

РЕАЛИЗАЦИИ АЛГОРИТМОВ КОЛЛЕКТИВНЫХ ОПЕРАЦИЙ В БИБЛИОТЕКАХ СТАНДАРТА MPI

| № | Операция | MPICH 3.3 | | Open MPI 4.0 | |
|---|-----------------|---|--|---|---|
| | | Blocking | Non-blocking | Blocking (base) | Non-blocking (libnbc) |
| 1 | MPI_IALLGATHER | 1. inter: local gather remote bcast 2. intra: bruck's 3. intra: recursive doubling 4. intra: ring | 1. inter: local gather remote bcast 2. intra: bruck's 3. intra: recursive doubling 4. intra: ring | 1. intra: linear 2. intra: bruck's 3. intra: recursive doubling 4. intra: ring 5. intra: neighbor exchange 6. intra: two procs | 1. inter: linear 2. intra: linear 3. intra: recursive doubling |
| 2 | MPI_IALLGATHERV | 1. inter: remote gather local bcast 2. intra: bruck's 3. intra: recursive doubling 4. intra: ring | 1. inter: local gather remote bcast 2. intra: bruck's 3. intra: recursive doubling 4. intra: ring pipeline | 1. intra: linear 2. intra: bruck's 3. intra: ring 4. intra: neighbor exchange 5. intra: two procs | 1. inter: linear 2. intra: linear |
| 3 | MPI_IALLREDUCE | 1. inter: reduce exchange bcast 2. intra: recursive doubling 3. intra: rabenseifner's | 1. inter: remote reduce local bcast 2. intra: naive 3. intra: recursive doubling 4. intra: rabenseifner's | 1. intra: linear 2. intra: non overlapping 3. intra: recursive doubling 4. intra: ring 5. intra: ring segmented 6. intra: rabenseifner's | 1. inter: linear 2. intra: linear 3. intra: dissemination 4. intra: recursive doubling 5. intra: ring segmented 6. intra: rabenseifner's |
| 4 | MPI_IALLTOALL | 1. inter: pairwise exchange 2. intra: bruck's 3. intra: pairwise exchange 4. intra: pairwise sendrecv replace 5. intra: scattered | 1. inter: pairwise exchange 2. intra: bruck's 3. intra: pairwise exchange 4. intra: batched sendrecv | 1. intra: linear 2. intra: linear sync 3. intra: inplace 4. intra: bruck's 5. intra: pairwise exchange 6. intra: two procs | 1. inter: linear 2. intra: linear 3. intra: dissemination 4. intra: pairwise exchange 5. intra: inplace |
| 5 | MPI_IALLTOALLV | 1. inter: pairwise exchange 2. intra: pairwise sendrecv replace 3. intra: scattered | 1. inter: pairwise exchange 2. intra: blocked 3. intra: inplace | 1. intra: linear 2. intra: inplace 3. intra: pairwise exchange | 1. inter: linear 2. intra: linear 3. intra: pairwise exchange 4. intra: inplace |
| 6 | MPI_IALLTOALLW | 1. inter: pairwise exchange 2. intra: pairwise sendrecv replace 3. intra: scattered | 1. inter: pairwise exchange 2. intra: blocked 3. intra: inplace | - | 1. inter: linear 2. intra: linear 3. intra: pairwise exchange 4. intra: inplace |
| 7 | MPI_IBARRIER | 1. inter: bcast 2. intra: dissemination | 1. inter: bcast 2. intra: recursive doubling | 1. intra: linear 2. intra: double ring 3. intra: recursive doubling 4. intra: bruck's 5. intra: two procs 6. intra: tree (another recursive doubling) | 1. inter: dissemination (?) 2. intra: dissemination |
| 8 | MPI_IBCAST | 1. inter: local bcast 2. intra: binomial 3. intra: scatter recursive doubling allgather 4. intra: scatter ring allgather | 1. inter: flat 2. intra: binomial 3. intra: scatter recursive doubling allgather 4. intra: scatter ring allgather | 1. intra: linear 2. intra: generic 3. intra: bintree 4. intra: split bintree 5. intra: pipeline 6. intra: chain 7. intra: binomial 8. intra: knomial 9. intra: scatter allgather (binomial tree scatter followed by a recursive doubling allgather) 10. intra: scatter allgather ring (binomial tree scatter followed by a ring allgather) | 1. inter: (?) 2. intra: linear 3. intra: chain 4. intra: binomial 5. intra: knomial |
| 9 | MPI_IEXSCAN | 1. intra: recursive doubling | 1. intra: recursive doubling | 1. intra: linear 2. intra: recursive doubling | 1. intra: linear 2. intra: recursive doubling |

