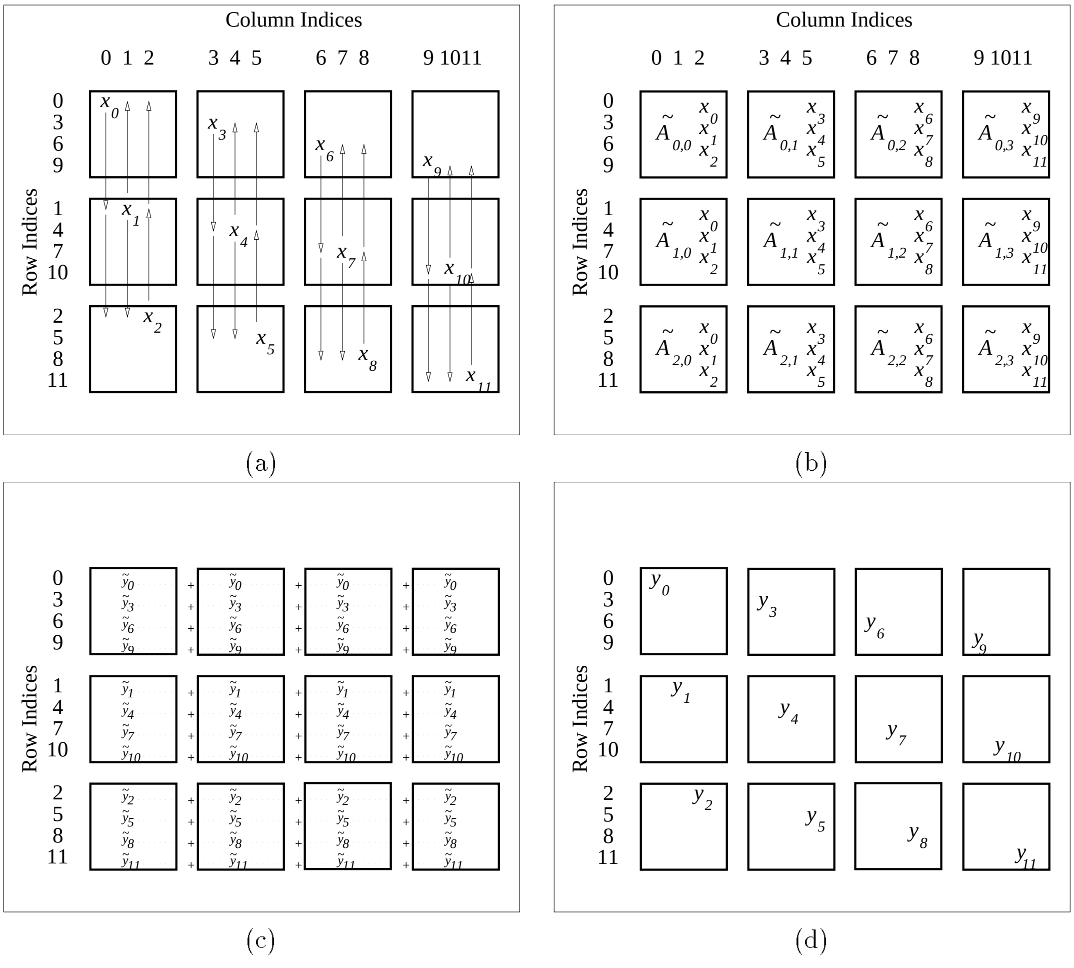

Figure 1.5
Parallel matrix-vector multiplication. (a) Sub-vectors of x start distributed among nodes in column-major order. (b) After spreading x within columns, local matrix-vector multiplication can commence. (c) Each node holds a partial contribution to the results that need to be summed within rows of nodes. (d) Sub-vectors of result y finish distributed among nodes in column-major order.