Kontakt behaved in practice with respect to:

1. Disk queue depth
2. Read block size

After collecting metrics during repeated Kontakt multi loads from both disks, I determined that:

1. Using PerfMon to examine the Current Disk Queue Length metric while Kontakt did its thing, the disk queue length rarely exceeded 1 during the sample loading. Occasionally it saw 2, but **by and large the queue depth remained at 1**.


2. Dividing Disk Read bytes/sec by Disk Reads/sec during active sample loading, both drives measured a mean block size of around 118KB. Consequently **benchmarks with a 128KB block size will be the most relevant.**

Given those observations, I ran the [Atto Disk Benchmark](#) tool against both disks with the following settings:


- Direct I/O: on
 - Otherwise the filesystem cache will warm up during the write stage inflating the read benchmark
- I/O Comparison: off
- Overlapped I/O: off
 - These ensure disk queue depth remains at 1 as was observed with Kontakt
- Transfer size: 32KB to 512KB
 - Wasn't much point in measuring too far outside that range as Kontakt primarily deals with ~128KB blocks while samples pulled in during initial load.

PerfMon collected data during Atto's execution so I could determine read latency and read operations per second (which Atto does not report itself).


Here are the results:



Read Operations Per Second



Read Latency (Microseconds)



Tabular summary:

Block size (KB)	32	64	128	256	512
850 EVO SATA MB/s	349.3	407.9	432.5	478.7	506.7
960 PRO NVMe MB/s	726.8	1151.9	1688.2	2080.6	2268.1
850 EVO SATA Reads/sec	11263	5938	3354	1838	981
960 PRO NVMe Reads/sec	24700	19636	13668	8003	4047
850 EVO SATA Latency usec	75.23	143.01	265.91	502.22	971.14
960 PRO NVMe Latency usec	27.28	37.37	57.2	101.7	204.85

Respectable enough performance difference between these two drives (unsurprisingly).

