
Subject: Re: GPULIB troubleshoot

Posted by [David Grier](#) on Mon, 25 Jun 2012 11:32:39 GMT

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

On 6/25/12 4:27 AM, alx wrote:

```
> This message is intended to aficionados of GPULIB (in particular
> Michael Galloy...)
>
> I was successfully using a Tesla C1060 over several monthes on some
> workstation running Windows XP 64 bits and IDL8.0. Very successfully.
>
> Recently, I upgraded the same machine to Windows 7 64 bits and IDL8.2,
> with the following unsatisfying result:
>
> IDL> print,!version
> { x86_64 Win32 Windows Microsoft Windows 8.2 Apr 10 2012    64
> 64}
>
> IDL> gpuinit,0, /HARDWARE, ERROR=err, /VERBOSE
> Welcome to GPULib 1.4.4 (Revision: 2107)
> % GPUINIT: using hardware
> Graphics card: Tesla C1060, compute capability: 1.3, memory: 4040 MB
> available, 4095 MB total
> Checking GPU memory allocation...no errors
>
> IDL> a = gpuputarr(findgen(10), ERROR=err) & print,err
>      0
> IDL> b = gpuputarr(findgen(10), ERROR=err) & print,err
>      0
> IDL> c = gpumult(a,b,LHS=c, ERROR=err) & print,err
>      8
> The c GPU variable is created, but is in error
>
> IDL> print, gpugetarr(c,ERROR=err) & print,err
> 1.43493e-042 1.73472e-018  0.000000 7.34684e-040  0.000000
> 0.000000  0.000000 6.01853e-036  0.000000 2.52435e-029
>      0
>
> Does someone have an explanation ?
> alain.
>
```

I see two possible issues

1. It appears that c is declared to be the LHS (left-hand side) of the gpumult() call without first having being allocated.
Have you tried getting gpumult to allocate the necessary GPU variable?

```
IDL> c = gpumult(a, b, ERROR=err) & print, err
```

You also might allocate the variable explicitly:

```
IDL> c = gpuftrarr(10)
```

```
IDL> c = gpumult(a, b, LHS = c)
```

2. Did you rebuild GPULib when you upgraded your OS and IDL installation? I've forgotten to rebuild GPULib on my linux machines after upgrades and have had strange results.

TTFN,

David

Subject: Re: GPULIB troubleshoot

Posted by [lecacheux.alain](#) on Mon, 25 Jun 2012 11:49:29 GMT

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

On 25 juin, 13:32, David Grier <david.gr...@nyu.edu> wrote:

> On 6/25/12 4:27 AM, alx wrote:

>

>

>

>

>

>> This message is intended to aficionados of GPULIB (in particular
>> Michael Galloy...)

>

>> I was successfully using a Tesla C1060 over several monthes on some
>> workstation running Windows XP 64 bits and IDL8.0. Very successfully.

>

>> Recently, I upgraded the same machine to Windows 7 64 bits and IDL8.2,
>> with the following unsatisfying result:

>

>> IDL> print,!version

>> { x86_64 Win32 Windows Microsoft Windows 8.2 Apr 10 2012 64

>> 64}

>

>> IDL> gpunit,0, /HARDWARE, ERROR=err, /VERBOSE

>> Welcome to GPULib 1.4.4 (Revision: 2107)

>> % GPUINIT: using hardware

>> Graphics card: Tesla C1060, compute capability: 1.3, memory: 4040 MB

>> available, 4095 MB total

>> Checking GPU memory allocation...no errors

>

>> IDL> a = gpuputarr(findgen(10), ERROR=err) & print,err

```

>>      0
>> IDL> b = gpuputarr(findgen(10), ERROR=err) & print,err
>>      0
>> IDL> c = gpumult(a,b,LHS=c, ERROR=err) & print,err
>>      8
>> The c GPU variable is created, but is in error
>
>> IDL> print, gpugetarr(c,ERROR=err) & print,err
>> 1.43493e-042 1.73472e-018 0.000000 7.34684e-040 0.000000
>> 0.000000 0.000000 6.01853e-036 0.000000 2.52435e-029
>>      0
>
>> Does someone have an explanation ?
>> alain.
>
> I see two possible issues
>
> 1. It appears that c is declared to be the LHS (left-hand side) of
>    the gpumult() call without first having being allocated.
>    Have you tried getting gpumult to allocate the necessary GPU
>    variable?
>
>    IDL> c = gpumult(a, b, ERROR=err) & print, err
>
>    You also might allocate the variable explicitly:
>    IDL> c = gpuputarr(10)
>    IDL> c = gpumult(a, b, LHS = c)
>
> 2. Did you rebuild GPULib when you upgraded your OS and IDL
>    installation? I've forgotten to rebuild GPULib on
>    my linux machines after upgrades and have had strange results.
>
> TTFN,
>
> David
>
>

```

Regarding issue 1: preallocation statement was missed in the message, but was done in actual processing.

