

Name: _____

Directions: **Work only on this sheet.** Put all your "fill the blank" answers in one place, say the lower-right part of this side, or on the back. Format:

```
#1a.  
x+y  
#1b  
if(u > v) w = 3;  
...
```

MAKE SURE TO COPY YOUR ANSWERS TO A SEPARATE SHEET FOR SENDING ME AN ELECTRONIC COPY LATER.

1. Below is an MPI version of the bucket sort in our OpenMP example (except that the bin boundaries are assumed known ahead of time, rather than calculated from a sample). The message-passing strategy is outline in the comments at the beginning of the code. Fill in the blanks.

```
1 // bucket sort, bin boundaries known in advance  
2  
3 // node 0 is manager, all else worker nodes; node 0 sends full data, bin  
4 // boundaries to all worker nodes; i-th worker node extracts data, for  
5 // bin i-1, sorts it, sends sorted chunk back to node 0; node 0 places  
6 // sorted results back in original array  
7  
8 // not claimed efficient; e.g. could be better to have manager place  
9 // items into bins  
10  
11 #include <mpi.h>  
12  
13 #define MAX_N 100000 // max size of original data array  
14 #define MAX_NPROCS 100 // max number of MPI processes  
15 #define DATA_MSG 0 // manager sending original data  
16 #define BDRIES_MSG 0 // manager sending bin boundaries  
17 #define CHUNKS_MSG 2 // workers sending their sorted chunks  
18  
19 int nnodes, //  
20 n, // size of full array  
21 me, // my node number  
22 fulldata[MAX_N],  
23 tmp[MAX_N],  
24 nbdries, // number of bin boundaries  
25 counts[MAX_NPROCS];  
26 float bdries[MAX_NPROCS-2]; // bin boundaries  
27  
28 int debug, debugme;  
29  
30 init(int argc, char **argv)  
31 {  
32     int i;  
33     debug = atoi(argv[3]);  
34     debugme = atoi(argv[4]);  
35     MPI_Init(&argc, &argv);  
36     MPI_Comm_size(MPI_COMM_WORLD, &nnodes);  
37     MPI_Comm_rank(MPI_COMM_WORLD, &me);  
38     nbdries = nnodes - 2;  
39     n = atoi(argv[1]);  
40     int k = atoi(argv[2]); // for random # gen  
41     // generate random data for test purposes  
42     for (i = 0; i < n; i++) fulldata[i] = rand() % k;  
43     // generate bin boundaries for test purposes  
44     for (i = 0; i < nbdries; i++) {  
45         bdries[i] = i * (k+1) / ((float) nnodes);  
46     }  
47 }  
48  
49 void managernode()  
50 {  
51     MPI_Status status;  
52     int i;  
53     int lenchunk; // length of a chunk received from a worker  
54     // send full data, bin boundaries to workers  
55     for (i = 1; i < nnodes; i++) {  
56         MPI_Send(BLANKa, BLANKb, MPI_INT, BLANKc, DATA_MSG, MPI_COMM_WORLD);  
