Subject: Re: Optimising A = B+C?
Posted by chase on Fri, 31 May 1996 07:00:00 GMT
View Forum Message <> Reply to Message

>>>> > "Karl" == Karl Glazebrook <kgb@aaoepp.aao.gov.au> writes:
In article <4oll84$318@pa.jach.hawaii.edu> kgb@aaoepp.aao.gov.au (Karl Glazebrook) writes:

Karl> Does anyone know how IDL optimises A=B+C where
Karl> A,B and C are arrays?

Karl> I did a test a while ago and it was several times faster
Karl> than C code along the lines of:

Karl> i=n;
Karl> while (i--)
Karl>   *a-- = *b-- + *c--

Karl> (This was on a 2048x2048 array and everything fiited into
Karl> physical memory.)

Karl> Is it just that it is done in assembler or something?

Karl> just curious...

I don't really know. But RSI may using a vector library for the particular machine architecture. For example, HP 7000 workstations (PA RISC architecture) come with a vector library optimized for performance on that machine. RSI may have used this or built an optimized version of their own. This type of library covers linear vector operations like those in the BLAS (Basic Linear Algebra Subprograms) collection.

Chris
--
==========
Bldg 24-E188
The Applied Physics Laboratory
The Johns Hopkins University
Laurel, MD 20723-6099
(301)953-6000 x8529
chris.chase@jhuapl.edu