Regarding issue 2: GPUlib cannot be rebuilt (from my side), since the binary DLL (GPUlib.1.4.4) file is distributed without any sources.

I forgot to say that, when upgrading OS and IDL, I kept same NVIDIA graphics driver and CUDA version than before... So, I do not think that rebuilding the library would be mandatory.

Thanks anyway for the tips.

alain.

Subject: Re: GPULIB troubleshoot
Posted by [David Grier](#) on Mon, 25 Jun 2012 13:11:14 GMT
[View Forum Message](#) <> [Reply to Message](#)

On 6/25/12 7:49 AM, alx wrote:

> On 25 juin, 13:32, David Grier <david.gr...@nyu.edu> wrote:

>> On 6/25/12 4:27 AM, alx wrote:

>>

>>

>>

>>

>>

>>> This message is intended to aficionados of GPULIB (in particular
>>> Michael Galloy...)

>>

>>> I was successfully using a Tesla C1060 over several monthes on some
>>> workstation running Windows XP 64 bits and IDL8.0. Very successfully.

>>

>>> Recently, I upgraded the same machine to Windows 7 64 bits and IDL8.2,
>>> with the following unsatisfying result:

>>

>>> IDL> print,!version

>>> { x86_64 Win32 Windows Microsoft Windows 8.2 Apr 10 2012 64

>>> 64}

>>

>>> IDL> gpuinit,0, /HARDWARE, ERROR=err, /VERBOSE

>>> Welcome to GPULib 1.4.4 (Revision: 2107)

>>> % GPUINIT: using hardware

>>> Graphics card: Tesla C1060, compute capability: 1.3, memory: 4040 MB
>>> available, 4095 MB total

>>> Checking GPU memory allocation...no errors

>>

>>> IDL> a = gpuputarr(findgen(10), ERROR=err) & print,err

>>> 0

>>> IDL> b = gpuputarr(findgen(10), ERROR=err) & print,err

>>> 0

>>> IDL> c = gpumult(a,b,LHS=c, ERROR=err) & print,err

>>> 8

>>> The c GPU variable is created, but is in error

>>

>>> IDL> print, gpugetarr(c,ERROR=err) & print,err

>>> 1.43493e-042 1.73472e-018 0.000000 7.34684e-040 0.000000

>>> 0.000000 0.000000 6.01853e-036 0.000000 2.52435e-029

>>> 0

>>

>>> Does someone have an explanation ?

>>> alain.

>>

>> I see two possible issues

```

>>
>> 1. It appears that c is declared to be the LHS (left-hand side) of
>>    the gpumult() call without first having being allocated.
>>    Have you tried getting gpumult to allocate the necessary GPU
>>    variable?
>>
>>    IDL> c = gpumult(a, b, ERROR=err) & print, err
>>
>>    You also might allocate the variable explicitly:
>>    IDL> c = gpuftrarr(10)
>>    IDL> c = gpumult(a, b, LHS = c)
>>
>> 2. Did you rebuild GPULib when you upgraded your OS and IDL
>>    installation? I've forgotten to rebuild GPULib on
>>    my linux machines after upgrades and have had strange results.
>>
>> TTFN,
>>
>> David
>>
>>
>
> Regarding issue 1: preallocation statement was missed in the message,
> but was done in actual processing.
> Regarding issue 2: GPULib cannot be rebuilt (from my side), since the
> binary DLL (GPULib.1.4.4) file is distributed without any sources.
> I forgot to say that, when upgrading OS and IDL, I kept same NVIDIA
> graphics driver and CUDA version than before... So, I do not think
> that rebuilding the library would be mandatory.
> Thanks anyway for the tips.
> alain.
>

```

Point 2 may be the issue. On Unix-like systems, GPULib is linked against libidl.so, which changes from version to version. An IDL 8.0 build of GPULib does not work properly under IDL 8.2 on linux or MacOS for this reason -- I know this from first-hand experience. I'd bet a similar thing is going on in your case too. Without access to the source code, upgrading GPULib to the 64-bit Windows 7 version might possibly do the trick. If not, downgrading IDL to version 8.0 seems to be the most likely remaining option.

