Subject: Would you consider this a bug?
Posted by David Kastrup on Wed, 30 Sep 1998 07:00:00 GMT
View Forum Message <> Reply to Message

After doing

a=[3,4,5]
a[[2,1,0]] = a

a will be set to [3,4,3]

One has to write
a[[2,1,0]] = a + 0

in order to get the more correct seeming [5,4,3].

Would you think this a bug?

--
David Kastrup Phone: +49-234-700-5570
Email: dak@neuroinformatik.ruhr-uni-bochum.de Fax: +49-234-709-4209
Institut fï¿½r Neuroinformatik, Universitï¿½tsstr. 150, 44780 Bochum, Germany

Subject: Re: Would you consider this a bug?
Posted by David Kastrup on Thu, 01 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

Kevin Ivory <Kevin.Ivory@linmpi.mpg.de> writes:

> Leaving the bug topic - switching to algorithms
>
> David Kastrup wrote:
>> Unfortunately, translating something like
>>
>> a[[5,2,3,0,4,1]] = a
>> into the equivalent
>> a = a[[3,5,1,2,4,0]]
>> is not easy to do in the general case.
>
> Actually it is: you need a second call to sort. See example below.
>
>> some data that needed to be processed in sorted order, the results
>> needed to be rearranged in original order.
>>
>> Something like
>> s = sort(a)

>> a = a[s]
>> process(a)
> ; a[s] = a   ; I still think this a dangerous thing to do.
>   r = sort(s)
>   a = a[r]

First, it is extremely impolite to send a copy of a followup in the newsgroup to the original poster without marking it as such. Now I have to type my reply all over again. Double work.

Second, I refuse to call a significantly time-consuming routine like "sort" (O(n log n)) twice without good reason. Especially when a simple workaround like writing
a[s] = a+0
will do the trick.

--
David Kastrup   Phone: +49-234-700-5570
Email: dak@neuroinformatik.ruhr-uni-bochum.de   Fax: +49-234-709-4209
Institut fï¿½r Neuroinformatik, Universitï¿½tsstr. 150, 44780 Bochum, Germany

---

Subject: Re: Would you consider this a bug?
Posted by korpela on Thu, 01 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

In article <m2zpbg4bi7.fsf@mailhost.neuroinformatik.ruhr-uni-bochum.de>,
David Kastrup <dak@mailhost.neuroinformatik.ruhr-uni-bochum.de> wrote:

> is not easy to do in the general case. In my particular case, I had
> some data that needed to be processed in sorted order, the results
> needed to be rearranged in original order.
>
> Something like
> s = sort(a)
> a = a[s]
> process(a)
> a[s] = a

I believe that using "a[s]=a[*]" will work properly here. Why "a[s]=a" does what it does, I'm not sure. But it kind of makes sense in that a is an array, and a[*] is the elements of the array. So "a[s]=a[*]" is an assignment of array elements, but "a[s]=a" is apparently something else.

Eric

--
Eric Korpela | An object at rest can never be
korpela@ssl.berkeley.edu | stopped.
Click for home page.

Subject: Re: Would you consider this a bug?
Posted by Kevin Ivory on Thu, 01 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

Leaving the bug topic - switching to algorithms

David Kastrup wrote:
> Unfortunately, translating something like
>
> a[[5,2,3,0,4,1]] = a
> into the equivalent
> a = a[[3,5,1,2,4,0]]
> is not easy to do in the general case.

Actually it is: you need a second call to sort. See example below.

> some data that needed to be processed in sorted order, the results
> needed to be rearranged in original order.
>
> Something like
> s = sort(a)
> a = a[s]
> process(a)
; a[s] = a ; I still think this a dangerous thing to do.
r = sort(s)
a = a[r]

Best regards
Kevin
--
Kevin Ivory Tel: +49 5556 979 434
Max-Planck-Institut fuer Aeronomie Fax: +49 5556 979 240
Max-Planck-Str. 2 mailto:Kevin.Ivory@linmpi.mpg.de
D-37191 Katlenburg-Lindau, GERMANY http://www.gwdg.de/~kivory2/

Subject: Re: Would you consider this a bug?
Posted by David Kastrup on Thu, 01 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

Kevin Ivory <Kevin.Ivory@linmpi.mpg.de> writes:

> David Kastrup wrote:
>> a=[3,4,5]
>> a[[2,1,0]] = a
>>
>> a will be set to [3,4,3]
>>
>> Would you think this a bug?
>
> I think this is extremely dangerous programming: you might get
> undefined or unpredictable results with almost any programming language
> in this type of case.

