

# Lecture 2: Parallel Programs

---

- Topics: parallel applications, parallelization process, consistency models

# Parallel Application Examples

---

- Simulating ocean currents
- Simulating evolution of galaxies
- Visualizing complex scenes using raytracing
- Mining data for associations

# Ocean

---

- Simulates motion of water currents, influenced by wind, friction, etc.
- We examine a horizontal cross-section of the ocean at a time and the cross-section is modeled as a grid of equidistant points
- At each time step, the value of each variable at each grid point is updated based on neighboring values and equations of motion
- Apparently high concurrency

# Barnes-Hut

---

- Problem studies how galaxies evolve by simulating mutual gravitational effects of  $n$  bodies
- A naïve algorithm computes pairwise interactions in every time step ( $O(n^2)$ ) – a hierarchical algorithm can achieve good accuracy and run in  $O(n \log n)$  by grouping distant stars
- Apparently high concurrency, but varying star density can lead to load balancing issues

# Data Mining

---

- Data mining attempts to identify patterns in transactions
- For example, association mining computes sets of commonly purchased items and the conditional probability that a customer purchases a set, given they purchased another set of products
- Itemsets are iteratively computed – an itemset of size  $k$  is constructed by examining itemsets of size  $k-1$
- Database queries are simpler and computationally less expensive, but also represent an important benchmark for multiprocessors

# Parallelization Process

---

- Ideally, a parallel application must be constructed by designing a parallel algorithm from scratch
- In most cases, we begin with a sequential version – the quest for efficient automated parallelization continues...
- Converting a sequential program involves:
  - Decomposition of the computation into *tasks*
  - Assignment of tasks to *processes*
  - Orchestration of data access, communication, and synchronization
  - Mapping or binding processes to processors

# Partitioning

---

- Decomposition and Assignment are together called partitioning – partitioning is algorithmic, while orchestration is a function of the programming model and architecture
- The number of tasks at any given time is the level of concurrency in the application – the average level of concurrency places a bound on speedup (Amdahl's Law)
- To reduce inter-process communication or load imbalance, many tasks may be assigned to a single process – this assignment may be either static or dynamic
- We will assume that processes do not migrate (fixed mapping) in order to preserve locality

# Parallelization Goals

---

Step	Architecture-Dependent?	Major Performance Goals
Decomposition	Mostly no	Expose enough concurrency
Assignment	Mostly no	Balance workload Reduce communication volume
Orchestration	Yes	Reduce communication via data locality Reduce communication and synch cost Reduce serialization at shared resources Schedule tasks to satisfy dependences early
Mapping	Yes	Put related processes on same processor Exploit locality in network topology



# Case Study: Ocean Kernel

---

- Gauss-Seidel method: sweep through the entire 2D array and update each point with the average of its value and its neighboring values; repeat until the values converge
- Since we sweep from top to bottom and left to right, the averaging step uses new values for the top and left neighbors, and old values for the bottom and right neighbors


# Ocean Kernel

---

```
Procedure Solve(A)
begin
  diff = done = 0;
  while (!done) do
    diff = 0;
    for i ← 1 to n do
      for j ← 1 to n do
        temp = A[i,j];
        A[i,j] ← 0.2 * (A[i,j] + neighbors);
        diff += abs(A[i,j] - temp);
      end for
    end for
    if (diff < TOL) then done = 1;
  end while
end procedure
```

# Concurrency

---



- Need synch after every anti-diagonal
- Potential load imbalance

# Algorithmic Modifications


---

- Red-Black ordering: the grid is colored red and black similar to a checkerboard; sweep through all red points, then sweep through all black points; there are no dependences within a sweep
- Asynchronous updates: ignore dependences within a sweep → you may or may not get the most recent value
- Either of these algorithms expose sufficient concurrency, but you may or may not converge quickly

# Assignment

---

- With the asynchronous method, each process can be assigned a subset of all rows



- What is the degree of concurrency?
- What is the communication to computation ratio

# Orchestration

---

- Orchestration is a function of the programming model and architecture
- Consider the shared address space model – by using the following primitives, the program appears very similar to the sequential version:
  - CREATE: creates  $p$  processes that start executing at procedure *proc*
  - LOCK and UNLOCK: acquire and release mutually exclusive access
  - BARRIER: global synchronization: no process gets past the barrier until  $n$  processes have arrived
  - WAIT\_FOR\_END: wait for  $n$  processes to terminate

# Shared Address Space Model

---

```
int n, nprocs;
float **A, diff;
LOCKDEC(diff_lock);
BARDEC(bar1);

main()
begin
  read(n); read(nprocs);
  A ← G_MALLOC();
  initialize (A);
  CREATE (nprocs,Solve,A);
  WAIT_FOR_END (nprocs);
end main
```

```
procedure Solve(A)
  int i, j, pid, done=0;
  float temp, mydiff=0;
  int mymin = 1 + (pid * n/nprocs);
  int mymax = mymin + n/nprocs -1;
  while (!done) do
    mydiff = diff = 0;
    BARRIER(bar1,nprocs);
    for i ← mymin to mymax
      for j ← 1 to n do
        ...
      endfor
    endfor
    LOCK(diff_lock);
    diff += mydiff;
    UNLOCK(diff_lock);
    BARRIER (bar1, nprocs);
    if (diff < TOL) then done = 1;
    BARRIER (bar1, nprocs);
  endwhile
```

# Message Passing Model

---

```
main()
  read(n); read(nprocs);
  CREATE (nprocs-1, Solve);
  Solve();
  WAIT_FOR_END (nprocs-1);

procedure Solve()
  int i, j, pid, nn = n/nprocs, done=0;
  float temp, tempdiff, mydiff = 0;
  myA ← malloc(...)
  initialize(myA);
  while (!done) do
    mydiff = 0;
    if (pid != 0)
      SEND(&myA[1,0], n, pid-1, ROW);
    if (pid != nprocs-1)
      SEND(&myA[nn,0], n, pid+1, ROW);
    if (pid != 0)
      RECEIVE(&myA[0,0], n, pid-1, ROW);
    if (pid != nprocs-1)
      RECEIVE(&myA[nn+1,0], n, pid+1, ROW);

    for i ← 1 to nn do
      for j ← 1 to n do
        ...
      endfor
    endfor
    if (pid != 0)
      SEND(mydiff, 1, 0, DIFF);
      RECEIVE(done, 1, 0, DONE);
    else
      for i ← 1 to nprocs-1 do
        RECEIVE(tempdiff, 1, *, DIFF);
        mydiff += tempdiff;
      endfor
      if (mydiff < TOL) done = 1;
      for i ← 1 to nprocs-1 do
        SEND(done, 1, i, DONE);
      endfor
    endif
  endwhile
```



# Message Passing Model

---

- Note that each process has two additional rows to store data produced by its neighbors
- Unlike the shared memory model, execution is deterministic -- two executions will produce the same result
- A send-receive match is a synchronization event – hence, we no longer need locks while updating the diff counter, or barriers while allowing processes to proceed with the next iteration

# Models for SEND and RECEIVE

---

- Synchronous: SEND returns control back to the program only when the RECEIVE has completed – will it work for the program on the previous slide?
- Blocking Asynchronous: SEND returns control back to the program after the OS has copied the message into its space -- the program can now modify the sent data structure
- Nonblocking Asynchronous: SEND and RECEIVE return control immediately – the message will get copied at some point, so the process must overlap some other computation with the communication – other primitives are used to probe if the communication has finished or not

# Title

---

- Bullet