Molecular Dynamics

Sathish Vadhiyar

Courtesy: Dr. David Walker, Cardiff University

Molecular Dynamics

- Application in many areas including biological systems (e.g. drug discovery), metallurgy (e.g. interaction of metal with liquids) etc.
- A domain consisting of number of particles (molecules)
- Each molecule, i is exerted a force, f_{ij} by another molecule, j
- Forces are of two kinds:
  - Non-bonded forces – computations of pairwise interactions.
  - Bonded forces – computations of interactions between molecules that are connected by bonds. Connectivities are fixed. Hence these forces depend on topology of the structure

Molecular Dynamics

- The sum of all the forces, $F_i = \sum_j f_{ij}$ makes the particles assume a new position and velocity
- Particles that are r distance apart do not influence each other
- Thus non-bonded forces are only computed between atoms that are within this cutoff distance
- Given initial velocities and positions of particles, their movements are followed for discrete time steps

MD Parallelization

- 3 methods
- 1. Atom decomposition
- 2. Space decomposition
- 3. Force decomposition

Atom Decomposition

- Each processor is assigned N/P atoms and updates their positions and velocities irrespective of where they move in the physical domain
- The computational work involved can be represented by the NxN matrix, F, where Fi,j is the non-bonded force on atom i due to atom j
- x and f are vectors that represent positions of and total force on each atom

Atom Decomposition

- For parallelization, F, x and f are distributed with 1-D block distribution across processors. i.e., every processor computes consecutive N/P rows
- Each processor will need the positions of many atoms owned by other processors; hence each processor stores a copy of all N atom positions, x
- Hence this algorithm is also called **replicated data** algorithm

RD Algorithm

- For each time step
- each processor computes forces on its atoms
- updates positions
- processors communicate their positions to all the other processors

- Different atoms have different neighbor entitites; hence the F matrix has to be load balanced
- The main disadvantage is the all-to-all communication of x; also causes memory overhead since x is replicated

Method 2 – Space decomposition

- Similar to Jacobi parallelization. Domain or space is decomposed
- In Jacobi iterations (2D), communication requirements are known in advance
- In a typical Molecular Dynamics simulation problem, the amount of data that are communicated between processors are not known in advance
- The communication is slightly irregular

Space Decomposition - Solution

- The cutoff distance, r is used to reduce the time for summation from $O(n^2)$

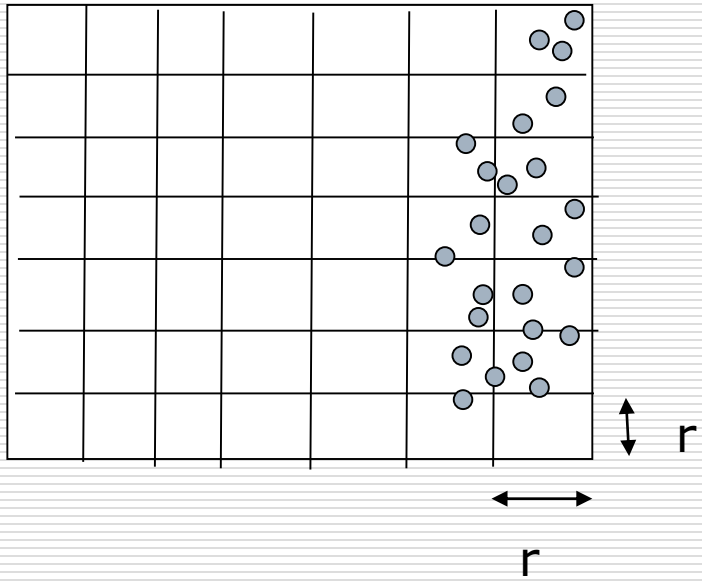

Domain decomposed into cells of size rxr

Particles in one cell interact with particles in the neighbouring 8 cells and particles in the same cell

Space Decomposition - Solution

Data structures:

An array of all particles. Each element holds <position, velocity>

A 2D array of linked lists, one for each cell. Each element of a linked list contains pointers to particles.

```
struct particle{
  double position[2];
  double velocity[2];
} Particles[MAX_PARTICLES];

struct list{
  particle* part;
  struct list* next;
}*List[MAX_CELLSX][MAX_CELLSY];
```

