

Data, Nonotify, or Silent?

Egon Willighagen 

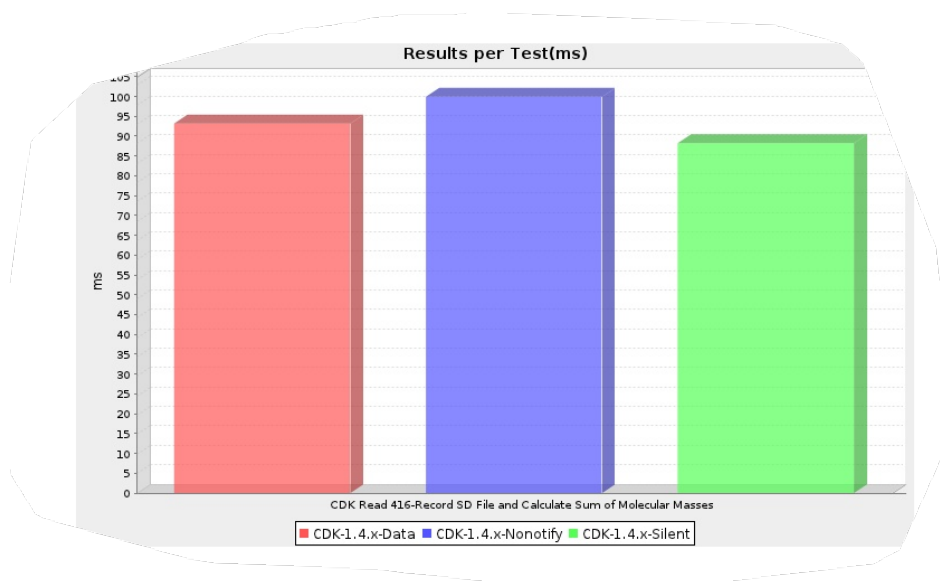
Published July 13, 2011

Citation

Willighagen, E. (2011). Data, Nonotify, or Silent?. *Chem-bla-ics*. <https://doi.org/10.59350/zq2m3-dxp07>

Keywords

Cdk



Copyright

Copyright © Egon Willighagen 2011. Distributed under the terms of the [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

chem-bla-ics

I cannot find the bug report just now, but the [CDK](#) has an open problem with change even notification, where the nonotify classes still caused change event to be sent around.

This was because the nonotify classes extended in a wrong way the data classes. So, I worked today on copying the data class implementations into a new implementation, not extending the data classes, while removing the listener code: the *silent* module. I'm not entirely done yet, but close enough to blog about it. While checking things, I ran the [cheminfbench](#) code on it, with these results:

So, removal of the notification listening improves the performance, when reading a 416 entry SD file. I think the difference will be more significant for other tasks, like ring finding.

But, but...?!?! Yeah, this is a rather weird plot indeed... the blue bar should also be lower than the red one! And it used to be too... :(Bad regression... hard to unit test too :(

OK, back to some final clean up.

Update: the clean up is done, and I have now run the fingerprint benchmark from [cheminfbench](#) using the new module and nonotify. In a situation when change events are much more used (as is with fingerprint calculation), we see that nonotify still improves speed, and that the new silent module shows about the same speed up. We also see that the 1.4.x classes are a bit slower than one classes of some 20 months ago. That probably reflects [bug 2992921](#) that was recently fixed. The full bar plot:

Red and blue are CDK 1.2.x (as the plot legend says), green and yellow the same for CDK 1.3.x (and both clearly faster than the 1.2 series, and purple and light blue the same for CDK 1.4.0. The last bar is the new silent module, a tid bit slower than nonotify.

Update 2: OK, one last update. The performance difference can actually be larger than this. The below screen shot shows the effect of the silent module (blue, yellow) on SMILES generation (without and with lower case formalism, red and green respectively):

If you did not get it yet, if you bring your system to production level, do not use the default implementation, **unless** you really need to change notifications.

References