Subject: Re: simple question (I hope)
Posted by JD Smith on Fri, 30 Mar 2007 18:30:01 GMT
View Forum Message <> Reply to Message

On Fri, 30 Mar 2007 12:20:28 -0700, David Fanning wrote:

> JD Smith writes:
>
>> The one difference which makes this distinction more than pedantic is
>> that true pass by value is very inefficient for large arrays. In a
>> pass-by-value scheme, all of that data (801x1000x100) would be copied
>> via the stack into the local address space of the routine MYPRO. It may
>> sound like a subtle difference, but it does represent a real gain in
>> efficiency, in particular when the temporary variable has a life outside
>> the called routine. Eventually, all temporary variables are harvested,
>> and their memory freed. So while you can't ever get at them yourself,
>> they do offer advantages.
>
> This is the kind of information I usually try to avoid,
> since it makes it VERY hard to teach IDL classes when
> you know it. I agree it is an important point, and I'll
> store it some place in the back of my head (or in an obscure
> corner of my web page), but I really think my explanation
> is a GREAT DEAL more useful in practice! :-)

You're probably right, but if you can make a mental model of IDL's operations in terms of temporary variables, many other issues relating to optimization of IDL memory usage, which have nothing to do with by-value or by-reference calling, become much clearer. You might also gain insight into those mysterious "temporary variables need cleaning up" messages which pop up from time to time ;).

JD