





<pre>133 { 133 } 133 } 133 } 134 Buffer[i-1]=al[i][1]; 135 } 135 Buffer[i-1]=al[i][1]; 135 } 135 } 135 } 135 } 136 Buffer[i-1]=set, HE_TCO 135 } 137 /* 138 formive measage 139 if(row i=500) 139 } 139 if(row i=500) 130 Buffer[i][1], Midth=3, HE_T 130 Buffer[i][1], Hight=1, HE_T 130 Buffer[i][1], Hight=1, Midth= 130 Buffer[i][1], Midth=3, ME_T 130 Buffer[i][1], Midth=3, ME_T 130 Buffer[i][1], Midth=3, ME_T 130 Buffer[i][1], Hight=1, Midth= 130 Buffer[i][1], Midth=3, Midth=3,</pre>	AT, MT_WORLD); /* Receive from North */ France, ONE_WORLD, setatus); /* Receive from East */ ONE, %4_WORLD, setatus); /* Receive from South */ 2. MFI FLOAT.
133         boffer[i-1]-val[i[1];           133         MFT_shouldbatces, MES[shot-2, MFT_rC           134         MFT_shouldbatces, MES[shot-2, MFT_rC           135         * Security messages           136         * Security messages           137         MFT_shouldbatces, MES[shot-2, MFT_rC           138         * Security messages           139         * Security messages           131         MFT_shouldbatces, MES[shot-2, MFT_rC           132         * Security messages           133         MFT_shouldbatces, MES[shot-2, MFT_rC           144         (fccol !=Right)           145         (fccol !=Right)           146         (fccol !=Right)           147         MFT_shouldbatces, MES[shot-2, MFT_rC           148         (fccol !=Right)           149         (statisfier)           149         (statisfier)           144         (statisfier)           145         (fccol !=Right)           146         (statisfier)           147         (statisfier)           148         (statisfier)           149         (statisfier)           149         (statisfier)           141         (statisfier)	NY, MY_WORLD); /* Bacelive from North */ /TLONY, ONE_WORLD, status}; /* Bacelive from East */ ONY, MY_WORLD, status; /* Bacelive from South */ 2, MFI FLOAY.
) ) ) ) ) ) ) ) ) )	AT SME_WORLD); /* Beenive from North */ //A Beenive from East */ ONF, /* Beenive from East */ ONF, /* Beenive from East */
133         WP_lend(buffer, led)t-3, WP_TCO           135         >           136         MesUPE(myr), tesp, MP_TCO           137         *           138         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           139         *           140         *           141         *           142         *           143         *           144         *           144         *           144         *           144         *           144         *           144         *           145         *           146         *           147         *	<pre>/* Heresive Eron North */ reaction From North */ reaction /* Becenive Eron East */ core, %%_DORLD, &amp;status); /* Becenive Eron East */ core, %%_DORLD, &amp;status); /* Becenive Eron Eosth */ core, c</pre>
<pre>133</pre>	<pre>/* Receive from North */ FLAT. FLAT. (* Receive from East */ ACT. (* Receive from East */ ACT. (*</pre>
<pre>133 ) 33 /* 34 /* 35 /* 36 /* 3</pre>	<pre>/* Bacsive from North */ FLAF, FLAF, MULTING, sature); /* Bacsive from East */ Ook, MS_NORLD, Astatus); /* Heresive from South */ 2, MSI FLAGE.</pre>
/*	/* Receive from North */ FLAAT, ONE_VORED, setatus); /* Receive from East */ ORE, ORED, setatus); /* Receive from South */ 2, MELFLAAT,
134         * Boscive messages           135         * Koncive messages           136         Xerchov (=>rpp)           136         Kerchov (=>rpp)           141         (PT_Recvent(0)[1]), Vichthez, HPT_C           144         )         NochDe(myTh), Kag, HPT_C           144         (Gcol !=Right)         (HPT_Becvent(0)[1]), Kag, HPT_C           145         (Gcol !=Right)         (HPT_Becvent(0)[1]), Kag, HPT_C           146         (Gcol !=Right)         (HPT_Becvent(0)[1]), Kag, HPT_C           147         (HPT_Becvent(0)[1]), Kag, HPT_C         (HPT_Becvent(0)[1]), Kag, HT_C           148         (HPT_Becvent(0)[1]), Kag, HT_C         (HPT_Becvent(0)[1]), Kag, HT_C           149         (HPT_Becvent(0)[1]), Kag, HT_C         (HPT_Becvent(0)[1]), Kag, HT_C	<pre>/* Receive from North */ FLOAT. FLOAT. ONE_WORKLO, status); /* Receive from East */ ONE_MORLD, setatus); /* Receive from Booth */ 2. MET FLOAT.</pre>
