
Subject: update stable branch...

Posted by [Johan Messchendorp](#) on Fri, 22 Jan 2010 00:08:52 GMT

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

Dear all,

I merged the present trunk code (revision 7522) with the stable branch. Furthermore, I fixed the fairbase externals to this revision as well in the stable branch. The genfit revision is set to 98 (same as for trunk). Note that this is a starting point for the next production release which I would like to bring online very soon. At the moment, not all the QA macros are passing the tests, which is the minimum requirement before we can announce the next production release. I would, therefore, like to urge all the developers to ensure that this is OK in the upcoming days. Please tell me which revision number for your piece of software I should use for the stable branch if it is different from revision 7522.

Best wishes,

Johan

Subject: Re: update stable branch...

Posted by [Johan Messchendorp](#) on Sat, 23 Jan 2010 12:18:08 GMT

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

Dear all,

Below I summarise the remaining problems which I noticed by looking at the dashboard and talking to some people. For the next production release, we should try to fix these items and incorporate them in the stable branch. Some of the problems only show up on 32-bit machines or somewhat older compilers. Nevertheless, our ambition is to keep compatibility with as many platforms and compilers as much as possible as part of the quality assurance and service to the collaboration. So, please take a regular look at the dashboard. In this respect, it might be good to have various virtual machines available and accessible for all the developers. Something we could discuss on the next meeting.

Tutorial macros: there seem to be crashes related to the PndEventReader and the Jul09 and Jan10 fairsoft ROOT releases. Klaus told me he will take a look at it.

On 32-bit machines, the fast simulation QA macro crashes with a floating point exception, for example <http://fairroot.gsi.de/CDash/testDetails.php?test=35414&build=15381> . I believe the problem is related to the track smearing function, which - sometimes - smears values to very large values causing a floating point error. Klaus, could you take a look at this as well?

qa_dch_macro2: on 32-bit machines this macro crashes as well with a floating point exception: <http://fairroot.gsi.de/CDash/testDetails.php?test=35403&build=15381> . There are also warnings of the type "caution: wrong drift time". Ola W., Arek or Garzia, would one of you still have some time to take a look at this? That would be really great! [UPDATE: noticed that the start momentum for the fitting is sometimes zero, which causes a floating point error in

evaluating Q/P. Added a check whether the start momentum is zero, before proceeding with the fit (line 252 in PndDchPreFitterTR.cxx , rev7537). I am not sure whether these situation should occur at all?!?]

qa_lhe_macro : on all machines it gives a segmentation fault and points to a problem in the PndMvdStripClusterTask. My guess is that this is related to the restructuring of the MVD code. Ralf, is there any hope on the short time for this, or shall we use for the stable release the "old" MVD structure?

qa_mvd_macro2 and qa_mvd_macroana: similar to the previous crash (I guess?).

qa_stt_macro3: on some machines it returns the statement "negative muon ptot resolution wrong" and "negative muon kalman ptot resolution wrong":

http://fairroot.gsi.de/CDash/testSummary.php?project=2&name=qa_stt_macro3&date=2010-01-23 . Lia/Susanna, do you have any idea what goes wrong here? Unfortunately, it doesn't happen on all machines. I attached as pdf the corresponding histograms which came out of the macro and which were used to perform the check.

reco_complete_stt/tpc: also related to the MVD (see above)... sorry, Ralf

Best wishes,

Johan.

File Attachments

1) [his_stt.pdf](#), downloaded 352 times

Subject: Re: update stable branch...

Posted by [Lia Lavezzi](#) on Mon, 25 Jan 2010 10:51:54 GMT

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

Dear Johan,

for the STT QA macro n.3: the histograms don't look so bad, I think the problem is that the resolution exceeds slightly the limits I chose to consider the test as passed (perhaps for the low statistics).

I just uploaded to svn a QAmacro_stt_3.C file with more debugging information to be able to understand. If it is due really to the low statistics I will increase a little the number of events in the simulation and see what happens. If the reason is something else I will try to understand. On my machine I don't see the failure

Cheers,
Lia.

Subject: Re: update stable branch...

Posted by [Johan Messchendorp](#) on Mon, 25 Jan 2010 17:56:51 GMT

Hi Lia,

Thanks for looking into it... Let me printout the information I now got from running the macro on a 32-bit machine:

Quote:

TEST 2: helix mu- resolution 0.0248863 (limits [0.025, 0.035])

TEST 6: genfit mu- resolution 0.0176549 (limits [0.018, 0.028])

Test Failed

Not Ok

negative muon ptot resolution wrong

negative muon kalman ptot resolution wrong

Looks like that the resolution is even better than the limits you take. If you think that this is OK, I would suggest to put the lower range to smaller values. What do you think?

Greetings,

Johan.

Subject: Re: update stable branch...

