

**CPM: A software tool for Communication Performance Modelling**  
Version 1.2.0 (Revision 308M)

Generated by Doxygen 1.7.1

Wed Jun 8 2011 18:06:19

## Contents

|          |  |           |
|----------|--|-----------|
| <b>1</b> | <b>Introduction</b>  | <b>1</b>  |
| 1.1      | Authors . . . . .  | 2         |
| <b>2</b> | <b>Installation</b>  | <b>3</b>  |
| <b>3</b> | <b>The software design</b>   | <b>4</b>  |
| <b>4</b> | <b>Communication performance models</b>                                      | <b>5</b>  |
| 4.1      | Library . . . . .  | 5         |
| 4.2      | Tools . . . . .  | 6         |
| 4.3      | Tests . . . . .  | 6         |
| 4.4      | Model builder . . . . .  | 6         |
| 4.4.1    | Collective benchmark (using models) . . . . .                                | 7         |
| 4.5      | Predictor of the execution time of p2p or collective communication . . . . . | 7         |
| 4.6      | Generator of hierarchy files . . . . .                                       | 7         |
| 4.7      | Hierarchical model builder . . . . .   | 8         |
| 4.7.1    | Collective benchmark (using models) . . . . .                                | 8         |
| 4.8      | Generator of rank-to-host mapping . . . . .                                  | 8         |
| <b>5</b> | <b>Model-based collectives</b>   | <b>8</b>  |
| 5.1      | Tools . . . . .  | 10        |
| 5.2      | Tests . . . . .  | 11        |
| 5.3      | Test for model-based collectives . . . . .                                   | 11        |
| 5.4      | Parallel matrix multiplication . . . . .                                     | 12        |
| <b>6</b> | <b>Module Documentation</b>  | <b>12</b> |
| 6.1      | Measurement: benchmarking specific communication experiments . . . . .       | 12        |
| 6.1.1    | Detailed Description . . . . .   | 12        |
| 6.1.2    | Function Documentation . . . . .   | 12        |
| 6.2      | Communication performance models . . . . .                                   | 13        |
| 6.3      | Prediction of communication time . . . . .                                   | 14        |
| 6.3.1    | Detailed Description . . . . .   | 14        |
| 6.3.2    | Function Documentation . . . . .   | 14        |
| 6.4      | The heterogeneous Hockney model . . . . .                                    | 15        |
| 6.4.1    | Detailed Description . . . . .   | 16        |
| 6.4.2    | Function Documentation . . . . .   | 16        |
| 6.5      | The linear interpolation model . . . . .                                     | 18        |

|          |  |           |
|----------|--|-----------|
| 6.5.1    | Function Documentation . . . . .   | 19        |
| 6.6      | LMO: an advanced heterogeneous communication performance model . . . . . | 19        |
| 6.6.1    | Detailed Description . . . . .   | 20        |
| 6.6.2    | Define Documentation . . . . .   | 20        |
| 6.6.3    | Function Documentation . . . . .   | 20        |
| 6.7      | The heterogeneous PLogP model . . . . .                                  | 23        |
| 6.7.1    | Detailed Description . . . . .   | 24        |
| 6.7.2    | Function Documentation . . . . .   | 24        |
| 6.8      | Generic model-based collectives . . . . .                                | 26        |
| 6.8.1    | Detailed Description . . . . .   | 27        |
| 6.8.2    | Function Documentation . . . . .   | 28        |
| 6.9      | Hockney-based collective operations . . . . .                            | 31        |
| 6.9.1    | Detailed Description . . . . .   | 33        |
| 6.9.2    | Function Documentation . . . . .   | 33        |
| 6.10     | LinInterp-based collective operations . . . . .                          | 38        |
| 6.10.1   | Detailed Description . . . . .   | 40        |
| 6.10.2   | Function Documentation . . . . .   | 40        |
| 6.11     | LMO-based collective operations . . . . .                                | 46        |
| 6.11.1   | Function Documentation . . . . .   | 46        |
| 6.12     | LogGP-based collective operations . . . . .                              | 46        |
| 6.12.1   | Function Documentation . . . . .   | 48        |
| 6.13     | PLogP-based collective operations . . . . .                              | 53        |
| 6.13.1   | Function Documentation . . . . .   | 55        |
| <b>7</b> | <b>Namespace Documentation</b>   | <b>61</b> |
| 7.1      | CPM::BRSG Namespace Reference . . . . .                                  | 61        |
| 7.1.1    | Detailed Description . . . . .   | 61        |
| 7.2      | CPM::hierarchy Namespace Reference . . . . .                             | 61        |
| 7.2.1    | Detailed Description . . . . .   | 61        |
| 7.3      | CPM::SGv Namespace Reference . . . . .                                   | 61        |
| 7.3.1    | Detailed Description . . . . .   | 61        |
| <b>8</b> | <b>Class Documentation</b>   | <b>62</b> |
| 8.1      | CPM::BRSG::BFS_binomial_builder Class Reference . . . . .                | 62        |
| 8.1.1    | Detailed Description . . . . .   | 62        |
| 8.2      | CPM::SGv::BFS_binomial_builder Class Reference . . . . .                 | 62        |
| 8.2.1    | Detailed Description . . . . .   | 63        |

|        |   |    |
|--------|---|----|
| 8.3    | CPM_predictor Struct Reference . . . . .                  | 63 |
| 8.3.1  | Detailed Description . . . . .                            | 63 |
| 8.3.2  | Member Data Documentation . . . . .                       | 63 |
| 8.4    | CPM_SGv_sorter Struct Reference . . . . .                 | 64 |
| 8.4.1  | Detailed Description . . . . .                            | 64 |
| 8.5    | CPM_triplet Struct Reference . . . . .                    | 64 |
| 8.5.1  | Detailed Description . . . . .                            | 65 |
| 8.5.2  | Member Data Documentation . . . . .                       | 65 |
| 8.6    | CPM_triplets Struct Reference . . . . .                   | 65 |
| 8.6.1  | Detailed Description . . . . .                            | 66 |
| 8.6.2  | Member Data Documentation . . . . .                       | 66 |
| 8.7    | CPM::SGv::DFS_binomial_builder Class Reference . . . . .  | 66 |
| 8.7.1  | Detailed Description . . . . .                            | 66 |
| 8.8    | CPM::BRSG::DFS_binomial_builder Class Reference . . . . . | 67 |
| 8.8.1  | Detailed Description . . . . .                            | 67 |
| 8.9    | H_Model< p2p_model > Class Template Reference . . . . .   | 67 |
| 8.9.1  | Detailed Description . . . . .                            | 67 |
| 8.9.2  | Member Function Documentation . . . . .                   | 68 |
| 8.10   | Hockney_model Struct Reference . . . . .                  | 68 |
| 8.10.1 | Detailed Description . . . . .                            | 69 |
| 8.10.2 | Member Data Documentation . . . . .                       | 69 |
| 8.11   | hockney_p2p_model Class Reference . . . . .               | 69 |
| 8.11.1 | Detailed Description . . . . .                            | 69 |
| 8.11.2 | Member Function Documentation . . . . .                   | 70 |
| 8.12   | LinInterp_model Struct Reference . . . . .                | 70 |
| 8.12.1 | Detailed Description . . . . .                            | 71 |
| 8.12.2 | Member Data Documentation . . . . .                       | 71 |
| 8.13   | LMO_model Struct Reference . . . . .                      | 72 |
| 8.13.1 | Detailed Description . . . . .                            | 72 |
| 8.13.2 | Member Data Documentation . . . . .                       | 72 |
| 8.14   | PLogP_model Struct Reference . . . . .                    | 74 |
| 8.14.1 | Detailed Description . . . . .                            | 74 |
| 8.14.2 | Member Data Documentation . . . . .                       | 74 |
| 8.15   | plogp_p2p_model Class Reference . . . . .                 | 75 |
| 8.15.1 | Detailed Description . . . . .                            | 75 |
| 8.15.2 | Member Function Documentation . . . . .                   | 75 |

|  |           |
|--|-----------|
| 8.16 CPM::second_greater Struct Reference . . . . .            | 75        |
| 8.16.1 Detailed Description . . . . .                          | 75        |
| 8.17 CPM::second_less Struct Reference . . . . .               | 75        |
| 8.17.1 Detailed Description . . . . .                          | 75        |
| 8.18 CPM::SGv::Traff_builder Class Reference . . . . .         | 76        |
| 8.18.1 Detailed Description . . . . .                          | 76        |
| 8.18.2 Member Function Documentation . . . . .                 | 76        |
| 8.19 TreePredictor Class Reference . . . . .                   | 77        |
| 8.19.1 Detailed Description . . . . .                          | 77        |
| 8.20 CPM::BRSG::UCS_binomial_builder Class Reference . . . . . | 78        |
| 8.20.1 Detailed Description . . . . .                          | 78        |
| 8.21 CPM::SGv::UCS_binomial_builder Class Reference . . . . .  | 78        |
| 8.21.1 Detailed Description . . . . .                          | 79        |
| <b>9 File Documentation</b>                                    | <b>79</b> |
| 9.1 models/cpm.h File Reference . . . . .                      | 79        |
| 9.1.1 Detailed Description . . . . .                           | 80        |

## 1 Introduction

Traditionally, communication performance models for high performance computing are analytical and built for homogeneous clusters. The basis of these models is a point-to-point communication model characterized by a set of integral parameters, having the same value for each pair of processors. The execution time of other operations (which are, in fact, collective), is expressed as a combination of the point-to-point parameters, and is analytically predicted for different message sizes and numbers of processors involved. The core of this approach is the choice of such a point-to-point model that is the most appropriate to the targeted platform, allowing for easy and natural expression of different algorithms of collective operations. For homogeneous clusters, the point-to-point parameters are found statistically from communication experiments between any two processors. Typical experiments include sending and receiving messages of different sizes, with the communication execution time being measured on one side.

A homogeneous communication model can be applied to a cluster of heterogeneous processors by averaging values obtained for every pair of processors. In this case, the heterogeneous cluster will be treated as homogeneous in terms of the performance of communication operations. If some processors or links in the heterogeneous cluster significantly differ in performance, predictions based on the homogeneous communication model may become inaccurate. More accurate performance models would not average the point-to-point communication parameters. The use of such heterogeneous communication models in model-based optimization of MPI collective operations on heterogeneous clusters do improve their performance.

The traditional models use a small number of parameters to describe communication between any two processors. The number of these parameters and their use in the model are always defined in a way that allows for their accurate estimation with a set of point-to-point communication experiments between these two processors. The price to pay is that such a traditional point-to-point communication model is not intuitive. The meaning of its parameters is not clear. Different sources of the contribution into the execution time are artificially and non-intuitively mixed and spread over a smaller number of parameters. This makes

the models difficult to use for accurate modelling of collective communications.

The alternative approach is to use original point-to-point heterogeneous models that allow for easy and intuitive expression of the execution time of collective communication operations such as this model designed for switched heterogeneous clusters. While more accurate, this heterogeneous model has a significantly larger number of parameters. This will result in a higher cost of their estimation. In particular, when applied to the heterogeneous communication model, the statistical methods of finding the point-to-point parameters, traditionally used in the case of homogeneous communication models, will require a significantly larger number of measurements. For our target architecture, which is a heterogeneous cluster based on a switched network, we can address this problem by performing most of the communication experiments in parallel, using the fact that the network switches provide no-contention point-to-point communications, appropriately forwarding packets between sources and destinations.

We present the software tool that automates the estimation of the heterogeneous communication performance models of clusters based on a switched network [8]. The software tool can also be used in the high-level model-based optimization of MPI collective operations. This is particularly important for heterogeneous platforms where the users typically have neither authority nor knowledge for making changes in hardware or basic software settings.

## 1.1 Authors

Alexey Lastovetsky, Vladimir Rychkov, Maureen O’Flynn, Kiril Dichev

Heterogeneous Computing Laboratory

School of Computer Science and Informatics, University College Dublin

Belfield, Dublin 4, Ireland

<http://hcl.ucd.ie>

{alexey.lastovetsky, vladimir.rychkov}@ucd.ie

## References

- [1] Albert Alexandrov, Mihai F. Ionescu, Klaus E. Schauser, and Chris Scheiman. LogGP: incorporating long messages into the LogP model – one step closer towards a realistic model for parallel computation. In *SPAA ’95: Proceedings of the seventh annual ACM symposium on Parallel algorithms and architectures*, pages 95–105, New York, NY, USA, 1995. ACM. 24
- [2] Prashanth B. Bhat, C. S. Raghavendra, and Viktor K. Prasanna. Efficient collective communication in distributed heterogeneous systems. *Journal of Parallel and Distributed Computing*, 63(3):251 – 263, 2003. 27, 33, 40
- [3] E.W. Chan, M.F. Heimlich, A. Purkayastha, and R.A. van de Geijn. On optimizing collective communication. *Cluster Computing, 2004 IEEE International Conference on*, pages 145–155, Sept. 2004. 27
- [4] David Culler, Richard Karp, David Patterson, Abhijit Sahay, Klaus Erik Schauser, Eunice Santos, Ramesh Subramonian, and Thorsten von Eicken. LogP: towards a realistic model of parallel computation. *SIGPLAN Not.*, 28(7):1–12, 1993. 24
- [5] J. Hatta and S. Shibusawa. Scheduling algorithms for efficient gather operations in distributed heterogeneous systems. *Parallel Processing, 2000. Proceedings. 2000 International Workshops on*, pages 173–180, 2000. 27, 33, 40

- [6] Roger W. Hockney. The communication challenge for MPP: Intel Paragon and Meiko CS-2. *Parallel Comput.*, 20(3):389–398, 1994. [13](#), [16](#)
- [7] Thilo Kielmann, Henri E. Bal, and Kees Verstoep. Fast measurement of LogP parameters for message passing platforms. In *IPDPS '00: Proceedings of the 15 IPDPS 2000 Workshops on Parallel and Distributed Processing*, pages 1176–1183, London, UK, 2000. Springer-Verlag. [14](#), [24](#)
- [8] A. Lastovetsky, V. Rychkov, and M. OFlynn. A software tool for accurate estimation of parameters of heterogeneous communication models. In A. Lastovetsky, T. Kechadi, and J. Dongarra, editors, *Recent Advances in Parallel Virtual Machine and Message Passing Interface: 15th European PVM/MPI User's Group Meeting*, volume 5205, pages 43–54, Dublin, Ireland, September 7-10 2008. Springer-Verlag Berlin Heidelberg. [2](#), [14](#)
- [9] Rajeev Thakur, Rolf Rabenseifner, and William Gropp. Optimization of collective communication operations in MPICH. *International Journal of High Performance Computing Applications*, 19(1):49–66, 2005. [27](#)

## 2 Installation

Installation  
=====

Included software (configured, built and installed together with CPM):

1. MPIBlib (MPI Benchmark library) - required (for benchmarking MPI communication operations)
2. logp\_mpi (The MPI LogP Benchmark, version 1.4) - required (for the PLogP model)

Required software:

1. any C/C++ and MPI (MPICH-1 does not support shared libraries)
2. GSL (GNU Scientific Library, version 1.11)
3. Boost (The Boost C++ libraries: Graph, MPI)

Optional software:

1. R (The R Project for Statistical Computing, version 2.6.1) - optional (for estimation of the LMO threshold parameters)
2. Gnuplot (An Interactive Plotting Program) - optional (for performance diagrams)
3. Graphviz (Graph Visualization Software: dot) - optional (for tree visualization)

GSL

If GSL is installed in a non-default directory  
\$ export LD\_LIBRARY\_PATH=DIR/lib:\$LD\_LIBRARY\_PATH

Boost

1. Boost should be configured with at least Graph, MPI and Serialization library.

bjam is the configuration script used for Boost.

2. The default is no component to be built. You can do:  
\$ ./bjam --with-<component> ... install  
to configure, compile and install the specified components.
3. Default installation directory is DIR=/usr/local
4. Default installation:

- DIR/include/boost
- DIR/lib/libboost\_library\_versions.\*

Create symbolic links:

```
$ cd DIR/include; ln -s boost_version/boost
$ cd DIR/lib; ln -s libboost_[library]_[version].[a/so] libboost_[library].[a/so]
```

```
$ export LD_LIBRARY_PATH=DIR/lib:$LD_LIBRARY_PATH
```

```

R
1. R should be configured as a shared library
$ ./configure --prefix=DIR --enable-R-shlib=yes
$ make install
2. Set up environment
$ export R_HOME=DIR/lib/R
3. Install required packages
$ DIR/bin/R
> install.packages(c("sandwich", "strucchange", "zoo"))
4. If R is installed in a non-default directory
$ export LD_LIBRARY_PATH=$R_HOME/lib:$LD_LIBRARY_PATH

For users
-----

Download the latest package from http://hcl.ucd.ie/project/cpm

$ tar -zxvf cpm-X.X.X.tar.gz
$ cd cpm-X.X.X
$ ./configure
$ make all install

Configuration
-----

Packages:
  --with-gsl-dir=DIR      GNU Scientific Library directory
  --with-boost-dir=DIR   The Boost C++ libraries directory
  --with-r-dir=DIR       The R Project for Statistical Computing directory

Check configure options:
$ ./configure -h

For developers
-----

Required software:
1. Subversion
2. GNU autotools
3. Doxygen, Graphviz and any TeX - optional (for reference manual)

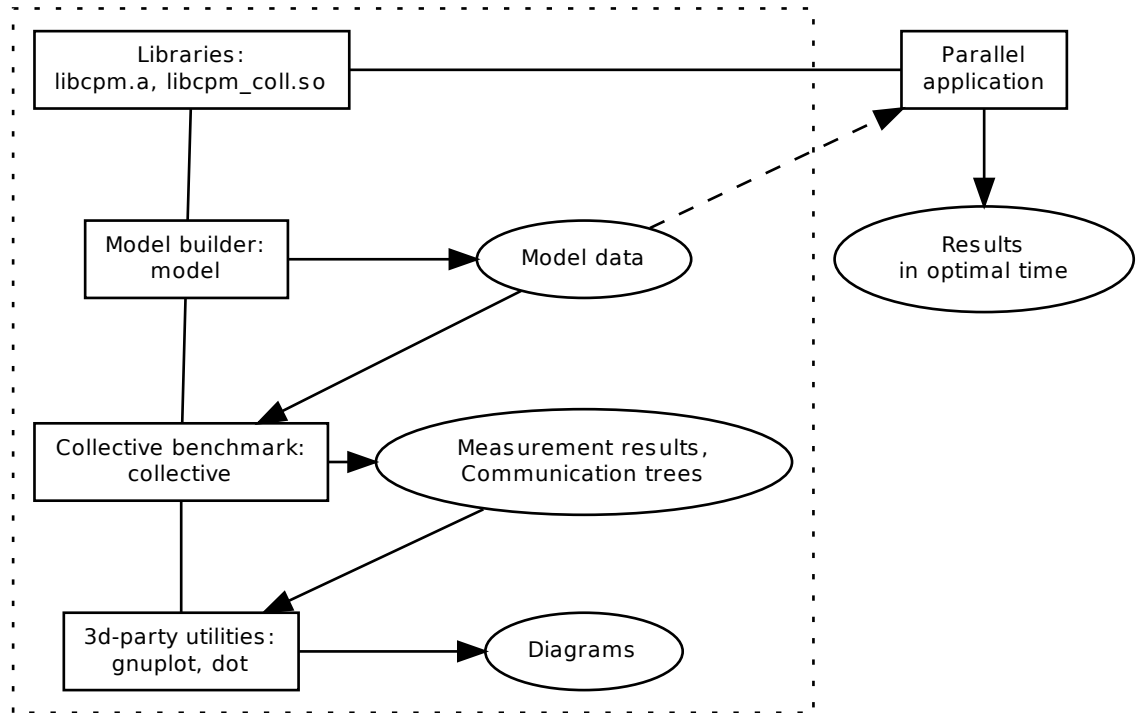
$ svn co https://hcl.ucd.ie/repos/CPM/trunk CPM
$ cd CPM
$ autoreconf -i
$ ./configure --enable-debug
$ make all

To create a package:
$ svn log -v > ChangeLog
$ make dist

```

### 3 The software design

CPM is implemented in C/C++ on top of MPI. The package consists of libraries, tools and tests. The libraries implements heterogeneous communication performance models and model-based collectives. The tools estimates the parameters of the models and evaluates the performance of the model-based collective communication operations.