<pre></pre>	<pre>/* Beenive from North */ // Beenive from East */ /* Beenive from East */</pre>
<pre>if(cov i=&gt;cop)     if(cov i=&gt;cop)     if(cov i=&gt;cop)     if(cov i=&gt;cov(+&gt;ci(i))), width&gt;-2, MFT_     if(col i=&gt;ci(i), i+&gt;, MFT_cov(+&gt;cov(+<cov(+>cov(+&gt;cov(+<cov(+>cov(+&gt;</cov(+></cov(+></pre>	<pre>// Beesive from North */ TrACAT. OPPLVORED, #Status); // Beesive from East */ OPPLVORED, #status); /* Beesive from South */ 2. MSE FLOAD.</pre>
<pre>141 (</pre>	PLONT. PLONT. V* Receive from East */ ONTLD, sotatus); /* Receive from South */ 2. MFI FLOAT.
142         H9:_[Mesv(aval[0][1], Width-2, M97_0]           143         Morth@fwighth, 4ag, M97_0]           144         Morth@fwighth, 4ag, M97_0]           145         (f(coll=stight))           146         (f(coll=stight))           147         (f(coll=stight))           148         (f(coll=stight))           149         (f(coll=stight))           149         (f(coll=stight))           149         (f(coll=stight))           149         (f(coll=stight))           149         (f(coll=stight))           150         (f(coll=stight))           151         (f(coll=stight))           155         (f(coll=stight))           155         (f(coll=stight))           155         (f(coll=stight))           155         (f(coll=stight))           156         (f(coll=stight))           157         (f(coll=stight))           158         (f(coll=stight))           159         (f(coll=stight))           150         (f(coll=stight))           150         (f(coll=stight))           150         (f(coll=stight))	FLONT, FLONT, /* Reserve from East */ OST, /* Reserve from South */ 2. MFL FLOAT.
143         BothPig(myID), tag, HP1_0           144         if(coll=Right)           145         if(coll=Right)           146         If(coll=Right)           147         (MF1_Developters, Bight-2, MF1_DL           148         Im1_Bevelopters, Bight-2, MF1_DL           149         (mf1_Bevelopters, Bight-1, MF1_DL           149         (mf1_Bevelopters, Bight-2, MF1_DL           149         (mf1_Bevelopters, Bight-2, MF1_DL           151         (mf1_Gevelopters, Bight-1, Bight-1	<pre>// Receive From East */</pre>
144 ) 146 if(col 1-RLight) 147 ( 148 North Starfford, Inc., North J., 149 149 ( 149 North Starfford, North Starfford, 149 149 ( 149 North Starfford, 149 149 North Sta	<pre></pre>
145         if(colisRight)           146         if(colisRight)           147         if(colisRight)           148         MP_EDec(Monter, Montpht-1, MP_ED)           149         if(colisRight)           149         if(colisRight)           149         if(colisRight)           149         if(colisRight)           151         if(colisRight)           155         if(colisRight)           155         if(colisRight)           155         if(colisRight)           155         if(colisRight)           155         if(colisRight)           156         if(colisRight)           157         Montparticle           158         if(colisRight)           159         Montparticle           150         if(colisRight)           151         if(colisRight)           155         if(colisRight)           156         if(colisRight)           157         if(colisRight)           158         if(colisRight)           159         if(colisRight)           150         if(colisRight)           151         if(colisRight)           152         if(colisRight)	<pre>/* Becnive from East */ OAF, MS[_MORLD, Astatus); /* Becnive from South */ 2, MSI FLOAT.</pre>
