

**Ben Foster** <foster@ucar.edu>

---

**TIE-GCM crash**

14 messages

---

**Cnossen, Ingrid** <inos@bas.ac.uk>  
To: Ben Foster <foster@ucar.edu>

Mon, Aug 3, 2015 at 4:22 AM

Hi Ben,

I'm doing a long run with the TIE-GCM at double resolution for the year 2010. It was going fine initially, but now the model has crashed on day 109. The error message suggests it is to do with the hnmf2 subroutine in diags.F. I think it's not finding a peak in the electron density profile, and because hmf2, nmf2, and fof2 are then all set to zero it may be trying to divide by zero. I've copied the error below. I seem to remember that you have encountered this before (I think we talked about it once), but I'm not sure what the best way is to get around it. Do you have any advice?

Thanks,

Ingrid

---

Step 7890 of 377280 mtime=109 11 30 0 secs/step (sys) = 2.37

GPI+IMF run: istep= 7890 GPI f107= 76.087 f107a= 78.333 power from bzimf,swvel= 33.657

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

>>> hnmf2: lat= 1 could not find kx -- setting hmf2=nmf2=fof2=0

forrtl: severe (174): SIGSEGV, segmentation fault occurred

Image	PC	Routine	Line	Source
libirc.so	00002AAAAE9DC961	Unknown		Unknown Unknown
libirc.so	00002AAAAE9DB0B7	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F9E4	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F7F6	Unknown		Unknown Unknown

tiegcm_trunk	00000000007EBD64	Unknown	Unknown	Unknown
tiegcm_trunk	00000000007F353D	Unknown	Unknown	Unknown
libpthread.so.0	00002AAAAEC3AB10	Unknown	Unknown	Unknown
tiegcm_trunk	00000000006F1ACA	Unknown	Unknown	Unknown
tiegcm_trunk	00000000006C8239	Unknown	Unknown	Unknown
tiegcm_trunk	0000000000435EC6	Unknown	Unknown	Unknown
tiegcm_trunk	000000000075EF91	Unknown	Unknown	Unknown

etc.

-----

Dr. Ingrid Cnossen

Natural Environment Research Council (NERC) Post-doctoral Fellow

British Antarctic Survey

Cambridge, UK

e-mail: [inos@bas.ac.uk](mailto:inos@bas.ac.uk)

phone: (+44)(0)1223 221647

---

This message (and any attachments) is for the recipient only. NERC is subject to the Freedom of Information Act 2000 and the contents of this email and any reply you make may be disclosed by NERC unless it is exempt from release under the Act. Any material supplied to NERC may be stored in an electronic records management system.

---

---

**Ben Foster** <[foster@ucar.edu](mailto:foster@ucar.edu)>  
To: "wbwang@ucar.edu" <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>

Sat, Aug 8, 2015 at 4:39 PM

Sent from my iPhone

Begin forwarded message:

**From:** "Cnossen, Ingrid" <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>  
**Date:** August 3, 2015 at 4:22:23 AM MDT  
**To:** Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>  
**Subject:** TIE-GCM crash

[Quoted text hidden]

---

**Ben Foster** <foster@ucar.edu>

Mon, Aug 10, 2015 at 1:34 PM

To: "Cnossen, Ingrid" <inos@bas.ac.uk>, Liying Qian <lqian@ucar.edu>, Wenbin Wang <wbwang@ucar.edu>

Ingrid,

Sorry for the delay responding - I've just returned from PTO.  
I suspect you are right that it's not finding the peak Ne.  
I am copying Liying and Wenbin so they can comment on that.

I presume you are using the diags.F, i.e., HNMF2 and/or NMF2, FOF2  
are in SECFLDS, which causes mkdiag\_HNMF2 to be called  
from sub elden (elden.F). It would be interesting to try continuing  
your run without these diags that search for the peak Ne, and see  
if it still seg-faults.