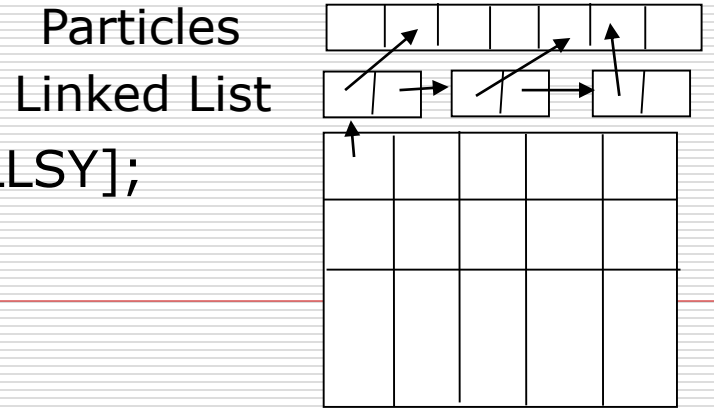

Space Decomposition – Sequential Logic

```
Initialize Particles and Lists;

for each time step
  for each particle i
    Let cell(k, l) hold i
    F[i] = 0;
     for each particle j in this cell and neighboring 8 cells, and
      are r distance from i{
        F[i]+= f[i, j];
    }
     update particle[i].{position, velocity} due to F[i];
     if new position in new cell (m,n) update Lists[k,l] and
      Lists[m,n]
```

MD – Space Decomposition

A 2D array of processors similar to Laplace

Each processor holds a set of cells

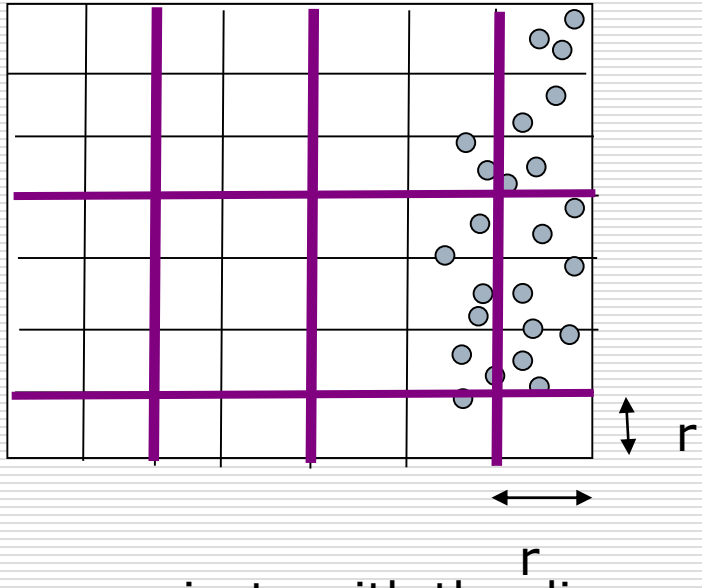

Differences:

- A processor can communicate with the diagonal neighbors
- Amount of data communicated varies over time steps
- Receiver does not know the amount of data

MDS – parallel solution

- Steps

1. Communication – Each processor communicates parameters of the particles on the boundary cells to its 8 neighboring cells

   Challenges – to communicate diagonal cells

2. Update – Each processor calculates new particle velocities and positions

3. Migration – Particles may migrate to cells in other processors

Other challenges:

1. Appropriate packing of data.
2. Particles may have to go through several hops during migration

Assumptions:

