
Subject: [FIXED] genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Wed, 28 Oct 2015 17:05:01 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear all,

I am performing some tests with genfit2, as some people tried and found memory leak problems.

The trunk that I am currently using is rev-28695, and OS Fedora19.

I simulated interactively 2000 events, using the standard `/macro/run/*complete.C`.

In the reco- macro, the number of iterations is set to 5 for both, barrel and fwd spectrometer. I made sure to run `genfit2, _not_ genfit`.

Up to now, no problem is found with the simulation `ppbar to Jpsi pi pi`.

Be sure that memory and CPU consumption increase when increasing the number of iterations, to have make converge the fit. Anyway, 5 is still a reasonable number, and memory leak problems should not occur.

I am currently running 10 000 events, interactively. I will do the same using the standard `genfit1`, and make the comparison.

I will give you an update by tomorrow.

I see that in the committed revision of genfit2, the class `WireMeasurement()` is used. An update is `WireMeasurementNew()`, that was never tried. However, in a chat with authors today I understood that they refer to the class `AbsMeasurement()` without passing through `WireMeaurement*` classes. This has to be tried, before coming to conclusions.

The advantage is that it will give you is correct Jacobian entries related to energy loss, which as we all know is fairly important for slow particles.

I am doing my home-works...

Best regards, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Thu, 29 Oct 2015 07:28:18 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello Stefano,

I run this night 10000 events, interactively. The standard simulation to check the new reconstruction tools is `ppbar to Jpsi pi pi`, as in the `/macro/run/*_complete.C`.

My jobs exit without any problem, and they are successfully. I attach here the log file of the reco- macro.

I modified the macro reco- and pid- in a way to run genfit2 and related classes (.e., I make use of `PndRecoKalmanTask2`, I set `recoKalman->SetNumIterations(5)`, and `corr->SetFlagCut(kFALSE)` in the pid- macro). Is it correct, or am I supposed to change something else? I updated the trunk to the last revision. I checked everything using my laptop (---->OS Fedora19). I will try with Suse, just to be sure, on my desktop. But I do not see any

trouble up to now.

Do you have any additional suggestions?

I would not run interactively more than 10000 events, just because when the sim- output is big, my experience is that quite often the params- output file is not properly closed, then useless. But this is not a problem in reconstruction...

cheers, Elisabetta

File Attachments

1) [logrec10000](#), downloaded 389 times

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Thu, 29 Oct 2015 21:19:24 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear Elisabetta,

this afternoon I tried with 10k events and I had no crash due to full memory.

Trying to go back in the code, since in the past the problem was seen not only in J/psipipi but also in llbar and also by Lu, I noticed that Tobias fixed some memory leaks in the Stt/Fts mapper. Maybe that was the reason of the memory filling, since the mappers are intensively used by the Stt/Fts Reco Hits.

In any case, since it seems WireMeasurementNew should be the best new class to use, I would suggest to try with this, in order to be coherent with the genfit2 developments.

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Sat, 31 Oct 2015 17:06:09 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello Stefano,

to my understanding the class WireMeasurement(), that right now is used in genfit2/measurement, is more general than the new class WireMeasurementNew(). For the STT detector I consider that more appropriate; however, in Belle II they use directly the class AbsMeasurement(). This is of course an attempt that I can still try. In any case, as all classes WireMeasurement* do not allocate dynamical resources, whatever problem of memory leak cannot be due to the latter classes. Indeed, I did not find problems of memory leak in my preliminary tests. I will further investigate.

Clearly the CPU time needed to run genfit2 depends also on how many number of iteration one sets. Let's remember that genfit(1) had a logic bug, and the numlter was forced to be =1. I tried with genfit2 to set numlter up to 10, and it worked. A good compromise between fit performance and CPU time consumption can be to run numlter = 3. It should be enough (my guess). I'll do my homework.

Please, schedule me for a genfit2 update talk during the next collaboration meeting. I will perform some new tests and comparisons in the last trunk revision, and show in Vienna. If you

have better suggestions, please feel free to say and I'll work on.

cheers, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Sun, 01 Nov 2015 12:45:16 GMT

[View Forum Message](#) <> [Reply to Message](#)

Maybe it is just matter to try and compare performances. Maybe with genfit2 the backpropagation using geane) is not needed anymore (with genfit w/o backpropagation the performances were a bit poorer).

In my tests, with only 1 iteration, I also noticed an increase in computing time compared to genfit1. It would be also nice to compare, with our new qa, the performances with different number of iterations.

I will schedule you a talk.

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Fri, 06 Nov 2015 09:10:37 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear Stefano, dear all,

I performed additional tests with genfit2, as there are some complains from the community. I summarize here the results of my preliminary tests, performed using a /development/branch where only genfit2 is in, `_not_genfit(1)`.