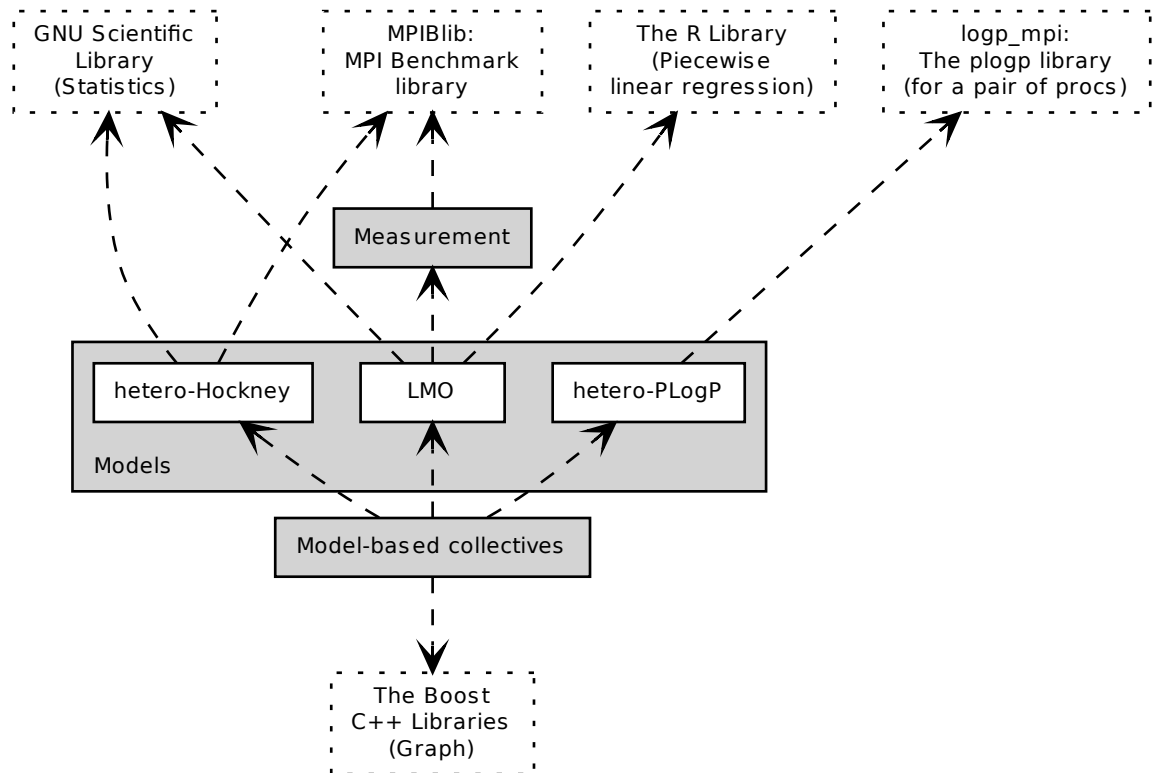


## 4 Communication performance models

### 4.1 Library

A static library `libcpm.a` implements heterogeneous communication performance models and consists of the following modules:

- [Measurement: benchmarking specific communication experiments](#)
- [Communication performance models](#)
- [Predictor of the execution time of p2p or collective communication](#)



## 4.2 Tools

- [Model builder](#)
- [Predictor of the execution time of p2p or collective communication](#)
- [Generator of hierarchy files](#)
- [Hierarchical model builder](#)

## 4.3 Tests

- [Generator of rank-to-host mapping](#)

## 4.4 Model builder

Builds the Hockney, PLogP and LMO models and estimates the execution time of point-to-point and collective communication operations. Arguments are described in the MPIBlib manual.

**Example:**

```
$ mpirun -n 16 model -C Hockney -o model.out
```

The model.out file can then be used by either of these tools:

- [Predictor of the execution time of p2p or collective communication](#)
- [Collective benchmark \(using models\)](#)

#### 4.4.1 Collective benchmark (using models)

The basic usage of this tool and a description how to plot the output file is described in MPIBlib.

**Example** using a model-based collective:

```
$ mpirun -n 16 collective -O Hockney_Scatter_dfs_binomial_min -o model=Hockney, file=model.out -l libcpm_coll.so -m 1024 -M 2049 > collective.out
```

will benchmark a binomial tree scatter operation with message sizes per process of 1024 and 2048. The binomial tree will be generated through depth-first approach choosing the node with minimum predicted communication by the Hockney model.

## 4.5 Predictor of the execution time of p2p or collective communication

Instead of performing p2p communication or model-based collective communication, this routine performs the model-based prediction how long the communication would take. This tool does not require a parallel MPI run.

**Examples:**

```
$ prediction -n <np> -O p2p -i model.out -C Hockney -m 1024 -M 2049
```

will print the predicted P2P communication time for message sizes 1024 and 2048 bytes for each of the <np> processes with each other process. It is expected that the model.out was generated for <np> processes as well.

```
$ prediction -n <np> -O MPIB_Scatter_binomial -i model.out -C Hockney -m 1024 -M 2049
```

will print the predicted communication time for the binomial tree implementation of MPI\_Scatter with message size of 1024 and 2048 bytes processes. It is expected that the model.out was generated for <np> processes as well.

## 4.6 Generator of hierarchy files

This tool generates a hierarchy file automatically. At the moment, it bases its hierarchy on following rule:

- if hostnames start with the same 3 characters, they are clustered under the same parent
- after that, nodes are clustered according to the domain name after the first "." It is assumed that the environment will return the fully qualified name. The tool can be extended easily to work further on the fully qualified name

Usage:

```
$ mpirun [mpi options] generate_hierarchy_input_file <prefix for cluster 1> <pre
fix for cluster 2> ... > h-file.out
```

The output of this tool can be used as input for the [Hierarchical model builder](#). The output file might need some manual modifications.

For visualizing the ranks on each host, see [Generator of rank-to-host mapping](#)

## 4.7 Hierarchical model builder

This tool builds a model file efficiently based on hierarchy information. The input file should be a hierarchical graphviz file (DOT format) with a single root element. Every element has two attributes:

- type (0 or 1) - to denote a leaf or non-leaf element
- model\_data (empty string or model-specific data/parameters/etc) - can be initially used for storing hostnames

The output file itself is not hierarchical but flat. **Example:**

```
$ mpirun -n 2 hierarchical_model -C Hockney -i h-file.out -o model.out
```

The model.out file can then be used by either of these tools:

- [Predictor of the execution time of p2p or collective communication](#)
- [Collective benchmark \(using models\)](#)

### 4.7.1 Collective benchmark (using models)

The basic usage of this tool and a description how to plot the output file is described in MPIBlib.

**Example** using a model-based collective:

```
$ mpirun -n 16 collective -O Hockney_Scatter_dfs_binomial_min -o model=Hockney, f
ile=model.out -l libcpm_coll.so -m 1024 -M 2049 > collective.out
```

will benchmark a binomial tree scatter operation with message sizes per process of 1024 and 2048. The binomial tree will be generated through depth-first approach choosing the node with minimum predicted communication by the Hockney model.

## 4.8 Generator of rank-to-host mapping

This program is useful in combination with [Generator of hierarchy files](#) for debugging and visualization of the actual ranks running on different hosts. It produces a list in Graphviz DOT format.

**Example:**

```
mpirun -n <np> ranks2hosts > ranks2hosts.out
```

## 5 Model-based collectives

Shared library `libcpm_coll.so` implements different model-based algorithms of collectives:

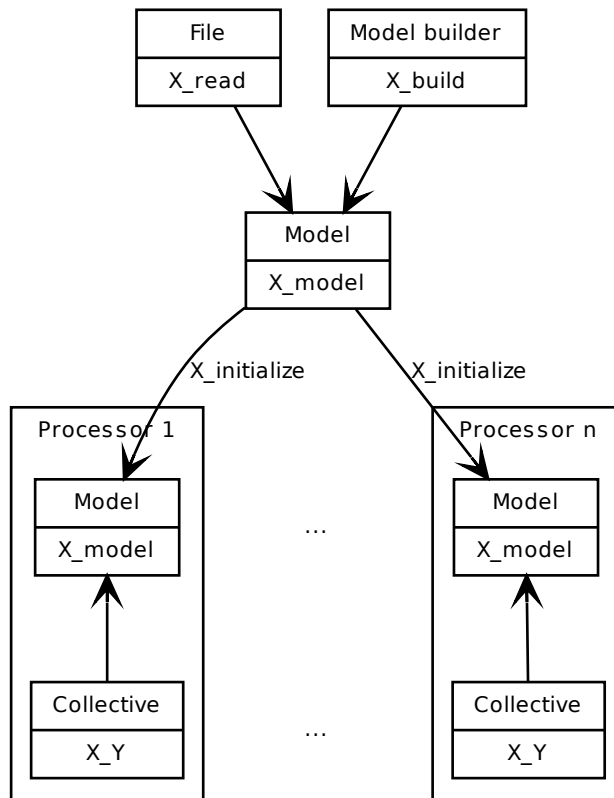
- [LinInterp-based collective operations](#)
- [Hockney-based collective operations](#)
- [LogGP-based collective operations](#)
- [PLogP-based collective operations](#)
- [LMO-based collective operations](#)

Command-line arguments:

- **verbose** verbose mode (default: no)
- **sgv** scatterv/gatherv mode
  - 0 propagation (default)
  - 1 broadcast
  - 2 allover
- **model** *S* communication model: Hockney, PLogP, LMO (required for model-based collectives)
- **file** *S* model data file (required for model-based collectives)

In order to preserve the original MPI interface, all model-based implementations use the global variables that provide model parameters. The interface of the implementation of a collective communication operation *Y* based on the model *X* includes the following components:

- `int X_Y(standard args)` is the model-based collective operation itself (for example, [Hockney\\_Scatter\\_bfs\\_binomial\\_min](#))
- `X_model* X_model_instance` is a global variable providing the model parameters (must be available at all processors in the MPI communicator)
- `int X_initialize(MPI_Comm, X_model* model), int X_finalize(MPI_Comm)` are functions responsible for allocation and deallocation of the model instance at all processors. The `model` argument encapsulates the model parameters obtained either from a file or the model builder.



All model-based algorithms are divided into two groups: model-specific and generic.

**Model-specific collectives** depend on certain communication performance models, using parameters specific for these models only. For example, [LMO\\_Gather\\_split\\_flat](#) directly uses the LMO model global variable and its threshold parameters to split the medium size messages and perform a series of linear gathers with small messages, in order to avoid escalations of the execution time on the clusters with the TCP/IP communication layer. A model-specific collective operation  $Y$  is implemented in the following way:

```

int X_Y(standard args) {
    if (condition with X_model_instance->param)
        return ...;
}

```

**Generic collectives** depend on the prediction of the execution time of some communication operation (communication primitive); they are parameterized by predictions, which can be provided by any model. The communication primitive can be either the collective operation itself or some other simple operation. See [Generic model-based collectives](#).

## 5.1 Tools

Part of [MPIBlib](#)

- [MPIBlib/tools/collective\\_test](#) - performs a universal or operation-specific collective benchmark with

given accuracy and efficiency

- MPIBlib/tools/collective - verifies implementations of collective communication operations

## 5.2 Tests

- [Test for model-based collectives](#)
- [Parallel matrix multiplication](#)

## 5.3 Test for model-based collectives

An example, showing how to work with communication models:

- input of the model file
- initialization and distribution of the model instance
- invocation of the model-based collective

```
#include "config.h"
#include "collectives/cpm_coll.h"
#include <getopt.h>
#include <malloc.h>
#include <stdio.h>

int main(int argc, char** argv) {
    MPI_Init(&argc, &argv);
    Hockney_model* model = NULL;
    int rank, size;
    MPI_Comm_rank(MPI_COMM_WORLD, &rank);
    MPI_Comm_size(MPI_COMM_WORLD, &size);
    if (rank == 0) {
        int c;
        if ((c = getopt (argc, argv, "i:")) != -1) {
            switch (c) {
                case 'i': {
                    FILE* file = fopen(optarg, "r");
                    if (file == NULL)
                        fprintf(stderr, "Cannot read input file %s\n", optarg);
                    else {
                        Hockney_read(file, &model);
                        fclose(file);
                    }
                    break;
                }
            }
        }
        else{
            fprintf(stderr, "You must specify an input argument\n");
        }
    }
    int exit = model == NULL;
    MPI_Bcast(&exit, 1, MPI_INT, 0, MPI_COMM_WORLD);
    if (exit) {
        MPI_Finalize();
        return 0;
    }
    Hockney_initialize(MPI_COMM_WORLD, model);
    if (rank == 0)
        Hockney_free(model);
    int count = 1024;
```

```

char* buffer = (char*)malloc(sizeof(char) * count);
Hockney_Bcast_ucs_binomial_min(buffer, count, MPI_CHAR, 0, MPI_COMM_WORLD);
free(buffer);
Hockney_finalize(MPI_COMM_WORLD);
MPI_Finalize();
return 0;
}

```

## 5.4 Parallel matrix multiplication

Heterogeneous parallel matrix product:

1. Estimates the execution time of `gsl_blas_dgemm` on all processors
2. Scatters the row blocks of matrix A proportionally to the speed of processors (`scaterv`)
3. Broadcasts matrix B
4. Gathers matrix C (`gatherv`)

Implemented with the following collective communications:

- native
- `Hockney_dfs_binomial_min`
- `Hockney_dfs_binomial_max`

## 6 Module Documentation

### 6.1 Measurement: benchmarking specific communication experiments

#### Functions

- void `CPM_measure_p2pp` (`MPI_Comm comm`, `int M`, `int parallel`, `MPIB_precision precision`, `MPIB_result *results`)
- void `CPM_measure_p2pp_p2p` (`MPI_Comm comm`, `int M`, `int parallel`, `MPIB_precision precision`, `MPIB_result *results_p2pp`, `MPIB_result *results_p2p[2]`)

#### 6.1.1 Detailed Description

This module extends the `MPILib` functionality in order to measure the execution time of some specific communication experiments required to estimate the parameters of communication performance models.

#### 6.1.2 Function Documentation

##### 6.1.2.1 void `CPM_measure_p2pp` ( `MPI_Comm comm`, `int M`, `int parallel`, `MPIB_precision precision`, `MPIB_result * results` )

Measures the 1-to-2 execution times.

Measures the execution times of

- $i \xleftrightarrow[0]{M} jk$ ,
- $j \xleftrightarrow[0]{M} ik$ ,
- $k \xleftrightarrow[0]{M} ij$

in the communicator,  $i < j < k$ .

#### Parameters

**comm** communicator, number of nodes  $\geq 3$

**M** message size

**parallel** several non-overlapped p2pp communications at the same time if non-zero

**precision** measurement parameters

**results** array of  $3C_n^3$  measurement results (significant only at root)

**6.1.2.2** void CPM\_measure\_p2pp\_p2p ( MPI\_Comm comm, int M, int parallel,  
MPIB\_precision precision, MPIB\_result \* results\_p2pp, MPIB\_result \* results\_p2p[2] )

Measures the 1-to-2 and 1-to-1 execution times.

Measures the execution times of

- $i \xleftrightarrow[0]{M} jk, i \xleftrightarrow[0]{M} j, i \xleftrightarrow[0]{M} k$ ,
- $j \xleftrightarrow[0]{M} ik, j \xleftrightarrow[0]{M} i, j \xleftrightarrow[0]{M} k$ ,
- $k \xleftrightarrow[0]{M} ij, k \xleftrightarrow[0]{M} i, k \xleftrightarrow[0]{M} j$

in the communicator,  $i < j < k$ .

#### Parameters

**comm** communicator, number of nodes  $\geq 3$

**M** message size

**parallel** several non-overlapped point-to-point communications at the same time if non-zero

**precision** measurement parameters

**results\_p2pp** array of  $3C_n^3$  measurement results (significant only at root)

**results\_p2p** 2 arrays of  $3C_n^3$  measurement results (significant only at root)

## 6.2 Communication performance models

This module provide the following models:

- [The linear interpolation model \[6\]](#)
- [The heterogeneous Hockney model \[6\]](#)

- [The heterogeneous PLogP model \[7\]](#)
- [LMO: an advanced heterogeneous communication performance model \[8\]](#)

For each model, this module provides the following interface (X stands for the name of the model: Hockney, PLogP or LMO):

- ```
struct X_model {
    CPM_predictor predictor;
    ... // parameters
}
```

is a data structure containing the parameters of the model and a set of the estimation functions defined by [CPM\\_predictor](#).

- ```
X_model* X_alloc(int);
void X_free(X_model*);
```

are functions to allocate/free the model data.

- ```
void X_estimate(MPI_Comm, MPIB_msgset, MPIB_precision, int parallel, X_model**);
```

is a function responsible for estimation of the model with given precision and message sizes. If the `parallel` argument is set to zero, the communication experiments will be performed consequently, otherwise in parallel.

- ```
void X_read(FILE*, X_model**);
void X_write(FILE*, const X_model*);
```

are functions for input/output of the model data.

This interface can be used directly in parallel applications. It is also a basis for optimized implementations of MPI collective communication operations (see [Model-based collectives](#)).

## 6.3 Prediction of communication time

### Functions

- double [CPM\\_predict\\_brs](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int size\_bytes, CPM\_coll\_ops operation)
- double [CPM\\_predict\\_sgv](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int \*size\_bytes, CPM\_coll\_ops operation)
- double [CPM\\_predict\\_flat\\_sgv](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int \*size\_bytes)
- double [CPM\\_predict\\_flat\\_sg](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int size\_bytes)
- double [CPM\\_predict\\_flat\\_sgv\\_parallel](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int \*size\_bytes)
- double [CPM\\_predict\\_flat\\_sg\\_parallel](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int size\_bytes)
- double [CPM\\_predict\\_flat\\_sgv\\_serial](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int \*size\_bytes)
- double [CPM\\_predict\\_flat\\_sg\\_serial](#) ([CPM\\_predictor](#) \*predictor, int size, int root, int size\_bytes)

### 6.3.1 Detailed Description

This module provides generic predict functions

### 6.3.2 Function Documentation

**6.3.2.1** `double CPM_predict_brs ( CPM_predictor * predictor, int size, int root, int size_bytes, CPM_coll_ops operation )`

Predicts the execution time of Bcast, Reduce, Scatter, Gather operations

**6.3.2.2** `double CPM_predict_sgv ( CPM_predictor * predictor, int size, int root, int * size_bytes, CPM_coll_ops operation )`

Predicts the execution time of Scatterv, Gatherv operations

**6.3.2.3** `double CPM_predict_flat_sgv ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

Predicts the execution time of flat-tree Scatterv, Gatherv operations

**6.3.2.4** `double CPM_predict_flat_sg ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

Predicts the execution time of flat-tree Scatter, Gather operations

**6.3.2.5** `double CPM_predict_flat_sgv_parallel ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

return maximum of all the P2P transfer times

**6.3.2.6** `double CPM_predict_flat_sg_parallel ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

return maximum of all the P2P transfer times

**6.3.2.7** `double CPM_predict_flat_sgv_serial ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

return sum of all the P2P transfer times

**6.3.2.8** `double CPM_predict_flat_sg_serial ( CPM_predictor * predictor, int size, int root, int * size_bytes )`

return sum of all the P2P transfer times

## 6.4 The heterogeneous Hockney model

### Classes

- struct [Hockney\\_model](#)

### Functions

- [Hockney\\_model](#) \* [Hockney\\_alloc](#) (int n)
- void [Hockney\\_free](#) ([Hockney\\_model](#) \*model)
- void [Hockney\\_read](#) (FILE \*stream, [Hockney\\_model](#) \*\*model)
- void [Hockney\\_write](#) (FILE \*stream, const [Hockney\\_model](#) \*model)

- void [Hockney\\_estimate](#) (MPI\_Comm comm, int M, MPIB\_precision precision, int parallel, [Hockney\\_model](#) \*\*model)
- void [Hockney\\_estimate\\_regression](#) (MPI\_Comm comm, MPIB\_msgset msgset, MPIB\_precision precision, int parallel, [Hockney\\_model](#) \*\*model)
- double [Hockney\\_predict\\_p2p](#) (void \*\_this, int i, int j, int M)
- double [Hockney\\_predict\\_sg\\_flat\\_serial](#) (void \*\_this, int root, int M)
- double [Hockney\\_predict\\_sg\\_flat\\_parallel](#) (void \*\_this, int root, int M)
- double [Hockney\\_predict\\_sg\\_binomial](#) (void \*\_this, int root, int M)
- double [Hockney\\_hpredict\\_sg\\_flat\\_serial](#) (void \*\_this, int M)
- double [Hockney\\_hpredict\\_sg\\_flat\\_parallel](#) (void \*\_this, int M)
- double [Hockney\\_hpredict\\_sg\\_binomial](#) (void \*\_this, int M)

#### 6.4.1 Detailed Description

This module provides building of the heterogeneous extension of the Hockney model and estimation of the execution time of point-to-point and collective communication operations.

