Assignment: Distributed Memory: representation and algorithm

1 Reduction

Consider the following three algorithms

```
reduce-star(p, P, val) {
  if (p == 0) {
    for (i=1; i<P; ++i) {
      recv vald from i;
      val += vald;
    }
  }
  else {
    send val to 0;
  }
}
```

```
reduce-chain(p, P, val) {
  if ( p != P-1) {
    recv vald from p+1;
    val += vald;
  }
  if (p != 0) {
    send val to p-1;
  }
}
```

```
//assume P is a power of 2
reduce-tree(p, P, val) {
  fakeP = P;

  while (p < fakeP) {
    if (p >= fakeP/2) {
      send val to p-fakeP/2;
    }else {
      recv valp from p+fakeP/2;
      val += valp;
    }
    fakeP = fakeP / 2;
  }
}
```

Consider the following two network structures

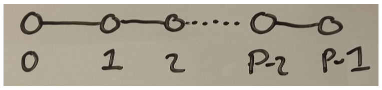

(a) chain

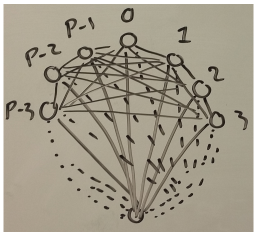

(b) clique

Question: Fill the following table. For each algorithm and each network structure, answer the following questions. Run a small example if you have difficulty seeing how communication happens; but express all answers for the case with P processors.

Case	Most loaded link (how much data)	Most loaded node (how much data)	Longest chain of communication (how long)
Reduce-star on chain			
Reduce-star on clique			
Reduce-chain on chain			
Reduce-chain on clique			
Reduce-tree on chain			
Reduce-tree on clique			

Question: What do you think is the best algorithm for each network structure? (One of the given algorithm or a different one.)

2 Heat Equation - 1D

One dimensional heat equation is the simplest example of a stencil computation. It computes iteratively the following equation for a stencil of size N.

$$Heat^k[0] = \frac{2Heat^{k-1}[0] + Heat^{k-1}[1]}{3}$$

$$Heat^k[N-1] = \frac{2Heat^{k-1}[N-1] + Heat^{k-1}[N-2]}{3}$$

$$Heat^k[i] = \frac{Heat^{k-1}[i-1] + Heat^{k-1}[i] + Heat^{k-1}[i+1]}{3}, \forall 0 < i < N-1$$

Consider the following partitionings of the data

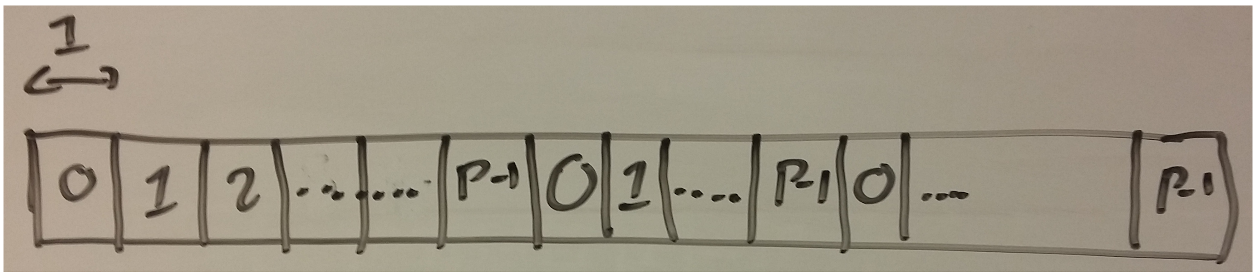

Figure 1: Round Robin

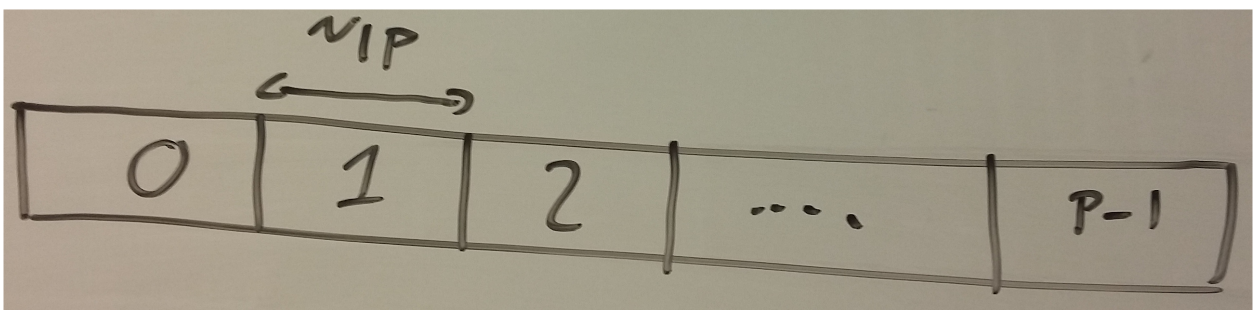

Figure 2: Block

(Assume network topology is a clique.)
Question: For each partitioning, write the algorithm that computes heat equation using this decomposition.
Question: For each partitioning, how much communication happen per iteration of the heat equation?
Question: What data partitioning would you use?

3 Dense Matrix Multiplication

Given a matrix A of size $N \times N$ and a vector x of size N, the value $y = Ax$ is given by $y[i] = \sum_j A[i][j]x[j]$. Or in other words, to compute $y[i]$ multiply element wise the ith row of the matrix by x and sum the values.

Consider the three data partitioning:

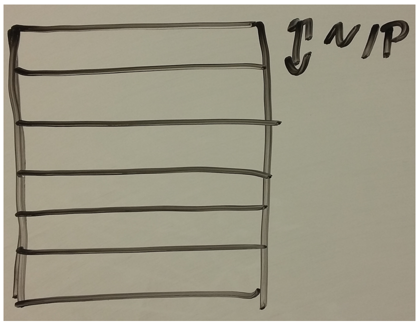

horizontal

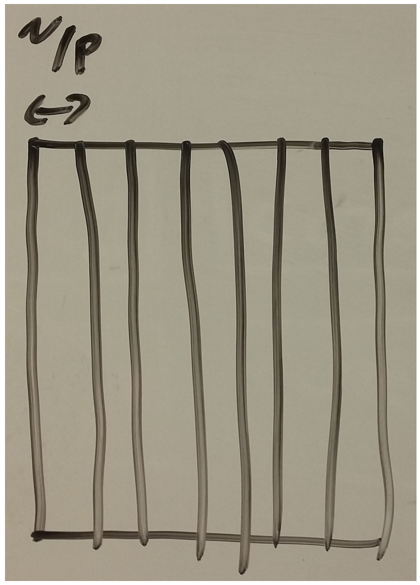

vertical

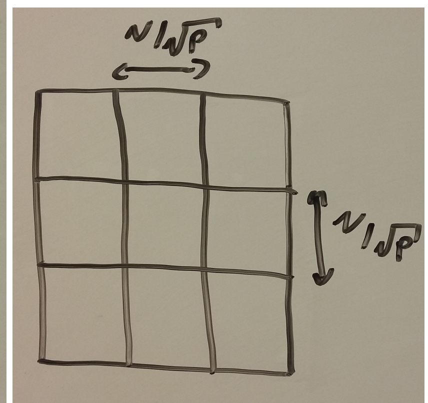

blocks

(Assume the network topology is a clique.)

For each data partitioning:

Question: Write the algorithm that performs $y = Ax; x = y$; 10 times in a loop.

Question: How much memory does each node need?

Question: How much communication does the algorithm do per iteration?