1. For simplicity, let us assume that particles are transported to only neighboring cells during migration

MDS – parallel solution – 1st step

- Communication of boundary data

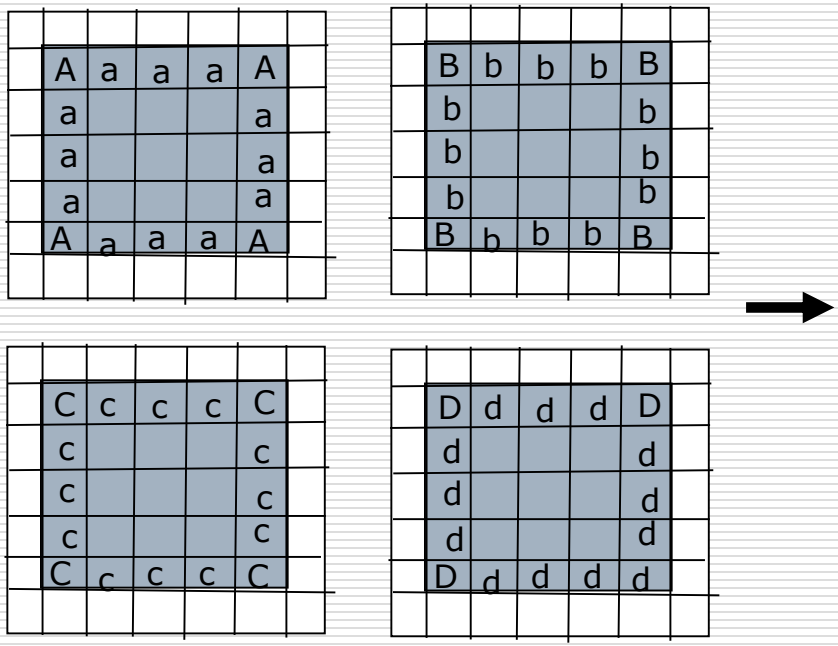

MDS – parallel solution – 1st step

- Communication of boundary data

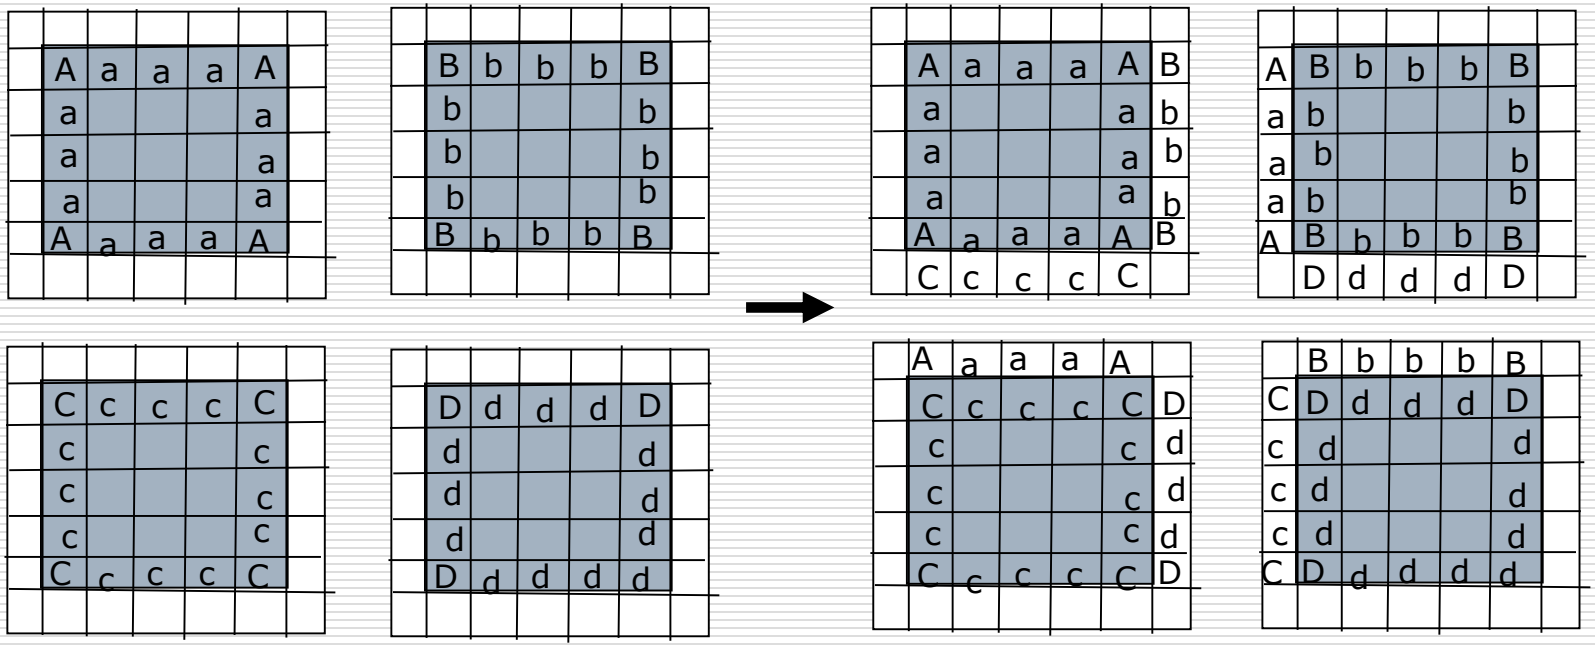

MDS – parallel solution – 1st step

- Communication of boundary data

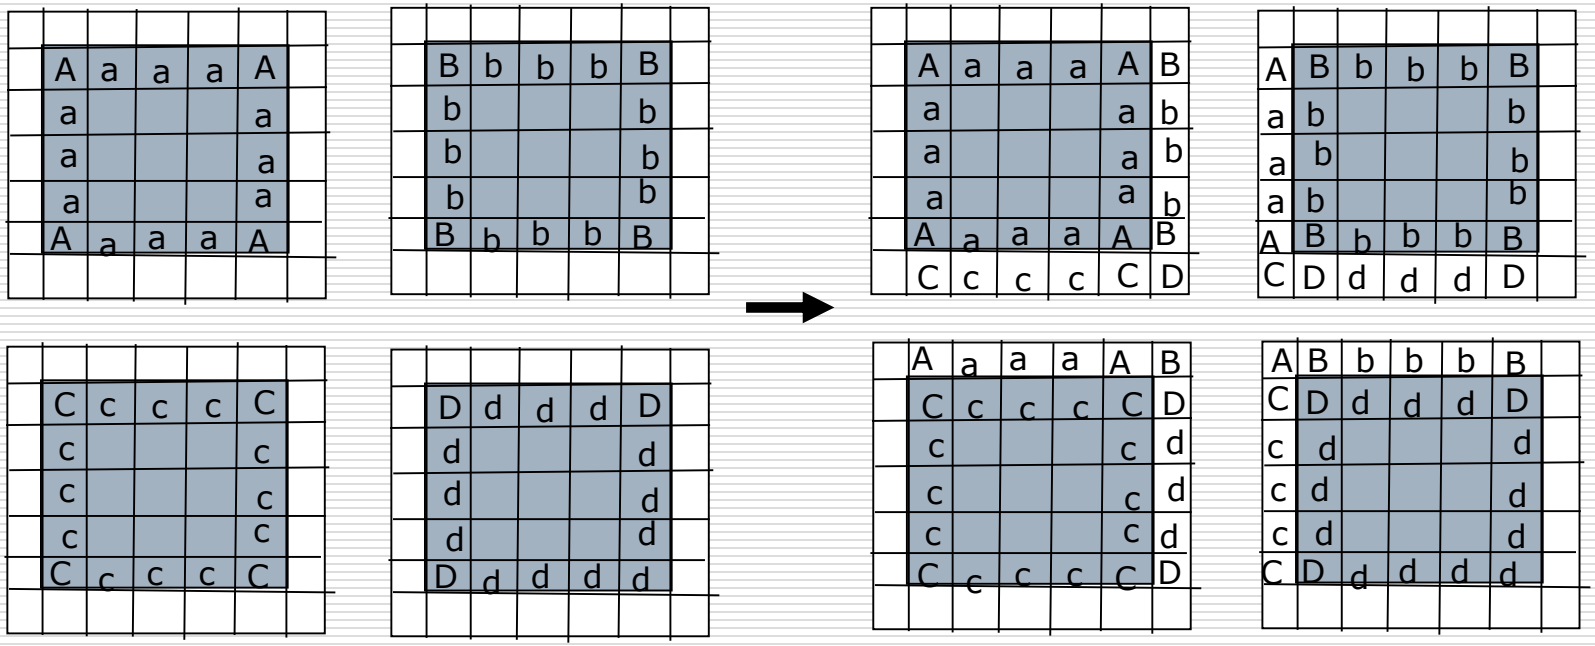

Can be achieved by ?

Shift left, shift right, shift up, shift down

MDS – parallel solution – 1st step

Left shift

```
nsend = 0;
for(i=0; i<local_cellsx; i++){
  for each particle p in cell (i, 1){
    pack position of p in sbuf
    nsend += 2
  }
}

MPI_Sendrecv(sbuf, nsend, …, left,..
                        rbuf, max_particles*2, …, right, &status);
MPI_Getcount(status, MPI_DOUBLE, &nrecv);
particles = nrecv/2;
for(i=0; i<particles; i++){
  read (x,y) from next 2 positions in rbuf;
  add (x,y) to particles[local_particle_count+i];
  determine cell k, l for the particle
  Add it to list (k, l);
}
```

MDS – parallel solution – 2nd step

Update:

- Similar to sequential program.
- A processor has all the required information for calculating F_i for all its particles
- Thus new position and velocity determined.
- If new position belongs to the same cell in the same processor, do nothing
- If new position belongs to the different cell in the same processor, update link lists for old and new cells.

MDS – parallel solution – 3rd step

- If new position belongs to the different cell in a different processor – **particle migration**