I have set the number of iterations to 1, 3, 5, 10, and check the memory consumption. The results are the following:

- a) file used: ppbar to Ds- Ds+ (Lu Cao analysis);
- b) number of generated events: 2000;

Using genfit2, when I run the macro/run/reco_complete.C, I get:

```
with num lter = 1 -----> 1045 MB
with num lter = 3 -----> 1070 MB
with num lter = 5 -----> 1094 MB
with num lter = 10 -----> 1109 MB
```

In the tests performed by Lu in the new trunk, the memory consumption when using genfit2 is 6559 MB. Clearly these numbers are not compatible each others.

Being the fact that genfit2 standalone code can run on my pc 500 000 events without any problem of memory leak, my conclusion is that - in case of problems in pandaroot- we have only 2 possibilities:

- 1) the interface GenfitTools is bugged;
or/and
- 2) the switch old-genfit vs genfit2 is bugged

My tests demonstrated that, in a branch where only genfit2 is in, the interface GenfitTools works fine. It is my opinion that in the new trunk the switch GF1 vs GF2 is corrupted. This brings me to ask you: what is the sense to have in the same trunk or release the old-genfit _and_ genfit2? genfit1 has clear problems, that genfit2 corrects. Once again I say, we should exclude genfit1 from the next trunks and releases, because this can cause troubles.

My plan is now to create a new /development/branch, where the release mar15 is adapted to use genfit2 only (it mean: I will eliminate any track on genfit1, and I will configure again the other packages where the genfit classes are called somehow), and I will show at the collaboration meeting the results of this last test.

A discussion on how to proceed at the coll meeting is needed. I would be in favor to forget about the old non-maintained genfit, and use only the new updated and maintained genfit2.

Best regards, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Fri, 06 Nov 2015 09:51:30 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear Elisabetta,

if there is a memory leak the number of iterations do not change the usage, only the number of events -> increasing the number of events the memory usage increases (which should not appear if there are no leaks). This test could be done.

Genfit1 and genfit2 are 2 completely separate packages and they do not interact at all, then your hypothesis cannot work. In particular, why in your pc you have only genfit2 and not both? Did you check the memory consumption using the trunk? If you have different results with standalone and with pandaroot version then we need to understand why. But I have not understood what you call "standalone" version. Are you using exactly the same genfit2 release which is in pandaroot?

A last general comment: in general the standard cores for grid have inside an average of 2GB memory. This is the basic requirements for farm machines.

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Fri, 06 Nov 2015 10:03:06 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello Stefano,

if people are saying that they find memory leak problems, I check every possibility. Then, it was not a bad idea to see what happens with different numIter, as I remember that this was source of troubles with genfit1 (otherwise, why is it forced to 1, in the old genfit)?

I do have in my development/branch only genfit2, as you know. This is how I performed every test since one year, simply because I prefer not to mix up things.

If every connection with genfit1 is excluded, no memory leak problems are seen at all. So, it is not /genfit2 or /GenfitTools generating troubles in pandaroot.

In the new trunk, up to 10 000 events, I did not see problems testing the channel Jpsi pi pi. But other people got it, with different decay channels.

In the new trunk, I do not know what was done. For sure, there are both, genfit and genfit2. Not my decision.

As I said before, what I can do is to set up the new release mar15 with only genfit2, and do my tests again. I do not want to mix up things, otherwise every debugging is difficult.

Remember that other 8 packages were slightly changed, when I had worked with genfit2. I can say what I did in /development/genfit2; I am not aware of development in other packages.

My best, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Fri, 06 Nov 2015 10:16:15 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hi,

Elisabetta Prencipe (2) wrote on Fri, 06 November 2015 11:03: Hello Stefano,

if people are saying that they find memory leak problems, I check every possibility. Then, it was not a bad idea to see what happens with different numlter, as I remember that this was source of troubles with genfit1 (otherwise, why is it forced to 1, in the old genfit)?

Since increasing the number of iterations there was no improvement in resolution, only a small reduction in efficiency. For the TDRs we used three iterations, but after a while we reduced such number to 1 since it was only increasing the computing time.

Quote:

I do have in my development/branch only genfit2, as you know. This is how I performed every test since one year, simply because I prefer not to mix up things.

If every connection with genfit1 is excluded, no memory leak problems are seen at all. So, it is not /genfit2 or /GenfitTools generating troubles in pandaroot.

But did you try the trunk? This is the starting point. If you do not try you cannot say that gf1 and gf2 mixing it is the reason of the problem.