Posted by [Lia Lavezzi](#) on Tue, 26 Jan 2010 11:16:05 GMT

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

Hi Johan,

thank you for giving me the print outs

Among the two limits in the resolution test, the upper limit is the really important one. The lower limit was just put there to check if too big statistical fluctuations were present (which might be related to some bugs), but it is also possible that while improving the code and fixing the bugs, the resolution gets better and so it turns out to be better than expected...

In this case I think it could be just a matter of statistics. The limits were tuned on 100 events, but I have seen that the number of events in macro 1 (which are then reconstructed in macro 2 and histogrammed in macro 3) has been turned from 100 to 50 some time ago by Mohammad.

Just to understand what to do, what was the reason for this? Are 100 events (with 6 track each, 3 mu+ and 3 mu-) too many for a qa macro? If so, no problem, I will change the limits for the tests to test only 50 events (I will enlarge them a little to account for bigger fluctuations from my values). If on the contrary it is possible to go back to 100 events, I would like to do it, just to lower the statistical fluctuations and see if everything runs fine. Mohammad, what do you think about this?

In any case, it seems to me that there is nothing we have to worry about for the code itself, it is just the macro that has to be fixed (in one of the two ways), since from the numbers: 0.0248863 is just like 0.025 and 0.0176549 is like 0.018

Please let me know about the number of events to use and I will fix the macro accordingly.

Cheers,
Lia.

Subject: Re: update stable branch...

Posted by [Mohammad Al-Turany](#) on Tue, 26 Jan 2010 11:26:55 GMT

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

Hi Lia,

I change the number because of some old machines which were too slow and I was getting some timeout, sorry for this. Any way if you see that 100 is better, then please change it, and maybe you can set the time in the cmake to a larger value (not to get a timeout).

best regards

Mohammad

Subject: Re: update stable branch...

Posted by [Lia Lavezzi](#) on Tue, 26 Jan 2010 13:06:15 GMT

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

Thank you Mohammad,
yes, I would try with 100 events.

Sorry, I didn't take too much care to the timeout value in the CMakeLists.txt : now I set it to 200 sec for macro 2 (which is the slowest) and to 50 for macro 3 (which only fills histo, so it is fast). This should work fine.

I also made some changes to macro 3 to fix some errors in efficiency calculation and to better be able to change the limits (if/when needed).

I hope this will fix the problems!

Cheers,
Lia.

Subject: Re: update stable branch...

Posted by [Johan Messchendorp](#) on Tue, 26 Jan 2010 20:23:35 GMT

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

Hi Lia,

Just watch the Dashboard, and see whether some problems still occur on some machines with your macro. The best is to use enough statistics, so I also support to go to 100 events and increase the time out limit...

Johan.

Subject: Re: update stable branch...

Posted by [StefanoSpataro](#) on Fri, 29 Jan 2010 09:33:58 GMT

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

Yesterday I have checked the qa macros, just to understand what is going good and what is wrong.

Here the list of wrong things:

a) DCH and forward EMC are overlapping. I don't know who is in the correct place, for sure I know that both DCH and FEMC geometries are obsolete. Maybe we could comment out the shashlick.

b) inside dipole geometry there are three overlaps

c) in my pc, sim_complete_tpc.C is stuck at event #6, and sometime it takes 300 seconds. Maybe the time there should be enlarged

d) Why in the reco_complete_tpc.C there are all the STT reconstruction tasks However they do not affect at all the reconstruction. I will correct it.

e) In several machines reco_complete_tpc.C fails, with the following error

Toggle Spoiler

Empty Track. No hits/digis inside

-l- PndMvdPixelClusterTask: 25 Mvd Clusters and 25 Hits calculated.

Program received signal SIGFPE, Arithmetic exception.

[Switching to Thread -1208527168 (LWP 17466)]

0x03fb484c in PndGemTrackFinderOnHits::RemoveCloneTracks (this=0xa9a6b00, nofRecoTracks=1) at /home/spataro/dart/trunk/gem/PndGemTrackFinderOnHits.cxx:433

433 if (meanPhi[itr]/nofTS[itr] < 5. && iterTS.trackPhi > 355.) { iterTS.trackPhi -= 360.; }

Current language: auto; currently c++

(gdb) bt

#0 0x03fb484c in PndGemTrackFinderOnHits::RemoveCloneTracks (this=0xa9a6b00, nofRecoTracks=1)

at /home/spataro/dart/trunk/gem/PndGemTrackFinderOnHits.cxx:433

#1 0x03fb3a9c in PndGemTrackFinderOnHits::DoFind (this=0xa9a6b00, hitArray=0xb32a370, trackArray=0xb47bb18)

at /home/spataro/dart/trunk/gem/PndGemTrackFinderOnHits.cxx:288

#2 0x03fbc00f in PndGemFindTracks::Exec (this=0xa9a3bf0, opt=0x72558e8 "") at /home/spataro/dart/trunk/gem/PndGemFindTracks.cxx:145

I think this is due to a not protected division by zero. But why that number is zero...

f) fsim is crashing due to a "wrong" smearing of the vertex for neutral particles. I have sent a mail to Klaus.

Now words to the "experts"...