Good luck,

David

Subject: Re: GPULIB troubleshoot

Posted by [Michael Galloy](#) on Mon, 25 Jun 2012 16:38:00 GMT

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

On 6/25/12 7:11 AM, David Grier wrote:

> On 6/25/12 7:49 AM, alx wrote:

>> On 25 juin, 13:32, David Grier <david.gr...@nyu.edu> wrote:

>>> On 6/25/12 4:27 AM, alx wrote:

>>>

>>>

>>>

>>>

>>>

>>>> This message is intended to aficionados of GPULIB (in particular
>>>> Michael Galloy...)

>>>

>>>> I was successfully using a Tesla C1060 over several monthes on some
>>>> workstation running Windows XP 64 bits and IDL8.0. Very successfully.

>>>

>>>> Recently, I upgraded the same machine to Windows 7 64 bits and IDL8.2,
>>>> with the following unsatisfying result:

>>>

>>>> IDL> print,lversion

>>>> { x86_64 Win32 Windows Microsoft Windows 8.2 Apr 10 2012 64

>>>> 64}

>>>

>>>> IDL> gpunit,0, /HARDWARE, ERROR=err, /VERBOSE

>>>> Welcome to GPULib 1.4.4 (Revision: 2107)

>>>> % GPUINIT: using hardware

>>>> Graphics card: Tesla C1060, compute capability: 1.3, memory: 4040 MB

>>>> available, 4095 MB total

>>>> Checking GPU memory allocation...no errors

>>>

>>>> IDL> a = gpuputarr(findgen(10), ERROR=err) & print,err

>>>> 0

>>>> IDL> b = gpuputarr(findgen(10), ERROR=err) & print,err

>>>> 0

>>>> IDL> c = gpumult(a,b,LHS=c, ERROR=err) & print,err

>>>> 8

>>>> The c GPU variable is created, but is in error

>>>

>>>> IDL> print, gpugetarr(c,ERROR=err) & print,err

>>>> 1.43493e-042 1.73472e-018 0.000000 7.34684e-040 0.000000

>>>> 0.000000 0.000000 6.01853e-036 0.000000 2.52435e-029

>>>> 0

>>>

>>>> Does someone have an explanation ?

>>>> alain.

>>>

```

>>> I see two possible issues
>>>
>>> 1. It appears that c is declared to be the LHS (left-hand side) of
>>>    the gpumult() call without first having being allocated.
>>>    Have you tried getting gpumult to allocate the necessary GPU
>>>    variable?
>>>
>>>    IDL> c = gpumult(a, b, ERROR=err) & print, err
>>>
>>>    You also might allocate the variable explicitly:
>>>    IDL> c = gpufltarr(10)
>>>    IDL> c = gpumult(a, b, LHS = c)
>>>
>>> 2. Did you rebuild GPULib when you upgraded your OS and IDL
>>>    installation? I've forgotten to rebuild GPULib on
>>>    my linux machines after upgrades and have had strange results.
>>>
>>> TTFN,
>>>
>>> David
>>>
>>>
>>
>> Regarding issue 1: preallocation statement was missed in the message,
>> but was done in actual processing.
>> Regarding issue 2: GPULib cannot be rebuilt (from my side), since the
>> binary DLL (GPULib.1.4.4) file is distributed without any sources.
>> I forgot to say that, when upgrading OS and IDL, I kept same NVIDIA
>> graphics driver and CUDA version than before... So, I do not think
>> that rebuilding the library would be mandatory.
>> Thanks anyway for the tips.
>> alain.
>>
>
> Point 2 may be the issue. On Unix-like systems, GPULib is linked
> against libidl.so, which changes from version to version. An IDL 8.0
> build of GPULib does not work properly under IDL 8.2 on linux or
> MacOS for this reason -- I know this from first-hand experience.
> I'd bet a similar thing is going on in your case too. Without access
> to the source code, upgrading GPULib to the 64-bit Windows 7 version
> might possibly do the trick. If not, downgrading IDL to version 8.0
> seems to be the most likely remaining option.

```

Yes, that's what I would suggest in the short term. Building on Windows is technically possible for users of GPULib, but I would not suggest it for the faint of heart (it's much easier on Linux and Mac OS X). An updated version of GPULib should be released shortly.

Mike

--

Michael Galloy

www.michaelgalloy.com

Modern IDL: A Guide to IDL Programming (<http://modernidl.idldev.com>)

Research Mathematician

Tech-X Corporation