```
for each particle p
  update {position, velocity}
  determine new cell
  if new cell # old cell
    delete p from list of old cell
    if(different processor)
      pack p into appropriate communication buffer
      remove p from particle array
```

Shift left
Shift right
Shift up
Shift down

MDS – parallel solution – 3rd step

- This shifting is a bit different from the previous shifting
- A processor may just act as a transit point for a particle
- Hence particles have to be packed with care

Shift left:

```
MPI_Sendrecv(leftbuf, nsend_left, …, left
                        rbuf, max_size*4, .., right, &status);
MPI_Getcount(status, MPI_DOUBLE, &nrecv);
particles = nrecv/4;

for(i=0; i<particles; i++){
  read next 4 numbers in {x, y vx, vy}
  if(particle in this process)
    add particle to particle array
    determine cell
    add particle to list for the cell
  else
    put data in the appropriate comm. buffer for the next up or down
      shifts
}
```

MDS – comments

- Generic solution
  - A particle can move to any cell
  - Force can be from any distance
- Load balancing

Force Decomposition

- For computing the total force on an atom due to all the other atoms, the individual force contributions from the other atoms are independent and can be parallelized
- Fine-grained parallelism
- Especially suitable for shared-memory (OpenMP) parallelization

Hybrid Decomposition

- Divide the domain into cells (spatial decomposition)
- Create a parallel thread whose responsibility is to compute interacting forces between every pairs of cells (force decomposition)