

# AN MPI IMPLEMENTATION OF A SELF-SUBMITTING PARALLEL JOB QUEUE

JOHN M. NEUBERGER, NÁNDOR SIEBEN, AND JAMES W. SWIFT

## MATRIX SQUARE EXAMPLE: `square.cpp`

- line 4: There are three types of jobs in this example. The job types need to be positive integers, so `NONE` is added to take the unused value of zero.
- line 6: The variable `matrix` contains the input matrix. The variable `square` contains the square of the input matrix, which is the result of the program.
- line 7–10: This data type is used to send back one row of the result matrix, together with the position of this row.
- line 12–16: The BSL requires a simple template function for the serialization of each structure.

### MPQswitch:

- line 18: This function is executed by the workers with the `job.type` variable set to `DATA` or `MULTIPLY`. It is also executed by the boss with `job.type` set to `RESULT`.
- line 19: The variable `data` contains one row of the result produced by a worker.
- line 20: Decide the type of the current job.
- lines 21–23: The input matrix is shared by all the workers. Each worker receives the matrix as a `DATA` job. The workers deserialize it and store it locally in the variable `matrix`.
- line 24–29: A worker calculates one row of the goal matrix.
- line 25: The input is serialized into `data.pos`. It tells the worker which row to compute.
- line 26: The calculated row is stored in `data.result`.
- lines 27–29: The row is calculated using standard matrix multiplication.
- lines 30–31: The worker sends the calculated row to the boss.
- line 33–36: The boss receives a row and puts it into the result matrix `square`.
- line 34: The row is serialized.
- line 35: The row is placed at the appropriate location.

### main:

- line 41: The MPI is initialized.
- line 42: The nodes are split into one boss and several workers.
- line 43: The example input matrix contains rows of length 10 containing a 1 at each position.
- line 44: The input variable `matrix` is initialized.
- line 45: The output variable `square` is resized to the correct dimensions.
- line 46: The input matrix is sent to all the workers
- lines 48–49: The job queue is filled with `MULTIPLY` jobs, each requesting the calculation of one row of the goal matrix.
- line 50: The boss starts the supervision of the workers. At the end of this work, the goal matrix `square` will have been calculated. No other result is created, so `outqueue` is empty.
- lines 51–55: The goal matrix is printed.
- line 56: The boss halts all the workers.

DEPARTMENT OF MATHEMATICS AND STATISTICS, NORTHERN ARIZONA UNIVERSITY, FLAGSTAFF, AZ 86011-5717, USA  
*E-mail address:* `john.neuberger@nau.edu`, `nandor.sieben@nau.edu`, `jim.swift@nau.edu`

```

1 #include "MPQueue.h"
2 using namespace std;
3
4 enum { NONE, DATA, MULTIPLY, RESULT };
5 typedef vector< int > Trow;
6 vector< Trow > matrix, square;
7 typedef struct {
8     int pos;                                // row index
9     Trow result;                            // output row
10 } Tdata;
11
12 template < class Archive > void           // needed to serialize Tdata
13 serialize (Archive & ar, Tdata & data, const unsigned int version) {
14     ar & data.pos;
15     ar & data.result;
16 }
17
18 void MPQswitch (Tjob & job) {
19     Tdata data;
20     switch (job.type) {
21     case DATA:                                // receive the input matrix
22         from_string (matrix, job.data);
23         break;
24     case MULTIPLY:
25         from_string (data.pos, job.data);      // get the row position
26         data.result = Trow (matrix.size (), 0);
27         for (int j = 0; j < matrix.size (); j++) // calculate one row
28             for (int k = 0; k < matrix.size (); k++)
29                 data.result[j] += matrix[data.pos][k] * matrix[k][j];
30         job = Tjob(RESULT, data);              // prepare the result
31         MPQtask (job);                      // send result to boss
32         break;
33     case RESULT:                                // receive one output row
34         from_string (data, job.data);
35         square[data.pos] = data.result;
36         job.data = "";                         // nothing to return
37     }
38 }
39
40 int main (int argc, char *argv[]) {
41     MPQinit (argc, argv);
42     MPQstart ();
43     Trow row (10, 1);                         // input matrix containing
44     vector< Trow > mat (row.size (), row); // 1 at every entry
45     square.resize (row.size ());               // container for the output
46     MPQsharedata (Tjob(DATA, mat));          // input matrix sent to workers
47     Tjobqueue inqueue, outqueue;
48     for (int i = 0; i < mat.size (); i++) // every row is a separate job
49         inqueue.push (Tjob(MULTIPLY, i));
50     MPQrunjobs (inqueue, outqueue);          // run the jobs
51     for (int i = 0; i < row.size (); i++) { // print the output matrix
52         for (int j = 0; j < row.size (); j++)
53             cout << square[i][j] << " ";
54         cout << "\n";
55     }
56     MPQstop ();
57 }
```