<pre>146 if(coll=inight) 146 if(coll=inight) 147 if(coll=inight) 148 if(coll=inight) 149 if(coll=inight) 149 if(coll=inight) 149 if(coll=inight) 149 if(coll=inight) 151 if(coll=inight) 151 if(coll=inight) 151 if(coll=inight) 152 if(coll=inight) 153 if(coll=inight) 154 if(coll=inight) 155 if(coll=inight) 155 if(coll=inight) 156 if(coll=inight) 157 if(coll=inight) 158 if(coll=inight) 1</pre>	<pre>/* Baselve from East */ OOF, MM_NORED, &amp;status); /* Baselve from South */ 2, MMI FLOOT.</pre>
<pre>147 {     (</pre>	OAF, MSK_WORLD, &status); /* Receive from South */ 2. MSI FLOAT.
140         HPI_Beev(beffer, Ins,H-2, HE]_CN           151         for(i=1)_iele(abc-1)_i+i+)           152         (           153         for(i=1)_iele(abc-1)_i+i+)           154         (           155         (           156         (           157         (           158         )           159         (           150         (           156         (           157         (           158         )           159         (           150         (           156         (           157         (           158         (           159         (           150         (           151         (           152         (           153         (           154         (           155         (           156         (           157         (           158         (           159         (           150         (           150         (           150         (	OAF, MORED, Sotatus); /* Receive from South */ 2. MFT FLOAT,
140         KaniF2(mylD), tag, NF2_CO           151         (reispin-1) i=1)           151         (val(i(kidb-1)=buffer[i-1]))           153         )           154         )           155         (ki(kidb-1)=buffer[i-1]))           156         )           157         (frow 1=Bottom)           158         )           159         SouthEr(mylD), tag, MF1_CO           159         SouthEr(mylD), tag, MF1_CO	994_WORLD, Astatus); /* Receive from South */ 2. MFT FLOAT.
<pre>11 for(1#1)/ztmidgfr=1/1++) 152 (val(1)(Width=1)=buffer(i=1); 153 ) 154 ) 155 (f(row i=Bottom) 157 ( 158 (RFT_Recv(sval(0)(Neight=1), Width=1); 157 ( 158 SoutDFE(mp1), tag, NFT_C 159 )</pre>	/* Receive from South */ 2, MPI FLOAT,
152 val[i][Width-1]=buffer[i=1]; 153 } 154 } 155 156 if(row i=botton) 157 { 157 { 159 NPI_Recv(&val[0][Beight-1], Width-1 159 SouthPE[myID], tag, MPI_C 160 }	/* Receive from South */ 2. MPI FLOAT,
<pre>153</pre>	/* Receive from South */ 2, MPI FLOAT,
154 ) 155 156 if(row !=Bottom) 157 { 158 HPI_Recv(&val[0][Height-1], Width-: 159 SouthPE(myID), tag, MPI_O 160 )	/* Receive from South */ 2. MPI FLOAT,
<pre>155 ' 156 if(row !=Bottom) 157 { 158 MPI_Recv(&amp;val(0)[Height-1], Width-' 159 SouthPE(myID), teg, MPI_C 160 }</pre>	/* Receive from South */ 2, MPI FLOAT,
<pre>156 if(row !=Bottom) 157 { 158 MPI_Recv(4val[0][Height-1], Width 159 SouthPE(myID), tag, MPI_O 160 }</pre>	/* Receive from South */ 2, MPI FLOAT,
<pre>157 { 158 NPI_Recv(&amp;val[0][Height-1], Width 159 SouthPE(myID), tag, MPI_C 160 }</pre>	2, MPI FLOAT,
<pre>158 NPI_Recv(&amp;val[0][Height-1], Width- 159 SouthPE(myID), tag, MPI_O 160 )</pre>	2, MPI PLOAT,
159 SouthPE(myID), tag, MPI_C 160 )	
160 )	OMM_WORLD, &status);
101	In manufact from more of
162 /	/ NUCLEVE LEON WERE -/
164 NDT Recy(Abuffer, Height-2, MDT FI	OPT -
165 WestPE(mvID), tag, MPI CO	MM WORLD, Astatus):
166 for(i=1; i <reight-1; i++)<="" td=""><td></td></reight-1;>	
167 (	
<pre>168 val(i)(0)=buffer(i=1);</pre>	
169 )	
170 }	
171	
172 delta=0.0; /* Calculate average, del	ta for all points */
173 for(i=1; i <height-1; i++)<="" td=""><td></td></height-1;>	
1/4 ( 1777	