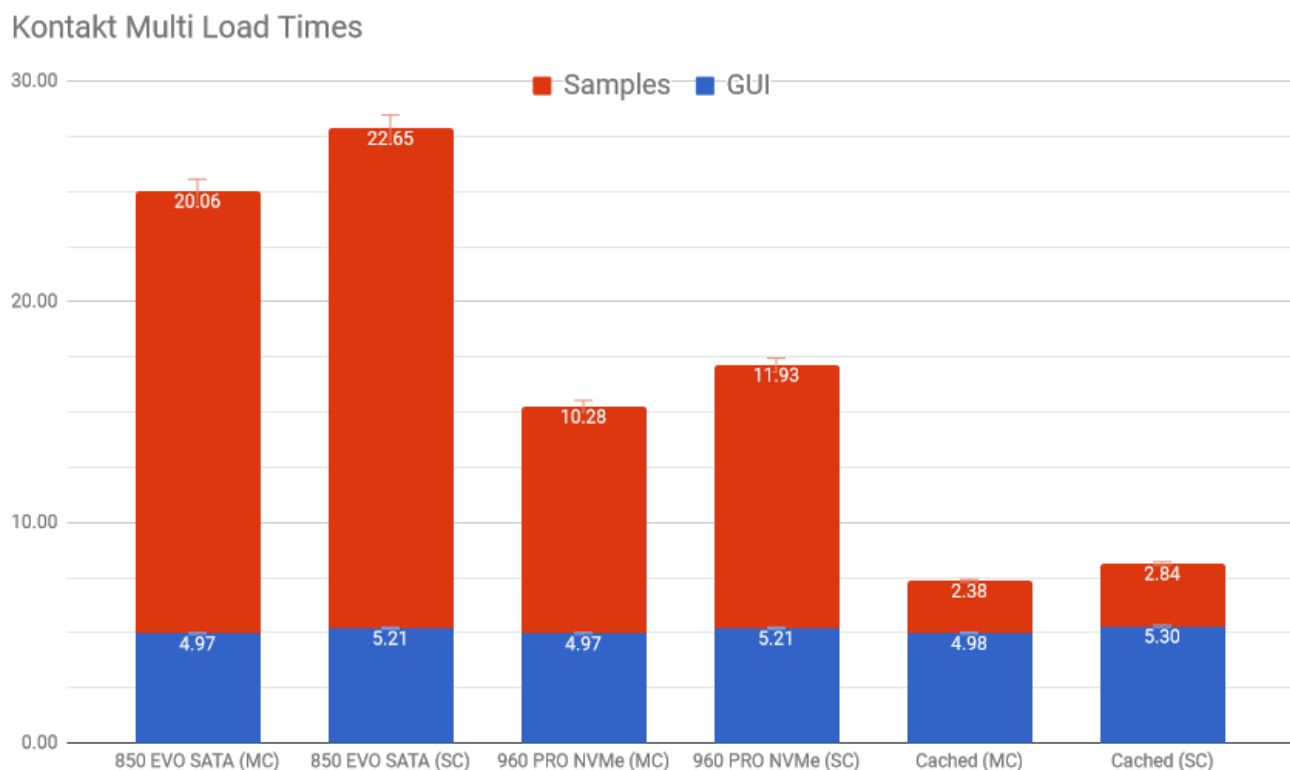
Kontakt Multi Load Times

Using a basic stopwatch test, I timed loading the multi 3 times with each drive. I measured the time for the GUI to become ready, and then again the additional time to finish loading all patches from disk.

I also tested loading without first evicting the filesystem cache, to effectively measure Kontakt's best possible load times with disk performance was completely taken out of the equation. (These results are labeled below as "Cached.")

Finally, I wanted to see what effect single core (SC) vs multicore (MC) had on load times, because prior testing led me to believe patch loading was almost entirely dependent on single core performance. In the SC case, I used Process Explorer to pin Kontakt to processor 0 (out of my 12 logical processors), while in the MC case it's assigned to all 12 (as is the default).

Here's the money shot:



My interpretation:

- The NVMe had absolutely zero benefit for the initial patch loading times that plague my overall project loads. My lowly 850 EVO is plenty fast enough to avoid bottlenecking this process, which is unambiguously CPU bound.
- More than one logical processor helps GUI loads by around 4%, and patch loading by around 16%. This supported my earlier observations that patch loading was almost entirely CPU bound, but frankly the extra cores helped more than I thought it would have. 16% isn't great, but it's not nothing. So I need to revise my position a bit on this one.

- Sample loading is drastically improved with NVMe. But don't we all run our templates purged anyway?


The graph indicates the measurement error bars, which were surprisingly low considering the human involvement in triggering the stopwatch. Across the GUI load times, the average relative standard deviation was only 0.43%. For sample loading, it was 1.51%.

Incidentally, these numbers are adjusted down to account for my own reaction time, which averaged to 252ms as measured by [Human Benchmark](#).


System Metrics

Let's get right into it:

850 EVO SATA: Throughput vs Disk / CPU Utilization %



960 PRO NVMe: Throughput vs Disk / CPU Utilization %



In each graph, the first spike in blue represents that initial ~5 second patch load time. This corroborates the observation that this is entirely CPU bound: very little disk activity registers, meanwhile CPU utilization is a smidge above 100%. This CPU utilization is explained by a single threaded task running full tilt, plus some additional minor work running in another thread, which is why we saw the marginal benefit in the multicore measurement earlier.

The interesting business happens after the GUI is ready and the samples start being read off disk:

- With both disks, we are well below the benchmarked throughput at the equivalent block size (128KB)
- With the 850 EVO, Windows reports around 20% headroom on disk utilization, but meanwhile with the 960 PRO, Windows reports 100% utilization.
 - Disk utilization being 100 - % Disk Idle as captured by PerfMon
- CPU is working 20-25% harder in the 960 PRO case, which makes sense given we are feeding it samples to decompress more quickly.
- Notice that even with the 960 PRO here we aren't really pushing *that* much beyond what SATA III can theoretically deliver (around 700MB/s while SATA III's limit is 600MB/s).


Why are we not able to saturate the 850 EVO? There's headroom to spare on both disk and CPU.

Why are we pegging disk utilization on the 960 PRO when we are running under half its benchmarked throughput? Is disk utilization (via % Disk Idle) misleading?


The obvious answer is that it's related to IOPS. Except that these read transactions are already measured to be around 118KB block size, and the threshold drawn in yellow is the throughput when using 128KB blocks.

Still, let's look at read operations anyway:

850 EVO SATA: Read operations vs Disk / CPU Utilization %



960 PRO NVMe: Read operations vs Disk / CPU Utilization %



These results from an IOPS perspective don't seem to help clarify matters. With the 850 EVO, we see read operations approaching the benchmark threshold, while disk utilization is around 80%. In contrast, as with the throughput measurement, the 960 PRO is measuring half the benchmarked read operations/second at equivalent blocksize (128KB) even as disk utilization was reported as running full tilt at 100%.

(In case you're wondering, there were no heavy writes going on in the background. The disk utilization was entirely reads, which is why I'm using IOPS and read operations/second interchangeably.)