In contrast to the original model [6], which is based on two point-to-point parameters, estimating the point-to-point execution time as  $T(M) = \alpha + \beta M$ , the heterogeneous model distinguishes the parameters of each pair of processors  $T_{ij}(M) = \alpha_{ij} + \beta_{ij} M$ .

#### 6.4.2 Function Documentation

##### 6.4.2.1 [Hockney\\_model\\*](#) [Hockney\\_alloc](#) ( int *n* )

Allocates memory for the Hockney model.

##### Parameters

*n* number of processors

##### 6.4.2.2 void [Hockney\\_free](#) ( [Hockney\\_model](#) \* *model* )

Frees the Hockney model.

##### Parameters

*model* the Hockney model

##### 6.4.2.3 void [Hockney\\_read](#) ( FILE \* *stream*, [Hockney\\_model](#) \*\* *model* )

Reads the Hockney model.

##### 6.4.2.4 void [Hockney\\_write](#) ( FILE \* *stream*, const [Hockney\\_model](#) \* *model* )

Writes the Hockney model.

#### 6.4.2.5 void Hockney\_estimate ( MPI\_Comm *comm*, int *M*, MPIB\_precision *precision*, int *parallel*, Hockney\_model \*\* *model* )

Estimates the parameters of the Hockney model in two series of roundtrips with empty and non-empty messages. (accuracy depends on the precision of measurements):

- Measures the execution time  $T_{ij}(0)$  of each  $i \xleftrightarrow{0} j$  roundtrip in the communicator,  $i < j$ , to find  $\alpha_{ij} = T_{ij}(0)$ . To obtain more accurate results performs a series of roundtrips and takes the average  $T_{ij}(0)$ .
- Measures the execution time  $T_{ij}(M)$  of each  $i \xleftrightarrow{M} j$  roundtrip in the communicator,  $i < j$ , to find  $\beta_{ij} = \frac{T_{ij}(M) - \alpha_{ij}}{M}$ . To obtain more accurate results performs a series of roundtrips and takes the average  $T_{ij}(M)$ .

##### Parameters

***comm*** communicator, number of nodes  $\geq 2$

***M*** message size

***precision*** measurement precision

***parallel*** several non-overlapped point-to-point communications at the same time if non-zero

***model*** Hockney model (significant only at root)

#### 6.4.2.6 void Hockney\_estimate\_regression ( MPI\_Comm *comm*, MPIB\_msgset *msgset*, MPIB\_precision *precision*, int *parallel*, Hockney\_model \*\* *model* )

Estimates the parameters of the Hockney model in a series of roundtrips with different message sizes (accuracy depends on the message set):

- Performs single communication experiments for different message sizes.
- Selects message sizes regularly TODO: Adaptive selection, comparing the result of measurement with the prediction based on the linear regression over all previous results.

##### Parameters

***comm*** communicator, number of nodes  $\geq 2$

***msgset*** message set

***precision*** measurement precision

***parallel*** several non-overlapped point-to-point communications at the same time if non-zero

***model*** Hockney model (significant only at root)

#### 6.4.2.7 double Hockney\_predict\_p2p ( void \* *\_this*, int *i*, int *j*, int *M* )

Predicts the execution time of a point-to-point communication as  $T_{ij}(M) = \alpha_{ij} + \beta_{ij}M$ .

**6.4.2.8 double Hockney\_predict\_sg\_flat\_serial ( void \* *\_this*, int *root*, int *M* )**

Heterogeneous prediction of flat-tree scatter/gather (sequential point-to-point communications)

$$\sum_{i=0, i \neq r}^{n-1} (\alpha_{ri} + \beta_{ri}M)$$

**6.4.2.9 double Hockney\_predict\_sg\_flat\_parallel ( void \* *\_this*, int *root*, int *M* )**

Heterogeneous prediction of flat-tree scatter/gather (parallel point-to-point communications)

$$\max_{i=0, i \neq r}^{n-1} (\alpha_{ri} + \beta_{ri}M)$$

**6.4.2.10 double Hockney\_predict\_sg\_binomial ( void \* *\_this*, int *root*, int *M* )**

Heterogeneous prediction of binomial scatter/gather (point-to-point communications are sequential within the same root or parallel otherwise)  $\alpha_{ri} + \beta_{ri}2^{\log_2(n-1)}M + \max S(\log_2(n-1) - 1)$

**6.4.2.11 double Hockney\_hpredict\_sg\_flat\_serial ( void \* *\_this*, int *M* )**

Homogeneous prediction of flat-tree scatter/gather (sequential point-to-point communications)  
 $(n-1)(\alpha + \beta M)$

**6.4.2.12 double Hockney\_hpredict\_sg\_flat\_parallel ( void \* *\_this*, int *M* )**

Homogeneous prediction of flat-tree scatter/gather (parallel point-to-point communications)  $\alpha + \beta M$

**6.4.2.13 double Hockney\_hpredict\_sg\_binomial ( void \* *\_this*, int *M* )**

Homogeneous prediction of binomial scatter/gather (point-to-point communications are sequential within the same root or parallel otherwise)  $(\log_2 n)\alpha + (n-1)\beta M$

**6.5 The linear interpolation model****Classes**

- struct [LinInterp\\_model](#)

**Functions**

- [LinInterp\\_model](#) \* [LinInterp\\_alloc](#) (int size, int num\_points)
- void [LinInterp\\_free](#) ([LinInterp\\_model](#) \*model)
- void [LinInterp\\_estimate](#) (MPI\_Comm comm, MPIB\_precision precision, MPIB\_msgset msgset, int parallel, [LinInterp\\_model](#) \*\*\_model)
- void [LinInterp\\_read](#) (FILE \*stream, [LinInterp\\_model](#) \*\*model)
- void [LinInterp\\_write](#) (FILE \*stream, const [LinInterp\\_model](#) \*model)
- double [LinInterp\\_predict\\_p2p](#) (void \*\_this, int i, int j, int M)

### 6.5.1 Function Documentation

#### 6.5.1.1 `LinInterp_model* LinInterp_alloc ( int size, int num_points )`

Allocates the linear interpolation model

#### 6.5.1.2 `void LinInterp_free ( LinInterp_model * model )`

Frees the linear interpolation model

#### 6.5.1.3 `void LinInterp_estimate ( MPI_Comm comm, MPIB_precision precision, MPIB_msgset msgset, int parallel, LinInterp_model ** _model )`

Estimates the linear interpolation model

#### 6.5.1.4 `void LinInterp_read ( FILE * stream, LinInterp_model ** model )`

Reads the linear interpolation model

#### 6.5.1.5 `void LinInterp_write ( FILE * stream, const LinInterp_model * model )`

Writes the linear interpolation model

#### 6.5.1.6 `double LinInterp_predict_p2p ( void * _this, int i, int j, int M )`

Predicts the p2p communication execution time using the linear interpolation model

## 6.6 LMO: an advanced heterogeneous communication performance model

### Classes

- struct [LMO\\_model](#)

### Defines

- #define [LMO\\_C3](#)(n) (n) \* ((n) - 1) \* ((n) - 2) / 6
- #define [LMO\\_JK2INDEX](#)(n, i, j, k) (2 \* (n) - (i) - 1) \* ((i) - 1) \* (i) / 4 + (2 \* (n) - (i) - (j) + 1) \* ((j) - (i)) / 2 + (k) - (j) - 1

### Functions

- [LMO\\_model](#) \* [LMO\\_alloc](#) (int n)
- void [LMO\\_free](#) ([LMO\\_model](#) \*model)
- void [LMO\\_read](#) (FILE \*stream, [LMO\\_model](#) \*\*model)
- void [LMO\\_write](#) (FILE \*stream, const [LMO\\_model](#) \*model)
- void [LMO\\_estimate](#) (MPI\_Comm comm, MPIB\_precision precision, MPIB\_msgset msgset, int parallel, [LMO\\_model](#) \*\*model)

- void [LMO\\_estimate\\_homogeneous](#) (MPI\_Comm comm, MPIB\_precision precision, MPIB\_msgset msgset, [LMO\\_model](#) \*\*model)
- double [LMO\\_predict\\_p2p](#) (void \*\_this, int i, int j, int M)
- double [LMO\\_predict\\_scatter\\_flat](#) (void \*\_this, int root, int M)
- double [LMO\\_predict\\_gather\\_flat](#) (void \*\_this, int root, int M)

### 6.6.1 Detailed Description

This module provides building of the LMO model and estimation of the execution time of point-to-point and collective communication operations.

### 6.6.2 Define Documentation

**6.6.2.1** `#define LMO_C3( n ) (n) * ((n) - 1) * ((n) - 2) / 6`

$$C_n^3$$

**6.6.2.2** `#define LMO_IJK2INDEX( n, i, j, k ) (2 * (n) - (i) - 1) * ((i) - 1) * (i) / 4 + (2 * (n) - (i) - (j) + 1) * ((j) - (i)) / 2 + (k) - (j) - 1`

The index of the  $(i, j, k)$  element in the array of  $C_n^3$  elements,  $i < j < k < n$ .

$$\frac{C_{n-1}^2 + C_{n-i}^2}{2}i + \frac{(n-i+1) + (n-j)}{2}(j-i) + (k-j-1).$$

### 6.6.3 Function Documentation

**6.6.3.1** `LMO_model* LMO_alloc ( int n )`

Allocates the LMO model.

#### Parameters

*n* number of processors

#### Returns

LMO model

**6.6.3.2** `void LMO_free ( LMO_model * model )`

Frees the LMO model.

#### Parameters

*model* LMO model

**6.6.3.3** `void LMO_read ( FILE * stream, LMO_model ** model )`

Reads the LMO model.

### 6.6.3.4 void LMO\_write ( FILE \* *stream*, const LMO\_model \* *model* )

Writes the LMO model.

### 6.6.3.5 void LMO\_estimate ( MPI\_Comm *comm*, MPIB\_precision *precision*, MPIB\_msgset *msgset*, int *parallel*, LMO\_model \*\* *model* )

Estimates the parameters of the LMO model.

1. Measures one-to-many execution time for the message sizes  $0 < M < max\_size$  and performs the Bai & Perron algorithm over the F statistic for the data to find the  $S$  breakpoint in the piecewise linear regression  $T \sim M$ .
2. Measures many-to-one execution time for the message sizes  $0 < M < max\_size$  and performs the Bai & Perron algorithm over the F statistic for the data to find the  $M_2$  breakpoint in the piecewise linear regression  $T \sim 1$ .
3. In the loop, measures many-to-one execution time for the message sizes  $0 < M < m$  with the stride reduced twice each time.  $m$  is a message size for which a tenfold escalation of the execution time has been observed at the previous step. As the stride reaches 1 byte,  $m$  is truncated to a kb value, which will be  $M_1$ .
4. Finds the fixed processing delays and latencies.

For each 3 nodes  $i < j < k$ , measures execution times and solves systems of equations:

$$\left\{ \begin{array}{ll} T_{ij}(0) = 2(C_i + L_{ij} + C_j) & i \xrightarrow{0} j \\ T_{jk}(0) = 2(C_j + L_{jk} + C_k) & j \xrightarrow{0} k \\ T_{ik}(0) = 2(C_i + L_{ik} + C_k) & i \xrightarrow{0} k \\ T_{ijk}(0) = 2(2C_i + \max_{x=j,k}(L_{ix} + C_x)) & i \xrightarrow{0} jk \\ T_{jik}(0) = 2(2C_j + \max_{x=i,k}(L_{jx} + C_x)) & j \xrightarrow{0} ik \\ T_{kij}(0) = 2(2C_k + \max_{x=i,j}(L_{kx} + C_x)) & k \xrightarrow{0} ij \end{array} \right.$$

averages solutions:

$$T_{ijk}(0) = 2(2C_i + \max_{x=j,k}(L_{ix} + C_x)) = 2C_i + \max_{x=j,k} T_{ix}(0)$$

$$\left\{ \begin{array}{l} C_i = (T_{ijk}(0) - \max_{x=j,k} T_{ix}(0))/2 \\ C_j = (T_{jik}(0) - \max_{x=i,k} T_{jx}(0))/2 \\ C_k = (T_{kij}(0) - \max_{x=i,j} T_{kx}(0))/2 \\ L_{ij} = T_{ij}(0)/2 - C_i - C_j \\ L_{jk} = T_{jk}(0)/2 - C_j - C_k \\ L_{ik} = T_{ik}(0)/2 - C_i - C_k \end{array} \right.$$

and checks confidence intervals.

5. Finds the variable processing delays and transmission rates.

For each 3 nodes  $i < j < k$ , solves systems of equations:

$$\left\{ \begin{array}{ll} T_{ij}(M) = 2(C_i + L_{ij} + C_j + M(t_i + \frac{1}{\beta_{ij}} + t_j)) & i \xleftarrow{\frac{M}{M}} j \\ T_{jk}(M) = 2(C_j + L_{jk} + C_k + M(t_j + \frac{1}{\beta_{jk}} + t_k)) & j \xleftarrow{\frac{M}{M}} k \\ T_{ik}(M) = 2(C_i + L_{ik} + C_k + M(t_i + \frac{1}{\beta_{ik}} + t_k)) & i \xleftarrow{\frac{M}{M}} k \\ T_{ijk}(M) = 2(2C_i + Mt_i) + \max_{x=j,k} (2(L_{ix} + C_x) + M(\frac{1}{\beta_{ix}} + t_x)) & i \xleftarrow{\frac{M}{0}} jk \\ T_{jik}(M) = 2(2C_j + Mt_j) + \max_{x=i,k} (2(L_{jx} + C_x) + M(\frac{1}{\beta_{jx}} + t_x)) & j \xleftarrow{\frac{M}{0}} ik \\ T_{kij}(M) = 2(2C_k + Mt_k) + \max_{x=i,j} (2(L_{kx} + C_x) + M(\frac{1}{\beta_{kx}} + t_x)) & k \xleftarrow{\frac{M}{0}} ij \end{array} \right.$$

averages solutions:

$$T_{ijk}(M) = 2(2C_i + Mt_i) + \max_{x=j,k} (2(L_{ix} + C_x) + M(\frac{1}{\beta_{ix}} + t_x)) = 2C_i + Mt_i + \max_{x=j,k} (T_{ix}(0) + T_{ix}(M))/2$$

$$\left\{ \begin{array}{l} t_i = (T_{ijk}(M) - \max_{x=j,k} (T_{ix}(0) + T_{ix}(M))/2 - 2C_i)/M \\ t_j = (T_{jik}(M) - \max_{x=i,k} (T_{jx}(0) + T_{jx}(M))/2 - 2C_j)/M \\ t_k = (T_{kij}(M) - \max_{x=i,j} (T_{kx}(0) + T_{kx}(M))/2 - 2C_k)/M \\ \frac{1}{\beta_{ij}} = (T_{ij}(M)/2 - C_i - L_{ij} - C_j)/M - t_i - t_j \\ \frac{1}{\beta_{jk}} = (T_{jk}(M)/2 - C_j - L_{jk} - C_k)/M - t_j - t_k \\ \frac{1}{\beta_{ik}} = (T_{ik}(M)/2 - C_i - L_{ik} - C_k)/M - t_i - t_k \end{array} \right.$$

and checks confidence intervals.

**Note**

R must be initialized by LMO\_init\_R at root beforehand.

**Parameters**

**comm** communicator

**precision** measurement precision

**msgset** message set

**parallel** several non-overlapped point-to-point communications at the same time if non-zero

**model** LMO model (significant only at root)

**6.6.3.6 void LMO\_estimate\_homogeneous ( MPI\_Comm comm, MPIB\_precision precision, MPIB\_msgset msgset, LMO\_model \*\* model )**

Estimates the parameters of the homogeneous LMO model (n = 1)

**6.6.3.7 double LMO\_predict\_p2p ( void \* \_this, int i, int j, int M )**

Predicts the execution time of point-to-point communication. The execution time of  $i \xrightarrow{\frac{M}{M}} j$  is equal to

$$C_i + L_{ij} + C_j + M(t_i + \frac{1}{\beta_{ij}} + t_j)$$

**Parameters**

*\_this* LMO model  
*i* index of the processor  
*j* index of the processor  
*M* message size

**Returns**

predicted execution time

**6.6.3.8 double LMO\_predict\_scatter\_flat ( void \* \_this, int root, int M )**

Predicts the execution time of flat-tree scatter.  $(n-1)(C_r + Mt_r) + \max_{i=1, i \neq r}^{n-1} (L_{ri} + C_i + M(\frac{1}{\beta_{ri}} + t_i))$

**Parameters**

*\_this* LMO model  
*root* root processor  
*M* message size

**Returns**

predicted execution time

**6.6.3.9 double LMO\_predict\_gather\_flat ( void \* \_this, int root, int M )**

Predicts the execution time of flat-tree gather.

$$(n-1)(C_r + Mt_r) + \begin{cases} \max_{i=1}^{n-1} (L_{ri} + C_i + M(\frac{1}{\beta_{ri}} + t_i)) & M < M_1 \\ \sum_{i=1}^{n-1} (L_{ri} + C_i + M(\frac{1}{\beta_{ri}} + t_i)) & M > M_2 \end{cases}$$

**Parameters**

*\_this* LMO model  
*root* root processor  
*M* message size

**Returns**

predicted execution time

**6.7 The heterogeneous PLogP model****Classes**

- struct [PLogP\\_model](#)

### Functions

- `PLogP_model * PLogP_alloc (int n)`
- `void PLogP_free (PLogP_model *model)`
- `void PLogP_read (FILE *stream, PLogP_model **model)`
- `void PLogP_write (FILE *stream, const PLogP_model *model)`
- `void PLogP_estimate (MPI_Comm comm, MPIB_precision precision, MPIB_msgset msgset, int parallel, PLogP_model **model)`
- `double PLogP_predict_p2p (void *_this, int i, int j, int M)`
- `double LogGP_predict_p2p (void *_this, int i, int j, int M)`
- `double PLogP_hpredict_sg_linear (void *_this, int M)`
- `double LogGP_hpredict_sg_linear (void *_this, int M)`

#### 6.7.1 Detailed Description

This module provides building of the heterogeneous extension of the parameterised LogP model and PLogP/LogGP predictions of the execution time of point-to-point and collective communication operations.

The PLogP model [7] is an extension of the LogP model [4]. The PLogP model is defined in terms of the following parameters:

- $L$  - latency,
- $o_s(M)$  and  $o_r(M)$  - sender and receiver overheads,
- $g(M)$  - gap per message (delay between consecutive communications),
- $P$  - number of processors.

Some of these parameters are functions, which makes this model "parameterised". They are represented by piecewise linear functions. The PLogP parameters are estimated with help of the `logp_mpi` library [7]. The heterogeneous extension is a set of the PLogP models built for each pair of processors:  $\{L_{ij}, o_{s_{ij}}(M), o_{r_{ij}}(M), g_{ij}(M)\}_{i \neq j=0}^{n-1}$ .

The parameters of the LogP and LogGP (another extension of the LogP model proposed in [1], which considers message size and includes an extra parameter,  $G$ , gap per byte) models can be expressed via the PLogP parameters:

- $L = L^p + g^p(1) - o_s^p(1) - o_r^p(1)$
- $2 * o = o_s^p(1) + o_r^p(1)$
- $g = g^p(1)$
- $G = g^p(M_{max})/M_{max}$

The heterogeneous extension of the LogGP model can be obtained by applying the above equations to the PLogP parameters of each pair of processors. Therefore, this module provides both the PLogP and LogGP predictions of the execution time of collective communication operations.

#### 6.7.2 Function Documentation

##### 6.7.2.1 PLogP\_model\* PLogP\_alloc ( int n )

Allocates memory for PLogP model

**6.7.2.2 void PLogP\_free ( PLogP\_model \* *model* )**

Frees the PLogP model.

**6.7.2.3 void PLogP\_read ( FILE \* *stream*, PLogP\_model \*\* *model* )**

Reads the PLogP model.