--Ben

[Quoted text hidden]

--

Ben Foster  
National Center for Atmospheric Research (NCAR)  
High Altitude Observatory (HAO)  
[303-497-1595](tel:303-497-1595)

---

**Cnossen, Ingrid** <inos@bas.ac.uk>

Tue, Aug 11, 2015 at 6:40 AM

To: Ben Foster <foster@ucar.edu>

Cc: Liying Qian <lqian@ucar.edu>, Wenbin Wang <wbwang@ucar.edu>

Hi Ben (cc to Liying and Wenbin),

No problem, I had already assumed you were probably on PTO.

Yes, I had HMF2 and NMF2 in SECFLDS. I've tried restarting the run without these fields selected, but it crashed again. So it looks like the seg-fault wasn't actually caused by mkdiag\_HNMF2. I have noticed that the crash occurred at a slightly different timestep this time, but no helpful clues otherwise; it just stops complaining about not finding kx. A section of the error message is copied below. Any idea what the problem could be?

Cheers,

Ingrid

---

Step 680 of 370080 mtime=109 11 20 0 secs/step (sys) = 1.31

GPI+IMF run: istep= 680 GPI f107= 76.082 f107a= 78.333 power from bzimf,swvel= 36.114

forrtl: severe (174): SIGSEGV, segmentation fault occurred

Image	PC	Routine	Line	Source
libirc.so	00002AAAAE9DC961	Unknown		Unknown Unknown
libirc.so	00002AAAAE9DB0B7	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F9E4	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F7F6	Unknown		Unknown Unknown
tiegcm_trunk	00000000007EBD64	Unknown		Unknown Unknown
tiegcm_trunk	00000000007F353D	Unknown		Unknown Unknown
libpthread.so.0	00002AAAAEC3AB10	Unknown		Unknown Unknown
tiegcm_trunk	00000000006F1ACA	Unknown		Unknown Unknown
tiegcm_trunk	00000000006C8239	Unknown		Unknown Unknown
tiegcm_trunk	0000000000435EC6	Unknown		Unknown Unknown
tiegcm_trunk	000000000075EF91	Unknown		Unknown Unknown
tiegcm_trunk	0000000000426D6E	Unknown		Unknown Unknown
libc.so.6	00002AAAAF6B9994	Unknown		Unknown Unknown
tiegcm_trunk	0000000000426C89	Unknown		Unknown Unknown

forrtl: error (78): process killed (SIGTERM)

Image	PC	Routine	Line	Source
libirc.so	00002AAAAE9DC961	Unknown		Unknown Unknown
libirc.so	00002AAAAE9DB0B7	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F9E4	Unknown		Unknown Unknown
tiegcm_trunk	000000000084F7F6	Unknown		Unknown Unknown
tiegcm_trunk	00000000007EBD64	Unknown		Unknown Unknown

etc.

**From:** Ben Foster [mailto:[foster@ucar.edu](mailto:foster@ucar.edu)]

**Sent:** 10 August 2015 20:34

**To:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>; Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>; Wenbin Wang

<[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>

**Subject:** Re: TIE-GCM crash

[Quoted text hidden]

[Quoted text hidden]

---

**Wenbin Wang** <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>

Tue, Aug 11, 2015 at 7:10 AM

To: "Cnossen, Ingrid" <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

Cc: Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>, Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

Hi, Ingrid

You can try to reduce the time step by half to see if the model runs fine. After you pass the particular date, you can set the time step back to the original one.

Good luck

Wenbin

Sent from my iPhone

[Quoted text hidden]

---

**Wenbin Wang** <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>

Tue, Aug 11, 2015 at 7:20 AM

To: "Cnossen, Ingrid" <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

Cc: Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>, Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

Also, if you run the model with a time step less than 3 minutes, you may want to change the Shapiro factor in cons.F from 3e-2 to 3e-3, this will give you a better defined ionosphere at lower latitudes.

Wenbin

Sent from my iPhone

[Quoted text hidden]

---

**Cnossen, Ingrid** <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

Tue, Aug 11, 2015 at 7:50 AM

To: Wenbin Wang <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>

Cc: Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>, Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

Thanks Wenbin! I will try your suggestions.