177 178 179 180 181 182 183 184 185 186 187 188	<pre>average=(val[i-1][j]+val[i][j+1]+ val[i+1][j]+val[i][j-1])/4; delta=Max(delta, Abs(average-val[i][j])); new[i][j]=average; } } /* Find maximum diff */ MPI_Reduce(idelta, iglobalDelta, 1, MPI_FLOAT, MPI_MIN, RootProcess, MPI_COMM_WORLD); Swap(val, new); ) while(globalDelta &lt; THRESHOLD);</pre>

# Question: Does this lead to deadlock?

# • No! Why not?

- Even though communication is blocking,
  - Send returns when MPI library on receiving end acknowledges that data has been received

  - The receiving process may not yet have received the data
  - So, a sequence of blocking Sends will not block

### • What about receives?

- Receiving process must wait until data arrives before proceeding
- A sequence of blocking receives must be received in order

11/09/10

UNIVERSITY

# Foster's methodology (Chapter 2)

- Partitioning: divide the computation to be performed and the data operated on by the computation into small tasks. 1. The focus here should be on identifying tasks that can be executed in parallel.
- Communication: determine what communication needs to be carried out among the tasks identified in the previous step. 2.
- Agglomeration or aggregation: combine tasks and communications identified in the first step into larger tasks. 3. For example, if task A must be executed before task B can be executed, it may make sense to aggregate them into a single composite task.
- Mapping: assign the composite tasks identified in the previous step to processes/threads. 4.

This should be done so that communication is minimized, and each process/thread gets roughly the same amount of work.

Copyright © 2010, Elsevier Inc. All rights Reserved







Calculate Force as a function of mass and positions Force between two particles (6.1)  $\mathbf{f}_{qk}(t) = -\frac{Gm_qm_k}{|\mathbf{s}_q(t) - \mathbf{s}_k(t)|^3} [\mathbf{s}_q(t) - \mathbf{s}_k(t)]$ Total force on a particle (6.2)  $\mathbf{F}_q(t) = \sum_{\substack{k=0\\k\neq q}}^{n-1} \mathbf{f}_{qk} = -Gm_q \sum_{\substack{k=0\\k\neq q}}^{n-1} \frac{m_k}{|\mathbf{s}_q(t) - \mathbf{s}_k(t)|^3} [\mathbf{s}_q(t) - \mathbf{s}_k(t)]$ Copyright © 2010, Elsevier Inc. All rights Reserved





Computation of the forces: Basic Algorith	m
For all pairs (except <q,q>) compute Total Force</q,q>	
<pre>for each particle q {   for each particle k != q {     x_diff = pos[q][X] - pos[k][X];     y_diff = pos[q][Y] - pos[k][Y];     dist = sqrt(x_diff*x_diff + y_diff*y_diff);     dist_cubed = dist*dist*dist;     forces[q][X] -= G*masses[q]*masses[k]/dist_cubed * x_forces[q][Y] -= G*masses[q]*masses[k]/dist_cubed * y_} }</pre>	diff; diff;
Copyright © 2010, Elsevier Inc. All rights Reserved	THE UNIVERS OF UTAH

Redu Com	ce calculation of Total Force, capitalizing on $f_{qk}$ = - $f_{kq}$ pute half the pairs
for e	each particle q
fo	$\operatorname{prces}[q] = 0;$
for e	ach particle q {
fo	$ ho {f r}$ each particle k > q $\{$
	$x_diff = pos[q][X] - pos[k][X];$
	$y_diff = pos[q][Y] - pos[k][Y];$
	<pre>dist = sqrt(x_diff*x_diff + y_diff*y_diff);</pre>
	dist_cubed = dist*dist*dist;
	force_qk[X] = G*masses[q]*masses[k]/dist_cubed * x_diff;
	<pre>force_qk[Y] = G*masses[q]*masses[k]/dist_cubed * y_diff</pre>
	<pre>forces[q][X] += force_qk[X];</pre>
	forces[q][Y] += force_qk[Y];
	forces[k][X] -= force_qk[X];
	<pre>forces[k][Y] -= force_gk[Y];</pre>
}	
}	







# Other topics on OpenMP parallelization

Data structure choice – group together in structure or separate arrays?