**Note**

Creates a temporary file.

**6.7.2.4 void PLogP\_write ( FILE \* *stream*, const PLogP\_model \* *model* )**

Writes the PLogP model.

**Note**

Creates a temporary file.

**6.7.2.5 void PLogP\_estimate ( MPI\_Comm *comm*, MPIB\_precision *precision*, MPIB\_msgset *msgset*, int *parallel*, PLogP\_model \*\* *model* )**

Estimates the PLogP model. Calls the logp\_mpi library.

**Parameters***comm* communicator*precision* measurement precision*msgset* message set*parallel* several non-overlapped point-to-point communications at the same time if non-zero*model* PLogP model**6.7.2.6 double PLogP\_predict\_p2p ( void \* *\_this*, int *i*, int *j*, int *M* )**Predicts the execution time of a point-to-point communication as  $L + g(M)$ **6.7.2.7 double LogGP\_predict\_p2p ( void \* *\_this*, int *i*, int *j*, int *M* )**Predicts the execution time of a point-to-point communication according to LogGP model as  $L + 2 * o + G(M - 1)$ **Parameters***\_this* PLogP model*i* index of the process*j* index of the process*M* message size

### 6.7.2.8 double PLogP\_hpredict\_sg\_linear ( void \* *\_this*, int *M* )

Homogeneous PLogP prediction of linear scatter/gather:  $L + (n - 1)g(M)$

### 6.7.2.9 double LogGP\_hpredict\_sg\_linear ( void \* *\_this*, int *M* )

Homogeneous LogGP prediction of linear scatter/gather:  $L + 2o + (n - 1)(M - 1)G + (n - 2)g$

## 6.8 Generic model-based collectives

### Functions

- int [CPM\\_Scatterv\\_sorted\\_flat](#) (CPM\_predictor \*predictor, MPIB\_sort\_order order, void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Gatherv\\_sorted\\_flat](#) (CPM\_predictor \*predictor, MPIB\_sort\_order order, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcnts, int \*displs, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Bcast\\_dfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- int [CPM\\_Bcast\\_bfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- int [CPM\\_Reduce\\_dfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- int [CPM\\_Reduce\\_bfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- int [CPM\\_Bcast\\_ucs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- int [CPM\\_Reduce\\_ucs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- int [CPM\\_Scatter\\_ucs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Gather\\_ucs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Scatter\\_bfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Scatter\\_dfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Gather\\_bfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)
- int [CPM\\_Gather\\_dfs\\_binomial](#) (CPM\_predictor \*predictor, CPM\_next\_node\_strategy next\_node, void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcnt, MPI\_Datatype recvtpe, int root, MPI\_Comm comm)

- `int CPM_Scatterv_dfs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Scatterv_bfs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Scatterv_ucs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Gatherv_dfs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcnts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Gatherv_bfs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcnts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Gatherv_ucs_binomial (CPM_predictor *predictor, CPM_next_node_strategy next_node, void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcnts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Scatterv_Traff (CPM_predictor *predictor, void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int CPM_Gatherv_Traff (CPM_predictor *predictor, void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcnts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`

### 6.8.1 Detailed Description

There are two typical examples of generic implementations: switch between algorithms and processor mapping.

- *Switch between algorithms* is an implementation that uses the prediction of the execution time of different algorithms of the collective operation, finds the fastest algorithm for given message size and number of processors, and switches between them. The prediction can be provided by any model. Different algorithms of the collective operations are communication primitives. [3], [9] .
- *Processor mapping* is an implementation of a tree-based algorithm of the collective operation that maps the processors to the nodes of the communication tree in accordance with the performance of the point-to-point communications. In this case, the point-to-point communication operation is a communication primitive, which execution time can be predicted by any model. [5], [2] .

The interface of a generic collective operation `Y` based on the prediction of the communication primitive `Z` consists of:

- a general, model-independent, implementation of collective operation:

```
int CPM_Y(CPM_predictor* predictor, standard args) {
    if (condition with predictor->predict_Z(predictor, args))
        return ...;
}
```

which includes an extra parameter, a model-based predictor `predictor`, and calls its function `predict_Z` to predict the execution time of the communication primitive. For example, `CPM_Scatter_bfs_binomial` is a generic binomial scatter based on the point-to-point predictions `CPM_predictor::predict_p2p`.

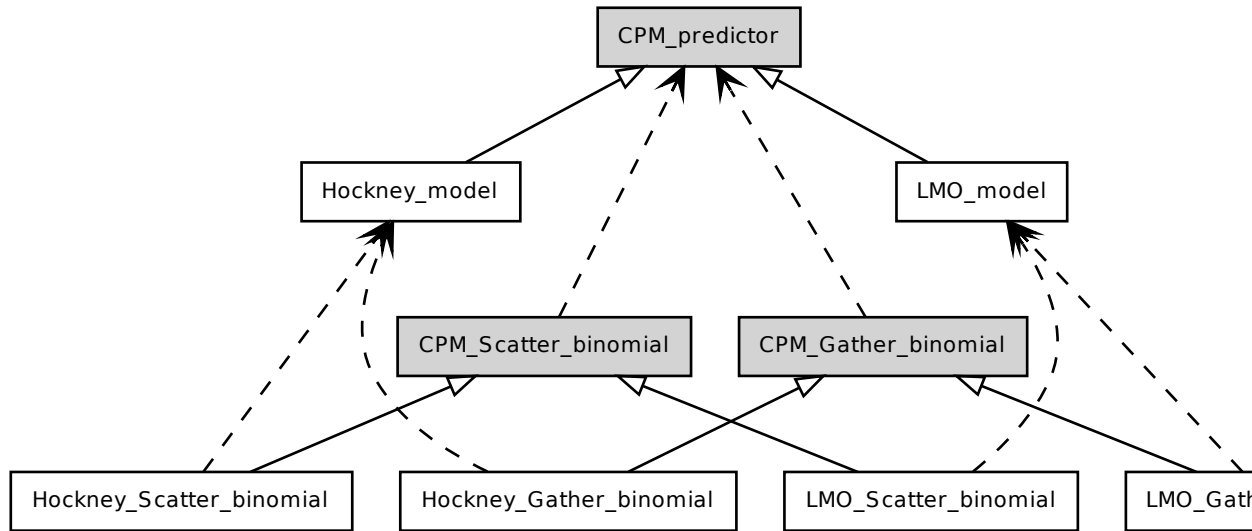
- particular model-based implementations (X stands for the name of the model):

```
int X_Y(standard args) {
    return CPM_Y(&X_model_instance->predictor, standard args);
}
```

For example, [Hockney\\_Scatter\\_bfs\\_binomial\\_min](#) is a derivative of the generic algorithm [CPM\\_Scatter\\_bfs\\_binomial](#) that is based on the Hockney prediction of the point-to-point execution time.

This approach provides flexibility by reusing the same (general) implementations of a collective operation with different communication performance models.

A design of the generic collectives:



Generic model-based functions are designed for use in both C/C++.

## 6.8.2 Function Documentation

**6.8.2.1** `int CPM_Scatterv_sorted_flat ( CPM_predictor * predictor, MPIB_sort_order order, void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic sorted flat-tree scatterv. Starts from the end/beginning (order = DESC/ASC) of the sorted list.

**6.8.2.2** `int CPM_Gatherv_sorted_flat ( CPM_predictor * predictor, MPIB_sort_order order, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic sorted flat-tree gatherv. Starts from the beginning/end (order = DESC/ASC) of the sorted list.

**6.8.2.3** `int CPM_Bcast_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

Generic DFS binomial bcst.

**6.8.2.4** `int CPM_Bcast_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

Generic BFS binomial bcst.

**6.8.2.5** `int CPM_Reduce_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

Generic DFS binomial reduce.

**6.8.2.6** `int CPM_Reduce_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

Generic BFS binomial reduce.

**6.8.2.7** `int CPM_Bcast_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

Generic UCS binomial bcst.

**6.8.2.8** `int CPM_Reduce_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

Generic UCS binomial reduce.

**6.8.2.9** `int CPM_Scatter_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic UCS binomial scatter.

**6.8.2.10** `int CPM_Gather_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic UCS binomial gather.

**6.8.2.11** `int CPM_Scatter_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic BFS binomial scatter.

**6.8.2.12** `int CPM_Scatter_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic DFS binomial scatter.

**6.8.2.13** `int CPM_Gather_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic BFS binomial gather.

**6.8.2.14** `int CPM_Gather_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic DFS binomial gather.

**6.8.2.15** `int CPM_Scatterv_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic DFS binomial scatterv.

**6.8.2.16** `int CPM_Scatterv_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic BFS binomial scatterv.

**6.8.2.17** `int CPM_Scatterv_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic UCS binomial scatterv.

**6.8.2.18** `int CPM_Gatherv_dfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic DFS binomial gatherv.

**6.8.2.19** `int CPM_Gatherv_bfs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic BFS binomial gatherv.

**6.8.2.20** `int CPM_Gatherv_ucs_binomial ( CPM_predictor * predictor, CPM_next_node_strategy next_node, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic UCS binomial gatherv.

**6.8.2.21** `int CPM_Scatterv_Traff ( CPM_predictor * predictor, void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic Traff scatterv.

**6.8.2.22** `int CPM_Gatherv_Traff ( CPM_predictor * predictor, void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Generic Traff gatherv.

## 6.9 Hockney-based collective operations

### Functions

- `int Hockney_initialize (MPI_Comm comm, Hockney_model *model)`
- `int Hockney_finalize (MPI_Comm comm)`
- `int Hockney_Scatterv_sorted_flat_asc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int Hockney_Gatherv_sorted_flat_asc (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcunts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int Hockney_Scatterv_sorted_flat_dsc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int Hockney_Gatherv_sorted_flat_dsc (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcunts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int Hockney_Bcast_dfs_binomial_min (void *buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm)`
- `int Hockney_Reduce_dfs_binomial_min (void *sendbuf, void *recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)`
- `int Hockney_Bcast_dfs_binomial_max (void *buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm)`
- `int Hockney_Reduce_bfs_binomial_min (void *sendbuf, void *recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)`
- `int Hockney_Reduce_bfs_binomial_max (void *sendbuf, void *recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm)`
- `int Hockney_Bcast_bfs_binomial_min (void *buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm)`

- `int Hockney_Bcast_bfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int Hockney_Reduce_dfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int Hockney_Bcast_ucs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int Hockney_Reduce_ucs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int Hockney_Bcast_ucs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int Hockney_Reduce_ucs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int Hockney_Scatter_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatter_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatter_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatter_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatter_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatter_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gather_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatterv_dfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gatherv_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatterv_dfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gatherv_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Scatterv_bfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int Hockney_Gatherv_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

- int [Hockney\\_Scatterv\\_bfs\\_binomial\\_max](#) (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Gatherv\\_bfs\\_binomial\\_max](#) (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Scatterv\\_ucs\\_binomial\\_min](#) (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Gatherv\\_ucs\\_binomial\\_min](#) (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Scatterv\\_ucs\\_binomial\\_max](#) (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Gatherv\\_ucs\\_binomial\\_max](#) (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Scatterv\\_Traff](#) (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- int [Hockney\\_Gatherv\\_Traff](#) (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

### 6.9.1 Detailed Description

The implementations of tree-based algorithms are similar to [5] and [2].

### 6.9.2 Function Documentation

#### 6.9.2.1 int Hockney\_initialize ( MPI\_Comm *comm*, Hockney\_model \* *model* )

Initializes the instances of the Hockney model at all processes in the communication.

##### Parameters

*comm* MPI communicator

*model* Hockney model (significant only at root)

#### 6.9.2.2 int Hockney\_finalize ( MPI\_Comm *comm* )

Destroys the instances of the Hockney model at all processes in the communication.

#### 6.9.2.3 int Hockney\_Scatterv\_sorted\_flat\_asc ( void \* *sendbuf*, int \* *sendcounts*, int \* *displs*, MPI\_Datatype *sendtype*, void \* *recvbuf*, int *recvcount*, MPI\_Datatype *recvtype*, int *root*, MPI\_Comm *comm* )

Sorted flat-tree scatterv based on the Hockney model.

#### 6.9.2.4 int Hockney\_Gatherv\_sorted\_flat\_asc ( void \* *sendbuf*, int *sendcount*, MPI\_Datatype *sendtype*, void \* *recvbuf*, int \* *recvcounts*, int \* *displs*, MPI\_Datatype *recvtype*, int *root*, MPI\_Comm *comm* )

Sorted flat-tree gatherv based on the Hockney model.

**6.9.2.5** `int Hockney_Scatterv_sorted_flat_dsc ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree scatterv based on the Hockney model.

**6.9.2.6** `int Hockney_Gatherv_sorted_flat_dsc ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree gatherv based on the Hockney model.

**6.9.2.7** `int Hockney_Bcast_dfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the Hockney model.

**6.9.2.8** `int Hockney_Reduce_dfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the Hockney model.

**6.9.2.9** `int Hockney_Bcast_dfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the Hockney model.

**6.9.2.10** `int Hockney_Reduce_bfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the Hockney model.

**6.9.2.11** `int Hockney_Reduce_bfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the Hockney model.

**6.9.2.12** `int Hockney_Bcast_bfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the Hockney model.

**6.9.2.13** `int Hockney_Bcast_bfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the Hockney model.

**6.9.2.14** `int Hockney_Reduce_dfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the Hockney model.

**6.9.2.15** `int Hockney_Bcast_ucs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the Hockney model.

**6.9.2.16** `int Hockney_Reduce_ucs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the Hockney model.

**6.9.2.17** `int Hockney_Bcast_ucs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the Hockney model.

**6.9.2.18** `int Hockney_Reduce_ucs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the Hockney model.

**6.9.2.19** `int Hockney_Scatter_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the Hockney model.

**6.9.2.20** `int Hockney_Gather_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the Hockney model.

**6.9.2.21** `int Hockney_Scatter_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the Hockney model.

**6.9.2.22** `int Hockney_Gather_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the Hockney model.

**6.9.2.23** `int Hockney_Scatter_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the Hockney model.

**6.9.2.24** `int Hockney_Gather_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the Hockney model.

**6.9.2.25** `int Hockney_Scatter_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the Hockney model.

**6.9.2.26** `int Hockney_Gather_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the Hockney model.

**6.9.2.27** `int Hockney_Scatter_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the Hockney model.

**6.9.2.28** `int Hockney_Gather_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the Hockney model.

**6.9.2.29** `int Hockney_Scatter_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the Hockney model.

**6.9.2.30** `int Hockney_Gather_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the Hockney model.

**6.9.2.31** `int Hockney_Scatterv_dfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the Hockney model.

**6.9.2.32** `int Hockney_Gatherv_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnt, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the Hockney model.

**6.9.2.33** `int Hockney_Scatterv_dfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the Hockney model.

**6.9.2.34** `int Hockney_Gatherv_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnt, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the Hockney model.

**6.9.2.35** `int Hockney_Scatterv_bfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the Hockney model.

**6.9.2.36** `int Hockney_Gatherv_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnt, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the Hockney model.

**6.9.2.37** `int Hockney_Scatterv_bfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the Hockney model.

**6.9.2.38** `int Hockney_Gatherv_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnt, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the Hockney model.

**6.9.2.39** `int Hockney_Scatterv_ucs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the Hockney model.

**6.9.2.40** `int Hockney_Gatherv_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the Hockney model.

**6.9.2.41** `int Hockney_Scatterv_ucs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the Hockney model.

**6.9.2.42** `int Hockney_Gatherv_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the Hockney model.

**6.9.2.43** `int Hockney_Scatterv_Traff ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Scatterv based on modified Traff using the Hockney model.

**6.9.2.44** `int Hockney_Gatherv_Traff ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Gatherv based on modified Traff using the Hockney model.

## 6.10 LinInterp-based collective operations

### Functions

- `int LinInterp_initialize (MPI_Comm comm, LinInterp\_model *model)`
- `int LinInterp_finalize (MPI_Comm comm)`
- `int LinInterp_Scatterv_sorted_flat_asc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int LinInterp_Gatherv_sorted_flat_asc (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcounts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int LinInterp_Scatterv_sorted_flat_dsc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)`

- `int LinInterp_Gatherv_sorted_flat_dsc` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_dfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_dfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_dfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_bfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_bfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_bfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_bfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_dfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_ucs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_ucs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Bcast_ucs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LinInterp_Reduce_ucs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatter_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gather_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

- `int LinInterp_Scatterv_dfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_dfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_bfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_bfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_ucs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_ucs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Scatterv_Traff` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LinInterp_Gatherv_Traff` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

### 6.10.1 Detailed Description

The implementations of tree-based algorithms are similar to [5] and [2] .

### 6.10.2 Function Documentation

#### 6.10.2.1 `int LinInterp_initialize ( MPI_Comm comm, LinInterp_model * model )`

Initializes the instances of the LinInterp model at all processes in the communicator.

#### Parameters

*comm* MPI communicator

*model* LinInterp model (significant only at root)

**6.10.2.2 int LinInterp\_finalize ( MPI\_Comm comm )**

Destroys the instances of the LinInterp model at all processes in the communication.

**6.10.2.3 int LinInterp\_Scatterv\_sorted\_flat\_asc ( void \* sendbuf, int \* sendcounts, int \* displs, MPI\_Datatype sendtype, void \* recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm )**

Sorted flat-tree scatterv based on the LinInterp model.

**6.10.2.4 int LinInterp\_Gatherv\_sorted\_flat\_asc ( void \* sendbuf, int sendcount, MPI\_Datatype sendtype, void \* recvbuf, int \* recvcounts, int \* displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm )**

Sorted flat-tree gatherv based on the LinInterp model.

**6.10.2.5 int LinInterp\_Scatterv\_sorted\_flat\_dsc ( void \* sendbuf, int \* sendcounts, int \* displs, MPI\_Datatype sendtype, void \* recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm )**

Sorted flat-tree scatterv based on the LinInterp model.

**6.10.2.6 int LinInterp\_Gatherv\_sorted\_flat\_dsc ( void \* sendbuf, int sendcount, MPI\_Datatype sendtype, void \* recvbuf, int \* recvcounts, int \* displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm )**

Sorted flat-tree gatherv based on the LinInterp model.

**6.10.2.7 int LinInterp\_Bcast\_dfs\_binomial\_min ( void \* buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm )**

DFS binomial bcast based on the LinInterp model.

**6.10.2.8 int LinInterp\_Reduce\_dfs\_binomial\_min ( void \* sendbuf, void \* recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm )**

DFS binomial reduce based on the LinInterp model.

**6.10.2.9 int LinInterp\_Bcast\_dfs\_binomial\_max ( void \* buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm )**

DFS binomial bcast based on the LinInterp model.

**6.10.2.10 int LinInterp\_Reduce\_bfs\_binomial\_min ( void \* sendbuf, void \* recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm )**

BFS binomial reduce based on the LinInterp model.

**6.10.2.11** `int LinInterp_Reduce_bfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the LinInterp model.

**6.10.2.12** `int LinInterp_Bcast_bfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the LinInterp model.

**6.10.2.13** `int LinInterp_Bcast_bfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the LinInterp model.

**6.10.2.14** `int LinInterp_Reduce_dfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the LinInterp model.

**6.10.2.15** `int LinInterp_Bcast_ucs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the LinInterp model.

**6.10.2.16** `int LinInterp_Reduce_ucs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the LinInterp model.

**6.10.2.17** `int LinInterp_Bcast_ucs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the LinInterp model.

**6.10.2.18** `int LinInterp_Reduce_ucs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the LinInterp model.

**6.10.2.19** `int LinInterp_Scatter_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the LinInterp model.

**6.10.2.20** `int LinInterp_Gather_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the LinInterp model.

**6.10.2.21** `int LinInterp_Scatter_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the LinInterp model.

**6.10.2.22** `int LinInterp_Gather_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the LinInterp model.

**6.10.2.23** `int LinInterp_Scatter_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the LinInterp model.

**6.10.2.24** `int LinInterp_Gather_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the LinInterp model.

**6.10.2.25** `int LinInterp_Scatter_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the LinInterp model.

**6.10.2.26** `int LinInterp_Gather_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the LinInterp model.

**6.10.2.27** `int LinInterp_Scatter_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the LinInterp model.

**6.10.2.28** `int LinInterp_Gather_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the LinInterp model.

**6.10.2.29** `int LinInterp_Scatter_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the LinInterp model.

**6.10.2.30** `int LinInterp_Gather_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the LinInterp model.

**6.10.2.31** `int LinInterp_Scatterv_dfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the LinInterp model.

**6.10.2.32** `int LinInterp_Gatherv_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounst, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the LinInterp model.