So, in both transactions and in throughput, why are we showing a disk utilization substantially higher than we would expect given the actual IOPS and throughput observed relative to the benchmark at equivalent IO block size?

Kontakt's File Access Pathology

Observing Kontakt through [Process Monitor](#) while performing this test only generates more questions.

Time of Day	Process Name	PID	TID	Operation	Path	Result	Detail
9:37:40.536471...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,825,058,399, Length: 120,928
9:37:40.536978...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.538011...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.538046...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,825,945,438, Length: 120, Priority: Normal
9:37:40.538059...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,825,945,552, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.538082...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,825,945,558, Length: 1,696
9:37:40.538197...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.538210...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,825,947,254, Length: 122,848
9:37:40.538260...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.538709...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.539451...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.539484...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.539507...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,826,755,829, Length: 120, Priority: Normal
9:37:40.539517...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,826,754,560, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.539742...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,826,755,949, Length: 1,708
9:37:40.539754...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.539807...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,826,757,687, Length: 122,304
9:37:40.540328...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.541019...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.541053...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.541076...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,827,585,528, Length: 120, Priority: Normal
9:37:40.541089...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,827,581,952, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.541222...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,827,585,648, Length: 1,760
9:37:40.541230...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,827,586,048, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.541345...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.541399...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,827,587,408, Length: 120,864
9:37:40.541759...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.542404...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.542452...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.542475...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,828,466,967, Length: 120, Priority: Normal
9:37:40.542488...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,828,466,688, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.542606...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,828,470,087, Length: 1,608
9:37:40.542616...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,828,470,784, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.542745...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.542804...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,828,471,695, Length: 121,728
9:37:40.543129...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.543834...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.543867...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.543889...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,829,252,686, Length: 120, Priority: Normal
9:37:40.543910...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,829,249,024, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.544026...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,829,252,806, Length: 1,776
9:37:40.544034...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,829,253,120, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.544150...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.544206...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,829,254,582, Length: 121,024
9:37:40.544534...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.545240...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.545273...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.545295...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,122,229, Length: 120, Priority: Normal
9:37:40.545307...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,121,472, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.545455...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,122,349, Length: 1,644
9:37:40.545467...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.545517...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,123,993, Length: 122,048
9:37:40.545525...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.546634...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.546667...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.546689...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,957,304, Length: 120, Priority: Normal
9:37:40.546798...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,957,056, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal
9:37:40.546931...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,957,424, Length: 1,672
9:37:40.546947...	Kontakt.exe	8032	14416	QueryStandardInformationFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	AllocationSize: 2,083,368,960, EndOffset: 2,083,368,660, NumberOfLinks: 1, DeletePending: False, Directory: False
9:37:40.547013...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,830,959,096, Length: 122,848
9:37:40.547528...	Kontakt.exe	8032	14416	CloseFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	
9:37:40.548258...	Kontakt.exe	8032	14416	CreateFile	E:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	REPARSE	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.548292...	Kontakt.exe	8032	14416	CreateFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Desired Access: Generic Read, Disposition: Open, Options: Sequential Access, Synchronous IO Non-Alert, Non-Directory File, Attributes: n/a, ShareMode: Read, Allocation
9:37:40.548314...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,831,750,263, Length: 120, Priority: Normal
9:37:40.548327...	Kontakt.exe	8032	14416	ReadFile	F:\Kontakt Libraries\Cinematic Studio Strings\Samples\CSStrings_007.nlx	SUCCESS	Offset: 1,831,747,584, Length: 4,096, I/O Flags: Non-cached, Paging I/O, Priority: Normal

Taking a closer look at that:

[illegible]

In this short 13ms snippet (representative of the entire process), we see Kontakt continuously opening, reading a bit, and closing the same file over and over again. It seems to close and reopen the file when it's done reading a sequential bit of data and wants to skip ahead into the file.

If I didn't know any better, I'd wonder if the developers at Native Instruments know about [SetFilePointer\(\)](#). Of course, I know how things can work in larger development shops: one team writes an abstraction layer used by another team that treats it as a black box, while the first team don't quite appreciate exactly how the second is using their code. Maybe something like that is at play here.

If it's the case that `CreateFile()` and `CloseFile()` generate even the tiniest amount of disk load, then the question posed at the end of the last section may have an answer.

In any case, surely this implementation is generating many needless context switches? (Granted, relative to context switches, disk access takes an eternity -- even with an NVMe -- so the value of optimization probably isn't very high from that angle.)