Thread creation

Alternative synchronization (reduced method)

Private data (reduced method)

Nowait

I/O (single)

UNIVERSITY

# What's different with Pthreads

- Must write more code to do basic OpenMP things:
  - Create threads, pass parameters, determine number and mapping of iterations and initiate "parallel loop"
  - Explicit barriers between "loops"
  - Explicit synchronization
  - Explicit declaration of global data that is shared, or thread-local declaration of private data





**MPI Parallelization: Basic version** 



MPI_Scatter()	
int MPI_Scatter( void *sendbuffer, int sendcount, MPI_Datatype sendtype, int destbuffer, int destbuffer, MPI_Datatype desttype,	<pre>// Scatter routine // Address of the data to send // Number of data elements to send // Type of data elements to send // Address of buffer to receive // Number of data elements to receive // Type of data elements to receive</pre>
Int root, MPI_Comm *comm );	// Rank of the root process // An MPI communicator
Arguments:	
The first three arguments specify the elements to send to each process. The the root process.	e address, size, and type of the data ese arguments only have meaning for
The second three arguments specify t elements for each receiving process. I and the receiving data may differ as a	he address, size, and type of the data The size and type of the sending data means of converting data types.
<ul> <li>The seventh argument specifies the r data.</li> </ul>	root process that is the source of the
<ul> <li>The eighth argument specifies the MI</li> </ul>	PI communicator to use.
Notes:	
This routine distributes data from the re including the root. A more sophisticated MPI_Scatterv(), allows the root proce data to the various processes. Details car	oot process to all other processes, l version of the routine, zss to send different amounts of n be found in the MPI standard.
Return value:	
An MPI error code.	
© 2009 Pearson Education. Inc	Publishing as Pearson Addison-Wesley



# MPI\_Gather is analogous

MPI\_Gather: Gathers together values from a group of processes int MPI\_Gather(void \*sendbuf, int sendent, MPI\_Datatype sendtype, void \*recvbuf, int recvent, MPI\_Datatype recvtype, int root, MPI\_Comm comm) Input Parameters sendbuf starting address of send buffer (choice) sendcount number of elements in send buffer (integer) sendtype data type of send buffer elements (handle) recvecount number of elements for any single receive (integer, significant only at root) recvtype data type of recv buffer elements (significant only at root) (handle) root rank of receiving process (integer) comm communicator (handle) Output Parameter recvbuf address of receive buffer (choice, significant only at root)

# Detour: MPI\_Allgather (a collective)

 $\ensuremath{\textbf{MPl\_Allgather:}}$  Gathers data from all tasks and distributes the combined data to all tasks

int MPI\_Allgather (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, MPI\_Comm comm)

#### **Input Parameters**

sendbuf sendcount	starting address of send buffer (choice) nt number of elements in send buffer (integer) addat byten of send buffer elements (handle)					
recvcount recvtype comm	data type of send buller elements (handle) number of elements received from any process (integer) data type of receive buffer elements (handle) communicator (handle)					
Output Parameter						
recvbufaddress of receive buffer (choice)						

UNIVERSITY

# I/O: Ch. 6, p. 291

if (my\_rank == 0) {
 for each particle, read masses[particle], pos[particle], vel[particle];
}

MPI\_Bcast(masses, n, MPI\_DOUBLE, 0, comm); MPI\_Bcast(pos, n, vect\_mpi\_t, loc\_vel, loc\_n, vect\_mpi\_t, 0, comm) MPI\_Scatter(vel, loc\_n, vect\_mpi\_t, 0, comm);

# **MPI Issues**

- Communicating standard data types performs much better than derived data types like structures
  - Separate fields (position, velocity, mass) into arrays that can be communicated independently
- Is it feasible for each processor to have all of the processor mass data locally? The position data?
  - $\boldsymbol{\cdot}$  Depends on how many bodies and capacity constraints













OF UTAH