**6.10.2.33** `int LinInterp_Scatterv_dfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the LinInterp model.

**6.10.2.34** `int LinInterp_Gatherv_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounst, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the LinInterp model.

**6.10.2.35** `int LinInterp_Scatterv_bfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the LinInterp model.

**6.10.2.36** `int LinInterp_Gatherv_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the LinInterp model.

**6.10.2.37** `int LinInterp_Scatterv_bfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the LinInterp model.

**6.10.2.38** `int LinInterp_Gatherv_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the LinInterp model.

**6.10.2.39** `int LinInterp_Scatterv_ucs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the LinInterp model.

**6.10.2.40** `int LinInterp_Gatherv_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the LinInterp model.

**6.10.2.41** `int LinInterp_Scatterv_ucs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the LinInterp model.

**6.10.2.42** `int LinInterp_Gatherv_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the LinInterp model.

**6.10.2.43** `int LinInterp_Scatterv_Traff ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Scatterv based on modified Traff using the LinInterp model.

**6.10.2.44** `int LinInterp_Gatherv_Traff ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnts, int * displs, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

Gatherv based on modified Traff using the LinInterp model.

## 6.11 LMO-based collective operations

### Functions

- `int LMO_initialize (MPI_Comm comm, LMO_model *model)`
- `int LMO_finalize (MPI_Comm comm)`
- `int LMO_Scatter_split_flat (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm)`
- `int LMO_Gather_split_flat (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm)`

### 6.11.1 Function Documentation

**6.11.1.1** `int LMO_initialize ( MPI_Comm comm, LMO_model * model )`

Initializes the instances of the LMO model on all processes in the communicator

#### Parameters

*comm* MPI communicator

*model* LMO model (significant only at root)

**6.11.1.2** `int LMO_finalize ( MPI_Comm comm )`

Destroys the instances of the LMO model on all processes in the communicator

**6.11.1.3** `int LMO_Scatter_split_flat ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

Split flat-tree MPI\_Scatter

**6.11.1.4** `int LMO_Gather_split_flat ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

Split flat-tree MPI\_Gather

## 6.12 LogGP-based collective operations

### Functions

- `int LogGP_Scatterv_sorted_flat_asc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm)`

- `int LogGP_Gatherv_sorted_flat_asc` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_sorted_flat_dsc` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_sorted_flat_dsc` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Bcast_dfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Reduce_dfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Bcast_dfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Reduce_bfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Reduce_bfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Bcast_bfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Bcast_bfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Reduce_dfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Bcast_ucs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Reduce_ucs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Bcast_ucs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int LogGP_Reduce_ucs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int LogGP_Scatter_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gather_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatter_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gather_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatter_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gather_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatter_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gather_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatter_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gather_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatter_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

- `int LogGP_Gather_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_dfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_dfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_bfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_bfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_ucs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_ucs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Scatterv_Traff` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int LogGP_Gatherv_Traff` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

### 6.12.1 Function Documentation

**6.12.1.1** `int LogGP_Scatterv_sorted_flat_asc ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree scatterv based on the LogGP model.

**6.12.1.2** `int LogGP_Gatherv_sorted_flat_asc ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree gatherv based on the LogGP model.

**6.12.1.3** `int LogGP_Scatterv_sorted_flat_dsc ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree scatterv based on the LogGP model.

**6.12.1.4** `int LogGP_Gatherv_sorted_flat_dsc ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree gatherv based on the LogGP model.

**6.12.1.5** `int LogGP_Bcast_dfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the LogGP model.

**6.12.1.6** `int LogGP_Reduce_dfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the LogGP model.

**6.12.1.7** `int LogGP_Bcast_dfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the LogGP model.

**6.12.1.8** `int LogGP_Reduce_bfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the LogGP model.

**6.12.1.9** `int LogGP_Reduce_bfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the LogGP model.

**6.12.1.10** `int LogGP_Bcast_bfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the LogGP model.

**6.12.1.11** `int LogGP_Bcast_bfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the LogGP model.

**6.12.1.12** `int LogGP_Reduce_dfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the LogGP model.

**6.12.1.13** `int LogGP_Bcast_ucs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the LogGP model.

**6.12.1.14** `int LogGP_Reduce_ucs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the LogGP model.

**6.12.1.15** `int LogGP_Bcast_ucs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the LogGP model.

**6.12.1.16** `int LogGP_Reduce_ucs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the LogGP model.

**6.12.1.17** `int LogGP_Scatter_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the LogGP model.

**6.12.1.18** `int LogGP_Gather_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the LogGP model.

**6.12.1.19** `int LogGP_Scatter_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the LogGP model.

**6.12.1.20** `int LogGP_Gather_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the LogGP model.

**6.12.1.21** `int LogGP_Scatter_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the LogGP model.

**6.12.1.22** `int LogGP_Gather_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the LogGP model.

**6.12.1.23** `int LogGP_Scatter_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the LogGP model.

**6.12.1.24** `int LogGP_Gather_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the LogGP model.

**6.12.1.25** `int LogGP_Scatter_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the LogGP model.

**6.12.1.26** `int LogGP_Gather_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the LogGP model.

**6.12.1.27** `int LogGP_Scatter_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the LogGP model.

**6.12.1.28** `int LogGP_Gather_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the LogGP model.

**6.12.1.29** `int LogGP_Scatterv_dfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the LogGP model.

**6.12.1.30** `int LogGP_Gatherv_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the LogGP model.

**6.12.1.31** `int LogGP_Scatterv_dfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatterv based on the LogGP model.

**6.12.1.32** `int LogGP_Gatherv_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gatherv based on the LogGP model.

**6.12.1.33** `int LogGP_Scatterv_bfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the LogGP model.

**6.12.1.34** `int LogGP_Gatherv_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the LogGP model.

**6.12.1.35** `int LogGP_Scatterv_bfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatterv based on the LogGP model.

**6.12.1.36** `int LogGP_Gatherv_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the LogGP model.

**6.12.1.37** `int LogGP_Scatterv_ucs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the LogGP model.

**6.12.1.38** `int LogGP_Gatherv_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the LogGP model.

**6.12.1.39** `int LogGP_Scatterv_ucs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the LogGP model.

**6.12.1.40** `int LogGP_Gatherv_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the LogGP model.

**6.12.1.41** `int LogGP_Scatterv_Traff ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Scatterv based on modified Traff using the LogGP model.

**6.12.1.42** `int LogGP_Gatherv_Traff ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Gatherv based on modified Traff using the LogGP model.

## 6.13 PLogP-based collective operations

### Functions

- `int PLogP_initialize (MPI_Comm comm, PLogP_model *model)`
- `int PLogP_finalize (MPI_Comm comm)`
- `int PLogP_Scatterv_sorted_flat_asc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int PLogP_Gatherv_sorted_flat_asc (void *sendbuf, int sendcount, MPI_Datatype sendtype, void *recvbuf, int *recvcounts, int *displs, MPI_Datatype recvtype, int root, MPI_Comm comm)`
- `int PLogP_Scatterv_sorted_flat_dsc (void *sendbuf, int *sendcounts, int *displs, MPI_Datatype sendtype, void *recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm)`

- `int PLogP_Gatherv_sorted_flat_dsc` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Bcast_dfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Reduce_dfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Bcast_dfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Reduce_bfs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Reduce_bfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Bcast_bfs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Bcast_bfs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Reduce_dfs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Bcast_ucs_binomial_min` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Reduce_ucs_binomial_min` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Bcast_ucs_binomial_max` (void \*buffer, int count, MPI\_Datatype datatype, int root, MPI\_Comm comm)
- `int PLogP_Reduce_ucs_binomial_max` (void \*sendbuf, void \*recvbuf, int count, MPI\_Datatype datatype, MPI\_Op op, int root, MPI\_Comm comm)
- `int PLogP_Scatter_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatter_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatter_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatter_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatter_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatter_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gather_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_dfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

- `int PLogP_Gatherv_dfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_dfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_dfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_bfs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_bfs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_bfs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_bfs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_ucs_binomial_min` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_ucs_binomial_min` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_ucs_binomial_max` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_ucs_binomial_max` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Scatterv_Traff` (void \*sendbuf, int \*sendcounts, int \*displs, MPI\_Datatype sendtype, void \*recvbuf, int recvcount, MPI\_Datatype recvtype, int root, MPI\_Comm comm)
- `int PLogP_Gatherv_Traff` (void \*sendbuf, int sendcount, MPI\_Datatype sendtype, void \*recvbuf, int \*recvcounts, int \*displs, MPI\_Datatype recvtype, int root, MPI\_Comm comm)

### 6.13.1 Function Documentation

#### 6.13.1.1 `int PLogP_initialize ( MPI_Comm comm, PLogP_model * model )`

Initializes the instances of the PLogP model at all processes in the communicator.

##### Parameters

*comm* MPI communicator

*model* PLogP model (significant only at root)

#### 6.13.1.2 `int PLogP_finalize ( MPI_Comm comm )`

Destroys the instances of the PLogP model at all processes in the communicator.

#### 6.13.1.3 `int PLogP_Scatterv_sorted_flat_asc ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree scatterv based on the PLogP model.

**6.13.1.4** `int PLogP_Gatherv_sorted_flat_asc ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree gatherv based on the PLogP model.

**6.13.1.5** `int PLogP_Scatterv_sorted_flat_dsc ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree scatterv based on the PLogP model.

**6.13.1.6** `int PLogP_Gatherv_sorted_flat_dsc ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcounts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Sorted flat-tree gatherv based on the PLogP model.

**6.13.1.7** `int PLogP_Bcast_dfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the PLogP model.

**6.13.1.8** `int PLogP_Reduce_dfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the PLogP model.

**6.13.1.9** `int PLogP_Bcast_dfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

DFS binomial bcast based on the PLogP model.

**6.13.1.10** `int PLogP_Reduce_bfs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the PLogP model.

**6.13.1.11** `int PLogP_Reduce_bfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

BFS binomial reduce based on the PLogP model.

**6.13.1.12** `int PLogP_Bcast_bfs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the PLogP model.

**6.13.1.13** `int PLogP_Bcast_bfs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

BFS binomial bcast based on the PLogP model.

**6.13.1.14** `int PLogP_Reduce_dfs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

DFS binomial reduce based on the PLogP model.

**6.13.1.15** `int PLogP_Bcast_ucs_binomial_min ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the PLogP model.

**6.13.1.16** `int PLogP_Reduce_ucs_binomial_min ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the PLogP model.

**6.13.1.17** `int PLogP_Bcast_ucs_binomial_max ( void * buffer, int count, MPI_Datatype datatype, int root, MPI_Comm comm )`

UCS binomial bcast based on the PLogP model.

**6.13.1.18** `int PLogP_Reduce_ucs_binomial_max ( void * sendbuf, void * recvbuf, int count, MPI_Datatype datatype, MPI_Op op, int root, MPI_Comm comm )`

UCS binomial reduce based on the PLogP model.

**6.13.1.19** `int PLogP_Scatter_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the PLogP model.

**6.13.1.20** `int PLogP_Gather_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the PLogP model.

**6.13.1.21** `int PLogP_Scatter_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial scatter based on the PLogP model.

**6.13.1.22** `int PLogP_Gather_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

DFS binomial gather based on the PLogP model.

**6.13.1.23** `int PLogP_Scatter_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the PLogP model.

**6.13.1.24** `int PLogP_Gather_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the PLogP model.

**6.13.1.25** `int PLogP_Scatter_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatter based on the PLogP model.

**6.13.1.26** `int PLogP_Gather_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gather based on the PLogP model.

**6.13.1.27** `int PLogP_Scatter_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the PLogP model.

**6.13.1.28** `int PLogP_Gather_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gather based on the PLogP model.

**6.13.1.29** `int PLogP_Scatter_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcount, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial scatter based on the PLogP model.

**6.13.1.30** `int PLogP_Gather_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

BFS binomial gather based on the PLogP model.

**6.13.1.31** `int PLogP_Scatterv_dfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

DFS binomial scatterv based on the PLogP model.

**6.13.1.32** `int PLogP_Gatherv_dfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnts, int * displs, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

DFS binomial gatherv based on the PLogP model.

**6.13.1.33** `int PLogP_Scatterv_dfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

DFS binomial scatterv based on the PLogP model.

**6.13.1.34** `int PLogP_Gatherv_dfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnts, int * displs, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

DFS binomial gatherv based on the PLogP model.

**6.13.1.35** `int PLogP_Scatterv_bfs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

BFS binomial scatterv based on the PLogP model.

**6.13.1.36** `int PLogP_Gatherv_bfs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcnts, int * displs, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

BFS binomial gatherv based on the PLogP model.

**6.13.1.37** `int PLogP_Scatterv_bfs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcnt, MPI_Datatype recvtpe, int root, MPI_Comm comm )`

BFS binomial scatterv based on the PLogP model.

**6.13.1.38** `int PLogP_Gatherv_bfs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

BFS binomial gatherv based on the PLogP model.

**6.13.1.39** `int PLogP_Scatterv_ucs_binomial_min ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the PLogP model.

**6.13.1.40** `int PLogP_Gatherv_ucs_binomial_min ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the PLogP model.

**6.13.1.41** `int PLogP_Scatterv_ucs_binomial_max ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial scatterv based on the PLogP model.

**6.13.1.42** `int PLogP_Gatherv_ucs_binomial_max ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

UCS binomial gatherv based on the PLogP model.

**6.13.1.43** `int PLogP_Scatterv_Traff ( void * sendbuf, int * sendcounts, int * displs, MPI_Datatype sendtype, void * recvbuf, int recvcunt, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Scatterv based on modified Traff using the PLogP model.

**6.13.1.44** `int PLogP_Gatherv_Traff ( void * sendbuf, int sendcount, MPI_Datatype sendtype, void * recvbuf, int * recvcunts, int * displs, MPI_Datatype recvtype, int root, MPI_Comm comm )`

Gatherv based on modified Traff using the PLogP model.

## 7 Namespace Documentation

### 7.1 CPM::BRSG Namespace Reference

#### Classes

- class [UCS\\_binomial\\_builder](#)
- class [BFS\\_binomial\\_builder](#)
- class [DFS\\_binomial\\_builder](#)

#### 7.1.1 Detailed Description

Communication tree builders for model-based algorithms of bcast/reduce/scatter/gather

### 7.2 CPM::hierarchy Namespace Reference

#### 7.2.1 Detailed Description

Hierarchy

### 7.3 CPM::SGv Namespace Reference

#### Classes

- class [BFS\\_binomial\\_builder](#)
- class [DFS\\_binomial\\_builder](#)
- class [Traff\\_builder](#)
- class [UCS\\_binomial\\_builder](#)

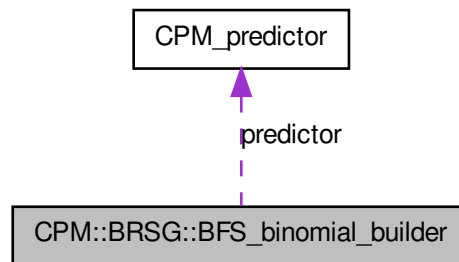
#### 7.3.1 Detailed Description

Communication tree builders for model-based algorithms of scatterv/gatherv

## 8 Class Documentation

### 8.1 CPM::BRSG::BFS\_binomial\_builder Class Reference

Collaboration diagram for CPM::BRSG::BFS\_binomial\_builder:



#### 8.1.1 Detailed Description

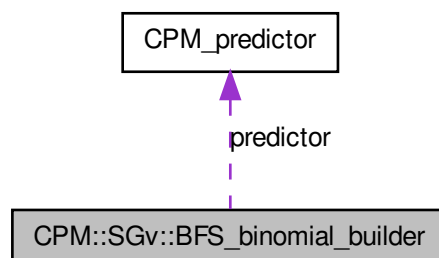
BFS binomial tree builder. The largest subtree on the left. If next\_node = MIN, the processor i with minimum p2p prediction is selected; otherwise, with maximum.

The documentation for this class was generated from the following file:

- models/cpm\_brsg\_tree\_builders.hpp

### 8.2 CPM::SGv::BFS\_binomial\_builder Class Reference

Collaboration diagram for CPM::SGv::BFS\_binomial\_builder:



### 8.2.1 Detailed Description

BFS binomial tree builder for scatterv/gatherv. The largest subtree on the left. If next\_node = MIN, the processor *i* with minimum p2p prediction is selected; otherwise, with maximum.

The documentation for this class was generated from the following file:

- models/cpm\_sgv\_bfs\_tree\_builder.hpp

## 8.3 CPM\_predictor Struct Reference

### Public Attributes

- double(\* [predict\\_p2p](#))(void \*\_this, int i, int j, int M)

### 8.3.1 Detailed Description

An predictor of the communication execution time. Contains a set of the predict functions:

- double (\*predict\_p2p)(void\* \_this, int i, int j, int M) is a function predicting the execution time of the point-to-point communication operation, where *i* and *j* are the indices of the processors, *M* is a message size.
- double (\*predict\_Y)(void\* \_this, int M, int root, int size) is a function predicting the execution time of the algorithm of collective communication operation *Y* (for example, Scatter\_binomial). The *root* parameter is significant for the collective operations with the root processor.

Each of the estimation functions includes a self pointer argument *\_this* to be used in the derived data structures (communication performance models).

### 8.3.2 Member Data Documentation

#### 8.3.2.1 double(\* CPM\_predictor::predict\_p2p)(void \*\_this, int i, int j, int M)

Predicts the execution time of the point-to-point communication

### Parameters

- \_this* a self pointer
- i* an index of a processor involved in point-to-point communication
- j* an index of a processor involved in point-to-point communication
- M* a message size.

### Returns

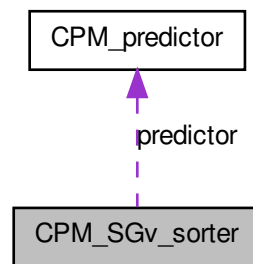
predicted execution time

The documentation for this struct was generated from the following file:

- models/cpm\_model.h

## 8.4 CPM\_SGv\_sorter Struct Reference

Collaboration diagram for CPM\_SGv\_sorter:



### 8.4.1 Detailed Description

Prediction sorter for generic sorted flat-tree scatterv/gatherv

The documentation for this struct was generated from the following file:

- collectives/cpm\_sgv\_flat.c

## 8.5 CPM\_triplet Struct Reference

Collaboration diagram for CPM\_triplet:



### Public Attributes

- int [values](#) [3]
- struct [CPM\\_triplet](#) \* [prev](#)
- struct [CPM\\_triplet](#) \* [next](#)

### 8.5.1 Detailed Description

List of triplets

### 8.5.2 Member Data Documentation

#### 8.5.2.1 `int CPM_triplet::values[3]`

values

#### 8.5.2.2 `struct CPM_triplet* CPM_triplet::prev`

previous triplet

#### 8.5.2.3 `struct CPM_triplet* CPM_triplet::next`

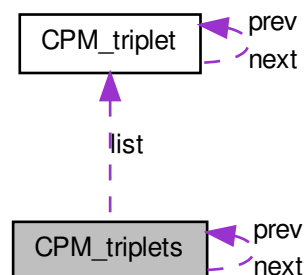
next triplet

The documentation for this struct was generated from the following file:

- `models/cpm_measurement.c`

## 8.6 CPM\_triplets Struct Reference

Collaboration diagram for CPM\_triplets:



### Public Attributes

- `CPM_triplet * list`
- `struct CPM_triplets * prev`
- `struct CPM_triplets * next`

### 8.6.1 Detailed Description

List of lists of triplets

### 8.6.2 Member Data Documentation

#### 8.6.2.1 CPM\_triplet\* CPM\_triplets::list

items

#### 8.6.2.2 struct CPM\_triplets\* CPM\_triplets::prev

previous list

#### 8.6.2.3 struct CPM\_triplets\* CPM\_triplets::next

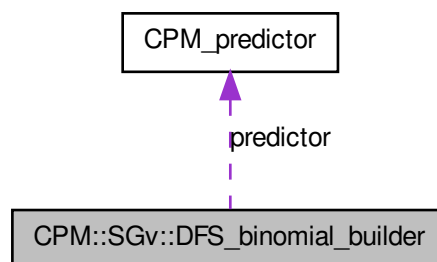
next list

The documentation for this struct was generated from the following file:

- models/cpm\_measurement.c

## 8.7 CPM::SGv::DFS\_binomial\_builder Class Reference

Collaboration diagram for CPM::SGv::DFS\_binomial\_builder:



### 8.7.1 Detailed Description

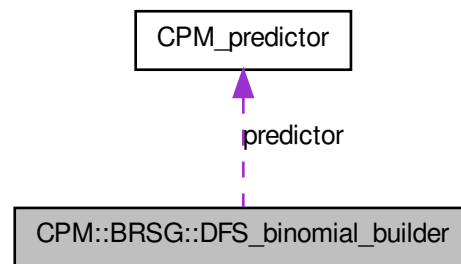
DFS binomial tree builder for scatterv/gatherv. The largest subtree on the right. If next\_node == MIN, the processor i with minimum p2p prediction is selected; otherwise, with maximum.