57         MPI_Send(BLANKd, BLANKe, MPI_FLOAT, BLANKf, BDRIES_MSG, MPI_COMM_WORLD);
```

```
58     }  
59     // collect sorted chunks from workers, place them in their proper  
60     // positions within the original array  
61     int currposition = 0;  
62     for (i = 1; i < nnodes; i++) {  
63         MPI_Recv(tmp, MAX_N, MPI_INT, BLANKg, CHUNKS_MSG, MPI_COMM_WORLD, &status);  
64         MPI_Get_count(&status, MPI_INT, BLANKh);  
65         // memcpy(d,s,nb) copies nb bytes from s to d  
66         memcpy(BLANKi);  
67         BLANKj;  
68     }  
69     if (n < 25) {  
70         for (i = 0; i < n; i++) printf("%d ", fulldata[i]);  
71         printf("\n");  
72     }  
73 }  
74  
75 // adds xi to the part array, increments npart, the length of part  
76 void grab(int xi, int *part, int *npart)  
77 {  
78     part[*npart] = xi;  
79     *npart += 1;  
80 }  
81  
82 int cmpints(int *u, int *v)  
83 { if (*u < *v) return -1;  
84   if (*u > *v) return 1;  
85   return 0;  
86 }  
87  
88 void getandsortmychunk(int *tmp, int n, int *chunk, int *lenchunk)  
89 {  
90     int i, count = 0;  
91     int workernumber = me - 1;  
92     for (i = 0; i < n; i++) {  
93         if (workernumber == 0) {  
94             if (tmp[i] <= bdries[0]) grab(tmp[i], chunk, &count);  
95         }  
96         else if (workernumber < nbdries-1) {  
97             if (tmp[i] > bdries[workernumber-1] &&  
98                 tmp[i] <= bdries[workernumber]) grab(tmp[i], chunk, &count);  
99         }  
100        else if (tmp[i] > bdries[nbdries-1]) grab(tmp[i], chunk, &count);  
101    }  
102    qsort(chunk, count, sizeof(int), cmpints);  
103    *lenchunk = count;  
104 }  
105  
106 void workernode()  
107 {  
108     int n, fulldata[MAX_N], // size and storage of full data  
109     chunk[MAX_N],  
110     lenchunk,  
111     nbdries; // number of bin boundaries  
112     float bdries[MAX_NPROCS-1]; // bin boundaries  
113     MPI_Status status;  
114     MPI_Recv(fulldata, MAX_N, MPI_INT, BLANKk, DATA_MSG, MPI_COMM_WORLD, &status);  
115     MPI_Get_count(&status, MPI_INT, BLANKl);  
116     MPI_Recv(bdries, MAX_NPROCS-2, MPI_FLOAT, BLANKm, BDRIES_MSG,  
117             MPI_COMM_WORLD, &status);  
118     MPI_Get_count(&status, MPI_FLOAT, BLANKn);  
119     getandsortmychunk(fulldata, n, chunk, &lenchunk);  
120     MPI_Send(chunk, lenchunk, MPI_INT, BLANKo, CHUNKS_MSG, MPI_COMM_WORLD);  
121 }  
122  
123 int main(int argc, char **argv)  
124 {  
125     int i;  
126     init(argc, argv);  
127     if (me == 0) managernode();  
128     else workernode();  
129     MPI_Finalize();  
130 }  
131
```