| | | | | | |
|----|---------------------------|--|---|--|---|
| 10 | MPI_GATHER | 1. inter: linear 2. inter: local gather remote send 3. intra: binomial | 1. inter: linear (for long messages) 2. inter: linear (for short messages) 3. intra: binomial | 1. intra: linear 2. intra: linear sync 3. intra: binomial | 1. inter: linear 2. intra: linear |
| 11 | MPI_GATHERV | 1. allcomm: linear | 1. allcomm: linear | - | 1. inter: linear 2. intra: linear |
| 12 | MPI_IREDUCE | 1. inter: local reduce remote send 2. intra: binomial 3. intra: reduce scatter gather | 1. inter: local reduce remote send 2. intra: binomial 3. intra: reduce scatter gather | 1. generic 2. intra: linear 3. intra: chain 4. intra: pipeline 5. intra: binary 6. intra: binomial 7. intra: in order binary 8. intra: rabenseifner's | 1. inter: linear 2. intra: linear 3. intra: chain 4. intra: binomial 5. intra: rabenseifner's |
| 13 | MPI_IREDUCE_SCATTER | 1. inter: remote reduce local scatter 2. intra: noncommutative recursive doubling (butterfly) 3. intra: pairwise exchange 4. intra: recursive doubling 5. intra: recursive halving | 1. inter: remote reduce local scatterv 2. intra: noncommutative recursive doubling (butterfly) 3. intra: pairwise exchange 4. intra: recursive doubling 5. intra: recursive halving | 1. intra: non overlapping 2. intra: recursive halving 3. intra: ring 4. intra: butterfly | 1. inter: (?) 2. intra: pairwise exchange (?) |
| 14 | MPI_IREDUCE_SCATTER_BLOCK | 1. inter: remote reduce local scatter 2. intra: noncommutative recursive doubling (butterfly) 3. intra: pairwise exchange 4. intra: recursive doubling 5. intra: recursive halving | 1. inter: remote reduce local scatterv 2. intra: noncommutative recursive doubling (butterfly) 3. intra: pairwise exchange 4. intra: recursive doubling 5. intra: recursive halving | 1. intra: linear 2. intra: recursive doubling 3. intra: recursive halving 4. intra: butterfly 5. intra: butterfly pof2 | 1. inter: (?) 2. intra: pairwise exchange (?) |
| 15 | MPI_JSCAN | 1. intra: recursive doubling | 1. intra: recursive doubling | 1. intra: linear 2. intra: recursive doubling | 1. intra: linear 2. intra: recursive doubling |
| 16 | MPI_ISCATTER | 1. inter: linear 2. inter: remote send local scatter 3. intra: binomial | 1. inter: linear 2. inter: remote send local scatter 3. intra: binomial | 1. intra: linear 2. intra: binomial | 1. inter: linear 2. intra: linear |
| 17 | MPI_ISCATTERV | 1. allcomm: linear | 1. allcomm: linear | - | 1. inter: linear 2. intra: linear |
| 18 | MPI_INEIGHBOR_ALLGATHER | | 1. allcomm: linear | - | 1. allcomm: linear |
| 19 | MPI_INEIGHBOR_ALLGATHERV | | 1. allcomm: linear | - | 1. allcomm: linear |
| 20 | MPI_INEIGHBOR_ALLTOALL | | 1. allcomm: linear | - | 1. allcomm: linear |
| 21 | MPI_INEIGHBOR_ALLTOALLV | | 1. allcomm: linear | - | 1. allcomm: linear |
| 22 | MPI_INEIGHBOR_ALLTOALLW | | 1. allcomm: linear | - | 1. allcomm: linear |

Intel MPI algorithms: <https://software.intel.com/en-us/mpi-developer-reference-windows-i-mpi-adjust-family-environment-variables>
 MPICH LibNBC sources: <https://github.com/pmodels/mpich/tree/master/src/mpid/common/sched>

© 2019 Александр Аненков
anekov.ru@gmail.com