The documentation for this class was generated from the following file:

- models/cpm\_sgv\_dfs\_tree\_builder.hpp

## 8.8 CPM::BRSG::DFS\_binomial\_builder Class Reference

Collaboration diagram for CPM::BRSG::DFS\_binomial\_builder:



### 8.8.1 Detailed Description

DFS binomial tree builder. The largest subtree on the right. If next\_node == MIN, the processor i with minimum p2p prediction is selected; otherwise, with maximum.

The documentation for this class was generated from the following file:

- models/cpm\_brsb\_tree\_builders.hpp

## 8.9 H\_Model< p2p\_model > Class Template Reference

### Public Member Functions

- void [read\\_hierarchy\\_file](#) (istream &is)

### 8.9.1 Detailed Description

**template<class p2p\_model> class H\_Model< p2p\_model >**

Hierarchical model

### Template Parameters

**p2p\_model** a class of P2P model implementing the following functions (on the example of [hockney\\_p2p\\_model](#)):

- [hockney\\_p2p\\_model::serialize](#)
- [hockney\\_p2p\\_model::estimate](#)
- [hockney\\_p2p\\_model::write\\_header](#)
- [hockney\\_p2p\\_model::write\\_elem](#)
- [hockney\\_p2p\\_model::write\\_footer](#)

### 8.9.2 Member Function Documentation

#### 8.9.2.1 `template<class p2p_model> void H_Model< p2p_model >::read_hierarchy_file ( istream & is ) [inline]`

Reads the hierarchy file in the followin format:

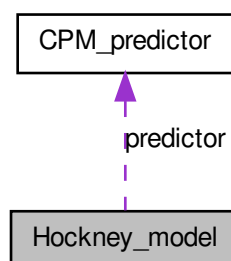
```
digraph G {
12 [type=0,model_data="csicluster01"];
2 [type=0, model_data="csicluster02"];
5 [type=0, model_data="csicluster03"];
6 [type=0, model_data="csicluster04"];
3 [type=1, model_data=""];
4 [type=1, model_data=""];
13 [type=1, model_data=""];
7 [type=1, model_data=""];
8 [type=1, model_data=""];
9 [type=1, model_data=""];
10 [type=1, model_data=""];
5->7;
6->8;
4->10;
9->10;
7->9 ;
8->9;
12->13 ;
2->3 ;
13->4;
3->4
}
```

The documentation for this class was generated from the following file:

- models/hierarchical/hierarchical\_model.hpp

## 8.10 Hockney\_model Struct Reference

Collaboration diagram for Hockney\_model:



**Public Attributes**

- [CPM\\_predictor predictor](#)
- `int n`
- `double * a`
- `double * b`

**8.10.1 Detailed Description**

The heterogeneous Hockney model. The point-to-point parameters for all pairs of processors.

**8.10.2 Member Data Documentation****8.10.2.1 CPM\_predictor Hockney\_model::predictor**

Predictor

**8.10.2.2 int Hockney\_model::n**

Number of nodes

**8.10.2.3 double\* Hockney\_model::a**

Array of  $C_n^2$  parameters:  $\{\alpha_{ij}\}_{i \neq j=0}^{n-1}$

**8.10.2.4 double\* Hockney\_model::b**

Array of  $C_n^2$  parameters:  $\{\beta_{ij}\}_{i \neq j=0}^{n-1}$

The documentation for this struct was generated from the following file:

- `models/hockney.h`

**8.11 hockney\_p2p\_model Class Reference****Public Member Functions**

- `template<class Archive >`  
void [serialize](#) (Archive &ar, const unsigned int version)
- void [estimate](#) (MPI\_Comm comm, MPIB\_msgset msgset, MPIB\_precision precision, int i, int j)
- void [write\\_header](#) (ostream &os)
- void [write\\_elem](#) (int i, int j, ostream &os)
- void [write\\_footer](#) (ostream &os)

**8.11.1 Detailed Description**

The Hockney P2P model

### 8.11.2 Member Function Documentation

**8.11.2.1** `template<class Archive > void hockney_p2p_model::serialize ( Archive & ar, const unsigned int version ) [inline]`

Serializes the model

**8.11.2.2** `void hockney_p2p_model::estimate ( MPI_Comm comm, MPIB_msgset msgset, MPIB_precision precision, int i, int j ) [inline]`

Estimates the model

**8.11.2.3** `void hockney_p2p_model::write_header ( ostream & os ) [inline]`

Writes the header to a stream

**8.11.2.4** `void hockney_p2p_model::write_elem ( int i, int j, ostream & os ) [inline]`

Writes the footer to a stream

**8.11.2.5** `void hockney_p2p_model::write_footer ( ostream & os ) [inline]`

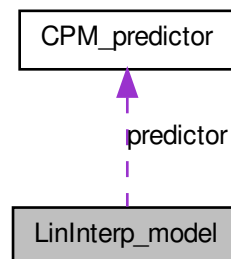
Writes the model to a stream

The documentation for this class was generated from the following file:

- models/hierarchical/hockney\_p2p\_model.hpp

## 8.12 LinInterp\_model Struct Reference

Collaboration diagram for LinInterp\_model:



**Public Attributes**

- [CPM\\_predictor](#) predictor
- int [n](#)
- int [num\\_points](#)
- gsl\_spline \*\* [splines](#)
- gsl\_interp\_accel \*\* [accels](#)

**8.12.1 Detailed Description**

The linear interpolation model

**8.12.2 Member Data Documentation****8.12.2.1 CPM\_predictor LinInterp\_model::predictor**

Predictor

**8.12.2.2 int LinInterp\_model::n**

Number of nodes

**8.12.2.3 int LinInterp\_model::num\_points**

Number of points

**8.12.2.4 gsl\_spline\*\* LinInterp\_model::splines**

Splines

**8.12.2.5 gsl\_interp\_accel\*\* LinInterp\_model::accels**

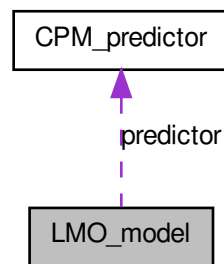
Accelerators

The documentation for this struct was generated from the following file:

- `models/lininterp.h`

## 8.13 LMO\_model Struct Reference

Collaboration diagram for LMO\_model:



### Public Attributes

- `CPM_predictor predictor`
- `int n`
- `double * C`
- `double * L`
- `double * t`
- `double * b`
- `int S`
- `double one2many_small [2]`
- `double one2many_large [2]`
- `int M1`
- `int M2`
- `double many2one_small [2]`
- `double many2one_large [2]`

#### 8.13.1 Detailed Description

The LMO model

#### 8.13.2 Member Data Documentation

##### 8.13.2.1 CPM\_predictor LMO\_model::predictor

Predictor

##### 8.13.2.2 int LMO\_model::n

number of nodes

**8.13.2.3 double\* LMO\_model::C**

array of n average fixed processing delays

**8.13.2.4 double\* LMO\_model::L**array of  $C_n^2$  average to/from latencies (  $L_{ij} = L_{ji}$  )**8.13.2.5 double\* LMO\_model::t**

array of n average variable processing delays

**8.13.2.6 double\* LMO\_model::b**array of  $C_n^2$  average to/from transmission rates (  $\beta_{ij} = \beta_{ji}$  )**8.13.2.7 int LMO\_model::S**

threshold between small and large message sizes for the one2many model

**8.13.2.8 double LMO\_model::one2many\_small[2]**

one2many linear model for small message sizes

**8.13.2.9 double LMO\_model::one2many\_large[2]**

one2many linear model for large message sizes

**8.13.2.10 int LMO\_model::M1**

threshold between small and medium message sizes for the many2one model

**8.13.2.11 int LMO\_model::M2**

threshold between medium and large message sizes for the many2one model

**8.13.2.12 double LMO\_model::many2one\_small[2]**

many2one linear model for small message sizes

**8.13.2.13 double LMO\_model::many2one\_large[2]**

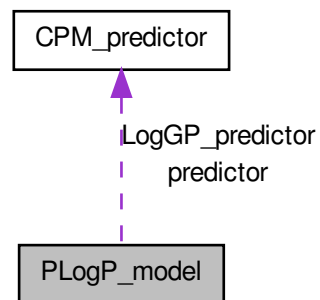
many2one linear model for large message sizes

The documentation for this struct was generated from the following file:

- models/lmo.h

## 8.14 PLogP\_model Struct Reference

Collaboration diagram for PLogP\_model:



### Public Attributes

- [CPM\\_predictor predictor](#)
- [CPM\\_predictor LogGP\\_predictor](#)
- `int n`
- `void ** logp`

### 8.14.1 Detailed Description

The heterogeneous PLogP model. The point-to-point parameters for all pairs of processors.

### 8.14.2 Member Data Documentation

#### 8.14.2.1 CPM\_predictor PLogP\_model::predictor

Predictor

#### 8.14.2.2 CPM\_predictor PLogP\_model::LogGP\_predictor

LogGP predictor

#### 8.14.2.3 `int PLogP_model::n`

Number of nodes

#### 8.14.2.4 void\*\* PLogP\_model::logp

Array of  $C_n^2$  pointers to the PLogP parameters, which correspond to all pairs of processors. Each pointer refers to the logp\_params data structure of the logp\_mpi library.

The documentation for this struct was generated from the following file:

- models/plogp.h

## 8.15 plogp\_p2p\_model Class Reference

### Public Member Functions

- template<class Archive >  
void [serialize](#) (Archive &ar, const unsigned int version)

#### 8.15.1 Detailed Description

The PLogP P2P model

#### 8.15.2 Member Function Documentation

##### 8.15.2.1 template<class Archive > void plogp\_p2p\_model::serialize ( Archive & ar, const unsigned int version ) [inline]

Serializes the model

The documentation for this class was generated from the following file:

- models/hierarchical/plogp\_p2p\_model.hpp

## 8.16 CPM::second\_greater Struct Reference

### 8.16.1 Detailed Description

Auxiliary function object that compares the second field of a pair.

The documentation for this struct was generated from the following file:

- collectives/cpm\_functional.hpp

## 8.17 CPM::second\_less Struct Reference

### 8.17.1 Detailed Description

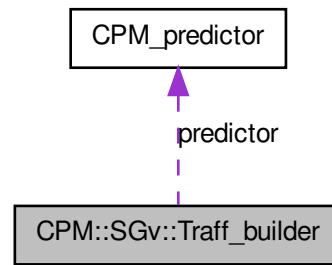
Auxiliary function object that compares the second field of a pair.

The documentation for this struct was generated from the following file:

- collectives/cpm\_functional.hpp

## 8.18 CPM::SGv::Traff\_builder Class Reference

Collaboration diagram for CPM::SGv::Traff\_builder:



### Public Member Functions

- void `build` (int size, int root, int rank, int \*counts, Graph &g, Vertex &r, Vertex &u, Vertex &v)

#### 8.18.1 Detailed Description

Model-based builder of Traff communication tree for scatterv/gatherv

#### 8.18.2 Member Function Documentation

**8.18.2.1** void CPM::SGv::Traff\_builder::build ( int size, int root, int rank, int \* counts, Graph & g, Vertex & r, Vertex & u, Vertex & v ) [inline]

The routine build is based on two phases which take place in a loop:

1. PHASE: in the first phase, the given set  $S_i$  is partitioned into subsets in following way:

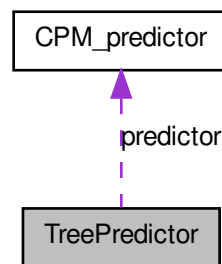
- if  $S_i$  is not empty:
  - sort the ranks from  $S_i$  in decreasing order according to the value of  $\text{predict\_p2p}(i, \text{parent}(S_i), m\_size(i))$
  - (\*) partition  $S_i$  into subsets so that for every subset  $W(S_{j1}) \geq \text{Sum}(j2=0..j1-1) W(S_{j2})$  and  $j1$  is minimal.  $\rightarrow W(S) := \text{sum}(i \text{ in } S) \text{ predict\_p2p}(i, \text{parent}(S), m\_size(i))$ ;
- 2. PHASE: in the second phase, determine the root of each new subset from 1(\*) in the following way:
  - in a routine `getRoot`, find the element of the subset to become root. For the optimal case, that routine would be:
    - find  $i$  from  $S_i$  so that  $\text{predict\_p2p}(i, \text{parent}(S_i), \text{Sum}(j \text{ in } S_i) m\_size(j)) = \text{MINIMAL}$
    - make  $i$  the root of  $S_i$  and continue from 1. with the children of  $S_i$  (remark: this root may be different than the first element (and so different than Traff's root))

The documentation for this class was generated from the following file:

- models/cpm\_sgv\_traff\_tree\_builder.hpp

## 8.19 TreePredictor Class Reference

Collaboration diagram for TreePredictor:



### 8.19.1 Detailed Description

The following collective prediction is based on our implementation with the recursive formula  $T(r) = \text{predict\_p2p}(r, \text{getChild}(r)) + \max(T(\text{getChild}(r)), T(r \setminus \langle \text{subtree with root } \text{getChild}(r) \rangle))$  if  $r$  is not a leaf where  $\text{getChild}(r) = \text{rightmost child of } r$  if R2L & TOP-DOWN = BCAST  $\text{getChild}(r) = \text{leftmost child of } r$  if L2R & DOWN-TOP = REDUCE  $T(r) = 0$  if  $r$  is a leaf (probably too idealistic) assumptions:

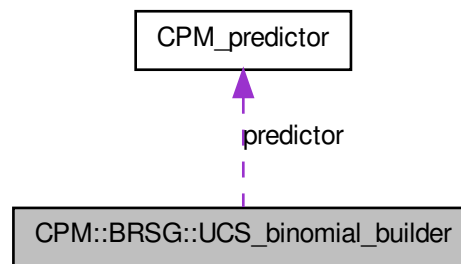
- a P2P communication (a,b) takes places exactly when both a and b are not involved in any other P2P communication
- any two pairs (a,b) and (c,d) can communicate in parallel
- symmetry -  $\text{predict\_p2p}(a,b) = \text{predict\_p2p}(b,a)$  - otherwise bcast and reduce would have different implementations (reverse directions)

The documentation for this class was generated from the following file:

- models/cpm\_predict\_trees.hpp

## 8.20 CPM::BRSG::UCS\_binomial\_builder Class Reference

Collaboration diagram for CPM::BRSG::UCS\_binomial\_builder:



### 8.20.1 Detailed Description

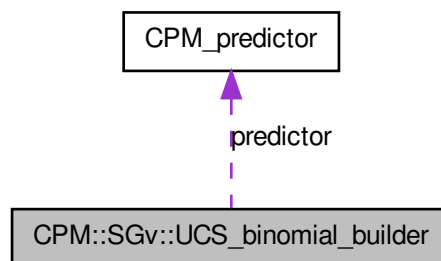
UCS binomial tree builder. The largest subtree on the left. If next\_node == MIN, the processor i with minimum p2p prediction is selected; otherwise, with maximum.

The documentation for this class was generated from the following file:

- models/cpm\_brsg\_tree\_builders.hpp

## 8.21 CPM::SGv::UCS\_binomial\_builder Class Reference

Collaboration diagram for CPM::SGv::UCS\_binomial\_builder:



### 8.21.1 Detailed Description

UCS binomial builder for scatterv/gatherv. The largest subtree on the left. Builds a binomial communication tree by subtrees of the same order, starting with the highest one. If `next_node == MIN`, the processor `i` with minimum p2p prediction is selected; otherwise, with maximum.

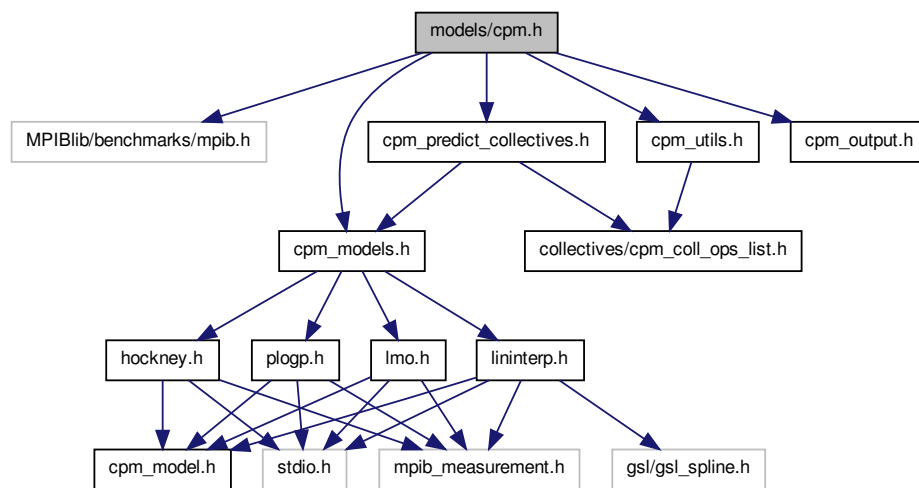
The documentation for this class was generated from the following file:

- `models/cpm_sgv_ucs_tree_builder.hpp`

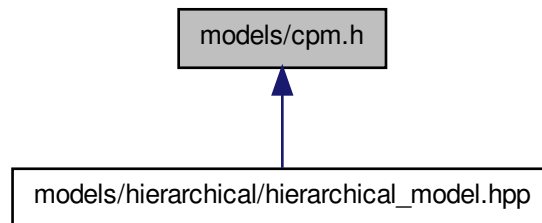
## 9 File Documentation

### 9.1 `models/cpm.h` File Reference

Include dependency graph for `cpm.h`:



This graph shows which files directly or indirectly include this file:



### 9.1.1 Detailed Description

For internal use

## Index

- a
  - Hockney\_model, 69
- accels
  - LinInterp\_model, 71
- b
  - Hockney\_model, 69
  - LMO\_model, 73
- build
  - CPM::SGv::Traff\_builder, 76
- C
  - LMO\_model, 72
  - Communication performance models, 13
  - CPM::BRSG, 61
  - CPM::BRSG::BFS\_binomial\_builder, 62
  - CPM::BRSG::DFS\_binomial\_builder, 67
  - CPM::BRSG::UCS\_binomial\_builder, 78
  - CPM::hierarchy, 61
  - CPM::second\_greater, 75
  - CPM::second\_less, 75
  - CPM::SGv, 61
  - CPM::SGv::BFS\_binomial\_builder, 62
  - CPM::SGv::DFS\_binomial\_builder, 66
  - CPM::SGv::Traff\_builder, 76
    - build, 76
  - CPM::SGv::UCS\_binomial\_builder, 78
  - CPM\_Bcast\_bfs\_binomial
    - cpm\_collectives, 29
  - CPM\_Bcast\_dfs\_binomial
    - cpm\_collectives, 28
  - CPM\_Bcast\_ucs\_binomial
    - cpm\_collectives, 29
  - cpm\_collectives
    - CPM\_Bcast\_bfs\_binomial, 29
    - CPM\_Bcast\_dfs\_binomial, 28
    - CPM\_Bcast\_ucs\_binomial, 29
    - CPM\_Gather\_bfs\_binomial, 30
    - CPM\_Gather\_dfs\_binomial, 30
    - CPM\_Gather\_ucs\_binomial, 29
    - CPM\_Gatherv\_bfs\_binomial, 30
    - CPM\_Gatherv\_dfs\_binomial, 30
    - CPM\_Gatherv\_sorted\_flat, 28
    - CPM\_Gatherv\_Traff, 31
    - CPM\_Gatherv\_ucs\_binomial, 31
    - CPM\_Reduce\_bfs\_binomial, 29
    - CPM\_Reduce\_dfs\_binomial, 29
    - CPM\_Reduce\_ucs\_binomial, 29
    - CPM\_Scatter\_bfs\_binomial, 29
    - CPM\_Scatter\_dfs\_binomial, 30
    - CPM\_Scatter\_ucs\_binomial, 29
    - CPM\_Scatterv\_bfs\_binomial, 30
    - CPM\_Scatterv\_dfs\_binomial, 30
    - CPM\_Scatterv\_sorted\_flat, 28
    - CPM\_Scatterv\_Traff, 31
    - CPM\_Scatterv\_ucs\_binomial, 30
  - CPM\_Gather\_bfs\_binomial
    - cpm\_collectives, 30
  - CPM\_Gather\_dfs\_binomial
    - cpm\_collectives, 30
  - CPM\_Gather\_ucs\_binomial
    - cpm\_collectives, 29
  - CPM\_Gatherv\_bfs\_binomial
    - cpm\_collectives, 30
  - CPM\_Gatherv\_dfs\_binomial
    - cpm\_collectives, 30
  - CPM\_Gatherv\_sorted\_flat
    - cpm\_collectives, 28
  - CPM\_Gatherv\_Traff
    - cpm\_collectives, 31
  - CPM\_Gatherv\_ucs\_binomial
    - cpm\_collectives, 31
  - CPM\_measure\_p2pp
    - measurement, 12
  - CPM\_measure\_p2pp\_p2p
    - measurement, 13
  - CPM\_predict\_brsg
    - prediction, 14
  - CPM\_predict\_flat\_sg
    - prediction, 15
  - CPM\_predict\_flat\_sg\_parallel
    - prediction, 15
  - CPM\_predict\_flat\_sg\_serial
    - prediction, 15
  - CPM\_predict\_flat\_sgv
    - prediction, 15
  - CPM\_predict\_flat\_sgv\_parallel
    - prediction, 15
  - CPM\_predict\_flat\_sgv\_serial
    - prediction, 15
  - CPM\_predict\_sgv
    - prediction, 14
  - CPM\_predictor, 63
    - predict\_p2p, 63
  - CPM\_Reduce\_bfs\_binomial
    - cpm\_collectives, 29
  - CPM\_Reduce\_dfs\_binomial
    - cpm\_collectives, 29
  - CPM\_Reduce\_ucs\_binomial
    - cpm\_collectives, 29
  - CPM\_Scatter\_bfs\_binomial

- cpm\_collectives, 29
- CPM\_Scatter\_dfs\_binomial
  - cpm\_collectives, 30
- CPM\_Scatter\_ucs\_binomial
  - cpm\_collectives, 29
- CPM\_Scatterv\_bfs\_binomial
  - cpm\_collectives, 30
- CPM\_Scatterv\_dfs\_binomial
  - cpm\_collectives, 30
- CPM\_Scatterv\_sorted\_flat
  - cpm\_collectives, 28
- CPM\_Scatterv\_Traff
  - cpm\_collectives, 31
- CPM\_Scatterv\_ucs\_binomial
  - cpm\_collectives, 30
- CPM\_SGv\_sorter, 64
- CPM\_triplet, 64
  - next, 65
  - prev, 65
  - values, 65
- CPM\_triplets, 65
  - list, 66
  - next, 66
  - prev, 66
- estimate
  - hockney\_p2p\_model, 70
- Generic model-based collectives, 26
- H\_Model, 67
  - read\_hierarchy\_file, 68
- hockney
  - Hockney\_alloc, 16
  - Hockney\_estimate, 16
  - Hockney\_estimate\_regression, 17
  - Hockney\_free, 16
  - Hockney\_hpredict\_sg\_binomial, 18
  - Hockney\_hpredict\_sg\_flat\_parallel, 18
  - Hockney\_hpredict\_sg\_flat\_serial, 18
  - Hockney\_predict\_p2p, 17
  - Hockney\_predict\_sg\_binomial, 18
  - Hockney\_predict\_sg\_flat\_parallel, 18
  - Hockney\_predict\_sg\_flat\_serial, 17
  - Hockney\_read, 16
  - Hockney\_write, 16
- Hockney-based collective operations, 31
- Hockney\_alloc
  - hockney, 16
- Hockney\_Bcast\_bfs\_binomial\_max
  - hockney\_collectives, 34
- Hockney\_Bcast\_bfs\_binomial\_min
  - hockney\_collectives, 34
- Hockney\_Bcast\_dfs\_binomial\_max
  - hockney\_collectives, 35
- Hockney\_Bcast\_dfs\_binomial\_min
  - hockney\_collectives, 35
- hockney\_collectives
  - Hockney\_Bcast\_bfs\_binomial\_max, 34
  - Hockney\_Bcast\_bfs\_binomial\_min, 34
  - Hockney\_Bcast\_dfs\_binomial\_max, 34
  - Hockney\_Bcast\_dfs\_binomial\_min, 34
  - Hockney\_Bcast\_ucs\_binomial\_max, 35
  - Hockney\_Bcast\_ucs\_binomial\_min, 35
  - Hockney\_finalize, 33
  - Hockney\_Gather\_bfs\_binomial\_max, 36
  - Hockney\_Gather\_bfs\_binomial\_min, 36
  - Hockney\_Gather\_dfs\_binomial\_max, 35
  - Hockney\_Gather\_dfs\_binomial\_min, 35
  - Hockney\_Gather\_ucs\_binomial\_max, 36
  - Hockney\_Gather\_ucs\_binomial\_min, 36
  - Hockney\_Gatherv\_bfs\_binomial\_max, 37
  - Hockney\_Gatherv\_bfs\_binomial\_min, 37
  - Hockney\_Gatherv\_dfs\_binomial\_max, 37
  - Hockney\_Gatherv\_dfs\_binomial\_min, 37
  - Hockney\_Gatherv\_sorted\_flat\_asc, 33
  - Hockney\_Gatherv\_sorted\_flat\_dsc, 34
  - Hockney\_Gatherv\_Traff, 38
  - Hockney\_Gatherv\_ucs\_binomial\_max, 38
  - Hockney\_Gatherv\_ucs\_binomial\_min, 38
  - Hockney\_initialize, 33
  - Hockney\_Reduce\_bfs\_binomial\_max, 34
  - Hockney\_Reduce\_bfs\_binomial\_min, 34
  - Hockney\_Reduce\_dfs\_binomial\_max, 34
  - Hockney\_Reduce\_dfs\_binomial\_min, 34
  - Hockney\_Reduce\_ucs\_binomial\_max, 35
  - Hockney\_Reduce\_ucs\_binomial\_min, 35
  - Hockney\_Scatter\_bfs\_binomial\_max, 36
  - Hockney\_Scatter\_bfs\_binomial\_min, 36
  - Hockney\_Scatter\_dfs\_binomial\_max, 35
  - Hockney\_Scatter\_dfs\_binomial\_min, 35
  - Hockney\_Scatter\_ucs\_binomial\_max, 36
  - Hockney\_Scatter\_ucs\_binomial\_min, 35
  - Hockney\_Scatterv\_bfs\_binomial\_max, 37
  - Hockney\_Scatterv\_bfs\_binomial\_min, 37
  - Hockney\_Scatterv\_dfs\_binomial\_max, 37
  - Hockney\_Scatterv\_dfs\_binomial\_min, 36
  - Hockney\_Scatterv\_sorted\_flat\_asc, 33
  - Hockney\_Scatterv\_sorted\_flat\_dsc, 33
  - Hockney\_Scatterv\_Traff, 38
  - Hockney\_Scatterv\_ucs\_binomial\_max, 38
  - Hockney\_Scatterv\_ucs\_binomial\_min, 37
- Hockney\_estimate
  - hockney, 16

- Hockney\_estimate\_regression
  - hockney, [17](#)
- Hockney\_finalize
  - hockney\_collectives, [33](#)
- Hockney\_free
  - hockney, [16](#)
- Hockney\_Gather\_bfs\_binomial\_max
  - hockney\_collectives, [36](#)
- Hockney\_Gather\_bfs\_binomial\_min
  - hockney\_collectives, [36](#)
- Hockney\_Gather\_dfs\_binomial\_max
  - hockney\_collectives, [35](#)
- Hockney\_Gather\_dfs\_binomial\_min
  - hockney\_collectives, [35](#)
- Hockney\_Gather\_ucs\_binomial\_max
  - hockney\_collectives, [36](#)
- Hockney\_Gather\_ucs\_binomial\_min
  - hockney\_collectives, [36](#)
- Hockney\_Gatherv\_bfs\_binomial\_max
  - hockney\_collectives, [37](#)
- Hockney\_Gatherv\_bfs\_binomial\_min
  - hockney\_collectives, [37](#)
- Hockney\_Gatherv\_dfs\_binomial\_max
  - hockney\_collectives, [37](#)
- Hockney\_Gatherv\_dfs\_binomial\_min
  - hockney\_collectives, [37](#)
- Hockney\_Gatherv\_sorted\_flat\_asc
  - hockney\_collectives, [33](#)
- Hockney\_Gatherv\_sorted\_flat\_dsc
  - hockney\_collectives, [34](#)
- Hockney\_Gatherv\_Traff
  - hockney\_collectives, [38](#)
- Hockney\_Gatherv\_ucs\_binomial\_max
  - hockney\_collectives, [38](#)
- Hockney\_Gatherv\_ucs\_binomial\_min
  - hockney\_collectives, [38](#)
- Hockney\_hpredict\_sg\_binomial
  - hockney, [18](#)
- Hockney\_hpredict\_sg\_flat\_parallel
  - hockney, [18](#)
- Hockney\_hpredict\_sg\_flat\_serial
  - hockney, [18](#)
- Hockney\_initialize
  - hockney\_collectives, [33](#)
- Hockney\_model, [68](#)
  - a, [69](#)
  - b, [69](#)
  - n, [69](#)
  - predictor, [69](#)
- hockney\_p2p\_model, [69](#)
  - estimate, [70](#)
  - serialize, [70](#)
  - write\_elem, [70](#)
  - write\_footer, [70](#)
  - write\_header, [70](#)
- Hockney\_predict\_p2p
  - hockney, [17](#)
- Hockney\_predict\_sg\_binomial
  - hockney, [18](#)
- Hockney\_predict\_sg\_flat\_parallel
  - hockney, [18](#)
- Hockney\_predict\_sg\_flat\_serial
  - hockney, [17](#)
- Hockney\_read
  - hockney, [16](#)
- Hockney\_Reduce\_bfs\_binomial\_max
  - hockney\_collectives, [34](#)
- Hockney\_Reduce\_bfs\_binomial\_min
  - hockney\_collectives, [34](#)
- Hockney\_Reduce\_dfs\_binomial\_max
  - hockney\_collectives, [34](#)
- Hockney\_Reduce\_dfs\_binomial\_min
  - hockney\_collectives, [34](#)
- Hockney\_Reduce\_ucs\_binomial\_max
  - hockney\_collectives, [35](#)
- Hockney\_Reduce\_ucs\_binomial\_min
  - hockney\_collectives, [35](#)
- Hockney\_Scatter\_bfs\_binomial\_max
  - hockney\_collectives, [36](#)
- Hockney\_Scatter\_bfs\_binomial\_min
  - hockney\_collectives, [36](#)
- Hockney\_Scatter\_dfs\_binomial\_max
  - hockney\_collectives, [35](#)
- Hockney\_Scatter\_dfs\_binomial\_min
  - hockney\_collectives, [35](#)
- Hockney\_Scatter\_ucs\_binomial\_max
  - hockney\_collectives, [36](#)
- Hockney\_Scatter\_ucs\_binomial\_min
  - hockney\_collectives, [35](#)
- Hockney\_Scatterv\_bfs\_binomial\_max
  - hockney\_collectives, [37](#)
- Hockney\_Scatterv\_bfs\_binomial\_min
  - hockney\_collectives, [37](#)
- Hockney\_Scatterv\_dfs\_binomial\_max
  - hockney\_collectives, [37](#)
- Hockney\_Scatterv\_dfs\_binomial\_min
  - hockney\_collectives, [36](#)
- Hockney\_Scatterv\_sorted\_flat\_asc
  - hockney\_collectives, [33](#)
- Hockney\_Scatterv\_sorted\_flat\_dsc
  - hockney\_collectives, [33](#)
- Hockney\_Scatterv\_Traff
  - hockney\_collectives, [38](#)
- Hockney\_Scatterv\_ucs\_binomial\_max
  - hockney\_collectives, [38](#)
- Hockney\_Scatterv\_ucs\_binomial\_min
  - hockney\_collectives, [37](#)
- Hockney\_write

- hockney, 16
- L
  - LMO\_model, 73
  - lininterp
    - LinInterp\_alloc, 19
    - LinInterp\_estimate, 19
    - LinInterp\_free, 19
    - LinInterp\_predict\_p2p, 19
    - LinInterp\_read, 19
    - LinInterp\_write, 19
  - LinInterp-based collective operations, 38
  - LinInterp\_alloc
    - lininterp, 19
  - LinInterp\_Bcast\_bfs\_binomial\_max
    - lininterp\_collectives, 42
  - LinInterp\_Bcast\_bfs\_binomial\_min
    - lininterp\_collectives, 42
  - LinInterp\_Bcast\_dfs\_binomial\_max
    - lininterp\_collectives, 41
  - LinInterp\_Bcast\_dfs\_binomial\_min
    - lininterp\_collectives, 41
  - LinInterp\_Bcast\_ucs\_binomial\_max
    - lininterp\_collectives, 42
  - LinInterp\_Bcast\_ucs\_binomial\_min
    - lininterp\_collectives, 42
  - lininterp\_collectives
    - LinInterp\_Bcast\_bfs\_binomial\_max, 42
    - LinInterp\_Bcast\_bfs\_binomial\_min, 42
    - LinInterp\_Bcast\_dfs\_binomial\_max, 41
    - LinInterp\_Bcast\_dfs\_binomial\_min, 41
    - LinInterp\_Bcast\_ucs\_binomial\_max, 42
    - LinInterp\_Bcast\_ucs\_binomial\_min, 42
    - LinInterp\_finalize, 40
    - LinInterp\_Gather\_bfs\_binomial\_max, 44
    - LinInterp\_Gather\_bfs\_binomial\_min, 43
    - LinInterp\_Gather\_dfs\_binomial\_max, 43
    - LinInterp\_Gather\_dfs\_binomial\_min, 42
    - LinInterp\_Gather\_ucs\_binomial\_max, 43
    - LinInterp\_Gather\_ucs\_binomial\_min, 43
    - LinInterp\_Gatherv\_bfs\_binomial\_max, 45
    - LinInterp\_Gatherv\_bfs\_binomial\_min, 44
    - LinInterp\_Gatherv\_dfs\_binomial\_max, 44
    - LinInterp\_Gatherv\_dfs\_binomial\_min, 44
    - LinInterp\_Gatherv\_sorted\_flat\_asc, 41
    - LinInterp\_Gatherv\_sorted\_flat\_dsc, 41
    - LinInterp\_Gatherv\_Traff, 45
    - LinInterp\_Gatherv\_ucs\_binomial\_max, 45
    - LinInterp\_Gatherv\_ucs\_binomial\_min, 45
    - LinInterp\_initialize, 40
    - LinInterp\_Reduce\_bfs\_binomial\_max, 41
    - LinInterp\_Reduce\_bfs\_binomial\_min, 41
    - LinInterp\_Reduce\_dfs\_binomial\_max, 42
    - LinInterp\_Reduce\_dfs\_binomial\_min, 41
    - LinInterp\_Reduce\_ucs\_binomial\_max, 42
    - LinInterp\_Reduce\_ucs\_binomial\_min, 42
    - LinInterp\_Scatter\_bfs\_binomial\_max, 44
    - LinInterp\_Scatter\_bfs\_binomial\_min, 43
    - LinInterp\_Scatter\_dfs\_binomial\_max, 43
    - LinInterp\_Scatter\_dfs\_binomial\_min, 42
    - LinInterp\_Scatter\_ucs\_binomial\_max, 43
    - LinInterp\_Scatter\_ucs\_binomial\_min, 43
    - LinInterp\_Scatterv\_bfs\_binomial\_max, 45
    - LinInterp\_Scatterv\_bfs\_binomial\_min, 44
    - LinInterp\_Scatterv\_dfs\_binomial\_max, 44
    - LinInterp\_Scatterv\_dfs\_binomial\_min, 44
    - LinInterp\_Scatterv\_sorted\_flat\_asc, 41
    - LinInterp\_Scatterv\_sorted\_flat\_dsc, 41
    - LinInterp\_Scatterv\_Traff, 45
    - LinInterp\_Scatterv\_ucs\_binomial\_max, 45
    - LinInterp\_Scatterv\_ucs\_binomial\_min, 45
  - LinInterp\_estimate
    - lininterp, 19
  - LinInterp\_finalize
    - lininterp\_collectives, 40
  - LinInterp\_free
    - lininterp, 19
  - LinInterp\_Gather\_bfs\_binomial\_max
    - lininterp\_collectives, 44
  - LinInterp\_Gather\_bfs\_binomial\_min
    - lininterp\_collectives, 43
  - LinInterp\_Gather\_dfs\_binomial\_max
    - lininterp\_collectives, 43
  - LinInterp\_Gather\_dfs\_binomial\_min
    - lininterp\_collectives, 42
  - LinInterp\_Gather\_ucs\_binomial\_max
    - lininterp\_collectives, 43
  - LinInterp\_Gather\_ucs\_binomial\_min
    - lininterp\_collectives, 43
  - LinInterp\_Gatherv\_bfs\_binomial\_max
    - lininterp\_collectives, 45
  - LinInterp\_Gatherv\_bfs\_binomial\_min
    - lininterp\_collectives, 44
  - LinInterp\_Gatherv\_dfs\_binomial\_max
    - lininterp\_collectives, 44
  - LinInterp\_Gatherv\_dfs\_binomial\_min
    - lininterp\_collectives, 44
  - LinInterp\_Gatherv\_sorted\_flat\_asc
    - lininterp\_collectives, 41
  - LinInterp\_Gatherv\_sorted\_flat\_dsc
    - lininterp\_collectives, 41
  - LinInterp\_Gatherv\_Traff
    - lininterp\_collectives, 45
  - LinInterp\_Gatherv\_ucs\_binomial\_max
    - lininterp\_collectives, 45
  - LinInterp\_Gatherv\_ucs\_binomial\_min
    - lininterp\_collectives, 45
  - LinInterp\_initialize