-Ingrid

---

**From:** Wenbin Wang [mailto:[wbwang@ucar.edu](mailto:wbwang@ucar.edu)]

**Sent:** 11 August 2015 14:20

**To:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

**Cc:** Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>; Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

**Subject:** Re: TIE-GCM crash

[Quoted text hidden]

[Quoted text hidden]

---

**Ben Foster** <foster@ucar.edu>

Tue, Aug 11, 2015 at 9:34 AM

To: "Cnossen, Ingrid" <inos@bas.ac.uk>

Cc: Wenbin Wang <wbwang@ucar.edu>, Liying Qian <lqian@ucar.edu>

You might want to back up at least a day before the crash  
when restarting with a shorter timestep.

[Quoted text hidden]

---

**Cnossen, Ingrid** <inos@bas.ac.uk>

Fri, Aug 14, 2015 at 3:19 AM

To: Ben Foster <foster@ucar.edu>

Cc: Wenbin Wang <wbwang@ucar.edu>, Liying Qian <lqian@ucar.edu>

Just to let you know: the reduction in timestep worked; it has run successfully past the point where it  
previously crashed. Thanks again,

Ingrid

**From:** Ben Foster [mailto:[foster@ucar.edu](mailto:foster@ucar.edu)]

**Sent:** 11 August 2015 16:35

**To:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

**Cc:** Wenbin Wang <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>; Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

**Subject:** Re: TIE-GCM crash

[Quoted text hidden]

[Quoted text hidden]

---

**Ben Foster** <foster@ucar.edu>

Fri, Aug 14, 2015 at 9:37 AM

To: "Cnossen, Ingrid" <inos@bas.ac.uk>

Out of curiosity, which resolution are you running, and what were the  
former and current timesteps? Thanks,

[Quoted text hidden]

---

**Cnossen, Ingrid** <inos@bas.ac.uk>

Fri, Aug 14, 2015 at 9:44 AM

To: Ben Foster <foster@ucar.edu>

Hi Ben,

I'm running the 2.5 degree resolution. The timestep was initially 60 seconds, now 30. Unfortunately it has  
just crashed again (now on day 115 instead of 109)...

-Ingrid

**From:** Ben Foster [mailto:[foster@ucar.edu](mailto:foster@ucar.edu)]  
**Sent:** 14 August 2015 16:37  
**To:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>  
**Subject:** Re: TIE-GCM crash

[Quoted text hidden]

[Quoted text hidden]

---

**Ben Foster** <[foster@ucar.edu](mailto:foster@ucar.edu)> Fri, Aug 14, 2015 at 9:49 AM  
To: "Cnossen, Ingrid" <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>, Wenbin Wang <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>, Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>

Ok, reduce the timestep to 15 secs, backup and restart again.  
Hopefully, you can gradually go back up to 30 secs, which is  
the recommended step for 2.5 deg resolution.

[Quoted text hidden]

---

**Cnossen, Ingrid** <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)> Fri, Aug 14, 2015 at 9:53 AM  
To: Ben Foster <[foster@ucar.edu](mailto:foster@ucar.edu)>

Yes, I have just backed up and reduced the timestep again, like you suggested. I had also noticed when I reduced it earlier that 60 seconds is recommended for single resolution, but hadn't thought of looking at that when I first set up this run. I guess I was lucky it ran for as long as it did without crashing.

-Ingrid

**From:** Ben Foster [mailto:[foster@ucar.edu](mailto:foster@ucar.edu)]  
**Sent:** 14 August 2015 16:50  
**To:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>; Wenbin Wang <[wbwang@ucar.edu](mailto:wbwang@ucar.edu)>; Liying Qian <[lqian@ucar.edu](mailto:lqian@ucar.edu)>  
**Subject:** Re: TIE-GCM crash

[Quoted text hidden]

[Quoted text hidden]

---

**Ben Foster** <[foster@ucar.edu](mailto:foster@ucar.edu)> Tue, Aug 18, 2015 at 9:24 AM  
To: Art Richmond <[richmond@ucar.edu](mailto:richmond@ucar.edu)>

Art,

Here is the conversation with Ingrid, Wenbin, [et.al.](#):  
I will send you the relevant source code next.

--Ben

----- Forwarded message -----

**From:** Cnossen, Ingrid <[inos@bas.ac.uk](mailto:inos@bas.ac.uk)>

[Quoted text hidden]



[Quoted text hidden]