

Sonic Lineup - Bug #1909

Extremely slow on Windows when doing more than one thing at once

2019-09-06 10:01 AM - Chris Cannam

Status:	Closed	Start date:	2019-09-06
Priority:	Urgent	Due date:	
Assignee:		% Done:	0%
Category:		Estimated time:	0.00 hour
Target version:			
Description			
On Windows:			
<ul style="list-style-type: none">- Load a number of substantial audio files- While alignment is in progress, switch to Melodic Spectrogram view- Observe that the window updates extremely slowly, and the application enters "not responding" state for long periods of time			
I think the application is still working and will eventually complete the job, but I haven't the patience to find out.			
There is no such problem on Linux, and I don't think there was such a bad problem on Windows either until quite recently.			

History

#1 - 2019-09-06 10:03 AM - Chris Cannam

Possible candidates:

- Some lock contention problem in c++11-mutex branch of vamp-plugin-sdk
- Excess locking on file access
- Lack of optimisation for Byld logic

?

#2 - 2019-09-11 06:09 PM - Chris Cannam

The origin of this appears to be that this machine is the only one I'm testing on on which enough free memory is reported for SpectrogramLayer to make a decision to use the whole-model cache. Disabling that and using either the default 8-column cache or the finer 4-column one makes things much faster (the 8-column cache especially).

#3 - 2019-09-11 08:35 PM - Chris Cannam

What is particularly problematic is "cascading" caches - generating one cache based on summaries of another cache which also needs to be filled. The whole-model cache and the normal cache were set up in this way (normal cache generated from whole-model cache). This causes massive memory pressure. Switching the normal cache to be generated directly from the FFT model essentially solves the problem.

I'm still uncertain whether the whole-model cache is useful in any circumstance.

#4 - 2019-09-12 02:00 PM - Chris Cannam

- Status changed from New to Closed

Also reduced cache sizes using the new `verticallyFixed` mode in `SpectrogramLayer` and `FFTModel::setMaximumFrequency`. With this change I think it is reasonable to leave the whole-model cache enabled.