Disagree. IDL is WRT to assignment a value oriented language, so one might expect first the right side to be evaluated and then the left side to be assigned (unless IDL can do this more efficiently without affecting the result).

If I really wanted a to be destroyed during the operation, I could have written
a[[2,1,0]] = temporary(a)

> The clean way of doing what you want is to
> index the right side of the expression:
> a = a[[2,1,0]]

Unfortunately, translating something like

a[[5,2,3,0,4,1]] = a
into the equivalent
a = a[[3,5,1,2,4,0]]
is not easy to do in the general case. In my particular case, I had some data that needed to be processed in sorted order, the results needed to be rearranged in original order.

Something like
s = sort(a)
a = a[s]
process(a)
a[s] = a

--
David Kastrup Phone: +49-234-700-5570
Email: dak@neuroinformatik.ruhr-uni-bochum.de Fax: +49-234-709-4209
Institut fï¿½r Neuroinformatik, Universitï¿½tsstr. 150, 44780 Bochum, Germany

Subject: Re: Would you consider this a bug?
Posted by Kevin Ivory on Thu, 01 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

David Kastrup wrote:
> a=[3,4,5]
> a[[2,1,0]] = a
>
> a will be set to [3,4,3]
>
> Would you think this a bug?

I think this is extremely dangerous programming: you might get
undefined or unpredictable results with almost any programming language
in this type of case. The clean way of doing what you want is to
index the right side of the expression:
a = a[[2,1,0]]

I don't think this is a bug.

Kevin
--
Kevin Ivory Tel: +49 5556 979 434
Max-Planck-Institut fuer Aeronomie Fax: +49 5556 979 240
Max-Planck-Str. 2 mailto:Kevin.Ivory@linmpi.mpg.de
D-37191 Katlenburg-Lindau, GERMANY http://www.gwdg.de/~kivory2/

Subject: Re: Would you consider this a bug?
Posted by Martin Schultz on Fri, 02 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

David Kastrup wrote:
>
> Kevin Ivory <Kevin.Ivory@linmpi.mpg.de> writes:
>
>> Leaving the bug topic - switching to algorithms
>>
>> David Kastrup wrote:
>>> Unfortunately, translating something like
>>>
>>> a[[5,2,3,0,4,1]] = a
>>> into the equivalent
>>> a = a[[3,5,1,2,4,0]]
>>> is not easy to do in the general case.
>>
>> Actually it is: you need a second call to sort. See example below.
>>

>>> some data that needed to be processed in sorted order, the results
>>> needed to be rearranged in original order.
>>>
>>> Something like
>>> s = sort(a)
>>> a = a[s]
>>> process(a)
>> ; a[s] = a ; I still think this a dangerous thing to do.
>> r = sort(s)
>> a = a[r]
>
> [netiquette comment cut]
>
> Second, I refuse to call a significantly time-consuming routine like
> "sort" (O(n log n)) twice without good reason. Especially when a
> simple workaround like writing
> a[s] = a+0
> will do the trick.
>

Although I really agree with Kevin in terms of the danger that lies in
using some pecularity of IDL (which makes it hard to understand a
program, too), I must admit that there is a significant difference in
execution speed if you can avoid the second sort command. For a 20000
floating point array (a=sin(findgen(20000)/360.*!PI)), I find
execution times of 0.01xx vs 0.06xx secs if once compares (in Eric's
'terminology') a[s]=a[*] with a=a[sort(s)]. Well, this is definitively
not worth risking anything. BUT if you go to 2M elements, the difference
becomes one between 2 secs and 22 secs, which does matter. In any case,
as always: a nice comment in the program would help ;-)

Regards,
Martin.