Quote:

In the new trunk, up to 10 000 events, I did not see problems testing the channel Jpsi pi pi. But other people got it, with different decay channels.

Did you check memory usage there? If you have 6GB used memory and you have 8GB RAM memory you do not see problems, but if you have 2GB ram then you have problems.

Quote:

In the new trunk, I do not know what was done. For sure, there are both, genfit and genfit2. Not my decision.

In the new trunk it was put the code your wrote, cleaned a bit. Did you check it?

Quote:

As I said before, what I can do is to set up the new release mar15 with only genfit2, and do my tests again. I do not want to mix up things, otherwise every debugging is difficult.

I suggest to start from the standard trunk, if not we will not be able to find which are the problems. genfit1 and 2 are written do that they do not mix in the code, if they are mixing this has to be demonstrated.

Quote:

Remember that other 8 packages were slightly changed, when I had worked with genfit2. I can say what I did in /developmentnet/genfit2; I am not aware of development in other packages.

What do you mean exactly? The code is a port of what you had in your development branch, and you told me which genfit2 version to use. Are you sure you are using exactly the same genfit2 version?

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Fri, 06 Nov 2015 10:23:50 GMT

[View Forum Message](#) <> [Reply to Message](#)

Stefano, as stated in the previous email:

Lu in the new trunk got: 6559 MB; with the same dec file, same p mom, same "everything", I got in *my* development/branch: ~ 1000 MB.

"My* development branch has got the same genfit2 revision; but it has not genfit1 in (my decision). I just deleted it.

As I said, there is a reason: I do not want to mix up things, for better understanding.

Here I see a problem.

Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Fri, 06 Nov 2015 10:25:18 GMT

[View Forum Message](#) <> [Reply to Message](#)

But your dev version starts from the exact trunk version of Lu, and from the same genfit2 version?

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Fri, 06 Nov 2015 12:22:04 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello Stefano,

the genfit2 revision used in my development/branch, and in the new trunk, it is the same. My point is that several packages (hyp, lmd,) made use of genfit classes directly; and some other made use of the Pnd-classes from GenfitTools (<----this is the correct way).

Now, in the new trunk I see:

/genfit /genfit2 .Correct!

/GenfitTools -----> with some classes *2 . Correct!

what about the other pandaroot packages as mentioned above? were they ported as I did in my development/branch? If not, here is the problem. If yes, some more deep debugging is needed. But the problem looks not in genfit2.

As first step, I will do what I just proposed in my previous email. If I get positive response (and I do believe so), I will update my development/branch in the repository, and ask you to check again. If you and me get the same output, clearly something was messed up in switching genfit/genfit2 in the new trunk. If you and me do not get the same results, we will come back to this tread again. I still say that it is better to have only genfit2, once everybody make sure that there are no troubles, of course.

Please, give me a couple of days and we can discuss again, later at the pandaroot meeting next week.

My best, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Stefano Spataro](#) on Sat, 07 Nov 2015 08:03:20 GMT

[View Forum Message](#) <> [Reply to Message](#)

In all those packages (lmd, hyp) the dependency from genfit packages was removed since it was wrong by design

Subject: Re: genfit2 update - bugs(?)

Posted by [Tobias Stockmanns](#) on Fri, 13 Nov 2015 17:34:20 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear all,

I was looking into memory leaks in the digitization stage and the reconstruction stage of PandaRoot and I checked for Genfit 1 and 2.

I found many in various parts of the code.

The most severe was still a problem in the PndSttMapCrator (and the corresponding FTS part). Others were in the GenFitTools both for Genfit1 and Genfit2 which often were only a copy of each other.

I could reduce the size of memory leaks for the reco stage from 93 MByte for 100 events down to 3 MByte for Genfit1 and to about 6 MB for Genfit2.

I am not finished with this task, but it would be nice if someone could test it if the situation has improved.

Cheers,

Tobias

Subject: Re: genfit2 update - bugs(?)

Posted by [Elisabetta Prencipe \(2\)](#) on Mon, 16 Nov 2015 07:53:05 GMT

[View Forum Message](#) <> [Reply to Message](#)

Hello Tobias,

thanks for this further check, and your fixes. Could you please say which trunk version shall we test?

Thanks again, Elisabetta

Subject: Re: genfit2 update - bugs(?)

Posted by [Tobias Stockmanns](#) on Mon, 16 Nov 2015 08:56:09 GMT

[View Forum Message](#) <> [Reply to Message](#)

Dear Elisabetta,

the trunk version is 28737. You can have a look at <https://subversion.gsi.de/trac/fairroot/browser/pandaroot> to see what the latest changes are and who has uploaded what. You can even search for changes from a specific author.

Cheers,

Tobias