Thread Processing

During the initial load (where the UI is blocked), there is an ephemeral thread named "Progress" spun up. This thread is destroyed once the initial patch load is completed and the UI becomes unblocked. Presumably this thread is responsible for precursor tasks such processing group setup, mapping individual samples, compiling scripts to platform-native bytecode, etc.

This thread consumes 100% of a logical processor until its work is complete, and then the thread is destroyed. This single thread is responsible for the most disruptive (and annoying) aspects of project loading.

From this point, the filesystem access occurs from a long-lived worker thread. The main thread consumes some CPU as well, however without Kontakt's symbol tables it's not possible to understand what this thread is doing short of walking through a disassembly.

Making some design inferences, it seems unlikely the main thread would be responsible for decompressing the samples being loaded in. Since the only other busy thread is the one accessing the filesystem, it's not unreasonable to assume this single thread is both reading data from disk and processing this data.

If that's the case, it could explain why neither disk tested nor CPU was saturated during sample loads: this is effectively equivalent to the [bandwidth-delay product](#) problem, where increased latency results in a reduction of overall throughput due to processing being stalled by round-trip times. Moreover, Kontakt's file I/O is synchronous (which explains why the disk queue depth largely stayed at 1).

Consider the following pseudocode:

```
while not samplefile.eof():
    block = samplefile.read(BLOCKSIZE)
    decompress(block)
```

With this code, when the disk is busy the CPU is idle, and when the CPU is busy, the disk is idle. The lower the disk access latency while reading a block, the busier we will be able to keep the CPU. (And similarly, the faster the CPU, the busier we will be able to keep the disk.) This may be the primary reason why the NVMe drive is doing so much better -- read latency -- because recall the throughput measured during the Kontakt multi load on the 960 PRO wasn't that much beyond what SATA III can deliver.

A better design -- and a fairly trivial one at that -- would be to move the decompression (and other sample processing) to a separate thread, using the producer/consumer pattern such that the disk access thread produces to a queue that is consumed by the data processing thread.

With this minor tweak (and a suitably sized intermediate queue), we should be able to easily saturate either disk or CPU (whichever resource is truly the bottleneck).

Conclusion

NVMe drives don't appear to help the worst part of Kontakt patch loading (the part in which the UI is blocked). However, they do significantly help in overall sample load times.

But if you run a purged template, you'll probably find the load times to be roughly similar to a reasonably spec'd SSD such as the 850 EVO.

DFD streaming could use a closer look in terms of the performance benefits at large scale, across a project consisting of 100+ tracks, and at aggressively lower preload buffer sizes.

Appendix

Kontakt Multi Load Times

Multicore Run	SATA GUI	SATA samples	NVMe GUI	NVMe samples	Cached GUI	Cached Samples
1	5.17	20.27	5.22	10.71	5.19	2.55
2	5.23	20.35	5.23	10.53	5.25	2.68
3	5.26	20.31	5.23	10.36	5.26	2.66
Mean	5.22	20.31	5.23	10.53	5.23	2.63
Stddev	0.046	0.040	0.006	0.175	0.038	0.070
RSD	0.88%	0.20%	0.11%	1.66%	0.72%	2.66%
Lag Adjusted	4.97	20.06	4.97	10.28	4.98	2.38

Pinned Run	SATA GUI	SATA samples	NVMe GUI	NVMe samples	Cached GUI	Cached Samples
1	5.45	23.16	5.45	12.24	5.56	3.05

2	5.5	22.73	5.5	11.96	5.62	3.06
3	5.43	22.83	5.43	12.34	5.49	3.16
Mean	5.46	22.91	5.46	12.18	5.56	3.09
Stddev	0.036	0.225	0.036	0.197	0.065	0.061
Lag Adjusted	5.21	22.65	5.21	11.93	5.30	2.84
% Delta from MC	4.83%	12.95%	4.69%	16.02%	6.49%	19.34%

System Specs

- Software
 - Kontakt standalone 5.73
 - Windows 10 x64 16299.248
- Hardware
 - i7 8700K @ 4.7GHz
 - Gigabyte Z370 AORUS Gaming 7
 - 64GB RAM
 - Nvidia 1080 Ti
 - RME Babyface Pro (ASIO buffer at 512 samples)
 - 1TB Samsung 960 PRO NVMe M.2
 - 2TB Samsung 850 EVO SATA SSD (non M.2)
 - (Plus other drives not involved in test)

Caveat

I'm a Linux systems engineer, feebly attempting to apply my understanding of operating systems to Windows, and therefore bound to overreach in certain assumptions about OS behavior. All data-driven cluebatting quite graciously accepted.