---------- -------
Dr. Martin Schultz
Department for Engineering&Applied Sciences, Harvard University
109 Pierce Hall, 29 Oxford St., Cambridge, MA-02138, USA

phone: (617)-496-8318
fax : (617)-495-4551

e-mail: mgs@io.harvard.edu
Internet-homepage: http://www-as.harvard.edu/people/staff/mgs/
---------- -------

Subject: Re: Would you consider this a bug?
Posted by muswick on Sat, 03 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

In article <m2ogrw9df6.fsf@mailhost.neuroinformatik.ruhr-uni-bochum.de>,
David Kastrup <dak@mailhost.neuroinformatik.ruhr-uni-bochum.de> wrote:
> Kevin Ivory <Kevin.Ivory@linmpi.mpg.de> writes:
>
>> Leaving the bug topic - switching to algorithms
>>
>> David Kastrup wrote:
>>> Unfortunately, translating something like
>>>
>>> a[[5,2,3,0,4,1]] = a
>>> into the equivalent
>>> a = a[[3,5,1,2,4,0]]
>>> is not easy to do in the general case.
>>
>> Actually it is: you need a second call to sort. See example below.
>>
>>> some data that needed to be processed in sorted order, the results
>>> needed to be rearranged in original order.
>>>
>>> Something like
>>> s = sort(a)
>>> a = a[s]
>>> process(a)
>> ; a[s] = a ; I still think this a dangerous thing to do.
>> r = sort(s)
>> a = a[r]
>
> First, it is extremely impolite to send a copy of a followup in the
> newsgroup to the original poster without marking it as such. Now I
> have to type my reply all over again. Double work.
>
> Second, I refuse to call a significantly time-consuming routine like
> "sort" (O(n log n)) twice without good reason. Especially when a
> simple workaround like writing
> a[s] = a+0
> will do the trick.
>
> --

Actually I don't think it is a bug. One needs to carefully review how IDL handles arrays and indices.

If we go back to the original posted question:

> a = [3,4,5]

> a[[2,1,0]] = a

> a will be set to [3,4,3]

Which is correct, based on 'a' being overwritten. This is avoid huge memory overhead in many cases where the right side equation is complex.

David Kastrup then suggested:
> One has to write
> a[[2,1,0]] = a + 0

> in order to get the more correct seeming [5,4,3].

Which does work, however this is not necessary to get the desired answer. The following works just fine and follows the access methods shown in

a[[2,1,0]] = a[*]

the IDL manuals, but not fully explained. In every example in the IDL manuals where left side subscripting is be used, the right hand is subscripted (even if with only the asteriks) except where left variable is not the same.

So, I do not consider the original question a bug, since the above works just fine by using [*]. It is also about 5 percent faster on large arrays versus adding zero.

Gary Muswick
University Hospitals of Cleveland.

Subject: Re: Would you consider this a bug?
Posted by David Kastrup on Mon, 05 Oct 1998 07:00:00 GMT
View Forum Message <> Reply to Message

muswick@uhrad.com writes:

> Actually I don't think it is a bug. One needs to carefully review
> how IDL handles arrays and indices.

An artifact not mentioned in the language description is a bug.
Imagining how this bug might come about does not make it go away.

> If we go back to the original posted question:
>

>> a = [3,4,5]
>> a[[2,1,0]] = a
>
>> a will be set to [3,4,3]
>
> Which is correct, based on 'a' being overwritten. This is avoid huge memory
> overhead in many cases where the right side equation is complex.

It is a reasonable optimization in general, but an optimization that
changes the results of the command is a bug. Being able to understand
how this bug comes about does not make it "correct". IDL needs to
find out where this optimization is inappropriate and not do it where
harm does ensue.

You cannot define correctness of a program based on the artifacts of
some optimizer. An optimizer has to work behind the scenes.

--
David Kastrup   Phone: +49-234-700-5570
Email: dak@neuroinformatik.ruhr-uni-bochum.de   Fax: +49-234-709-4209
Institut fï¿½r Neuroinformatik, Universitï¿½tsstr. 150, 44780 Bochum, Germany