CISC 811: High Performance Computing, Assignment 2

Set January 29th due Feb 12th
Office: 308A Stirling Hall. Please email if you wish to set a time for an appointment.
Q1. Which of the following loops vectorize and which do not? If they do not vectorize state why, and whether they can be made to vectorize by a code change.

(a)
```
do i=1,n
a(i)=b(i)+c(i)
d(i)=a(i+1)**2
end do
```
(b)
```
do i=2,n-2,2
a(i+1)=a(i)+b(i+1)
end do
```
(c)
```
do i=1,n,2
b(i)=b(i-2)*3.
end do
```
(d)
```
do i=1,n
j=j+1
k=1-k
a(i)=b(i*k+j,l)+c(i)
end do
```
(e)
```
do i=1,n
if (a(i).gt.1.0) then
f(i)=f(i)+a(i)
else
f(i)=f(i)+1.0
end if
end do
```
(f)
```
do i=1,n
c=a(i)
b(i)=a(i-1)*c
end do
```

Q2. Interpolation of particles on to grids is common problem in both graphics ("splatting") and HPC. Suppose we have a vector of positions r(3,N) and a 3dimensional grid of size L^3 (in the question L=64), where the spatial positions of the particles are commensurate with the grid. For each particle we wish to assign weights to the nearest neighbours on the grid according to the following equation:

$$W(d\mathbf{r}) = W(dx,dy,dz) = \begin{cases} (1-|dx|)(1-|dy|)(1-|dz|) & |dx| \le 1, |dy| \le 1, |dz| \le 1, \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

where dx, dy, dz are the distances from the particle to the grid point being considered. If all the distances to the nearest 8 grid points are different then clearly 8 different weights will be necessary. Further, assume that the grid is triply periodic so that weights that would normally extend beyond the grid wrap around and appear on the other side of the grid (for example, if the grid is on the interval $[1:65)$ the 65th position on a grid 64 across would wrap around to 1).

(a) Write and optimize an algorithm to perform this operation for the list of 64^3 particles given on the class website on a grid of size 64^3. Note that the particle positions are defined on [0:1] and you should scale them up to be on the same interval as the grid cells. Follow the examples for optimization we examined in lecture 2 (for example try not to use IF statements). Provide a hard copy of your source and also email it to me (thacker@astro.queensu.ca)

(b) (Difficult) Vectorize as much of this routine as possible. You will have to find some way to deal with the data-dependence caused by two particles writing to the same grid cell. (HINT: think about the spatial extent of weighting algorithm, the geographic location of the particles and order in which you address the particles.) You can test your speed-up by logging on to the www.cray-cyber.org machines To get the password go to http://www.cray-cyber.org/access/obtain_guestpwd.php and then request a password using the form at the bottom of the page. Login instructions are given on this page. The guest account gives you access to a Cray YMP-EL. Both C (cc) and FORTRAN77 (cf77) compilers are available on the system. The only editor on this system is vi. To get a vectorization report from the compiler: cf77 -Wf"-em" and for the c compiler use cc -h report=fisvt. Alternatively I can give you an account on my workstation which has the Intel compilers which will perform a limited amount of vectorization (but perhaps more importantly they will provide a vectorization report). If you prefer this option then email me. As for the previous question provide both a hard copy and email version of your source.

Q3. Global Arrays (`http://www.emsl.pnl.gov/docs/global/ga.html`) is a programming API that provides a shared-memory like model on distributed memory machines. Categorize Global Arrays according to the parallel programming models we discussed in class. Specifically, describe how it deals with parallelism, decomposition, mapping, communication and synchronization.

Q4. (Q 1.5 from "Parallel Programming" by Wilkinson & Allen) Determine the route taken in a five-dimensional hypercube network from node 7 to node 22 using the e-cube routine algorithm described in the lecture. Repeat for an 8x8 mesh, assuming that nodes are numbered in row order (across the rows starting at the top left corner).

Q5. (Research) For each of the following machines find out the connection topology, diameter, degree and bisection width (for machines where the number of processors is not specified give results as a function of N). In cases where the bisection width is poorly defined quote the minimum and maximum:

(a) NEC Earth Simulator

(b) Cray XT3

(c) Intel Touchstone Delta

(d) BlueGene/L