Solutions:

```
1 // bucket sort, bin boundaries known in advance
2
3 // node 0 is manager, all else worker nodes; node 0 sends full data, bin
4 // boundaries to all worker nodes; i-th worker node extracts data for
5 // bin i-1, sorts it, sends sorted chunk back to node 0; node 0 places
6 // sorted results back in original array
7
8 // not claimed efficient; e.g. could be better to have manager place
9 // items into bins
10
11 #include <mpi.h>
12
13 #define MAX_N 100000 // max size of original data array
14 #define MAX_NPROCS 100 // max number of MPI processes
15 #define DATA_MSG 0 // manager sending original data
16 #define BDRIES_MSG 0 // manager sending bin boundaries
17 #define CHUNKS_MSG 2 // workers sending their sorted chunks
18
19 int nnodes, //
20     n, // size of full array
21     me, // my node number
22     fulldata[MAX_N],
23     tmp[MAX_N],
24     nbdries, // number of bin boundaries
25     counts[MAX_NPROCS];
26 float bdries[MAX_NPROCS-2]; // bin boundaries
27
28 int debug, debugme;
29
30 init(int argc, char **argv)
31 {
32     int i;
33     debug = atoi(argv[3]);
34     debugme = atoi(argv[4]);
35     MPI_Init(&argc, &argv);
36     MPI_Comm_size(MPI_COMM_WORLD, &nnodes);
37     MPI_Comm_rank(MPI_COMM_WORLD, &me);
38     nbdries = nnodes - 2;
39     n = atoi(argv[1]);
40     int k = atoi(argv[2]); // for random # gen
41     // generate random data for test purposes
42     for (i = 0; i < n; i++) fulldata[i] = rand() % k;
43     // generate bin boundaries for test purposes
44     for (i = 0; i < nbdries; i++) {
45         bdries[i] = i * (k+1) / ((float) nnodes);
46     }
47 }
48
49 void managernode()
50 {
51     MPI_Status status;
52     int i;
53     int lenchunk; // length of a chunk received from a worker
54     // send full data, bin boundaries to workers
55     for (i = 1; i < nnodes; i++) {
56         MPI_Send(fulldata, n, MPI_INT, i, DATA_MSG, MPI_COMM_WORLD);
57         MPI_Send(bdries, nbdries, MPI_FLOAT, i, BDRIES_MSG, MPI_COMM_WORLD);
58     }
59     // collect sorted chunks from workers, place them in their proper
60     // positions within the original array
61     int currposition = 0;
62     for (i = 1; i < nnodes; i++) {
63         MPI_Recv(tmp, MAX_N, MPI_INT, i, CHUNKS_MSG, MPI_COMM_WORLD, &status);
64         MPI_Get_count(&status, MPI_INT, &lenchunk);
65         memcpy(fulldata+currposition, tmp, lenchunk*sizeof(int));
66         currposition += lenchunk;
67     }
68     if (n < 25) {
69         for (i = 0; i < n; i++) printf("%d ", fulldata[i]);
70         printf("\n");
71     }
72 }
73
74 // adds xi to the part array, increments npart, the length of part
75 void grab(int xi, int *part, int *npart)
76 {
77     part[*npart] = xi;
78     *npart += 1;
79 }
80
81 int cmpints(int *u, int *v)
82 { if (*u < *v) return -1;
83   if (*u > *v) return 1;
```

```

84     return 0;
85 }
86
87 void getandsortmychunk(int *tmp, int n, int *chunk, int *lchunk)
88 {
89     int i,count = 0;
90     int workernumber = me - 1;
91     if (me == debugme) while (debug) ;
92     for (i = 0; i < n; i++) {
93         if (workernumber == 0) {
94             if (tmp[i] <= bdries[0]) grab(tmp[i],chunk,&count);
95         }
96         else if (workernumber < nbdries-1) {
97             if (tmp[i] > bdries[workernumber-1] &&
98                 tmp[i] <= bdries[workernumber]) grab(tmp[i],chunk,&count);
99         } else
100             if (tmp[i] > bdries[nbdries-1]) grab(tmp[i],chunk,&count);
101     }
102     qsort(chunk,count,sizeof(int),cmpints);
103     *lchunk = count;
104 }
105
106 void workernode()
107 {
108     int n,fulldata[MAX_N], // size and storage of full data
109         chunk[MAX_N],
110         lchunk,
111         nbdries; // number of bin boundaries
112     float bdries[MAX_NPROCS-1]; // bin boundaries
113     MPI_Status status;
114     MPI_Recv(fulldata,MAX_N,MPI_INT,0,DATA_MSG,MPI_COMM_WORLD,&status);
115     MPI_Get_count(&status,MPI_INT,&n);
116     MPI_Recv(bdries,MAX_NPROCS-2,MPI_FLOAT,0,BDRIES_MSG,MPI_COMM_WORLD,&status);
117     MPI_Get_count(&status,MPI_FLOAT,&nbdries);
118     getandsortmychunk(fulldata,n,chunk,&lchunk);
119     MPI_Send(chunk,lchunk,MPI_INT,0,CHUNKS_MSG,MPI_COMM_WORLD);
120 }
121
122 int main(int argc,char **argv)
123 {
124     int i;
125     init(argc,argv);
126     if (me == 0) managernode();
127     else workernode();
128     MPI_Finalize();
129 }
130

```