- lininterp\_collectives, 40
- LinInterp\_model, 70
  - accels, 71
  - n, 71
  - num\_points, 71
  - predictor, 71
  - splines, 71
- LinInterp\_predict\_p2p
  - lininterp, 19
- LinInterp\_read
  - lininterp, 19
- LinInterp\_Reduce\_bfs\_binomial\_max
  - lininterp\_collectives, 41
- LinInterp\_Reduce\_bfs\_binomial\_min
  - lininterp\_collectives, 41
- LinInterp\_Reduce\_dfs\_binomial\_max
  - lininterp\_collectives, 42
- LinInterp\_Reduce\_dfs\_binomial\_min
  - lininterp\_collectives, 41
- LinInterp\_Reduce\_ucs\_binomial\_max
  - lininterp\_collectives, 42
- LinInterp\_Reduce\_ucs\_binomial\_min
  - lininterp\_collectives, 42
- LinInterp\_Scatter\_bfs\_binomial\_max
  - lininterp\_collectives, 44
- LinInterp\_Scatter\_bfs\_binomial\_min
  - lininterp\_collectives, 43
- LinInterp\_Scatter\_dfs\_binomial\_max
  - lininterp\_collectives, 43
- LinInterp\_Scatter\_dfs\_binomial\_min
  - lininterp\_collectives, 42
- LinInterp\_Scatter\_ucs\_binomial\_max
  - lininterp\_collectives, 43
- LinInterp\_Scatter\_ucs\_binomial\_min
  - lininterp\_collectives, 43
- LinInterp\_Scatterv\_bfs\_binomial\_max
  - lininterp\_collectives, 45
- LinInterp\_Scatterv\_bfs\_binomial\_min
  - lininterp\_collectives, 44
- LinInterp\_Scatterv\_dfs\_binomial\_max
  - lininterp\_collectives, 44
- LinInterp\_Scatterv\_dfs\_binomial\_min
  - lininterp\_collectives, 44
- LinInterp\_Scatterv\_sorted\_flat\_asc
  - lininterp\_collectives, 41
- LinInterp\_Scatterv\_sorted\_flat\_dsc
  - lininterp\_collectives, 41
- LinInterp\_Scatterv\_Traff
  - lininterp\_collectives, 45
- LinInterp\_Scatterv\_ucs\_binomial\_max
  - lininterp\_collectives, 45
- LinInterp\_Scatterv\_ucs\_binomial\_min
  - lininterp\_collectives, 45
- LinInterp\_write
  - lininterp, 19
- list
  - CPM\_triplets, 66
- lmo
  - LMO\_alloc, 20
  - LMO\_C3, 20
  - LMO\_estimate, 21
  - LMO\_estimate\_homogeneous, 22
  - LMO\_free, 20
  - LMO\_IJK2INDEX, 20
  - LMO\_predict\_gather\_flat, 23
  - LMO\_predict\_p2p, 22
  - LMO\_predict\_scatter\_flat, 23
  - LMO\_read, 20
  - LMO\_write, 20
- LMO-based collective operations, 46
- LMO: an advanced heterogeneous communication performance model, 19
- LMO\_alloc
  - lmo, 20
- LMO\_C3
  - lmo, 20
- lmo\_collectives
  - LMO\_finalize, 46
  - LMO\_Gather\_split\_flat, 46
  - LMO\_initialize, 46
  - LMO\_Scatter\_split\_flat, 46
- LMO\_estimate
  - lmo, 21
- LMO\_estimate\_homogeneous
  - lmo, 22
- LMO\_finalize
  - lmo\_collectives, 46
- LMO\_free
  - lmo, 20
- LMO\_Gather\_split\_flat
  - lmo\_collectives, 46
- LMO\_IJK2INDEX
  - lmo, 20
- LMO\_initialize
  - lmo\_collectives, 46
- LMO\_model, 72
  - b, 73
  - C, 72
  - L, 73
  - M1, 73
  - M2, 73
  - many2one\_large, 73
  - many2one\_small, 73
  - n, 72
  - one2many\_large, 73
  - one2many\_small, 73
  - predictor, 72
  - S, 73

- t, [73](#)
- LMO\_predict\_gather\_flat
  - lmo, [23](#)
- LMO\_predict\_p2p
  - lmo, [22](#)
- LMO\_predict\_scatter\_flat
  - lmo, [23](#)
- LMO\_read
  - lmo, [20](#)
- LMO\_Scatter\_split\_flat
  - lmo\_collectives, [46](#)
- LMO\_write
  - lmo, [20](#)
- LogGP-based collective operations, [46](#)
- LogGP\_Bcast\_bfs\_binomial\_max
  - loggp\_collectives, [49](#)
- LogGP\_Bcast\_bfs\_binomial\_min
  - loggp\_collectives, [49](#)
- LogGP\_Bcast\_dfs\_binomial\_max
  - loggp\_collectives, [49](#)
- LogGP\_Bcast\_dfs\_binomial\_min
  - loggp\_collectives, [49](#)
- LogGP\_Bcast\_ucs\_binomial\_max
  - loggp\_collectives, [50](#)
- LogGP\_Bcast\_ucs\_binomial\_min
  - loggp\_collectives, [49](#)
- loggp\_collectives
  - LogGP\_Bcast\_bfs\_binomial\_max, [49](#)
  - LogGP\_Bcast\_bfs\_binomial\_min, [49](#)
  - LogGP\_Bcast\_dfs\_binomial\_max, [49](#)
  - LogGP\_Bcast\_dfs\_binomial\_min, [49](#)
  - LogGP\_Bcast\_ucs\_binomial\_max, [50](#)
  - LogGP\_Bcast\_ucs\_binomial\_min, [49](#)
  - LogGP\_Gather\_bfs\_binomial\_max, [51](#)
  - LogGP\_Gather\_bfs\_binomial\_min, [51](#)
  - LogGP\_Gather\_dfs\_binomial\_max, [50](#)
  - LogGP\_Gather\_dfs\_binomial\_min, [50](#)
  - LogGP\_Gather\_ucs\_binomial\_max, [51](#)
  - LogGP\_Gather\_ucs\_binomial\_min, [51](#)
  - LogGP\_Gatherv\_bfs\_binomial\_max, [52](#)
  - LogGP\_Gatherv\_bfs\_binomial\_min, [52](#)
  - LogGP\_Gatherv\_dfs\_binomial\_max, [52](#)
  - LogGP\_Gatherv\_dfs\_binomial\_min, [52](#)
  - LogGP\_Gatherv\_sorted\_flat\_asc, [48](#)
  - LogGP\_Gatherv\_sorted\_flat\_dsc, [48](#)
  - LogGP\_Gatherv\_Traff, [53](#)
  - LogGP\_Gatherv\_ucs\_binomial\_max, [53](#)
  - LogGP\_Gatherv\_ucs\_binomial\_min, [53](#)
  - LogGP\_Reduce\_bfs\_binomial\_max, [49](#)
  - LogGP\_Reduce\_bfs\_binomial\_min, [49](#)
  - LogGP\_Reduce\_dfs\_binomial\_max, [49](#)
  - LogGP\_Reduce\_dfs\_binomial\_min, [49](#)
  - LogGP\_Reduce\_ucs\_binomial\_max, [50](#)
  - LogGP\_Reduce\_ucs\_binomial\_min, [50](#)
  - LogGP\_Scatter\_bfs\_binomial\_max, [51](#)
  - LogGP\_Scatter\_bfs\_binomial\_min, [51](#)
  - LogGP\_Scatter\_dfs\_binomial\_max, [50](#)
  - LogGP\_Scatter\_dfs\_binomial\_min, [50](#)
  - LogGP\_Scatter\_ucs\_binomial\_max, [51](#)
  - LogGP\_Scatter\_ucs\_binomial\_min, [50](#)
  - LogGP\_Scatterv\_bfs\_binomial\_max, [52](#)
  - LogGP\_Scatterv\_bfs\_binomial\_min, [52](#)
  - LogGP\_Scatterv\_dfs\_binomial\_max, [52](#)
  - LogGP\_Scatterv\_dfs\_binomial\_min, [51](#)
  - LogGP\_Scatterv\_sorted\_flat\_asc, [48](#)
  - LogGP\_Scatterv\_sorted\_flat\_dsc, [48](#)
  - LogGP\_Scatterv\_Traff, [53](#)
  - LogGP\_Scatterv\_ucs\_binomial\_max, [53](#)
  - LogGP\_Scatterv\_ucs\_binomial\_min, [52](#)
  - LogGP\_Gather\_bfs\_binomial\_max
    - loggp\_collectives, [51](#)
  - LogGP\_Gather\_bfs\_binomial\_min
    - loggp\_collectives, [51](#)
  - LogGP\_Gather\_dfs\_binomial\_max
    - loggp\_collectives, [50](#)
  - LogGP\_Gather\_dfs\_binomial\_min
    - loggp\_collectives, [50](#)
  - LogGP\_Gather\_ucs\_binomial\_max
    - loggp\_collectives, [51](#)
  - LogGP\_Gather\_ucs\_binomial\_min
    - loggp\_collectives, [51](#)
  - LogGP\_Gatherv\_bfs\_binomial\_max
    - loggp\_collectives, [52](#)
  - LogGP\_Gatherv\_bfs\_binomial\_min
    - loggp\_collectives, [52](#)
  - LogGP\_Gatherv\_dfs\_binomial\_max
    - loggp\_collectives, [52](#)
  - LogGP\_Gatherv\_dfs\_binomial\_min
    - loggp\_collectives, [52](#)
  - LogGP\_Gatherv\_sorted\_flat\_asc
    - loggp\_collectives, [48](#)
  - LogGP\_Gatherv\_sorted\_flat\_dsc
    - loggp\_collectives, [48](#)
  - LogGP\_Gatherv\_Traff
    - loggp\_collectives, [53](#)
  - LogGP\_Gatherv\_ucs\_binomial\_max
    - loggp\_collectives, [53](#)
  - LogGP\_Gatherv\_ucs\_binomial\_min
    - loggp\_collectives, [53](#)
  - LogGP\_hpredict\_sg\_linear
    - plogp, [26](#)
  - LogGP\_predict\_p2p
    - plogp, [25](#)
  - LogGP\_predictor
    - PLogP\_model, [74](#)
  - LogGP\_Reduce\_bfs\_binomial\_max
    - loggp\_collectives, [49](#)
  - LogGP\_Reduce\_bfs\_binomial\_min

- loggp\_collectives, 49
- LogGP\_Reduce\_dfs\_binomial\_max
  - loggp\_collectives, 49
- LogGP\_Reduce\_dfs\_binomial\_min
  - loggp\_collectives, 49
- LogGP\_Reduce\_ucs\_binomial\_max
  - loggp\_collectives, 50
- LogGP\_Reduce\_ucs\_binomial\_min
  - loggp\_collectives, 50
- LogGP\_Scatter\_bfs\_binomial\_max
  - loggp\_collectives, 51
- LogGP\_Scatter\_bfs\_binomial\_min
  - loggp\_collectives, 51
- LogGP\_Scatter\_dfs\_binomial\_max
  - loggp\_collectives, 50
- LogGP\_Scatter\_dfs\_binomial\_min
  - loggp\_collectives, 50
- LogGP\_Scatter\_ucs\_binomial\_max
  - loggp\_collectives, 51
- LogGP\_Scatter\_ucs\_binomial\_min
  - loggp\_collectives, 50
- LogGP\_Scatterv\_bfs\_binomial\_max
  - loggp\_collectives, 52
- LogGP\_Scatterv\_bfs\_binomial\_min
  - loggp\_collectives, 52
- LogGP\_Scatterv\_dfs\_binomial\_max
  - loggp\_collectives, 52
- LogGP\_Scatterv\_dfs\_binomial\_min
  - loggp\_collectives, 51
- LogGP\_Scatterv\_sorted\_flat\_asc
  - loggp\_collectives, 48
- LogGP\_Scatterv\_sorted\_flat\_dsc
  - loggp\_collectives, 48
- LogGP\_Scatterv\_Traff
  - loggp\_collectives, 53
- LogGP\_Scatterv\_ucs\_binomial\_max
  - loggp\_collectives, 53
- LogGP\_Scatterv\_ucs\_binomial\_min
  - loggp\_collectives, 52
- logp
  - PLogP\_model, 74
- M1
  - LMO\_model, 73
- M2
  - LMO\_model, 73
- many2one\_large
  - LMO\_model, 73
- many2one\_small
  - LMO\_model, 73
- measurement
  - CPM\_measure\_p2pp, 12
  - CPM\_measure\_p2pp\_p2p, 13
- Measurement: benchmarking specific communication experiments, 12
- models/cpm.h, 79
- n
  - Hockney\_model, 69
  - LinInterp\_model, 71
  - LMO\_model, 72
  - PLogP\_model, 74
- next
  - CPM\_triplet, 65
  - CPM\_triplets, 66
- num\_points
  - LinInterp\_model, 71
- one2many\_large
  - LMO\_model, 73
- one2many\_small
  - LMO\_model, 73
- plogp
  - LogGP\_hpredict\_sg\_linear, 26
  - LogGP\_predict\_p2p, 25
  - PLogP\_alloc, 24
  - PLogP\_estimate, 25
  - PLogP\_free, 24
  - PLogP\_hpredict\_sg\_linear, 25
  - PLogP\_predict\_p2p, 25
  - PLogP\_read, 25
  - PLogP\_write, 25
- PLogP-based collective operations, 53
- PLogP\_alloc
  - plogp, 24
- PLogP\_Bcast\_bfs\_binomial\_max
  - plogp\_collectives, 56
- PLogP\_Bcast\_bfs\_binomial\_min
  - plogp\_collectives, 56
- PLogP\_Bcast\_dfs\_binomial\_max
  - plogp\_collectives, 56
- PLogP\_Bcast\_dfs\_binomial\_min
  - plogp\_collectives, 56
- PLogP\_Bcast\_ucs\_binomial\_max
  - plogp\_collectives, 57
- PLogP\_Bcast\_ucs\_binomial\_min
  - plogp\_collectives, 57
- plogp\_collectives
  - PLogP\_Bcast\_bfs\_binomial\_max, 56
  - PLogP\_Bcast\_bfs\_binomial\_min, 56
  - PLogP\_Bcast\_dfs\_binomial\_max, 56
  - PLogP\_Bcast\_dfs\_binomial\_min, 56
  - PLogP\_Bcast\_ucs\_binomial\_max, 57
  - PLogP\_Bcast\_ucs\_binomial\_min, 57
  - PLogP\_finalize, 55
  - PLogP\_Gather\_bfs\_binomial\_max, 58

- PLogP\_Gather\_bfs\_binomial\_min, 58
- PLogP\_Gather\_dfs\_binomial\_max, 57
- PLogP\_Gather\_dfs\_binomial\_min, 57
- PLogP\_Gather\_ucs\_binomial\_max, 58
- PLogP\_Gather\_ucs\_binomial\_min, 58
- PLogP\_Gatherv\_bfs\_binomial\_max, 59
- PLogP\_Gatherv\_bfs\_binomial\_min, 59
- PLogP\_Gatherv\_dfs\_binomial\_max, 59
- PLogP\_Gatherv\_dfs\_binomial\_min, 59
- PLogP\_Gatherv\_sorted\_flat\_asc, 55
- PLogP\_Gatherv\_sorted\_flat\_dsc, 56
- PLogP\_Gatherv\_Traff, 60
- PLogP\_Gatherv\_ucs\_binomial\_max, 60
- PLogP\_Gatherv\_ucs\_binomial\_min, 60
- PLogP\_initialize, 55
- PLogP\_Reduce\_bfs\_binomial\_max, 56
- PLogP\_Reduce\_bfs\_binomial\_min, 56
- PLogP\_Reduce\_dfs\_binomial\_max, 57
- PLogP\_Reduce\_dfs\_binomial\_min, 56
- PLogP\_Reduce\_ucs\_binomial\_max, 57
- PLogP\_Reduce\_ucs\_binomial\_min, 57
- PLogP\_Scatter\_bfs\_binomial\_max, 58
- PLogP\_Scatter\_bfs\_binomial\_min, 58
- PLogP\_Scatter\_dfs\_binomial\_max, 57
- PLogP\_Scatter\_dfs\_binomial\_min, 57
- PLogP\_Scatter\_ucs\_binomial\_max, 58
- PLogP\_Scatter\_ucs\_binomial\_min, 58
- PLogP\_Scatterv\_bfs\_binomial\_max, 59
- PLogP\_Scatterv\_bfs\_binomial\_min, 59
- PLogP\_Scatterv\_dfs\_binomial\_max, 59
- PLogP\_Scatterv\_dfs\_binomial\_min, 59
- PLogP\_Scatterv\_sorted\_flat\_asc, 55
- PLogP\_Scatterv\_sorted\_flat\_dsc, 56
- PLogP\_Scatterv\_Traff, 60
- PLogP\_Scatterv\_ucs\_binomial\_max, 60
- PLogP\_Scatterv\_ucs\_binomial\_min, 60
- PLogP\_estimate
  - plogp, 25
- PLogP\_finalize
  - plogp\_collectives, 55
- PLogP\_free
  - plogp, 24
- PLogP\_Gather\_bfs\_binomial\_max
  - plogp\_collectives, 58
- PLogP\_Gather\_bfs\_binomial\_min
  - plogp\_collectives, 58
- PLogP\_Gather\_dfs\_binomial\_max
  - plogp\_collectives, 57
- PLogP\_Gather\_dfs\_binomial\_min
  - plogp\_collectives, 57
- PLogP\_Gather\_ucs\_binomial\_max
  - plogp\_collectives, 58
- PLogP\_Gather\_ucs\_binomial\_min
  - plogp\_collectives, 58
- PLogP\_Gatherv\_bfs\_binomial\_max
  - plogp\_collectives, 59
- PLogP\_Gatherv\_bfs\_binomial\_min
  - plogp\_collectives, 59
- PLogP\_Gatherv\_dfs\_binomial\_max
  - plogp\_collectives, 59
- PLogP\_Gatherv\_dfs\_binomial\_min
  - plogp\_collectives, 59
- PLogP\_Gatherv\_sorted\_flat\_asc
  - plogp\_collectives, 55
- PLogP\_Gatherv\_sorted\_flat\_dsc
  - plogp\_collectives, 56
- PLogP\_Gatherv\_Traff
  - plogp\_collectives, 60
- PLogP\_Gatherv\_ucs\_binomial\_max
  - plogp\_collectives, 60
- PLogP\_Gatherv\_ucs\_binomial\_min
  - plogp\_collectives, 60
- PLogP\_hpredict\_sg\_linear
  - plogp, 25
- PLogP\_initialize
  - plogp\_collectives, 55
- PLogP\_model, 74
  - LogGP\_predictor, 74
  - logp, 74
  - n, 74
  - predictor, 74
- plogp\_p2p\_model, 75
  - serialize, 75
- PLogP\_predict\_p2p
  - plogp, 25
- PLogP\_read
  - plogp, 25
- PLogP\_Reduce\_bfs\_binomial\_max
  - plogp\_collectives, 56
- PLogP\_Reduce\_bfs\_binomial\_min
  - plogp\_collectives, 56
- PLogP\_Reduce\_dfs\_binomial\_max
  - plogp\_collectives, 57
- PLogP\_Reduce\_dfs\_binomial\_min
  - plogp\_collectives, 56
- PLogP\_Reduce\_ucs\_binomial\_max
  - plogp\_collectives, 57
- PLogP\_Reduce\_ucs\_binomial\_min
  - plogp\_collectives, 57
- PLogP\_Scatter\_bfs\_binomial\_max
  - plogp\_collectives, 58
- PLogP\_Scatter\_bfs\_binomial\_min
  - plogp\_collectives, 58
- PLogP\_Scatter\_dfs\_binomial\_max
  - plogp\_collectives, 57
- PLogP\_Scatter\_dfs\_binomial\_min
  - plogp\_collectives, 57
- PLogP\_Scatter\_ucs\_binomial\_max

- plogp\_collectives, 58
- PLogP\_Scatter\_ucs\_binomial\_min
  - plogp\_collectives, 58
- PLogP\_Scatterv\_bfs\_binomial\_max
  - plogp\_collectives, 59
- PLogP\_Scatterv\_bfs\_binomial\_min
  - plogp\_collectives, 59
- PLogP\_Scatterv\_dfs\_binomial\_max
  - plogp\_collectives, 59
- PLogP\_Scatterv\_dfs\_binomial\_min
  - plogp\_collectives, 59
- PLogP\_Scatterv\_sorted\_flat\_asc
  - plogp\_collectives, 55
- PLogP\_Scatterv\_sorted\_flat\_dsc
  - plogp\_collectives, 56
- PLogP\_Scatterv\_Traff
  - plogp\_collectives, 60
- PLogP\_Scatterv\_ucs\_binomial\_max
  - plogp\_collectives, 60
- PLogP\_Scatterv\_ucs\_binomial\_min
  - plogp\_collectives, 60
- PLogP\_write
  - plogp, 25
- predict\_p2p
  - CPM\_predictor, 63
- prediction
  - CPM\_predict\_brsg, 14
  - CPM\_predict\_flat\_sg, 15
  - CPM\_predict\_flat\_sg\_parallel, 15
  - CPM\_predict\_flat\_sg\_serial, 15
  - CPM\_predict\_flat\_sgv, 15
  - CPM\_predict\_flat\_sgv\_parallel, 15
  - CPM\_predict\_flat\_sgv\_serial, 15
  - CPM\_predict\_sgv, 14
- Prediction of communication time, 14
- predictor
  - Hockney\_model, 69
  - LinInterp\_model, 71
  - LMO\_model, 72
  - PLogP\_model, 74
- prev
  - CPM\_triplet, 65
  - CPM\_triplets, 66
- read\_hierarchy\_file
  - H\_Model, 68
- S
  - LMO\_model, 73
- serialize
  - hockney\_p2p\_model, 70
  - plogp\_p2p\_model, 75
- splines
  - LinInterp\_model, 71
- t
  - LMO\_model, 73
  - The heterogeneous Hockney model, 15
  - The heterogeneous PLogP model, 23
  - The linear interpolation model, 18
  - TreePredictor, 77
- values
  - CPM\_triplet, 65
- write\_elem
  - hockney\_p2p\_model, 70
- write\_footer
  - hockney\_p2p\_model, 70
- write\_header
  - hockney\_p2p\_model, 70